

Applying mSpace interfaces to the Semantic Web

Nicholas Gibbins¹, Stephen Harris¹, Alan Dix², and m.c. schraefel¹

¹ School of Electronics and Computer Science,
University of Southampton,
Southampton, United Kingdom
nmg@ecs.soton.ac.uk

² Computing Department,
University of Lancaster,
Lancaster, United Kingdom

Abstract. Ontologies can represent large, multidimensional spaces: classical music, research in computer science in the UK, health care for breast cancer are examples of rich domains. There have been no easy ways to represent meaningful slices through these multidimensional spaces to privilege the parts of the domain that are of interest to a given user. mSpace, an interaction model we describe here, is particularly suited to ontology-based interaction because it is designed to expose and support exploration of relations in a domain. In this paper we propose the formalism for this interaction model to support mapping this kind of user-determined interaction onto a high dimensional space represented by an ontology. The model provides semantic web designers with a means for rapidly prototyping and interrogating the data represented by an ontology. It also and provides a fast, effective UI alternative to keyword search and browsing for users to explore the domain space while maintaining domain context.

1 Introduction

The most popular method for searching the Web at present is by keyword search; it has proven to be an effective way to *retrieve* information from the Web. The effectiveness of the technique improves with the precision of the keywords used. As such, keyword searches rely on one's domain expertise to retrieve appropriate information. When domain expertise is less certain, one can use category resources like Yahoo or the Internet Project to search by subject and subcategories of subject, and narrow one's search in this way. An advantage of category search is that it communicates a sense of the range of data available in a given domain. Numbers besides category labels also help communicate the scope of the instances within that part of the domain. Indeed, there have been efforts to apply categorization to keyword search results [1], but these alert users to the categories only after the fact - based on the results of the retrieved information. One of the limitations of popular Web-based representations of category

searches is that they usually rely on fixed hierarchies of categories. That is, both the order and the number of subcategories in the tree is fixed. This means that a user cannot reorient the space to support their focus. They must approach the domain from the organizational bias of the designers.

The Semantic Web's use of ontologies to structure domains presents new opportunities to represent domains to users for their interaction with the domain space. This means that users can have greater options in how they engage with Web-based information, thus affording more diverse approaches than keyword or category search for building knowledge from the information they discover. The semantics of ontologies, for instance, privilege relationships; ontologies also afford multiple paths to view the same data instance from multiple perspectives. Exposing these relationships and paths provides knowledge seekers with additional information to inform their knowledge building tasks. Within the Semantic Web community, however, there has been little research to date on how to exploit ontologies to formalize and automate such interaction design opportunities. The closest work in this space has been in information visualization. Understandably, in that case, most work has emphasized how extant visualization techniques can be applied to the Semantic Web [2]; interaction issues have been, again understandably, secondary. Our interest, however, is specifically how interaction affordances can be enhanced by leveraging the properties of the Semantic Web. To that end, we present a formal representation of an interaction model informed by the Semantic Web. Though we present one visualization to describe the model's interaction characteristics, the model itself is visualization agnostic. Designers can use whichever visualization they choose - tree maps [3], self-organizing maps [4], polyarchies [5], lists. Our interest here is to formalize the behaviours of the data given the set of interactions defined.

In this paper, therefore, we propose a formalism for mSpace [6], an interaction model which leverages the advantages of the Semantic Web in order to facilitate user-determined exploration of a domain. In brief, the model lets users arrange an n-dimensional information space such that they can determine both a slice through the space, and then the scope, orientation, and arrangement of the attributes in that slice. A slice is determined first by the selection of attributes (effectively, class expressions) within the ontology. This selection acts like a projection through an n-dimensional space, which is then flattened. The result is a hierarchical representation of the dependencies of attributes in that hierarchy, based on the ordering of the selections. The first attribute in the order represents a query for all the instances matching that attribute/expression. Selection of instances within that listing then act as a constraint in populating the instances of the next attribute in the hierarchy and so on. The model therefore supports two levels of user interaction: manipulation of the ontology representation itself and selection of the instances of the data associated with those configurations. The logic of the model also provides for automatic reasoning across the domain to ensure that only meaningful attribute orderings/selections can occur. Which of these affordances of the interaction model a system designer wishes to implement are up to the designer.

There are two key advantages to this interaction model for domain interaction. First, mSpace provides system designers with a way to support fast visual data inspection in the domain from numerous perspectives. It also gives designers an automatic way to leverage their ontologies to present the domain to a user such that users first, can readily perceive the scope and relations within the domain from the available attributes, and second can then explore the domain from an orientation of the information that suits their interests.

In the following section, we describe the mSpace interaction model in terms of the interaction design itself. We situate the related work within this discussion. Next we describe the formal model for the behaviours associated with each interaction. We follow this with a service-level characterization of the model. We conclude with our plans for future work, including a brief sketch of an API instantiating the mSpace model.

2 Interaction Design of the mSpace model

The goal of an mSpace is to support users' *exploration* of a domain. By exploration, we mean something other than the Web sense of browsing or surfing. Surfing or browsing Web pages suggests moving among discrete Web pages which potentially have only very loose associations between them. The only association several pages may have in common, for instance, is that a user has visited each page within the past thirty seconds. By exploration, we mean that the user is making information selections within a structured information domain, based on ontological associations among its parts. The interaction design for exploration would provide mechanisms for the user to take advantage of these associations in exploring the domain. In exploration, context and the availability of contextual information to support users' exploration is a critical component of the interaction design. An mSpace, designed to support exploration, therefore privileges associations and contexts in the domain interaction and representation.

While mSpace was initially designed as an abstract interface model for exploring information domains, it seems particularly well suited to the Semantic Web, since its semantics themselves foreground the relations within a domain's ontology. Foregrounding these relations within an interaction design in turn provides users with an alternative method for interrogating an information source: while keyword searches focus on retrieving a set of matching instances from a information source, exploration focuses on representing the context of that information. Within the Human Computer Interaction community, this dichotomy between instance and relations is generally referred to as Focus + Context [7, 8]. The Web's emphasis on keyword searches has so far privileged the Focus part of the equation, if for no other reason than the unstructured nature of Web data is best searched against keyword indexing, which returns in a list of discrete instances/Web pages. The affordances of the Semantic Web raise the possibility of facilitating the Context side of this interaction equation. mSpace is designed

to leverage such semantic affordances. The remainder of this section describes the specific interaction affordances of an mSpace.

2.1 mSpace Representation and Interaction

An ontology can represent a vast, multidimensional domain that can support many ways of focusing on its data. an mSpace facilitates both the representation of the space and the interaction with this representation in a variety of ways. These are:

- slices:** projections through the n-dimensional space,
- context:** spatial rather than temporal layout of a slice
- dimensional sorting:** the organization of attributes within a slice, and how the organization constrains the data associated with each attribute.
- substitution:** the revision of a slice by changing one attribute in a slice with another not previously in the slice
- expansion/contraction:** the addition of a new attribute to a slice; the removal of a previous attribute from a slice.

In the following sections we use the example of the Classical Music domain to ground our description of these features.

Slices, Columns and Context As described above, slices represent projections through an n-dimensional space that are flattened. The flattened projection results in a hierarchical ordering of the attributes in slice. Slices mean that certain views of the domain are privileged over others. For example, the classical music domain may be sliced to privilege information about compositions in a variety of ways. One approach is to foreground information about compositions through the attributes Period (Romantic or Baroque) — Composer (Mozart or Bach) — Form (symphony or serenade) — Arrangement (solo instrument, orchestra) — Composition. Likewise the domain might be sliced to privilege information on compositions through recordings: Record Company — Orchestra — Conductor — Artist — Year Recorded — Composition. We refer to each of these attributes as Columns, in which the label of the column is the attribute name and the body of the column is the list of instances which are associated with the column - depending on its hierarchical location within the slice

Rationale: Flattened projections are not uncommon treatments for managing representations of n-dimensional spaces. Beyond making an n-dimensional space representationally manageable, flattened projections offer other benefits over visualizations attempting to approximate more than two dimensions. First, the resulting hierarchies can be readily represented in a variety of 2-dimensional views easily mapped onto the native affordance of a 2D computer screen. Three dimensions can at best only be simulated on such screens. Research has shown that 25% of the population cannot conceptualize information in 3-dimensions [9]. The same study showed that users performed significantly better in information retrieval tasks when navigating text-based hierarchies rather than either 2D map

views of information clusters of hierarchies or 3D representations of these clusters. Other research has also challenged claims about the value of 2.5 or greater dimensional views for information organization and access [10]. Therefore 1-2D representations of slices through n-dimensional spaces provide both accessibility and performance benefits over more complex dimensional simulations.

While hierarchies can be represented in a variety of visualizations, our default visualization is a spatial, multipane or multicolumn view. An example of a spatial multicolumn view for the first Classical Music slice is shown in Figure 1.




Fig. 1. Spatial multicolumn view for classical music slice

A spatial view in general means that the associated context of the information is persistently maintained in the interface. The complement of spatial views are temporal views where context is maintained in memory. For instance, the standard Web representation of information is temporal: a click on a link generally replaces the originating page with a new page in the browser window, erasing the previously visible information. While mSpaces are visualization-neutral, we strongly recommend that mSpace visualizations support spatial context: in previous work we have shown that supporting spatial rather than temporal context results in better navigation performance by users [11]. Indeed, in the same study, we found that temporal interfaces had a significant negative correlation with age. With our simple default multipane view, we get a variety of benefits for low implementation cost: there is little overhead in rendering what are effectively contiguous list views of a data space.

Dimensional Sorting. In each of the above slices, the organization of the slice imposes an ordering on the hierarchy. Thus, In a domain slice, the arrangement of the attributes within the slice act as constraints on the information associated with the attributes presented. In the first slice, if Composer is placed after Form, the selected form will determine which composers are listed: only those who set works in the selected form will be listed under composer. Putting Composer before Period allows one to see quickly which period is associated with a given

composer. This arrangement of the slice turns out to be an interesting opportunity for learning more than a single association of a period with a composer, since in certain cases, the composer is associated with multiple periods (Bach with Classic and Baroque, for instance). Thus, this one selection retrieves both the requested information, and potentially reveals meta-information about the construction of the classical music domain.

Hierarchical rearrangement is not to be confused with sorting order within a fixed hierarchy. That is, in a fixed hierarchy such as Period — Composer — Form — Arrangement — Composition, one might choose to sort the composer column alphabetically or by date of birth. This sorting has no effect on the current instances associated with each attribute. Reordering the attributes in a hierarchy, as we explain in more detail in the formal model description below, causes instances associated with each attribute to be reassessed, both in terms of the data associated with that attribute, as well as how that rearrangement effects it neighbours on either side of it, since a dimensional sort - a rearrangement of attributes in the hierarchy - effectively creates a new query within the slice.

Rationale: Dimensional sorting lets the user organize the domain in a manner that suits their current interests/knowledge. Providing multiple perspectives on an information space has long been valued by hypertext researchers not only for access for but the value of learning more about a domain [12, 13] of building knowledge by seeing associations that would not otherwise be apparent. It also lets the explorer move into the domain from a locus of familiarity. For someone less familiar with a domain (a domain-naive user), dimensional sorting may provide an approach into a domain that would not otherwise be available. Someone who knows little about classical music, but who once took piano lessons may find it easier to access the domain by privileging Instrument as the first attribute in the hierarchical slice. By supporting dimensional sorting to improve access from multiple perspectives, we are not explicitly modelling users for a particular version of a domain, but are letting users determine the domain version for their needs/tasks.

Domain Substitution, Expansion and Contraction Substitution, Expansion and Contraction are variations on Dimensional Sorting. The key difference conceptually is that unlike Dimensional Sorting, these other techniques alter the slice through the space. Substitution can fundamentally alter the slice by potentially replacing each attribute with another in the domain. Expansion/Contraction can also have a considerable effect on the slice by increasing or decreasing the scope of the slice.

Rationale: The refinements of substitution, expansion and contraction provide one more lightweight mechanisms by which users can explore a domain. Substitution provides ways to engage what-if scenarios within a domain space: what if the user extends the slice Period — Composer — Form — Arrangement — Composition with Recording — Instrumentation — RecordingYear. From exploring this extended version of the data, users may learn that period instrumentation for Bach’s keyboard work would not include its most popular form of

representation, the piano. Similarly, by substituting dimensional values one can gradually reformulate the slice to support new perspectives on an attribute. By translating an Historical slice gradually into a Recordings slice, one might also see that some very familiar works by Bach, like the solo cello sonatas, have been rearranged for a variety of instruments repeatedly over the centuries, into one of the latest for five string electric bass but have not been frequently recorded, whereas the same arrangement of Beethoven's Fifth symphony has been rerecorded by countless orchestras. In other words, by facilitating manipulation of the slices through substitution, expansion and contraction, we facilitate more ways to explore and therefore provide more ways to understand the domain and its inter-relations - and perhaps the biases informing its organization.

3 Related Work

mSpaces are largely informed by work in Human Computer Interaction in direct manipulation interfaces [14] and query previews [15]. Direct Manipulation interfaces support immediate feedback of manipulations in an interface. Thus, moving a slider up in an onscreen volume control and hearing the volume level increase is an example of direct manipulation. Changing the attribute in an mSpace column and immediately seeing the change and its effect is an example of direct manipulation. This style of interaction design is in contrast to user interfaces which change based on inferences about user interactions. Adaptive Interfaces [16], for instance, might adjust the ordering of elements in a menu based on frequency of use. Evaluations of such adaptive models have rarely been positive [17].

Query previews are a branch of direct manipulation user interfaces specifically designed for query formulation. They allow users to construct queries visually rather than textually. An interface on recipes, for instance, may present a variety of sliders with values associated with each so that one can select the degree of spiciness in a dish, whether the dish is fish, meat, or vegetarian; how long it takes to prepare and so on. Based on these constraints, the search engine returns a list of appropriate recipes. The design of available query attributes values eliminates the possibility of returning an empty set. In this way, users gain an easy method to quickly construct potentially more complex queries than a keyword search would allow, and equally can evaluate the results, tweak the parameters of the query, and run the revised query again just as quickly

Like query previews, mSpaces are direct manipulation, visual queries. Unlike most instantiations of such queries, mSpaces extend the preview query space, foreground domain manipulation affordances. Not only can users determine the attributes in the domain to query, they can organize them to privilege their evolving explorations of the domain.

Most direct manipulation interfaces are hand-crafted for a given application: the application adopts the heuristics for direct manipulation or for preview queries and creates an implementation to support that approach. In the following section, we present the formal model for the interaction design described above.

Our goal in doing so is to provide Semantic Web designers with an interaction model they can first formally evaluate, and next instantiate to rapidly provide the described interaction affordances.

Our description of mSpace using Semantic Web technologies bears some similarity to work in the Formal Concept Analysis community, particularly in the area of ontology-aware browser interfaces for information retrieval and discovery [18, 19]. These systems construct a concept lattice which is used to provide support for navigation through an information space; this concept lattice represents the different hierarchies that would result from ordering entities according to their attributes. mSpace also provides such a hierarchical view of the entities in the system, but in a more structured manner that affords the user greater control over their exploration of the information space.

The CS AKTiveSpace (CAS) Semantic Web application [20] which recently won the Semantic Web Challenge for best Semantic Web Application ³ was very much informed by the interaction design expressed above. The goal of the application was to support user-determined exploration of the domain Computer Science Research in the UK. Underneath the UI, the application exploits a wide range of semantically heterogeneous and distributed content. The content currently comprises around ten million RDF triples, with information constantly being harvested and updated. The content in the application is mediated through an ontology constructed for the application domain, the AKT Reference Ontology [21], and incorporates components from other published ontologies as well. The challenge for the UI which an mSpace approach addressed, was to make the space manageable enough to support meaningful interrogation by users. The UI would afford users the ability to explore the domain such that they could readily compare, for instance, what areas of the domain are researched in different areas of the country, who the top rated researchers are in a given domain; who do these researchers work with the most, what have they published. These are each results of queries that would be complex and potentially too cumbersome to construct to facilitate real-time shifting between one formation of a query and another. By modeling the CAS UI on the mSpace model, we were able to hide the complexity of the queries from the user, and support real-time query manipulation and ontology exploration, where data instances are revealed in context.

A large goal of the mSpace work has been to formalize the interaction model in order to create a generalizeable model for semantically informed information spaces. CAS is the first application of the mSpace model mapped into a Semantic Web back end. Through the development process of implementing mSpace affordances, such as slices and dimensional sorting, against the AKT ontology, we have been better able to determine where the issues would be in formalizing the model for general application deployment. Indeed, the effort to deploy the mSpace interaction design has reinforced the value of formalizing the model. It will make building mSpace-informed applications like CAS more principled. This is important on a practical level, since based on the demonstrations we have

³ <http://challenge.semanticweb.org/>




Fig. 2. CS AKTiveSpace: an example of an mSpace interface

given of CAS, we have been approached to design other CAS-inspired applications for granting councils, engine manufacturers and a space agency to name a few. Below, we touch on two of the ways that the formal mSpace model will improve mSpace-informed application development.

Query generation: The query generation performed by the CAS application is based on programmatic query constructions performed in the application logic of the program. This can make extracting and debugging the queries unnecessarily complex, whereas the constraint matrix suggested by this model, described below, will provide a cleaner and more general abstraction for the query generation algorithm.

Inter-column constraints: The current version of the CAS application is based on adjacent constraints, which has the consequence that as columns are rearranged (one of the more common interactions) the semantics of the column positions changes, causing undesirable and confusing changes in the contents of the columns. With the benefit of the model we see that privileged column constraints would have been more appropriate for this application (the user is attempting to narrow down selections, approaching a target group of individuals). This approach would have the advantage that re-arranging the columns would not affect the semantics of column selection.

In the next section, therefore, we describe the formal representation of the interaction model described above.

4 Modelling mSpace interfaces

As described in the previous section and in [6], an mSpace interface consists of a series of user interface controls, which we call *columns*, each of which presents a selection of objects for the user to choose from. A user's interaction with these columns progresses from the left to the right. The selection that they make in a column affects the selections which are offered in the subsequent columns that lie to its right. The selection offered by each column is a collection of objects of the same conceptual type; a list of people, or of institutions, or of publications, for example. Each of the columns represents some facet of the objects being searched, and a selection within a column can be thought of as specifying some dimension in the sparse multidimensional space of a knowledge base (where each class is treated as a dimension).

In this respect, the interface resembles the menu-based interfaces used for exploring databases; a user chooses from lists of options that lead the user through the formulation of a query. The mSpace method differs in that a context (or *column layout*) may be dynamically reconfigured so as to place one dimension at a higher importance than others.

Our basic approach in constructing this model is to consider each column as representing some class of objects. For each column i , we have three sets. T_i is the complete set of instances of the class associated with that column, D_i is the set of instances which are displayed at a given moment, and S_i is the set of instances which have been selected by the user.

$$T_i \subseteq D_i \subseteq S_i \quad (1)$$

The set T_i is the extension of the class associated with the column i . This class could be specified by a class expression in a description logic-based language such as OWL [22] or DAML+OIL [23], and might vary from a simple named class to a more complex class expression (for example, 'all higher education institutions within the UK that were rated 5* in the 2001 Research Assessment Exercise').

The relationships between columns are described in terms of constraints which determine the displayed instances in each column and the effect that a selection has on other columns. For example, the constraint C_i^j , which relates column i to column j would be:

$$C_i^j \subset T_i \times T_j \quad (2)$$

These constraints should relate the classes associated with each column, and so could also be expressed as class expressions in the description logic language. In this case, the class expressions would take the form of a restriction on a property which relates the classes associated with the two columns. These restrictions are typically (but not always) `owl:hasValue` restrictions, written in description logic terms as $\exists R.\{s\}$, where s is the selected value and R is the relating property). Multiple selections in a column are handled by forming the union of the constraints for the individual selections, or by adding both selections to the enumerated class, as in $\exists R.\{s_1, s_2\}$.

The set of constraints for a layout must provide a means to relate each column to every other column, either directly, or indirectly by chaining more than one constraint together. For a layout with n columns, the *basis set* of constraints contains $n - 1$ constraints. It is possible to construct systems with more than this minimum number of basic constraints, but this raises the possibility of inconsistency when there is more than one way to relate one column to another, possibly by chaining constraints, and those relations are incompatible or have different meanings. If we restrict ourselves to basis sets of constraints, there are two common approaches to forming constraints which we can use in mSpace systems, each with its advantages and disadvantages.

4.1 Distinguished column constraints

In this approach, the layout contains a distinguished results column to which all other columns are constrained. As shown in Figure 4.1, the constraints in such an approach are ‘star-like’.




Fig. 3. Distinguished column constraints

The instances displayed in the results column, D_{res} , will be those that satisfy the constraints from the selections on other columns:

$$D_{res} = \bigcap_i S_i C_i^{res} \quad (3)$$

where i ranges over the columns

The notation $S_i C_i^j$ is used to mean the set of values in T_j that correspond to values in S_i , according to the constraint C_i^j :

$$S_i C_i^j = \{y \in T_j : \exists x \in S_i \wedge \langle x, y \rangle \in C_i^j\} \quad (4)$$

This approach to constructing basis sets for constraints allows us to suppress the display of ‘dead end’ selections. A dead end selection in a column is one that, when made, will produce no options to choose from in the final column.

For example, if we have a column containing a list of countries that includes the Vatican City, and our final column contains a list of computer science researchers, but we have no information about the state of computer science research in the Vatican, then we remove Vatican City as a choice in the country column. In effect, we are limiting the system in such a way as to only let the user answer questions that it can answer. The instances displayed in column i (the members of the set D_i) will be those instances that satisfy the constraint C_i^{res} for some value of D_{res} :

$$D_i = C_i^{res} D_{res} \quad (5)$$

4.2 Adjacent column constraints

In the second approach, each column provides a constraint, which we call an *adjacent column constraint*, to its immediate successor in the layout.




Fig. 4. Adjacent column constraints

$$S_{i+1} \subseteq D_{i+1} = S_i C_i^{i+1} \cap C_{i+1}^{i+2} C_{i+2}^{i+3} \dots C_{n-1}^n T_n \quad (6)$$

where n is the number of columns

In Figure 4.2, this type of constraint is denoted by C_n^A where n is the column from whose selection the constraint is derived (the constraint is applied to column $n+1$). Since this constraint set is a basis, we can determine the pairwise effect of a selection on any other column by composing the basis constraints. For example, in Figure 4.2 the constraint on column three from column one is $C_1^A \circ C_2^A$.

In practice, these composed constraints are not used in implementations (where the constraints are evaluated column by column from left to right); we describe them in order to illustrate the parallels between the two approaches to building a basis set.

Constraint example As an example of the above constraint schemes, consider a mSpace interface for exploring countries, the institutions located in those countries, the projects in which those institutions are involved, and the people who




Fig. 5. mSpace column constraints

work on the projects. We choose an initial layout that consists of a column of institutions, then of projects, and finally of people, as shown in Figure 4.2. The ontology on which the interface is based defines the four named classes *Country*, *Institution*, *Project* and *Person*, each of which is used as the type constraint on the relevant column. Institutions are related to projects by the *contributes-to* property, people to projects by the *works-on* property, people to institutions by the *works-at* property, institutions to countries by the *located-in* property and people to countries by the *born-in* property.




Fig. 6. mSpace example

The adjacent selection basis constraint from countries to institutions is of the form $\exists \text{located-in}.\text{Country}$, while that from institutions to projects is $\exists \text{contributes-to}^-\text{Institution}$, and that from projects to people is $\exists \text{works-on}.\text{Project}$. Initially, the only constraints applied to each column are the type constraints and the initial forms of the selection constraints given above. As selections are made in columns, the selection constraints are modified so that the (named) class in the local range restriction is replaced by an enumerated class whose members are the selected instances (we refer to such post-selection constraints as *concrete selection constraints*).

The distinguished column selection constraint basis set consists of \exists *works-at.Institution*, which runs from institution to person, \exists *works-on.Project*, which runs from project to person, and \exists *born-in.Country*, which runs from country to person. As in the adjacent case, these constraints are modified as selections are made. The concrete selection constraint between the institution and project columns, which takes into account the existence of people which satisfy the selections made in these columns, is:

$$\exists \text{ works-on}^- . (\exists \text{ works-at} . \{s_1, s_2, \dots\})$$

where s_1, s_2 and so on are the individuals which have been selected in the *Project* column.

In this example, where the selection constraints are straightforward owl:hasValue restrictions on single properties, the pairwise selection constraint is an owl:hasValue restriction on the composition of one of the properties with the inverse of the other (eg. \exists (*works-on*⁻ \circ *works-at*).{*s_{inst}*}).

So far, we have described layouts which use owl:hasValue restrictions as selection constraint class expressions, but these are not the only possibility. We can also form constraints of the form $\forall R.\{s_1, s_2 \dots\}$ (in OWL terms, an owl:allValuesFrom restriction to an enumerated class defined using owl:oneOf), and can incorporate other restrictions which are not parameterised by selections.

4.3 Layout operations

The description of intra- and inter-column constraints above assumes a fixed layout of columns, whereas an mSpace interface allows the user to modify the layout at will. We define two primitive operations on a layout: the *addition* of a column within the layout, and the *removal* of a column from the layout. As they are, these primitive operations are unlikely to be supported in a given mSpace interface because they are at too low a level to fit well within the paradigm of dimensional exploration introduced by mSpace. Instead, we define three distinct types of higher-level operation which can be expressed as a composition of the primitive operations, and which an mSpace interface would be expected to support (illustrated in Figure 7):

Expansion: A new column may be added to the right of the existing columns

Substitution: A column within the layout may be switched with a column of a different type

Transposition: The positions of two columns within a layout may be exchanged. This is a specific example of a dimensional sorting operation – there are others.

The operations that may be performed on a column affect the constraints on a given layout to a greater or lesser extent depending on the selection constraint scheme being used. In turn, the change in the constraint set caused by the operation may affect the selections which the user has made (by affecting the




Fig. 7. mSpace operations

set of options from which they made the selection), which in turn may change more instantiated constraints, and so on.

We describe operations which cause the invalidation of selections as *disruptive*. For example, if a selection has been made in a column, and the column is substituted for one of a different kind, the selection is lost, which affects any presented options or selections in subsequent columns. Similar behaviour is to be expected in the event that two columns are transposed, although it may be possible to retain the selections in the columns provided they do not violate the new selection constraints.

In the general case, we cannot easily characterise this disruptive propagation of changes, but we can determine what the change to the constraint set will be in response to the initial operation for the two selection constraint schemes. We describe the changes brought about by an operation in an abstract fashion, by counting the number of selection constraints which must be changed. These are listed in Table 1; $+n$ indicates that n constraints are added to the set, $-n$ that n are removed and $\pm n$ that n are changed.

	Adjacent		Distinguished column	
	first/last	other	distinguished	other
add	+1	-1, +2	$\pm(n-1)$	+1
remove	-1	-2, +1	$\pm(n-1)$	-1
expand	+1	—	$\pm(n-1)$	+1
transpose	± 3	± 4	$\pm(n-1)$	0
substitute	± 1	± 2	$\pm(n-1)$	± 1

Table 1. Operation cost on column based on constraint scheme

Since the observed flow of information in an mSpace interface runs from left to right, an operation which changes a column in a layout using adjacent selection constraints is likely to change those columns which lie to its right. For example, the expansion of an existing layout under the adjacent scheme by the addition of a new final column is a non-disruptive operation which does not involve the removal or replacement of existing constraints (shown as +1 in Table 1).

Conversely, under a distinguished column scheme (where the final column is the distinguished column), all the selection constraints are replaced with ones suitable for use with the new final column. In this case, the change in the type of the final column leads to the introduction of a new set of selection constraints, which may be violated by the existing selections. Since constraints are applied in a pairwise manner under this scheme (going from one column of a pair to the distinguished column, and then back to the other column in the pair), a change in the type of the distinguished column affects all selection constraints, which raises the prospect that the selections in all existing columns may be invalidated (shown as $\pm(n - 1)$ in Table 1).

Similar behaviours exist for the operations of transposition and substitution under either of the constraint schemes described above. In general, adjacent selection constraints yield more disruptive operations for small layouts with less than five columns (there are no distinguished columns, so the cost is mostly constant across the columns in a layout, the first or last columns being an obvious exception).

In contrast, distinguished column selection constraints designate one column as special and express all selection constraints in terms of that column. Operations on columns other than the distinguished column are less disruptive than their counterparts under adjacent selection constraints (note the values for transposition in Table 1). However, operations on the distinguished column become significantly more expensive because they change the entire constraint set.

Thus, a rule of thumb which informs a designer's choice of selection constraint scheme is this. For interfaces where dead end suppression is considered essential, or where layouts have fewer than five columns, or where the expansion operation is not to be supported, distinguished column selection constraints are the appropriate approach, otherwise adjacent selection constraints should be chosen. One proviso to this is that expansion under distinguished column constraints, where the final column is not the distinguished column, are no more expensive than expansion under adjacent constraints.

4.4 Permitted layouts

In certain circumstances, an arrangement of columns may not make sense. For example, a column which required the user to pick from a list of university departments should not be followed by a column which requires the user to pick from a list of universities; there is a dependency between the types of these columns which makes the second column irrelevant. This is particularly true of layouts which use the adjacent selection constraint scheme, where a column directly depends on its predecessor.

We express this behaviour in a successor matrix which relates each column type to the column types after which it may appear; the column types used in this matrix are the type constraint class expressions described in the previous section. When a user is presented with a layout, the set of permissible operations which they may use to rearrange the layout is constructed by consulting this matrix. For example, the user is presented with a list of column types which may be used to expand the current layout that is generated by taking the set of all possible column types and removing those which are prohibited by the columns already in the layout.

Similarly, the transposition of two columns is a two part operation. The user selects the first of the columns to be transposed, then the columns with which that column may be swapped (which would yield a permissible layout when swapped) are indicated. The user then selects the second column from those indicated, and the layout is rearranged.

4.5 Column presentation

The model of mSpace outlined above concerns itself primarily with the generation of the abstract choices offered to the user in each column and the propagation of the effects of making a selection through the system, but this is not the only concern. If each column represents an set of objects from which the user must make some choice, the manner in which that set is rendered for presentation to the user is also of importance.

In our prototype mSpace system, CS AKTiveSpace, we have implemented a column view for selecting UK higher education institutions which is rendered as a map of the UK showing the geographical distribution of HEIs. This column need not have been rendered in this fashion; it could have been shown as a simple multipick widget containing the names of the institutions.

In each case, the information required to render the column is obtained from the knowledge base (human-readable labels and latitude/longitude coordinates for the HEIs). The design of an mSpace interface must therefore not only include the set of column types, their permitted combinations and the constraints which relate them, but also the alternate rendering styles for each column.

Our CS AKTiveSpace system also illustrates the use of ordering and limiting constraints on columns. The final column contains a list of people (filtered by the preceding columns) which may be ordered according to various criteria (we allow ordering by total grant income and by the research rating of their institution). Similarly, the number of people shown in this column may be limited to the first five, ten, etc. This behaviour is separate from the selection-based behaviour described earlier in this section and is largely a presentational issue; we implement it as a set of filters which apply to the column after the application of inter-column constraints.

We treat the issue of column presentation as a level above that of the abstract column model, on which it has no direct bearing. In this respect, mSpace interfaces fit well within the model-view-controller paradigm, the model can be

effectively decoupled from the view/controller pairs through the adoption of a clean API by which the view interacts with the model.

4.6 Detail view presentation

As with column presentation above, we consider the generation of the detail view provides the user with contextual information as they interact with the system to be a purely presentational issue that has little to no impact on the interactions between the columns described in our formal model. The detail view provides information (at a fairly coarse granularity) about the entity that was most recently the focus of the user’s attention, namely the most recent selection that they made in a column.

5 Implementation

The characterisation of an mSpace user interface in terms of description logic expressions given above should be considered to be an abstract description of such a user interface, and not a detailed specification of an implementation itself.

In our prototype Semantic Web application, CS AKTiveSpace, we chose to limit the class expressions used as selection constraints to be `owl:hasValue` restrictions. We further chose to implement these constraints in a naïve manner which does not require recourse to a DL reasoner; constraints were represented as triple patterns containing variables which were bound as selections are made. For example, the constraint $\exists \textit{works-at.Institution}$ (used as a distinguished column selection constraint in the example) would be replaced with the triple patterns $(?person, \textit{works-at}, ?institution)$ and $(?institution, \textit{rdf:type}, \textit{Institution})$. These triple patterns were expressed in the RDQL language, and used as queries to generate the values with which panels were populated.

This approach has the advantage of simplicity, because a constraint and its inverse have exactly the same form. The triple patterns do not assume that the constraint will be read in the direction of a relation or of its reverse; the reading of the constraint depends on which variable is bound (by a selection), and which is free (and generates potential selection with which a column is populated). In addition, the RDQL query language is supported by a number of RDF stores and inference engines, including 3store [24] (on which we have built our CS AKTiveSpace application), Sesame [25] and Jena [26], so providing some flexibility in our choice of infrastructure.

The disadvantage of this approach is that it permits only one particular type of constraint, and that the construction and composition of disjunctive query patterns (as would be the case when multiple selections are made in a column), is clumsy in RDQL.

6 Conclusion and Further Work

In this paper, we have presented a treatment of the mSpace style of user interfaces which builds on Semantic Web technologies. In doing so, we open up

possibilities for more strongly correlating the interaction behaviours with the semantics inherent in the information represented by the system.

Our plans for future work include the development of a framework to facilitate the construction of mSpace-based systems (drawing on our experiences of implementing our prototype CS AKTiveSpace system), and the investigation of techniques for automating the development of mSpace systems from ontologies. From this, we intend to develop an API and toolset to make deploying mSpace in OWL-based Semantic Web applications a plug and play operation. To this end, part of our goal with this paper is to generate discussion about and feedback for our model, to better inform the development of an API.

7 Acknowledgements

This work was supported by the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC). The AKT IRC is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01 and comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

The authors would like to thank Nigel Shadbolt and Srinandan Dasmahapatra for their comments on this work.

References

1. Vivisimo: Vivisimo document clustering engine. <http://www.vivisimo.com/> (1970)
2. Geroimenko, V., Chen, C., eds.: Visualizing the Semantic Web. Springer (2002)
3. Bederson, B.B., Shneiderman, B., Wattenberg, M.: Ordered and quantum treemaps: Making effective use of 2d space to display hierarchies. *ACM Trans. Graph.* **21** (2002) 833–854
4. Chen, H., Dumais, S.: Bringing order to the web: Automatically categorizing search results. In: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'00). (2000) 145–152
5. Robertson, G., Cameron, K., Czerwinski, M., Robbins, D.: Polyarchy visualization: visualizing multiple intersecting hierarchies. In: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press (2002) 423–430
6. m.c. schraefel, Karam, M., Zhao, S.: mspace: interaction design for user-determined, adaptable domain exploration in hypermedia. In: Proceedings of the AH2003 Workshop on Adaptive Hypermedia and Adaptive Web-Based Systems, Nottingham, UK (2003) 217–235
7. Ridsen, K., Czerwinski, M., Munzner, T., Cook, D.: An initial examination of ease of use for 2d and 3d information visualizations of web content. *IJHCS* **53** (2000)
8. Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., Roseman, M.: Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer Human Interaction (TOCHI)* (1996) 162–188

9. Modjeska, D., Waterworth, J.: Effects of desktop 3d world design on user navigation and search performance. In: Proceedings of Information Visualization 2000, IEEE (2000) 215–220
10. Cockburn, A., McKenzie, B.: Evaluating the effectiveness of spatial memory in 2d and 3d physical and virtual environments. In: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press (2002) 203–210
11. m.c. schraefel, Karam, M., Zhao, S.: Preview cues for exploring domain hierarchies. In: Proceedings of Interact 2003, Switzerland (2003)
12. Neuwirth, C., Chimera, D.K.R., Gillespie, T.: The notes program: A hypertext application for writing from source texts. In: Proceeding of the ACM Conference on Hypertext. (1987) 121–141
13. Shipman, F., Moore, J., Maloor, P., Hsieh, H., Akkapeddi, R.: Spatial hypertext: Semantics happen: Knowledge building in spatial hypertext. In: Proceedings of the Thirteenth Conference on Hypertext and Hypermedia. (2002) 25–34
14. Chimera, R., Shneiderman, B.: An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents. ACM Transactions on Information Systems (TOIS) **12** (1994)
15. Plaisant, C., Shneiderman, B., Doan, K., Bruns, T.: Interface and data architecture for query preview in networked information systems. ACM Transactions on Information Systems (TOIS) **17** (1999) 320–341
16. Sukaviriya, P.N., Foley, J.D.: Supporting adaptive interfaces in a knowledge-based user interface environment. In: Proceedings of the 1st international conference on Intelligent user interfaces, ACM Press (1993) 107–113
17. McGrenere, J., Baecker, R.M., Booth, K.S.: An evaluation of a multiple interface design solution for bloated software. In: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press (2002) 164–170
18. Kim, M., Compton, P.: A web-based browsing mechanism based on conceptual structures. In: Conceptual Structures: Extracting and Representing Semantics: Contributions to the 9th International Conference on Conceptual Structures (ICCS2001). (2001) 47–60
19. Eklund, P., Cole, R.: Structured ontology and information retrieval for email search and discovery. In: Foundations of Intelligent Systems, 13th International Symposium (ISMIS 2002). (2001) 75–84
20. Shadbolt, N., monica schraefel, Gibbins, N., Glaser, H., Harris, S.: CS AKTiveSpace: Representing Computer Science in the Semantic Web. <http://triplestore.aktors.org/tmp/www2004-cas.pdf> (2003)
21. : The AKT Reference Ontology. <http://www.aktors.org/publications/ontology/> (2002)
22. McGuinness, D.L., van Harmelen, F.: OWL Web Ontology Language Overview. Candidate recommendation, World Wide Web Consortium (2003) <http://www.w3.org/TR/owl-features/>.
23. Connolly, D., van Harmelen, F., Horrocks, I., McGuinness, D.L., Patel-Schneider, P.F., Stein, L.A.: DAML+OIL (March 2001) Reference Description. W3C Note, World Wide Web Consortium (2001)
24. Harris, S., Gibbins, N.: 3store: Efficient Bulk RDF Storage. In: Proceedings of the First International Workshop on Practical and Scalable Semantic Systems (PSSS2003), Sanibel Island, Florida, USA (2003) 3–17
25. Broekstra, J., Kampman, A., van Harmelen, F.: Sesame: An Architecture for Storing and Querying RDF Data and Schema Information. In: Spinning the Semantic Web. MIT Press (2003) 197–221

26. Labs, H.P.: The Jena Semantic Web Toolkit. Web page, Hewlett-Packard Labs (2003) <http://www.hpl.hp.com/semweb/jena.htm>.