

Opening the information bottleneck in complex scheduling problems with a novel representation: STARK diagrams

Peter C-H. Cheng¹, Rossano Barone^{1,2}, Peter I. Cowling², Samad Ahmadi^{1,2}

¹ESRC Centre for Research in Development Instruction and Training,
School of Psychology

²ASAP Group, School of Computer Science and Information Technology,
University of Nottingham, Nottingham, NG7 2RD, UK.

peter.cheng@nottingham.ac.uk rb@psychology.nottingham.ac.uk
pic@cs.nott.ac.uk s.ahmadi@cs.nott.ac.uk

Abstract. This paper addresses the design of representational systems for complex knowledge rich problems, focussing on scheduling in particular. Multiple tables are ubiquitous in representations of schedule information, but they impose large cognitive demands and inhibit the comprehension of high-level patterns. The application and evaluation of representational design principles in the development of STARK diagrams, a novel system for scheduling problems, is reported. STARK diagrams integrate conceptual dimensions, principal relations and individual cases into a single diagrammatic structure. An experiment compared performance on STARK diagrams and a conventional representation with features typical of current commercial scheduling software interfaces. Subjects using the STARK diagram performed better at improving an examination schedule by minimising constraint violations. This provides support for the validity and utility of the design principles.

1 Introduction

The critical role of problem representations has been well established by research in cognitive science [1, 10, 15]. A problem can be more than an order of magnitude more difficult to solve with a poor representation than a good representation [11]. There are potentially substantial benefits of diagrammatic representations in contrast to sentential representations [12]. Mental effort and computation can be off-loaded onto external representations [14]. By matching the informational dimensions with appropriate visual dimensions it is possible to make representations that support the discovery of patterns and efficient inferences [18].

Although there has been research on how these and other findings on the nature of representations can be applied to the design of effective representations [5, 7, 16], the science of representational design is in its infancy. Given a particular domain, how should a representation be designed to support the different tasks that need to be performed and to support users with different levels of knowledge and experience of the domain? One of the grand challenges to this growing area is the design of effective representational systems for information intensive and conceptually

demanding domains. In previous studies with knowledge rich domains, including mechanics, electricity and probability theory, we have designed novel representational systems to support problem solving and learning in instructional contexts [6, 7]. By comparing the new representations with the existing conventional notations, including the empirical evaluation of the representations, principles for the design of representations for complex conceptual domains have been discovered [8].

This paper describes work on a project that is applying the principles to the analysis and design of novel representations for real-world informationally intensive scheduling problems. Automated systems are typically used to solve such problems, but the way they operate tends to be inflexible and difficult for users to understand. The solutions they produce are only as good as the model of the problem they are given, so rare circumstances or idiosyncratic requirements are not handled. The context of this project is to attempt to humanise such automated systems by designing new representations that will allow the flexibility and creativity of humans to be integrated with the computational power of automated systems. The aim is to get the best of both worlds by using the particular advantages of the approaches to overcome each others' specific limitations.

Examination timetabling is the first scheduling problem that has been addressed by the project. A novel representation, STARK diagrams (Semantically Transparent Approach to Representing Knowledge), has been designed for this class of scheduling problem. This paper describes the principles and how they were used to design STARK diagrams. It also reports on the design of a conventional representation that served as the basis for comparison in an experiment that has been conducted to compare how well users can manually improve an examination timetable using the two representations. The outcomes of the experiment are discussed and the implications for the validity and utility of the principles considered. The nature of the examination scheduling domain will first be outlined.

2 Examination scheduling

Examination scheduling is a complex organisational problem which occurs in many educational institutions across the world. It is often solved by modelling as an NP-hard combinatorial optimisation problem that demands the allocation of exams to rooms and time periods under a high density of constraints [4]. These problems are large, typically involving tens of rooms and periods, hundreds of exams, and thousands of students. Solutions to the problems are typically generated by automated software systems, with the user defining a fixed set of rooms, days and periods within each day. Together with data for students, exams and constraints this information is used to generate solutions that the user may then edit.

Several different types of constraints exist in examination scheduling problems. Individual universities may differ in the kinds of constraints they employ but common to most are two categories that we term *resource* and *intersection* constraints. For a description of other constraints used in examination scheduling problems see [3]. Resource constraints concern the availability of space and time for the allocation of an exam to a given room and period. Room capacity violations occur when there is insufficient space to hold the number of students in the exams

allocated to a room. Examination schedules are organised into daily periods of unequal size. The inequality of period size demands the specification of the second type of resource constraint, which maintains that an exam should not be allocated to a period that is shorter than the duration of the exam. A violation of this constraint is termed a period excess violation.

Intersection constraints are concerned with the temporal proximity of allocated exams that involve the same students. They are termed intersection constraints because they occur as a consequence of the intersection, or sharing, of students between two exams. Clash exam violations occur when intersecting exams are allocated to the same period because students cannot take more than one exam simultaneously. Consecutive exam violations result when intersecting exams are allocated to adjacent periods. In this case students would have to take two exams in succession. Clash and room capacity constraints are often treated as inviolable or *hard* in optimisation approaches, with solution quality evaluated with respect to the degree of violation of *soft* constraints such as consecutive exams. However, this distinction is somewhat artificial in practice, and we treat all constraints as being soft, giving appropriate weights depending upon the relative importance of each one.

Despite the raw computational power of automated systems the simplified models of the problem they employ leaves much scope for human intervention. The solutions produced by automated systems will often satisfy the clash and room constraints but leave many soft constraints still violated, in particular those which were not present in the original model. The user must try to find new allocations for the violated exams that satisfy all the constraints. Resolving these violations is not straightforward otherwise the automated software would satisfy them. To effectively improve a solution the user needs to access and integrate many sources of different kinds of information. The manner in which this information is represented can substantially determine the users capacity to manually improve the solution.

Conventional ways of representing examination scheduling problems typically involve organising the problem data into numerous lists and displaying these lists in different windows or frames of reference. Advanced interfaces may provide *a la carte* tables that the user can construct to show particular perspectives of interest. Whilst conventional scheduling interfaces may be both sophisticated and flexible, the presence of multiple tabular representations tends to inhibit, rather than facilitate, user intervention. There are three common limitations of traditional scheduling interfaces: (a) they impose large demands on cognitive resources; (b) they support local inspection of information but do not provide global overviews; (c) they play little or no role in constraining user behaviour due to their profound lack of conceptual structure.

The principles of representational design have been used to design a novel representation that attempt to overcome these limitations.

3 Principles of representational design

The principles were derived in our previous work on the design of novel representations to enhance problem solving and learning by re-codifying knowledge in conceptually demanding educational domains including probability theory and

electricity [6, 7, 8]. Six principles have been formulated to date and have been classified according to (1) whether they prescribe that the conceptual structure of the domain should be made apparent in the structure of the representation, *semantic transparency*, or (2) whether they concern efficient problem solving operators and procedures, *syntactic plasticity*.

The semantic transparency principles consider how the meaning of a domain can be made clearly apparent in three ways:

(1) **Integrate levels of abstraction.** Different levels of abstraction should be integrated to reduce the conceptual gulf between (a) overarching laws or general relations that govern a domain and (b) specific cases or instance at a concrete level. In a representation that integrates levels of abstraction extreme cases will help interpret the general nature of the laws and the laws will explain the form of typical cases.

(2) **Integrate alternative perspectives.** Alternative perspectives, including alternative ontologies [9], can be used to describe a domain. Perspectives at the same level of granularity or abstraction should be integrated in an effective representation, to allow the alternative perspectives to act as mutual contexts for each others' interpretation.

(3) **Combine a globally homogenous with a locally heterogeneous representation of concepts.** An interpretative framework should be provided that simultaneously combines (a) a globally coherent interpretative scheme based on the principal conceptual dimensions of the domain with (b) local representational features that make specific conceptual distinctions clear. A principle conceptual dimension is a property or aspect that is universal to the domain. Time and space are such dimensions for many domains. Under a given global dimension, one way a specific conceptual distinction may be identified is by the use of alternative scales of measurement for the different things under the global dimension (i.e., ratio, interval, ordinal or categorical scales). A representation with such an interpretative framework should support the making of valid generalisations whilst reducing the chance of over generalising specific concepts.

There are three ways in which the syntactic plasticity principles can make problem solving with a representation easier:

(1) **Malleable expressions.** The expressions of a representation should not be too rigid nor too fluid, they should be malleable. A rigid representation lacks the procedures to allow all the meaningful expressions needed in problem solving to be generated. This will cause dead ends to occur during problem solving. A fluid representation allows many arbitrary expressions to be generated at each potential solution step, so making the space of possible expression impracticably large to consider. A malleable representation sails a middle course between these extremes.

(2) **Compact procedures.** The procedures for solving problems in a representation should be compact, in the sense that the typical number of operations needed to find a solution from the initial state should be relatively small. The fewer the operators the less computation that is needed and the less the chance of making errors.

(3) **Uniform set of operations.** A representation should have a small variety of consistent and uniform operators making up its problem solving procedures. The fewer the types of operators the simpler the overall approach to problem solving is likely to be.

4 Applying the principles to scheduling

The application of the principles to examination scheduling requires the conceptual structure of the domain and the nature of the problem solving activities to be specified. In doing this, the term *slot* is taken to mean the conjunction of a particular room in a specific period to which an exam may be assigned.

In examination timetabling there are two aspects over which levels of abstraction should be integrated. First, there are the implicit relations due to the underlying physical nature of the domain. At the abstract level this means that arithmetic operations apply to the sizes of slots and to non-intersecting exams (e.g., free room capacity = total room capacity - \sum sizes of exams allocated to the room). Also at the abstract level set theoretic notions apply to exams as "members" of a slot and to students as "members" of exams (e.g., students taking both exam-A and exam-B = (exam-A \cap exam-B)). At the concrete level is information about the size and duration of actual exams and slots. The second aspect concerns the quality of solutions. At the abstract level the quality of a solution is assessed by a mathematical model called the evaluation function. At the concrete level are actual distributions of allocations and constraints over which the evaluation function computes the quality of a solution.

The different perspectives in the domain that should be integrated include: space and time resources that together define available slots; demands in the form of exams to be allocated; constraints that are to be satisfied and violations to be eliminated; intersections amongst exams at the level of students, in particular the *intersection set* for a target exam that comprises the group of all other exams sharing students with the target exam.

The global conceptual dimensions of the domain to be incorporated into an overarching interpretative scheme include: time, space, constraints, and types of entities. Time is obviously uni-dimensional. For timetabling, space is also essentially one dimensional, because it is sufficient to distinguish different rooms without specifying their location in space and because size and capacity, number of students, are one-dimensional properties. On the time dimension local conceptual distinctions to be made include: days, periods, duration of periods, duration of exams. On the space dimension rooms, room capacity, exam location, and exam size are to be distinguished. Constraints and violations can be distinguished on the basis of the things to which they apply. With respect to students there may be exam clashes or consecutive exams. For exams themselves there may be preferences for certain orders of exams or for holding some exams simultaneously. For periods and rooms there may be preferred times or locations, such as a laboratory for a practical exam. Finally, exams and slots are different entities to be distinguished.

What are the problem solving procedures that must be supported by an effective representation for exam scheduling? The class of examination scheduling problems that was chosen for the experiment was the manual improvement of an exam schedule solution by the reallocation of exams. In this context, a specification of a solution constitutes an expression and a new expression is generated whenever an exam is moved to a new slot. The nature of the problem and the kinds of problem solving procedures that human problem solvers typically use were considered. Users' choices in solution improvement should be guided by the information present in the current solution. It would be ideal for a user to be able to: (a) clearly identify

the constraint violations in the solution; (b) prioritise those exams that are causing the most problems; (c) estimate with some degree of precision how easy it is to satisfy the problem exams; (d) make an informed decision about which exams to reallocate based on some estimation of the cost of moving the exam relative to the gain expected in solution improvement; (e) clearly evaluate the improvement or lack of it following each reallocation.

In examination scheduling much of the cognitive work relies on the capacity to access, compare and integrate detailed information. An examination scheduling representation that is a bottleneck to such information is likely to result in greater errors, inhibit exploration and the generation of novel strategies, and may ultimately frustrate the user. Our interviews of professional examination schedulers suggest that this is precisely the case. These schedulers expressed little confidence in their capacity to manually improve a solution with current systems and appeared to depend entirely on automated solution generation. For effective examination scheduling ease of access to relevant information appears to be an important factor which is missing from the current generation of decision support software for the problem.

The next two sections consider how the structure of the new STARK diagram representation and the conventional representation were designed according to the principles discussed above.

5 STARK diagram representation

Figure 1 shows a screen span shot of the STARK diagram interface designed with the principles and as used in the experiment. For ease of inspection a magnified view covering three days and nine rooms is shown, but a full size schedule can be reasonably viewed on a standard computer monitor. STARK diagrams combine globally homogeneous and locally heterogeneous representations of concepts. For global homogeneity, a general interpretive scheme is provided that includes space, time, constraints and entities. Time is represented along the horizontally axis and space is represented along the vertical axis. Basic timetable entities are represented by rectangular icons. Most constraints are represented by lines connecting different graphic components and others by specific configurations of icons. The specific conceptual distinction between exams and slots (a room for a given period) is made using light yellow icons and dark blue icons, respectively. On the time dimension periods are shown by columns of slot icons. Days consist of three periods and are represented by a group of three columns. The duration of a period or an exam is given by the width of its icon. Under the space dimension each row of slot icons represents a room (available over time) and the exams allocated to a particular slot are shown by exam icons contained within its slot icon. The capacity of a room or the size of an exam is shown by the height of the respective icon. This representational structure provides an interpretative scheme for identifying exam and room quantities and also provides a frame of reference for locating and making inferences about the temporal proximity of specific exams, rooms, periods and days.

The specific conceptual distinctions under the global constraints dimension also exploit this scheme, as shown in Figure 2. The schemes allows three types of information to be encoded: (1) the type of constraint violation; (2) the exams

involved and their present allocation; (3) the number of students involved. Resource constraints such as room capacity and period excess are represented as exam icons overflowing the slot icons containing them, Figure 2a. Such constraints are distinguishable from intersection constraints, which are represented by lines connecting offending exams, Figure 2b. Types of intersection exam are themselves differentiated by the colour and orientation of the connecting lines.

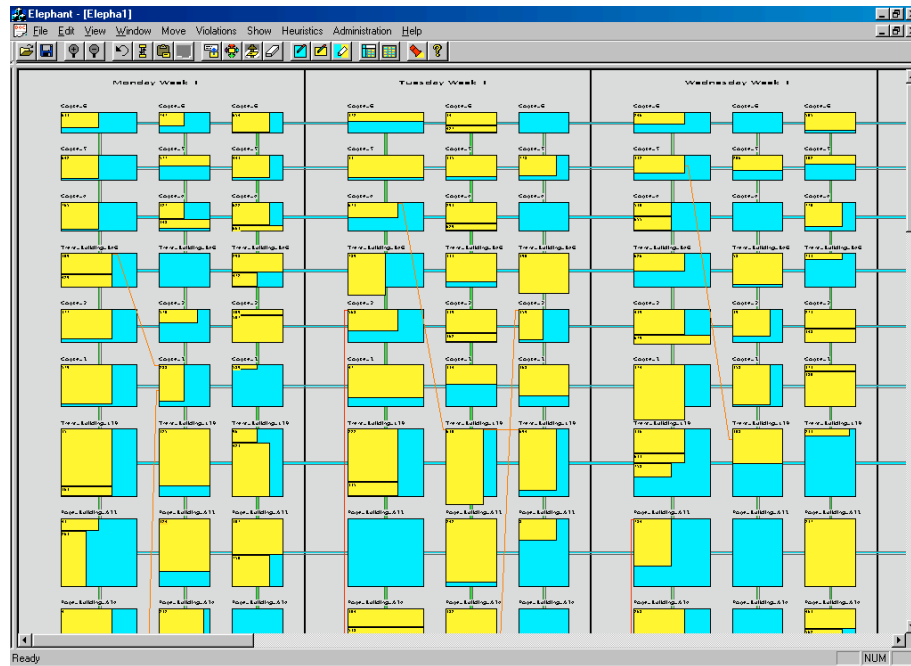


Fig. 1. A section of a STARK diagram examination schedule




Fig. 2. Some constraints in the STARK diagrams: (a) top left – over capacity, bottom right – period exceeded; (b) vertical line – exam clash, diagonal connecting line – consecutive exam

With respect to the underlying physical regularities of the domain, levels of abstraction are integrated in STARK diagrams by: (1) using geometric configurations to capture arithmetic relations and spatial arrangements to capture set theoretic relations; (2) encoding quantities as sizes of icons. The assignment of a number of students to an exam and the allocation of an exam to a slot are represented by viewing exam and slot icons as containers. For example, the exam icons for two exams

that are allocated to the same slot are not permitted to overlap. The total time a period is exceeded is given by the amount that the offending exam icon horizontally extends beyond the boundary of the slot icon. Connections between the underlying physical regularities of the domain and actual cases are thus directly visualised.

Another important aspect of abstraction is the evaluation function. The STARK diagram allows the user to view the solution and evaluate its quality, at some level of precision, using his/her own model of the evaluation function rather than relying solely on a given mathematical model. This is feasible because STARK diagrams not only support the recognition of individual instances but allow distributions of allocations and constraints to be judged. For instance in Figure 1 the distribution of exams is fairly uniform, but some slots are empty .

It is not surprising that the perspectives of resources, demands and constraints are integrated in STARK diagrams, because the global interpretive scheme combines space, time, entities and constraints. The resource perspective is represented by the space-time matrix of slots. The demands perspective is represented by the layer of exam icons distributed over that matrix. The constraints perspective is shown by lines connecting exam icons and by the overflow of exam icons over slot icon boundaries. The intersection set perspective is also integrated with the others under the interpretative framework, but information from this perspective is only made available when a target exam is selected. Figure 3 shows how when a target exam is selected all exams that share students with that exam are highlighted in dark red. The black region at the top of each of these intersecting exams denotes the number of students that are also taking the selected exam. The integration of the intersection set perspective with the other perspectives has a substantial impact on the nature of the problem solving procedures used with STARK diagrams.



Fig. 3. Intersection set for a selected target exam coloured white (row 3, column 5)

STARK diagrams constitute a malleable representation for examination scheduling, because they make problem relevant information salient that greatly constrains

the choice of manual changes that are meaningful to make. Consider two examples of this. First, a user can also readily judge whether the target exam will violate resource constraints in a period by comparing the size of the target exam icon with the size of the rooms considered for potential reallocation. Second, being able to identify those exams that share the same students with a target exam, intersection set information, is particularly important. Suppose the selected exam in Figure 3 is to be reallocated. If it is reallocated to the same period as any of its intersecting exams (in red), such as the immediately preceding or following period, a clash violation would occur. If reallocation is to any period immediately before or after a period with intersecting exams, such as column 3 in Figure 3., then consecutive violations would again occur. The only periods free from such potential constraint violations are the 10th and 11th columns. By supporting such meaningful judgements the representation is malleable because it restricts the user's options to ones that are good.

The accessing of information is critical to problems in exam scheduling. If the procedures for finding relevant information are not simple the overall problem solving procedures are likely to be complex and involved, so not compact. In this respect STARK diagrams are compact in three ways. First, there is a single frame of reference for all the major classes of information, so information needed for solution improvement can be read directly from the diagram. Second, many useful relations that are not usually stated explicitly in examination schedules appear as emergent features in STARK diagrams. Third, the representation exploits the built-in discriminatory power of the human visual system to make relevant information "pop out"; for example the red exam icons for the intersection set, as in Figure 3. These representational devices mean that laborious searches for information, in the form of indistinct textual labels spread across multiple lists, are avoided.

This ease of information access allows sophisticated problem solving strategies to be adopted. For example, in many solutions it is impossible to reallocate a given exam without clearing space for them by reallocating other exams in the intersection set of the given exam. In turn these other exams will each have their own intersection set. This requires a recursive approach that users of the STARK diagram in the experiment were seen to execute successfully. This would not be possible unless the representation has relatively uniform procedures.

The principles were successfully used to design the STARK diagram so that it possessed semantic transparency and syntactic plasticity. The design process initially focussed on the semantic transparency principles and satisfaction of the syntactic plasticity principles appeared to follow naturally.

6 Conventional representation

There are many commercial examination scheduling systems with sophisticated user interfaces. The principles can be used to analyse the merits of the representations behind those interfaces. However, for the purpose of the present study a "conventional" representation, which reflects the current state of the art of existing systems, was designed. It is informationally equivalent to STARK diagrams, largely diagrammatic, but does not satisfy the principles well, so as to give a good basis of comparison for STARK diagrams.

Figure 4 shows the conventional representation. It uses a tabular format that is common to traditional timetables, with time represented on both axes of the 2-D plane. Days are represented along the x-axis and for each day there are three periods shown by three rectangular areas down the y-axis. Within each rectangle rooms are shown by the dark blue icons on the left and the exams allocated to the rooms shown by the light yellow icons to the right.

The spatial and temporal information for each exam and room are shown by the numbers on the icons. The identity of the room is written above its icon along with the time that the room is available. On the room icons the room capacity is given and along with the current free space in parentheses (negative values denote how many students exceed the room capacity). For exams the number of students taking the exam and the duration of the exam are shown.

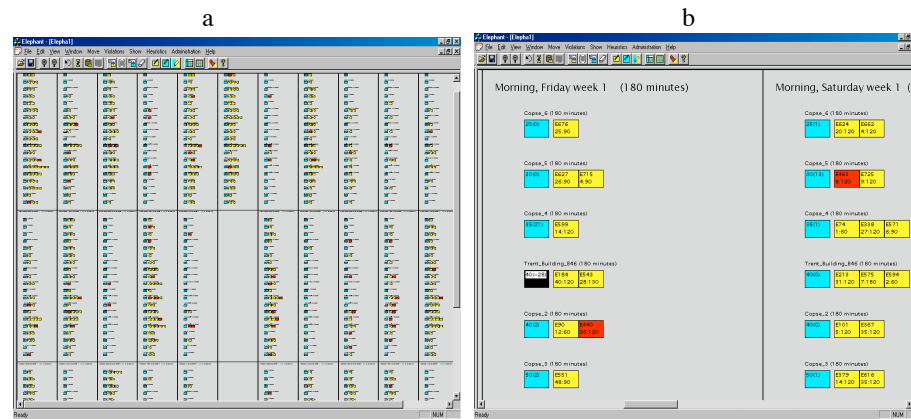


Fig. 4. Conventional schedule representation: (a) overview, (b) detailed view

Constraint violations are shown by changing the icon colour. Black room icons indicate some resource constraint is violated. The type and reason for the violation must be inferred from the values, or alternatively users can display the details of the violation as listed in a separate window. A red exam icon indicates an intersection constraint violation. To see if it is a clash or consecutive exam violation the user must inspect a list in a separate window that gives relevant details.

The conventional representation does a poor job of combining a globally homogenous with a locally heterogeneous representation of concepts. It does not use the principal visual dimensions in a coherent way to encode space, time, constraint and entity dimensions. The y-axis encodes both time (period) and spatial (room) information. The specific conceptual distinction between periods and days is made using principal spatial dimensions, but other specific temporal distinctions are associated with numbers attached to icons. The same is true for information about sizes of entities.

Levels of abstraction are not well integrated in the conventional representation. Set theoretic relations and arithmetic relations are not directly encoded in the representation. Judgements on these bases require the user to make mental inferences and computations rather than perceptual ones. In relation to the evaluation function, the distribution of the exam allocations is apparent but information about

spare capacity and the distribution of types of constraints requires deliberate search across separate lists.

The conventional representation does not integrate alternative perspectives well. Although resources and demands are shown together, detailed information about violations and intersection sets for particular exams are not available in the main diagrammatic window. Combining the information from different perspectives is a task left to the user. For example, in the simple scenario where there are two consecutive periods that do not have intersecting exams, the user must infer this by working through the lists for intersection constraints for all the exams in the periods.

The need to search for information across separate lists means the representation is not compact and tends to be rigid (not malleable). To make a meaningful move many separate operations are needed to find relevant information about intersection sets and suitable slots. This information bottleneck in turn makes it difficult to take a solution in one state and transform it into some better state. The conventional representation does not have particularly uniform operators, because there are a greater number of information lookup functions than in the STARK diagram.

The conventional representation and the STARK diagram have been considered in detail to show how the principles can be used for the design and analysis of representations. This also provides a good sense of the complexity of the task faced by the participants in the experiment.

7 Experiment

The problem used in the experiment was a real full scale examination scheduling problem for the University of Nottingham, which involved: 800 exams, 20 rooms and 32 periods (i.e., 640 slots), and 10113 instances of exam intersections. Two schedule solutions were generated by automated software of comparable power to commercial systems. One was used for a practice session and one was used for the test session.

Participants were six research students of the Automated optimisation and planning group (ASAP) at the University of Nottingham. Two were conducting research on automated approaches to examination timetabling and were allocated to different conditions (STARK vs Conventional representation). The others had little knowledge of the domain so were randomly allocated. Participants were paid for doing the experiment, which lasted between three and four hours.

The experiment involved an independent subjects design comprising of the two interface conditions. The training and test sessions were almost identical for both groups and involved three sections: (1) learning about the examination scheduling problem; (2) learning to use the scheduling software; (3) a practice session. In the 20 minute practice session participants attempted to improve a timetable solution by minimising the number of constraint violations in a simplified problem with approximately 50% of the exam intersections randomly removed. Before the test session began a brief screening interview established that subjects' level of understanding was sufficient to proceed. The test problem was a full Nottingham University data set without simplification and so was substantially harder than the practice problem. The test session lasted an hour and participants were told to improve the solution to the best of their ability. They were also told a simple weighting scheme for the

importance of different violated constraints and instructed to use the weightings to prioritise their work. The weightings were: 5 points for clashes and room capacity violations; 2 points for period excess violations; 1 point for consecutive exam violations. The software automatically logged all the actions performed by the subjects.

Table 1. Frequencies of operation types for each subject in the test session

Participants	STARK			CON			STARK	CON
	1	2	3	1	2	3	mean	mean
Show Intersections	110	171	58	15	59	70	113	48
Reallocate exam	67	105	41	8	28	52	71	29.3
Swap exam	0	1	0	0	0	0		
Exam to clipboard	0	0	0	0	20	0		
Allocate from clipboard	0	0	0	0	16	0		
Undo	24	17	13	1	33	36	18	23.3
TOTAL	294	201	112	158	24	156	202.3	112.6

Six types of operations performed by the subjects on the systems are of interest here: (1) showing the intersection set of an exam; (2) reallocating an exam; (3) swapping exam allocations; (4) placing an exam on the clipboard; (5) moving an exam from the clipboard to the schedule; (6) undoing operations 2 to 5. Table 1 shows their frequencies for all six subjects. There is considerable variability between subjects in the number of operations performed, but for the STARK group the mean was nearly twice that of the conventional group. The mean number of exam reallocations and show intersection operations were also substantially greater for the STARK group suggesting that in general these participants were more productive or at least more adventurous. Only one subject performed clipboard operations (Con 3) and one subject swapped a single pair of exams (STARK 1). On average the STARK group needed to do fewer undo operations relative to the number of reallocations made.




Fig. 5. Mean number of constraint violations over time for each group

The number of constraint violations present in the solution were recorded at the end of each minute. Figure 5 shows how the number of violations changed over time for the two groups as a whole. On average the maximum number of constraint violations resolved by STARK interface group was nearly three times greater than

Conventional interface group (STARK=32.4, Con=11.7). Figure 6 shows how individuals in each group performed. All the STARK interface subjects removed more violations than any of the conventional interface subjects. It is noteworthy that two of the STARK group succeeded in less than half the time to eliminate as many violations as the best Conventional group participants.




Fig. 6. Change in the number of constraint violations over time for each subject

8 Discussion

The pilot experiment reported here provides further evidence for the validity and utility of the representational design principles and suggests that the STARK diagram gives superior support for manual exam schedule improvement compared to the conventional interface. In general the STARK group was more productive than the Conventional group. Analyses are currently underway to reveal detailed differences in cognitive behaviour that result from the differences in the representations.

The success of the STARK diagram design and the outcomes of the experiment are notable because they extend previous studies addressing the principles in at least two ways. First, the domains previously studied were conceptually demanding topics in science and mathematics education that had substantial abstract components [8]. The examination scheduling problem is more concrete, with no high level laws referring to intangible properties. However, the design of a new representation using the principles still conferred a substantial advantage to its users compared to the conventional representation. Second, exam scheduling is a very information intensive domain compared to the previous domains studied. There are many more types of data and the sheer quantity of data is great. Nevertheless a new successful representation was designed using the principles. Together these points provide support for the claim that the principles are valid and have utility.

The benefits of satisfying the semantic transparency principles are directly apparent from the comparison of the structure of STARK diagrams and the

conventional representation. The syntactic plasticity principles must be considered in the light of the experimental outcomes. The substantially greater productivity and extent of constraint violations removed by the STARK diagram group, relative to the Conventional group, supports the hypothesis that the STARK diagram interface has more compact procedures comprising fewer operations. The basic analysis reported here does not provide direct evidence that the STARK diagram representation is more malleable than the conventional representation. To address this it will be necessary to examine how well the representations enhance the ability of users to sample information resulting in the selection of effective operations. Further analysis will attempt to assess how effective the different users' choices were in the context of the information available.

The development and evaluation of the principles may be viewed as an extension of previous work on the cognitive implications of the nature of the relation between represented domains and representing symbolic systems [2, 13, 17]. Such work has addressed relatively information lean and knowledge sparse domains compared with the richness and complexity of the real-world domains to which we are attempting to apply the principles considered here.

The application of the design principles to the information rich domain of scheduling has revealed some interesting new aspects about the design of effective representations. First, one important difference between STARK diagrams and the conventional representation is the major reduction in the number of discrete expressions required to represent the problem. This appears to be a consequence of the demands of semantic transparency principles that information be organised and integrated in a conceptually coherent manner. This confers obvious benefits to users in computer based tasks. In Figure 1 a portion of the full schedule is shown, but a user can view all the data for a full scale examination schedule on a standard computer monitor, which has a positive impact on memory, cognitive loads and the complexity of operations. Second, STARK diagrams appear to allow users to view examination schedules in a way that is abstracted from the domain. It is often unnecessary to think about the actual details of the domain itself. Problem solving can occur at the level of relations that underpin the domain by considering the containment and spatial arrangement of exam icons and room icons along with the constraint violation links that connect them. Such an approach is arguably less demanding and depends upon how well the representational structure reflects the conceptual structure of the problem. It is possible that this might allow schedule generation to be undertaken by individuals with less experience and knowledge of the domain. Third, the scheduling domain involves many different levels of granularity, ranging from the individual student through to complete schedules. Such levels are distinct from levels of abstraction and alternative perspectives, because they focus on concrete aspects of solutions and alternative levels of granularity can be identified under each perspective. STARK diagrams support the integration of these levels of granularity in a manner that does not occur with the conventional representation. In the previous instructional domains where the principles were applied, levels of granularity were not as important in the present domain. This suggests that an additional semantic transparency principle – integrate levels of granularity – should be proposed. Work on other information intensive scheduling domains in the current project will consider whether such a principle is necessary.

Acknowledgements

The research was supported by an ESRC/EPSRC research grant (L328253012) under the PACCIT programme and by the ESRC through the Centre for Research in Development, Instruction and Training CREDIT).

References

1. Anderson, M., Cheng, P C-H., & Haarslev, V. (eds.): Theory and Application of Diagrams: First International Conference, Diagrams 2000. Berlin: Springer (2000)
2. Barwise, J., & Etchemendy, J.: Heterogenous Logic. In: Glasgow, J., Narayanan, N. H., & B. Chandrasekaran, B., (eds.): Diagrammatic Reasoning: Cognitive and Computational Perspectives. AAAI Press Menlo Park, CA (1995) 211-234
3. Burke, E.K., Elliman, D., Ford, P., Weare, R.: Examination Timetabling in British Universities: A Survey, Practice and Theory of Automated Timetabling I, Vol. 1153. Springer LNCS (1996) 76-90
4. Carter, M. W., Laporte, G.: Recent Developments in Practical Examination Timetabling, Practice and Theory of Automated Timetabling I, Vol. 1153. Springer LNCS (1996) 3-21
5. Casner, S. M.: A task-analytic approach to the automated design of graphic presentations. ACM Trans. on Graphics, 10(2), (1991) 111-151
6. Cheng, P. C-H.: Law encoding diagrams for instructional systems. Journal of Artificial Intelligence in Education, 7(1), (1996) 33-74
7. Cheng, P. C-H.: Interactive law encoding diagrams for learning and instruction. Learning and Instruction, 9(4), (1999) 309-326
8. Cheng, P. C-H.: Unlocking conceptual learning in mathematics and science with effective representational systems. Computers in Education, 33(2-3), (1999) 109-130
9. Chi, M. T. H.: Conceptual change within and across ontological categories: examples from learning and discovery in science. In: Giere, R. N. (ed.): Cognitive models of science. Minneapolis: University of Minnesota Press (1992) 129-186
10. Glasgow, J., Narayanan, N. H., & Chandrasekaran, B. (eds.): Diagrammatic Reasoning: Cognitive and Computational Perspectives. Menlo Park, CA: AAAI Press (1995)
11. Kotovsky, K., Hayes, J. R., & Simon, H. A.: Why are some problems hard? Cognitive Psychology, 17, (1985) 248-294
12. Larkin, J. H., & Simon, H. A: Why a diagram is (sometimes) worth ten thousand words. Cognitive Science, 11, (1987) 65-99
13. Palmer, S. E.: Fundamental aspects of cognitive representation. In: Rosch, E., & B. B. Lloyd, B. B. (eds.): Cognition and Categrorization. Hillsdale, N.J.: Lawrence Erlbaum (1978) 259-303
14. Scaife, M., & Rogers, Y.: External cognition: how do graphical representations work? International Journal of Human-Computer Studies, 45, (1996) 185-213
15. Simon, H. A.: Models of Discovery and other topics in the methods of Science. Dordrecht: Reidel (1977)
16. Vincente, K. J.: Improving dynamic decision making in complex systems through ecological interface design: A research overview. Systems Dynamics Review, 12(4), (1996) 251-279
17. Zhang, J.: A representational analysis of relational information displays. International Journal of Human Computer Studies, 45, (1996) 59-74
18. Zhang, J.: The nature of external representations in problem solving. Cognitive Science, 21(2), (1997) 179-217