Collaborative Environmental Planning with GeoMed

Nikos Karacapilidis, Dimitris Papadias, Thomas Gordon and Hans Voss

GMD FIT, Artificial Intelligence Research Division
German National Research Center for Information Technology
Schloss Birlinghoven, D-53754 Sankt Augustin, Germany
e-mail: {karacapilidis, papadias, gordon, hvoss}@gmd.de

Abstract: Environmental planning usually involves a large number of decision makers with different backgrounds and interests. Appropriate decision making procedures are needed in order to jointly consider their individual approaches and achieve collaboration. This paper discusses issues involved in Collaborative Environmental Planning and reports work on GeoMed, a group decision support system for Geographical Mediation on the World Wide Web. We first present an argumentation framework that constitutes the core of the system and supports fair, rational and efficient decision making. In the sequel, we discuss the necessary enhancements in order to deal with spatial applications, and environmental planning in particular. Finally, we give a description of GeoMed and an overview of the services provided.

Keywords: Environment; Group Decision Making; Argumentation; Computer-Supported Cooperative Work; Geographical Information Systems.

1 Introduction

During the last five years, implementation of Information Systems has witnessed a movement towards data types of increasing complexity. The simple information expressed in numbers and character strings, while still important, has been joined by large numbers of multimedia "documents" and complex data forms (Kemp, 1995). This is widely facilitated by recent advances in computing (i.e., low-cost, high-speed hardware components) and communication (i.e., proliferation of World-Wide Web and Internet services). At the same time, the *information highway* becomes a two-way street; users should not only have improved access to existing information, but also the appropriate means to contribute their own. The distinction whether a user is an expert or not becomes more and more unclear. For instance, in a city planning procedure where a variety of organizations and individual users is involved, each participant (including common citizens) is eventually expert in his respective field or in expressing his own interests. As a result of the above, the core task of an Information System for collaborative work nowadays is to simultaneously support users, not only with the mere retrieval facilities, but also, with capabilities of adding multimedia information (Densham et al., 1995).

Collaborative Environmental Planning (CEP) is an area where the above observations should be taken into account (Daniel et al., 1996). It requires an interactive, real-time environment for various groups of agents (thereafter, the terms *decision maker* and *agent* will be used interchangeably), usually with different interests and interpretations on a environmental planning problem. The well-tried interaction of decision makers with spatial analysis tools has to be enhanced with (Karacapilidis et al., 1995):

- Methods for a fair, rational and efficient handling of debates in a group decision-making environment: The aim is to provide assistance to the parties involved with an argumentation framework that structures the discussion in an hierarchical way, while managing the dependencies between important elements of the underlying argumentation (Karacapilidis, 1996).
- Tools providing the appropriate group-based user interfaces: Agents, as often happens in real life with people who are responsible for decision making (e.g. politicians, managers), are not necessarily expert users of Information Systems. They probably use linguistic terms to express their opinions and they need intuitive, easy-to-use tools that facilitate the decision making process without having to interpret large amounts of spatial data themselves.
- Methods to elicit and capture data from related Geographical Information Systems (GISs), Data and Knowledge Bases: Each form of multimedia data type requires implementation of the appropriate data structure and access methods. It would be impossible to cope with the explosion of multimedia information, unless organized in efficient ways. In the case of Spatial Database Management Systems, for instance, a number of alternative data structures has been proposed for the manipulation of spatial data (Papadias et al., 1995).

The above features are included in GeoMed, a EU-funded project for Geographical Mediation on the World Wide Web (WWW). CEP is intended to be one of the system's basic application areas. The rest of the paper proceeds as follows: Section 2 illustrates the Argumentation Framework. The basic elements are discussed and a comprehensive example is given in order to describe the system's features. Section 3 presents the enhancements needed to deal with spatial applications and environmental planning in particular. Several issues such as semantic models, interfaces, interoperability and spatial inference mechanisms are discussed in detail. Section 4 reports on GeoMed, addressing the types of services to be integrated with the WWW platform. Finally, Section 5 concludes the paper discussing the importance of the users' participation in a CEP procedure.

2 The Argumentation Framework

Collaborative environmental planning has to be performed through debates, negotiations, and argumentation among various agents (Jelassi and Foroughi, 1989). Predominant conditions in real world planning paradigms are (Karacapilidis and Gordon, 1995):

- Planning agents assert arguments supporting or against alternative solutions. Conflicts are rather unavoidable and a rational planning framework is required.
- Reasoning is defeasible. Further information may cause another alternative to be more preferable than what seems optimal at the moment.
- Factual knowledge is not always sufficient. Usually, value judgments on aspects like weighing of attributes that alternative solutions pose, appear as the most critical issues. In addition, these judgments should be derivable and subject to debate.
- Agents have to manage the existence of both not enough and too much information, as well as of limited resources for finding a solution.

The ubiquitous task in all planning problems is the identification and selection among alternative courses of action. CEP systems require adequate tools for supporting and reasoning about rational and effective decision-making, even in the presence of ill-structured information and conflicts of interest. The Argumentation Framework should act as an assistant and advisor, by recommending solutions and leaving the final enforcement of decisions and actions to the agents.

2.1 Elements

Our model is a formal variant of Rittel's informal *Issue-Based Information System (IBIS)* model of argumentation (Kunz and Rittel, 1970; Rittel and Webber, 1973). The discussion forms a hierarchical tree which consists of an issue (at the top), positions and constraints. An *issue* corresponds to decisions to be made or goals to be achieved. It contains a set of alternative options, the goal being to choose the most appropriate one(s) (e.g., Issue: "choose the most appropriate airport site given three alternatives").

Positions are the basic elements in our framework; they may contain data that have been brought up to declare alternative options, justify a claim, advocate the selection of a specific course of action, or avert one's interest from it (e.g., position P1: "choose site 1"; position P1_1: "site 1 is public"; position P1_2: "site 1 is environmentally sensitive"). An *argument* links together two positions of the debate. The framework allows for *supporting arguments* (e.g., P1_1 is a supporting argument to P1) or *counter-arguments* (e.g., P1_2 is a counter-argument to P1). Any combination of arguments can be applied.

Constraints represent preference relations between two positions. The relations involved can be of the type: "more important than" and "equally important to" denoted by the symbols ">" and "=", respectively. For instance: "P1_2 (site 1 being environmentally sensitive) is *more important* than P1_1 (site 1 being public land)".

The sub-tree below each position or constraint affects its activation status. A position/constraint with no argumentation underneath (e.g., a recently inserted position) is by default *active*. When there has been some discussion about it, its activation status is determined according to the *burden of proof (BoP)* used (the concept is motivated by law procedures). The following burden of proofs have been implemented (other BoPs, that match specific application needs, can be easily incorporated to the system):

- *Scintilla of Evidence* (SoE): According to this BoP a position/constraint is active, if there is at least one active position supporting it (to be more specific, there is a supporting argument that links it with an active position).
- *Beyond Reasonable Doubt* (BRD): A position/constraint is active if there are no active positions against it (once again, if there are no counter-arguments at all, or all the counter-arguments link it with inactive positions). Even a single counter-argument, when linked with an active position, suffices to inactivate it (independently of the number of supporting arguments).
- *Preponderance of Evidence* (PoE): Unlike the previous cases (where only the arguments pro and con are taken into account), constraints are used in this proof standard in order to determine the activation status. A *scoring mechanism* (based on *topological sort*) calculates the weights of the positions involved. In the current implementation, weights fall within the range [0 .. 10], where 0 (10) denotes minimum (maximum) importance for a position. In the absence of constraints, each position is given the average value of 5. The scoring mechanism refines the weights when a new constraint is inserted. Every time a position is more important than another one (either in an implicit or explicit constraint), its weight is increased (and

vice versa). A position/constraint with PoE as burden of proof becomes active if the sum of weights of the supporting arguments is larger than that of the counter-arguments (to be specific, the sum of weights of the active positions that are linked with supporting arguments is larger than that of the active positions that are linked with counter-arguments).

PoE involves some topics of special interest, the most important being constraint consistency checking. Whenever a new constraint is added, a path consistency algorithm (Macworth and Freuder, 1985) determines if it contradicts with the other active (implicit or explicit) constraints. If this is the case, the new constraint is labeled *inconsistent* and not taken into account by the scoring mechanism. The constraint will become consistent only if some of the conflicting constraints get inactivated. In general, argumentation in our framework is a complicated process where a change in some part of the tree may have consequences in other parts. Note that path consistency (which is polynomial) suffices here because constraints involve only the relations ">", "<" and "=". If arbitrary disjunctions were allowed (e.g., "≠", "≥", etc.) exponential algorithms would be required for satisfiability.

2.2 An argumentation example

This subsection demonstrates the features of the argumentation framework by elaborating the "airport site" example mentioned above. Figure 1 illustrates an instance of the discussion. The initial issue is at the top, the alternative positions at the second level, whereas the pro (+) and con (-) positions for each alternative are at the bottom. For the shake of the example, we assume that the BoP of P1 and P3 (and their sub-trees) is SoE, while the BoP of P2 is BRD. All leaf positions are (by default) active. Positions P1 and P3 are active because they have at least one active supporting argument, while P2 is inactive since it has one active counter-argument.

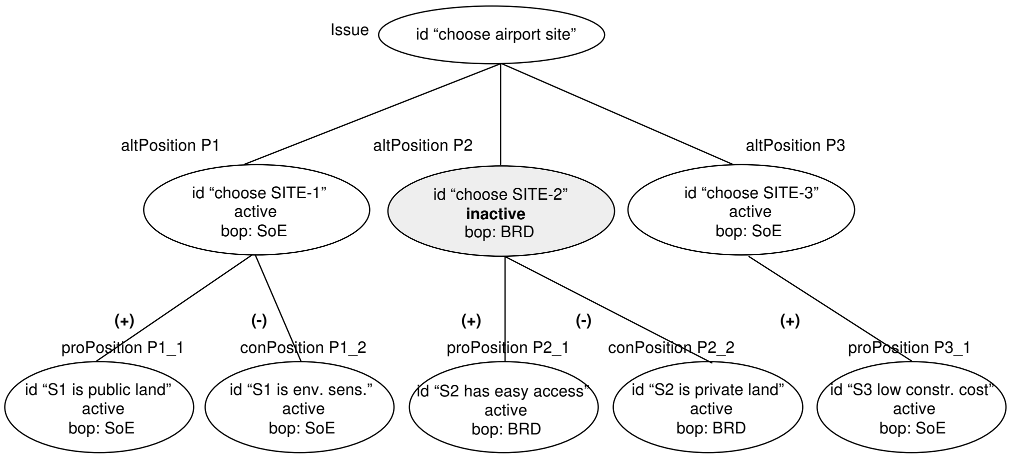

Figure 1.

Activation is a recursive procedure; the change of the activation status of a position propagates upwards until the root of the tree. Figure 2(a) illustrates the insertion of position P2_2_1 as a supporting argument to P2_2. At this stage the activation status of P2_2 does not change. Therefore, there is no need to propagate the update. However, as shown in Figure 2(b), when P2_2_2 is inserted, P2_2 becomes inactive (BRD and an active counter-argument). This change of activation status is propagated to P2, which now becomes active (since it does not have any active counter-arguments).

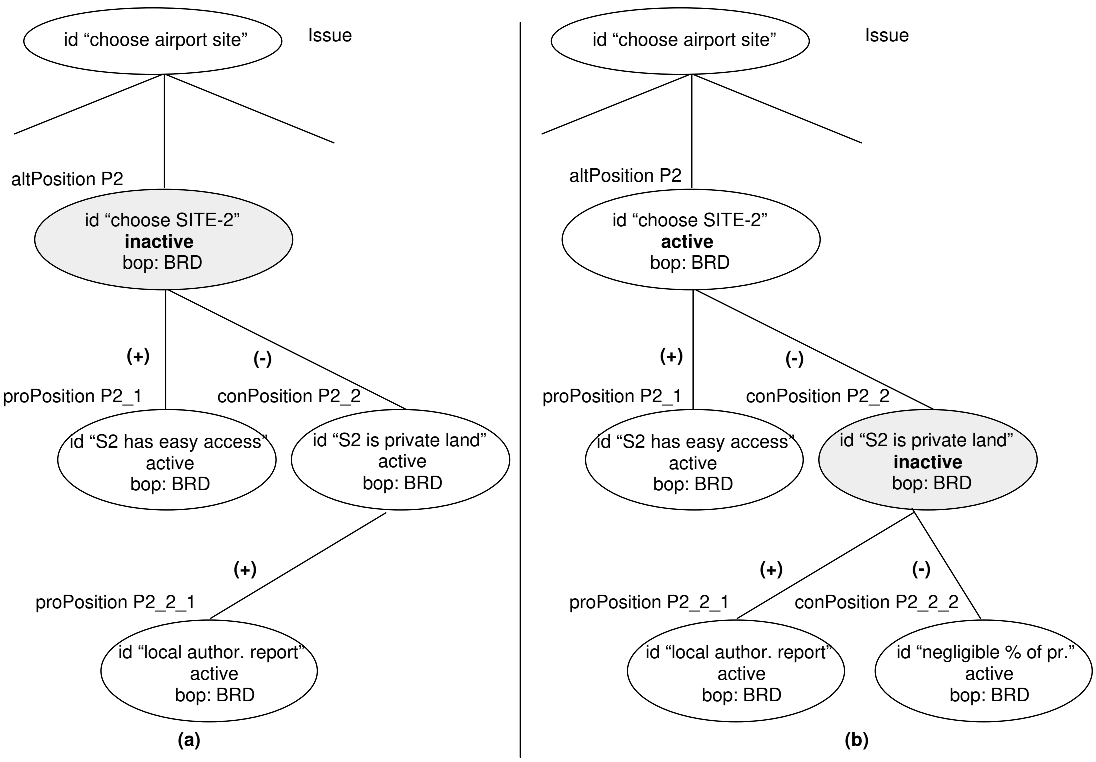

Figure 2.

The BoP can be changed at any phase during the discussion. Figures 3(a) and 3(b) illustrate the transition from SoE to PoE for position P1. At this point, there are no constraints about the relative importance of the pro and con positions. Therefore, P1_1 and P1_2 have equal weights. *PosScore* denotes the sum of weights for the supporting positions, and *conScore* for the counter-positions. At the next step, as shown in Figure 3(c), a constraint stating that P1_2 is more important than P1_1 is inserted, causing the weights to be updated by the scoring mechanism. P1 becomes inactive because *conScore* is now greater than *posScore*.

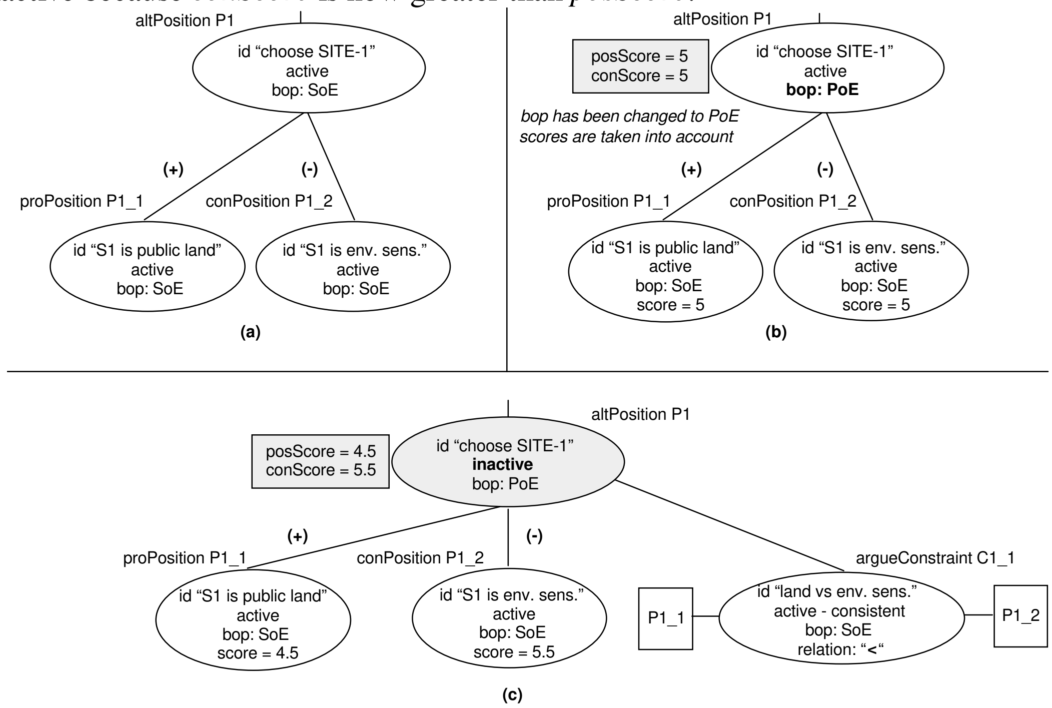

Figure 3.

The framework includes algorithms for *redundancy* and *consistency* checking. When a new constraint is about to be added, the system checks if both its left and right parts exist in another, previously imposed, constraint. If yes, the new constraint is *redundant* (Figure 4, constraint C1_3) and ignored. If not, the system checks whether it is consistent with the previously inserted (active) constraints in which case, it is taken into account by the scoring mechanism. Otherwise, it is labeled *inconsistent* (Figure 4, constraint C1_4).

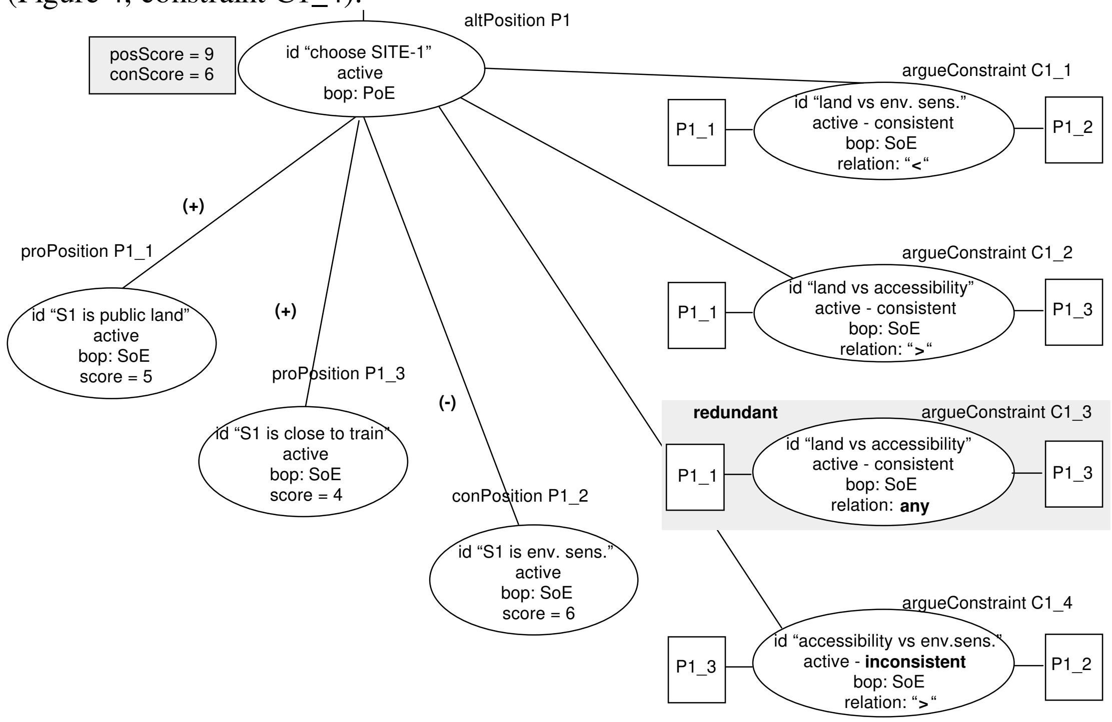

Figure 4.

The activation status of constraints can change due to two factors: the discussion underneath and the activation status of their consistuent positions. As shown in Figure 5, position PC1_2_1 inactivates constraint C1_2 (SoE with no supporting arguments). In addition, position P1_1_1 inactivates position P1_1, which in turn inactivates constraint C1_1 (we assume that C1_2 is already inactive). C1_4 remains active because it does not involve P1_1. Notice that C1_4 now becomes consistent (*path consistency* runs whenever there is a change to the activation status of a constraint), and is the only constraint taken into account by the scoring mechanism.

Finally, with the aid of similar procedures, the issue is resolved by choosing among the (active) alternative positions the one(s) with the best score(s) (the BoP for the initial issue is always PoE). Although the above example refers to environmental planning, the argumentation framework is general, in the sense that it can be applied to any type of discussion. Environmental planning applications however, need a number of additional features to enhance the usability of the system and facilitate collaboration between the involved parties.

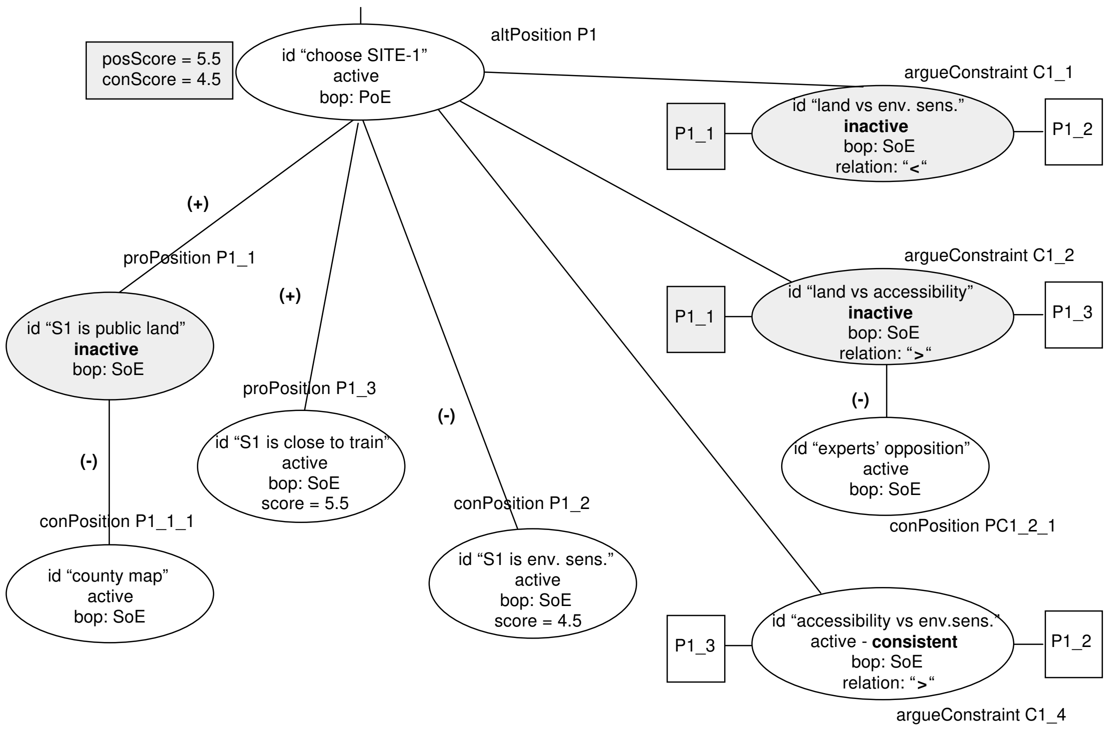

Figure 5.

3 Enhancements for Collaborative Environmental Planning

Environmental planning usually involves large amounts of information which is stored in multiple systems. These systems may be geographically distributed, follow numerous data formats, and provide different types of services. Since most decision makers are not expert users, they should have on-line access to these data through user-friendly interfaces. Therefore, *interface issues* and *interoperability* play a very important role in CEP systems. Furthermore, in addition to preference relations, environmental planning often involves spatial constraints (constraints that describe location in space). The processing of such constraints requires *spatial inference mechanisms* which are usually more complex than the ones used for preference relations. In the rest of the section, we discuss enhancements to the original argumentation framework that deal with these issues.

3.1 Semantics, interfaces and interoperability

A system for CEP should be able to answer advanced queries involving numerous types of data (e.g. maps, images, text). Usually, multimedia information can be described in many different ways. Thus, specification of the appropriate *semantics* is of high importance. Semantics may be keywords or any other kind of meta-data (information about the structure and the content of data). Domain experts may be the best *annotation tool* available for the specification of semantics. Instead of trying to develop automatic annotation tools, we should rely on human expertise and intelligence, whenever we can. Annotations can be obtained from the usage of data and should be dynamic rather than assigned just at the input time. Semantics should be used in *content-based retrieval* of multimedia information. Meta-data have an important role in content-based retrieval: they provide very important information for browsing data and can be used to reduce the search space. Techniques to introduce and manage semantics are essential because queries require assignment of semantics to data. However, the fact that semantics are domain-dependent as well as subjective has to be taken into account. An efficient system should provide an environment in which data interpretation from different perspectives is facilitated.

CEP involves both *machine-based* activities, where a database or other information resource is used, and *human-based* activities, where a person or a group is required to intervene, interacting with the information system. The user evaluates items (e.g. with regard to relevance) as he proceeds through the search, and may change the search strategy at any point of this process (Tague-Sutcliffe, 1996). User acts, such as navigation, selection, and refinement of the search, help greatly in reducing the burden of the related algorithms.

Besides, CEP systems require special *interfaces*. For example, spatial data are often best queried through a display of the space, upon which the user can outline regions that may be hard or impossible to describe as an SQL query. Multimedia information repositories should maintain an evolving set of representations of the same or similar information (both for data and metadata). Relationships among these representations must be also maintained. In this way, changes to some representation of an object can be propagated to related representations.

Since existing GISs and other types of data repositories are based on proprietary formats and/or structures, data conversion and integration is another basic requirement for the system. In general, some integrating notation and model must stand in the middle among all these sources. Each source is *wrapped* by a component that translates between the viewpoint of the source and the shared, global viewpoint. Higher-level products may then be built from these wrapped sources. Wiederhold uses the concept of *mediators* as a means to combine disparate information sources (Wiederhold, 1992). He considers mediators to be tools able to perform customized integration of the information, perhaps with additional filtering or processing. According to him, we have to model information processing in future information systems as an interaction between data and knowledge. Data may be gathered automatically or clerically, while experts are needed to gather and formalize knowledge.

The enormous growth of the GIS community advocates for establishing *open* approaches on issues such as spatial data models, system architectures, transmission protocols, and user interfaces. Interfaces between existing GISs and international geospatial data standards, such as the *NSDI Q geospatial metadata* and spatial data transfer standards, need to be built. Heading to an *open* GIS, interoperability issues between different platforms have also to be solved.

3.2 Knowledge representation and reasoning techniques

In addition to preference (and other non-spatial) constraints, environmental planning problems may involve spatial constraints. Such constraints may be explicit (e.g., the airport should be *near* a train station or a highway in order to minimize cost of transportation) or implicit (e.g., the airport cannot be *in* a mountainous area or wetland). As in the case of preference relations, spatial constraints may contain inconsistencies that depend only on the nature of the constraints; their early detection helps avoid unnecessary extensive access to stored data. If for example, an agent asserts that the airport should be *east* of the major urban area and another *west*, this leads to an inconsistency independently of the position of the urban area (assuming that the relations *east* and *west* are mutually exclusive). In real applications, that involve a large number of spatial constraints imposed by numerous agents, the detection of inconsistencies requires sophisticated spatial inference mechanisms.

Spatial inference mechanisms should remove inconsistencies and provide a set of constraints that corresponds to actual configurations of objects in space. They may have multiple goals: find one solution that satisfies all constraints, find all possible solutions, a number of possible solutions etc. It is difficult to develop a model that can express all types of spatial constraints and is capable of performing efficient inconsistency detection. Work on spatial constraint satisfaction problems has concentrated mostly on *homogeneous constraints*, that is constraints of the same form (only topological or only distance). Even for such cases, a large class of problems is intractable (Grigni et al., 1995).

The first step of a spatial inference mechanism is *path consistency checking* using some composition table. *Composition* infers the relation between two objects X and Y when their relation with a third object Z is known. If, for example, X is *inside* Z and Z is *north* of Y, we can conclude that X is *north* of Y. The path consistency algorithm explicates all constraints between every pair of objects and removes relations that yield inconsistencies. Work in the area of spatial reasoning has produced composition tables for homogeneous spatial constraints (e.g., topological in (Grigni et al., 1995), direction in (Papadias and Sellis, 1994), distance in (Topaloglou, 1994)) and for combinations of several types of constraints (topological and direction in (Hernandez, 1994); direction and distance in (Frank, 1992)).

Path consistency algorithms have a polynomial complexity (Egenhofer and Sharma, 1993) but in general do not suffice for achieving satisfiability. They reveal only local inconsistencies and search-based constraint satisfaction algorithms (such as *backtracking*) have to be applied to locally consistent networks in order to ensure global satisfiability. Although such algorithms have an exponential worst case complexity, efficient heuristics provide good results in many cases (Dechter, 1990).

4 The GeoMed system

World-Wide Web provides an interesting platform for the implementation of certain types of Multimedia Information Systems. Such an approach is becoming quite popular in various applications, ranging from scientific collaboration to political negotiation and from medical diagnosis to home shopping. However, special attention must be paid to indexing, organization and management of information in the unstructured world of the Web.

GeoMed aims at improving both the accessibility of geographical information in heterogeneous, distributed GISs, and the users' participation in related decision making processes, such as city, regional or environmental planning. Multimedia information is represented within electronic documents and messages, not only as passive data but as graphical user interfaces to external Information Systems (Karacapilidis et al., 1996). The GeoMed server is being written in Java, which allows its use on a variety of operating systems. It supports three types of services, namely information, documentation and mediation services:

- *Information services* provide efficient access to multimedia information in distributed databases over wide-area networks. This includes services for finding relevant multimedia data, converting proprietary data to standard formats for data interchange (by *wrapping* their appropriate objects), as well as ways of viewing and browsing multimedia information from within general purpose hypermedia systems such as the Web.

- *Documentation services* provide a shared workspace for storing and retrieving multimedia documents. For this purpose, GeoMed uses the *Basic Support for Cooperative Work (BSCW)* system (Bentley et al., 1995), originating at GMD and being further developed in the CoopWWW EEC project. The above workspace is actually a multimedia information repository which stores and manages both data and metadata.
- *Mediation services* provide a new kind of issue-based conferencing and group decision support system. Users can propose and discuss alternative solutions to some problem or issue by sending electronic messages to a GeoMed server. While receiving these messages, the system builds up a *discussion forum* on the Web by classifying the inputs according to their context. Users are able to use hyperlinks to access a structured protocol of the discussion, which is based on the argumentation framework discussed in Section 2. Hyperlinks between the argumentation elements and the original messages allow quick and easy access to the original full contents of the cited messages.

In order to use GeoMed one requires only a Web browser and Internet access. Users' messages may include any HTML hypermedia elements, such as links to other resources on the Web, images, movies and sound files. Many of these links may depend on the content. Thus, clicking on an object leads to information related to it. The above structure provides a view on the messages which substantially facilitates the browsing and retrieval of the required past contributions to the discussion. As shown in Figure 6, this schema is more precise than the typical *thread* mechanism of *newsgroups*.

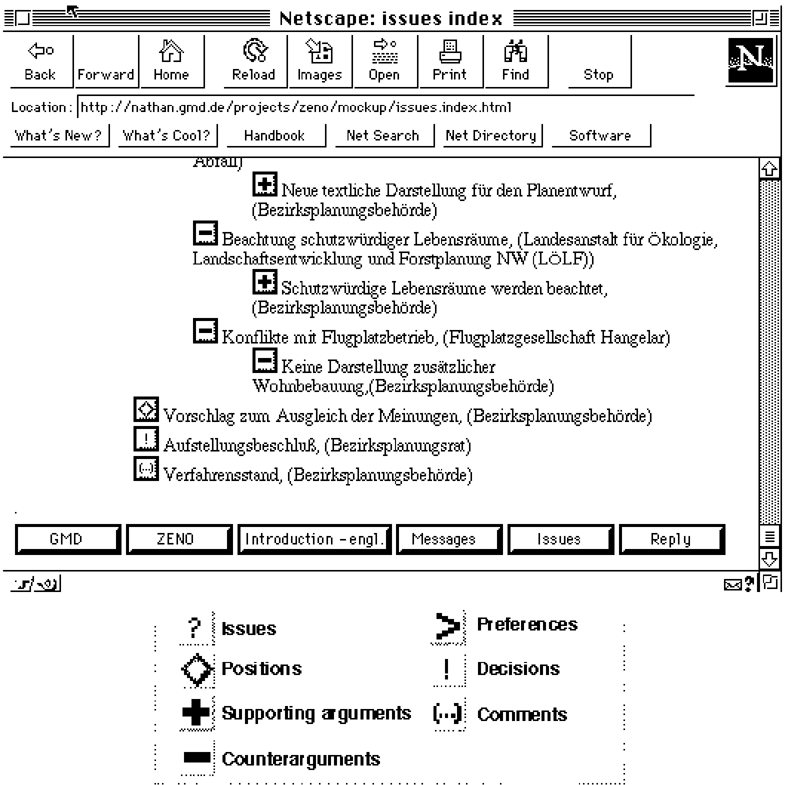

Figure 6. Structure of the discussion in GeoMed.

The index is stored in a relational database, allowing messages and argumentation elements within messages to be selected, filtered and sorted using standard SQL as a database query language. Web-based user interfaces provide a simple way for users to insert their input and/or make common queries, eliminating the need of familiarity with SQL. For instance, the interface of Figure 7 provides the users a means of arguing (replying) to an already asserted position. The types of reply allowed are “Argument Pro”, “Argument Con”, “Comment” and “Decision”. The first two types, as discussed in Section 2, may alter the status of the discussion protocol according to the reasoning mechanisms. The third type allows a user to insert any kind of input (e.g., additional information on the issue which is being discussed) without influencing the above status. The last type of reply “enforces” the system to accept the position under discussion without involving the reasoning mechanisms. In the case of alternative positions, this can be viewed as a decision for the corresponding position to be the strongest or the most important one.

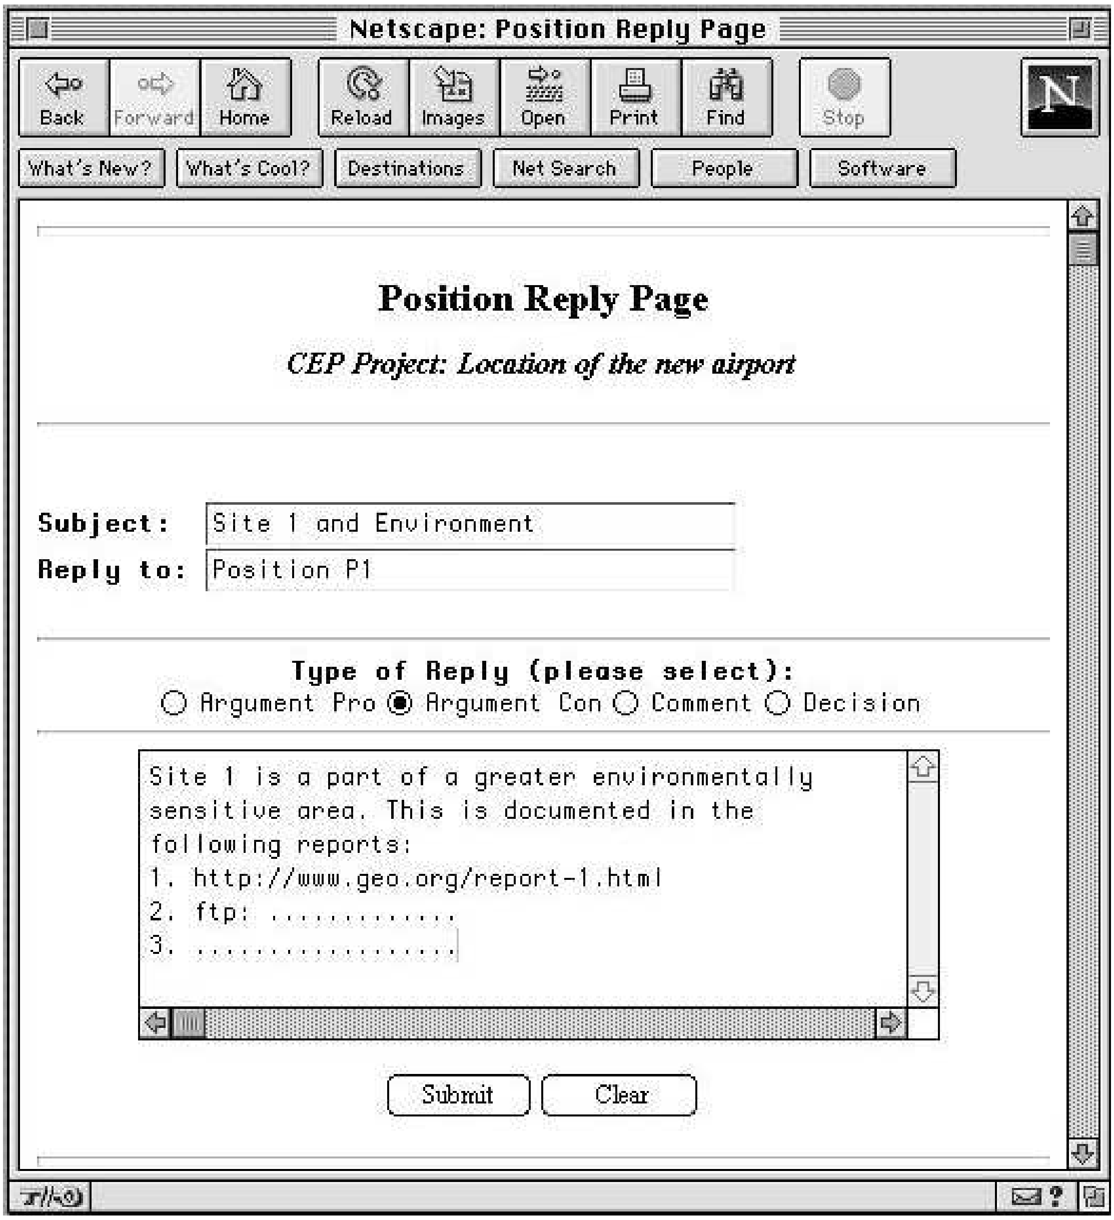

Figure 7.

A future version of GeoMed will make use of the roles of participants in the planning process, and models of procedural norms, to advise users about their rights and obligations in the proceeding. GeoMed is not, however, intended to play the role of a cop or judge. If a planned action would violate the norms of the proceeding, according to the model, GeoMed's task is limited to advising users of this fact. The users remain free to decide for themselves just how rigidly they wish to follow the model. Finally, the use of established database technology in the future will provide transaction control and improved security, by means of optional user accounts and password protection.

5 Discussion

As further documented in (Voss, 1996), it is rather impossible to completely automate the processes described in the previous sections. Users involved usually have diverse and conflicting interests, power and positions. Nevertheless, the use of information technology may assist them in various ways. One important goal is to provide easy access to the current knowledge, at any time. Referring to GeoMed, this goal would be greatly facilitated if all relevant data and documents, including cartographic and thematic maps, decisions, minutes, evaluations, etc., are made available and maintained in electronic form, in a well-structured and organized way. Another possibility, and this is our primary goal, is to provide direct computer support for the argumentation, negotiation and mediation process. A computer network can be used as a medium for such a process, in which documents are structured and indexed according to their role and function, using a model of argumentation.

Three major practical requirements are essential for the development of an efficient CEP system (they are almost ideally met by exploiting the Internet and the World-Wide Web):

- the system must be available on all prominent operating systems and hardware platforms;
- it must provide relatively inexpensive access to a broad public, and
- it must have a very intuitive and easy user interface.

In procedures like CEP, the quality and acceptability of geographical planning decisions depends not only on the availability and distribution of accurate information, but also on the fairness and openness of the planning procedure. For example, when planning the path of a high-speed train through communities and natural spaces, or the location of a site for storing hazardous wastes, the interests and perspectives of the affected communities and citizens, the responsible regional or federal governments, environmental protection groups and industry representatives are likely to be in conflict. Involving representatives of these diverse interests in the planning process, at the earliest possible stage, can facilitate the extraction of a better plan and its acceptability by the interested parties, helping to avoid long, expensive delays or even legal battles.

Acknowledgements: This work was partially funded by the European Commission (DG XIII, Information Engineering Program, Project IE-2037). The GeoMed consortium consists of Intecs Sistemi (IT), GMD (D), Intrasoft (GR), Vrije Universiteit Brussel (B), TNO-FEL (NL), TNO-Bouw (NL), City of Bonn (D), City of

Tilburg (NL), Tuscany Region (IT) and the Technical Chamber of Greece (GR). Dimitris Papadias is financed by the Commission of the European Communities through the ERCIM Fellowship Programme.

References

Bentley, R., Horstmann, T., Sikkel, K., and Trevor, J. (1995), "Supporting collaborative information sharing with the World-Wide Web: The BSCW Shared Workspace system", in: *Proceedings of the 4th International WWW Conference*, Boston, Massachusetts, 63-73.

Daniel, S.E., Diakoulaki D.C., and Pappis, C.P. (1996), "Operational Research and Environmental Planning", Position Paper, XIV EURO Summer Institute on Environmental Planning, Athens, Greece.

Dechter , R. (1990), "Enhancement Schemes for Constraint Processing: Backjumping, Learning and Cutset Decomposition", *Artificial Intelligence* 41, 273-312.

Densham, P.J., Armstrong, M., and Kemp, K. (1995), Collaborative Spatial Decision-Making, Scientific Report for the Initiative 17 Specialist Meeting, NCGIA Technical Report 95-14, UC Santa Barbara.

Egenhofer, M., and Sharma, J. (1993), "Assessing the Consistency of Complete and Incomplete Topological Information", *Geographical Systems* 1/1, 47-68.

Frank, A.U. (1992), "Qualitative Spatial Reasoning about Distances and Direction in Geographic Space", *Journal of Visual Languages and Computing* 3, 343-371.

Grigni, M., Papadias, D., and Papadimitriou, C. (1995), "Topological Inference" in: *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95),* Montreal, Canada, 901-906**.**

Hernandez, D. (1994), *Qualitative Representation of Spatial Knowledge*, Lecture Notes in AI 804, Springer-Verlag, Berlin.

Jelassi, M.T., and Foroughi, A. (1989), "Negotiation Support Systems: An Overview of Design Issues and Existing Software", *Decision Support Systems* 5, pp. 167-181.

Karacapilidis, N., and Gordon, T. (1995), "Dialectical Planning", in: *Proceedings of IJCAI-95 Workshop on Intelligent Manufacturing Systems*, Montreal, Canada, 239-250.

Karacapilidis, N., Papadias, D., and Egenhofer, M. (1995), "Collaborative Spatial Decision Making with Qualitative Constraints", in: *Proceedings of the 3rd ACM International Workshop on Advances in Geographic Information Systems*, Baltimore, Maryland, 53-59.

Karacapilidis, N., Gordon, T., Papadias, D., and Voss, H. (1996), "Building an Interactive Multimedia Information System: The GeoMed case study", in: *Proceedings of the ECAI'96 Workshop on Knowledge Representation for Interactive Multimedia Systems: Research and Experience*, Budapest, Hungary, 10-16.

Karacapilidis, N. (1996), "An Argumentation Based Framework for Defeasible and Qualitative Reasoning", in: *Proceedings of the ECAI'96 Workshop on Gaps and*

Bridges: New Directions in Planning and Natural Language Generation, Budapest, Hungary, 37-42.

Kemp, Z. (1995), "Multimedia and Spatial Information Systems", *IEEE Multimedia* 2/4, 68-76.

Kunz, W., and Rittel, H.W.J. (1970), "Issues as Elements of Information Systems", Working Paper No 0131, Universität Stuttgart, Institut für Grundlagen der Planung.

Mackworth, A., and Freuder, E. (1985), "The Complexity of some Polynomial Network Consistency Algorithms for Constraint Satisfaction Problems", *Artificial Intelligence* 25, 65-74.

Papadias, D., and Sellis, T. (1994), "Qualitative Representation of Spatial Knowledge in Two-Dimensional Space", *Very Large Data Bases Journal, Special Issue on Spatial Databases* 3/4, pp. 479-516.

Papadias, D., Theodoridis, Y., Sellis, T., and Egenhofer, M. (1995), "Topological Relations in the World of Minimum Bounding Rectangles: a Study with R-trees", in: *Proceedings of the ACM Conference on the Modeling of Data (SIGMOD),* San Jose, CA, 92-103.

Rittel, H.W.J., and Webber, M.M. (1973), "Dilemmas in a General Theory of Planning" *Policy Sciences* 4, 155–169.

Tague-Sutcliffe, J.M. (1996), "Some Perspectives on the Evaluation of Information Retrieval Systems", *Journal of the American Society for Information Science* 47/1, 1-3.

Topaloglou, Th. (1994), "Qualitative Reasoning about Imprecise Spatial Information", in: *Proceedings of the ACM Workshop on Advances in Geographic Information Systems*, Baltimore, Maryland, 122-127.

Voss, H. (1996), "Dimensions of KBS Applications", in: Fr. Maurer (ed.), *Proceedings of the 2nd Knowledge Engineering Forum*, Karlsruhe, Germany, Universität Kaiserslautern, SFB 501 Bericht 01/96.

Wiederhold, G. (1992), "Mediators in the Architecture of Future Information Systems", *IEEE Computer* 25, 38-49.