Geo-spatial Tactical Decision Aid systems: fuzzy logic for supporting decision making

Raffaele Grasso
NURC, NATO Undersea Research Centre
La Spezia, Italy
grasso@nurc.nato.int

Simone Giannecchini
NURC, NATO Undersea Research Centre
La Spezia, Italy
giannecchini@nurc.nato.int

Abstract - *Tactical decision aid systems support the military or civilian operation planning process providing the decisional authorities with a simplified view of the environmental conditions over a theatre of operations. Methods for multi-source data fusion and support for disseminating and managing geospatial data are key factors for a successful system implementation. This paper describes a tactical decision aid system based on fuzzy logic reasoning for data fusion and on current Open Geospatial Consortium specifications for interoperability, data dissemination and geo-spatial services support. Results from system implementation tests during live exercises are reported and discussed showing the flexibility and reliability of the proposed architecture. Future directions are provided and discussed as well, including web processing services, context fuzzy reasoning and group decision making.*

Keywords: Geospatial data fusion, decision support systems, interoperability, OGC specifications, geospatial services.

1 Introduction

1.1 Decision Support Systems

TDA systems implement data fusion schemes for supporting military as well as civilian operations such as disaster prevention and search and rescue, by giving a simplified and meaningful picture of a particular situation described by a series of data coming from different sources. This analysis is the last step, after environmental data collection, before presenting an assessment of the situation to the end user.

Data fusion plays a fundamental role in reducing the complexity of the environmental picture by producing value-added products, storing synthetic information that could be easily interpreted by a non expert end user with no strong technical and scientific background. Such a necessity is becoming increasingly important due to the growing number of sources that could be employed in characterising the environment including in situ sensors, meteorological and oceanography models, and remote sensing data.

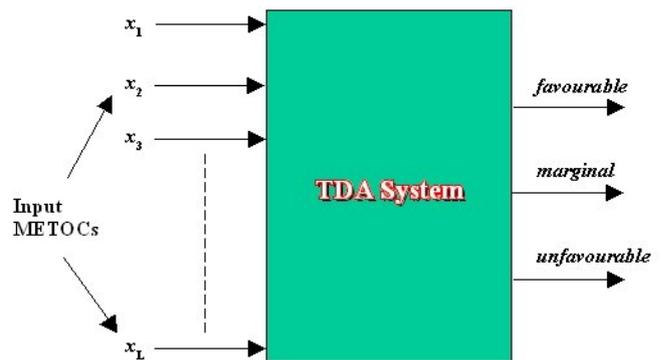TDA systems provide automatic decisions for launching an operation and calculating percentages of

Figure 1: TDA System.

risk in relation to decisions for a given forecast time and spatial position, based on forecasts of the environmental parameters affecting that operation. They belong to the more general class of Decision Support Systems (DSS) that are used in a variety of fields (finance, economy, tracking systems and medical diagnosis, for instance) for assessing the risk associated to an action.

The architecture of above mentioned systems usually include a prior knowledge base that is used for inferring the decision for a given input. Typical approaches to DSSs include Bayesian methods, evidential reasoning and neural networks, but in general, they can be categorised into two main classes: normative and descriptive [1]. Normative DSSs (Bayesian methods, for instance) are based on objective prior knowledge, while descriptive systems are more oriented toward problems involving non-objective knowledge and vague concepts. The more general the prior information to be incorporated, the more we will be driven toward a descriptive decision making scheme.

Objective prior knowledge is not always available for decision making applications. For this reason, TDA systems in general should deal with both objective and subjective prior information. These systems should be able to treat approximate knowledge and linguistic terms (such as very high or very low), which are usually used by human reasoning to make assertions and draw decisions [2]. A very general architecture which is capable of managing both types of information is needed in our case.
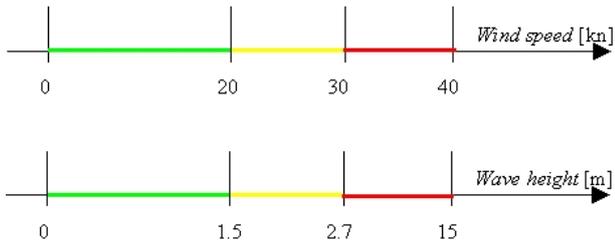
Figure 2: METOC's thresholds for Naval Refuel operations. The green, yellow and red colours indicate good, medium and bad environmental conditions for the operation.

1.2   Past Experiences at NURC

A TDA system for a particular operation can be considered as a *black-box* accepting as input a vector $\bar{X} \in R^L$ of meteorological and oceanography parameter (METOCs) values, given by models and/or estimated from remotely sensed data and in situ measurements (see fig. 1). A vector of parameters stores the METOCs for a given spatial position and forecast time. The output is a decision among three different categories, favourable, unfavourable and marginal, which may be translated to 3 decisions, to launch or not to launch the operation, or to wait and observe as the situation develops.

Systems making use of classic logic, such as the METOC Impact Matrix Summary (MIMS) system developed at the NATO Undersea Research Centre (NURC) in the recent past, infers the final decision by means of a set of thresholds defined for each input METOC (see, for instance, fig (2) for wind speed and wave height affecting Naval Refuel). The thresholds partition each METOC domain along the real axes, producing a series of subsets representing a particular environmental condition. The system checks which subset the input value of a METOC belongs to, and then makes a decision based on a lookup table. This implementation has a few limitations:

1. the decision thresholds are crisp values and the system does not take into consideration possible uncertainties arising especially at the borders between two subsets along the METOC axes.

2. the system can process input crisp values only; it is not possible to take into account the uncertainty due to measurement noise and/or modelling errors in the input values.

3. the system does not generalise, i.e. all entries in the decision lookup table should be specified in order to take into account all the possible situations that may be encountered; this may not be feasible for large L-dimensional input vectors and large number of METOC subsets.

4. it is not possible to work with vague concepts and linguistic terms directly.

5. the system does not provide any measure of the cost associated with making a decision among the three possible risk categories favourable, unfavourable and marginal.

6. it is not trainable using sets of input output pairs.

The above listed shortcomings provided motivation for moving toward an approach based on fuzzy logic rather than on classic crisp logic as we had done in the past. The next section will introduce and discuss such an approach.

1.3   TDAs based on fuzzy experts.

Fuzzy logic adds much more flexibility than a classic logic system, like the one introduced and described above. As such, it is able to overcome the issues described in the MIMS system. A fuzzy system manipulates sets using fuzzy logic operators, which are an extension of classic theory of sets and logic. The membership function is a key concept in fuzzy logic. In classic logic, the value of the membership function is 1 for each member of a set, indicating that these elements belong totally to that set, and 0 for those elements not inside the set.

The classic theory of sets results as a particular case of fuzzy set theory. In fact, fuzzy logic allows the membership function to take values in the range [0, 1], so that an element could belong to different sets, but with a certain degree of membership. This is used to give analytical meaning to vague concepts, which are typical of human reasoning (such as the wave height is very high or the wave height is not so high) [2].
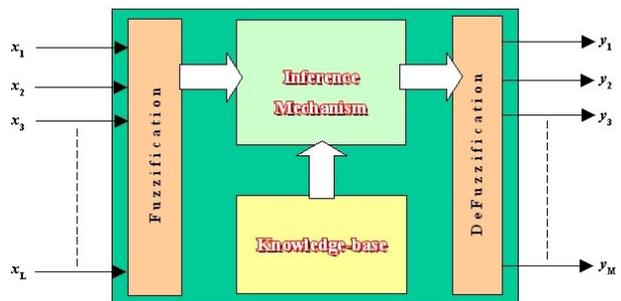

Figure 3: Fuzzy inference engine basic architecture.

The fuzzy TDA system developed at NURC is an expert system which exploits the features of a standard fuzzy inference engine [2] as depicted in fig. 3. The system maps vectors in R^L to vectors in R^M, with L and M being the number of inputs and outputs, respectively. The main components are the interfaces with the crisp world, namely the fuzzification and the de-fuzzification blocks, the knowledge base and the inference mechanism.

The fuzzification step converts a numerical input value (crisp input), x_i, into a fuzzy set with its own membership function. The membership function is defined over the universe of discourse (UOD) of the input and its shape, for example, may be set in order to model the uncertainty affecting the input.

The knowledge-base block stores all the prior knowledge needed for the inference system to draw conclusions and to provide the output values for a given input. The stored prior knowledge consists of two parts: a) input and output linguistic variables and terms with their membership functions and b) fuzzy rules in the form of *if-then* statements [2]. A linguistic variable is a variable whose values are words from a natural or artificial language. A variable is characterised by its name, *wave height* for instance, and a set of linguistic values or terms, such as *low*, *medium* or *high*, which are fuzzy sets covering the variable UOD. The meaning of each linguistic term is determined by a specific membership function defined on the variable UOD (see fig.4). A
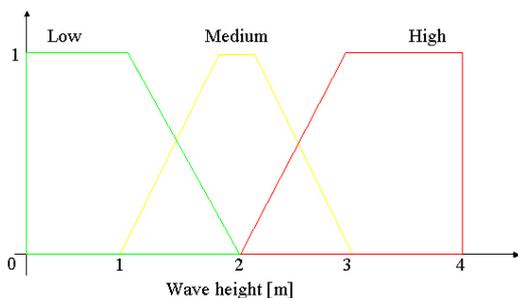

Figure 4: Example of linguistic terms with trapezoidal membership functions defined for the variable *wave height* with UOD ranging between 0 and 4 m.

fuzzy rule is an *if-then* sentence having a premise, the *if* part, which connects the input fuzzy variables and terms through AND-OR logical operators, and a consequence (the *then* part) determining the fuzzy value of the outputs in response to a particular combination of the fuzzy terms in the premise, e.g., "*if* (wave height = High) AND (wind speed = High) *then* (favourable = Low) AND (marginal = Low) AND (unfavourable = High)".

The inference mechanism transforms the input fuzzy sets into the output fuzzy sets by using fuzzy logic operators and taking into account the premises and the consequences of the rules stored in the knowledge base [2].

The defuzzifier converts the sets from the outputs of the inference mechanism into crisp numbers. In TDA applications, for example, the user may be interested in a measure of belief of the risk categories favourable, marginal and unfavourable, in order to make a decision based on a maximum criterion.

A fuzzy system used as the constitutive block for the new TDA system prototype is the valid solution for overcoming the main issues related to the MIMS system developed in previous works. In particular:

1. the user may configure the rule base by using linguistic variables, terms and expressions from the natural language,

2. the system is flexible in being able to deal with both uncertain reasoning, by exploiting fuzzy logic tools, and classic logic by simply constraining the membership functions in the rule base to be of the crisp type,

3. the system is able to generalise so that there is no need to extract all the possible fuzzy rules, which can be obtained by combining the linguistic terms in the rule antecedent,

4. the fuzzifier allows the system to take into account the numerical uncertainties of the inputs by simply using non-singleton fuzzy sets during fuzzification,

5. it is possible to have a measure of the final decision beliefs by looking at the numerical values of the fuzzy engine outputs.

A fuzzy system is placed in the final TDA prototype architecture which also supports data pre-processing, geo-referencing, polygon vector geometry extraction and data visualisation. The overall system architecture is described in the following section.

2   System Description

2.1   System's Requirements.

Systems for the fusion of data from different sources are in general very complex. They have to support several applications such as parameter estimation, target detection, tracking and classification, situation assessment and high level inference [5]. A variety of tasks have to be performed, including sensor data pre-processing, data base management, data ingestion and dissemination, communication between data fusion nodes and sensors, geographical transformations and effectual data visualisation.

During the first half of 2005 a prototype of a general purpose TDA for the support of military and civilian operations was developed at NURC and it was validated during the Coalition Warrior Interoperability Demonstration (CWID) live exercise.

Requirements which guided the design and development were as follows,

- support for multiple inputs and outputs.

- usage of a fuzzy inference engine for supporting human reasoning.

- automatic processing of data time series for event tracking.

- full configurability (choice of the inference mechanism, type and parameters of the membership functions, linguistic variables and terms, data pre-processing scheme, selectable bounding box and forecast time interval, selectable data source format, etc.).

- ability to process geo-spatial input data arranged on regular grids, in a series of sparse points or as vector-polygons,

- possibility of multi-layer data visualisation for spatial querying of stacked input and output layers using public, wide-accepted standards from the Open Geospatial Consortium (OGC) [6].

Having defined the base requirements of the proposed system, in next section we will introduce and dissect its architecture.

2.2 Overall Architectural Design

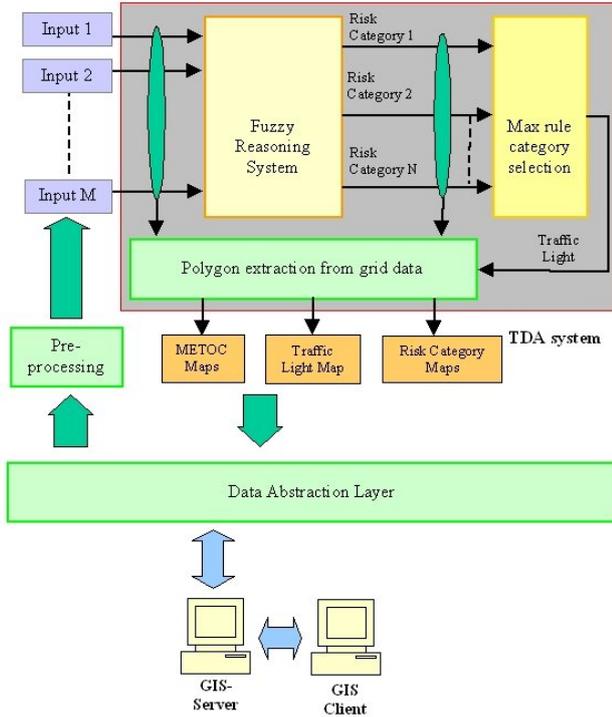The system architecture designed and implemented at NURC is depicted in fig. (5).

Figure 5: TDA system basic architecture with maximum decision rule.

The system was developed for the processing of data arranged both in a regular geo-referenced grid or as sparse points. Future versions will include the possibility of processing polygon vector data [11]. Grided data METOCs, in situ measurements and/or remotely sensed data, for instance, are stored in a data base ready to be pre-processed depending on the scheduled operation. The inputs to the system are grids, which are segmented and converted into polygons before entering the inference engine. The reader may refer to [3] for an overview of the algorithms for connected component labelling and contour tracking which are the fundamental processing blocks used in extracting the polygons from the data grids.

Data in the inference engine is processed on sample basis and the outputs are arranged on regular grids one for each category of risk (favourable, unfavourable and marginal, for instance). The outputs provide the percentage of belief of each risk category, in response to a particular value of the input METOCs. The final decision is inferred by using a max rule: the category having the maximum output value is selected for each

sample in the data grid. A numeric label indicating the chosen risk category is stored in the final output grid (a traffic light map or risk map indicating the areas where the conditions are likely to be favourable, unfavourable or marginal for launching the operation). Polygon vectors are extracted from the fuzzy engine outputs and the final risk map as well and stored in a format suitable for vectorial data like, as an instance Shapefile or WKT.

The end user can access the results from a GIS client to retrieve the inputs and the outputs, and display them one on top of the other. The GIS allows the user to make spatial queries on the layers in order to display the information related to a specific site and stored as attributes of the shape files.

An alternative implementation consists of a fuzzy system with one risk output. The output risk map is obtained through a multi-threshold decision rule depending on the number of categories the traffic light map should have (see fig. 2.2).
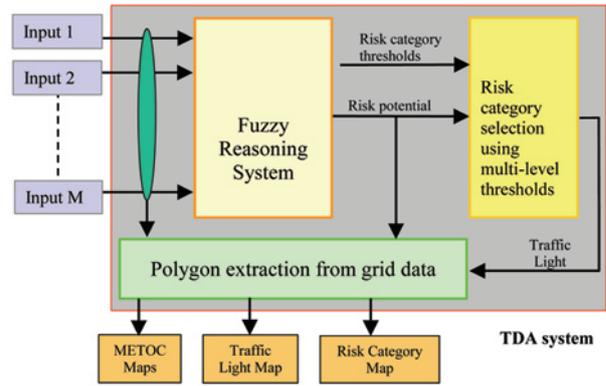
Figure 6: TDA system architecture with multilevel threshold decision rule.

During CWID05, the former implementation was used to configure each output fuzzy linguistic variable with three terms low, medium and high. For the input linguistic variables and terms the same configuration was used with membership function parameters depending on the METOCs affecting the particular operation. Using the alternative configuration in 2.2 implies setting the linguistic terms and the membership functions for one output linguistic variable. In general, the number of linguistic terms for each variable (input or output) is configurable without limits allowing a finer categorisation of the variable domain. In any case, the time needed to run the inference engine depends among other reasons, on the number of fuzzy linguistic variables and terms, so a trade-off between the processing burden and this number has to be considered during system configuration. The current version of the system allows the application manager to set the parameters determining the shape of the membership functions heuristically based on his own knowledge of the problem at hand. Future versions will include automatic algorithms for estimating the

membership parameters from training data sets, when available, making this step less subjective.

2.3 Data Abstraction Layer and Interoperability support

The design of the system aims at achieving maximum support for extensibility at all possible levels as well as a high degree of customization. The inference engine is completely configurable through a series of text files specifying the number of input and output variables, the linguistic terms of each variable with their membership functions, the rule base and the parameters of the fuzzy engine, i.e. the fuzzification and defuzzification methods and the inference mechanism.

It is also worth to introduce and discuss an aspect which, although not related to the reasoning part of the TDA proposed here, is gaining more and more importance throughout the geoinformation and geodata fusion community, support for *interoperability* by adoption of OGC and ISO standards.

In figure (5) a layer called *Data Abstraction Layer* is introduced. Its role is to decouple the system itself from the management of underlying geospatial data sources in order to achieve superior robustness through the exploitation of the concepts of *feature* and *coverage*, as introduced by the OGC, for correct data abstraction and architectural tiers' separation [11].

By employing a common format, the data abstraction tier presents the higher level with a uniform view of the underlying sources based on the concepts of feature and coverage (see 7) relieving them from the burden of dealing with issues intrinsic to I/O management such as handling different formats and conventions, allowing them to concentrate on their key requirements. In addition, such an approach ensures that support for new data sources can be added and improvements can be made independently from higher tiers, provided that everything in features or coverages is always converted, depending on the intrinsic nature of the treated dataset. It is worth to point out that a pluggable architecture could be easily developed in order to support this approach. An interesting added value of using this

them in an interoperable way. This solution constitutes a flexible and efficient way to address the main issues related to data fusion systems since end users can access all of the services available on the above servers through suitable clients application developed in compliance with OGC specifications. Standard capabilities provided include the ability to retrieve maps built by superimposing layers from different data sources, the ability to retrieve coverages and features by specifying complex refined queries and the ability to search for metadata about geospatial data [11].

In order to build the above mentioned abstraction layer, we exploited the Geotools library which implements many OGC and ISO specifications, providing strong capabilities where data abstraction and extensible architecture are concerned. Data produced with the developed system was delivered to our customers during the CWID05 exercise by employing a custom version of the Geoserver, an OGC compliant server.

3 TDA operational use during CWID05 and results

CWID 2005 provided a good opportunity to demonstrate new systems for data processing and exploitation, and the capabilities of NURC data fusion laboratory to deliver near real time value-added products for environmental assessment. In particular, the TDA system and related products contributed by building the virtual operational scenario of the geographic area selected for the theatre of the exercise. From the fusion centre at NURC, grids of six hours forecasts of the risk related to a set of selected operations were routinely sent through the internet and stored in the locally installed geo-server. Participants sent requests through a client to the geo-server to retrieve results, using them during daily briefings. The main theatre of the exercitation was sited in an area of the Florida peninsula facing the Gulf of Mexico and the Atlantic (see fig. 8(a)). A second area on the Adriatic Sea was chosen to exhibit remote sensing products available at NURC [4] (see fig. 8(b)). TDA products were provided for both the chosen sites.

(a) main area, Florida, bounding box 90.2W, 23.6N, 74.6W,37.4N (b) Adriatic Sea area, bounding box 12E, 40N, 20E, 46N

Figure 8: CWID05 theatre of operations

The list of supported operations included:
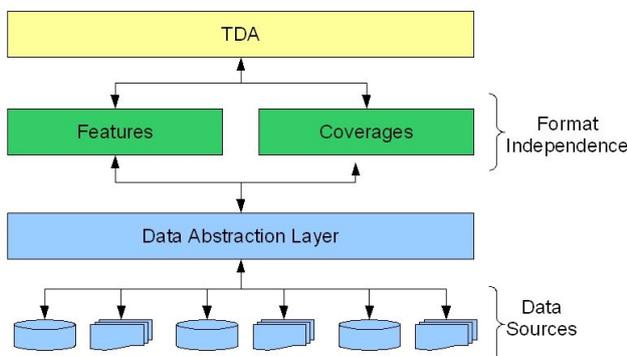
1. Naval refuel (NR),
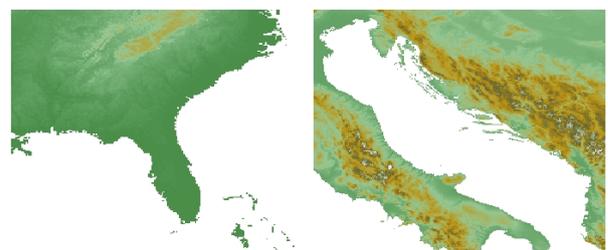

Figure 7: Data Abstraction Level.

approach is that the outputs of the system can be easily exported to servers implementing specifications from the Open GIS Consortium (OGC) in order to expose

2. Amphibious landing,

3. Diver.

4. Heat Impact (HI) on the ground

The Diver operation was chosen as an example of fusion between remote sensing data and models. For this operation, one example is available for the area of the Adriatic Sea. Water visibility [7] and temperature estimated from MODIS satellite sensor data [4] were fused with oceanographic models of wave height and current speed. A day which has all models data and remotely sensed images simultaneously available from four six hours forecasts was selected as reference.

It is worthwhile mentioning that of the two areas, only the Adriatic Sea was chosen for testing during the exercitation planning phase at NURC, before CWID05 took place. However, the Florida area was the main area of interest for the other partners participating in CWID05, who asked for TDA products on this particular region. For that reason, during the exercise, Florida was included in the activity plan of NURC and the TDA system was set up on the fly to support the selected operations for this area. This was a demonstration of the flexibility of the implemented system, which can be easily adapted to a changing scenario, and of the short response time of the data processing and distribution system within the NURC Fusion Centre.

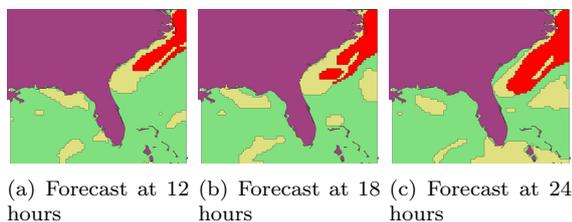(a) Forecast at 12 hours (b) Forecast at 18 hours (c) Forecast at 24 hours

Figure 9: Naval refuel time series for Florida area, Base time 2005-06-20 00:00:00GMT. Green: favourable; yellow: marginal; red: unfavourable.

Figure 9 presents an example of the results of the Naval Refuel operation on the Florida area for 3 successive forecasts. The base time for the two examples is 2005-06-20 00:00:00GMT and the forecasts were provided every 6 hours starting at noon (12 hours later the base time). The NR traffic light maps are dominated by favourable areas in green with wave height roughly between 0.15 and 0.45 m and wind speed of about 5 m/sec. A red area (unfavourable) is located near the US Atlantic coast, with a size growing in time, where the wave height is about 2.5 m and the wind speed with maximum intensity of 13 m/sec. Marginal areas present intermediate METOCs values.

It is worth to remark that the referred pictures has been taken using the uDig OGC compliant client which was used to connect to a running instance of the OGC compliant server GeoServer leveraging the capabilities introduced in section 2.3.

4   Future Directions

Future works will address several issues in order to expand the capabilities of the system. In particular, web processing services (WPSs), context based fuzzy systems and group decision making are of great interest in the framework of operation support.

A major discussion is taking place inside the OGC in order to leverage the experience of ISO for supporting *remote geospatial data fusion*, where the term *remote* refers to the ability of requesting, by mean of a standard protocol, the execution of computations procedures on a remote server. This is a critical capability, especially when performing computationally-intensive procedures on very large datasets where data transfer can be cumbersome and counterproductive (for instance, in terms of bandwidth consumption).

Referring to [8], "a *Web Processing Service[1] (WPS) would provide access to pre-programmmed calculations and/or computations models that operate on spatially referenced data that can be delivered across a network, or available at the server*". This approach does not aim to specify a standard to follow for implementing processes and algorithms but rather, it aims at providing a standard mechanism for enabling remote execution of generic geospatial data fusion computations, regardless to their nature.

The WPS interface[2] comprises of three base operations, *GetCapabilities*, *DescribeProcess* and *Execute*.

The *GetCapabilities* operation provides the user with an XML document describing basic service metadata like service indentification, service version and operations metadata as well as a brief description for each of the offered processes.

The *DescribeProcess* operation retrieves a detailed description of one or more offered processes which should extend the one retrieved from the *GetCapabilities* response. This description provides information on how to trigger the execution of a process in terms of provision of input parameters and retrieval of output values, as well as descriptive information about the process itself.

The *Execute* operation allows users to trigger the execution of one and only one offered process.

It is worthwile introducing some additional capabilities provided by the WPS service. *Synchronous* and *asynchronous* executions are supported. Each process, by means of the *statusSupported* option present in its *ProcessDescription*, can declare through the *DecribeProcess* response whether it supports (or not) asyncronous executions. In case of an *asynchronous* execution the *Execute* response will be sent immediately without waiting for the invoked process to terminate and it will contains a *Status* URL. Clients may exploit the returned URL in order to obtain updated information on the status of the process' execution. Moreover a WPS may be able to *store* the output of an executed process to allow successive web access. By

---

[1]A *service* is a distinct part of the functionality that is provided by an entity through interfaces.

[2]An *interface* is a well-recognizable set of operations[3] characterizing the behavior of an entity.

means of the *storeSupported* option in the *ProcessDescription*, a WPS can make available, through the web, complex results. A client, when requesting execution with store option, will be provided with a URL for retrieving the results of the invoked process.

In figure (10) the execution of an asynchronous process supporting the store option is depicted. We would like to describe further the final phase of the interaction. The client requests the execution of a remote asynchronous process, receiving an XML file which contains a *Status* URL. Once the execution has started the client can check the status of the pending process by exploiting the *Status* URL he has been provided with. Once the pending process has ended, the client obtains an URL indicating where to get the ouputs. In figure (10) usage of an FTP server is depicted.
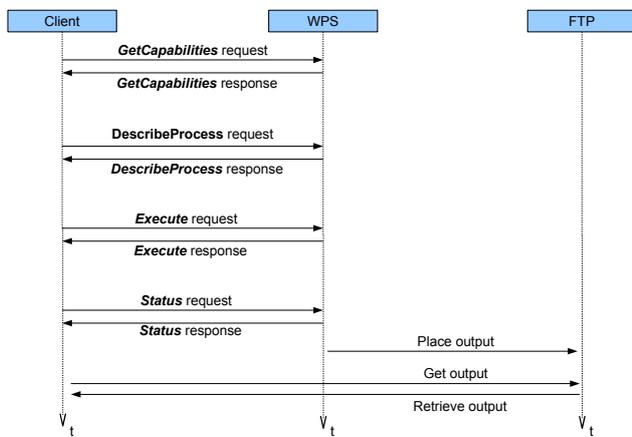

Figure 10: Interaction with a WPS.

In real world applications, the meaning of a linguistic term, *i.e.* its membership function, strongly depends on the context in which the system has to operate. The meaning of *deep water*, for example, may be different for open ocean or coastal areas. Again, the term *high*, referred to *wave height*, depends on the tonnage of vessels employed in a particular maritime operation (naval refuel, for instance). The number of situations characterising a theatre of operations may be extremely wide, and managing all these cases becomes as more complex as the number of system input variables and linguistic terms increases. The simple solution consisting of using different knowledge bases (fuzzy sets and rules) for each of the possible situations may be inefficient or even impracticable.

Context-sensitive reasoning is the area where this type of problems are addressed in order to find tractable solutions. The combined use of context and fuzzy reasoning systems aims at adjusting dynamically, based on the current situation, the shape of the membership function associated with a linguistic value, by using prior information stored in a context knowledge base [9]. This requires finding those parameters characterising context and a way to use this knowledge for extracting the meaning of the linguistic terms.

Investigating the use of these techniques in the framework of operation support is really needed in order to cope with the complexity added to the system architecture by considering context information such as geographic area location, temporal execution and type of vehicles employed.

The management of cooperating operations, with several involved institutions, is another important issue to be addressed in future investigations. Group decision-making (GDM) [10] provides the framework to efficiently manage those situations in which partners from several institutions participate in a joined operation. An agreement among the partners has to be reached if they use different rules for inferring the decision on launching or not launching an operation from the analysis of the environmental conditions. In this case, a centralised TDA system (see fig. 11), providing support for GDM problems, collects decisions from local TDAs and interacts with them in order to achieve a final decision which may be considered as a reasonable compromise by the participants.
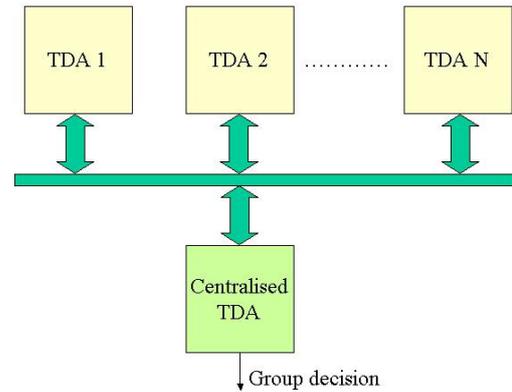

Figure 11: Data fusion architecture for group decision making.

A GDM problem arises when a decision making process involves several expert systems each one providing its own solution to the same problem. The GDM system aims at automatically achieving a common solution maximising the agreement among the experts. Typically, a GDM system is composed of two processes [10]: a consensus process and a selection process (see fig. 12). The consensus process tries to obtain the maximum consensus among the experts on the alternative solutions of a specific problem. In our case, the problem refers to a particular operation to support and the alternative solutions are the risk categories at the output of the TDA expert systems of each cooperating institution. The process is iterative and is coordinated by a *moderator*. The moderator checks the opinion on the alternatives offered by each expert and provides advice to the experts on the base of a consensus measure in order to bring their opinions closer. If the degree of consensus reaches a predefined value, the process stops and the selection process takes place in order to choose the final solution among the alternatives provided under consensus.
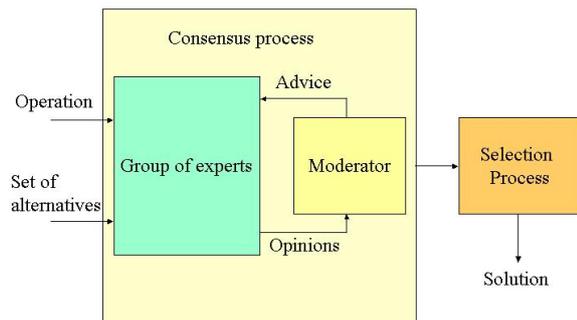
Figure 12: Group decision-making problem schematic.

5 Conclusions

The architecture of a geo-spatial TDA system based on fuzzy logic inference was described. A system prototype based on this architecture was implemented and tested with success during CWID05 [4]. The system prototype is able to process data arranged on geo-referenced regular grids or as sparse points and to produce results in polygon vector form. Even though the prototype contains a minimum set of requirements, the test demonstrates the high flexibility of the chosen architecture. In particular, the fuzzy logic approach to decision making, the fully configurable and adaptable architecture to changing scenarios (type of operations required by the end user, geographical area, different input data formats), the exploitation of vector geometry, OGC-based clients and servers for data visualization, and the possibility of stacking layers with different degrees of information content for characterising the scenario at different levels of complexity are the strong points of this implementation.

Possible improvements as well as future directions has been introduced and extensively discussed.

References

[1] Blackman, S., Popoli, R.,
Design and Analysis of Modern Tracking Systems, Artech House Publishers, August, 1999.

[2] Mendel, M., M.,
"Fuzzy Logic Systems for Engineering: a Tutorial", Proc. of the IEEE, Vol. 83, No. 3, March, 1995.

[3] Chang, F., Chen, J., C.,
A Component-Labelling Algorithm Using Contour Tracing Technique, IEEE Proc. 7th International Conference on Document Analysis and Recognition (ICIDAR 2003), 0-7695-1960-1/03, 2003 IEEE.

[4] Ranelli, P., et. al.,
Report on CWID05 NATO exercise: technical description, NURC internal report, to be published.

[5] Hall, D., L., McMullen, S., A., H., "Mathematical Techniques in Multi-sensor Data Fusion", Artech House Information Warfare Library, Artech House Publishers; 2nd edition, March 1, 2004.

[6] Open Geospatial Consortium,
http://www.opengeospatial.org.

[7] Pennucci, G., Zanasca, P., "A comparison of in situ measurements and satellite remote sensing of absorption coefficient, beam attenuation coefficient and horizontal diver visibility", NURC internal report, to be published.

[8] Schut, P., Whiteside, A., "OpenGIS Web Processing Service, Version 0.4", OGC Document 05-007r4.

[9] Turner, M., R., "Determining the context-dependent meaning of fuzzy subsets", in Proceedings of the First International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97), Rio de Janeiro, February 4-6, 1997, pp. 233-242.

[10] Errera, V., E., Martinez, L., Mata, F., Chiclana, F., "A consensus support system model for group decision-making problems with multigranular linguistic preference relations", IEEE Transactions on Fuzzy Systems, vol. 13, No. 5, Oct. 2005.

[11] Giannecchini, S., Spina, F., Nordgren, B., Desruisseaux, M., "Supporting interoperable geospatial data fusion by adopting OGC and ISO TC 211 standards, submitted to Fusion 2006.