Improving Situational Ontologies to Support Adaptive Crisis Management Knowledge Architecture

Gary Berg-Cross, Ph.D.
Engineering, Management & Integration
Semantic Technology
Herndon VA 20170
gbergcross@gmail.com

ABSTRACT

There is considerable interest in advance technologies to support crisis and disaster management as they face the challenges of designing, building, and maintaining large-scale distributed systems able to adapt to the dynamics and complexity of crises. Candidate technologies include Service Oriented Architecture (SOA), related Semantic Web technology, agent-based architecture and cognitive architectures. Each embodies some principles of the Adaptive Architecture – including modularity, openness, standards-based development, runtime support and importantly explicitness. However, truly adaptive architectures for crisis management will require some deepening the knowledge architecture's content and not just its representation. Light and more robust ontological models of situations are discussed to show how better formalization of conceptual patterns like "participation" can be developed to support cognitive architectures. The feasibility of an ontological design pattern approach is described as an avenue for future research and development describing specific types of situations.

Keywords

Adaptive architecture, ontologies, crisis situations, design patterns.

INTRODUCTION

It is now recognized that workflow approaches, using predefined process logic is inadequate for the inherent complexity and dynamics needed to cope with extreme events of disaster, crises and related emergency & relief situations. As noted by Harrald (2006) "Extreme events present unforeseen conditions and problems, requiring a need for adaptation, creativity, and improvisation while demanding efficient and rapid delivery of services ." It is understandable then that in the aftermath of events like hurricane Katrina there was considerable interest in advanced technologies to support more adaptive crisis and management including emergent Public Health (PH_ needs for integrated PH information or proving early warnings to health and safety officials about potential bioterrorism attacks. Some notable, but technically easy efforts, focus on user interfaces while other efforts combine technologies coupling grid-based emergency planning & crisis response to provide emergency planners/responders timely forecasts(Gadgil et al, 2004). Other approaches apply advanced knowledge technologies from AI and Semantic Web research (Potter et al 2007) for crisis response. Many design and development challenges remain to improve on the static, reactive nature of current applications leading to large-scale distributed, enterprise-wide systems that are able to adapt to the dynamics and complexity of crises in a timely manner providing real-time visibility into disaster and crises situations. There is is hope that new architectural research opportunities are presented by emerging architectures designed but truly adaptive architectures will require some particular advancements and detailed work in knowledge architectures as well as processing architectures. My discussion of the issues is divided into four parts. First I summarize some maturing IT technologies which can work together to support adaptive functionality. In the second section, a hybrid cognitive architecture is described to frame some of the knowledge content needed to populate systems capable of crisis situation understanding. A third section discusses extant models of situations and how concepts like "participation" can be formalized in situational ontologies to support a cognitive architecture. The paper concludes by summarizing the feasibility of an improved knowledge architecture approach and describes how modular patterns may be used for practical research and development.

MATURING TECHNOLOGY TO SUPPORT ADAPTIVE ARCHITECTURE

Several related technologies and architectures are being leveraged to build more agile systems. This maturation grows out of earlier research investigating core concepts and the framing of critical questions from which preliminary applications have clarified underlying ideas, generalized approach to flesh out important for system solutions (Redwine and Riddle, 1985). Two interesting approaches are to build systems that include some agility in their architecture (such as SOA) and those based on enough intelligence to make systems adaptive in a knowledgeable way. A principle example of the first, from broad mainstream of candidate technologies, is Service Oriented Architecture (SOA) and related Semantic Web technology, while the 2nd is represented by cognitive architectures and agent-based architecture[1]. Each architecture offers some advantages and together they embody various aspects of agile architecture including an emphasis on modularity, openness, and standards-based development supported by the increased use of formal models. These aspects can be seen in service architectures which offers a degree of "agility" using explicit, standards-based descriptions of service interface. This requires building in and enabling a degree of flexibility so changes in service use or re-implementation do not require modification of how a requester invokes the service . Currently most service architectures have important agility limitations because they are built on middleware using standards primarily for efficient transport and orchestration of messages rather than handling different message content. For real agility independently designed systems must be able to exchange and coordinate meaningful information, not just data, across application areas as well as coordinate the processes that link them together. Better informational "semantic representation" standards are emerging from Semantic Web work, for software agents. This has lead to improved semantic representations using RDF (Resource Description Framework**)** and a formal ontology language called OWL. RDF triples and some ontology models have been used in emergency response systems as discussed by Potter et al (2007) and to co-ordinate medical responses (Bloodsworth and Greenwoodl, 2005) . As discussed later, improved "semantic" representation does not always equal improved semantics, since semantics relies on the content of models and not just its representation.

Multi-agent-based systems provide an adaptive basis for information fusion technology through an architecture that models collaboration between agent teams. When combined with perceptual and reasoning models of mobile intelligent agents, these characteristics match up well against disaster management requirements (Buford et al. 2006). Generally, multi-agent architectures include adaptive infrastructures that can be used to make problem solving responsive to real-time constraints, available network resources, and coordination requirements. They also often include specialize components for adaptive functions such as local agent scheduling, multi-agent coordination, organizational design, detection and diagnosis, and on-line learning. These can be configured to interact so that a range of different situation-specific coordination strategies can be implemented and adapted as problem "situations" evolves (Sims, 2007).

Making such systems situation/context-aware systems is an important step towards greater agility. This work involves fusing sensor-based data into situational information (Nicholas et al, 2006) . In these multi-agent architectures sensor agents are responsible for an initial degree of processing, mid-level agents associate these inputs with object concepts, while still "higher" agents are responsible for assembly into "situations". Automation of "situation assessment" has proven increasingly important for complex diverse designs (Endsley and Garland, 2000). Some systems generate operating pictures using precision geospatial environment information layers (modifiable digital overlays) that can support decision-making based on detailed "knowledge" shared by the agents sensing a physical environment. An example of such a system called SAPPHIRE (Situation Awareness and Preparedness for Public Health Incidents using Reasoning Engines) which utilizes OWL and RDF(S) to represent medical taxonomies and emergency room data from eight hospitals in the Houston metropolitan area (Mirhaji 2006). SAPPHIRE provides syndrome classification, information visualization and cross-domains navigation of information and demonstrated real adaptability to changing situations, in response to hurricane Katrina. A modified SAPPHIRE dynamically absorbed new data feeds to make new interpretations and produced new results (Mirhaji 2006).

COGNITIVE ARCHITECTURE AND SITUATIUONAL KNOWLEDGE

A third type of adaptive architecture is called Cognitive Architectures because it uses sets of cognitive mechanisms and a larger pool of structured knowledge thought to be typical of human cognition. The history of cognitive architectures use in disasters and emergency response starts with "intelligent" robots to provide assistance as well as support for intelligent simulations to understand situations. Cognitive based analysis is illustrated in MacEachren et al (2006), which developed an understanding of cognitive processes involved in crisis management work. Use of a

[1] Besides the Semantic Web new levels of autonomy are emerging in infrastructures like the Grid computing and pervasive computing environments which are not covered here.

realistic cognitive model enabled practical distribution of information access in support of geo-semantic interoperability for mobile devices and map-based web portal to support international humanitarian relief logistics. More recently, real applications have been developed using computational models by Mendonça and Wallace (2007) supporting improvisation in a simulated emergency response situation. This application combines declarative knowledge, represented in an ontological form, and procedural knowledge in the form of decision logic. The overall cognitive architecture, is implemented with ACT-R, (Anderson et al. 2004) which leverages advances in machine learning, distributed agent technology and semantic representation. Cognitive architectures improve on general architecture like SOA which often operate without an explicit cognitive model. Lack of explicit cognitive models exist in applications monitoring filtered information streams, where it is important to understand the factors that humans use to select subsets of task-relevant information. Such cognitive tasks should include explicit commitments to help make service solution more adaptable. When combined with agents producing suitable input, cognitive, agent architecture could be used to reason about "crises situations" where a large number of inter-dependent, dynamic objects change their states in time and space, and engage each other in fairly complex relations. Cognitive architectures like ACT-R can serve as an adaptive higher agent but need a provide proper knowledge base for the system to handle the complex situations previously described since ACT-R's machine learning has never dealt with such a large domain of knowledge. No machine learning based system is allowed to run for childhood lengths of time to become as intelligent is real world concepts as a child. While rule-based knowledge can be learned by a cognitive agent like ACT-R, these can't be across cognitive agents, in coordinated fashion unless there are common "concepts/terms" in the rules. Typically, discourse about disaster situations will generate explicit content that presupposes a larger pool of implicit, higher-level, "background" information needed to understand the situation properly. For example, we know that a "bus stop" is a place to wait for scheduled bus transport and that buses use roads that are made of concrete, blacktop, etc. However, there is implicit background information that is crucial for making relevant inferences and becomes particularly relevant in disasters. We know that roads can be washed out and that road makeup in low elevations and nearness to streams are important factors in determining if a road is "useable". For proper reasoning background information must be made explicit to support the inferences needed to support adaptive queries such. Ontologies can provide explicit representations of this type of higher-level background information, and are thus crucial to the realization of the adaptive systems in general and Cognitive Architectures in particular. But it is the content of the ontologies not just the representational formalism that is important. Multi-agent systems, like COSMOA (Bloodsworth and Greenwoodl, 2005) which uses ontological knowledge, are a useful development, but the actually quality of the ontology content is more important than its representational formalism[2]. These issues and the characteristics of quality ontologies are discussed in the next section starting with the use of light-weight models/taxonomies and how they may be formalized in foundational ideas, heterarchical concepts, distinctions between parts and components as well as modular patterns that can be reused over time as articulated in the SOA vision.

SITUATIONAL UNDERSTANDING MODELS AND ONTOLOGIES

To understand ontology models for disaster situations it is useful to discuss two levels needed for situational knowledge – mid-model patterns of situations and more foundational models that captures the participatory aspects of situations. The initial aspects of a mid-level ontology can be illustrated by the Situation Awareness (SAW) "model" of Matheus et al (2003). SAW is a "light" model that captures some of the core elements and relations needsed to describe situations. SAW provides a principled way to depict crises and disasters situations . using two primary classes **-SituationObjects**, and **Events** (Figure 1). The organizing point of the ontology is the concept of **Situation** defined as a 3-fold relationship of Goals, SituationObjects and Relations. Goals, obviously provide a top-down, rational basis of situational understanding starting with an agent's goals & objectives that guide its reasoning about events, relations and objects. A Goal might be to evacuate a section of the city using buses for people without cars. Situation Objects are entities in a situation that participate in Relations and can have Attributes such as Volume, Position and Velocity - characteristics of PhysicalObject in the model. Some dynamics are handled by object attributes, such as expected/unexpected weight, color etc. that can change as events unfold. We can think of richer situational characteristics to represent bus transportation systems (a **Physical Object**) which in a **Flood Situations** are being reported on by environmental sensors. For this situation we need to include characteristics such as realistic street information such as Bus Stops (**Physical Location**). In a flood more abstract entities like timetables become critical along with circumstances that are salient in emergencies (e.g. sensors reports on water

[2] In this case it is hard to tell since what knowledge is captured in the ontologies since it is not specified in the publication.

height, closed streets). Geospatial information is important in disasters and in SAW these are represented as first-order objects using geometrical and positional concepts as in the example above dealing with streets.

While the SAW ontology was not designed specifically for disaster situations, it is applicable and includes more than an ontological hierarchical representation of some situationObjects for emergency response (ER) as discussed in Mendonça et al. (2007). That ER ontology focuses not on **Situations** but **Resource** as a top-level concept with 4 major sub-types- vehicle, boat, equipment and personnel. However, there are several ontology engineering ideas that could improve on such hierarchies needed for extreme events starting with generalizing the idea of "vehicle" and "boat" to make them more useful across a variety of situations. Rather than treating these as separate classes or categories it would be useful to recognize them as "transportation devices". This not only allows us to include the ground transportation (ambulance, pumper truck, bus and gravel truck) and air transport (helicopter) listed in Mendonça et al.'s ER model but also allows for "water devices" boat and many other items that might be appropriate to different circumstances. By adding a deeper conceptualization and categorization of domain content using distinguishing properties we achieve a more systematic relations between ontological levels. This moves the ontology beyond a mere representational formalization (Frank, 2006) to reflect a concentration on the nature of the information content rather than on how to represent information . Rich representation without the conceptual analysis needed to produce adequate content does not solve the problem of semantics. It merely makes them

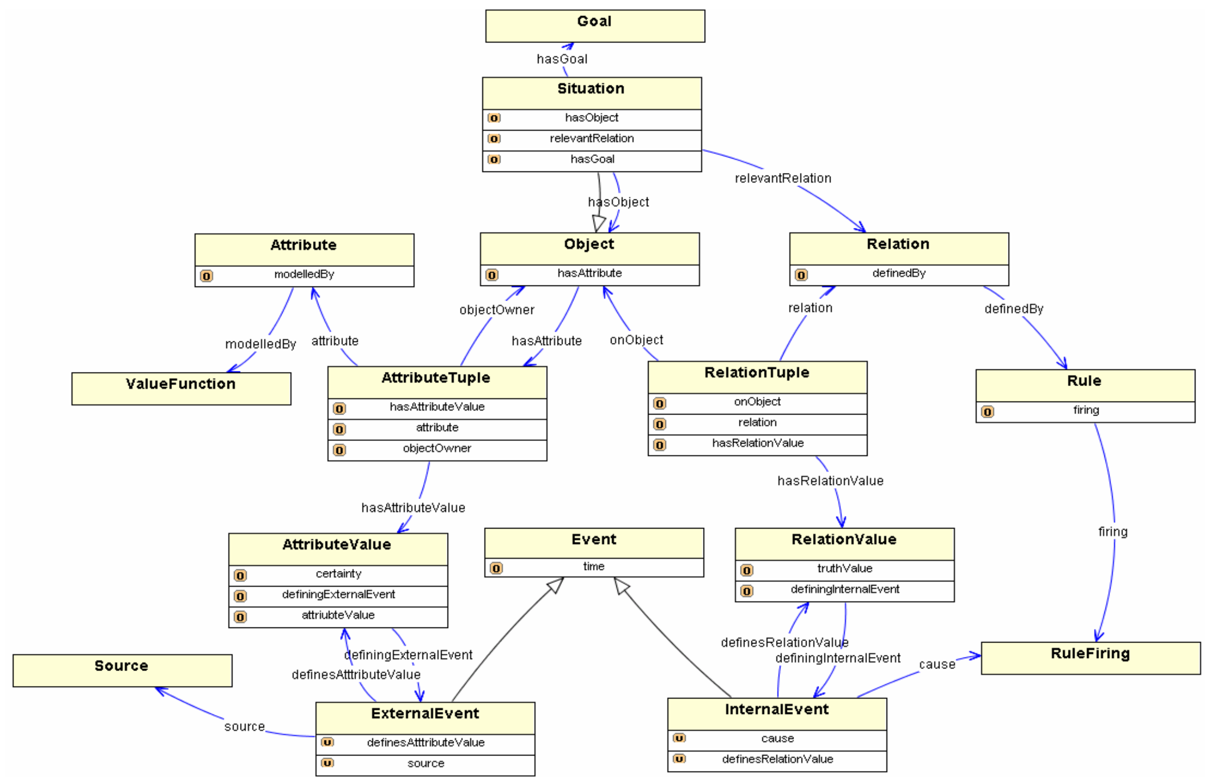

Figure 2. A Light Situation Awareness Ontology (after Matheus et al. 2003)

formal in the sense that they be put into a form amenable to automated processing. The COSMOA ontology-centric multi-agent system uses ontologies in this sense, but does not make them meaningful in the sense that all named classes can have instances, or that they correctly capture intuitions of domain experts can stand up to rational analysis[3]. Lack of ontological quality remains one of the ongoing issues for the semantic web and SOA models

[3] Other desirable characteristics are that ontologies be minimally redundant with no unintended synonyms and sufficiently axiomatized with detailed constraining descriptions for the design goals of the work.

which use its representational standards. This is true because shared representation (language) is not a guarantee of shareable content. Issues like this make it clear to ontologists that even partial standardization of the content of semantic web ontologies will be an order of magnitude more difficult than standardization of ,an ontological language like OWL. Adequate “semantics” remains a problem in part because most ontology development is done by computer scientists/engineers who are not sufficiently trained in ontological analysis. A simple example is illustrated by considering an improved “transportation device” ontology in which land-water and air are not the only important ontological distinctions. We can also distinguish motorized devices from manual devices, such as bikes (land) and row-boats (water). The order in which we make these distinctions produces differently looking hierarchies each of which becomes more relevant in certain situations. But humans know both and act as if they have merged two such hierarchies into a heterarchy. Unlike a hierarchic system, which starts from a single root, humans see a heterarchic system with many root points sending out branches (Frank, 2006). It is this type of content model, rather than a simple controlled vocabulary that is needed for more adaptable knowledge.

Special roles are important in crises but are modeled differently in particular ontologies. “Role” is a property in Mendonça et al’s ontology but it treated more as a relation in SAW. Thus in Mendonça et al.’s ontology a particular resource like “pumper truck” may play the role of “fire department”. Intuitively, we can say that social concepts like role are like properties, but for rigorous modeling they should be treated as first class citizens within an ontological framework. In SAW a social role can be established as a relation rather than a property. This seems a more flexible way to model changing roles. One can see how to do this by adapting the SAW model so that Property Values are distinguished from Attributes and mediated by Relations. Still better for representing the management aspects of disaster management is integrating mid-level ontologies into foundational concepts. An example of this is the ontological modeling of social roles using a conceptualization of them as a subclass of social concepts. The Ontoloclean methodology (Masolo et al. 2004) distinguishes social role by 2 features –“anti-rigidity” and “foundation” that can be represented as axiomatic commitments. Anti-rigidity expresses the fact that roles have dynamic properties on objects so that for the period of time an entity is classified with the role concept, there exists some time at which the entity is present but not classified under the concept (Bottazzi & Ferrario, 2006). This is used in a later example. Foundation, on the other hand, is a property used to show the relational nature of roles as previously mentioned. Thus the model states that a concept like X “bus unable to move” is founded when its definition involves one or more concept such that “for each entity classified by X, there is an entity classified by Y which is external to it (Bottazzi & Ferrario, 2006).

Something missing in both the ER and SAW models is adequate composition of part-wholes which is so important to SOA. As a start we need to make better distinction between parts of physical units (finger is a proper part of hand) and geographic regions (Maryland is contained in the USA). This is an important distinction because containment is NOT parthood. For example, we can say that Amber is part of the Maryland disaster management team and Amber’s head is part of Amber. But Amber’s head is NOT part of the Maryland disaster management team. To handle these many of these distinctions something like the SAW ontology needs to be improved by explicitly extending it, as we did for the transportation devices with more explicit relations and grounding it in a more foundational ontology where these distinctions are axiomitized. A good candidate is DOLCE developed within the WonderWeb Project (Masolo et al, 2003). DOLCE is a cognitively based, “reference” ontology, consisting of about 30 classes, 80 properties and a much larger set of axioms. It is designed to provide a sufficiently broad base to map, integrate, and build domain ontologies, such as an improved SAW ontology. DOLCE includes the idea of a high-level participation pattern for objects taking part in Events in the SAW model but conceptualizes endurants (Objects or Substances) and perdurants (Events, States, or Processes) as distinct types linked by the relation of “participation”. This is illustrated by a response team (Perdurant) participating in disaster management (endurant event). General participation patterns help us understand the structure of repeating events that occur for types of situation. The general pattern in Figure 2 expressed in a version of UML that indexes the temporal location of an event at a time interval, while the respective spatial location use the idea of “space region” that is provided by the participating object. We can illustrate this portion of participation pattern by an example which specializes the “participates-in” relation between objects and events, with location for objects and a time indexing for events. An easy to read propositional structure similar to UML is used to represent the information, with relations treated as formal properties. Objects have parts and in our example a “response team’ is part of the object is a “personA” who might play the role of nurse. To be explicit the proposition states that nurse constantly participate in a response.

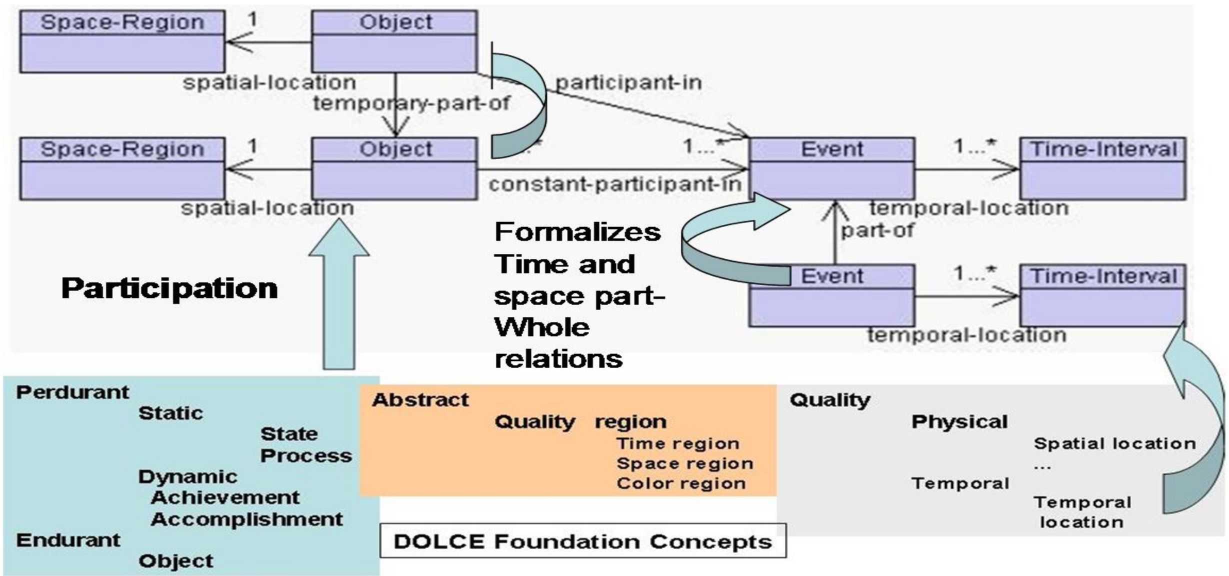

Figure 2. Ontology Design Pattern for Participation in a Situation (Based on Gangemi et al. 2004)

[object: care team]->(participates-in)->[event: care]-(Part)-[event: nursing]- (duration)->[time-interval@2days]

[object:care team]-(part)->[object: personA]->(location)-[space-region: hospital]

-(constant-participates)->[event: care]-(Part)->[event: nursing]-(loc)->[space-region: bedside]

This captures a simple intuition that as a refinement we can reify the Figure 2 abstract participation pattern features into a kind of "situation" which Gamgemi et al (2005) calls "time-indexed-participation". That is, participation, formerly a relation in our examples, is now reified as an object. This provides a "unit" setting for exactly one object, one event, and one time interval. From this unit reification pattern others can be made as complex as needed, by adding parameters, more participants, places, etc. Thus we can make the SituartionDuration 2 days and the setting "bedside". Such explicitness provides a degree of robustness beyond typical lightweight domain ontologies.

Ontology Design Patterns

The situation participation pattern discussed above in one of example of a class of models called Ontology Design Patterns (ODPs) by Gangemi et al (2004). ODPs are based on the idea that some parts of general entities defined in foundational ontologies are meaningfully interrelated with some concepts from specific domains. ODPs bring together foundational models and lighter models. They can be used to capture abstract ideas such as the relationship between abstract descriptions and situations instantiated by data. The combination can be used adaptively by hybrid agents using abstract reasoning and lower level rules. Methodologically one might say a "situation" is observed and conceptualized so that domain experts can reuse this pattern of ontology pieces as a scalable piece of an intended model. To do this the pattern is expressed in a simpler fashion, so as to organize subsequent modeling work. Because an ODP is part of a more complex model (in the previous Figure as part of DOLCE's foundation) it implicitly represents and maintains the larger, original "context" on which the conceptualization was based. An attractive feature is that ODPs improve flexible, informal schemes in the light of foundational notions (such as DOLCE's Perdurant, Abstract, Quality etc.) and can be applied to intuitive, broad conceptualizations such as "participation" previously illustrated. ODPs allow us to start with informal schemas and formalize these without sacrificing flexibility. Adding Role and Task concepts to the situation pattern allows us to see these and develop the distinction between descriptions and situations (D&S) of these situations using "concepts". The ontology pattern shown in Figure 3 allows for the temporary roles that objects can play, as well as the task executions that events afford. In our previous example we expressed the knowledge that the "object" PersonA provided care constantly. Here a role as part of the pattern allows us to specify that Person A is in the nursing role. Most importantly, however, the ODP includes the conceptualization that descriptions exist for "situations" (such as the participating situation). This is a way that satisfies formal semantics. Thus, role and task are not confused as entities, relations or attributes but are concepts used to describe Person. This is directly useful to knowledge supporting cognitive architecture. Sensor information can be represented by event and object instances shown in the lower half of Figure 3, while high level "concepts" are represented separately in the top half of the diagram. This is a very general pattern with 2 parts that together specify a match between situations, such as "water is 2 feet above floor stage", and descriptive knowledge of these. Such D&S ontologies (Gangemi & Mika, 2003) are based on a conceptualization that supports a first-order manipulation of descriptive objects (such as emergency plans, diagnoses, norms, institutions, etc.) -in effect theories and situations (such as cases, facts, settings)[4]. D&S's explicit committed conceptualization is an important distinction between an unstructured world (some might say context) and an intentional description that recognizes (some say constructs) a structure (situation) in that world or context. One nice thing about the ontological commitment using D&S is that it supports organizing domain theories for areas like

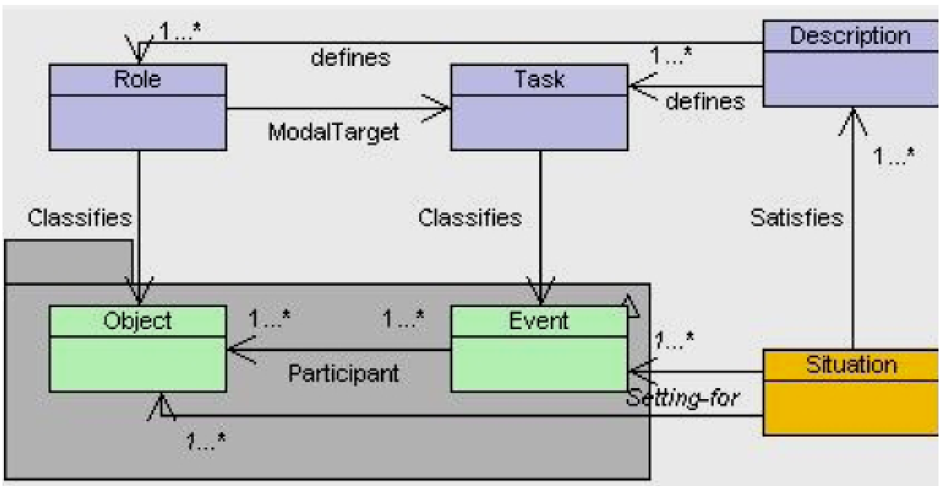

Figure 3. Situations relating Roles, Tasks and Descriptions (Based on Gangemi et al. 2005)

disaster management into different ontologies as well as into different descriptions or situations. This provides a good degree of flexibility to knowledge content relating different situations that may be related. For example, "a flood situation" is a disaster entity whose conceptualization is realized in a flood ontology module, but whose conceptualization is also realized in other modules - disaster, transportation, hydrology, geography and as a health emergency concept. This shows a general way in which the situations may be related. In the terms of the D&S diagram above flood as a condition satisfies disaster as a situational diagnosis. Thus the D&S conceptualization adds important distinctions often confused in simpler SOA, semantic web and agent technologies.

DISCUSSION AND CONCLUSIONS

I have discussed several practices that can be employed to move informal knowledge towards more formal, rigorous versions in support of adaptive systems. The "semantic" issues of adaptive architectures supporting emergency situations have been presented in terms of situational ontologies and patterns. Several particularly salient factors have been identified to support development and use of high quality information derived from sources including sensor inputs. Simple examples of improvements include the use of foundational ontologies, adequate part relations and more complex heterarchies reflecting alternative classification schemes. ODPs, in particular, may help provide an intuitively acceptable, modest effort and practical way to leverage broader and more abstract formal ontologies to support particular adaptability requirements. ODPs should be a good way to integrate heterogeneous information flows involving, multiple formats, multiple database schemata and communication protocols. The strategy of building ontologies such as espoused by Bloodsworth and Greenwood (2005) offers the hope of producing an integrated knowledge layer to support reasoning about various heterogeneities, but the value offered here is how to achieve more explicit semantic content in a stepwise, modular fashion. ODPs support this goal by leveraging intuitive conceptualizations, such as depicted in SAW capturing core notions of domains. Reuse of ODPs supports more efficient knowledge gathering efforts in general and across

[4] When D&S plugs into DOLCE o description are a non-physical endurant & situation are added as a top level.

specific types of disaster and emergency situations as discussed previously. Finally ODPs support modular use such as enshrined in SOA principles. By "extracting" fragment of foundational ontologies as "background" ODPs leverage a piece of a formal ontology without having to re-conceptualize all of them or take all of their parts. Putting these together ODPs may start from foundations and from informal or simplified, intuitive schemes which are be built up incrementally by domain experts using a methodology for domain ontology analysis. Informal schemes become more formal by specializing existing, abstract ODPs that use part of related to a foundational ontology. There remain many technical challenges to enable systems to do such advanced adaptive things as reconfigure themselves, without interruption, as their environment changes, support graceful degradation and error recovery at different application level. However, the knowledge content level will prove every bit as interesting an area to support adaptive systems.

REFERENCES

1. Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. Psychological Review, 111(4), 1036–1060.
2. Berg-Cross, Gary Fu, Wai-Tat and Kwon Augustine (2007) "A Cognitive-based Agent Architecture for Autonomous Situation Analysis" , Performance Metrics for Intelligent Systems (PerMIS) Workshop, NIST.
3. P. Bloodsworth and S. Greenwood, COSMOA: An Ontology-Centric Multi-Agent System For Co-ordinating Medical Responses To Large-Scale Disasters, extended paper in the European Journal on Artificial Intelligence - AI Communications, 2005.
4. Bottazzi, E. and Ferrario, R. (2006) 'Preliminaries to a DOLCE Ontology of Organizations', Int. J. Business Process Integration and Management, Vol. 1, pp.64–74
5. Buford, J. Jakobson G. and Lewis, L. (2006) Multi-Agent Situation Management for Supporting Large-Scale Disaster Relief Operations**.** INTERNATIONAL JOURNAL OF INTELLIGENT CONTROL AND SYSTEMS, VOL.11, NO. 4, DECEMBER 2006, 284-295.
6. Endsley M. and Garland. D. (2000) (Eds.) Situation Awareness Analysis and Measurement, Lawrence Erlbaum Associates: Mahwah, New Jersey.
7. Frank, A. (2006) "Distinctions Produce a Taxonomic Lattice", : Edited by B. Bennett and C. Fellbaum Proceedings of the Fourth International Conference Formal Ontology in Information Systems (FOIS).
8. Gadgil, H., Choi, J., Engel, B., Fox, G., Ko, S., Pallickara, S. and Piercem, M. (2004) Management of Data Streams for a Real Time Flood Simulation Community Grids Projects Technical Report June.
9. Gangemi, A. & Mika, P. (2003). "Understanding the Semantic Web through Descriptions and Situations." Proceedings of the DOA/CoopIS/ODBASE 2003 Confederated International Conferences. LNCS 2888. Springer Verlag, 2003.
10. Gangemi A, Catenacci C, Battaglia M. (2004) "The Inflammation Ontology Design Pattern: an Exercise in Building a Core Biomedical Ontology with Descriptions and Situations", in Pisanelli D (ed.), *Biomedical Ontologies*, IOS Press, Amsterdam, 2004
11. Harrald, John R. (2006) Agility and Discipline: Critical Success Factors for Disaster Response, *The ANNALS of the American Academy of Political and Social Science*, Vol. 604, No. 1, 256-272.
12. MacEachren, A. M., Cai, G., McNeese, M., Sharma, R., & Fuhrmann, S. (2006). GeoCollaborative Crisis Management: Designing Technologies to Meet Real-World Needs. 7th Annual National Conference on Digital Government Research: Integrating Information Technology and Social Science Research for Effective Government, San Diego, CA, May 21-24, 2006.
13. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., and Oltramari, A. (2003). Wonderweb deliverable d18. Technical report, CNR.
14. Mendonça, D. and W.A. Wallace (2007). "A Cognitive Model of Improvisation in Emergency Management." *IEEE Transaction on Systems, Man and Cybernetics (Part A)* 37(4) 547-561.
15. Mendonça, D. Hu, Yao & Gu, Qing, (2007) Cognitive-level Support for Improvisation in Emergency Response, Intelligent Human-Computer Systems for Crisis Response and Management**,** ISCRAM May 13-16.

16. Mirhaji, P. (2006), SAPPHIRE: A Semantic Web Solution for Public Health Preparedness, Semantic Technology Conference, March, 2006.
17. Potter, S. and Kalfoglou, Y. and Alani, H. and Bachler,. Buckingham S and M., and Carvalho, R and Chakravarthy, A. and Chalmers, S. and Chapman, S. and Hu, B. and Preece, A.. and Shadbolt, N. and Tate, A. and Tuffield, M. (2007) The Application of Advanced Knowledge Technologies for Emergency Response. In *Proceedings 4th International Information Systems for Crisis Response and Management (ISCRAM 07) Conference.*
18. Redwine, S. and Riddle, W. (1985) Software technology maturation. *Proc 8th International Conference on Software Engineering*, May, pp. 189-200.
19. Sims, M.; Corkill, D.; and Lesser, V. (2007) Knowledgeable Automated Organization Design for Multi-Agent Systems UMass Computer Science Technical Report 07-43 August 2007. Submitted to JAAMAS.