

Web-based Configuration Assistants

Giuseppe Attardi
Antonio Cisternino
Maria Simi

Dipartimento di Informatica, Università di Pisa, Italy
{attardi, cisterni, simi}@di.unipi.it

Abstract. Configuration assistants are tools for guiding the final user in simple configuration tasks, such as product assembling and customisation or study plans generation. For their wide availability, Web-based configuration assistants are valuable in fields such as electronic commerce and information services. We describe a general approach for building web-based configuration assistants: from a high level description of the configuration constraints and of the basic items, given in a declarative language, the hypertext files for user guidance and the Java code for constraint checking are generated. We claim that the general approach of *process oriented* configuration, where the user is guided through the configuration process by an explanatory hypertext, as opposed to *product oriented* configuration, where one starts from a high level description of the product of the configuration, is better suited for many application domains.

1 Introduction

In the terminology of expert systems, configuration systems are a subclass of design systems, whose task is to assemble a product out of a set of predefined parts according to problem specific constraints [Hayes-Roth, Waterman & Lenat 1983]. Examples of configuration problems are computer equipment configuration [McDermott J. 1982, Barker, O'Connor 1989]), software configuration, timetables generation and scheduling, and product configuration in different domains [Tiihonen et al. 1996].

The complexity of a configuration task arises from having to cope with many and interacting design decisions, whose consequences cannot readily be assessed, and with constraints of different nature. Configuration problems are therefore a challenging domain for expert system technology.

For simpler configuration tasks we can envision interactive configuration assistants which guide the user step by step through design decisions and exploit their knowledge of domain constraints and of constraints deriving from previous choices. A *configuration assistant* (CA) is an interactive tool for configuring systems, which does not make choices on user's behalf but assists the user in assembling a consistent product. Ideal targets for this kind of system are applications where the user wants to retain control on the configuration choices, yet he must be made aware of the available alternatives and supplied with necessary information to perform the right choices. A configuration assistant typically addresses sales persons or end-users rather than experts in the configuration domain.

The applicability of this kind of systems is restricted by some simplifying assumptions, reflected in the configuration model we adopt. Nevertheless the range of configuration problems that can be dealt in this way is significant. For example, configuration tasks that are amenable to this simplified vision are study plans compilation or the assembling of a product from a catalogue of components, like a complex piece of modular furniture (kitchen furniture for instance), or a personal computer.

Configuration assistants of this kind running on the Web open opportunities in many fields, including electronic commerce, for the wide availability and the possibility of remote access and use (see <http://www.volvocars.com> for an example). The system does not need to be installed on the user computer in order to be used and portability problems need not to be addressed. This is also an application domain where a Java solution [Gosling 1996] has clear advantages over a server-based solution with constraint checking done on the server side. All the job of the configuration assistant can be done locally, at the client's side, by downloading the necessary Java code. Communication with the server can be reduced to tasks such as user validation, statistics gathering or archiving.

The major problem with product configuration systems is to maintain product information up-to-date. The cost of long term management and maintenance of product knowledge, as product models and configuration constraints evolve, may seriously impair the advantages of configuration systems. For this reason *single-use* configuration systems, which do not clearly separate configuration knowledge from the configuration programme are doomed to failure. A *general-use* configuration model, which facilitates the development of configuration applications in different domains, and for different product classes, is the only viable alternative [Männistö, T. et al. 1996, Tiuhonen, Soininen, 1997].

In this paper we describe our general configuration model, characterised as *process oriented*, in contrast to a *product-oriented* approach. In fact the aim is guiding the user step by step through the configuration process rather than starting from a high level description of the product to be configured.

Most configuration models in the literature are *product-oriented*: their task is described as the synthesis, more or less automatic, of a set of objects satisfying a set of predefined constraints [Stumper 1997]. In [Klein, Buchheit & Nutt 1994] the configuration task is described as model construction, the constraints being expressed as a set of logic formulas defining a theory; any model of the theory is a legitimate configuration. An alternative approach is described in [Faltings & Weigel 1994], where configuration is described as a constraint satisfaction problem.

Our approach differs from these because, given the interactive nature of a configuration assistant, and the fact that the user performs all the choices, the configuration process does not do any search in a solution space; the CA it is strictly deterministic and thus very efficient.

Our model has been used as the basis for the implementation of a generic tool for developing domain specific configuration assistants running on the Web. A configuration application is generated starting from a high level description of the basic components and the constraints expressed in a declarative form. The hypertext files for user guidance and the Java code for constraints checking are automatically generated from this high level description. The reduced cost of delivering a specific configuration assistant and the efficiency and compactness of the generated code are clear advantages of this approach.

The tool has been used to generate specific configuration assistants in the domain of study plans compilation and submission (the Compass application). Compass has been used to generate study plans by students of the Faculty of Letters and Philosophy, and of the Faculty of Sciences at the University of Pisa.

2 The configuration model

2.1 Process oriented configuration

The standard way to think of a configured product is as an assembly of basic components matching a set of structural and functional requirements. As an alternative approach, we focus on the process of posing a series of relevant questions, with the goal of completely identifying a correctly configured object fulfilling user's requirements.

The metaphor we have in mind is an experienced salesperson posing a series of questions to the customer, according to some standard pattern. According to the answers received, the vendor gains an increasingly accurate idea of the customer's needs, until he is able to propose a specific product. If the product is complex and needs to be assembled from a catalogue of components, the vendor uses his knowledge of the configuration constraints to guide further choices among the available options.

We can identify two sorts of questions in this dialogue: general questions concerning the strategy in the selling process, functionality of the product, or user's needs, preferences and profile; specific questions whose purpose is to make the user perform selections from a catalogue. Both of them, together with configuration constraints, contribute to the resulting configured object.

The interaction going on between vendor and customer can be seen as the analogue of the configuration process.

The *process oriented* configuration model relies on a *choice graph* (CG), which is a directed acyclic rooted graph. The *choice graph* can be seen as a representation of the vendor strategy in posing questions to the customer. Each node of the graph represents a possible answer to a general question proposed by the vendor. Configuration constraints are associated to each node and are used to suggest possible selections from the catalogue.

The *configuration process* starts in the root of the graph, with an empty configuration. Following an edge of the graph corresponds to choosing a possible user alternative in the configuration process, and moving to a successive choice node. Within a choice node, in order to fulfil a configuration constraint the user may be asked to select among a set of possible items. Traversing the graph, the user builds a configuration, accumulating choices and selecting items within choice nodes.

A *configuration domain* is fully defined by the set of basic components, *the item set*, and *configuration constraints* specific to a configuration application. The basic items we consider are structured objects with attributes.

A *valid configuration* is a set of basic items, and a set of user choices, matching a given set of constraints.

The model provides two different notions of valid configuration, *partial* and *final*. A partial configuration can only be checked for correctness but not for completeness; a final configuration must be correct and complete, for example it must include all the required items. Correspondingly, a system implementing this model can check partial configurations during the configuration process, warning the user as soon as some constraint is violated. At the end of the process the system checks the final configuration. –

The choice graph is a representation of the configuration knowledge, which is very convenient for generating a graphic interface to the CA. A possibility is a hypertextual interface where a node of the graph is displayed as a single hypertext unit, with selection buttons for available options which the user is prompted to select, and links to other pages, corresponding to successor nodes.

The possibility of automatic generation of the CA, including its user interface and documentation, has been successfully exploited in the CompAss system. This feature of the model is especially important for dynamic configuration domains, where the catalogue of components or the constraints change frequently.

2.2 Formalisation of the configuration model

In the following we will introduce formal definitions of choice graph, configuration and configuration process.

Let I be the set of basic items of the configuration domain.

A choice graph can be formally defined as

$$G = \langle S, N, \mathfrak{F}, \mathfrak{N} \rangle$$

with S is the set of nodes in the graph (choice nodes); N is a binary relation among the nodes in the graph denoting the edges of the graph. \mathfrak{F} and \mathfrak{N} correspond to two different classes of *constraint functions* associated to the nodes of the graph: built-in functions of general use (called *composition functions*) and domain specific functions (called *custom functions*). Constraint functions are predicates used to check that the items selected in a node fulfil configuration constraints. We will denote with \mathfrak{F}_j and \mathfrak{N}_j the functions associated to node s_j .

Composition predicates take as input a set J of items, and return the value *true* or *false*. They check a configuration with respect to constraints local to a choice node; for example, they are used for checking whether a required item has been included.

Custom functions take as input a set J of items, and check global properties of the set of items selected so far; for example the number of items in a configuration or the need to avoid duplicate items. These functions are usually domain dependent. The result, in addition to *true* and *false*, may also be a *warning*.

Constraint functions may be invoked with parameter *partial* or *final* for incremental or final constraint checking. The validation component of the CA may thus operate according to two different modes: in incremental mode, it gives immediate feedback on wrong user actions during the configuration process; in final mode, it provides for final validation of all the constraints.

A configuration $K = [C_1, C_2, \dots, C_n]$, is a sequence of choice components. A choice component C_i is a pair $\langle s_i, I_i \rangle$, where $s_i \in S$ is a choice node of the graph and I_i is the set of items, possibly with repetitions, selected in node s_i during the configuration process.

A configuration is incrementally built as a result of user or system actions of different types during the configuration process.

Let us formally define *validity* of a partial and final configuration.

Definition 1 (*Validity of partial configuration*)

A partial configuration $K = [C_1, C_2, \dots, C_n]$ is valid if the following properties are satisfied:

1. The sequence $[s_1, s_2, \dots, s_n]$ is a path on the choice graph, starting from the root node.
2. Let $\mathfrak{F}_j = \{f_1, \dots, f_k\}$ the set of composition constraint functions associated to node s_j . For each $j = 1, \dots, n$, it must exist a partition $P_j = \{p_1, \dots, p_k\}$ of the set of items I_j such that $f_h(p_h, \text{partial}) = t$, for each $h = 1, \dots, k$, i.e. any f_h must be satisfied by a set of items in the partition.
3. Let $H_j = I_j \cup \dots \cup I_n$ i.e. the set of items in K selected in node s_j and its successor nodes. For each $j = 1, \dots, n$, and each g in \mathfrak{S}_j it must be the case that $g(H_j, \text{partial}) = t$ or $g(H_j, \text{partial}) = w$.

Definition 2 (*Validity of final configuration*)

A final configuration K is valid if the following properties are satisfied (with \mathfrak{F}_j, H_j , defined as before):

1. The sequence $K = [C_1, C_2, \dots, C_n]$ forms a complete path on the choice graph, starting from the root node and ending in a leaf node.
2. For each j , it must exist a partition $P_j = \{p_1, \dots, p_k\}$ of the set of items I_j such that $f_h(p_h, \text{final}) = t$, for $h = 1, \dots, k$.
3. For each j and each g in \mathfrak{S}_j it must be the case that $g(H_j, \text{final}) = t$

Three conditions are thus required to be met: (1) selected choice nodes must lie on a path on the choice graph, originating in the root node; (2) for each node, composition functions are checked on the items selected in the node; (3) for each node, custom functions are checked on the items selected in the node itself and its successors.

Note that composition predicates in \mathfrak{F}_j are not checked independently on the set of items I_j : an item used to satisfy a constraint cannot be used to satisfy another constraint; this is the reason why we look for a partition of the items I_j satisfying all composition functions.

3 The configuration language

To build a configuration application we need a language for describing the various aspects of the configuration domain: the structure of the items, the item database, the custom constraint functions, and the *choice graph*, with associated configuration constraints.

3.1 The components database

The *item structure* consists in an item definition (similar to a `struct` of the language C). The items themselves are described in the *item database*, according to the defined item structure. Some of the attributes may be dynamically filled by the user during the configuration process.

3.2 The constraints language

The choice graph, with associated configuration constraints, is defined in a special declarative language designed for this purpose.

The language supports the definition of two types of blocks: *list blocks* and *choice blocks*.

A list block is simply a way to define a group of items so that it can be referred by name. The operator ‘#’ is the way to refer to all the items belonging to a defined *list block*; for example #ref refers to a block named ref, which could be defined as

```
ref [item 1, item 2, ..., item n]
```

An item in a list block can in turn be a reference to another list block.

A choice block corresponds to a node in the choice graph and defines composition constraints for a node. A choice block can have for example the following structure:

```
Personal computer {
  [Motherboard, Case],                               /* 1 */
  [CPU 200Mz, CPU 333Mz] (1),                         /* 2 */
  [8Mb RAM] (2+),                                     /* 3 */
  OR(Game PC, Multimedia PC)=>                       /* 4 */
  [CD-ROM, Sound card, Speakers, 3D Video Card],
  [Power UPS] (1-),                                   /* 5 */
  CHOICE(CRT monitor, LCD monitor)                   /* 6 */
} Max_cost(3000);                                     /* 7 */
```

A block description includes a list of *constraints*, specifying items that are necessarily needed for the block and items that are needed depending on the configuration state, i.e. the presence of other items or previous choices. For example line 1 says that the items Motherboard and Case are always required for a personal computer. Line 4 says that items on the right of the ‘=>’ operator are to be included in the configuration only if the user has declared to be interested in a PC for playing computer games or in a Multimedia PC. The condition for inclusion are expressed by means of logical operators: AND (all of them), OR (some of them), NOT (none of them).

Line 2 is an example of a construct that prescribes the selection of a number of items out of a list of items; in particular the example says that exactly one item out must be selected from the given list of available CPU’s. Lines 3 and 5 are similar with different number restrictions: in the first case *two or more* items are required, in the second case *at most one* item is accepted. Line 6 defines the successors of the current node in the choice graph, i.e. the available choices at this level.

Thus a choice block defines the items required for the node in the configuration. The type of standard controls that are generated concern the admissibility of items (i.e. answers the question “is it correct that this item is in the current configuration”?) or the presence of items (i.e. answers the question “is this required item present in the current configuration”?). These constraints are translated in a set of built-in composition constraint functions, as detailed below.

Other kinds of constraints which are often needed, such as “The cost of the configuration must be at most XXX\$”, are implemented by *custom constraints functions*, which are typically application dependent: these can be defined by the user or supplied by a library. Custom constraint functions appear at the end of a choice block (like in line 7 above) and, as specified above, apply to all the items in the block and successor blocks.

The first choice block, appearing in the constraint file, corresponds to the root of the choice graph.

Each constraint in a choice block has the following structure¹:

$$\langle \text{item list} \rangle [(\langle n \rangle [+ | -])]]$$

that is an *item list*, optionally followed by a numeric restriction such as (n) , $(n+)$, or $(n-)$. A *item list* may be conditioned or unconditioned, and has the following form:

$$[\text{AND} | \text{OR} | \text{NOT} (\langle \text{item list} \rangle \Rightarrow)] [\langle \text{item list} \rangle]$$

The antecedent part, when present, expresses a condition under which the *item list* in the consequent is returned. If the condition is false the empty list is returned; if the condition is missing the item list is always returned. The items in an item list may include single items, item lists (possibly referred to by name) and also conditional item lists. The item list in the antecedent may also be a list of names of choice nodes.

The evaluation of a Boolean condition applied to an item list proceeds as follows:

1. the item list is evaluated with respect to the current configuration and choice block; this results in a list of items which may be different in different configuration states;
2. the Boolean condition is evaluated on the resulting list according to the standard meaning of the Boolean operators.

The constraint function generated is one of a set of built-in control functions: *Nec* (for simple item lists without number restrictions, meaning that all the items are necessary), *Atleast*, *Atmost*, *Exactly* (for numerically restricted selection from the item list).

~~Here are the details of the evaluation of an item list. We use the function $S(K, s, \text{list})$ for list evaluation with respect to a configuration K and choice block s ; $B(K, s, \text{cond})$ for the evaluation of a condition with respect to a configuration K and choice block c .~~

~~$$S(K, s, \text{condition} \Rightarrow \text{item_list}) = \text{If } B(K, s, \text{condition}) \text{ then } S(K, s, \text{item_list}) \text{ else } []$$~~

~~$$S(K, s, [i_1, i_2, \dots, i_k]) = S(K, s, i_1) \circ S(K, s, i_2) \circ \dots \circ S(K, s, i_k)$$~~

~~$$S(K, s, i) = i, \text{ when } i \text{ is an item identifier}$$~~

~~$$S(K, s, \# \text{list_block_name}) = S(K, s, l), \text{ where } l \text{ is the item list referred by } \text{list_block_name}$$~~

~~Here is how conditions are evaluated:~~

~~$$B(K, s, \text{AND}(i_1, i_2, \dots, i_k)) = t \text{ if all the items in } S(K, s, [i_1, i_2, \dots, i_k]) \text{ are among the items in the}$$~~

~~$$\text{configuration up to the previous choice node, i.e. they are in } I_{s1} \cup I_{s2} \cup \dots \cup I_{sj}, \text{ given that } A = [s1, \dots, sj, s, \dots, sn];$$~~

~~$$f \text{ otherwise.}$$~~

~~$$B(K, s, \text{OR}(i_1, i_2, \dots, i_k)) = t \text{ if some item in } S(K, s, [i_1, i_2, \dots, i_k]) \text{ is in } I_{s1} \cup I_{s2} \cup \dots \cup I_{sj};$$~~

~~$$f \text{ otherwise.}$$~~

~~$$B(K, s, \text{NOT}(i_1, i_2, \dots, i_k)) = t \text{ if no item in } S(K, s, [i_1, i_2, \dots, i_k]) \text{ is in } I_{s1} \cup I_{s2} \cup \dots \cup I_{sj};$$~~

~~$$f \text{ otherwise.}$$~~

¹ The terminal symbols of the language are in bold. Square brackets are used for optional syntactic constructs. The '|' is used for alternatives.

$B(K, s, \text{AND}(s_1, s_2, \dots, s_k)) = t$ if all the choices in $[s_1, s_2, \dots, s_k]$ are in A ; f otherwise.

$B(K, s, \text{OR}(s_1, s_2, \dots, s_k)) = t$ if some choice in $[s_1, s_2, \dots, s_k]$ is in A ; f otherwise.

4 Generation of a Configuration Application

A configuration application is generated by using a compiler, which takes as input the following data files: the *item structure* file, the *item data* file, the *custom functions* file, and the *constraint file*, which defines the choice graph. All these specifications are given in input, as separate text files in human readable form, to the CA generator.

The compiler is divided in two modules: C1 and C2 (Figure 1). The first module is needed for generating the Java code for the applet and the items database. The second module of the compiler generates a compact binary representation of the constraints and the HTML files for documentation and user guidance.

The HTML files generated are to be used in user interface of the configuration assistant. In particular a set of HTML skeleton files are generated out of the choice graph: for each node in the graph a file is generated with selection icons for the items in the node and hypertextual links to the items descriptions. The file also contains a few lines of text, which synthetically describe the node constraints (for example “Choose at least three items out of the following:”), which can be enriched with additional text deemed useful to guide the user during the configuration process. In addition the file contains choice icons and hyperlinks to other HTML files in correspondence of available choices.

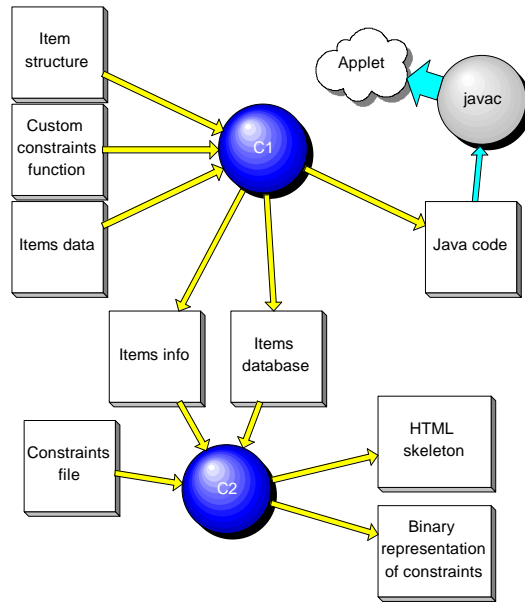


Figure 1. Generation of a Configuration Application

5 Communication with the Server

The system uses a custom server to store configurations in a database. The server, written in Java, uses the JDBC interface to interact with the database.

The configuration assistant runs on the client machine after the downloading of the CA applet from the http server. When the user has completed the configuration process and the final configuration has been validated, the CA applet submits the configuration to the server; the server accepts the connection, generates a unique identifier for the configuration, and generates a HTML page that contains the configuration and its unique identifier. The client receives back this HTML page as result of the configuration process and can decide to print it locally or simply take note of the identifier for future use.

At the moment we provide this simple mechanism of authentication of the configurations generated by our CA. To check that a configuration is really generated by our system it is sufficient to find a configuration stored in the server database with a given identifier.

6 CompAss: a Configuration Assistant for Plans of Study Compilation

CompAss (COMPilazione ASSistita di piani di studio) is a system to assist students in the task of producing a study plan. CompAss and its associated support tools have been developed in the context of a pilot project for the Faculty of Letters and Philosophy of the University of Pisa and have been further developed and extended for use in other Courses of Study within the university.

Plans of study approval is a time consuming job for all the courses of study in the university, due to the high number of submissions each year (around 4000 for the Faculty of Letters and Philosophy) and the high rate of incorrect submissions. One of the requirements was that students could use any computer located in the various departments of the faculty to compile plans of study; data had to be collected in one single place for archival. The Java solution was the obvious choice and offers additional advantages such as the possibility of using the system from home.

The Web page of the CompAss configuration assistant is vertically divided in two parts (Figure 2). The right part contains the *navigation frame*. The left part contains the *configuration frame* and an *application specific tool bar*.

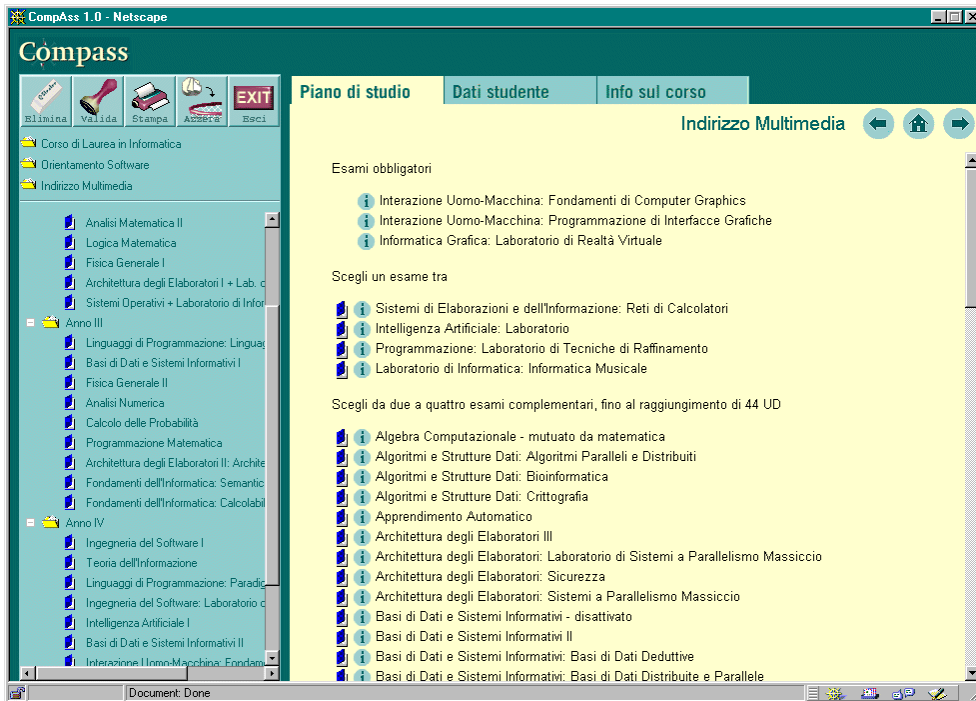


Figure 2. The User Interface of Compass

The navigation frame is a HTML frame displaying a normal hypertextual document; it displays the available choices together with any informative text deemed useful to guide the user to do the right choices during the configuration process. Documentation on the items can be obtained by clicking on the “i” round icon.

Special icons associated to choice points and to items are used to perform configuration actions: intermediate choices or item selections; when these icons are selected they send messages to the configuration program, the *validator*, which is a Java Applet associated to the configuration frame.

The configuration frame on the left contains the Java applet managing the configuration. The applet receives input by direct interaction in its client area or by selection of special icons in other frames. Whenever a configuration action is performed the applet reacts by checking the current partial configuration, accepting the change or prompting the user if any configuration constraints is violated.

Tool icons in the toolbar denote general utility or application specific actions available to operate on the partial configuration displayed in the configuration frame (i.e. item deletion, final configuration validation, abortion of the configuration process, printing or submission of the final configuration).

Interaction between HTML pages and the Java applet is implemented by using JavaScript to post the events of icon selection to the configuration applet. With this solution the interface

can exploit all the power of the HTML language and standard browsing capabilities, while still allowing user interaction with the Java program.

In this configuration application the basic items are the courses offered by the faculty; the constraints file implements the rules for plan of study formation; it includes a choice graph where nodes correspond to choices such as the course of study, the orientation, the field of specialisation and so on, together with the necessary constraints. A configuration is a legitimate plan of study, i.e. a list of courses which a student plans to take, fulfilling all the requirements imposed by the faculty.

CompAss saves a lot of work for secretaries who previously had to type in the plans from the paper forms submitted by students and eliminates the routine work of the faculty committees which have to verify and approve the plans.

The plan of study manager running on the server accepts communications from several CompAss clients, receives data from plans of study, generates HTML pages, stores data in a database, and gathers statistics on the number of users and on the pattern of use of the system. An instance of CompAss can be seen at the Web address <http://compass.di.unipi.it>.

7 Conclusions and future work

We have described a model for process oriented configuration and a general tool for generating configuration assistants for the Web; the strategy has been successfully exploited in the specific configuration domain of study plans compilation.

We believe that other configuration applications are amenable to this simple paradigm. Future work will be to experiment in different domains in order to define exactly the range of applications that are worth to be tackled with this approach and to come out with a general enough configuration language.

We also plan to enhance the configuration language by introducing specialisations of the basic item data structure and thus the possibility to define different kinds of items through inheritance. We also plan to introduce a support for 3D visualisation of configured products, using as a basis the current tree structured display of the configuration. We are developing a graphic editor for the choice graph, to give to the configuration expert a higher level tool to define configuration constraints.

Finally, we want to add provisions for security, thus develop new models for submission of the configurations and user authentication, to make the system suitable for electronic commerce.

Acknowledgements

Many people have contributed in various ways to the development of the Compass application for plans of study and to the Compass project in general. We wish to thank in particular the project manager Vincenzo Macrì, and the many people who contributed the expert knowledge about study plans. Theo van Boxel for the graphics design and Andrea Riboldi for the icons of the graphical interface.

Comment [MS1]: Ci vuole un po' di bibliografia su Java.

References

- [Hayes-Roth, Waterman & Lenat 1983] Hayes-Roth, F., Waterman, D.A. & Lenat, D.B. Eds., Building Expert Systems, Addison Wesley Publishing Company, Reading, Ma.
- [McDermott J. 1982] McDermott J., R1: a rule-based configurer of computer systems. In *Artificial Intelligence*, vol 19, no 1, 39-88.
- [Barker & O' Connor 1989] Barker, V.E. & O'Connor, D.E. Expert Systems for Configuration at Digital: XCON and beyond. *Communications of the ACM*, 32(3):298-318.
- [Tiihonen et al. 1996] Tiihonen, J., Soininen, T., Männistö, T., Sulonen, R., State of the practice in product configuration – a survey of 10 cases in the Finnish industry. In *Knowledge Intensive CAD, First Edition*, T. Tomiyama, M. Mäntylä, S. Finger, editors., Chapman & Hall.
- [Gosling 1996] Gosling, A.J. *The Java Programming Language*. Addison Wesley Publishing Co., Reading, Ma.
- [Männistö, T. et al. 1996] Männistö, T., Peltonen, H., Sulonen, R., View to Product Configuration Knowledge Modelling and Evolution, in *Configuration Papers from the AAAI Fall Symposium*, Faltings and Freuder Eds., AAAI Press, 111-118.
- [Tiihonen, Soininen, 1997] Tiihonen, J., Soininen, T., Product Configurators – Information System Support for Configurable Products, Technical Report TKO-B137, Helsinki University of Technology, Laboratory of Information Processing Science. Also published in *Increasing Sales Productivity through the Use of Information Technology during the Sales Visit, A Survey of the European Market*, Hewson Consulting Group.
- [Stumper 1997] Stumper, M., An Overview of Knowledge-based Configuration, *AI Communications* 10, 111-125.
- [Klein, Buchheit & Nutt 1994] Klein, R., Buchheit, M., Nutt, W., Configuration as model construction: the constructive problem solving approach, in *Proc. Of the 3rd International Conference on Artificial Intelligence in Design*, AID '94.
- [Faltings & Weigel 1994] Faltings, B., Weigel, R. Constraint-based knowledge representation for configuration systems, Technical Report No. TR-94./59, Département d'Informatique, Laboratoire d'Intelligence Artificielle, EPFL, Lausanne, Switzerland.