

# AIMS: Agent-based Information Management System

**Srikanth Kallurkar, Donald Steiner and,**

**Robert Flo**

**Apperson Johnson**

Quantum Leap Innovations, Inc  
3 Innovation Way Suite 100 Newark DE 19711 USA  
{svk, dds, ahj}@quantumleap.us

Air Force Research Lab/IFSB  
525 Brooks Road Rome, NY 13441-4505  
Robert.Flo@rl.af.mil

## Abstract

Information Management is an important component in the development of viable information sharing environments and systems. Information published in shared environments may be manipulated to generate or derive subsequent information. The qualities of such derived information depend on qualities of the contributing units of information and of the entities that published them. The ability to capture and propagate data about such manipulations through successive generations of information is essential to the calculation of the qualities of information. Such qualitative and pedigree data about the information can provide a basis for making application specific decisions and attributing confidence to those decisions. The cost of propagating pedigree data increases as the volume of information increases through generations. Without loss of generality, most users of an information sharing environment could be satisfied with approximate and quick assessments of qualities and an abbreviated history of related information. However, not all assessments need to be performed in a time critical environment and not all approximate and quick information assessments can satisfy every quality need. In this paper, we present the architecture for an Agent-based Information Management System (AIMS) that provides for mechanisms for establishing quick and approximate qualities and related provenance (to accommodate time-critical decision making, and limited data storage), complemented by a mechanism that supports exhaustive treatment of data provenance, and parallel calculation of precise qualities. We present the novel concept of a Variable Information Chromosome per unit of information that propagates only the most representative provenance across generations. The multi-agent framework provides extensive support for provenance related services.

## Introduction

Information Management is an important component in the development of viable information sharing environments and systems because it provides for determining and propagating qualities of information. Information published in shared environments maybe manipulated to

generate/derive new information whose qualities depend on the qualities of the contributing information and of the sources that published them. The ability to capture and propagate data about such manipulations through and across successive information generations is crucial to the calculation of derived qualities.. A user of such a managed environment should be able to quickly obtain qualities of information and relevant history about its derivation. Such qualitative and pedigree data about the information can provide a basis for making application dependent decisions and attributing confidence to those decisions. In high volume, high throughput and highly dynamic information sharing environments, users can publish and obtain information from myriad mechanisms that create, publish, subscribe and query information. The information obtained to satisfy an information need can be from multiple sources and be complementary, supplementary, divisive or contradictory. Various information analysis and generation tools and mechanisms can be used to derive new, or to update existing information. Such information is made available through the environment to be accessible and (possibly) useful to other related users. The published items may be highly informative but is devoid of qualitative attributes that can provide an arbitrary user with confidence, or any other mechanism of inferring the information's usefulness. Examples of qualitative attributes are trustworthiness, reliability, currency etc. Additionally, pedigree of information can provide valuable insights about trail of its derivation. Such pedigree data may include contributing information items, their qualities and also the qualities of the sources that published or produced the items.

## Motivation

As the volume of information increases through generations of derivation, the volume of accompanying pedigree data increases proportionally. Subsequently, storage and computational costs of calculating qualities increase in proportion to the accuracy required and the size of the data over which they are calculated. A brute force method that calculates qualities over all available pedigree is therefore very costly and can require an unwarranted degree of storage overhead. In addition, quality and

pedigree about information many generations old is not always required in quality calculations for new information. Such a requirement of quality attributes about aged information should commit resources proportional to the precision desired. In addition, as new information is added into the system and old information is updated, the additions and changes may have to be propagated to preceding and succeeding generations, depending on the generations are linked to the changed/new information.

We posit that there is a need for an Information Management solution that provides an effective and efficient determination and propagation of qualities of information. We address the problem for Information Management within an information sharing environment by providing such an effective and efficient mechanism in order to offer a basis of information assurance and assistance in informed decision making.

Without loss of generality, most users of an information sharing environment could be satisfied with approximate and quick assessments of qualities and a brief history about the information derivation. Such generalization is supported by the fact that in real time deployment of information sharing environment, such as for military purposes and real time tracking of financial markets, the users may not have enough time to navigate or assess detailed descriptions of the derivation history or to wait for the precise quality computations to be completed over exhaustive provenance records. Users, in other domains, such as heavy equipment operators, plant managers and operations managers, require quick quality assessments to make judgments about resource allocation, failure identification and reconfiguration of plans. However, not all assessments need to be performed in a time critical environment and not all approximate and quick information assessments can satisfy every quality need. Thus, an Information Management system that provides for quick and approximate qualities to save on time and storage costs should also provide for a mechanism that can support complete, different and more accurate qualities of information. The problem of determining the provenance of information has been studied by the research community with the focus ranging from metrics of trust to trust management [Buneman, Golbeck, Guha]. A core problem that has not been addressed concretely in the research so far is that of solutions that scale reasonably in terms of the storage and computational costs of maintaining and exploiting provenance. The issue is directly related to a suitable representation of pedigree data that can be configured to provide cost-apportioned provenance information. Specifically, the users of the system that require the provenance of information must be able to specify the scope of the provenance based on the costs associated with both procuring and subsequent reasoning over the data. Based on these requirements, we present a generalized architecture of an Agent-based Information Management System (AIMS). We introduce a novel concept of Variable Information Chromosomes for the propagation of the only most relevant provenance across

generations. The Related Works Section surveys the literature, the AIMS overview Section presents an overview of the system along with discussion on Pedigree, Provenance and Quality of information and how they are related. The Architecture Section describes the core components of AIMS Architecture. We conclude in the Integrated View Section with AIMS summary and future direction.

## Related Work

Buneman et al [Buneman] characterize provenance as the reason (why-provenance) a particular value is present in database query and the source that supplied it (where-provenance). Such a characterization is useful in multiple and heterogeneous databases where the views may be constructed from deep rooted sources. The provenance establishes the lineage, accuracy and currency of the data. The work develops a syntactic provenance approach for semi-structured data. This work provides a clear distinction between the why-provenance and the where-provenance which we draw upon to characterize information provenance at a higher level of information abstractness, not necessarily grounded to structured or semi-structured data models. Determination of data provenance, per this research is deterministic, exhaustive, and potentially very expensive in terms of computation and storage. At the same time, availability of exhaustive data provenance does not, by itself, answer the primary concern of users – “is this data appropriate for a particular use?”

In contrast to the data provenance approach, the semantic web community has focused on the issues of establishing and maintaining trust, for example about web services or simply information published on websites. In particular, Golbeck et al. [Golbeck] characterize trust networks on the semantic web by drawing from social networks as a loose formulation of collaborating evidence. The FOAF project is an example of such a network where the trustworthiness of a node is a function of the trustworthiness of the nodes connected to it. The premise of trust networks is propagation of trust about the sources and the claims made by them. Trust ontologies assist in annotating the nodes with measures of trustworthiness. The key aspect of this approach is that different users can have different views about the levels of trustworthiness of a node and that these can be resolved by using simple trust calculi such as least trustworthiness or average trustworthiness.

Guha et al [Guha] suggest a computational model for propagating trust and distrust in an online community of users. In this model, a matrix of trust and distrust between users with respect to certain information is computed. The trust and distrust values are the propagated under a transitivity paradigm,  $i$  trusts  $j$ ,  $j$  trusts  $k$ , therefore,  $i$  trusts  $k$  and, co-citation paradigm,  $i$  trusts  $j$ ,  $k$  trusts  $l$  and  $m$ , therefore,  $i$  trusts  $l$  (where  $i$ ,  $j$ ,  $k$ ,  $l$  and  $m$  represent users). The main feature of this work is that of inclusion of distrust between users. The model is relevant for

determining trust for ecommerce applications, where the reviews about objects as well as reviews about the reviews can be expressed as a singleton numeric value. The model also assumes taxonomy of topics describing the objects to allow for topic related trust values and maintaining the size of the matrix.

In addition to the models and mechanisms of trust, the representation of the provenance that is amenable to distribution and complex reasoning is important. The International Document Object Identifier (DOI) Foundation developed a DOI identification system [Wiley]. The system consists of a framework managing intellectual property in the digital environment. It works with a centralized database that allows the tracking of the identifiers as the documents are exchanged online. We draw upon the DOI representation scheme as it is similar in construction and association with the digital information as the information chromosomes.

### AIMS Overview

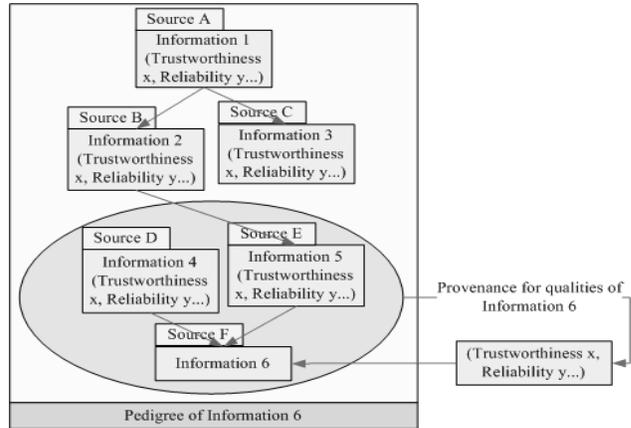
AIMS provides for qualities of information and related provenance through a provenance storage structure called the Variable Information Chromosome (VChrome). The concept of VChrome is motivated by biological chromosomes [Lowrance J, Matching DNA, Molecular Genetics, Petsko G] that carry only a portion of the original information in their constituent genes through successive generations. Generalizing this ability to the information provenance domain, a VChrome is a lossy storage structure that carries the most representative provenance about the most common or salient qualities of information. In addition, AIMS provides provenance streams to establish complete provenance for more accurate quality calculations. The two provenance establishment mechanisms are supported by a multi-agent framework that provides flexibility in adapting to changing provenance needs and is scalable to growth of information volume and the number of users of the system.

The AIMS architecture is developed for implementation in publish-subscribe information sharing environments. Clients of such an environment have three main operations – publish, subscribe, and query. The AIMS architecture provides for Information Management by providing requisite interaction mechanisms between the various elements of the environment and multi-agent framework. The objectives of the Architecture for Agent-based Information Management System are as follows:

- A **provenance representation structure** that provides fast access to provenance information, and that satisfies the most common situations of provenance requests.
- A **provenance management system architecture** that provides a robust framework for enabling configurable provenance requests

### Information Pedigree, Provenance and Qualities

Provenance is the data required to compute a quality of Information. It can be obtained from the pedigree of information. A quality of information is a specific assessment of information (Figure 1). Examples are its reliability and currency.



**Figure 1: Information Pedigree in relation to Provenance and Qualities. Figure shows Pedigree of Information 6 with its entire derivation history, entities that produced contributing information, and qualities (of information and entity).**

We classify information pedigree as all the external characteristics of information such as the entities (sources, authors, sensors, etc) that produced information, their qualities, their location, time of production and so on. Such a classification allows quantifying, to the best extent possible, the characteristics such that quantitative measurement of qualities can be made. We note that pedigree refers to only the external characteristics of information and does not describe internal characteristics such as meaning of the information content or any other inferred knowledge based on the content. While both internal and external characteristics are required to get a full assessment of a quality, we define the scope of the Information Management problem as to manage the external characteristics (It is an obvious and useful extension to combine the extrinsic and intrinsic characteristics). Table 1 illustrates classification of qualities. For example, capability of an entity relates to the entity’s (such as a source) ability to produce information about a topic. This is extremely relevant to Intelligence Community.

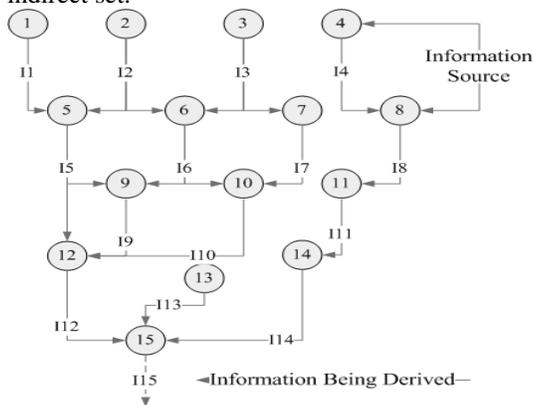
**Table 1: Categories of Qualities of Information**

Class	Name
Quality about an entity	Trust
	Capability
Quality about information	Consistency
	Currency
	Reliability

Approximate provenance is a subset of the complete provenance required to calculate a quality. The basis for

creating the subset is that approximate provenance should be able to propagate enough representative data about the derivation of information to future generations. Consequently, a quality that is computed from approximate provenance is an approximate quality.

Figure 2 depicts a subgraph of information derivation from multiple entities. Such a subgraph can be a Directed Acyclic Graph (DAG) on the assumption that information is immutable and changes to a unit of information results in a new unit of information. In all other cases, the subgraph can have cycles that will need to be resolved. The nodes of the subgraph are entities (or sources) that produced a unit of information, represented by the edges. An edge between a pair of nodes represents flow of information in the direction of the pointed edge. The figure depicts the subgraph with respect to the edge labeled **I15** indicating that I12, I13, and I14 are the units of information that directly led to the derivation of I15 and the corresponding entities are its direct contributing entities (or contributors). The remainder of the edges represents indirect information affecting I15 and their entities, the indirect contributors. I1, I2, I3 and I4 are the original units of information in its pedigree and the corresponding entities are its root entities. The figure does not depict other edges from this set of nodes because those edges (i.e. the units of information) do not impact the derivation of I15. The figure however, does consist of the pedigree for every edge in the subgraph because an edge pointing to a node indicates the information that contributed to the derivation of the unit of information. Furthermore, the distinction between direct and indirect contributors is made explicit to facilitate establishment of provenance establishment. For example, the set of direct contributors for I15 would include (12, 13 and 14), while the set of indirect contributors would include (1, 2 ... 11). The sets are non-intersecting and provide for establishing approximate provenance. For example, approximate provenance of I15 can include entities from direct-contributors set and some from the indirect set.



**Figure 2: Information Pedigree Subgraph depicting the flow of Information between entities (sources).**

AIMS, through its VChromes and IPS services provides quality and related provenance by managing pedigree graphs for units of information that are published and accessible in the information sharing environment. It

allows the application to configure defaults for establishing approximate provenance and the types of qualities most desired. It also allows for changes in the type and scope of qualities through the extensive provenance streaming capabilities. The multi-agent framework offers a service oriented framework for all provenance related tasks. In addition to the provenance capabilities, the framework also provides for the clients to interact with the system to specify the types of qualities and the type and scope of provenance they desire. AIMS architecture was developed under these assumptions:

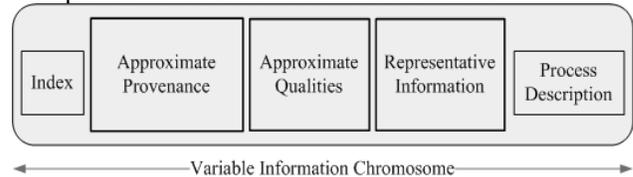
- The architecture does not provide quality calculation methods/functions.
- Given a quality computation function, AIMS establishes provenance for the quality.

## Architecture

### Variable Information Chromosome

A Variable Information Chromosome (VChrom) is a self-contained structure (Figure 3) associated with an information object with storage for:

- Approximate and representative provenance about information that contributed to the derivation of information.
- Approximate and representative Qualities of information that contributed to the derivation of Information.
- Representative Metadata about information.

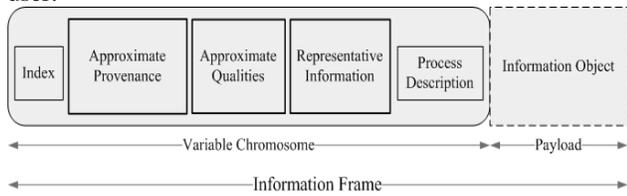


**Figure 3: Structure of a Variable Information Chromosome**

VChrom is self contained in that a quick, immediate and approximate sense of provenance and qualities can be achieved by observing the structure without any other supporting data or data generation tools. Such a structure is referred to as a *chromosome* as it associates with the information, the most critical and representative provenance while losing unnecessary and redundant provenance. As in biological chromosomes, the loss of provenance through successive generation provides for compact storage while retaining only the most impact-full and important provenance. Unlike biological chromosomes, VChromes can represent lineage trees with any number of intermediate parents. A VChrom is also characterized by its variable length adjustable to the variations in the type and volume of provenance. In addition, the VChrom structure provides storage for representative qualities.

AIMS defines Information Frame (Figure 4) as a structure comprising the VChrom and the payload. The quality and related provenance about the information payload is carried in a VChrom. The information frame

provides for the ability for the VChrome associated with the particular information payload to be delivered to the user.



**Figure 4: Structure of an Information Frame.**

A Variable Chromosome Architecture (VChrome) comprises the following components:

**Index:** A VChrome provides an index for easy access to its content.

**Approximate Provenance:** A VChrome stores approximate provenance about Information in the payload computed as a representative subset of pedigree of the most prominent contributors to the derived information. Approximate provenance in a VChrome comprises the deriving client,  $k$  of  $K$  direct contributors (parent entities),  $m$  of  $M$  pedigree nodes (indirect contributors, i.e. those that contributed to direct entities) and representative or summarization of the Information in the payload. The structure is:

*One deriving client:*

- Identifier: Globally unique identifier of the entity.
- Quality (Trust, Capability): Obtained from a look up table.
- Time of publication.

*k immediate entities* (user selected value,  $k \ll K$ ): Each of the  $k$  entities is represented by

- Identifier: Globally unique identifier of the entity.
- Quality (Trust, Reliability, Capability): Obtained from a look up table.
- Time of publication of Information.

*m pedigree entities* (user selected value,  $m \ll M$ ): Each of the  $m$  entities is represented by

- Identifier: Globally unique identifier of the entity.
- Quality (Trust, Reliability, Capability): Obtained from a look up table.
- Time of publication of Information.
- Number of Children: Contributors to current derivation that were derived from this entity.
- Representative Information: Information dependent and may include suitable metadata, data values (integer, double, strings) and/or summarized text.

**Approximate Qualities:** Trust, Reliability, Consistency, Currency

**Process Descriptions:** A VChrome stores descriptions about the quality computing processes and related provenance establishment methods. The processes are available through AIMS libraries

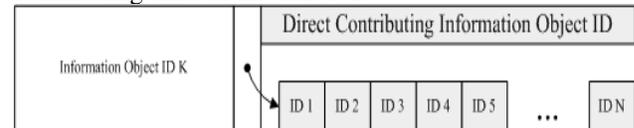
### Information Provenance Stream

An Information Provenance Stream (IPS) is a set of pedigree nodes that satisfy a complete-provenance determination request. The set consists of those pedigree

nodes (direct and indirect information contributors) whose published information led to the derivation of new information. Establishing such a set provides the provenance required for computing accurate and complete qualities of Information.

IPS is established by an agent on request for provenance by another entity (client or agent) and is based upon the pedigree subgraph of a unit of information. The stream refers to obtaining provenance from provenance storage. The architecture consists of IPS services as part of the Provenance Management System (PMS) subsystem.

The IPS is constructed from the set of pedigree nodes that contributed to the derivation of the particular unit of information. The set of nodes form a subgraph with a single sink at the particular derived unit of information. The graph-sources are the nodes that contributed the original information. A derivation of new information triggers storage of a new pointer object (conceptual) in a Pedigree-base. The pointer object lists the direct contributing units of information.



**Figure 5: Pointer Object about an Information Object depicting linkage to the directly contributing information objects.**

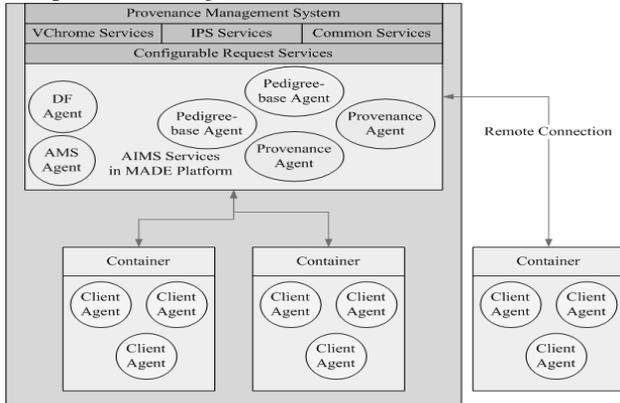
The construction of a pedigree subgraph for a particular unit of information is done by traversing the pointers and setting the contributing information as the graph edges and the contributing entities as the graph nodes. Figure 5 depicts a traversal through the subgraph nodes during construction of pedigree subgraph. The directed edges are in the direction of traversal. The flow of information however, is in the reverse direction.

### Multi-Agent Framework

The Agent framework provides communication and coordination for establishing provenance. The framework consists of special provenance agents and infrastructure agents. The provenance agents establish provenance through various services that are made available through a service stack in the agent framework. All aspects of VChrome(s) and IPS are managed by provenance agents through these services.

The AIMS multi-agent framework comprises set of support or infrastructure agents and three basic provenance related agents. The multi-agent framework is implemented using our Multi-Agent Development Environment (MADE). MADE is built on open source agent communication framework Jade [1]. MADE agents are easy to construct, requiring minimal programming language use. MADE incorporates pre-defined collaboration sequences based on common workflow and business process management standards, decentralized policy management, modular addition and deletion of behaviors, advanced subscription mechanisms, and rich service descriptions. MADE

currently extends the widely used FIPA standards-compliant JADE agent framework.



**Figure 6: MADE Platform and Container depicting various AIMS agents and services.**

The two basic provenance related agents are Client Agent, Provenance Agent and Provenance-base Agent. Figure 6 describes them in relation to deployment in a MADE platform. The platform hosts the services from the Provenance Management System and the Provenance and the Pedigree-base agents. The client agents can be distributed across various containers depending upon the network or resource-use optimization and utilization constraints.

### Client Agent

A Client Agent provides the primary interface capability with the clients of the information sharing system. The interface is implemented through either an agent or a protocol or via standard agent-access API. The client agent uses the interface to represent the client in AIMS to request and receive provenance related services.

### Provenance Agent

A Provenance Agent provides access to all the provenance related functionality in AIMS. It is an independent actor in AIMS. Clients do not extend control over it for a non-intrusive execution of provenance-related tasks. These tasks are carried out and coordinated by the agent using the services offered by the provenance management system (PMS). It has two main interactions with the Client Agent:

- Establishing provenance and calculating qualities for a new derived Information Object.
- Request for provenance about a particular quality of an information object.

The Provenance Agent establishes provenance about new Information derived by the client by invoking the various services. For such a task, the agent has access to the following:

- The derived information.
- Derivation process.
- Contributing information.

The agent can establish provenance for a particular quality both from VChromes of the contributing information and an IPS. Independent of the how the provenance has been established, the agent executes the task of creating a VChrome for the new information. This task is carried by using the Create VChrome and Publish/Associate VChrome service. The system defaults also specify the type of qualities and how they should be computed. The agent also uses the system default in selecting the provenance establishment. The application can set the defaults based on its needs. In the absence of an application configured default, the agent uses the VChromes to establish provenance. To establish provenance from an IPS, the agent invokes the establish IPS Service based on the input about the contributing information.

### Pedigree-base Agent

A Pedigree-base Agent provides for storage and retrieval of pedigree data about the information being exchanged in the system. It either interfaces with the existing information repository of the information sharing environment or implements its own information base. In either of the approaches, the storage and retrieval of the pedigree data is abstracted by the Pedigree-base agent. The services in the Provenance Management System interact with the Pedigree-base Agent for

- Querying for pedigree about units of information.
- Storing VChromes per unit of information.
- Querying for VChromes about units of information.

### Provenance Management System

Provenance Management System is an AIMS subsystem that provides a robust framework enabling management of VChromes and Information Provenance Stream. The different services can be provided by different agents or by a single agent depending upon the needs of the application.

### Variable Information Chromosome Management Services

VChrome management services are:

- Create VChrome Service: This service allows an agent to establish provenance required for a quality of Information. The service is invoked by the agents after the client has derived new information and is ready to publish it into the information sharing environment. Once the provenance has been established for the given quality, the service performs
  - Calculating the most representative subset from the established provenance and storing it in the VChrome as approximate provenance. This is done based on the particular method of calculating a quality (obtained from Library of Quality Computing Processes) and using the methods from Library of Statistical Measures and other related mechanisms such those offered by Pedigree data classifiers.

- Calculating the given quality (s) from the established provenance and storing it in the VChromes as Qualities. This is performed by processes available through the Library of Quality Computing Processes.
- Calculating additional metadata and storing it as representative data. This is done by processes available through the Library of Information Processing and Analysis Processes.
- Storing the methods used to calculate a and b.
- Creating an index for the contents of the VChrome.
- **Manage VChrome Service:** This service allows agents to proactively manage VChromes about published Information such as Merge, Reconcile, and Diverge VChromes.
- **Publish/Associate VChrome Service:** This service allows the agent to publish VChrome along with the information payload as an Information Frame into the information sharing environment. The publish service is invoked by the agent after the derivation of information.

### Information Provenance Stream Management Services

The IPS management services are:

- **IPS Establishment Service:** This service establishes an IPS by creating a pedigree subgraph sink-ed at the client that derived information (i.e. information for which the provenance is being sought). The service interacts with the IPS management service by providing pertinent parameters to identify the pedigree nodes. Once the pedigree subgraph is created, the Provenance Agent retrieves pedigree data from the encapsulated stream.
- **IPS Management Service:** This service provides for storage and retrieval of pointers for each derived of information. A stream is constructed by recursively traversing the pointers starting from the sink-ed information. The depth of recursion may be exhaustive to the graph-sources or up-to a horizon-delimited depth. The service provides for automated pointer updates to maintain consistency with changes in the pedigree subgraph. It also provides for removal of pointers from storage for information that has been determined to be discontinued from the shared environment. The dissolution of a pedigree subgraph is also performed under this service.

### System Services

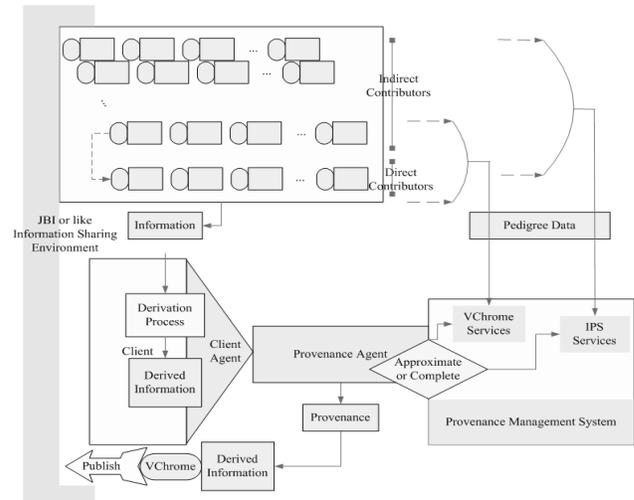
- **Pedigree Data Lookup Service:** This service allows the agents to perform look up operations for data from the following tables/databases about Entity-Trust, Entity-Reliability and Entity-Capability.
- **Library Service for Quality Computing Processes:** The service provides a library of various quality computing processes as part of the service to create VChrome. These refer to externally developed capabilities supplied through AIMS API. Example processes include Corroboration, Dispute, Referral, Propagation of

Minimum, Maximum, Pessimistic, Optimistic value of a quality.

- **Library Service for Statistical methods.**
- **Configurable Request Service:** This service provides for calculation of new qualities not present in the VChrome, change in the horizon over which provenance was established or change in the method of computing an existing quality.

### Integrated View of AIMS Components

The main architectural components and their basic interactions are shown in Figure 12.



**Figure 7:** Integrated View of AIMS Components.

The Client obtains information from the information sharing environment and derives new information. The Client Agent passes the derived information and a list of contributing information to the Provenance agent. The Provenance Agent establishes either complete or approximate provenance to calculate various qualities by invoking various services in the Provenance Management System. The mechanics of establishing provenance are based on the system defaults. The service invocation parameters govern the scope of provenance to be established, i.e. by some subset of direct and indirect contributing information in the pedigree subgraph. An entire set of direct and indirect contributors are is also a subset – referring to establishing complete provenance. New information is published into the system along with a VChrome thus providing for immediate and quick assessments of qualities and related provenance.

### Performance Metrics

We have identified the following performance metrics:

- Effectiveness of approximate V/ complete provenance in assessing a quality of Information.
- Fast approximation V/ Effectiveness of approximation.
- Scalability to growth in volume of information.

## Use Case Scenario

AIMS can be integrated into a news aggregation and derivation system. The system extracts news articles from a number of RSS news feeds specified by the user and enters them into an information repository. The user can view articles (Figure 8) which contains information about the article such as title, author and date, followed by the content, and the article's VChrome. Users derive new articles based on the articles read, which triggers the generation of VChrome by a Provenance Agent. The user can also view the contents of the VChrome for new derived information.

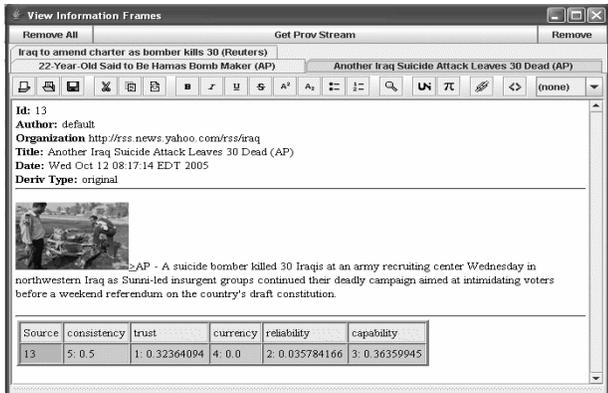


Figure 8: View of results to a query. Each tab is a document from the results shown along with its VChrome.

## Discussion and Future Work

We presented the AIMS architecture and introduced the concepts of VChromes and Provenance Streams. The architecture is developed to be abstract enough to be ported to any application domain and the types of information representation. We also abstracted the actual methodologies of quality computation, provenance establishment and provenance approximation to for the same reason of domain independence. We note that for certain quality calculations that are based linear methods, the qualities may be qualities may be conserved over the tree of calculations, and thus produce exact calculations. For instance, exact, non-weighted minimum or average can easily be conserved through any number of derivations. Some other statistical qualities can be maintained exactly at relatively low cost, if intermediate aggregates are maintained at each node. The AIMS architecture is currently under implementation and testing. We are currently working on implementing a trust computing process based on corroboration and dissention and will evaluate the system performance based on the metrics defined. A key feature that we are currently investigating is the extension of AIMS to learning qualities from data.

## Acknowledgements

The work presented in this paper was funded in part by the Air Force Research Laboratory (AFRL), Rome, NY, under Contract FA8750-05-C-0200. We gratefully acknowledge Robert Flo (AFRL), our colleagues Lewis Fishgold, for providing a first implementation of the described concepts in a Proof of Concept demonstrator and, Dr. Ganesh Vaidyanathan and David Johnson for their extensive contributions and advice during this project.

## References

- Bellifemine F., Poggio A. and Rimassa G. JADE – A FIPA-compliant agent framework. <http://sharon.cselt.it/projects/jade/papers/PAAM.pdf>
- Buneman P., Khanna S. and Tan W. Why and Where: A Characterization of Data Provenance. In *Lecture Notes in Computer Science* (1973) 2001.
- Golbeck J., Parsia B. and Hendler J. Trust Networks on the Semantic Web. In the *Proceedings of Workshop on Cooperative Information Agents* 2003.
- Guha R., Kumar R., Raghavan P. and Tomkins A. Propagation of Trust and Distrust. In the *Proceedings of the Thirteenth International World Wide Web Conference* 2004.
- Lowrance J. Evidential Intervals. <http://www.ai.sri.com/~gister/int-prob.html>
- Matching DNA. <http://science.howstuffworks.com/dna-evidence1.htm> 2004
- Molecular Genetics [http://www.ncbi.nlm.nih.gov/About/primer/genetics\\_molecular.html](http://www.ncbi.nlm.nih.gov/About/primer/genetics_molecular.html) 2004
- Petsko G. Size Doesn't Matter. In *Genome Biology* 2(3) 2001.
- Wiley. Digital Object Identifier. <http://www3.interscience.wiley.com/doiinfo.html>