A Distributed Event-triggered Knowledge Sharing System

Seema Degwekar[1], Jeff DePree[1], Stanley. Y. W. Su[1], Howard Beck[2]
[1]Database Systems R&D Center, Computer and Information Science and Engineering, University of Florida
[2]Agricultural and Biological Engineering Department, Institute of Food and Agricultural Sciences, University of Florida
Gainesville, Florida 32611
1-352-392-2693

spd@cise.ufl.edu, jdepree@cise.ufl.edu, su@cise.ufl.edu, hwb@ufl.edu

ABSTRACT
Government agencies globally are facing problems such as illegal immigration, terrorism, and disease diagnosis and control. Solutions to these problems rely heavily on collaborating organizations' ability to effectively and efficiently share not only ***data*** but also ***knowledge*** embedded in organizational and inter-organizational policies, regulations and constraints. Responding to an emergency often requires organizational and inter-organizational processes and complex operating procedures to be followed. We focus on the sharing of data associated with events of interest to collaborating organizations. Condition-action-alternative-action rules, logic/derivation rules, and constraint rules are used to define organizational and inter-organizational policies, regulations, and data and security constraints. Structures of these heterogeneous rules are used to capture organizational and inter-organizational processes and operating procedures. In this demonstration, operational procedures developed by collaborating organizations in USDA's National Plant Diagnostics Network (NPDN) will be used to show the knowledge definition facilities and the distributed event-triggered knowledge sharing strategy.

Categories and Subject Descriptors
H.3.4 [**Information Storage and Retrieval**]: Systems and Software – d*istributed systems, information networks.*

H.4.2 [**Information Systems Applications**]: Types of Systems – *decision support.*

General Terms
Management, Documentation, Performance, Design, Reliability, Experimentation, Security, Human Factors, Standardization.

Keywords
Knowledge representation, knowledge sharing, web services, decision support, collaborative federation, event-and-rule based systems.

1. INTRODUCTION
Establishing effective collaboration among government agencies holds the key for solving complex problems such as illegal immigration, terrorism, and disease diagnosis and control. One important form of collaboration is for these organizations to share data as well human and organizational **knowledge.** As of yet, an effective system for this form of collaboration is lacking. In this work, we focus on capturing the knowledge embedded in organizational and inter-organizational policies, regulations, and constraints by using three popular types of *knowledge rules* [1, 2, 4]: integrity constraints [5], logic-based derivation rules [6], and action-oriented rules [7]. Organizational and inter-organizational processes and operating procedures are specified by *rule structures*. A rule structure is a directed graph with different types of rules as nodes and edges that represent the relationships between these nodes. Using these rules and rule structures, we can effectively capture the **multifaceted knowledge** of collaborating organizations.

An organization typically does not open up its database to others, nor is it interested in processing all of the data residing in collaborating organizations' databases. Rather, collaborating organizations are interested in obtaining only those data that are pertinent to the occurrence of an event of common interest (i.e., **event data**) and in processing only those knowledge rules that are applicable to the event data. An event is anything of significance to collaborating organizations (e.g., an arrest, a terrorist incident, the detection of a disease, a special state of a database, a signal from a sensor, etc.) that occurs at a particular point in time. An event-triggered knowledge sharing system that facilitates event subscription, event notification, delivery of event data, and processing and interoperation of applicable knowledge rules and rule structures would be ideal for achieving data and knowledge sharing.

Since we allow knowledge to be specified in different types of rules, we need to have an effective approach to process them and to achieve their interoperability. Using multiple rule engines with wrappers to process heterogeneous rules would lead to a highly complex and unwieldy system. Instead, we provide a user-interface tool for organizations to define their rules and rule structures, translate these rules and rule structures into program code, wrap them as web services at definition time, and deploy the code at their own sites for processing in a web service infrastructure at runtime. By doing so, rules and rule structures can now interoperate programmatically as web services without using different types of rule engines.

2. SYSTEM DEMONSTRATION
This demonstration accompanies a research paper to be presented at the same conference. It covers two major components (a user interface tool and an event and rule processing subsystem) of a distributed, event-triggered knowledge sharing system being developed for deployment in the National Plant Diagnostic Network (NPDN) environment.

NPDN is a multi-year national project launched by the United States Department of Agriculture (USDA) for strengthening the homeland security protection of food and agriculture by connecting five regional plant diagnostic centers with a national

center to facilitate quick and accurate detection of and response to economically threatening diseases and pest outbreaks in crops. Such outbreaks can occur due to intentional acts of bioterrorism. Complementing USDA's effort, our research team funded by NSF has been developing a web-based, distributed system for event-triggered knowledge sharing among NPDN organizations.

The implemented user interface provides the facilities for each collaborating organization to define events of interest and publish them in a global registry at a host site for browsing, querying and event subscription and notification. It is also used to define an organization's action-oriented rules, derivation rules and constraint rules, as well as structures of these rules. Additionally, triggers that link distributed events to distributed rules and rule structures can be specified by each organization. The main features of the user interface are: 1) it provides a unified tool for each collaborating organization to define and publish its events, rules, rule structures, and triggers, 2) it provides usage instructions by means of help buttons to guide the user in performing different event, rule and trigger definition tasks, 3) it provides data-type and error checking facilities to prevent the user from making mistakes in event, rule and trigger specifications, 4) it provides facilities to hide/view complex constructs of rules so that complex rules can be defined/viewed more easily, and 5) it allows rule structures to be defined and viewed graphically.

The defined heterogeneous rules and rule structures are automatically converted to web services at definition time for their uniform and efficient processing in a web service infrastructure at runtime. To enable distributed event, trigger and rule processing, the implemented event and rule processing subsystem, which contains an event server and a rule server, is replicated and installed at all collaborating sites to form a peer-to-peer server architecture. The occurrence of an event at one site will cause the data associated with the event occurrence (i.e., ***event data***) to be sent to the replicas of the event server of those sites that contain applicable rules and rule structures. These replicas will activate their corresponding replicas of the rule server to process the applicable rules and rule structures. Each site may add new data to or alter the content of the initial event data. The aggregation of data returned from applicable sites with the original event data may make some other distributed rules and rule structures become applicable. Multiple rounds of event data transmission and aggregation, and rule and rule structure processing may take place until all the data that are relevant to the event occurrence have been generated and received by all relevant organizations for their decision-making, problem solving and activity coordination. Event data are ***dots*** that are ***connected across organizational boundaries*** through the interoperation of knowledge rules and rule structures. The main features of the event and rule processing subsystem are: 1) it provides facilities to translate rules and rule structures into web services and to register them with the web service registry, 2) it enables the processing of heterogeneous rules and rule structures uniformly as web services to achieve their interoperability without having to use heterogeneous rule processing engines, 3) it provides algorithms for aggregating event data returned from sites that contain applicable rules, and for handling inconsistencies and cyclic conditions, 4) it provides XML-based specification languages for formal representations of events, rules, rule structures and triggers defined through the user interface, and 5) it provides facilities for event management, subscription, notification, and event data delivery and aggregation.

Some of the standard operational procedures (SOPs) developed by collaborating organizations in USDA's National Plant Diagnostics Network [3] are used in this demonstration to explain and illustrate the knowledge definition facilities and the distributed event-triggered knowledge sharing strategy. These SOPs outline the tasks to be carried out and the lines of communication to be followed when an outbreak of a "pest of concern" is detected. The demo system will run on three computers over the Internet: one serving as the host site and two as the collaborating sites. The demonstration will be of interest to all conference attendees as the developed distributed, event-triggered, knowledge sharing technology is a general technology that can be applied not only in the NPDN environment, but also in other collaborative environments such as e-business, e-learning, and e-government.

3. ACKNOWLEDGMENTS

This project is supported by the National Science Foundation (Grant No. IIS-0534065)

4. REFERENCES

[1] Business Rules Group. *Defining Business Rules – What Are They Really?* Final report, http://www.businessrulesgroup .org/first_paper/BRG-whatisBR_3ed.pdf, 2000.

[2] Rouvellou, I., et al. Combining Different Business Rules Technologies: A Rationalization. In *Proceedings of the OOPSLA 2000 Workshop on Best-practices in Business Rule Design and Implementation*, Minnesota, USA, Oct. 15, 2000.

[3] The National Plant Diagnostic Network Standard Operating Procedure for APHIS-PPQ Pest of Concern Scenario – General SOP, June 17, 2005.

[4] Sowa, J. *Knowledge Representation: Logical, Philosophical and Computational Foundations*. Brooks Cole, Pacific Grove, CA, 2000.

[5] Ullman, J. *Principles of Database Systems, 2nd Ed,* Computer Science Press, Rockville, MD, 1982.

[6] Ullman, J. *Principles of Database and Knowledge-Base Systems*, Computer Science Press, Rockville, MD, 1988.

[7] Widom, J., and Ceri, S. *Active Database Systems, Triggers and Rules for Advanced Database Processing*. Morgan Kaufmann, San Mateo, CA, 1996.