Application of Fuzzy Logic in Federated Trust Management for Pervasive Computing

Zhengping Wu and Alfred C. Weaver
Department of Computer Science, University of Virginia
151 Engineer's Way, P.O. Box 400740, Charlottesville, VA 22904
{zw4j,weaver}@cs.virginia.edu

Abstract

In federated and pervasive networks, trust management has become a cornerstone for information security and privacy. Although people have recognized the importance of privacy and security for their personal information, they remain uncertain when they have to define and enforce their own access control rules or have to handle indirect information. Indirect information and subjective judgment are the major sources of uncertainty in federated trust management. This paper introduces fuzzy logic into the definition and evaluation of trust, and then provides a formal representation of fuzzy rules. It also offers a set of derivation rules for analyzing and reasoning among fuzzy rules in order to enforce these rules with a certain level of uncertainty. Application of this model to a healthcare environment with pervasive computing devices across trust domains provides a new method to handle uncertainty in trust management for federated and pervasive networks.

1. Introduction

Trust management is defined by M. Blaze, et al. [1] as "a unified approach to specifying and interpreting security policies, credentials, and relationships which allow direct authorization of security-critical actions" for the first time. With the application of trust management in research for network security, a more general definition is proposed by Grandison [2]: "Trust management is the activity of collecting, encoding, analyzing and presenting evidence relating to competence, honesty, security or dependability with the purpose of making assessments and decisions regarding trust relationships." Expanded from network security research, trust management has been studied in the context of access control [3], public key architecture [4], and reputation systems for peer-to-peer networks [5,6]. Meanwhile industries that require collaboration and sharing (e.g. healthcare, manufacturing, financial services, government services) will also benefit from trust management. For instance, HIPAA [7] requires rigorous privacy and security protection for medical data in healthcare systems. Thus trust management for hospital administration requires being able to create and enforce certain rules to secure healthcare data, such as "public web sites are permitted anonymous access," "patients may access their own records with password protection," "modification of patient data requires fingerprint verification of the physician," "deletion of any patient data can only be done by a medical records employee and only after iris scan verification authentication." Trust management enables us to increase the security and privacy of our shared resources and collaborated activities without increasing our workload. For instance, trust management systems can enforce the tasks stated in the above hospital example and yet not interfere with normal operations.

Federation is the current and future direction of trust management. It is an expansion of local infrastructure to an enterprise-wide and even global one and intensifies the demand for integrating inter-domain networks and services. With the increased flexibility that distributed yet interconnected networks such as the Internet and pervasive computing environments require, trust management becomes more and more important for federation among networks. Federated trust management incorporates not only internal factors such as intra-domain regulations, but also external factors such as reputations and recommendations into the process of formation of trust and enforcement of trust. Federated trust management can be defined as a unified approach to managing a collection of trust-related activities across multiple and heterogeneous security domains and autonomous systems. Because federation lacks central control and the partners in federation are not all predetermined, federated trust management is required.

To manage a collection of trust-related activities across multiple networks or domains, we need to provide flexibility in our enforcement mechanism for trust policies. For example, we cannot simply reject a cross-domain access request if we cannot find a matched policy. In this situation, we need to discover the intention of that request. If that intention can be proved to comply with another policy in a changed format or some combination of policies, we should allow that request and add a new policy for that request. This introduces fuzziness into federated trust management, because the intention of the request may be fuzzy, or the request itself may be fuzzy, or the policies to be enforced may be fuzzy also. Applying fuzzy logic into federated trust management can

help us handle trust-related activities with a certain level of uncertainty in a federated or pervasive network environment.

2. Related work

Trust is a complex subject relating to belief in the honesty, trustfulness, competence, and reliability of a person or service. In D. Harrison McKnight et al.'s "The Meanings of Trust" [8], the most tangible aspects of trust are trusting behavior and trusting intention. In the context of pervasive computing, trust is usually specified in terms of a relationship between a resource or service requester and a resource or service provider. Trust forms the basis for allowing a requester to use services or manipulate resources owned by a service provider, or it may influence a requester's decision to use a service or resource from a provider. So trust is an important factor in the decision-making process [9]. In many business relationships, trust is based on a combination of judgments or opinions from face-to-face meetings and recommendations of colleagues, friends, and business partners, which involve uncertainties.

Trust management can be considered as a collection of trusting behaviors, which includes capture, evaluation and enforcement of trusting intentions, and trust management systems enable us to increase the security and privacy of our federation activities without increasing our workload. After Blaze et al. introduced the trust management concept for the first time, they designed and developed two trust management systems, PolicyMaker [1] and KeyNote [10], with different emphases. These trust management systems are restricted to intra-domain trusting behavior or else only partially solve inter-domain trust related problems. It is these inter-domain problems that require further research in federated trust management systems.

Beth et al. [11] categorize the inter-domain trust relationships into two classes: direct trust and recommended trust. Based on the expectation for an entity being able to finish a task, the system can calculate the probability of whether the entity will finish the task based upon positive and negative experience, measure the trustworthiness using this probability, and create a formula for calculating a number value of the trustworthiness with a set of derivation and integration rules. But this mechanism simplifies real life by modeling trustworthiness based only on probability, and equals the subjectivity and uncertainty to the randomness. At the same time, it uses the mean value of multiple sources of trustworthiness as the indicator of the integrated trust and final trust value number, which omits possible weight on each trust source. In [12], Josang proposed a trust model based on subjective logic [13], which introduces the concepts of evidence space and opinion space to describe and measure trustworthiness. Based upon the Beta distribution function that describes the posteriori probability for binary events, the author calculates the trustworthiness for every possible event from every entity. Meanwhile, Josang defines a set of operators for the calculation of trustworthiness. Josang's model literally equals the subjectivity and uncertainty to the randomness also. But as a cognitive activity, the subjectivity and uncertainty of trust is mainly expressed in its fuzziness. How to model this fuzziness and apply this model to federated trust management is the problem.

In the following sections of this paper, we will identify different kinds of trust, find a suitable categorization for uncertainty from subjective judgment and indirect information sources, and represent uncertainty by fuzzy rules with a set of operations and derivation rules for decision-making and enforcement processes. We will also apply this fuzzy logic approach to a healthcare environment to handle uncertainties from real life. We assume that the authenticity of identities and the integrity of trust-related information can be guaranteed by the application itself, and this is outside the scope of this paper.

3. Classification of trust

3.1. Direct and indirect trust

To manage a collection of trust-related activities across domains, we need to understand trust itself. From different points of views, trust can be categorized into different classes. Following the categorization described by Beth et al. [11], we categorize trust into two classes - direct trust and indirect trust. A trust relationship formed from direct experience or negotiations can be characterized as direct trust; a trust relationship or a potential trust relationship built from recommendations by a trusted third party or a chain of trusted partied, which create a trust path, is called indirect trust. For example, suppose Dr. Jones needs to perform a clinical test, and she asks Dr. Smith for his advice about where to find a good hospital technician. Smith is thus directly trusted by Jones to know about a good technician and to provide his honest opinion. If in another scenario, Smith actually trusted the technician based on his own experience, but the technician happened to know very little about the test Jones needed, then Jones's trust of the recommended technician will not be so positive, because this indirect trust path is linked by a very positive trust from Smith to the technician and a not-so-positive trust from Jones to Smith. Figure 1 illustrates these two scenarios.

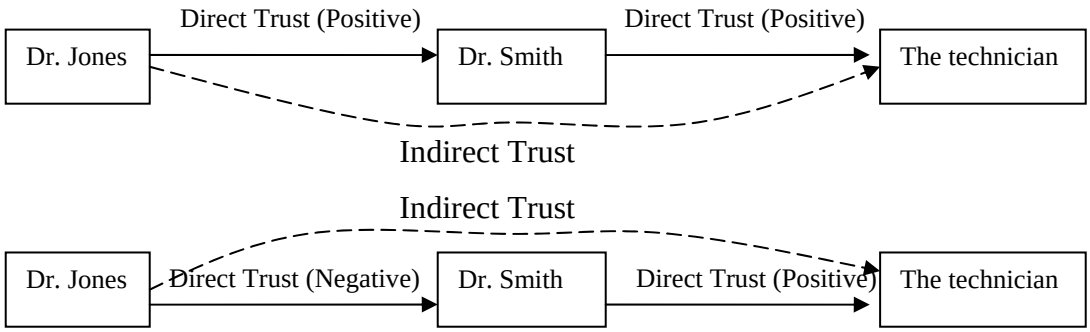

Figure 1. Direct Trust and Indirect Trust

In this example, we see that indirect trust is derived from direct trust. Indirect trust is a function of direct trust(s). It may add new values to direct or indirect trust from a trusted party. The added values are uncertain at some level and tend to be fuzzy. For example, if Jones is not sure about what Smith recommends, and Smith is not sure about the expertise of the technician, Jones will be totally uncertain about the expertise of the recommended technician. And the decision Jones makes is based on a fuzzy idea or we can say fuzzy evidence.

3.2 Objective and subjective trust

From another point of view, trust is a concept everybody understands at some personal level, but most people will have trouble providing a specific definition of the concept. Some people will have objective measures they use to evaluate their level of trust in a person or company, while others rely on a more subjective feeling for determining whether to trust somebody. So trust can be either derived from one's belief/feeling or based on an evaluation of certain measurements.

An entity's trustworthiness is associated with the quality of services it provides to others. If the quality of a service can be objectively measured, then the trust relationship relying on that service is objective. If the quality of services provided by that entity could be objectively measured, then the trustworthiness of that entity is objective. For example, in a telemedicine system the accuracy of a digital image provided by another healthcare provider can be indisputably checked against the original data. So the quality of that information service can be measured objectively, and the trust based on that information service is objective.

For some other services, their quality cannot be objectively measured. For example, different doctors may have different opinions about the interpretation of certain digital images. That result is informed from the doctors' subjective judgments after reading the image, combined with their previous experience and cases. It depends heavily upon each doctor's experience, feel and other subjective factors. The trust relationships relying on that kind of subjective quality information are subjective, and are called subjective trusts.

Intuitively, if the quality of a service can be objectively measured, then an entity's trustworthiness for that service reflects some intrinsic property of that entity, which should be independent of the source of the trust evaluation. An entity's subjective trust, however, may vary greatly when different sources of trust evaluation are considered. Due to this variation, subjective trust is uncertain at some level, and therefore needs special representations and enforcement processes to handle this aspect of federated trust management.

4. Fuzzy representation

4.1 Representation of uncertainty in trust

The trust relationships in pervasive computing environments and cross-trust domains are hard to assess due to the uncertainties involved. Figure 2 illustrates the comparison of different sources of uncertainties in federated trust management. If a trust relationship relies upon a subjective judgment based on indirect information, it will be very uncertain and any operations related to that trust relationship may cause unexpected results.

	Direct Information	Indirect Information
Objective Judgment	Certain	Uncertain
Subjective Judgment	Uncertain	Most uncertain

Figure 2. Comparison of Different Sources of Uncertainties

Fuzzy logic is a suitable way to represent uncertainties, especially when they need to be handled quantitatively. Two advantages of using fuzzy logic to quantify uncertainty in federated trust management are:

(1) Fuzzy inference is capable of quantifying imprecise data or uncertainty in measuring the level of trust.

(2) Different membership functions and inference rules could be developed for different trust relationships, without changing the fuzzy inference engines and enforcement mechanisms.

L. Zadeh first introduced fuzzy logic in the development of the theory of fuzzy sets [2]. The theory of fuzzy logic extends the ontology of mathematical research to be a composite which leverages of quality and quantity, which contains certain fuzziness. Introducing fuzzy logic into the research of trust management, we try to solve the issues associated with uncertainty in federated trust management. First, we need to identify the subjects of those issues. These subjects are either the sources of trust-related information needed in federated trust management or the entities with which trust relationships are built. This subject set can be defined as follows.

Definition 4.1 Set of subjects in federated trust management

The set of subjects in federated trust management is all the subjects that are either the sources of trust-related information or are the entities with which trust relationships are built. This set is represented as X *in this paper.*

Then we need to define a general fuzzy set in federated trust management.

Definition 4.2 Fuzzy set for federated trust management

For every element x *in the set of subjects* X*, there is a mapping* $x \mapsto \delta(x)$*, in which* $\delta(x) \in [0,1]$*. The set* $\Delta = \{(x, \delta(x))\}$ *for* $\forall x \in X$ *is defined as a fuzzy set for federated trust management.* $\delta(x)$ *is defined as the membership function for every* x *in* Δ*.*

All the fuzzy sets on X are represented as $Z(X)$. Then we can use a group of fuzzy sets from $Z(X)$ to group all the elements of X into several sets with different levels of uncertainty. For example, we can use a group of five sets $z_i \in Z(X)$ to categorize of uncertainty in federated trust management.

z_1 represents definitely uncertain;

z_2 represents probably uncertain;

z_3 represents equivocal;

z_4 represents probably certain;

z_5 represents definitely certain.

In real life, the level of uncertainty cannot be limited to only one set, and the degrees to these sets are not simply 'total' or 'none'; additionally, it is sometimes difficult to determine which set or sets should be used for certain kind of uncertainty. In other words, these sets are not exclusive to each other. So when we deal with certain kinds of uncertainty, a vector consisting of the degrees of belongingness to each set $D = \{d_1, d_2, d_3, d_4, d_5\}$ is more appropriate for describing the actual judgment from daily life, in which $d_i (i = 1,2,...,5)$ is the degree of belongingness to set $z_i (i = 1,2,...,5)$. Meanwhile, there are several ways to determine or calculate the degrees d_i. One way is direct judgment that determines the degree from direct experience or evaluation. Another one is indirect inference that determines the degree via an analysis of an indirect source such as reputation or recommendation. The first one is relatively subjective while the evaluation method may be very objective; and the second one is relatively objective while the source of information may be subjective. Other ways to determine the degrees also exist, which will not be discussed in this paper.

4.2 Formal representation

To reason among the degrees of uncertainty in federated trust management for further inference or decision-making, we need to represent uncertainty formally. Direct trust is formally described as $a \xrightarrow{D} b[Z]$, which means entity a is willing to rely upon entity b to degree D for the categorized uncertainty Z. D is a vector with corresponding degrees of belongingness for each set in categorization Z. Direct trust is from direct experience of the trustworthiness of the other entity or from a judgment with subjective/objective evaluation. Indirect trust is described as $a \xrightarrow[P]{D} b[Z]$, which means entity a is willing to rely upon b to degree D following P's recommendation for the categorized uncertainty Z. P is one or more entities constructing a path that gives a

recommendation to entity a for entity b. D is a vector with corresponding degrees of belongingness for each set in categorization Z. Indirect trust is derived from the recommendation passed through one or more intermediate entities. There are also two types of recommendations. One type is that the recommender had direct experience with the recommended entity so that the P has only one entity; the other is that the final recommender formed the recommendation from further recommendations of other recommenders so that the P has more than one entity constructing a chained recommending path or a compound recommending graph. But from the recommendee's (entity a's) point of view, there is no big significance related to with the number of entities forming the recommending path; the recommendee (entity b) only cares about the final recommender's capability to make accurate recommendation based on its own experience and trustworthiness.

5. Fuzzy enforcement

5.1 Fuzzy operations with adjustable parameters

Currently, most people use Zadeh operators $\wedge$ and $\vee$ to perform calculation and analysis with fuzzy logic. But these operators are too imprecise that too much information will be lost if using them only. Thus several general class fuzzy operators are proposed [14]. To adapt to different sources of uncertainties in federated trust management, a parameterized general intersection operator and union operator are needed. They are also called T-norm and S-norm. With different values of the parameters, these operators can maximize the expressiveness and flexibility of the system to capture people's intentions towards these uncertainties. Here we choose a general class of parameterized fuzzy operators proposed by Dubois and Prade [15] to perform further calculation and analysis. Because these operators are suitable for policy analysis and have clear semantic meanings [15], the intention embedded in fuzzy sets can be easily enforced. So we define T-norm and S-norm as follows.

Definition 5.1 T-norm

For fuzzy set $A, B \in Z(X)$ *and* $\alpha \in [0,1]$,

$$(A \cap B)(x) = T(A(x), B(x), \alpha) = \frac{A(x)B(x)}{\max\{A(x), B(x), \alpha\}}$$

, in which $A(x)$ *and* $B(x)$ *represent* x *'s degrees of member function to fuzzy sets* A *and* B.

Definition 5.2 S-norm

For fuzzy set $A, B \in Z(X)$ *and* $\alpha \in [0,1]$,

$$(A \cup B)(x) = S(A(x), B(x), \alpha) = \frac{A(x) + B(x) - A(x)B(x) - \min\{A(x), B(x), (1-\alpha)\}}{\max\{1 - A(x), 1 - B(x), \alpha\}}$$

, in which $A(x)$ *and* $B(x)$ *represent* x *'s degrees of member function to fuzzy sets* A *and* B.

Then we can define two calculators on vectors of fuzzy values. Suppose we have two fuzzy value vectors $D_1 = \{d_{11}, d_{12}, ..., d_{1P}\}$ and $D_2 = \{d_{21}, d_{22}, ..., d_{2P}\}$. We define "connection" and "union" calculators as follows.

Definition 5.3 Connection calculator

$$D_1 \otimes D_2 = \{T(d_{11}, d_{21}, \alpha), T(d_{12}, d_{22}, \alpha), ..., T(d_{1P}, d_{2P}, \alpha)\}$$

Definition 5.4 Union calculator

$$D_1 \oplus D_2 = \{S(d_{11}, d_{21}, \alpha), S(d_{12}, d_{22}, \alpha), ..., S(d_{1P}, d_{2P}, \alpha)\}$$

After we define the above calculators, we can perform formal analysis on fuzzy sets and fuzzy rules used for uncertainty expressions. Here we define two sets of derivation rules (deduction rules and consensus rules) to handle different types of uncertainty from different trust relationships and different recommenders.

Deduction rules are used for a recommendation's connection to construct a whole recommendation chain that allows the trust to be transferred from one end to the other end. For the trust relationships from the same categorization, deduction rules can form a new connection using the trust relationship between the recommender and the recommendee and embed the content of that recommendation into the new connection. Below are the formal descriptions of deduction rules.

Definition 5.5 Deduction rules

$$a \xrightarrow{D} b[Z] \wedge b \xrightarrow{D'} c[Z] \Rightarrow a \xrightarrow[P']{D'} c[Z] \wedge (P' = \{b\}) \wedge (D' = D \otimes D')$$

$$a \xrightarrow{D} b[Z] \wedge b \xrightarrow[P']{D'} c[Z] \Rightarrow a \xrightarrow[P']{D'} c[Z] \wedge (P' = \{b, P'\}) \wedge (D' = D \otimes D')$$

$$a \xrightarrow[P]{D} b[Z] \wedge b \xrightarrow[P']{D'} c[Z] \Rightarrow a \xrightarrow[P']{D'} c[Z] \wedge (P' = \{P, P'\}) \wedge (D' = D \otimes D')$$

Consensus rules are used for combining of multiple recommendations for the same kind of categorization. When two or more recommendation paths appear simultaneously, consensus rules can synthesize the opinions to form a comprehensive recommendation. Below are the formal descriptions of consensus rules.

Definition 5.6 Consensus rules

$$a \xrightarrow{D_1} b[Z] \wedge a \xrightarrow{D_2} b[Z] \wedge .. \wedge a \xrightarrow{D_n} b[Z] \Rightarrow a \xrightarrow{D'} c[Z] \wedge \left(D' = D_1 \oplus D_2 \oplus .. \oplus D_n\right)$$

$$a \xrightarrow[P_1]{D_1} b[Z] \wedge a \xrightarrow[P_2]{D_2} b[Z] \wedge ... \wedge a \xrightarrow[P_n]{D_n} b[Z] \Rightarrow a \xrightarrow[P']{D'} c[Z] \wedge$$

$$\left(P'' = \left\{P_m \mid |P_m| = \min\left\{|P_i| (i = 1...n)\right\}\right\}\right)$$

$$\wedge \left(D'' = D_1 \oplus D_2 \oplus ... \oplus D_n\right)$$

The shortest recommending path is the easiest path to verify that indirect information, even if the value of the trust degree vector is not as high as others. So we keep that path as the recommending path for the comprehensive recommendation in case we need to follow that path to verify that recommendation. But more likely we will only use the unified trust degree vector alone after the composition.

5.2 Decision-making process

With the help of the fuzzy operations and rules defined above, we can form a formal decision-making process to handle uncertainty in federated trust management. The diagram of the process is illustrated in figure 3. Users need to define the categorization of uncertainties in their mind first. Then the decision-making process uses fuzzy operations to combine uncertain information from different sources. After defuzzification of the trustworthiness degrees, users or system administrators need to judge whether the final degree is consistent with the users' intentions or comply with the practical application environment. If not, the parameters of the fuzzy operations need to be adjusted.

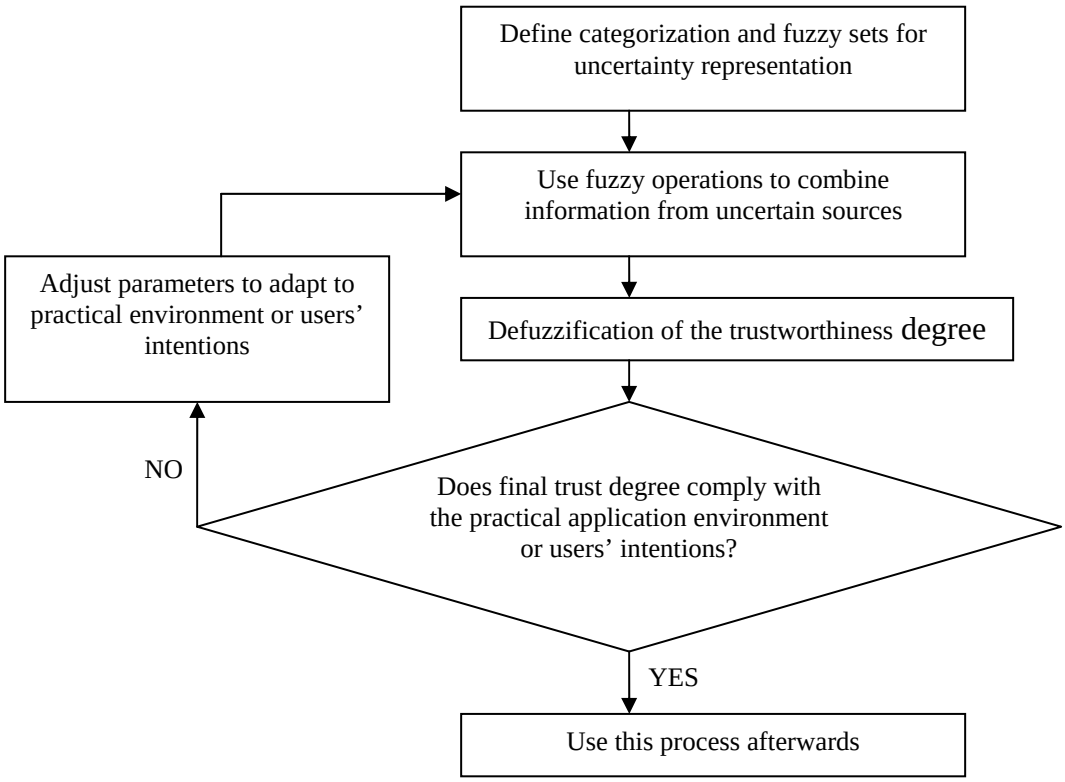

Figure 3. The Diagram of Decision-making Process for Uncertainty in Federated Trust Management

This decision-making process can solve the issues with uncertain information or judgment in federated trust management. For example, the diagnosis process for detecting uterine fibroids in a telemedicine system needs to coordinate display of digital images from another healthcare provider, response from doctor's analysis, and other activities across trust domains. If the network transmission is trustworthy, reading the digital image from another healthcare provider still raises concern about the technician's skill and the technical resolution of the remote sensor. Because this is indirect information, uncertainty from these factors is unavoidable. Meanwhile the doctor's response always includes some level of fuzziness, because the doctor's diagnosis is based on subjective judgment from medical knowledge, previous experience, and possibly further test results. The categorization of uncertainty in section 4 is the practical response used in medical diagnoses [16]. Using that

categorization can provide a practical way to deal with uncertainty in medical practices. Thus following the proposed decision-making process, a telemedicine system can allow fuzzy input of indirect information and subjective judgment from different trusted entities with appropriate categorization, combine fuzzy input, and reach a final decision for a fast consult to the doctors or other users. So we not only provide a model for uncertainty in federated trust management but also propose a practical way to handle uncertainty. Further this decision-making process can incorporate users' fuzzy definitions of policies or rules into federated trust management systems to provide more expressiveness and flexibility.

6. An application in healthcare environment

Following the example described above, we illustrate the practical fuzzy policies, the user interface to input fuzzy policies, and the enforcement mechanism to enforce these policies for a healthcare environment. Since the diagnosis of uterine fibroids involves both indirect information and subjective judgment, we have two sets of fuzzy policies to describe corresponding fuzzy rules. We also have a regular policy set without fuzziness. The fuzzy policy for indirect information is illustrated below.

- The confidence level of the technician's skill at the remote sensor is high/medium/low.

High, medium and low are membership functions to describe the level of uncertainty. The fuzzy policies for subjective judgment are illustrated below.

- The existence of focal fibroid tumors is definitely uncertain/probably uncertain/ equivocal/probably certain/definitely certain.

- If focal fibroid tumors are (probably/definitely) certain, more than one tumor is definitely uncertain/probably uncertain/equivocal/probably certain/definitely certain.

- If focal fibroid tumors are (probably/definitely) certain, their locations being within or bordering the endometrial canal is definitely uncertain/probably uncertain/equivocal/probably certain/definitely certain.

Definitely uncertain, probably uncertain, equivocal, probably certain and definitely certain are membership functions too. And the overall policy is described below.

- If more than one focal fibroid tumor is within or bordering the endometrial canal, the patient needs a hysteroscopy (treatment).

After we have defined the policies, we design and implement a user interface to assist doctors to input these policies. As illustrated in figure 4, we allow doctors to change the flexible parts in fuzzy policies like 'more than one focal fibroid tumor', 'within or bordering' and 'certain' according to their diagnositic needs.

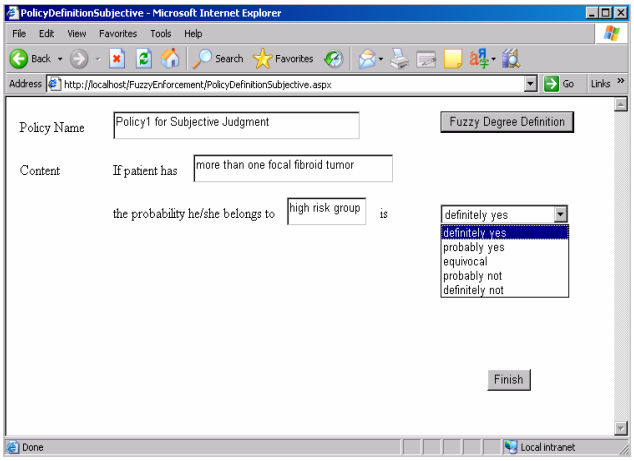

Figure 4. Fuzzy Policy Input

Also we allow doctors to define all the membership functions. As illustrated in figure 5, we provide a set of default sampling point values for every membership function. If doctors do not satisfy with those definitions, we allow them to input new sampling point values. And after doctors enter the new values, we will visualize the curve of a new membership function in a new window as the feedback mechanism to refine the membership function step by step.

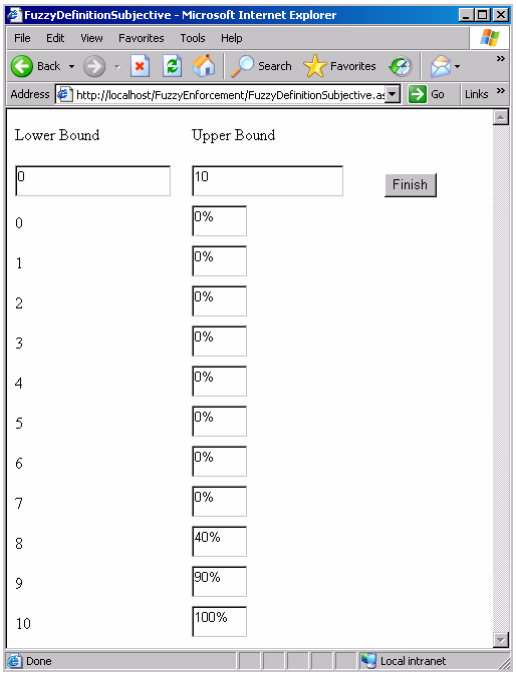

Figure 5. Membership Function Input

Once the definitions of fuzzy polices are finally determined, we use a policy generator to translate the fuzzy policies into XACML [17] format, and store them in a policy database. Then once the digital images from a remote sensor are present and the doctor's judgment has been input, the system can tell the patient what treatment they will need from a web service interface. And the patient can use that service anytime, anywhere. Figure 6 illustrates the system architecture. The enforcement engine is triggered when a request from a patient is received, and it will use a fuzzy policy filter to go through the decision-making process to reach the final decision for the patient.

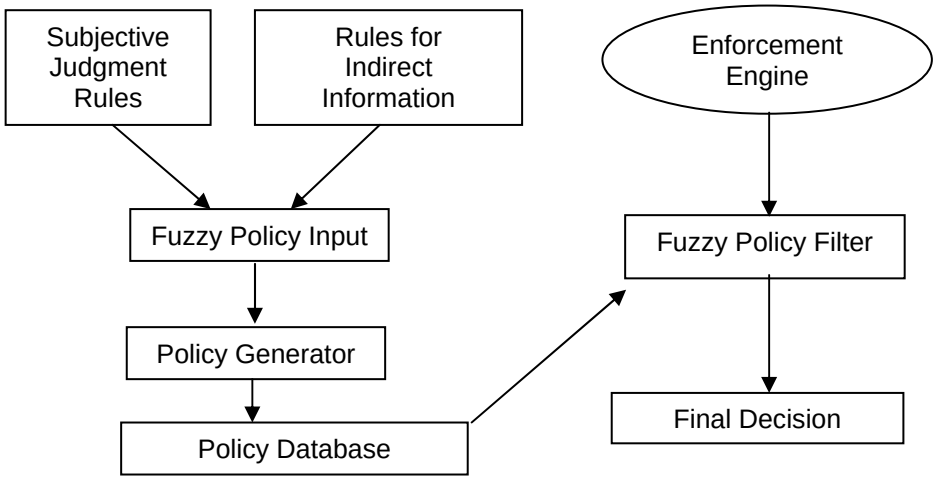

Figure 6. System Architecture

7. Conclusion

This paper proposed a model of uncertainty based on fuzzy logic to handle uncertainty and fuzziness in trust management. Compared with the trust management model proposed by Josang [13], this paper identifies different sources of uncertainty and fuzziness in trust management, and finds that this uncertainty cannot be treated as a probability and thus cannot be described by a probability model. This paper introduces the membership function from fuzzy logic to describe uncertainty and fuzziness in trust, and defines a trust degree vector to evaluate level of trustworthiness. This paper also introduces a general categorization to describe various types of trust in daily life and practical application environments. In addition, the derivation rules proposed in this paper incorporate a parameter to allow users to adjust the membership function through a feedback mechanism in order to make the system adapt better to actual application environments, which solves the inadequacies in the model proposed by Josang [13] and the model proposed by Beth et al. [11]. The model proposed in this paper can be used in evaluation, analysis and derivation of policies in trust management directly. As illustrated in section 6, application of this model in a healthcare environment can help doctors provide diagnosis online with pervasive computing devices operating across trust domains. This provides effective support for policy and decision making in trust management for federated and pervasive networks.

8. References

[1] M. Blaze, J. Feigenbaum, and Jack Lacy, "Decentralized Trust Management", *Proc. of the 1996 IEEE Symposium on Security and Privacy*, 1996, pp. 164 -173.

[2] T. Grandison, "Trust Management for Internet Applications", PhD thesis, Imperial College London, 2003.

[3] N. Li and J. C. Mitchell, "Datalog with Constraints: A Foundation for Trust-management Languages", *Proc. of the Fifth International Symposium on Practical Aspects of Declarative Languages*, 2003, pp. 58 - 73.

[4] G. Caronni, "Walking the web of trust", *Proc. of 9th IEEE International Workshops on Enabling Technologies*, 2000, pp. 153 - 158.

[5] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "Eigenrep: Reputation management in p2p networks", *Proc. of 12th International WWW Conference*, 2003, pp. 640 - 651.

[6] C. Duma, N. Shahmehri, and G. Caronni, "Dynamic trust metrics for peer-to-peer systems", *Proc. of 2nd IEEE Workshop on P2P Data Management, Security and Trust*, August 2005.

[7] S. J. Dwyer III, A. C. Weaver and K. K. Hughes, "Health Insurance Portability and Accountability Act," *Security Issues in the Digital Medical Enterprise*, Society for Computer Applications in Radiology, second edition, April 2004.

[8] D. Harrison McKnight and Norman L. Chervany, "The Meanings of Trust," *MISRC Working Papers Series*, 2000.

[9] D. W. Manchala, "Trust Metrics, Models and Protocols for Electronic Commerce Transactions," *18th International Conference on Distributed Computing Systems*, 1998.

[10] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis, "KeyNote: Trust Management for Public-Key Infrastructures," *Proc. 1998 Security Protocols International Workshop*, Springer LNCS vol. 1550, April 1998, pp. 59 - 63.

[11] T. Beth, M. Borcherding, and B. Klein, "Valuation of Trust in Open Networks," *Proc. 1994 the European Symposium on Research in Security*, Springer-Verlag, 1994, pp. 3 – 18.

[12] A. Josang, "Trust-based Decision Making for Electronic Transactions," *Proc. of 4th Nordic Workshop on Secure Computer Systems*, 1999.

[13] A. Josang, "A logic for uncertain probabilities," *International Journal of Uncertainty, Fuzziness and Knowledge-based Systems*, 2001, 9 (3), pp. 279 – 311.

[14] S. Weber, "A general concept of fuzzy connectives, negations, and implications based on t-norms," *Fuzzy Sets System*, 1983(11), pp. 115-134.

[15] D. Dubois and H. Prade, "New results about properties ans semantics of fuzzy set theoretic operators," *Fuzzy sets: theory and applications to policy analysis and information systems*, Plenum, New York, 1980, pp. 59-75.

[16] G.. A. DeAngelis, et al., "Diagnosis efficacy of compressed digitized real-time sonography of uterine fibroids," *Academic Radiology*, Volume 4, Number 2, February 1997, pp. 83-90.

[17] OASIS standard, "eXtensible Access Control Markup Language (XACML) Version 2.0", http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf.