Using a Trust Inference Model for Flexible and Controlled Information Sharing During Crises

Qian Yang*, Danfeng Yao**, James Garnett*** and Kaitlyn Muller****

*Division of Computer and Information Sciences, Rutgers-The State University of New Jersey, 110 Frelinghuysen Road, Piscataway, NJ 88854-8019, USA. E-mail: qianyang@cs.rutgers.edu

**Department of Computer Science, Virginia Polytechnic Institute and State University 2202 Kraft Dr. KWII, Blacksburg, VA 24060, USA. E-mail: danfeng@cs.vt.edu

***Department of Public Policy and Administration, Rutgers –The State University of New Jersey, 401 Cooper Street, Camden, NJ 08102, USA. E-mail: garnett@camden.rutgers.edu

****Refugee and Immigration Services, Catholic Charities, Diocese of Camden Inc., 1845 Haddon Avenue Camden, NJ 08103, USA. E-mail: Kaitlyn.Muller@Camdendiocese.org

The fluid, urgent nature of crises requires flexible, responsive information sharing. Recent studies show, however, that in business catastrophes and other kinds of crises conventional access control mechanisms favor security over flexibility. Our work addresses these seemingly contradictory needs for security and flexibility and designs a trust inference model based on fuzzy logic, a model that can be used with pervasive computing technologies using sensors and mobile devices. Drawing upon research on trust, we design a trust inference model using attributes of affiliation, task performance, and urgency; apply the model to a known crisis; discuss implementation issues; and explore issues for further research.

This article is dedicated to Alan Jarman, a founding influence in the *Journal of Contingencies and Crisis Management* who died in Canberra 15 July 2010. Alan's quantitative, engineering background and his long standing commitment to improving crisis decision making prompted him to encourage our applying fuzzy logic to crisis information sharing. We are grateful for Alan's encouragement and advice.

1. Introduction

Crises are characterized by urgency, rapid change, unpredictability and complexity (Alink, Boin, and 't Hart, 2001; Boin and Lagadec, 2000; Comfort, Ko, and Zagorecki, 2004; Drabek, 1986). Information technologies have been utilized to cope with these aspects of crises, particularly for the communication during a crisis (Newkirk, 1993; Fischer, 1999; Torrieri, Concilio, and Nijkamp, 2002; Garnett & Kouzmin, 2007; Palm & Ramsell, 2007). Much of the existing research on information sharing during crises has extensively addressed the need for secure access, focusing on comprehensive policy designs and analysis and efficient management of users' privileges and privacy. Recent studies found that in mission-critical systems, e.g., military, firefighting or supervisory control and data acquisition (i.e., computerized monitoring and controlling system), conventional access control mechanisms may be too rigid for urgent information-sharing scenarios and

DOI: 10.1111/j.1468-5973.2010.00621.x

often fail to provide adequate support for access in non-routine, critical situations (Cheng, Rohatgi, Keser, Karger, Wagner, & Reninger, 2007; MITRE Corp., 2004; Swarup, Seligman, and Rosenthal, 2006; Keppler, Swarup, and Jajodia, 2006; Singh, Sanders, Nicol, and Seri, 2006). In critical infrastructures such as utility networks, oil and gas pipelines, and disaster and anti-terrorist communications, there is an increasing need to secure the information collected from and about the infrastructure, and yet to be able to allow flexible data sharing to facilitate problem-solving.

Networked computers, sensors and mobile devices are pervasively used for data collection and storage. To utilize effectively the vast amount of data generated, information needs to be shared across organizational and administrative boundaries. Facilitating trust and cooperation in crisis situations can aid in successful interventions and recovery efforts. The purpose of authorization is to control and facilitate the access to shared resources by entities (people or devices) belonging to different autonomous domains. The challenge for security research on distributed authorization is twofold: (1) secure and accountable access: how to guard the integrity and confidentiality of shared resources; (2) flexible adaptation: how to facilitate flexible and dynamic information sharing.

Research in crisis management (Boin and Lagadec, 2000; Waugh & Sylves, 2002; Wise, 2006; Comfort, 2007; Derthick, 2007; Garnett & Kouzmin, 2007), however, shows that in crisis situations (e.g., natural and technological disasters, terrorism, firefighting), traditional central command and control models are either unavailable or too inflexible for urgent information sharing, and often fail to provide adequate supports for data access across organizational boundaries. Although the majority of causes of failures found in the above studies relate to the administrative issues, the lack of any technical infrastructure that can be used to facilitate and enable cross-domain information sharing also exacerbates the problem. There is an increasing need to secure the information collected from distributed and mobile devices (e.g., location information), and yet to be able to allow flexible sharing to facilitate problem-solving and decision-making. Cross-domain information sharing also requires high accountability, so that misuses of data can be discovered and malicious users can be identified and held accountable for their behaviors. These problems are unique and challenging in emergency and crisis situations because of the dynamic nature of shared data and users. Several notable papers have proposed interesting solutions to the problem of flexible and controlled information sharing (Cheng et al. 2007; MITRE Corp., 2004; Swarup et al., 2006; Keppler et al., 2006; Tamassia, Yao, and Winsborough, 2004; Yao, Frikken, Atallah, and Tamassia, 2006; Yao, Tamassia, and Proctor, 2005). The MITRE Corp. (2004) report presented a tokenized access framework and an economic model for regulating the tokens. In their proposed approach, tokens may be viewed as cash and can be spent to access sensitive information. A fuzzy multi-level security (MLS) model based on probability was proposed by Cheng et al. (2007). Despite the name, the work is not based on fuzzy logic, but rather on a new probabilistic formulation of MLS model that supports quantified access decisions. Keppler et al. (2006) developed a Flexible Authorization Framework that redirects mission-related denied requests to corresponding entities who may serve as an override authority, enabling dynamic information sharing.

An alternative to these approaches is a trust inference mechanism that (1) is based on a comprehensive profile of a requester, (2) utilizes the digital credential infrastructure, (3) adapts to environments and (4) is rule-based. In this paper, we propose such a contingency trust inference model for crisis communication that supports flexible and secure information sharing across different administrative domains. Our goal is to support the automatic prediction of a requester's trustworthiness based on what is learned about the requester, including affiliation, identification, history and context. The information owner then determines the corresponding access privileges for the requester. The main strength of our proposed model in comparison with existing access control solutions is that we support *ad hoc* trust establishment by dynamically inferring access privileges without requiring any prior trust relationship between the requester and the resource owner. Our model allows a requester to obtain partial access to a resource belonging to another organization during emergency situations, while preserving the integrity of the shared resource. Our technique based on fuzzy logic facilitates cross-organizational information sharing and the completion of critical missions. For example, under our model, during crisis situations, a Federal Emergency Management Agency (FEMA) employee does not need to obtain the authorization letter from his supervisor, the process of which may be time-consuming, in order to gain (partial) access to U.S. Coast Guard (USCG) data or collaborate with local emergency management units.

Studies from crisis and emergency management have found that technologies can sometimes cause communication problems during crises (Chartrand, 1985; Korac-Boisvert & Kouzmin, 1994; Mitroff, 1994; Eriksson, 2001; Garnett & Kouzmin, 2007; US Senate, Committee on Homeland Security and Governmental Affairs, 2006; Heegaard & Trivedi, 2009). Conventional authorization systems are substantially designed to meet the need for intra-domain information access, e.g., requesters are typically employees of the organization. In crisis situations, however, the need for inter-organizational information sharing sharply increases, and access requests for sensitive data may come from outside the organization and from people who are not previously known. To meet the cross-domain information-sharing requirements, the *status quo* is that one or multiple administrators are usually needed to be involved to give specific permissions to the outside users.

Ideally, in crisis situations, exceptions may be made to normal access rules according to the specific conditions and scenarios. This step involves a logic process to evaluate the tradeoffs of the associated risks and benefits and is conventionally performed by a human administrator. For example, a USCG official will assess the urgency of crisis conditions and the benefits of sharing information with FEMA staff, in order to decide whether or not to share location information to FEMA. There may not be access rules defined for this unique situation; therefore, logical human judgement is typically required according to the following patterns. If the FEMA employees are trustworthy and the rescue missions are urgent, then FEMA personnel are allowed to access the location information of USCG units. Fuzzy logic systems can be used to define and automate this logic process and therefore is particularly useful for controlling information sharing in these open systems. A challenging aspect of this problem is to enable sharing in dynamic collaboration environments, such as sharing among first-response teams who are not previously known to each other. The sharing and control of access need to be established dynamically in response to the need of crisis communications.

In this paper, we design a fuzzy logic system for inferring trustworthiness for cross-domain information sharing in crisis situations. We identify and describe the key attributes involved in evaluating the trustworthiness of a requester, and define concrete membership functions for each fuzzy variable in the system. We illustrate the operations of aggregation and defuzzification for obtaining the final trust scores. We also design an audit mechanism for identifying cheating users (e.g., taking advantage of or manipulating context information) and fold the information into the trustworthiness computation to improve accountability. We propose to use a simple logging and auditing mechanism to monitor and adjust the accuracy of long-term trustworthiness predictions. Testing the model would come as a next step.

2. Fuzzy systems and logic

Fuzzy logic, unlike conventional crisp logic, is defined as the logic system that uses imprecise or uncertain inputs to infer outputs (Munakata and Jani, 1994). Fuzzy systems collectively refer to: *fuzzy sets* (sets whose elements have degree of membership), *logic* (a form of multi-valued logic), *algorithm* (an ordered set of instructions that yield an approximate solution to a specified problem) and *control* (a process including an input stage, a process stage and an output stage). The fundamental idea behind all fuzzy systems is: the transition from one output state (e.g., 0) to the other (e.g., 1) is gradual and continuous, which is contrary to abrupt and crisp changes between 0 and 1. The value of fuzzy systems was first proposed by Zadeh in 1965 (Zadeh, 1993, 1999). Fuzzy systems became widely used in commercial applications such as train operation systems (Yasunobu & Miyamoto, 1985), electronic appliances (Lee, 1990) and trading systems (Deboeck, 1994) in the late 1980s and early 1990 (Song, Hwang, and Kwok, 2005). Fuzzy systems have also been applied to terrorism and other crisis and homeland security issues (Ren & Liang, 2005).

In general, fuzzy systems can be used for approximate reasoning where the inputs and the parameters of a system are incomplete, inaccurate or imprecise. Fuzzy logic makes estimated decisions with inputs that have degrees of fuzziness, rather than trying to model the system mathematically. Intuitively, modelling problems with uncertainty is very costly, even if it is possible. By focusing what the system should do, existing applications of fuzzy logic take advantage of its efficacy and are usually less costly to compute than non-fuzzy methods. To develop a fuzzy logic system, one needs to identify the inputs and outputs and their ranges, define membership functions for the variables, construct fuzzy rule sets and fine-tune the systems. In this paper, we describe the application of fuzzy logic in access control that can increase the flexibility of access policies and enable information sharing across organizational domains. This advantage of fuzzy systems allows a requester to gain access to critical information controlled by another organization in emergency and crisis situations, avoiding the delay in requesting authorization in a conventional model. We also describe the mechanism in ensuring the trustworthiness of the requester in the process.

3. A contingency trust inference model

In this section, we design a contingency trust inference model that allows an information owner to infer the trustworthiness of a request. There are two main players in our contingency trust inference model: an *information owner* and a *requester.* We assume that the requester may be malicious and submitting false information to the information owner in order to gain access. We do not assume any prior trust relationships between the information owner and the requester, i.e., they may not know each other. The key point in our contingency trust inference model is that the trustworthiness is computed based on the profile of a requester, rather than from a single attribute. The profile of a requester captures several facets of the user or his or her organization. The elements in the user's profile are integrated using fuzzy logic rules and are collectively evaluated to make access decisions. As a result, the final access decision does not depend solely on any single input. The less rigid structure of fuzzy inference rules allows flexible-yet-controlled access decisions, namely, a partial access decision that is between 0 and 1.

3.1. Overview and set-up of trust inference model

Before an information owner performs contingency trust inference, it needs to go through a *set-up phase*. In the set-up phase, the information owner defines several important components of the fuzzy logic system including attributes, fuzzy variables, membership functions and fuzzy rules. The details of how to define these fuzzy logic components are described in the following sections.

(1) Define attributes from which trustworthiness may be inferred.
(2) Define the fuzzy variables associated with each attribute. See Table 1.
(3) For each fuzzy variable, define a membership function. See Section 3.3.
(4) Define the output membership function for the output variable (i.e., degrees of trustworthiness).
(5) Define fuzzy rules to specify the logic used to infer the trustworthiness score from attributes.

Before we proceed, a brief overview on the procedure of contingency trust inference may be helpful. Our contingency trust inference procedure is run by the information owner and consists of five main steps: Fuzzification, Rule Application, Aggregation, Defuzzification and Authorization. The inputs are *n* crisp values $(x_1,\ldots, x_n)$ where x_i in the interval [0,1], $1 \leq i \leq n$ is a numerical attribute value defined in Section 3.2. For the output, a crisp numerical value, in the interval [0,1], is computed representing the inferred trustworthiness score.

(1) *Fuzzification*: For each input, compute the degrees of membership based on the membership functions.
(2) *Rule Application*: Apply fuzzy logic rules to the inputs and obtain a conclusion for each applicable rule.
(3) *Aggregation*: Combine the conclusions into a logical sum.
(4) *Defuzzification*: Compute a firing strength for each output membership function. Combine these logical sums in a defuzzification process to produce a crisp trust score.
(5) *Authorization*: Determine the requester's information access level based on the computed trust score and the sensitivity of requested information.

In the following sections, we describe our contingency trust inference system in detail.

3.2. Attributes and fuzzy variables

Studies have been conducted across a variety of organizational and management situations that help to improve our understanding of trust as a component of decision-making. There is consensus that although the nature of trust varies across sectors and relationships, the underlying research and learning can be applied to practical applications of trust regardless of the industry or the sector. Therefore, we can examine trust as it applies to crisis and emergency situations by using the research conducted on the ways in which humans make decisions on whom to trust, the nature of trusting environments (Schoorman, Mayer, & Davis, 2007; Hurley, 2006) and how to develop a system to evaluate trust-related issues (Hofstede, 2007; Huber & McDaniel, 1986). Research has identified a number of attributes linked to trust in different organizational and inter-organizational contexts. We focus on three that have particular relevance for crisis management: *affiliation* (Hurley, 2006; Rousseau, Sitkin, Burt, & Camerer, 1998; Amason, 1996; Mayer, Davis, & Schoorman, 1995); *performance* (Davis, Schoorman, Mayer, & Tan, 2000); and *urgency* (Zand, 1972; Scott, 1980). The detailed description of these three attributes will be given below. Our trust inference model work is also related to the existing work on recommendation or reputation systems in decentralized models (Kohlas & Maurer, 2000). Trust evidences that are generated by recommendations and past experiences have been used for establishing trust in both *ad hoc* and ubiquitous computing environments (Eschenauer, Gligor, & Barras,

Table 1. Input and Output Attributes in the Contingency Trust Inference Model For Crisis Communication

Attribute	Type	Ranges	Fuzzy variables	Authentication method
Affiliation	Input	[0, 1]	Very high, high, medium, low, very low	Digital credentials
Task performance	Input	[0, 1]	Very high, high, medium, low, very low	Transaction monitoring
Urgency level	Input	[0, 1]	Very high, high, medium, low, very low	Audit mechanism
Trust score	Output	[0, 1]	Very high, high, medium, low, very low	–

Authentication methods refer to how to verify the correctness of attribute values.

2002; Shand, Dimmock, & Bacon, 2004; Theodorakopoulos & Baras, 2004; Covington, Ahamad, Essa, and Venkateswaran, 2004). This section will focus on trust through attributes and variables, which will be evaluated to assess an attribute value or trust score.

Our inference model for computing trustworthiness is based on three attributes associated with a request as shown in Table 1.

Definition 1: In our contingency trust inference model, an *attribute* describes a property of a request or the person who submits the request. In our model, an attribute takes on a numerical value (e.g., 0, .5 or 1) and is associated with several fuzzy variables, which are defined below. For example, attribute *urgency* may have a value of 1, which indicates a high degree of urgency. Besides a numerical value, attributes may be assigned a level, e.g., high, medium or low urgency.

Definition 2: In our contingency trust inference model, a *fuzzy variable* is a linguistic value (i.e., a word or a phrase and usually an adjective) that describes and characterizes the numerical attribute value. An attribute may have multiple fuzzy variables. Our contingency trust inference system has five fuzzy variables for all attributes and for the output: very high, high, medium, low and very low.

For example, our attribute *urgency* has five fuzzy variables: very high, high, medium, low and very low. As it will soon become clear, an attribute value will be mapped to several fuzzy variables (e.g., high, medium and low) according to membership functions.

The attributes we use in our model were derived from a meta-analysis of research on antecedents of trust in management situations (Mezgar & Kinsces, 2003; Rousseau et al., 1998; Mayer, Davis, & Schoorman, 1995; Jarvenpaa & Leidner, 1999) including those involving computer information-sharing applications (Kerr & Hiltz, 1982; Hart & Saunders, 1997). We have intentionally limited the attributes to represent the core aspects of our model for clarity of representation.

Affiliation is an attribute representing the trustworthiness of an organization to which an information requester belongs (Hurley, 2006; Rousseau et al., 1998; Amason, 1996; Mayer et al., 1995). Higher scores mean higher trustworthiness or a trustworthy relationship in the past. The score is determined based on the home organization (i.e., main affiliation) of a requester and the relationship standing of that organization with the information owner, according to a scale set by that owner. This attribute can be authenticated with digital credentials (e.g., role credentials) submitted by the requester. A default score may be given if the requester's home organization is unknown to the information owner. Audit results may be used to dynamically adjust affiliation scores assigned to organizations, and will be discussed in more detail later.

Complementary attributes to affiliation may be added to signify the trustworthiness of identity factors of an individual requester, such as rank (Jarman, Sproats, & Kouzmin, 2000). The sensitivity level of the requested information is not included in these attributes as it is independent of a request. However, the information owner should assess the sensitivity level when determining a requester's access authorization.

Task performance contains the information about a requester or his organization that is derived from the history of interactions with the information owner (Davis, Schoorman, Mayer, & Tan, 2000). Higher attribute values indicate higher or better previous performance. There are several methods to evaluate previous performance. For example, a simple approach is to compute (number of good transactions)/(number of bad transactions). Because of space constraints, we do not delve into this topic in our paper. If there is no prior interaction between the information owner and the requester, the information owner may assign a default value to this attribute. The previous transaction history is usually retained by the information owner and is not submitted by the requester. Therefore, the attribute value is computed by the information owner and there is usually no need to validate the attribute value.

This area may be expanded to capture more historical information on the requester. Individual relationship record and expertise level are two of the commonly discussed attributes that complement task performance (Mayer et al., 1995; Jarman, 2001). Individual relationship record (or personal reputation) is computed from previous experiences with the individual requester or gathered from peers of the information owner (regardless of organization affiliation), and is computed by the information owner. This attribute, along with expertise level, further evaluates the individual requester's past interactions and capabilities.

Urgency level is an attribute whose value is specified by the requester and defines how urgently a requester needs the information that has been requested (Zand, 1972; Scott, 1980). Higher attribute values mean higher urgency. Because the urgency level is *self-claimed*, it may or may not reflect the real situation (e.g., a user may falsely claim that his request is extremely urgent in order to receive a higher trust score and authorization), although prior research states that creating a trusting environment helps to curb possible abuses of the vulnerability of the information owner (Zand, 1972). To catch this type of exaggeration, our model requires an audit mechanism to monitor the truthfulness of self-claimed urgency levels and provides feedback to the trust inference process. For example, if a user or a group of users has been consistently exaggerating the urgency levels of requests, then this information will be incorporated into a previous performance attribute and affiliation score attribute. Thus, in future requests, prior exaggeration will be factored into the information request. Factors – such as the magnitude that releasing information would have on the crisis situation, the inclusion of the spatial relationship with the crisis and the creation of a dual-structure of routine and non-routine emergency situations (Huber & McDaniel, 1986) – may be included in computing the access decisions.

Our model does not rely on the organizations of the requester and resource owner having a prior trust relationship or interaction. In conventional access control models, an unknown requester will not be granted any access at all. In comparison, in our model, although the attribute values of *affiliation* and *task-performance* may be unknown, the final trust score may be non-zero as the computation integrates the requester's self-claimed and audited urgency-level attribute. This flexibility can significantly improve the information accessibility in crisis situations.

Formal policies in conventional access control models can be used to specify access requirements based on the aforementioned attributes. However, such a model would not be efficient, because the task of enumerating all possible access cases is nontrivial even for a small number of attributes. The advantage of using fuzzy logic for aggregating various attribute values in the trust computation is its simplicity and efficiency in modelling the access-control logic of resource owners.

The above attributes are factors to be used to determine a requester's trustworthiness. The output of a trust inference model is a *trust score*. The output is also associated with multiple fuzzy variables (e.g., {very high, high, medium, low, very low} ← in our model). As shown in Table 1, the score for the attribute affiliation can be a value between 0 and 1, and can be mapped to five fuzzy variables according to the membership functions of the fuzzy variables. The range is chosen arbitrarily in this paper but the range of five has common usage and interpretation among research scales. Membership functions are defined in order to fuzzify an attribute value to multiple fuzzy sets. A fuzzy rule set is also defined to infer a set of trustworthiness values of a requester from the fuzzified attribute values. The inferred trustworthiness values are then aggregated and defuzzified to obtain the final crisp score. More details of this process are described next.

3.3. Membership functions

In fuzzy theory, a membership function defines to what degree a variable belongs to a fuzzy set. Formally, a fuzzy set is defined as follows: the process of mapping a fuzzy variable to its membership of a fuzzy set is called *fuzzification*.

Definition 3: A fuzzy set is a pair (X, m) where X is a set and $m{:}X \rightarrow [0,1]$. For each $x \varepsilon X$, $m\ (x)$ is the degree of membership of x.

If an element is not included in the fuzzy set, then $m\ (x) = 0$; if it is a fully included member, then $m\ (x) = 1$. Fuzzy members are characterized by values that are between 0 and 1.

In our trust inference model, a membership function is defined for each fuzzy variable. Our model has five fuzzy variables {*very high, high, medium, low, very low*} for our three attributes.

There are several commonly used membership functions. For the ease of illustration, we choose a triangular-shaped membership function with a height of one as shown in Figure 1. Bell-shaped membership functions are also widely used in many fuzzy logic systems that yield nonlinear (e.g., quadratic) transitions between 0 and 1.

Once membership functions are defined for fuzzy variables, a crisp input can be fuzzified to obtain degrees of membership for all the fuzzy variables. While cross-domain, inter-organizational information sharing is needed in different types of crises such as natural and technological disasters, health crises, terrorism and other types, we can illustrate our methodological approach with a scenario drawn from the Hurricane Katrina crisis. While Katrina was extraordinarily complex, we choose a sensitive but more direct example to illustrate our model. A number of different organizations

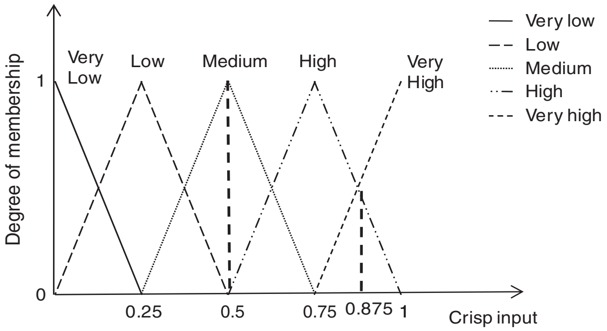

Figure 1. An example of triangular-shaped membership functions for five fuzzy variables {very high, high, medium, low, very low}. Bold dashed lines show that the centre point for medium's membership function is .5; for very high, it is .875.

were involved in search and rescue operations to find people trapped in homes, other buildings, on bridges and other places. Search and rescue operations are typically facilitated by effective communication about where search and rescue units are, where they are headed, who they are going to rescue, what their transportation capacity is and related information. Information sharing during Katrina, however, was often another victim. According to Cooper and Block, 'Throughout the disaster, state and federal agencies worked independently, under their own initiative, sometimes at cross-purposes. The Coast Guard was one of the worst offenders: FEMA officials would later say the agency [Coast Guard] did almost nothing to keep other units up to speed on its activities' (2006, p. 230). Allowing a FEMA official to access the real-time location information of USCG units without having to get bureaucratic clearances from both organizations would clearly save precious time and resources.

Example: USCG is the information owner. Davis is a FEMA employee who requests the location of certain USCG boats to coordinate rescue efforts. USCG assigns FEMA members the affiliation score of .8, indicating a high trust level presumably expected for another national government agency. If FEMA affiliation proves lest trustworthy as the crisis continues, this score could be readjusted lower. To illustrate another feature of our model, in the example, we assume that Davis has never requested information from USCG before (i.e., no previous transaction history). The USCG therefore assigns a default value .4 as the previous performance attribute. Davis claims that his request is urgent; the corresponding Coast Guard official agrees on 1 for the urgency-level attribute. Adjustments in the affiliation rating would be performed by Coast Guard central staff charged with auditing information transactions. Such adjustments could be made periodically as feedback about the quality of the transaction gets back to the auditors. For highly urgent requests, the Coast Guard would need to grant FEMA personnel access without waiting for adjusted ratings that would follow as soon as feasible.

Table 2. Examples of Membership Degrees

Attribute	Value	Very low	Low	Medium	High	Very high
Affiliation	.8	0	0	0	.8	.1
Task performance	.4	0	.3	.6	0	0
Urgency level	1.0	0	0	0	0	1.0

Using the membership functions in Figure 1, by looking up, we obtain the degrees of membership for each attribute, shown in Table 2. Once the degrees of membership of each crisp input are computed, fuzzy rules are to be applied as presented next.

3.4. Fuzzy rule sets

Fuzzy rule sets are defined in the IF-AND-THEN form as follows, where for each rule R^i, input fuzzy variables $x_1,\ldots,x_n$ are compared with pre-defined values $A^i{}_1,\ldots,A^i{}_n$, respectively, and the fuzzy output y is derived:

$$R^i : \text{IF}\, x_1 = A^i_1 \text{ AND}\, x_2 = A^i_2 \text{ AND } \ldots \text{ AND}\, x_n = A^i_n \text{ THEN}\, y = B_i$$

The rules and the number of rules to be defined may be based on the specific applications and administrative policies of the information owner. To illustrate how fuzzy rules can be defined for our attributes, we give several examples of contingency trust inference rules in a table format in Table 3.

In our set-up, each attribute including the output contains five fuzzy variables. In order to enumerate all the combinations of fuzzy variables, it requires a number of fuzzy rules. However, fuzzy logic systems do not require all possible rules to be explicitly defined. A very complex system may contain just a hundred rules. For our model, because the number of attributes is small, we expect that the number of rules in an actual prototype authorization system is manageable. We discuss this topic further in our future work in section 4.

Compared with conventional predicate-based logic rules, fuzzy rules are simple to define and intuitive to understand as they follow human logic. Such simplifications can reduce the management difficulty for large complex systems, which in turn reduces administrative mistakes. For example, plug in our membership values in Table 2 to the example rules given in Table 3; the fuzzy outputs are *very high* and *medium*. *Fuzzy output* refers to the output fuzzy variable corresponding to a rule that is fired or has a non-zero result. Next, we describe the

Table 3. Examples of Fuzzy Rules

Attribute name	Affiliation	Task performance	Urgency level	Output trust score
Rule R1	Very high	Medium	Very high	Very high
Rule R2	High	Low	Very high	Medium
Rule R3	Medium	High	Medium	High
Rule R4	Low	Low	Very high	Low
Rule R5	Very low	Medium	Very high	Low

aggregation and defuzzification steps, which produce the final output of our trust inference process.

3.5. Aggregation and defuzzification

The aggregation step is to determine the firing strength (degree of fulfilment) of each rule and to combine the logical products for each rule. There exist several aggregation methods and the choice may be up to the information owner. We illustrate the root-sum-square method in our paper because it is the most common and simplest aggregation method used in fuzzy systems. Because only rules R_1 and R_2 yield nonzero results, the output fuzzy variable can be very high or high, correspondingly. For fuzzy variable *very high*, the firing strength denoted by P_1 is computed as $(0.1^2+0.6^2+1.0^2)^{1/2}=1.1$. For fuzzy variable *medium*, the firing strength denoted by P_3 is computed as $(0.8^2+0.3^2+1.0^2)^{1/2}=1.3$.

For completeness, we give the general formula for computing the firing degree P_i of a fuzzy variable f_i using rootsum-square method in Equation 1, where P_i denotes the firing strength of f_i, n is the number of (input) fuzzy variables, k is the number of rules that yield f_i as the output response, and d_{ij} denotes the degree of membership of input variable x_i in rule R_j.

$$P_i = \sqrt{\sum_{j=1}^{k}\sum_{i=1}^{n} d_{ij}^2} \tag{1}$$

The defuzzification step is to compute a crisp output by combining inference results using a fuzzy centroid algorithm, as specified in Equation 2, where C_i denotes the center point of f_i's membership function (shown in Figure 2 by bold dotted lines), P_i denotes the firing strength of a fuzzy variable f_i.

$$\text{Output} = \frac{\sum_{i=1}^{n} C_i \times P_i}{\sum_{i=1}^{n} P_i} \tag{2}$$

Using the values from our example in the above formula, we obtain the output $(0.875\times 1.1+0.5\times 1.3)/(1.1+1.3)=0.6$ as the crisp output. Thus, the inferred final trust score is .6 in our example. This score would not indicate that automatic information transfer is warranted but would indicate a higher than average trust level. The information resource owner would then need to decide whether to grant the information based on the fuzzy logic result. The resource owner would need to define trust 'clearance levels' that specify what information could be shared with a requester whose crisp trust score or a trust range meets the specified level. For example, accessing a full document according to stated policy may require a trust score in the range of [.8, 1.0], but a redacted version of the document may be accessed by those whose trust scores are in the range of [.6, .8]. This type of partial access is not supported in conventional access control models making binary allow-or-deny decisions. We next turn to some issues involved in implementing this trust inference model.

3.6. Implementation issues

3.6.1. Auditing mechanism

How a user judges a transaction as a bad or a good transaction is usually specific to applications. For example, in peer-to-peer file-sharing applications, a correct download from a peer in a timely fashion can be counted as a good transaction. For access control and information-sharing scenarios such as those we study, judging a transaction as good or bad is based on whether a requester is truthful in submitting his or her attributes. We propose to use an auditing mechanism to selectively monitor the transactions and provide the feedback to the inference process.

Each administrative domain will deploy a domain-wide auditor that is capable of monitoring all the transactions associated with the resources controlled by the domain. Our contingency trust inference model requires an auditing component that aims to (1) deter requesters from lying about their environment attributes, (2) catch inconsistencies between the self-claimed urgency level and (3) propagate the auditing results back to identity and history attribute values.

The main task of the auditor is to monitor whether a requester exaggerates the urgency level associated with a request, which can be realized by the manual verification on randomly selected transactions. Crises have varying degrees of urgency from those that are fast terminating and fast developing to those that are slow developing and slow terminating ('t Hart & Boin, 2001). While it can be difficult to distinguish between levels of urgency, it is important to perceive patterns of exaggeration. The goal is to identify those who abuse our access mechanism, i.e., those who intentionally exaggerate their urgency level when requesting cross-organizational access. Whenever there are major or minor crisis events, the information associated with the event, including time, severity and location, is given to the auditor. The event's information will be used to map an urgency level that will be then used to compare with the self-claimed urgency level associated with past transactions. In general, the auditing service only needs to check transactions whose urgency levels are relatively high to catch any inconsistencies.

3.6.2. Transaction monitoring

Currently, our model considers the transaction history that contains only transactions of the information owner. In order to also consider the transaction history carried out with other nodes, a reputation model may be utilized and the computation for trust values needs to be adjusted accordingly. In principle, more data on the transaction history of a requester will provide higher accuracy in trustworthiness prediction. In decentralized environments, however, it is infeasible to gather all the available transaction history from all possible sources. One simple approach is to have a collaborative filtering mechanism where several organizations form a compact or a consortium to share the transaction histories of previous interactions. In the trust inference computations, additional attributes may be introduced to capture these factors. Such arrangements raise a privacy issue as the access history of an individual may be traced and analysed by clique members to infer additional knowledge, which would be impossible to obtain if the transaction histories are not shared. How to achieve privacy-preserving collaborative filtering in reputation systems remains a problem for further exploration. Because of the cost, effort and privacy sensitivities, highly systematized information sharing collaborations like this are only practical when interorganizational information transactions are high and/or of a critical nature. As disasters and crises are by their nature critical and organizational actors are often the same, our model has applicability in this context.

4. Conclusions and future work

We have described a contingency trust inference model where decisions about information access are adaptive to a requester's affiliation, past performance and request urgency. Our trust inference model is built on fuzzy logic that has the flexibility appropriate to crisis situations. Most importantly, fuzzy logic systems lend themselves to more balanced and comprehensive decision making that mimics the process of human thinking. Using soft computing techniques is a promising direction for flexible and controlled information sharing. There are exciting directions to pursue.

We intend to pursue two directions: fine-tuning our contingency trust inference model and making it more understandable and usable for organizational use. For future work, we plan to study the sensitivities of fuzzy logic components on the trust score computation. For example, other attributes such as organizational rank and connection security could broaden the analysis. It is also interesting to investigate and experiment various membership functions and fuzzy rule definitions in our model, in order to identify the impacts of parameter changes on the final decision-making process. Another important problem to study is how to integrate the contingency trust inference system with predicate logic-based access control systems, in order to achieve smooth transitions between the two systems under normal and crisis situations. We would also like to explore how the model would withstand a stress test – an attack by purposefully injecting wrong attribute values. We expect to use a combination of statistical and cryptographic techniques to address this issue. In terms of improving usability, we plan to develop spreadsheet-like programmes that allow users to enter data on a terminal or a handheld device and have the trustworthiness score computed for quick and easy use. As this model is designed to mimic the process of human thinking, it would be a natural move to devise a user interface that can interactively collect information from users and give back a score with a more elaborate explanation. This could be an application that can run on a hand-held device, and, by applying technologies, such as text-to-speech and speech recognition, respond on the phone for the convenience of people in the field. Applying methodologies already used in other contexts can aid in managing crises, arguably the most challenging form of management.

Acknowledgements

This research was supported by a Pervasive Computing grant from Rutgers University.

References

Alink, F., Boin, A. and 't Hart, P. (2001), 'Institutional Crises and Reforms in Policy Sectors: The Case of Asylum Policy in Europe', *Journal of European Public Policy*, Volume 8, Number 2, pp. 286–306.

Amason, A.C. (1996), 'Distinguishing the Effects of Functional and Dysfunctional Conflict on Strategic Decision Making: Resolving a Paradox for Top Management Teams', *Academy of Management Journal*, Volume 39, Number 1, pp. 123–148.

Boin, A. and Lagadec, P. (2000), 'Preparing for the Future: Critical Challenges in Crisis Management', *Journal of Contingencies and Crisis Management*, Volume 8, Number 4, pp. 185–191.

Chartrand, R.L. (1985), 'The many Potentials of Information Technology for Emergency Management', *The Information Society*, Volume 3, Number 4, pp. 275–289.

Cheng, P., Rohatgi, P., Keser, C., Karger, P.A., Wagner, G.M. and Reninger, A.S. (2007), *Fuzzy Multi-Level Security: An Experiment on Quantified Risk-Adaptive Access Control*. Proceedings of the 2007 IEEE Symposium on Security and Privacy (May 20–23, 2007). SP. IEEE Computer Society, Washington, DC, pp. 222–230.

Comfort, L.K. (2007), 'Crisis Management in Hindsight: Cognition, Communication, Coordination, and Control', *Public Administration Review*, Volume 67, Suppl. 1, pp. 189–197.

Comfort, L.K., Ko, K. and Zagorecki, A. (2004), 'Coordination in Rapidly Evolving Disaster Response Systems: The Role of Information', *American Behavioral Scientist*, Volume 48, Number 3, pp. 295–313.

Covington, M.J., Ahamad, M., Essa, I. and Venkateswaran, H. (2004), *Parameterized Authentication*. Lecture Notes in Computer Science, pp. 276–292.

Davis, J.H., Schoorman, F.D., Mayer, R.C. and Tan, H.H. (2000), 'The Trusted General Manager and Business Unit Performance: Empirical Evidence of a Competitive Advantage', *Strategic Management Journal*, Volume 21, Number 5, pp. 563–76.

Deboeck, G. (1994), *Trading on the Edge: Neural, Genetic, and Fuzzy Systems for Chaotic Financial Markets*, New York: Wiley.

Derthick, M. (2007), 'Where Federalism Didn't Fail', *Public Administration Review*, Volume 67, Suppl. 1, pp. 36–47.

Drabek, T.E. (1986), *Human System Responses to Disaster: An Inventory of Sociological Findings*, New York: Springer-Verlag.

Eriksson, J. (2001), 'Cyberplagues, IT, and Security: Threat Politics in the Information Age', *Journal of Contingencies and Crisis Management*, Volume 9, Number 4, pp. 211–222.

Eschenauer, L., Gligor, V.D. and Barras, J. (2002), *On Trust Establishment in Mobile Ad-Hoc Networks*, Proceedings of the Security Protocols, 10th International Workshop, Cambridge, UK, April 17–19, 2002.

Fischer Iii, H.W. (1999), 'Enhancing Disaster Mitigation Planning and Response Through the Use of Cyberspace: Suggestions and Issues to Consider', *Journal of Contingencies and Crisis Management*, Volume 7, Number 1, pp. 48–54.

Garnett, J.L. and Kouzmin, A. (2007), 'Communicating throughout Katrina: Competing and Complementary Conceptual Lenses on Crisis Communication', *Public Administration Review*, Volume 67, Suppl. 1, pp. 171–188.

't Hart, P. and Boin, A. (2001), 'Between Crisis and Normalcy: The Long Shadow of Post-Crisis Politics', in Rosenthal, U., Boin, R.A. and Comfort, L.K. (Eds), *Managing Crises: Threats, Dilemmas, and Opportunities*, Springfield, Illinois.

't Hart, P. and Saunders, C. (1997), 'Power and Trust: Critical Factors in the Adoption and Use of Electronic Data Interchange', *Organization Science*, Volume 8, Number 1 (January–February 1997), pp. 23–42.

Heegaard, P.E. and Trivedi, K.S. (2009), 'Network Survivability Modeling', *Computer Networks*, Volume 53, Number 8, pp. 1215–1234.

Hofstede, G.J. (2007), 'Trust and Transparency in Netchains: A Contradiction?', in Wang, W.Y.C., Heng, M.S.H. and Chau, P.Y.K. (Eds), *Supply Chain Management: Issues In The New Era of Collaboration and Competition.*, Idea Group Publishing, Hershey, PA, USA, pp. 105–126.

Huber, G.P. and McDaniel, R.R. (1986), 'The Decision-Making Paradigm of Organizational Design', *Management Science*, Volume, 32, Number 5, pp. 572–589.

Hurley, R.F. (2006), 'The Decision to Trust', *Harvard Business Review*, Volume 84, Number 9, pp. 55–62.

Jarman, A. (2001), 'Reliability' Reconsidered: A Critique of the HRO-NAT Debate', *Journal of Contingencies & Crisis Management*, Volume 9, Issue 2, pp. 98–107.

Jarman, A., Sproats, K. and Kouzmin, A. (2000), 'Crisis Management: Toward a New Informational "Localisim" in Local Government Reform', *International Review of Public Administration*, Volume 5, Number 2, pp. 81–97.

Jarvenpaa, S.L. and Leidner, D.E. (1999), 'Communication and Trust in Global Virtual Teams', *Organization Science*, pp. 791–815.

Keppler, D., Swarup, V. and Jajodia, S. (2006), *Redirection Policies for Mission-Based Information Sharing*, Proceedings of the Eleventh ACM Symposium on Access Control Models and Technologies, Lake Tahoe, CA, USA, pp. 210–218.

Kohlas, R. and Maurer, U. (2000), *Confidence Valuation in a Public-Key Infrastructure Based on Uncertain Evidence*, Proceedings of the Third International Workshop on Practice and Theory in Public Key Cryptography (PKC '00). Lecture Notes in Computer Science, Melbourne, Victoria, Australia, pp. 93–112.

Korac-Boisvert, N. and Kouzmin, A. (1994), 'The Dark Side of Info-Age Social Networks in Public Organizations and Creeping Crises', *Administrative Theory and Praxis*, Volume 16, Number 1, pp. 57–82.

Lee, C.C. (1990), 'Fuzzy Logic in Control Systems: Fuzzy Logic Controller. II', *IEEE Transactions on Systems, Man and Cybernetics*, Volume 20, Number 2, pp. 419–435.

Mayer, R.C., Davis, J.F. and Schoorman, F.D. (1995), 'An Integrative Model of Organizational Trust', *The Academy of Management Review*, Volume 20, Number 3, pp. 709–734.

Mezgár, I. and Kincses, Z. (2003), 'The Role of Trust in Information Technology Management', in Gunasekaran, A., Khalil, O. and Rahman, S.M. (Eds), *Knowledge and Information Technology Management: Human and Social Perspectives*, Idea Group Inc., Hershey, pp. 283–304.

MITRE Corp. (2004), *Horizontal Integration: Broader Access Models for Realizing Information Dominance*, JASON Program Office, http://www.fas.org/irp/agency/dod/jason/classpol.pdf (accessed 2 July 2009).

Mitroff, II. (1994), 'The Role of Computers and Decision Aids in Crisis Management: A Developer's Report', *Journal of Contingencies and Crisis Management*, Volume 2, Number 2, pp. 73–84.

Munakata, T. and Jani, Y. (1994), 'Fuzzy Systems: An Overview', *Communications of the ACM*, Volume 37, Number 3, pp. 68–76.

Newkirk, R.T. (1993), 'Extending Geographic Information Systems for Risk Analysis and Management', *Journal of Contingencies & Crisis Management*, Volume 1, Number 4, pp. 203–206.

Palm, J. and Ramsell, E. (2007), 'Developing Local Emergency Management by Co-Ordination between Municipalities in Policy Networks: Experiences from Sweden', *Journal of Contingencies and Crisis Management*, Volume 15, Number 4, pp. 173–182.

Ren, Q. and Liang, Q. (2005), *Fuzzy Logic-Optimized Secure Media Access Control (FSMAC) Protocol Wireless Sensor Networks*, Proceedings of the 2005 IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety, 2005. CIHSPS 2005, Orlando, FL, USA, pp. 37–43.

Rousseau, D.M., Sitkin, S.B., Burt, R.S. and Camerer, C. (1998), 'Not so Different After All: A Cross-Discipline View of Trust', *Academy of Management Review*, Volume 23, Number 3, pp. 393–404.

Schoorman, F.D., Mayer, R.C. and Davis, J.H. (2007), 'An Integrative Model of Organizational Trust: Past, Present, and Future', *The Academy of Management Review (AMR)*, Volume 32, Number 2, pp. 344–354.

Scott, C.L. (1980), 'Interpersonal Trust: A Comparison of Attitudinal and Situational Factors', *Human Relations*, Volume 33, Number 11, pp. 805–812.

Shand, B., Dimmock, N. and Bacon, J. (2004), 'Trust for Ubiquitous, Transparent Collaboration', *Wireless Networks*, Volume 10, Number 6, pp. 711–721.

Singh, S., Sanders, W.H., Nicol, D.M. and Seri, M. (2006), *Automatic Verification of Distributed and Layered Security Policy Implementations*, 3rd Midwest Security Workshop (MSW), Purdue University, USA.

Song, S., Hwang, K. and Kwok, Y.K. (2005), 'Trusted Grid Computing with Security Binding and Trust Integration', *Journal of Grid Computing*, Volume 3, Number 1, pp. 53–73.

Swarup, V., Seligman, L. and Rosenthal, A. (2006), 'Specifying Data Sharing Agreements', *MITRE Paper MP 06W0000065*, The MITRE Corporation, McLean, VA 22101.

Tamassia, R., Yao, D. and Winsborough, W.H. (2004), *Role-Based Cascaded Delegation*, Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies (SACMAT '04), ACM Press, Yorktown Heights, NY, June 2004, pp. 146–155.

Theodorakopoulos, G. and Baras, J.S. (2004), *Trust Evaluation in Ad-Hoc Networks*, Proceedings of the 3rd ACM Workshop on Wireless Security, (Philadelphia, PA, USA, October 01, 2004). WiSe '04 ACM, New York, NY, pp. 1–10.

Torrieri, F., Concilio, G. and Nijkamp, P. (2002), 'Decision Support Tools for Urban Contingency Policy: A Scenario Approach to Risk Management of the Vesuvio Area in Naples, Italy', *Journal of Contingencies and Crisis Management*, Volume 10, Number 2, pp. 95–112.

Turoff, M., Hiltz, S.R. and Kerr, E.B. (1982), *Controversies in the design of computer-mediated communication systems: A Delphi study.* Proceedings of the 1982 Conference on Human Factors in Computing Systems (Gaithersburg, Maryland, United States, March 15–17, 1982). CHI '82. ACM, New York, NY.

US Senate, Committee on Homeland Security and Governmental Affairs. (2006), *Hurricane Katrina: A Nation Still Unprepared*, Washington, DC: Government Printing Office.

Waugh, W.L. Jr. and Sylves, R.T. (07/02 2002), 'Organizing the War on Terrorism', *Public Administration Review*, Volume 62, Number 4, pp. 145–153.

Wise, C.R. (May/June 2006), 'Organizing for Homeland Security After Katrina: Is Adaptive Management what's Missing?', *Public Administration Review (Washington, D.C.)*, Volume 66, Number 3, pp. 302–318.

Yao, D., Frikken, K.B., Atallah, M.J. and Tamassia, R. (2006), "Point-Based trust: Define How Much Privacy is Worth', in Ning, P., Qing, S. and Li, N. (Eds), *Proceedings of the International Conference on Information and Communications Security (ICICS). Lecture Notes in Computer Science 4307*, Springer, Raleigh, NC, USA, p. 190.

Yao, D., Tamassia, R. and Proctor, S. (2005), *On Improving the Performance of Role-Based Cascaded Delegation in Ubiquitous Computing*, First International Conference on Security and Privacy for Emerging Areas in Communications Networks, 2005. SecureComm 2005, IEEE Press, Athens, Greece, pp. 157–170.

Yasunobu, S. and Miyamoto, S. (1985), 'Automatic Train Operation System by Predictive Fuzzy Control' in Sugeno, M. (Ed.), *Applications of Fuzzy Control*, Elsevier: North-Holland, pp. 1–18.

Zadeh, L.A. (1993), 'Fuzzy Logic, Neural Networks and Soft Computing', *In Communications of the ACM*, Volume 37, Number 3, pp. 77–84.

Zadeh, L.A. (1999), 'Fuzzy Logic = Computing with Words', *Computing with Words in Information/Intelligent Systems*, Volume 1, pp. 3–23.

Zand, D.E. (1972), Trust and Managerial Problem Solving, *Administrative Science Quarterly*, Volume 17, Number 2, pp. 229–239.