Automated Service Integration for Crisis Management *

Alan Berfield, Panos K. Chrysanthis, Alexandros Labrinidis
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260, USA
{alandale, panos, labrinid}@cs.pitt.edu

In the Proceedings of the First Workshop on Databases In Virtual Organizations (DIVO 2004)
Held in conjunction with the ACM SIGMOD/PODS 2004 Conference
http://dais.cs.uiuc.edu/divo2004

Abstract

The integration and coordination of different emergency service personnel is crucial to Crisis Management. Crisis centers create plans of action to deal with the various situations that arise during an event. These plans require different personnel with different expertise to execute them within a given set of constraints. In this paper, we show that such plans can be represented as workflows, and that the discovery and integration of personnel can use a scheme we previously developed for the establishment of virtual enterprises. An important extension to our previous work is the introduction of Intelligent Personal Assistants, or iPAs. iPAs aid in the discovery of appropriate personnel and provide the coordination between them and the crisis centers. iPAs also assist their users during plan execution by interfacing with different information sources, such as sensor networks, and by helping to dynamically modify the current plan of action as necessary.

1 Introduction and Motivation

Emergency service agencies create various plans of action for different disaster scenarios. Recently, much work has focused on improving information gathering and dissemination during a crisis. One example is project RESCUE[11], which details the numerous technical challenges in supporting crisis management, and proposes different ways to handle them. Another one is our project, S-CITI[3], which aims to provide a robust infrastructure to support the data gathering and decision making before, during, and after emergencies or disasters.

When a crisis occurs, through the use of such information collection methods an appropriate plan is dynamically instantiated with the specific details of the crisis. These plans may often involve or require coordination and interaction with personnel from other agencies or centers. This coordination could be due to policy or because of a lack of expertise. As an example, a fire station could have various plans for handling different kinds of fires and evacuations. Some of these plans may require the expertise of police officers or medical technicians.

In order to commit such plans for execution, the agency must locate and coordinate with personnel having the necessary skills and being able to meet the constraints and requirements of the tasks. Finding such personnel

*Funded in part by NSF ITR Medium Award (ANI-0325353)

during a crisis can be difficult, as communication may be unreliable, and the personnel may be scattered over different locations (and areas) and engaged in different activities. Many would be detained or unavailable entirely. Instead of locating and contacting individuals directly, we envision an environment where the agency can broadcast electronically its plan of action. This broadcast will be received by both the appropriate centers as well as any available personnel through their mobile devices (e.g., PDAs), and those able to contribute to the execution of the plan can respond electronically. In our example, the fire station can broadcast a requirement for a number of police officers to go to a building to aid in its evacuation. Any officers that receive the broadcast and are able to go can reply to the call. The agency chooses those most suited from all responders, and confirms their participation. In this way, teams of emergency personnel are dynamically formed based on the requirements of executing different plans of action for the emergency situation at hand.

During the actual execution of these plans, it may become necessary to modify them. This can be due to many reasons, such as environmental changes (e.g., a building collapses or fire erupts), team members being drawn off for more critical tasks, unforeseen obstacles or delays, and so on. In order to handle such a change, all team members must be alerted, along with any monitoring centers.

Using an electronic broadcast to communicate with emergency personnel provides many advantages, with improved security being one of them. With "smart" PDAs like the proposed iPAs, secure data communication becomes very inexpensive (since it is software-based), compared to the purchasing of specialized "secure" voice communication devices (e.g., military-grade walkie-talkies), which are very expensive. In this way, sensitive information and details about the emergency at hand can be transmitted to all those that could help. In addition to the improved security, using iPAs allows for the emergency broadcast to reach a much greater pool of people, which could include for example off-duty personnel or doctors who may not be monitoring a radio or carrying around a bulky walkie-talkie, but can conveniently keep a small PDA that also has other uses. This allows such people to become "first responders", if they are located closer to the site of the emergency, and thus drastically improve the quality of the emergency response.

In this paper, we discuss representing plans of action and coordination as workflows. Integrating the various personnel to fulfill the goals and requirements of these workflows is similar to the integration of different businesses in the establishment of a virtual enterprise. In our previous work, we developed a scheme to automate this process[8, 5]. We will show that we can extend this framework to also support emergency service integration. One key modification is the introduction of our mobile Intelligent Personal Assistants[4], or iPAs, to be used by all emergency personnel. We utilize the iPAs to properly and effectively coordinate with individuals. They can also significantly aid in the performing of tasks, by providing additional information, as needed.

In section 2, we discuss our workflow representation of crisis management plans. We then describe the Intelligent Personal Assistants in section 3. In section 4, we present an overview of the automated discovery and integration of emergency services personnel. Workflow execution, coordination, and monitoring with the help of the iPAs is detailed in section 5. We conclude in section 6 with a summary and future work.

2 Workflow Model

Workflow Representation Plans of action for crisis management can be represented as a set of tasks and the relationships among them. Workflows allow us to provide a clear specification of a service that is easily understood. While many workflow models[7, 13, 10] include only tasks and relative orderings, we require more detailed information for accurate service integration. For this reason, we assume an enriched model of workflows, which includes the following:

- *Pre- and post-conditions*, which specify what must be true before a task can be executed, and what will be made true as a result of the task's performance.
- *Causal links*, which relate each task that establishes a condition (listed in its post-conditions) to the task that requires it (listed in its pre-conditions).
- *In- and out-parameters*, which are used in the evaluation of pre-conditions and post-conditions. They carry information and engender data flow during execution. For example, an address could be an in-parameter to a "go-to-location" task.
- *Temporal Constraints* that specify the earliest and latest start and end times of a task, as well as the minimal and maximal durations of the task. They can be absolute times or relative to the execution of other tasks.
- *Resource Constraints*, which specify the equipment, material, or agent resources required for the task. These resource constraints can be further classified based on their particular properties. This classification is very similar to the National Incident Management Resource Typing system[9].
- *Significance*, which indicates whether the task is *vital* to the workflow and therefore must be executed, or whether it is *non-vital*, and need only be executed if feasible[8].

Workflow Properties In addition to specifying the representation, our workflow model includes three properties of workflows which are important for scheduling and integration. These are: *correctness*, *completeness*, and *compatibility* with existing commitments.

Definition 1 Workflow Correctness: *A workflow is* correct *if and only if*

1. *it has no conflicting temporal or resource constraints,*
2. *for each goal/precondition P, there is a task that achieves P (the producer task), and it is ordered before the task that requires it (the consumer task), and*
3. *for each goal/precondition P, no task that may negate P can possibly be ordered in between the producer and the consumer.*

This notion of correctness is important because only correct workflows can possibly be executed. Note that some workflows may contain preconditions that are assumed to be established independently of the workflow itself. We will call such preconditions *open* with respect to the workflow. A simple example of such an open precondition is a workflow for renting a car that assumes the precondition of having a driver's license. Workflows with such open preconditions are incorrect until they have been combined with other workflows that establish all open preconditions.

Definition 2 Workflow Completeness: *A* complete *workflow is a workflow that specifies all tasks needed to achieve its goals and preconditions.*

Definition 3 Workflow Compatibility: *A workflow is* compatible *with another if none of its tasks conflicts with any of the other's (and vice versa).*

This means that the temporal constraints, resource usage, and postconditions of a workflow's tasks do not prevent the execution of the tasks in the other workflow (though they may place limits on the times when those tasks can be executed). So, for example, a compatibility conflict between workflows arises if two tasks that use the same resource (e.g., equipment) are set to execute at the same time.

Example A simple example workflow is shown in Figure 1. It contains steps for locating and retrieving a victim, requiring both search & rescue and medical expertise. We use a simple graph notation to depict workflows. The nodes/boxes in the graph represent tasks to be carried out, whereas the arrows show temporal constraints. This workflow has no conditional tasks; resource constraints and other features are also not shown for simplicity. In the example of Figure 1, a dispatch agency must contact a search & rescue team to locate and retrieve a victim, and to send an ambulance to take him/her to the hospital. Victim retrieval varies with the nature of the incident, e.g., retrieval from a collapsed building or retrieval from malicious captors.

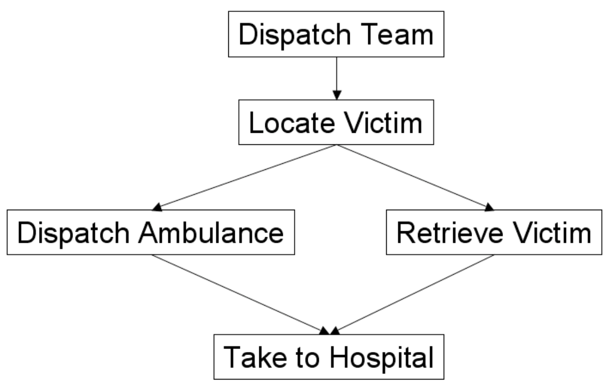

Figure 1: Example Workflow

We propose to use an XML representation of such workflows, to allow for customization among different emergency response agencies and scenarios. We are closely monitoring standardization efforts within the context of crisis management, in order to utilize such standards in our workflow representation. One such related effort is the National Incident Management System (NIMS) proposal[9], which emphasizes the need for the development of national standards of communication protocols and data formats. We plan to keep our XML representation compatible with the NIMS standard (which does not currently support workflows however). Specifically, we will encode the NIMS resource typing system into our workflow representation of resource constraints.

The NIMS resource typing system describes resources using:

- *Category*, a function for which a resource would be most useful.

- *Kind*, broad classes that characterize like resources.

- *Components*, when a resource comprises multiple parts.
- *Metrics*, measurement standards describing capability and/or capacity.
- *Type*, level of resource capability (e.g., Type I has a greater level of capability than Type II).

Returning to our search & rescue example, Figure 2 shows an example of our NIMS-compliant XML description of the resource types available to search & rescue task forces. This description captures statistics, expertise, and equipment information for a search & rescue task force. It also illustrates some differences between Type I and Type II versions of the resource. Using such XML allows for easy workflow requirement specifications, and facilitates matching to resources.

```
<resource name=“Search & Rescue Task Force” category=“Search and Rescue” kind=”Team”>
    <component name=“Personnel”>
        <metric name=“Number in Team”>
            <type1> 70 </type1>
            <type2> 28 </type2>
        </metric>
        <metric name=“Training”>
            <type1> NFPA 1670 </type1>
            <type2> NFPA 1670 </type2>
        </metric>
        <metric name=“Areas of Specialization”>
            <type1> high angle rope rescue, confined space rescue </type1>
            <type2> basic rope rescue </type2>
        </metric>
    </component>
    <component name=“Equipment”>
        <metric name=“Rescue Equipment”>
            <type1> Pneumatic Tools, Electric Tools, Hand Tools, Safety </type1>
            <type2> Pneumatic Tools, Electric Tools, Hand Tools, Safety </type2>
        </metric>
    </component>
</resource>
```

Figure 2: Example of NIMS-compliant XML Resource Description

While the NIMS scheme defines resource typing, it does not propose the use of XML, nor does it consider any notion of workflows. Therefore, we believe that our XML workflow model can be a helpful influence on the emerging standards.

3 Intelligent Personal Assistant

We envision all emergency service personnel to be equipped with specialized PDAs or cell phones which will act as intelligent personal assistants. Towards this, we propose to adopt and adapt our iPA (Intelligent Personal Assistant) framework. iPA was developed to aid its user in a ubiquitous computing environment[4]. We designed it to run on a mobile platform (PDA or cell phone) in order to stay with its user at all times, even when she/he is mobile. Over time, an iPA “learns” about its user, building a profile to improve its own behavior with the goal of becoming a better personal assistant. It can gather, sort, and present information tailored to its user's current activities. To help protect its user's privacy, it never shares this profile with anyone.

An iPA does not require its user to ask for anything. Instead, it "learns" from what it observes from its user's actions and from anything it was "told" explicitly. Using this information, it will let its user know when something important or relevant is found, even when not asked. In this way, it becomes knowledgeable about everything going on and knows what its user wants to hear about at the right time and place.

In addition to the profile, an iPA also tailors its information gathering and delivery based on its user's current context. Context information includes current day and time, the user's location and possibly the current task being performed by the user. Location can be determined via GPS or similar positioning system if available, or from queries to nearby sensor networks. Additionally, context can incorporate other conditions supplied by the user, the current environment, or even different users and their activities and statistics.

An iPA gathers information from a variety of sources. These can include servers, both query-response and broadcast, and other iPA peers. Servers could include standard database and web servers, as well as sensor networks. Broadcast servers would ideally consist of two main channels[1, 2]. The first channel is a push-based channel that continuously broadcasts a data stream of XML documents. The other channel is a pull-based channel where the user or iPA can submit queries to the server requesting more detailed information about an item in the broadcast.

In our proposed scheme, it is the iPAs that will receive broadcast workflows from various crisis centers. Based on its profile, context, and user's current activities, the iPA will filter these requests and only present those that its user could participate in.

The iPA must know its user' abilities and preferences in order to be able to match requests accordingly. Originally, the iPA was targeted towards a commerce environment. This involved monitoring the habits and preferences of the user regarding restaurant choices and product purchases, tailored to different contexts. Any XML advertisement received was only given to the user depending on how well it matched these preferences, given the current user context.

In a crisis management environment, the iPA instead needs to identify the capabilities and skills of the user. One way to do this is to keep statistics on the experience of the user by monitoring the incidents the user participates in. In this way, when an XML workflow is received, its skill requirements can be matched to the expertise of the user, again taking into account the current context.

4 Automated Discovery and Integration

In the event of a crisis, a new workflow is instantiated at a crisis center in order to respond to the situation at hand. These workflows may be based on those kept in database repositories of workflow templates for different situations and are possibly distributed across different agencies. Instantiation involves adding all specific details for the current crisis to the template and attempting to commit it for execution. The workflow must then be *outsourced* for execution to appropriate personnel. Locating the most appropriate personnel based on their characteristics and the workflow requirements is done by the *integration* process.

We use the notion of views to express the outsourcing. Any subgraph of a workflow graph defines a *view* of the workflow. Formally, a workflow view can be defined as a *projection* on the graph based on some criteria. The

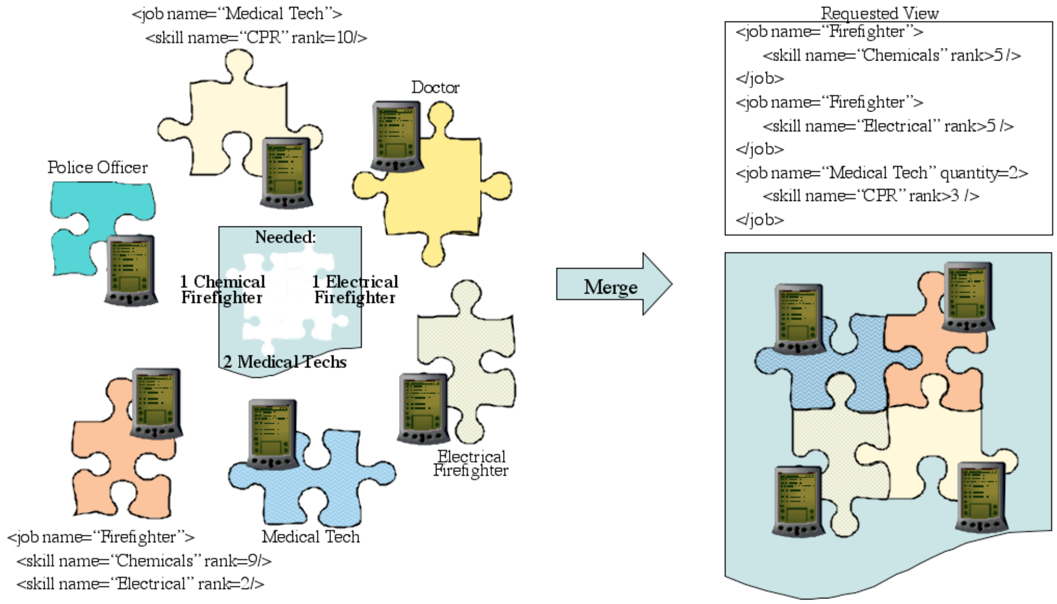

Figure 3: Discovery and Integration

nodes in a view retain all the tasks and information of the original nodes, including all constraints.

A workflow view can represent any activities performed by a *service provider* on behalf of a *service requester*. Consequently, workflow views can be used to express both service requests and the contributions of different emergency service providers. For example, a request has the following structure:

Request: $Rq = (P, V)$, where P = Requester Profile and V = Requested View

Given such requests, integration takes place in a distributed environment consisting of multiple centers and emergency agencies (acting as requesters) and various personnel (acting as providers). The requester profile can contain various information about the requester, such as name, site identification, credentials, etc.. While we are not concerned with issues of trust in this paper, the profile and request could be enhanced with schemes like Trust-X[6] and TrustBuilder[14].

When an iPA receives the broadcast view request, it must determine if its user is capable of contributing to its execution. It must consider the specified goals of the request and its constraints and requirements with respect to its user's abilities, context, and current activities and commitments. In many cases, the iPA would immediately eliminate a request as being infeasible for its user. Otherwise, if its user could execute all or part of the request view within the associated constraints, it will provide its user the details and seek approval. If the user accepts, the iPA relays to the requester the portion of the view (a view of the view) its user can achieve, along with any added constraints.

At the requester, it is necessary to find a valid combination of providers. This combined workflow must adhere to our definitions of correctness, completeness, and compatibility (with any already-scheduled workflows). It must achieve all goals of the outsourced view, all temporal constraints must be satisfiable, and the combined workflow must not contain any resource conflicts. To do this, we have developed and implemented a process for merging workflow views using Epilitis[12], a temporal constraint-satisfaction solver. Epilitis will generate a correct, complete, and compatible workflow if possible, and when not possible will indicate problems and suggest solutions.

An overview of the integration process is depicted in Figure 3. The workflow view indicates needed resources and constraints. The puzzle pieces on the left represent different personnel that could possibly fit the request. We have included some sample XML to illustrate the kind of data being used. All pieces that match any of the view respond to the requester, and those chosen are merged into a team with a single overall workflow.

5 Coordination, Monitoring, and Execution

Coordination and monitoring during the crisis involves the interaction of many different systems and individuals. Crisis centers must manage their overall workflows. This involves monitoring the progress of execution, receiving notifications from the personnel involved, and forwarding revisions as necessary.

The centers can coordinate with the individuals executing the tasks in their workflows via their iPAs, which can receive updated instructions, revised workflows, and queries. IPAs notify their users only as necessary, so as not to be too distracting. They can report to the center the current status and context of their users, and notify them of any failures or constraint violations. In some cases they can even anticipate violations or failures given current conditions and known properties. They will also report successful completions of tasks.

In the event an iPA or its user determines changes to the current workflow are necessary, the iPA can help by providing alternatives that still satisfy all requirements. If the iPA's user has sufficient authority, he/she can select one of these alternatives. The iPA will automatically notify the center and any other personnel involved. Otherwise, alternatives will be forwarded to the center for selection or approval.

In addition to helping with coordination, iPAs provide other valuable assistance during workflow execution. They provide an interface to information sources and fellow emergency personnel. Queries to various sensor networks can be initiated through an iPA. This can help with tasks such as locating victims, avoiding obstructions, and detecting dangerous environmental conditions. Since iPAs are highly personalized, information is tailored to the current specific needs of the user.

The greatest advantages of using an iPA for such applications are its own "initiative" and context awareness. Appropriate assistance can be initiated by the iPA itself, as it constantly monitors data streams for relevant information, thus freeing its user to work more effectively. Knowledge of what to monitor can be based on properties of the workflow being executed, or on observed information (e.g., from queries made by the user in the past). As an example of this initiative, imagine a firefighter searching for people in a burning or collapsing building. His/her iPA is constantly monitoring data for her/him from various sensors in the building and from other firefighters and personnel. It will warn the firefighter if he/she is approaching a dangerous area and can guide him/her to any detected people while avoiding obstacles. To prevent duplicate effort, it will notify others of his/her location and activities, and inform the firefighter if others are already performing the task he/she intends to do. Finally, combined with physical sensors (like a heart monitor), the iPA can also automatically "sense" when the firefighter needs medical assistance and contact someone nearby.

6 Conclusions

We have presented a workflow representation of action plans for crisis management. Using workflows we can adapt our scheme for the establishment of virtual enterprises to the discovery and integration of appropriate emergency services personnel for a given crisis management plan. We use a global broadcast of the needed workflow which includes goals to be achieved and constraints and requirements of execution. To discover personnel, we make use of mobile Intelligent Personal Assistants (iPAs) that receive the broadcast and in part determine if their users are able to participate. These iPAs also aid in the coordination between personnel and crisis centers, as well as between individual personnel. iPAs can aid their users during plan execution by interfacing with different information sources, such as sensor networks, and by helping to dynamically modify the current plan of action as necessary.

Our future work involves modifying our implementations of the integration scheme and the iPAs to better support crisis management outlined in this paper with special emphasis on power awareness and communication reliability.

References

[1] Acharya S., M. Franklin, and S. Zdonik. Disseminating Updates on Broadcast Disks. *VLDB*, 1996.

[2] Acharya S., M.Franklin, and S.Zdonik. Balancing Push and Pull for Data Broadcast. *ACM SIGMOD*, 1997.

[3] Amer A., J. Brustoloni, P.K. Chrysanthis, M. Hauskrecht, A. Labrinidis, R. Melhem, D. Mosse, K. Pruhs, L. Comfort. Secure-CITI: A Secure Critical Information Technology Infrastructure for Disaster Management. *Hazard Reduction and Response in Metropolitan Regions: An Interdisciplinary Model, Workshop and Interactive Videoconference among Metropolitan Regions*, 2003.

[4] Berfield A., J. Beaver, and P.K. Chrysanthis. Profile and Context Filtering of Streaming Data for a Mobile Personal Assistant. *ACM Symposium on Applied Computing*, 2004.

[5] Berfield A., P.K. Chrysanthis, I. Tsamardinos, M. Pollack, and S. Banerjee. A Scheme for the Establishment of Virtual Enterprises. *Research Issues in Data Engineering*, 2002.

[6] Bertino E., E. Ferrari, and A.C. Squicciarini. Trust-X: An XML Framework for Trust Negotiations. *Communications and Multimedia Security CMS*, 2003.

[7] Chrysanthis P.K.. Guest Editors Introduction to Special Issue on Workflow Systems. *Distributed Systems Engineering*, 1996.

[8] Chrysanthis P.K., T. Znati, S. Banerjee, and S. Chang. Establishing Virtual Enterprises by means of Mobile Agents. *Research Issues in Data Engineering*, 1999.

[9] Department of Homeland Security. National Incident Management System. 2004.

[10] Kwak M., D. Han, and J. Shim A framework for dynamic workflow interoperation using multi-subprocess task. *Research Issues in Data Engineering*, 2002.

[11] Mehrotra S., et al. Project Rescue: Challenges in Responding to the Unexpected. *IS&T/SPIE Annual Symposium*, 2004.

[12] Tsamardinos I. Constraint-Based Temporal Reasoning Algorithms with Applications to Planning. *Ph.D. Thesis. University of Pittsburgh Intelligent Systems Program*, 2001

[13] Yang J., M.P. Papazoglou, and W. Van den Heuvel. Tackling the challenges of service composition in e-marketplaces. *Research Issues in Data Engineering*, 2002.

[14] Yu T., M. Winslett, and K.E. Seamons. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM TISSEC*, 2003.