INFORMATION INVASION IN ENTERPRISE SYSTEMS

Modelling, simulating and analysing system-level information propagation

Peter Henderson, Stephen Crouch, Robert J Walters
Electronics and Computer Science, Southampton University, University Road, Southampton, UK

Keywords: Enterprise information systems, Systems-level modelling, System simulation, Distributed systems, Dynamic systems

Abstract: With the proliferation of internet-based technologies within and between organisations, large-scale enterprise systems are becoming more interconnected than ever before. A significant problem facing these organisations is how their information systems will cope with inconsistency being introduced from external data sources. Major problems arise when low quality information enters an authoritative enterprise system from these external sources, and in so doing gains credibility. This problem is compounded by the propagation of this information to other systems and other enterprises, potentially 'invading' an inter-enterprise network. In this paper we will introduce and examine this behaviour, which we term 'information invasion'. Characterisation of systems that are most vulnerable from such an occurrence is provided, and details of an experiment are given which simulates information invasion on an example network topology.

1 INTRODUCTION

Information systems within large-scale enterprises are becoming increasingly large, complex, interconnected and essential to the operation and co-operation of businesses. With the advent of this information revolution, and as these enterprises open up their systems to e-commerce, the volume of information that enterprise systems are empowered to manage is increasing dramatically. Such an increase in information and interconnectivity between enterprise systems introduces new problems to which methods and solutions have yet to be found (Bichler, Segev & Zhao, 1998; Gray, 1996).

A major problem faced by such enterprises is how to deal with information inconsistency within their systems, or put another way, how to survive despite this inconsistency existing within their systems (Henderson, Walters & Crouch, 2001; Anderson et al. 1998). For example, where information is introduced to an enterprise system which conflicts with information that is already held, how is it decided which contribution to adopt? The new information that has been received may be more accurate than the information that is currently held, or vice versa. Knowing which value represents the truth would render the problem trivial. However, determining which value is preferred with respect to the real world may not be practical given the quantity of information handled by enterprise systems. We use the term ‘preferred’ since determining the ‘correct’ value may not be possible.

The cause of this inconsistency might not only be information that is introduced into a system from within its own organisational process, but also from information received from external sources which can also pollute the information pool. It is certainly possible that an enterprise could pass on received low quality information to other interconnected enterprise systems. A factor that significantly compounds this situation is that of authority. Where enterprise systems are considered an ‘authoritative’ information source, their ability to spread potentially low quality information to other enterprises should not be underestimated.

This paper will illustrate how information propagation and authority can together facilitate ‘information invasion’; a phenomenon that can cause debilitating problems within and across inter-enterprise systems. Previously, we have investigated the effects of various behaviours interacting in e-commerce systems (Henderson 2002; 2003) including negotiation (Henderson et al. 2003), and in this paper we demonstrate the effects

Henderson P., Crouch S. and J Walters R. (2004).
INFORMATION INVASION IN ENTERPRISE SYSTEMS - Modelling, simulating and analysing system-level information propagation.
In *Proceedings of the Sixth International Conference on Enterprise Information Systems*, pages 473-481
DOI: 10.5220/0002621204730481

of information invasion with a specified behaviour in an experiment.

2 THE PHENOMENON OF INFORMATION INVASION

2.1 Problems within Modern Enterprise Systems

To illustrate the ways in which modern enterprise systems can cause real-world problems, a certain case provides compelling evidence (Gerth, 2000a; 2000b). Because of a bureaucratic mistake, the individual, whom we will refer to as John Doe, was declared dead. Because of this, his Social Security stipend was stopped, and subsequently his bank account frozen. Additionally, the insurance company paying his medical benefits also stopped paying claims. This simple dependency is shown in Figure 1.

Although the bank account was immediately unfrozen when the error was brought to light, the medical benefits company was unable to restart paying claims due to a lack of proper verification from the government department. Interestingly, the reason for the mistake was a death certificate for a John Doe arriving at the department, although the origin behind the erroneous certificate was unknown. Subsequently, this information was passed to Doe's bank and his medical insurance company. However, according to government officials at the time, a likely cause for this mistake was the death of a John Doe in a nearby county, and "that could be a source of confusion". A spokesman for the insurance company's regional office stated when commenting about Doe's case that "We can't put him back on until they [the government department] reinstate him," and "yes, it does happen from time to time." In effect, they lacked the local authority to correct the information. The matter was eventually resolved when the government department corrected its system, and the insurance company received the update.

Such incidents are not uncommon. An official working in the retirement benefits department for a real estate company elaborated on a similar occurrence (Haga, 2002), explaining that such mistakes happen occasionally when computers at the government department "compare their data base to ours". She added, on the day of interview, that she had had "four calls this morning from people who were told they're listed as dead".

2.2 Characterising this Behaviour

The key issue that initially causes these problems, and also prevents them from being quickly resolved, is the authority of the initial carrier. Information that is received and accepted by such organisations, correct or not, is awarded instant credibility due to their authoritative position. This information then propagates to other systems dependent on this information, and they propagate this information, and so on. When this occurs to the degree that the old information is no longer effectively represented within an inter-enterprise network, the new information has effectively *invaded* the network; hence we term this phenomenon information invasion.

After information invasion has occurred, attempting to revert any subservient system's information to the previous value (assuming this is possible) is very difficult, for two main reasons:

- **Authoritative top-down propagation:** future updates, or database synchronisations, from the authoritative source will filter down through the inter-enterprise network, eventually overwriting any deviation from the authoritative system's spurious view of information.
- **Communal reinforcement:** neighbouring enterprise systems are in the position to reinforce the belief that a certain value (possibly the new one) is correct.

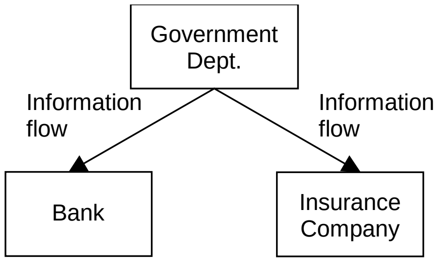

Figure 1: Information dependency between the government department, the insurance company, and the Bank

Communal reinforcement can be achieved in many ways. For example, databases may be synchronised and inconsistent data queried, potentially overwriting any attempt to revert. This is a common feature in replication systems (Thompson 1997). Another example is an informal query. A system that is attempting a reversion may query neighbouring systems directly (or even the authoritative source itself) to determine the 'official' value. This may be achieved using methods other than database synchronisation, such as a telephone or email. In both examples, reverting to a previous value can prove difficult, since the network (and crucially the authoritative source) collectively believes the newly invaded value. Of course, this can make fixing a problem such as Doe's difficult.

2.3 Influences on an Authoritative Source

Systems that contribute information to an authoritative source clearly have a great influence as a potential instigator of information invasion. With many contributing information sources, a critical factor that contributes to information invasion is time.

Let us consider Figure 2, which depicts an enterprise system topology in which three sources (*A*, *B* and *D*) contribute to a single authoritative source (*C*). Assuming these systems adopt the common policy that the latest data received is accepted, a significant problem can result: the *time* that information arrives from one of *A*, *B* and *D* dictates what *C* will believe, and possibly pass on to other systems. If, for example, many systems were dependent on *C* for information, information invasion may occur.

What a system in a network believes, therefore (*C* in our example), is dependent on three key factors of its information sources:

- What information they believe
- Their level of authority
- In what order (as a group) they pass their information on to the receiving system

Time has an obvious impact on the adoption of a certain value by an authoritative source, and hence on information invasion. But in essence, if something is repeated often enough by the source systems, it may eventually be perceived as true by the target system. From a larger perspective, if something is repeated often enough within an enterprise network, it is likely to be perceived as true by the enterprise.

2.4 Properties of Enterprise Systems that Contribute to Information Invasion

In addition to authority, there are other properties of interconnected enterprise information systems that promote the spread of low quality information due to information invasion.

The first of these is interconnectivity:

- **Loose coupling:** enterprise systems are often loosely coupled, where asynchronous methods of information transfer are used. This asynchronous functionality is provided by many middleware solutions, such as Microsoft Message Queue (MSMQ) (Microsoft Corporation, 2004), IBM's WebSphere MQ (formerly MQSeries) (IBM Corporation, 2004), and Sun's Java Messaging Service (JMS) (Sun Microsystems Incorporated, 2004). Asynchronous methods greatly increase the possibility that while data is in transit, it becomes out of date. The possibility exists that out of date information will be favoured over existing information simply because it arrives later. Information can also arrive out of order, unless transaction methods are used. For example, if two sequential updates of someone's address arrived out of order, an information system would end up believing the incorrect one.
- **Intermittent availability:** enterprise information systems can be intermittently

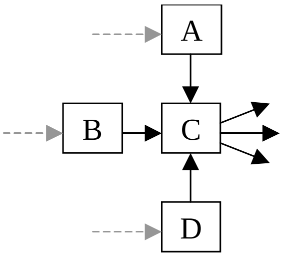

Figure 2: Example inter-enterprise system topology

unavailable, perhaps because of workload or problems with their connectivity (e.g. the internet), either of which may be out of their control. In addition, working practices may contribute to system availability (e.g. working hours). Coupled with the first property, there is no way to guarantee a system is being reliably informed and kept up to date.

- **Responsibility:** there is the ethical problem of responsibility. An enterprise is often exclusively concerned with maintaining the quality of its own information, and not that of others. They often have little regard, or knowledge, of other external information systems that house the same information. They may also have no obligation to ensure their information is consistent with every data source, and they often have no responsibility to ensure external data sources to which they are interconnected are also consistent. Such mentality is predominant, and greatly increases the likelihood of inconsistent information arising within, and between, such enterprise systems.

An issue associated with interconnectivity that promotes inconsistency is that of dependency. Within a network of interconnected enterprises, an enterprise may depend on many other organisations for information. This may mean that one enterprise receives several different values for an information item. This increases the possibility of inconsistency being introduced and information invasion being initiated. Where authority is not a concern, the level of trust awarded to source organisations may form part of the selection process for this information, although such trust relationships require constant re-evaluation (Grandison & Sloman 2000; 2002). Similarly, there may be many enterprises that are dependent on this organisation for information. Inconsistent information that exists within the supplying system may then be passed to those dependent systems, triggering information invasion.

The methods of information-passing between and into enterprises, ‘channels’, can form a contributing factor. It is conceivable that an enterprise may receive many different values for a data item from a single enterprise through *many channels*: such as email, database synchronisation, and other forms of data feed which are not computer system-oriented such as postal mail and the telephone. Unfortunately, more informal channels are often more unreliable, such as a human conversation over the telephone. Interestingly, however, this issue can work both ways. In such situations, humans are able to spot inaccuracies in the information and correct them during the telephone call, which may not be possible with information systems.

As well as interconnectivity, properties of the information that flows through the interconnections can promote information invasion. The quantity and frequency of information received from and passed to other enterprises increases the possibility of inconsistency, and the onset of information invasion. The quality of such information can also prove to be an issue. There may be no way to ensure that information that is received from other enterprises is of the quality required by the receiving organisation.

An issue that can contribute significantly to information invasion is the ownership of information systems across enterprises. Although enterprises may be interconnected, such that one or both are dependent on the other for information, they are not obligated to correct their data if found incorrect by the other enterprise. An enterprise may realise a data partner’s information is partially incorrect, but simply does not possess the authority to correct the information in their partner’s systems. Unfortunately, this can lead to the incorrect information entering their systems over and over again; continually overwriting corrected data (McLaughlin & Krishnamurthy 2003).

3 THE EXPERIMENT

3.1 Overview

In order to understand how information invasion affects information networks where authoritative sources are present, we constructed an environmental framework and corresponding implementation in which networks of different information handling behaviours could be specified and simulated. This enabled the simulation of scenarios more complex than that of Doe. Many simulations were executed, and the results of one of these experiments will be covered in this paper.

Networks of individual information systems or, abstractly, whole enterprises, are represented as individual information ‘nodes’. This allows us to conduct information invasion experiments at either the intra-enterprise or inter-enterprise level. In addition, each node is assigned a numerical measure that indicates the relative authority of that node with respect to the rest of the network.

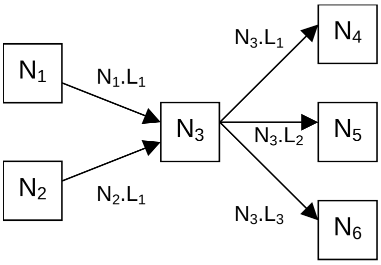

Figure 3: Graph representation of an example information network

Connections between nodes are modelled as queues. This asynchronous model provides a number of advantages. From an intra-enterprise perspective, queuing systems are commonly used on enterprise servers to log incoming requests. From an inter-enterprise viewpoint, the queues effectively simulate the delay between information being sent to an enterprise (via various channels) and being received and handled by that enterprise. Moreover, an asynchronous model provides another realistic degree of uncertainty in that messages can be received in an order other than the order in which they were sent. Therefore, a queuing model awards a realistic abstraction of both perspectives.

In addition to specified behaviour (see section 3.3), each node defined in this framework requires the following parameters to be specified:

- **Authority** the relative strength of influence the node has on other nodes it informs of its belief.
- **Initial belief** what the node believes before the simulation is initiated.
- **Information dependents** list of other nodes that a node can elect to inform of its belief.
- **Reactive rate** represents a node's maximum time in-between handling message(s) in its queue (in milliseconds).
- **Proactive rate** represents a node's maximum time in-between sending message(s) to others (in milliseconds).

A node's reactive and proactive rates represent maximum times between each respective activity; in other words, the maximum intervals between two reactive tasks or two proactive tasks. The time that event t_{n+1} occurs is defined as:

$$t_{n+1} = t_n + random(1\ldots r)$$

Where t_n represents the time at which task n occurred, r represents the maximum reactive rate for a node, and $random(a,b)$ is a function that returns a random value between a and b inclusive. Specifying reactive and proactive intervals in this manner models a node responding to received input and sending data to others within time boundaries.

The structure of an information network in the simulation is a non-negative, unweighted, directed and partially connected graph. Nodes are vertices $\{N_1, N_2, \ldots N_m\}$ where m is the number of nodes, and each node contains a set of links to dependent nodes, represented as $\{N_a.L_1, N_a.L_2, \ldots N_a.L_{Na.d}\}$ where a is the number of the node and $Na.d$ represents that node's total number of dependent nodes. An

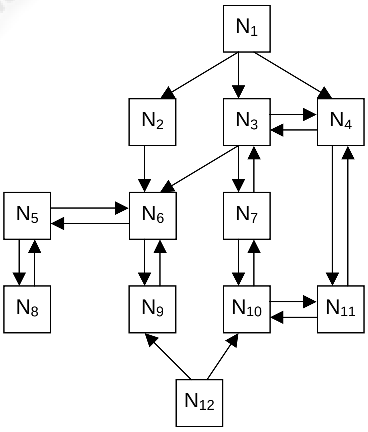

Figure 4: A network topology

example network graph is shown in Figure 3.

Data within a network is represented as a single data item portrayed as a colour, where each colour represents a single possible value for that data item. This enables us to easily visualise the state of a network at any time. We are able to observe clearly the progress of information as it propagates throughout a network.

Each node recorded the state of its belief at every reactive iteration, providing data for analysis. This enables us to determine a node's overall percentage of belief for each colour following a simulation, and determine the overall impact for information invasion. Each of the simulations was executed a number of times and these results averaged to mask anomalies caused by random effects.

A number of experiments were conducted on a number of network topologies, however detailing all of these is beyond the scope of this paper. Figure 4 illustrates the topology upon which the experiments described in this paper were performed. The topology was designed to approximate the complexity and structure typically found in modern computer systems.

3.2 Structure of the Experiment

The experiment was to ascertain how different maximum reactive and proactive rates affected information invasion. N_{12} was given yellow as its initial belief, with authority 2, with all other nodes set to authority 1 and a null belief. A null belief is held with zero authority and therefore can be overwritten by any arriving belief. Each simulation was executed for one minute, and after thirty seconds blue was introduced into N_1 with authority 3. Therefore, yellow invades first from the bottom of the topology, then blue invades with a higher authority from the top of the topology.

Three different simulations were performed on the experiment based on different reactive and proactive rates. The initial starting rate for both (2000 milliseconds) is represented as *x*:

- **Simulation 1** maximum reactive rate = *x*, maximum proactive rate = *x*
- **Simulation 2** maximum reactive rate = *x*, maximum proactive rate = 2*x*
- **Simulation 3** maximum reactive rate = 2*x*, maximum proactive rate = *x*

3.3 Specification of Node Behaviour

For both experiments detailed in this paper, the following behaviour was specified:

```
int belief, authority
collection dependents, queue

on proactive_process {
   select d from dependents
   send(belief, authority) to d
}

on reactive_process {
   select (bel, auth) from queue
   if (auth >= authority) {
      authority = auth
      belief = bel
   }
}
```

According to a node's proactive rate, *proactive_process()* is triggered which selects a dependent at random to which the node's belief (and authority) is sent. Correspondingly, according to a node's reactive rate, *reactive_process()* is triggered which takes a message at random from the node's queue if one exists. Upon receipt, if the authority of the sender's belief is greater than or equal to the authority at which we believe our current belief, our belief and authority becomes equal to the sender's belief and authority.

4 RESULTS

Each figure in each table for the Sim1, Sim2 and Sim3 columns represents the proportion of time during a simulation a node believed blue. During the course of a simulation, the node may have changed its belief many times. For this experiment, however, starting with a null belief that can be overwritten by anything else, nodes will change their mind a maximum of twice: once for yellow introduced by N_{12} and once for N_1 introducing blue, in that order. If a node believes blue before they encounter any yellow, they will remain believing blue. There is also the possibility that they remain with a null belief within the duration of a simulation since they do not receive any belief from anyone.

For each node across all three simulations there are average figures in the last column. The bottom two rows indicate the average and standard deviation for the belief of all nodes throughout each simulation. References to individual table cells are given as *(row, column)*.

4.1 Experiment Results

Table 1 shows the results for this experiment.

Table 1: Results from the experiment

Node	Sim1	Sim2	Sim3	Avg.
N1	0.53	0.53	0.53	0.53
N2	0.47	0.42	0.48	0.46
N3	0.46	0.45	0.42	0.44
N4	0.49	0.45	0.38	0.44
N5	0.28	0.45	0.29	0.34
N6	0.39	0.38	0.20	0.32
N7	0.16	0.37	0.10	0.21
N8	0.13	0.32	0.06	0.17
N9	0.18	0.38	0.10	0.22
N10	0.31	0.32	0.09	0.24
N11	0.41	0.38	0.20	0.33
N12	0	0	0	0
Avg.	0.32	0.37	0.24	0.31
StDev.	0.17	0.13	0.18	0.15

By looking at the (Avg., Sim1-Sim3) cells we can determine the overall impact of introducing blue into each simulation scenario. Simulation one provides a baseline to which we can compare the other two simulations, since the proactive and reactive rates are equal. The results are counter to what we may expect.

In simulation two, we might naively expect blue to propagate more slowly throughout the network, due to the lower proactive rate, but this is not what transpires. We can see that the network as a whole believed the invading value 37% of the time. This is 5% more than in the first simulation, where the proactive rate was as fast as the reactive rate. This counter-intuitive behaviour is due to the queues of the nodes in the network filling up *slower* as a consequence of all nodes' slower proactive rate. Yellow, therefore, does not exist in any great quantity in any node's queue, if at all, before blue is introduced. As a result, when blue enters a node's queue, it has a greater chance of being selected by the reactive function. After which, of course, it will not revert due to blue's higher authority. However, although initially counter-intuitive, this is true to the notion of information invasion. In simulation one, communal reinforcement of yellow at the start of the simulation occurred at a greater rate than in simulation two, due to the higher proactive rate, therefore increasing the chance that yellow was selected. In real terms, this means an enterprise system that conforms to our specified behaviour and consumes less potentially conflicting information from other sources is more prone to information invasion. Such a system is more likely to believe the newly invading value.

In simulation three, the reactive rate is slower than the proactive rate. As a consequence, the queues effectively 'stockpile' received yellow beliefs before blue is introduced. This means that if a blue arrives in a node's queue, it has a reduced chance of being read and set as believed. Therefore, blue becomes little more than 'noise' among the dominant yellow. However, when blue is eventually received and believed by a node, it is propagated as in simulation one. Information invasion occurs as expected, although in the final analysis, with 8% less impact. An enterprise is less likely to contribute to information invasion if it responds to its queues infrequently.

4.2 Other Results

Other experiments have yielded interesting results which hold true to our observations. In an experiment where both N_1 and N_{12} in our topology are equally authoritative and possess different belief values, a state of eternal conflict can result. Each node in the network will believe either value, and so information invasion cannot successfully occur. The beliefs of the network will never acquiesce to a single value. Where simulation three was executed in this scenario, the behaviour of each node's belief over time became chaotic. Initially, information invasion is clearly observed originating from both authoritative sources, but because of the slow reactive rates of the nodes, the queues become filled with both possible beliefs. Even when the inconsistency in beliefs appears to have resolved to a single value over the entire network, and communal reinforcement is maintaining this value, the other value still exists within the queues. This can mean that inconsistency can still potentially arise at any time. However, in simulation two, where the queues are less full, and occasionally empty, the propagation of beliefs is more predictable.

5 CONCLUSIONS AND FURTHER WORK

With information system network topologies becoming more populated and interconnected (Bichler, Segev & Zhao 1998), information invasion will become more widespread and harder to fix. In addition, as information becomes more a commodity, the possibility of low quality information being introduced into such networks becomes a greater risk. We have introduced the concept of information invasion as a natural phenomenon that needs to be recognised and understood if problems such as those experienced by Doe are to be properly resolved.

We have detailed an experiment in information invasion that has presented some unexpected results. Of particular interest is the finding in our experiments that decreasing the rate of belief propagation within a topology of enterprise systems actually increases the rate of information invasion. This is provided that a dominant belief existed within the network before the new belief was introduced. Information is able to invade an information domain more efficiently. The simulation counter to this where queue reading rates were faster confirms this finding. In such a scenario, the faster belief propagation rate hindered the invasion of the new belief.

Inevitably, reality places some constraints upon solutions to alleviate some of the discussed problems. Consider the eternal conflict experiment as outlined in section 4.2. If we stop proactivity (the sending of information) altogether during an execution of simulation two until eventually the queues become empty, its execution becomes less chaotic and more predictable. The same would be true if we had just emptied the queues by throwing away the messages. Of course, in reality this is both impossible and undesirable. It is not always possible to prevent other enterprises sending an enterprise information, and we obviously have no control over all information in transit, especially non-machine oriented methods such as postal mail. Naturally this may not be desirable in any case, since throwing away new information can mean throwing away business.

Exacerbating the phenomenon of information invasion is the issue of the lack of responsibility and accountability for information system accidents that needs to be addressed (Nissenbaum, 1996; 2001). Currently, in the large, enterprises are not obliged to ensure their data is accurate before passing it on. Nor are they accountable in many cases for originating inaccurate information. Such occurrences are simply considered unfortunate mistakes. A lack of local authority can also mean remedying problems caused by information invasion are impossible. In Doe's case, he was able to convince the insurance company of his 'alive' status yet they were unable to fix the problem despite this.

However, various means to address these current attitudes and approaches to information present themselves, and require consideration and investigation.

Firstly, it is not difficult to imagine a common process that is adopted by all those participating in an information domain that is triggered when there is a query regarding the accuracy of some information upon discovering an inconsistency. When this occurs, all participants, including the authoritative source, engage in a discussion to agree on the value of the disputed information. This process could even be automated, utilising a common negotiation framework (Jennings et al, 1998; 2000).

A second approach concerns the traceability of information, in order to identify the original, perhaps authoritative, source. One method of achieving this traceability is to tag all passed-on information with historical metadata. Each time an item of information is received and accepted by an enterprise, it appends it's own identity. Any amendments to the data are also added to the metadata with time information. In effect, this metadata encapsulates the information's entire lifecycle. By examining this metadata, observers not only have access to the origin of potentially inconsistent information, but are also able to ascertain other properties of the information, such as how regularly and frequently it changes. However, this approach is not without obvious issues that need to be addressed, such as security (altering metadata to avoid or implicate responsibility), information size (over time, the size of the metadata could become unmanageable), and metadata inconsistency (the potentially huge task of reconciling inconsistent metadata). A more simple, but less powerful, method of introducing traceability would be for organisations to specifically maintain their database audit information for this purpose. This information is then provided subject to an appropriate and validated request. This could also be an automated process, perhaps provided by a secure web service.

An approach to address the issue of responsibility consists of the notion of an *information responsibility contract*, where enterprises that are involved in sharing information between themselves in a loosely-coupled manner agree to be directly responsible for information they pass on to their contract partners. Such a contract would highlight areas of responsibility for different types of information, depending on origin. This would encourage individual partners to be more proactive in ensuring their information is up-to-date before they pass it on, provide accountability, perhaps even liability, for mistakes, and help to ameliorate the 'denial of accountability' attitude (Nissenbaum, 1996) that can exist in such collectives.

REFERENCES

Anderson, T. et al., 1998. Replication, Consistency, and Practicality: Are These Mutually Exclusive? In *Proceedings of the ACM-SIGMOD 1998 International Conference on Management of Data*, Seattle, WA, pp.484--495.

Bichler, M. Segev, A. & Zhao, J.L., 1998. Component-based E-Commerce: Assessment of Current Practices and Future Directions. *SIGMOD Record*, 27 (4), pp.7-14.

Gerth, J., 2000. Man declared dead by U.S. 'not quite ready to go'. *The Courier Journal*, Louisville, Kentucky, Jan. 3rd.

Gerth, J., 2000. Medicare officials agree with man that he is, in fact, alive. *The Courier Journal*, Louisville, Kentucky, Jan. 6th.

Grandison, T. Sloman, M., 2000. A Survey of Trust in Internet Applications. *IEEE Communications Surveys and Tutorials*, 3 (4), October-December.

Grandison, T. Sloman, M., 2002. Specifying and Analysing Trust for Internet Applications. In *Proceedings of the 2nd IFIP Conference on E-Commerce, e-Business, and e-Government (I3e 2002)*, Lisbon, Portugal, October 7th - 9th.

Gray, J. et al., 1996. The Dangers of Replication and a Solution. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pp.173-182.

Haga, C., 2002. Twin Cities Journal: There's life after death in New Hope, Golden Valley. *The Star Tribune*, Minneapolis, Minnesota, August 5th.

Henderson, P., 2002. Reasoning about Asynchronous Behaviour in Distributed Systems. In *Proceedings of the 8th IEEE International Conference on Engineering of Complex Computer Systems (ICECCS'02)*.

Henderson, P., 2003. Modelling Architectures for Dynamic Systems. In McIver, A., Morgan, C. eds. *Programming Methodology*. Springer-Verlag New York Inc.

Henderson, P. et al., 2003. A Comparison of some Negotiation Algorithms. In Kowalczyk, R. et al. eds. *Agent Technologies, Infrastructure, Tools, and Applications for E-Services*, Lecture Notes for Computer Science, 2592, Springer-Verlag Berlin Heidelberg, pp.137-150.

Henderson, P. Walters, R.J. & Crouch, S., 2001. Inconsistency Tolerance across Enterprise Solutions. In *Proceedings of the 8th IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS'01)*.

IBM Corporation, 2004. WebSphere MQ. Available from: http://www-306.ibm.com/software/integration/wmq/ [Accessed January 29th, 2004].

Jennings, N.R. et al., 1998. On Argumentation-Based Negotiation. In *Proceedings of the International Workshop on Multi-Agent Systems*, Boston, USA.

Jennings, N.R. et al., 2000. Automated Negotiation. In *Proceedings of the 5th International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM 2000)*, Manchester, UK, pp.23-30.

McLaughlin, S. Krishnamurthy, V., 2003. Managing Data Incest in a Distributed Sensor Network. In *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.

Microsoft Corporation, 2004. Microsoft Message Queue Homepage. Available from: http://www.microsoft.com/windows2000/technologies/communications/msmq/default.asp [Accessed January 29th, 2004].

Nissenbaum, H., 1996. Accountability in a Computerized Society, *Science and Engineering Ethics*, 2, pp.25-42.

Nissenbaum, H., 2001. How Computer Systems Embody Values, *IEEE Computer*, March 2001.

Sun Microsystems Incorporated, 2004. Java Messaging Service. Available from: http://java.sun.com/products/jms/ [Accessed January 29th, 2004].

Thompson, C., 1997. Database Replication. *Intelligent Enterprise (formerly DBMS) Magazine*, May. Available from: http://www.dbmsmag.com/9705d15.html [Accessed January 29th, 2004].