

Topic Maps For Improving Services In Disaster Operations Management

Christopher W. Zobel, Virginia Polytechnic Institute & State University, USA

G. Alan Wang, Virginia Polytechnic Institute & State University, USA

ABSTRACT

Disaster operations management is an increasingly important application area for the developing techniques of service science. This paper examines the use of topic maps, a semantic technology, within this environment, and provides a preliminary discussion of the benefits that its implementation can provide in the capture and exchange of contextual information. The discussion is motivated by a look at the different phases of disaster operations management in a services context, and focuses on the need for effective and relevant information exchange as an important part of the services process. As the amount and complexity of information increases within such processes, semantic technologies are becoming increasingly important as a means representing and managing contextual information. This paper seeks to help further the understanding of the relevance of such tools as part of the study of service science.

INTRODUCTION

Whether natural or man-made in origin, disasters play an ever more significant and visible role in today's increasingly global society. Managing the effects of such disasters is an ever more important service provided both by professional emergency personnel and by volunteers and volunteer-based organizations like the Red Cross, and service science has a very important role to play in helping to improve this process. By studying how analytical tools from science, engineering, and mathematics can improve the effectiveness of disaster operations, service science researchers have the opportunity not just to reduce the costs and improve the efficiency of such operations, but also to have a real impact on people's lives and livelihoods.

In particular, the ability to share information in a timely and contextually relevant manner is crucial to providing effective service delivery within the domain of disaster operations management. Different organizations (both governmental and non-governmental) frequently need to work together before, during, and after the occurrence of a disaster, and information sharing is critical for supporting the multi-user and multi-task decision making in this collaborative environment (Fuerth & Baker, 1997). Although there are well-recognized managerial considerations associated with coordinating such information (Mayer-Schoenberger, 2002; Lund, 2002; Timmons, 2007), there are also technical approaches available for making the exchange of relevant information more effective. This paper discusses the use of data semantics in this environment, focusing on *topic maps* as an example of a semantic technology that can represent the context of information and thus help to support the process of information exchange. The general idea of topic maps for service science was first introduced in Zobel (2007); the discussion in this current paper is more specifically focused on their application to services in disaster operations management.

We begin our discussion with an in-depth look at disaster operations management as a particularly important application area for the study of service science. We then discuss the idea of semantic knowledge exchange in the context of this services environment, and provide a look at the characteristics of topic maps which make them an important tool in this situation. As part of this discussion, we propose and describe several examples of situations in which the use of topic maps to represent contextual knowledge can be of great benefit to improving disaster management services. Finally, we offer some conclusions and suggestions for further work.

DISASTER OPERATIONS MANAGEMENT

Disaster operations management (DOM) can be defined as the management of "activities that are performed before, during, and after the occurrence of a disaster with the goal of preventing loss of human life, reducing its impact on the economy, and returning to a state of normalcy" (Altay & Green, 2006). This choice of terminology is more accurate than the related term *disaster management* in that it acknowledges managing the activities associated with the disaster (or emergency, or crisis), and not managing the disaster itself. The term *disaster* is typically used to refer to a situation which has escalated out of control, and thus under these conditions one can only truly manage the situation through appropriate decision-making and by taking corresponding actions.

As discussed in the recent report by the National Academies (Rao, Eisenberg, & Schmitt, 2007), there are typically considered to be four different phases associated with the managing the effects of a disaster: mitigation, preparedness, response, and recovery (Figure 1). Each of these phases differs in terms of the time-frames involved and the urgency with which information is needed, but they all share a need for building the most complete and accurate picture of a disaster situation as it currently stands, so that appropriate and timely services can be provided. The following discussion provides a brief overview of each of these phases from a services standpoint.

Phases Of DOM

Disaster mitigation is the process of designing and implementing procedures for reducing the risk associated with the occurrence of a disaster, typically by reducing either the likelihood or the impact of a potential disaster event (Ridge & United States Dept. of Homeland Security, 2004). Typical examples of disaster mitigation practices include strengthening the foundation of a bridge and establishing redundant command and control systems. The process of defining and instituting a disaster mitigation plan can be characterized as a service operation in that continuous maintenance and knowledge-sharing between the solution provider and the client are often necessary in order to establish and maintain the effectiveness of a solution (Zobel, Martin, & Olgun, 2008).

Disaster preparedness measures are aimed at improving the expected capabilities of response and recovery efforts in advance of an actual disaster event, often by providing relevant information to help individuals and organizations prepare themselves for its occurrence (Rao et al., 2007). This can involve such services as setting up and maintaining emergency shelters in advance of a storm, and broadcasting information about traffic patterns or emergency evacuation routes. As illustrated by the May 2008 cyclone in Myanmar, neglecting to provide such services, whether for political, logistical, or historical reasons, can have catastrophic implications (Casey, 2008).

In contrast to the previous two phases, disaster response operations typically take place only during a disaster event and in its immediate aftermath; such operations typically seek to contain the event and reduce property loss, and to minimize loss of life and injuries (Rao et al., 2007). Service providers involved in this stage of the process would typically include local emergency personnel such as firefighters or public safety officials, but may also include regional, national, and even international organizations. The types of services provided at this stage would include such things as emergency medical services, evacuation services, and immediate post-disaster relief.

The final phase of DOM, disaster recovery, is the process of minimizing the long-term effects of a disaster situation and facilitating restoration to conditions that are as good or better than those in place before the disaster occurred (Rao et al., 2007). Recovery typically begins immediately after urgent health and safety issues have been addressed, and it is the phase towards which the largest share of costs in the disaster management process are directed (DITF, 1997). With this in mind, it can be very important to incorporate mitigation practices into recovery and rebuilding efforts in order to avoid incurring the same losses in future disasters. Humanitarian relief services are considered an important part of disaster recovery, and even more so than in the response phase the potential involvement of a large number of different governmental and non-governmental organizations can lead to a spectrum of different information sources and needs to be addressed.

Services In DOM

Although there are many different definitions of “service” in the literature, most of them include some notion of an activity, process, or interaction which creates value both for the customer and for the service provider (Chesbrough & Spohrer, 2006; Dietrich & Harrison, 2006; IBM, 2006). Dietrich and Harrison (2006) provide a good overview of some of the characteristics of service systems that help to differentiate them from traditional manufacturing or supply chain problems. The perishable nature of resources is one such characteristic. As opposed to a standard inventory problem, in which a resource often may be set aside for future use, resources in a service system may lose all value to that system after a certain point in time. This can be particularly true in a disaster operations management context, since the success of delivering services such as evacuation procedures or emergency medical care depends a great deal on when and how quickly they are provided.

Resources in a service system also may have changing capabilities, such as a Red Cross volunteer who receives additional training or a search and rescue team that is provided with new equipment. Dynamically substituting one such resource for another in certain situations may be not only possible but also desirable, depending on the relevance of those capabilities (Dietrich & Harrison, 2006). Particularly in a disaster environment, where an evolving situation may be associated with a great deal of change and uncertainty, the flexibility to make such substitutions may have a significant impact on the effectiveness of the services provided.

Communication and the exchange of knowledge are also important aspects of many services processes (Chesbrough & Spohrer, 2006; Tien & Berg, 2003). Partnerships in a disaster situation may often be ad hoc and dynamic in nature, and it can be important to have a framework in place, either technical or policy-related, within which to share relevant knowledge about available resources and capabilities. Longer-term service relationships, such as those established in the disaster mitigation stage, have a particular need for extended knowledge exchange (Zobel et al., 2008). The levees in New Orleans that failed during Hurricane Katrina are an excellent example of this. Although they were originally built to protect the population that existed at the time of construction and to guard against the known disaster risks at that time, there was no long-term analysis provided in support of re-evaluating and updating the design of the levees (Christian, 2007; USACE, 2007). The failures occurred in part because the inspection and maintenance program that was implemented did not appropriately account for changes in the initial assumptions. More active involvement and communication may have been able to ensure that these changes were recognized and that the level of protection provided by the levees was modified accordingly.

The exchange of knowledge in a service transaction, however, is not always simple and straightforward. Although some knowledge may be easy to articulate, in that it is well-structured and easily described, like the physical characteristics of an ambulance or a helicopter, other types of knowledge are recognized as being more tacit in nature (Hitt et al., 2001). Such tacit knowledge can be difficult to describe explicitly, such as the complex business realities that can lead to investing in certain mitigation practices. The complexity of such tacit knowledge tends to complicate the exchange of information when a service is being provided, and it often makes it difficult for each party to understand the other’s requirements and qualifications (Chesbrough & Spohrer, 2006).

This problem of exchanging complex, and often tacit, knowledge is an important one that needs to be addressed, both for short-term relationships and for longer-term services arrangements that are associated with managing the effects of a disaster. One step towards accomplishing this is to explicitly capture the relative context within which information and knowledge exists and is understood, and to present it in such a way as to be meaningful and accessible. Expressing such context can help to codify tacit knowledge and thus can help to support more effective information exchange within a services relationship. The following discussion explores the idea of capturing and expressing the context of such knowledge in more detail, and provides a detailed look at *topic maps* as an information technology-based technique that has a number of important characteristics in support of this objective.

SEMANTIC KNOWLEDGE EXCHANGE

In Euzenat et al. (2008), the authors argue that well-accepted standards should be the basis for the representation of contextual information, in order to ensure interoperability across various service providers. They also point out, however, the importance of recognizing that the context of information is often dynamic in nature, rather than static (Euzenat et al., 2008). This can be seen in complex environments such as DOM, where new types of context information may appear over time and are not well anticipated at the initial stages of defining a representation (Coutaz, Crowley, Dobson, & Garlan, 2005). In establishing standards, therefore, not only should they be capable of handling the new contexts, but they should relate new information to the existing knowledge so that all parties can take advantage of the entire body of knowledge.

There are two primary groups associated with developing standards to represent the context of information and semantic knowledge exchange: the computer science community, and the information and library sciences community. The Semantic Web (SW) is a term coined by Tim Berners-Lee, the Director of the World Wide Web Consortium (W3C), to refer to (Web-based) documents that not only incorporate machine-readable content, but also incorporate the semantics through which that content can be put into context by automated software applications (Berners-Lee, Hendler, & Lassila, 2001). In the context of knowledge representation within the Semantic Web, the computer science community has focused on a semantic framework known as RDF (Resource Description Framework) in support of this vision.

RDF, which became a W3C Recommendation in 1994 (Manola & Miller, 2004), provides an underlying syntax and structure for describing simple relationships among objects, by representing such relationships in Extensible Markup Language (XML) format as "triples" of subject, object, and associating predicate. The ultimate intent of RDF is to provide a framework which will enable the information in a document to be automatically processed by a computer (McGuinness & van Harmelen, 2004). The computer science community has invested a significant amount of research into realizing this capability, and there are a number of related technologies, such as RDF Schema (RDFS) and the Web Ontology Language (OWL), which depend on and extend RDF in order to support automated reasoning techniques developed in the Artificial Intelligence (AI) community (W3C, 2004).

Topic maps, which were developed separately by the information and library science community, provide a different (and, in many ways, complementary) approach to representing contextual knowledge. Formally classified as an ISO standard (ISO, 2002), rather than as a W3C Recommendation, topic maps were originally created in order to provide support for the semantic indexing of documents. The subsequent development of an XML-based syntax for the topic maps standard, XML Topic Maps (Pepper & Moore, 2001), allows them also to be implemented within a web-based environment.

Topic maps play a similar role to that of RDF, in that their purpose is to provide a means for representing concepts and the relationships (context) between them, and both topic maps and RDF provide abstract yet powerful knowledge representation techniques. As discussed by Pepper (2002), one of the major differences between the two initiatives lies in the origin and perspective of the developing community. Whereas RDF is focused on supporting automated information management, topic maps are more focused on knowledge representation for humans. As a result, RDF is very streamlined and efficient, but topic maps contain a richer array of semantic elements in their abstractions and can be more intuitive for humans to interact with. For example, in contrast to RDF, topic maps allow more than two topics to be simultaneously associated with one another (Pepper, 2002). This can be useful for representing situations where complex, ad hoc working relationships can arise between different organizations, such as during disaster response and relief operations.

Pepper (2002) views topic maps and RDF as complementary technologies, rather than as competitors, and there are several ongoing research initiatives (e.g., RDF/Topic Maps Interoperability Task Force) that are seeking to bridge the two technologies and combine their advantages. From the standpoint of emergency and disaster operations management applications, which often have a high degree of human involvement in the decision-making process, topic maps are a more human-centric approach to knowledge representation than RDF, and we believe that there is significant value in looking at them in more detail. With this in mind, the following discussion provides a

more in-depth look at topic maps, beginning with a general overview of the idea and moving to a discussion of the characteristics which can have a significant impact on the effectiveness of various aspects of the emergency management process.

TOPIC MAPS

Pepper's "The TAO of Topic Maps" (Pepper, 2000) is an excellent introduction to the structure and implementation of topic maps. A typical topic map consists of *topics*, which represent individual concepts within a given domain; *associations*, which represent the relationships between each of those topics; and *occurrences*, which are separate sources of information (like a book, an image, or a web page) that describe a topic in some way. A given topic can also have *facets*, or properties, which can be assigned values to more completely characterize the topic. Figure 2 provides a small portion of a topic map related to the rescue and relief operations of the 2008 earthquake in China that includes each of these elements. In this topic map, topics and associations are represented using ovals and circles respectively, and facets are represented using rectangles.

The underlying basis for using a topic map to represent knowledge is the unique identification of each individual topic in the topic map with a text string (frequently in the form of a URL) called a Published Subject Indicator (PSI). Two topics that share the same PSI are considered to be identical, whether they coexist in the same topic map or are contained within separate topic maps. When two different topic maps share at least one topic, as identified by a common PSI, those topic maps can be automatically merged to create a larger, and potentially more comprehensive, topic map. This provides a very powerful approach for merging the contents of two disparate knowledge bases by simply merging the topic maps that represent each of them.

Although identifying and merging identical concepts can be fairly straightforward when discussing physical resources such as ambulances or relief supplies, the intangible aspects of a service process can be difficult to clarify and thus all topics may not be equally susceptible to merging because of uncertainty that two similar topics actually refer to the same concept. Because a topic map explicitly represents context, however, through its associations and occurrences, it provides a means by which the relative perspective of each contributor can be captured and represented locally. Even though automatic merging of such topics (and their associated concepts) may not be possible without additional discussion, the ability to represent these varying perspectives, together with the process of clarifying the differences in perception, can provide valuable insight into the services relationship.

An important consequence of the origins of topic maps in document indexing is that a topic map exists independently from the knowledge that it represents. Although identifying the occurrence of a given topic within a data source, such as a new type of rescue vehicle, may necessitate that information about that vehicle be added into the topic map, it does not require that any changes or additions be made to the original data source to indicate that a topic map is associated with it in some way. This can be very important from the standpoint of implementing topic maps within a new organization which already has an established information infrastructure, since it requires no changes to the format or composition of the original documents or data stores. This separation also allows topic maps to be used to combine and simultaneously represent information that exists in different formats and in different locations. For example, a training video, a web site, and a corporate database could all be linked together by a single topic map that represents the relationships between the individual data items (and the knowledge they represent) from each source (See Figure 3).

In practice, the generation and exchange of knowledge using topic maps is simplified by the fact that topic maps, as with most tools associated with the Semantic Web, have a standardized formal representation in XML format. This allows them to be exchanged and shared using standard internet protocols, and avoids the need for developing complicated and expensive interfaces between existing information systems. It also allows the exchange to take place within the context of existing security procedures and policies, since all documents being exchanged are simply text files which can easily be encoded or encrypted as necessary.

Topic Maps For Disaster Operations Management

There are several characteristics of topic maps that make them particularly useful in the context of providing disaster-related services. For example, because of time pressure and the need to interact with a large number of different individuals and organizations when responding to a disaster, new information about that disaster may often become available only a little bit at a time. Because a topic map is able to represent just the subset of topics and occurrences that are known, and because it is relatively simple to merge existing topic maps, the knowledge associated with a given portion of a disaster, or with a given perspective on that disaster, can be gathered in independent topic maps and then easily merged with a dynamically growing map of the larger situation. Although the associated higher-level task of defining and merging new topics and relationships would be very difficult in such a disaster response situation, the flexibility afforded by the modularity of the topic map specification would also support this level of distributed and ad hoc knowledge generation during a more time-extended process such as disaster mitigation.

From a slightly different standpoint, this ability to combine knowledge in a piece-wise fashion can also be very useful for training purposes in this environment, particularly for an organization such as the International Red Cross which relies heavily on a volunteer workforce. Given the high turnover rate of such organizations and the resulting loss of acquired knowledge, a tool such as topic maps could be beneficial in its ability to help gather and combine expertise from distributed locations, and to represent and characterize the organization's response to different events. When incorporated into a web-based learning tool, a topic map would allow for contextually-relevant searches for existing information from the organization's knowledge base, while supporting revisions and additions as new situations are encountered and new knowledge is developed.

Because it can be important to identify the provenance of individual pieces of knowledge in a merged topic map, the concept of *scope* is also a relevant and important aspect of the topic map specification. Particularly if there are differences between the knowledge contributed by multiple individuals or groups, assigning a given scope to subsets of that knowledge can help to better organize them and to clarify the context within which they were provided. For example, in a disaster relief situation, information about potential resource availability may come from both registered aid organizations and from local groups or individuals. It can be important to differentiate between the two sources of information on many levels, but particularly from the standpoint of both information accuracy and the level of understanding of local cultural needs. Similarly, as a dynamic knowledge base is developed in a piecemeal fashion, new information will likely need to be validated before being permanently added to the main body of knowledge. Scope could be used to manage such control, without preventing the information from being made available for use on a provisional basis.

Due to the often international nature of disaster operations, such as those that take place in Europe where there is significant cross-border collaboration, it is also significant that the topic map specification supports assigning a given topic multiple descriptive names, potentially from several different languages. Since it is the Published Subject Indicator that uniquely identifies a topic, the name by which it is represented in a particular topic map does not reduce its generality. This allows a topic map associated with flooding risks that is generated in France to be automatically merged with one on the same topic that is generated in Italy, assuming only that there is a common underlying ontology (a specification of the set of terms and relationships) that is shared by the two, and that there is agreement on the underlying PSI values that are used for each topic.

Although, as discussed, human interaction and discussion may be necessary to clarify the understanding of the topics and their context as it exists, particularly if different languages are represented, the process of storing and merging knowledge from multiple topic maps can be automated behind the scenes. To further support decision making, many existing topic map software applications take advantage of the fact that topic maps have a natural visual representation as a series of connected nodes of various types. The ability to visualize such information graphically can make it much simpler to understand the relationships being described, which can be of significant benefit in an environment such as disaster management services which transcends technical ability, cultural differences, and socio-economic boundaries.

CONCLUSIONS

Our brief discussion has concentrated primarily on the use of topic maps for capturing and representing the context of knowledge. Although the ability of topic maps to make information more relevant and accessible is one of its greatest strengths, the topic map specification also supports such extensions as data validation (via the Topic Map Constraint Language (TMCL)), and the derivation of new knowledge through inferencing (Rath, 2001). Like many technologies associated with the Semantic Web, topic maps are still in the beginning stages of being widely accepted by the general public. There are, however, a number of different companies that create advanced software platforms for authoring topic maps (See, for example: (Ontopia, 2008), (Mondeca, 2008), (NetworkedPlanet, 2008), and (Empolis, 2008)), and thus a fair amount of potential support exists for applying them to new interdisciplinary, service-driven areas such as disaster operations management.

Within the disaster operations management community, there are a number of initiatives underway, including a W3C incubator group (W3C, 2007) and a strong sub-community within the larger Information Systems for Crisis Response and Management (ISCRAM) community (ISCRAM, 2008), that are focused on the problem of developing formal ontologies in support of semantic information representation and exchange. The development and use of topic maps will benefit a great deal from this work because it has the potential to lead to a standard set of terms and relationships for a variety of different disaster and emergency-related situations. This would support the ability to automatically generate a topic map from an ontology, and thus to more easily and consistently define the initial knowledge upon which decision-making ultimately can be based. Topic maps, in turn, will be able to provide a user-friendly representation of an ontology that can be examined, compared, and updated by the wide range of people associated with a particular disaster operations management process. Because techniques to improve communication are vitally important in order to share knowledge, particularly in a services environment, we believe that technologies such as topic maps, which leverage the semantics associated with human decision-making, can play a significant role in enabling and supporting such knowledge transfer.

REFERENCES

1. Altay, N., & Green, W. G. (2006). OR/MS research in disaster operations management. *European Journal of Operational Research*, 175(1), 475-493.
2. Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. *Scientific American*, 284(5), 34-43.
3. Casey, M. (2008) Why the Cyclone in Myanmar Was So Deadly, accessed May 15, 2008, [available at <http://news.nationalgeographic.com/news/2008/05/080508-AP-the-perfect.html>].
4. Chesbrough, H., & Spohrer, J. (2006). A research manifesto for services science. *Communications of the ACM*, 49(7), 35-40.
5. Christian, J. T. (2007). Lessons from Hurricane Katrina. *The Bridge, Quarterly Journal of the National Academy of Engineering*, 37(1), 5-11.
6. Coutaz, J., Crowley, J., Dobson, S., & Garlan, D. (2005). Context is Key. *Communications of the ACM*, 48(3), 49-53.
7. Dietrich, B., & Harrison, T. (2006). Serving the services: The emerging science of service management opens opportunities for operations research and management science. *OR/MS Today*, 33(3), 42-49.
8. DITF (1997). Harnessing Information and Technology for Disaster Management: The Global Disaster Information Network. Washington D.C.: Disaster Information Task Force Report.
9. Empolis (2008), accessed June 1, 2008, [available at <http://www.empolis.com/>].
10. Euzenat, J., Pierson, J., & Ramparany, F. (2008). Dynamic Context Management for Pervasive Applications. *The Knowledge Engineering Review*, 23(1), 21-49.
11. Fuerth, L., & Baker, D. J. (1997). Harnessing Information and Technology for Disaster Management. In *The Global Disaster Information Network* (Ed.).
12. Hitt, M. A., Bierman, L., Shimizu, K., & Kochhar, R. (2001). Direct and Moderating Effects of Human Capital on Strategy and Performance in Professional Service Firms: A Resource-Based Perspective. *The Academy of Management Journal*, 44(1), 13-28.
13. IBM (2006) Services Sciences, Management and Engineering, accessed September 30, 2006, [available at <http://www.research.ibm.com/ssme/>].

14. ISCRAM (2008) Information Systems for Crisis Response and Management, accessed June 1, 2008, [available at <http://www.iscram.org/>].
15. ISO (2002). ISO/IEC 13250 Topic Maps.
16. Lund, D. A. (2002). Learning to Talk: The Lessons on Non-Interoperability in Public Safety Communication Systems. Durham, NH: University of New Hampshire, JusticeWorks
17. Manola, F., & Miller, E. (2004) RDF Primer, accessed June 1, 2008, [available at <http://www.w3.org/TR/REC-rdf-syntax/>].
18. Mayer-Schoenberger, V. (2002). Emergency Communications: The Quest for Interoperability in the United States and Europe. Cambridge, Mass.: John F. Kennedy School of Government, Harvard University.
19. McGuinness, D. L., & van Harmelen, F. (2004) OWL Web Ontology Language Overview, accessed June 1, 2008, [available at <http://www.w3.org/TR/owl-features/>].
20. Mondeca (2008), accessed June 1, 2008, [available at <http://www.mondeca.com/>].
21. NetworkedPlanet (2008), accessed June 1, 2008, [available at <http://www.networkedplanet.com/technology/topicmaps/intro.html>].
22. Ontopia (2008), accessed June 1, 2008, [available at <http://www.ontopia.net/>].
23. Pepper, S. (2000) The TAO of Topic Maps: Finding the Way in the Age of Infoglut, accessed June 1, 2008, [available at <http://www.ontopia.net/topicmaps/materials/tao.html>].
24. Pepper, S. (2002) Ten Theses on Topic Maps and RDF, accessed June 1, 2008, [available at <http://www.ontopia.net/topicmaps/materials/rdf.html>].
25. Pepper, S., & Moore, G. (2001) XML Topic Maps (XTM) 1.0, accessed June 1, 2008, [available at <http://www.topicmaps.org/xtm/>].
26. Rao, R. R., Eisenberg, J., & Schmitt, T. (Eds.) (2007). *Improving Disaster Management: The Role of IT in Mitigation, Preparedness, Response, and Recovery*. Washington, DC: The National Academies Press.
27. Rath, H. H. (2001). Semantic Resource Exploitation with Topic Maps. GLDV-Spring Meeting 2001. Giessen University. 3-15.
28. Ridge, T. J., & United States Dept. of Homeland Security (2004). *National Incident Management System*. [Washington, D.C.]: U.S. Dept. of Homeland Security.
29. Tien, J. M., & Berg, D. (2003). A case for service systems engineering. *Journal of Systems Science and Systems Engineering*, 12(1), 13-38.
30. Timmons, R. P. (2007). Interoperability: Stop Blaming the Radio. *Homeland Security Affairs*, 3(1).
31. USACE (2007). Performance Evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System - Final Report of the Interagency Performance Evaluation Task Force (IPET). *US Army Corps of Engineers*, Volumes I-IX, March 2007.
32. W3C (2004) Resource Description Framework (RDF), accessed June 1, 2008, [available at <http://www.w3.org/RDF/>].
33. W3C (2007) Emergency Information Interoperability Framework Incubator Group, accessed June 1, 2008, [available at <http://www.w3.org/2005/Incubator/eiif/>].
34. Zobel, C. W. (2007). Topic Maps for Services Science. Proceedings of the 37th Annual Meeting of the Southeast Decision Sciences Institute. Savannah, GA.
35. Zobel, C. W., Martin, J. R., II, & Olgun, C. G. (2008). Disaster Risk Management for Critical Infrastructure: A Services-Based Viewpoint. *International Journal of Services Sciences - Special Issue on Disaster Operations Management (forthcoming)*.



Figure 1: The lifecycle of disaster operations management

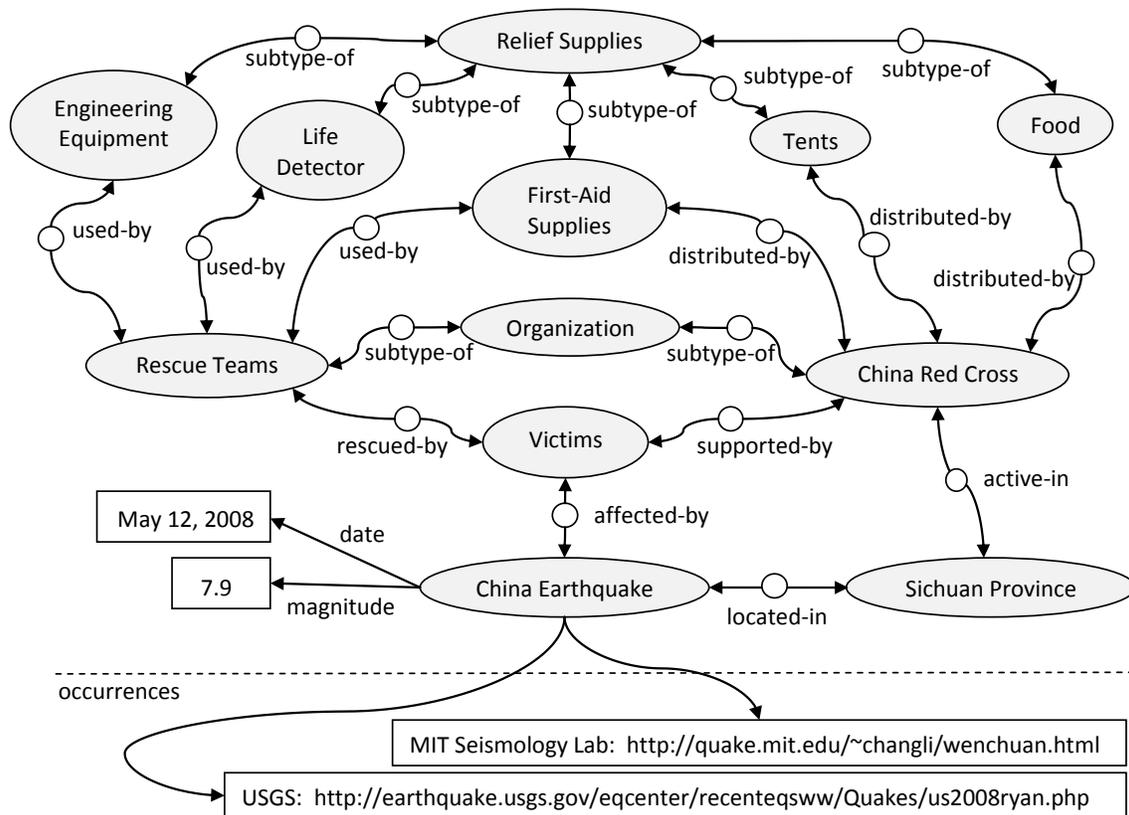


Figure 2: Example topic map – 2008 earthquake in Sichuan Province, China

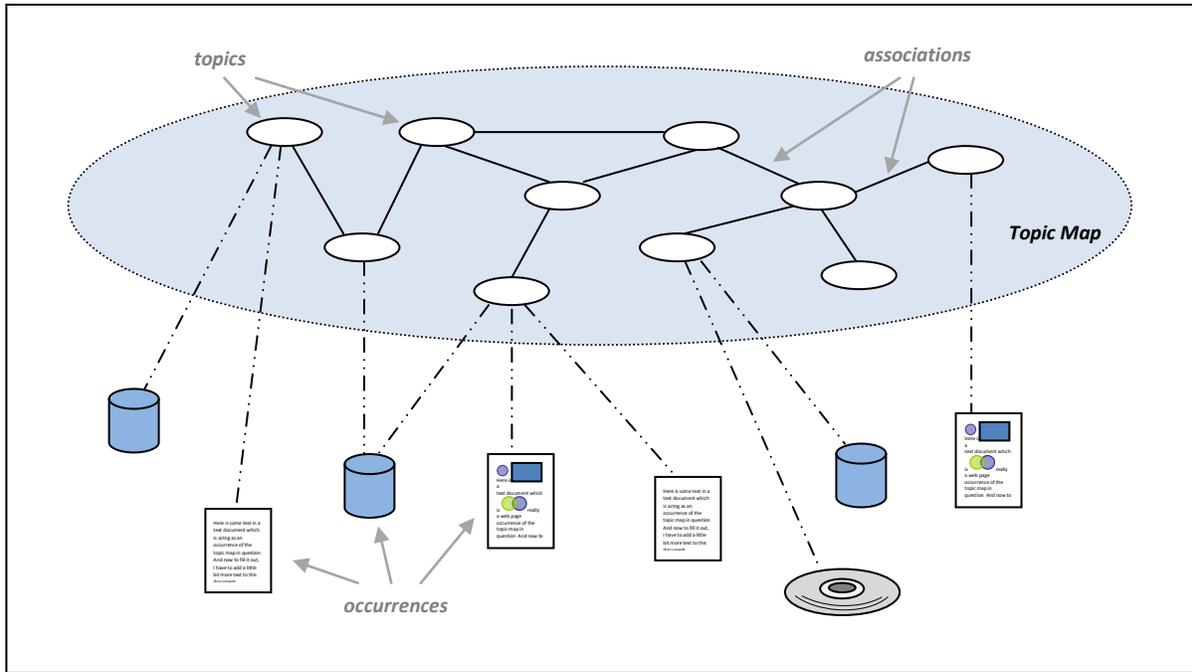


Figure 3: Topic map with multiple occurrence types