

Ontology-based Rescue Operation Management

Stasinos Konstantopoulos¹, Jens Pottebaum², Jochen Schon³, Daniel Schneider³, Thomas Winkler³, Georgios Paliouras¹, Rainer Koch²

¹ Institute of Informatics and Telecommunications,
NCSR ‘Demokritos’, Ag. Paraskevi 153 10, Athens, Greece
{konstant,paliourg}@iit.demokritos.gr

² University of Paderborn, C.I.K., Pohlweg 47-49, 33098 Paderborn
{pottebaum,r.koch}@cik.uni-paderborn.de

³ Fraunhofer IAIS, Schloss Birlinghoven, Sankt Augustin, Germany
{jochen.schon,daniel.schneider,thomas.winkler}@iais.fraunhofer.de

Abstract The focus of this paper is ontology-based knowledge management in the framework of a mobile communication and information system for rescue operation management. We present a novel ontology data service, combining prior domain knowledge about large-scale rescue operations with dynamic information about a developing operation. We also discuss the integration of such a data service into a service-oriented application framework to reach high performance and accessibility, and offer examples of SHARE applications to demonstrate the practical benefits of the approach chosen.

1 Introduction: The Fire Fighting Domain

Mobile information technology is a key technology in the emergency-response domain, as emergency forces deployed at a disaster site have very limited access to conventional IT infrastructure. Currently, IT is limited to operation control centres in the fire stations and even IT-based emergency dispatch systems are, typically, not linked to other information systems.

As a result, on-site operation management is conducted over paper maps, magnetic boards, analogue radio, hand-written message forms, and fax. Especially for large-scale operations these tools are often not sufficient for the complex emergency management task at hand. As a result, commands processed via hand-written message forms might, in multi-level command hierarchy operations, need as much as 20 minutes to reach their destination, and staff personnel involved in decision-making lacks a comprehensive, integrated view of the various aspects of the operation.

By contrast, mobile information technology can be a valuable tool in the hands of emergency professionals to increase speed, precision, efficiency and effectiveness of their operations. Mobile IT, however, has to overcome immense entry barriers prior to its widespread use in the rescue-operation domain; most importantly immaturity and instability of the technology, higher cost of equipment, and the heterogeneity of the organizational structure of SaR departments.

The gap between the great benefit that usable mobile IT could yield in the domain of emergency response and the specific challenges for such technologies in this particularly demanding domain, results in a situation which calls for a new strategy to release this immense potential with a sustainable impact. The FP6-IST SHARE project⁴ carried out extensive research and development work in the domain, and proposed introducing a few small but extremely efficient solutions that substantially help rescue units, instead of building an overall solution that is meant to replace all existing tools at once. Another important step towards usability and robustness was the introduction of semantic representation technology to support a number of SHARE end-user applications.

The rest of this paper is organized as follows: we first provide an overview of the SHARE system (Section 2) and then describe SHARE-ODS, its central knowledge representation and management component (Section 3). We proceed to present some key SHARE applications and discuss how they capitalize on SHARE-ODS (Section 4), and finally conclude (Section 5).

2 A Mobile Information and Communication System

The SHARE system is a service-oriented application framework, supporting management tasks in large-scale rescue operations. The main applications available to rescue team members cover the following functionalities: interactive operation map (MAP), Push-to-Share communications (PTS), digital message forms (DMF), interactive resource management (IRM) and, finally, communications indexing and retrieval. Multimodal interfaces using automatic speech recognition and text-to-speech synthesis are available for most of the applications to enable hands-free access to the services offered. Figure 1 shows the graphical user interfaces of the applications. Depending on the task (strategical, tactical, operational) the appropriate application is selected by the users; depending on the user's operational role, different functionality is provided by each application, matching the role's responsibilities and access rights.

Although fire brigade personnel are highly motivated with respect to using new technologies, working with new and complex IT-based systems while heavily involved in an operation is poses extra difficulties. With this in mind, we identified the applications that are in the centre of operation management and that show a great potential as entry points for mobile IT into the domain. The two main applications we identified are the digital operation map on the one hand and the interactive resource management on the other.

Domain knowledge about rescue operations is logically represented as an ontology and accessed during the operation via an ontology data service. The main advantages of interlinking SHARE applications by an ontological knowledge base will be described in the following sections. Logical inference is mostly used in resource planning and decision making. The dynamic characteristics of ontology-based knowledge management allow a flexible configuration of group-based com-

⁴ See also <http://www.ist-share.org/>

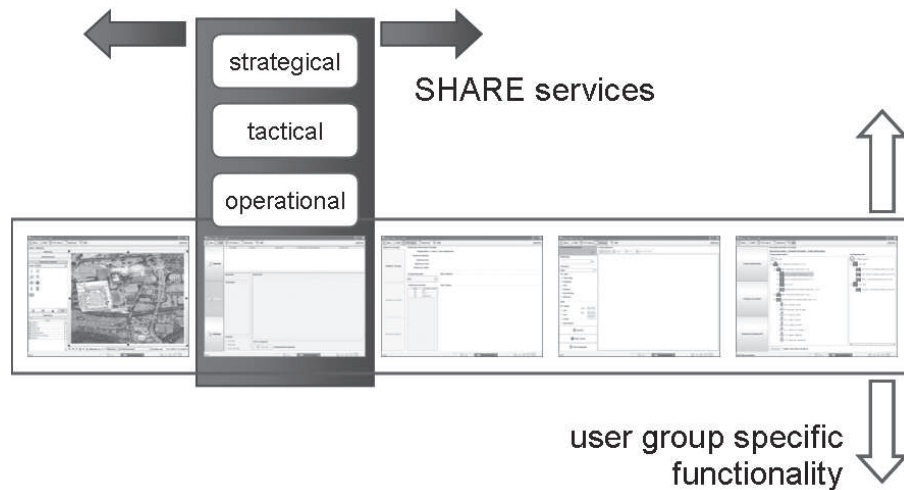


Figure 1. Usage of SHARE applications in the operation

munication. Finally, robust indexing and retrieval of multimedia communication messages benefit from the usage of an ontology in the SHARE system.

3 Knowledge Management in SHARE

The knowledge management system relies on Semantic Web technologies in order to model the operation, derive inferences from the model, and provide for the interaction between the (inferred) model and the end-user operation management applications. More specifically, the operation is modelled as an *ontology*, an abstract representation often used in the areas of knowledge representation, artificial intelligence and the Semantic Web as a way of structuring and representing knowledge.

3.1 Conceptual view: ontology structure

The elementary pieces of information in the ontology – corresponding to the individuals of the domain of discourse – are called *instances*. In the SHARE Ontology, instances correspond to officers, vehicles, units, audio and video documents exchanged, geographical sections of the operation, and so on. These instances are organized in a conceptual hierarchy where each *concept* or *class* groups together a set of conceptually similar instances.

The SHARE conceptual model comprises two main hierarchies, one involving concepts from the search-and-rescue domain and one involving multimedia objects. In addition, there are auxiliary concepts represents spatio-temporal references.

SaRThing is the top concept of the search-and-rescue hierarchy and subsumes concepts such as: Formation, OperationalRole, Personnel, Equipment, Vehicle, Station, etc. Formation, OperationalRole, and Personnel, in particular, are further specialized into A, B, and C-Level, to reflect the three-level operational structure followed in fire brigade operations. MultimediaThing, on the other hand, subsumes logical representations of the metadata of all documents (DMF, audio, and video) generated and transmitted during an operation.

Personnel instances are not directly related to Formation instances, but only an OperationalRole instance which specifies the each officer’s role in the operation. This approach allows modelling multiple roles being assigned to the same person, as is, for example, the case with A-Level staff members in smaller operations.

Both main sub-ontologies as well as the auxiliary ones are tightly integrated in a comprehensive model of the operation and extensively cross-linked using ontological *relations* between their instances. Some of the most prominent relations are *partOf* relations, which link Formation instances into a command structure and Section and Subsection geo-reference instances into a geographical decomposition of the various tasks and sub-tasks carried out. It should be noted that the geographical structure does not necessarily match the operational one: for instance, a supply or rescue-service formation might be assigned a section (area of responsibility) which overlaps with several fire-fighting formations’ sections.

3.2 Logical view: ontological reasoning

The SHARE Ontology is represented in OWL-DL [1], a web semantics representation language that is compatible with Description Logics (DL). DLs are a family of formal logics falling inside a *decidable* fragment of first-order predicate logic. The *SHOIN* description logic is of particular interest, as it is the minimal DL that covers OWL-DL [2].

DL reasoners, like Pellet⁵ used in SHARE, are used to deduce knowledge that is implicit in the model, based on explicit facts and axioms present in the ontology. In SHARE, for example, the A, B, and C-Level Formation concepts are refined into sub-concepts which include *well-formed* formations at each alarm level. Fire brigade rules and practices with respect to operational structure are represented as axioms concerning membership in the well-formed sub-class of each formation class. So, for instance, for alarm level 4 we have provided the following axioms:

$$\begin{aligned}
 \text{ALF4} &\equiv \text{ALevelForm} \sqcap \forall \text{hasProSub} . \text{BLF4} \sqcap \geq 2 \text{ hasProSub} \sqcap \leq 2 \text{ hasProSub} \sqcap \\
 &\quad \forall \text{hasResSub} . \text{MANV2} \sqcap \geq 1 \text{ hasResSub} \sqcap \leq 1 \text{ hasResSub} \\
 \text{BLF4} &\equiv \text{BLevelForm} \sqcap \forall \text{hasProSub} . \text{CLF4} \sqcap \geq 2 \text{ hasProSub} \sqcap \leq 2 \text{ hasProSub} \\
 \text{CLF4} &\equiv \text{CLevelForm} \sqcap \geq 3 \text{ hasVehicle} \sqcap \leq 5 \text{ hasVehicle}
 \end{aligned}$$

to represent the following rules about alarm level 4 operations: the A-Level Formation must have exactly 2 professional fire brigade B-Levels and 1 MANV2

⁵ See <http://pellet.owldl.com/> for more information.

rescue formation. The two B-Level Formations must have 2 full C-Level Formations, which have 3 to 5 vehicles each.

Given these axioms, an operation's compliance with alarm level 4 guidelines is checked by logically verifying that all `Formation` instances of the operation are subsumed under `ALF4`, `BLF4`, or `CLF4`, depending on the formation level they belong to. Compliance with the rest of the alarm levels is checked in a similar fashion.

3.3 Technical view: embedding the ontology

One of the main concerns while designing the system architecture was how to make the knowledge, stored in the ontology, available and useful to the client applications. By analysing the requirements of those applications towards the knowledge base it was possible to identify a set of hierarchically structured services called SHARE-ODS.

3.3.1 Hierarchical Service Structure The SHARE Ontology Data Service [3], is a set of comprehensive data and knowledge services for the SHARE system. It provides access to the ontology and the reasoner's conclusions through web services, presenting a specifically customized API populating, updating, and querying the SHARE knowledge base to each client application. SHARE-ODS provides additional extra-logical functionality, like logging and client-specific composite functionality. Composite functionality groups together commonly-recurring service calls into a single call, providing a high-level interface to the ontology.

A high-level API might not only add, but also restrict functionality by blocking access to low-level functions. For instance, the API presented to the retrieval application permits querying about document meta-data and related operational structure, but does not permit updating the knowledge base. The PTS API, on the other hand, provides a method for adding new audio and video documents and for relating them to SaR instances, but does not permit the addition of new units; structural changes in the operation (committing and relieving units, moving units around) can only be performed through the IRM API, and so on.

This also provides limited access-rights management capabilities, by presenting to each client an API which offers only functionality matching the access rights allocated to the user of the client application. In order that a fire officer can use the client application he has to be assigned a functional role within the regular workflow of the operation. Starting the client application the officer is authenticated by the SHARE system; the person is matched to the role assignments. Each role may execute a different set of applications, affecting (a) the available top level use cases, (b) the information displayed and (c) the functionalities and their options. Actions out of a role's responsibility are not available as application options, and the officer filling the role has to request such actions following the chain of command.

3.3.2 Synchronization of critical information In the course of an operation, new facts about the operation and its various agents get added to the SHARE-ODS knowledge base. Some of this information is critical as it has a direct impact on the system configuration or need to be reacted upon immediately. Examples for this kind of information are changes in the communication group structure or updates of the operation map.

The challenge in this context is efficiently and timely notifying client applications that they need to synchronize their internal information structures with the ODS. In the first SHARE prototype this was approached with *polling* in fixed time intervals. This approach proved sub-optimal, as short intervals were very resource demanding and had a negative impact on the overall response time of the system. Longer intervals, on the other hand, did not honour the time critical nature of the information.

This was addressed by replacing polling with a synchronization mechanism based on the *Java Message Service* (JMS). JMS provides for *topics*, thematically separated messaging systems, where JMS clients participate as subscribers, publishers, or both. In SHARE, topics correspond to client applications (i.e. MAP-Topic, IRM-Topic a.s.o) that potentially need to react to the update. Any application that updates the knowledge base will subsequently send message to the appropriate topic, so that subscribed clients know to contact SHARE-ODS to also update their internal information structures. Following this approach information is timely updated without any noticeable delays in the system response time.

4 Ontology-Driven Services

Several services in SHARE interact with the ontology via the SHARE-ODS. Generally, these services provide the ontology with information gained from user input or extracted from media data, and vice versa request collected and derived data from the ontology. Hence, the decision support, the group-based communication and the media indexing and retrieval described in this section are enabled or enhanced by some means or another using the SHARE Ontology.

4.1 Decision Support for Resource Management

In the field of emergency response, resource management tasks are critical to the efficiency of an operation: commanders base their decisions on the availability of resources, the operational and geographical situation, and the various regulations and practices governing fire-brigade operations. These three aspects must correspond to each other, thus the hierarchy of the command must be congruent to the geography of the operations theatre and in accordance to fire-brigade rules.

The most important resources are personnel and vehicles, which might arrive at the operation site already organized into small units or might be allocated to units on-site. At the time of their arrival, they are available, but *non-operative*

resources. Before being deployed, units might be organized into larger B or C-level formations, in which case they are *prepared resources*. To become *operative*, non-operative, loose units might be allocated to existing formations, or prepared formations might be wholesome deployed.

Especially for larger operations, this involves a complex coordination job where formations are organized, assigned a commanding officer, assistants and staff, and integrated into the command and communication structure. The coordination of this structure is the foundation of the documentation of an operation and an important instrument concerning fire-brigade regulations. The fundamental use cases for a resource management tool is to view and manage personnel, vehicles, and other equipment available on-site and support starting, escalating, de-escalating, and terminating fire brigade operations, relieving and replacing units or formations, and re-allocating resources; all abiding by the established command structure patterns permitted by the current alarm level (cf. Section 3).

Interactive Resource Management (IRM) [4] is one of the most important functionalities of the SHARE system, presenting resources to the A-level staff in an intuitive and usable way and supporting them in:

- command & control, offering comprehensive and cross-indexed information about the operation;
- visualizing the operation through tactical symbols on a digital map, including annotations about available resources and cross-referencing geographical and operational items;
- quickly and easily retrieving past communications.

Furthermore, of great benefit is information derived from the relations between the various decision-support tools: different applications can update each other and refer to a consistent knowledge base, complementary data pools help build more complete view, and implicit information can be logically inferred from explicit facts. Fire-brigade practices regarding operation structuring and resource allocation and utilization can be easily implemented (cf. Section 3.2 above) and the interplay between geo-references and resource-based structural relations more easily managed [5].

4.2 Dynamic Configuration of Group-based Communication

One of the core components of the SHARE system is the *push-to-share* (PTS) voice and video communication system. PTS replaces currently-used radio devices which are limited to broadcasting with a group-based communications system, where audio and video messages are only received by automatically-inferred recipients. This functionality is an imperative requirement for emergency operations, as it retains the simplicity of push-to-talk while at the same reducing the clutter of irrelevant messages caused by radio broadcasting.

The PTS system establishes communication groups according to the command-and-communications structure of the operation. At the beginning of an operation, the IRM tool is used to setup the initial command structure and assign units

and personnel to formations. SHARE-ODS infers the communication groups from the command structure and logical rules modelling actual fire brigade communication practices. PTS devices query SHARE-ODS for communication group membership information, completing the initial configuration.

Changes in the command structure during the operation (unit or personnel re-allocations, operation escalation or de-escalation, etc.) are also performed through the IRM application. SHARE-ODS re-calculates communication groups to comply with the new command structure, and the PTS system is accordingly re-configured. IRM messages the PTS clients which are directly or indirectly connected to those changes, i.e., are in the same communication group, and the PTS system is reconfigured. Ongoing voice conferences at the time of the update are, naturally, preserved regardless of whether the relevant communication group exists in the new configuration.

4.3 Robust Media Indexing and Retrieval

Retrieval of voice and text communications is a necessary functionality during a rescue operation, but also important for later analysis of large-scale operations or training events. While querying DMF archives is straightforward, automatic speech recognition (ASR) techniques are needed to spot and index keywords in the voice communications. Rescue operations, however, cover very diverse acoustic environments, with a wide range of noise of varying in type and level; a challenging environment which significantly decreases the performance of ASR.

In SHARE, indexing relies on detecting keywords in unconstrained speech, based on a holistic statistical paradigm integrating various knowledge sources such as acoustic, language and pronunciation models. A phoneme-based garbage model [6] is used to detect a set of keywords after applying Wiener Filtering to reduce noise. Training is performed on domain-specific speech data collected during exercises of the Dortmund fire brigade. Acoustic models have been adapted to the commanding tone of the fire fighters, local accent, and the transmission characteristics of the headsets used in SHARE [7].

To further improve the accuracy of the recogniser and increase the robustness of the indexing, minimizing the vocabulary of the ASR is of paramount importance. SHARE-ODS supports ASR by predicting a vocabulary which is as small as possible without lacking important keywords, based on the hierarchical command structure of the operation. This list includes vehicle radio names, personnel names and roles, section codes, and street names that are relevant to the operation. As a rescue operation is a dynamic process, operating units and their subordination as well as responsibilities for sections and streets might change continuously, hence the keyword list for each PTS device user will be dynamic.

SHARE-ODS infers the abstract `SARThing` instances (operational roles, formations, units, etc.) that are pertinent each PTS communication session, based on operation-structural relations between the members of the session and the rest of the operation. `PTSSessions` instances are related to the inferred `SARThing` instances at the time of addition of each `PTSSessions` to the ontology; in this

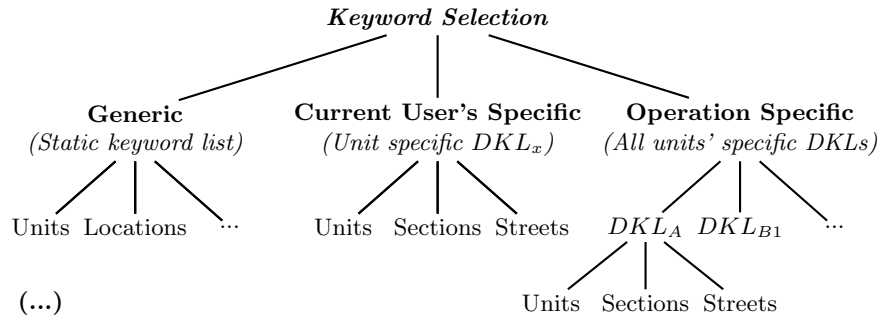


Figure 2. Example for clustering keywords for a guided keyword selection

manner these relevance relations persist subsequent structural changes in the operation, so that SHARE-ODS provides a keyword list which reflects the status of the operation at the time when the PTS session took place.

This list of keywords is finally merged with a small static list of user independent, generic keywords and relevant geographical keywords (street and place names) extracted from the map application. The combined list is used as vocabulary for the ASR module, significantly improving keyword spotting results.

Using one uniform keyword list for all possible commanding units would require a list of several thousand words, as all possible street names, personnel names, vehicle radio names, and section specifiers would have to be covered. Generally, it is easy to see that the number of words in a shared static keyword list will be larger than the number of words in the dynamically-created keyword lists used in SHARE.

SHARE-ODS dynamic keyword lists are also used to provide user-specific clustering and selection for information retrieval. Queries for keywords in a voice communication can be formulated in several ways: a single keyword can be chosen from a list or typed directly into the query field; but also groups or subgroups of keywords—e.g. *all units* and *all sections* or the units and the sections subordinated to a specified user—can be selected. In this latter case, dynamic keyword lists provide a intuitive selection of keyword sets that are very likely to match a user's intentions.

In Figure 2 an example of a hierarchical keyword structure is presented, guiding the user to groups and subgroups of keywords. Generic keywords and specific (dynamic) keywords for the current user are the main clusters. For higher commanding levels and for operation and training analysis it is also possible to access the dynamic keyword lists of all involved commanding units.

The dynamic keyword list for each unit is managed and updated by the ontology. The ontology regularly provides each client application with the update of all dynamic keyword lists which can be selected for a query. So it is assured that every keyword which has been used for indexing during the operation can also be used for retrieval.

5 Conclusions

In this paper we discussed SHARE, a powerful platform, integrating communication and information services to support large-scale rescue operations. After introducing the domain and its challenges, we discussed the SHARE ontology data service (ODS) and

We have also shown how SHARE-ODS acts as the central information and knowledge management for SHARE, extending the explicitly stored information and allowing the seamless coupling of a variety of heterogeneous applications. This coupling provided the basis for new features, but also allowed for the modular design of SHARE, which we have argued to be an important factor for the introduction of IT technologies in the mobile response domain.

To demonstrate the above, we have presented three SHARE applications (resource management, communications, and multimedia indexing/retrieval) and discussed the considerable benefits of SHARE-ODS support, a discussion further supported by the evaluation of the SHARE system at Dortmund Fire Brigade exercises.

References

1. Smith, M.K., Welty, C., McGuinness, D.L., eds.: OWL web ontology language. W3C Recommendation (February 2004), <http://www.w3.org/TR/owl-guide/>.
2. Horrocks, I., Patel-Schneider, P.F., van Harmelen, F.: From *SHIQ* and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics* **1**(1) (2003) 7–26
3. Konstantopoulos, S., Paliouras, G., Chatzinotas, S.: SHARE-ODS: an ontology data service for search and rescue operations. Technical Report DEMO-2006-1, NCSR ‘Demokritos’, Athens (January 2006)
4. Pottebaum, J., Konstantopoulos, S., Koch, R., Paliouras, G.: SaR resource management based on Description Logics. In Löffler, J., Klann, M., eds.: Proc. of 1st International Workshop on Mobile Information Technology for Emergency Response (Mobile Response 2007), Bonn, Feb. 2007. LNCS 4458, Berlin/Heidelberg, Springer (2007)
5. Löffler, J., Schon, J., Hernandez-Ernst, V., Pottebaum, J., Koch, R.: Intelligent use of geospatial information for emergency operation management. In van der Walle, B., Burghardt, P., Nieuwenhuis, K., eds.: Proceedings of the 4th International Conference on Information Systems for Crisis Management (ISCRAM 2007), Brussels, Academic and Scientific Publishers NV (2007)
6. Wilpon, J.G., Rabiner, L.R., Lee, C.H., Goldman, E.R.: Automatic recognition of keywords in unconstrained speech using HMMs. *IEEE Transactions on Acoustics, Speech and Signal Processing* **38** (1990) 1870–1878
7. Schneider, D., Winkler, T., Löffler, J., Schon, J.: Robust audio indexing and keyword retrieval optimized for the rescue operation domain. In Löffler, J., Klann, M., eds.: Proc. of 1st International Workshop on Mobile Information Technology for Emergency Response (Mobile Response 2007). LNCS 4458, Berlin/Heidelberg, Springer (2007)