

Using Event Representation and Semantic Enrichment for Managing and Reviewing Emergency Incident Logs

Symeon Papadopoulos
Informatics & Telematics Inst.
57001, Thessaloniki, Greece
papadop@iti.gr

Ansgar Scherp
University of Koblenz
Koblenz, Germany
scherp@uni-koblenz.de

Neil Ireson
University of Sheffield
Sheffield, United Kingdom
n.ireson@dcs.shef.ac.uk

Ioannis Tsampoulatidis
Informatics & Telematics Inst.
57001, Thessaloniki, Greece
itsam@iti.gr

Yiannis Kompatsiaris
Informatics & Telematics Inst.
57001, Thessaloniki, Greece
ikom@iti.gr

ABSTRACT

During an emergency incident, several different log files are created by members of the Emergency Response (ER) personnel to document the emergency events that occur throughout the incident. Managing and reviewing these logs is a critical task for understanding and improving the implemented ER actions. A major challenge arising in this task is the merging of log files that are created by the different members of the ER personnel for the incident under study. Extensive manual effort is necessary to identify critical information, such as person names and locations, in order to align and merge the incoming log entries to make them suitable for review.

In this paper, we present the WeKnowIt ER Log Manager (WERL), a web-based application that facilitates the task of ER log merging and management by automatically aligning multiple log files and extracting ER-relevant semantic event information from log entry text. WERL makes use of the representation patterns of Event-Model-F in order to facilitate information sharing and reuse. Furthermore, WERL enables interactive exploration of the collected log files by means of temporal, location and semantic filters. Preliminary evaluation of WERL by members of the Sheffield City Council Emergency Planning team confirm that the application provides them with enhanced support during the ER log management and reviewing process.

Categories and Subject Descriptors

H.4.m [Information Systems Applications]: Miscellaneous

General Terms

Design, Experimentation, Human factors

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EiMM'10, October 25, 2010, Firenze, Italy.

Copyright 2010 ACM 978-1-4503-0176-3/10/10 ...\$10.00.

Keywords

Emergency response, Events, Incidents

1. INTRODUCTION

An important task in the function of professional Emergency Response (ER) organisations is the creation, management and reviewing of log files that document the actions taken by members of the ER personnel during the course of an incident. The resulting log files constitute valuable source of information for members of the ER personnel for two functions: (a) **decision making** during the course of an incident: based on the collected log entries in the control room of the organisation, tactical and strategic decisions need to be made until the incident is resolved, (b) **evaluation** of the implemented actions after the incident has been resolved based on reviewing the collected log entries.

Log files are generated by members of the ER personnel that are dispatched at the site(s) where the incident takes place. Due to the complexity and scale of several emergency situations, it is a frequent case to have numerous individuals dispatched at different locations. Each of them deals with a different aspect of the incident at hand and thus provides an isolated view of the incident through his/her log entries. In the end, multiple log files end up in the control room of the organisation and there is a profound need for assessing them together in order to form a complete awareness of the incident as a whole and determine the next actions in an optimal way.

Currently, this task is performed manually at the time the log entries arrive in the control room. Obviously, this is a labour intensive process that allocates precious human resources that would otherwise be exploited for the analysis of the incident. Furthermore, custom reports are created from the log files that summarise the main entities participating in the event: persons, locations and objects. The extraction of such entities is also a manual process requiring additional human effort.

To this end, we present a novel ER Log Merging and Management application, called the WeKnowIt ER Log Manager (WERL). WERL has been developed in the EU project WeKnowIt¹ and addresses the aforementioned problems arising in the management of ER log files by automatically merg-

¹<http://www.weknowit.eu/>

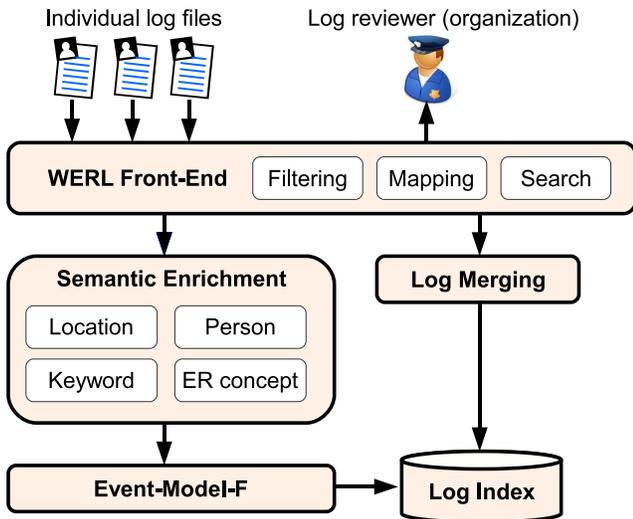


Figure 1: Application architecture.

ing and aligning log entries produced by different members of the ER personnel during an incident. In addition, it automatically extracts semantic information from the text of each log entry, i.e., an event that happens in the course of an incident. By this, WERL enables a concise view of the important content in log files. It allows its users to interactively explore the analysed log files by online temporal, location-based and semantic filtering, resulting in enhanced reviewing capabilities for decision making and evaluation. Finally, the application makes use of the Event-Model-F [8] as a representation framework for the structure and content of the events described by the ER log files. As the Event-Model-F provides a formal specification of events and event relations, it also makes log file content shareable and reusable. Figure 1 presents a high-level view of the WERL application. The individual components depicted in the figure (Event-Model-F representation, Semantic enrichment and WERL front-end) are described in the corresponding paper sections that follow.

The rest of the paper is structured as follows. Section 2 documents the representation model for ER logs. Section 3 describes the process of extracting semantic information from text in order to enrich the log entries. The application front-end and supported functionalities are detailed in Section 4. Some preliminary evaluation of the system is presented in Section 5. Section 6 provides a discussion of existing systems used for emergency incident management.

2. ER EVENT REPRESENTATION

Before describing the adopted representation formalism, it is first necessary to describe the typical format of log files used within modern ER organisations. Although there is no standardised format among all ER organisations, the basic elements of ER log files are commonly present independent of the particular organisation or unit. An example log file, which was produced by the Emergency Planning Team of the Sheffield City Council, is illustrated in Figure 2². The file comprises a header with information regarding the whole

²All log snapshots presented in this paper have been fully

log file and a table of log entries that document in detail the communication details and actions of the ER personnel members involved in the incident. The log file header contains authoring information, such as log file creator, unit to which the log creator belongs, incident date and title. The individual log entries contain four fields, namely (a) the time that the particular log entry was recorded, (b) the source and/or destination of the message between members of the ER personnel, (c) the content of the recorded message, and (d) an action associated with the message.

To represent the entries in the emergency response log files, we employ the Event-Model-F [8]. The Event-Model-F is a pattern-based core ontology that bases on the foundational ontology DOLCE+DnS Ultralight [2]. It provides a formal representation of events and event relations such as participation, composition, causality, and others and has been successfully applied to the development of the SemaPlorer++ application [9] for the distributed creation and sharing of event information in emergency response. As such, it is a suitable means to represent the event information stored in the incident log files and log entries. We consider log entries as events that happen in an emergency incident. To represent these entries we employ the participation pattern of the Event-Model-F to model the participation of living and non-living objects in the events and the documentation pattern to attach further metadata and documentary evidence to support the event. Each entry is considered part of a larger incident that is covered by the log file itself. Thus, we employ the composition pattern of the Event-Model-F to express that the events described by the single log entries are component events of a larger emergency incident event.

A graphical depiction of this representation is shown in Figure 3. Both the log files and the log entries are represented by a URI (incidentURI), a date/time of occurrence (dateTime), an identifier of their location (locationURI) and a set of documentation properties (additionalERProperties). The properties describing an incident log file are:

- Incident title: a short description of the incident that triggered the emergency response procedure,
- Logger: the author of the log,
- Service area: the division of emergency response organisation and
- Job role: the logger’s role in the particular incident.

while the ones needed for a log entry description are:

- From-To: the entities identified as the sender and receiver of the message of this entry,
- Message: the message to be delivered and
- Action: the action that needs to be done.

Below we provide an example of how to represent a log file using the Event-Model-F. Figure 4(a) shows the modelling of the incident log file itself, an acid incident at a school in Sheffield. The representation of the first two records of the log file, i.e., the first two sub-events of the incident are depicted in Figures 4(b) and 4(c).

anonymised to protect the identity and privacy of the people mentioned in the respective log files.

3.1 Location

For emergency incidents there is a requirement to identify relatively fine-grained locations (i.e. at the street level), thus a specific Sheffield-centric gazetteer was developed. In the application discussed in this paper, a number of high-quality, freely available (for non-commercial use), professionally produced resources were used, however these resources are UK-specific. In order to provide generic coverage it would be possible to use the collaborative constructed OpenStreetMap³ data; whilst this was considered at the time of development, it did not provide comparable levels of accuracy and completeness, for the Sheffield area. Given the continual and intense development of OpenStreetMap, it is envisaged that it will develop into a comparable resource to professionally produced resources, and indeed it was widely used by a number of professional organisations during the recent Haiti Disaster⁴.

During the location identification process the gazetteer place names are matched with the log text, taking into consideration the use of common abbreviations such as, rd=road, st=street, etc. However, even after identifying a location name, ambiguity may still exist due to multiple occurrences of the name, e.g. there are 12 locations called School Lane. In order to display the locations on a map, a disambiguation process attempts to uniquely identify the location. The process constructs a probability distribution of the potential locations based on the other (contextual) non-ambiguous locations in the log. The probability is increased if these contextual locations do not conflict with the potential location, i.e. they are within some distance threshold or the potential location is within the non-ambiguous location, if it refers to a wider area. If a single location receives the highest probability, given the context, it is selected as the actual location, if no unique location is identified, it is then a task of the reader to disambiguate the location. A more ideal situation would be to ensure that potential ambiguities were identified to the writer of the log, so that they could disambiguate the location ensuring the correct interpretation of the log.

3.2 Person names

The identification of person names applies a gazetteer of first names and common titles (e.g. Mr, Mrs, Dr, etc.) and case-sensitive patterns to identify likely entities. The two main issues with the process are to ensure that the gazetteer has the appropriate coverage of names/titles used in the logs, and dealing with ambiguous names. Name ambiguity can be caused by the use of common words as names, e.g. Heather, Rose, Hector, however it is presumed that the frequency of such ambiguous entities in the logs will be low. In addition names are commonly used in location names, e.g. John Street, in such cases the location takes precedence.

3.3 Key words/phrases

There are two processes employed in the identification of the key words and phrases in the incident logs; a general term extractor utilises freely available web-services, whilst specific emergency incident terminology is extracted using an ER thesaurus.

Three web services which provide general terminology ex-

³<http://www.openstreetmap.org/>

⁴http://wiki.openstreetmap.org/wiki/WikiProject_Haiti

traction were examined: Zemanta⁵, OpenCalais⁶ and Yahoo! Term Extraction (YTE)⁷. Both Zemanta and OpenCalais provide semantic annotations of text, extracting concepts and linking these to specific web content (from Wikipedia, Youtube, IMDB, etc.) identified by their unique URI. As these services are aimed at writers of general web content their response is limited to Wikipedia related concepts, to limit the amount of noise. YTE simply provides a list of extracted terms/phrases. There is no clear indication of the data and technologies used to derive these key terms/phrases; however examining the results on a number of logs indicated that it tends to produce more comprehensive and pertinent extraction.

There is no standardised thesaurus for ER terminology. A number of organisations involved in Emergency Management publish glossaries of terms and acronyms which can be used to provide more specific terminology extraction. The current implementation is limited to the identification of ER acronyms, such as FLO standing for Forward Liaison Officer.

4. USER INTERFACE OF WERL

WERL was developed after discussions with members of the Emergency Planning team of Sheffield City Council who made clear the need for an automated means of merging disparate log files and searching in them by use of pertinent criteria, such as location, person name and ER-specific keywords.

The front-end of WERL provides online filtering capabilities for facilitating the interactive exploration of the available log entries. A snapshot of the application main screen is provided in Figure 5. At the top, a slider-based time filter is available that enables the examination of a particular time interval of the incident. In addition, standard full text search capabilities are provided for retrieving only the subset of log entries that are relevant to the input query. Most importantly, there is a series of four semantic filters that summarise the main entities found in the log files by the text annotation component described in Section 3. Thus, it is possible to view only the log entries that are related to a particular location, person name, significant keyword or ER acronym. The presentation of all identified semantic entities in these lists can provide the ER user with a quick overview of the semantic content of the log file. The filtering and search operations are performed on the client side so that the log reviewing process is highly responsive and the application scales to many users by limiting server-side processing to the semantic enrichment of the input log files.

A significant feature of WERL is the presentation of provenance information and the possibility to filter based on the provenance of log entries. As illustrated in Figure 6, beside each log entry there is a marker indicating its origin. At the bottom of the log entry list, there is an associated legend, which can also be used for filtering based on the log entry provenance. In that way, it is possible to inspect only the log entries produced by a particular log creator, thus gaining insight into his/her perspective of the incident. In larger scale incidents, involving many members of ER personnel coming from different organisations (e.g. fire department, police), it

⁵<http://www.zemanta.com/>

⁶<http://www.opencalais.com/>

⁷<http://developer.yahoo.com/search/content/V1/termExtraction.html>

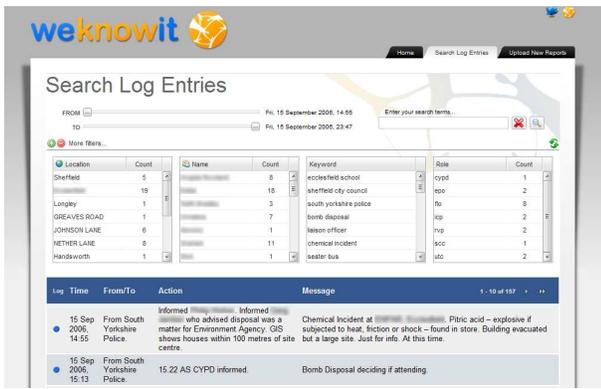


Figure 5: Snapshot of WERL front screen

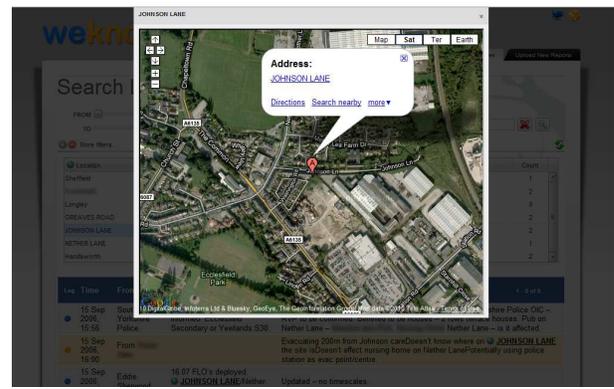


Figure 7: Snapshot of WERL map view

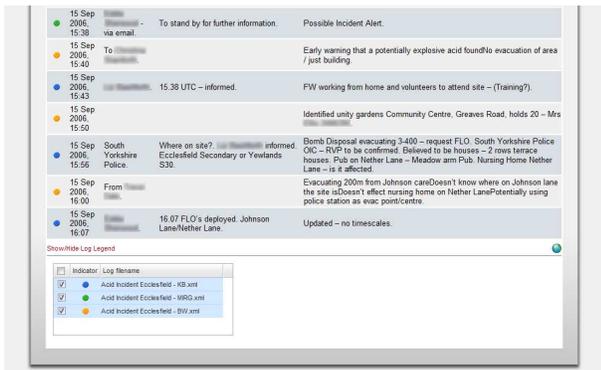


Figure 6: Snapshot of WERL provenance filtering

is expected that more sophisticated provenance mechanisms will be necessary, e.g. provenance by organisation, unit, role in the organisation, etc.

Finally, a valuable WERL feature is the possibility to link from the log entries directly on the map of the area where the ER incident has taken place. Figure 7 depicts this capability. Such a feature is possible due to the automatic localisation of text providing geo-coordinates, as described in subsection 3.1. Such visualisations aim to improve the ER personnel's overall situational awareness of the incident and thus allow them to make more informed decisions with respect to the next actions.

5. EVALUATION

During the development of the WERL application in the WeKnowIt project, there have been two experts from the Emergency Planning team of the Sheffield City Council involved. The experts had access to the online application and we provided them with basic usage instructions such that they can work with the application. We have discussed the application with the experts in face-to-face meetings during project meetings and the experts gave us feedback on a regular basis.

For the current implementation of the WERL application, we have received an experience report by one expert of the Emergency Planning team documenting several attributes of the application as well as points for improvement and desirable features. Other than that, the pilot test was conducted in a completely unsupervised way. The general impression

was that the application was easy to use and presented a number of advantages both during and post incident phase of log management. The expert expressed a desire to test the system with a larger dataset (for testing a single incident covered by three log files was available). With respect to the search facilities, the expert noted that they worked well, but there were cases where the search produced wider results than expected. On the other hand, it would be desirable for the search to pick up misspellings and typing errors, which are not uncommon in ER logs, since they are typed under time pressure.

In terms of features, the expert found useful the possibility to jump from a location reference in the text to its map-based view. In addition, he noted that it would be useful to view on the map all locations contained in the log entries under examination. An additional desirable feature for future versions of the application is the possibility to create summary reports from the log (which is supported by state-of-the-art incident management systems).

6. RELATED WORK

In generic terms the work relates to the identification of similarities, and differences, amongst related documents; an area explored in the document summarisation field [5]. The approach adopted in this work does not use summarisation as the number of documents is relatively low and there is a desire for the user to have all the information available to maintain control. Instead the work facilitates the user's exploration of the documents via the automatic identification of generic and domain specific concepts and the use of faceted browsing [4] to filter documents according to those concepts.

At a system level, the work is related to the general domain of C4I software, which stands for Command and Control Systems and Components. Amongst the numerous solutions available in the market, which target a wide range of applications, such as military operations and surveillance, WERL is most closely related to ER incident management solutions, three of which we briefly describe here: the Atlas Incident Management System (AIMS)⁸, the Emergency Command System⁹, and the Bristol City Council map-based application [1]. The main focus of such software solutions is

⁸<http://www.atlasops.com/products/aims.php>

⁹<http://www.emergencycommandsystem.com/>

the support for information sharing and communication, as well as task and asset management.

AIMS is the emergency incident management system available from Ultra Electronics Command & Control Systems. The system is a complete IT solution targeted at the needs of ER organisations and provides extensive infrastructure features for message logging and distribution among the members of the organisation. However, AIMS does not support automatic log file merging, nor does it support automatic information extraction from the log text. Thus, we consider that WERL offers complementary capabilities on top of AIMS.

Emergency Command System by VectorCommand Ltd. is an alternative integrated software solution to the management of ER incidents. It supports a wide range of command and control features such as GIS mapping, live and still imagery, asset management and instant messaging. Still, this system does not offer automatic extraction of semantic information from log text, nor does it provide an underlying formal representation model for the recorded log entries.

The Bristol City Council Emergency Planning team has built a suite of tools that brings together ER-pertinent geographic information. Their main source of information comes from Ordnance Survey datasets [1]. The focus of this suite of tools is the improvement of situational awareness by means of overlaying different kinds of information on the same map. However, it does not include information coming from the logs of the ER personnel nor does it perform any semantic analysis on the log text.

Developed at the University of Maryland, project 911.gov aims at developing a Web 2.0 platform supporting the collaboration of organisational entities for ER and citizens [10, 12]. A recent Open Source system, Sahana [6], provides a Web 2.0 platform for connecting ER organisations with volunteers. This platform is aimed at the setup of an online community by an organisation for a specific (large scale) incident.

Finally, a pertinent trend in the ER domain is the exploitation of massive user contributed content for detecting emergency incidents from streams of online information, e.g. micro-blogging messages. In fact, for a number of recent earthquakes in Japan it has been demonstrated in [7] that user messages from microblogging services can be used to detect seismic events with high accuracy and notify people before the national broadcast services or even professional ER organisations. Other ongoing works exploit microblogging messages to create spatio-temporal visualisations of large-scale emergency incidents, e.g. swine flu [11]. Although the results of such works are currently far from being integrated in the formal ER incident management solutions used by professional ER organisations, it is not hard to imagine that event detection and situational awareness based on user contributed content will eventually constitute an integral part of professional ER management systems.

7. CONCLUSIONS

We have presented WERL, the WeKnowIt ER Log Manager application for the management of ER log files. The log files are processed and enriched with semantic metadata by means of information extraction. Using the Event-Model-F as representation mechanism, we store the event information contained in the log files and make it available to the ER entities for review and post-incident analysis.

In the future, we plan to enrich WERL with additional capabilities following the suggestions received after the pilot test. First, we are going to refine the search facilities in two ways: (a) return more specific results, (b) support retrieval of misspellings and typing errors. In addition, we plan to support the possibility for simultaneously viewing all log referenced locations on the map. Finally, we are also going to provide support for semi-automatic summary report creation.

8. ACKNOWLEDGMENTS

This work was supported by the FP7 project WeKnowIt, partially funded by the EC under contract number 215453. We would like to thank Keith Bradley and Peter Whitham from Sheffield City Council - Emergency Planning Team for their valuable feedback during the problem formulation and application development and evaluation.

9. REFERENCES

- [1] Using geography can help you to meet your flood management responsibilities. <http://www.ordnancesurvey.co.uk/oswebsite/business/sectors/government/local/news/flood-management.html?cmp=floodmgt>, June 2010.
- [2] S. Borgo and C. Masolo. *Handbook on Ontologies*, chapter Foundational choices in DOLCE. Springer, 2nd edition, 2009.
- [3] M. Greenwood, J. Iria, and F. Ciravegna. Saxon: an extensible multimedia annotator. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, May 2008.
- [4] M. A. Hearst. Clustering versus faceted categories for information exploration. *Commun. ACM*, 49(4):59–61, 2006.
- [5] I. Mani and E. Bloedorn. Summarizing similarities and differences among related documents. *Information Retrieval*, 1(1-2):35–67, May 1999.
- [6] Sahana Software Foundation. Sahana - home of the free and open source disaster management system. <http://sahanafoundation.org/>, 2010.
- [7] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 851–860, 2010.
- [8] A. Scherp, T. Franz, C. Saathoff, and S. Staab. F—a model of events based on the foundational ontology DOLCE+DnS Ultralight. In *K-CAP*, pages 137–144, 2009.
- [9] A. Scherp, S. Papadopoulos, A. Kritikos, F. Schwagereit, C. Saathoff, T. Franz, D. Schmeiss, S. Staab, S. Schenk, and M. Bonifacio. D5.2.1 - prototypical knowledge management methodology. <http://www.weknowit.eu/sites/default/files/D5.2.1.pdf>, Jan. 2010.
- [10] B. Shneiderman and J. Preece. Public health: 911.gov. *Science*, 315(5814), Feb. 2007.
- [11] V. K. Singh, M. Gao, and R. Jain. Situation detection and control using spatio-temporal analysis of microblogs. In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 1181–1182, 2010.
- [12] U.S. Department of Transportation. 911.gov - when to call 911. <http://911.gov/>, 2010.