Utility-Driven Proactive Management of Availability in Enterprise-Scale Information Flows

Zhongtang Cai¹, Vibhore Kumar¹, Brian F. Cooper¹, Greg Eisenhauer¹,
Karsten Schwan¹, and Robert E. Strom²

¹ College of Computing, Georgia Institute of Technology, Atlanta, GA 30332
{ztcai, vibhore, cooperb, eisen, schwan}@cc.gatech.edu

² IBM Watson Research Center, Hawthorne, NY 10532
robstrom@us.ibm.com

Abstract. Enterprises rely critically on the timely and sustained delivery of information. To support this need, we augment information flow middleware with new functionality that provides high levels of availability to distributed applications while at the same time maximizing the utility end users derive from such information. Specifically, the paper presents utility-driven ‘proactive availability-management’ techniques to offer (1) information flows that dynamically self-determine their availability requirement based on high-level utility specifications, (2) flows that can trade recovery time for performance based on the ‘perceived’ stability of and failure predictions (early alarm) for the underlying system, and (3) methods, based on real-world case studies, to deal with both transient and non-transient failures. Utility-driven ‘proactive availability-management’ is integrated into information flow middleware and used with representative applications. Experiments reported in the paper demonstrate middleware capability to self-determine availability guarantees, to offer improved performance versus a statically configured system, and to be resilient to a wide range of faults.

1 Introduction

Modern enterprises rely critically on timely and sustained delivery of information. An important class of applications in this context is a company’s operational information system, which continuously acquires, manipulates, and disseminates information across the enterprise’s distributed sites and machines. For applications like these, a key attribute is their availability - 24 hours a day, 7 days a week. In fact, system failures can have dire consequences, including loss of productivity, unhappy customers, or serious financial implications. For example, the average cost of downtime for financial companies, as reported in [1], is up to 6.5 million dollars per hour and hundreds of thousands of dollars per hour for retail companies. This has resulted in a strong demand for operational information systems that are available almost continuously.

Providing high availability in widely distributed operational information systems is complex for multiple reasons. First, because information flows are distributed, they are difficult to manage, and failures at any of a number of distributed components or sites can reduce availability. Second, multiple flows may use the same distributed resources, thereby increasing the complexity of the system and the difficulty of managing and preventing failures. Third, such systems often have high data rates and intensive processing

requirements, and there are frequently insufficient system resources to replicate all this data and processing to achieve high reliability. Fourth, information flows must have negligible recovery times to limit losses to the enterprise. Finally, based on our experience working with industry partners like Delta Air Lines and Worldspan, systems must recover not only from transient failures but also from non-transient ones (e.g., failures that will recur unless some root cause is addressed) [2].

How can we provide high availability for information flows, given all of these requirements? Traditional techniques such as recovery from disk-based logs [3] may have recovery times that are unacceptable for the domain in question. Using active replicas [4] imposes high additional communication and processing overheads (since all data flow and processing is replicated) and therefore, may not be an economically viable option. Another option is to use an active-passive pair [4], where a passive replica of a component can be brought up to date by retransmitting messages that had gone to the failed, active one. This option reduces communication costs, since messages are only sent to the passive component at failure time. Unfortunately, this may result in long recovery times. A better solution would be a hybrid of the above approaches, accepting dynamically determined levels of additional processing and communication during normal operation in order to reduce recovery times when failures occur.

In this paper we extend the active-passive approach to dynamically tune the tradeoff between normal operation cost and recovery time. In particular, the passive replica will be periodically refreshed with ‘soft-checkpoints’: these checkpoints transfer the current state from the active node to the passive node (passive standby), but are not required for correctness (hence, they are ‘soft’). If the passive replica has been recently brought up to date by a soft-checkpoint, then recovery will be relatively fast. The tradeoff between cost and recovery is tuned by changing the frequency at which soft-checkpoints are transmitted during normal operation. Such tuning is based on user-provided expressions of information utility, and it takes advantage of the following methods for failure prediction:

- *Availability-Aware Self-Configuration* – a user-supplied per information flow ‘benefit-function’ drives the level of additional resources used to guarantee availability. This ensures preferential treatment of flows that offer more benefit to the enterprise, with the aim of maximizing benefit across the system.
- *Proactive Availability-Management* – during its execution, a system may be at different levels of stability (e.g., a heavy memory load could mean an imminent failure). In many cases, the ‘current stability’ of the system can be quantified in order to increase or decrease the resources expended to ensure desired levels of availability.
- *Handling Non-Transient Failures* – some failures will recur if the same sequence of messages that caused the failure is resent during recovery. In this case, we must use application-level knowledge to avoid fault recurrence. We present several techniques, based on real-world case studies, to deal with such faults.

Proactive availability-management techniques have been integrated into IFLOW, a high performance information flow middleware described in [5]. The outcome is a flexible, distributed middleware for running large-scale information flows and for managing their availability. In fact, experimentation shows that proactive availability-management not only imposes low additional communication and processing overheads on

Fig. 1. Information Flow-Graph and Operator State

distributed information flows, but also, that proactive fault tolerance is an effective technique for recovering from failures, with a low recovery time of 2.5 seconds for an enterprise-scale information flow running on a representative distributed computing platform. Experiments further show that utility-based availability-management offers 1.5 times the net-utility of the basic active replica approach.

1.1 Example: Operational Information System

An operational information system (OIS) [2] is a large-scale, distributed system that provides continuous support for a company or organization's daily operations. One example of such a system we have been studying is the OIS run by Delta Air Lines, which provides the company with up-to-date information about all of their flight operations, including crews, passengers and baggage. Delta's OIS combines three different sets of functionality:

- *Continuous data capture* - for information like crew dispositions, passengers, airplanes and their current locations determined from FAA radar data.
- *Continuous status updates* - for low-end devices like airport flight displays, for the PCs used by gate agents, and even for large databases in which operational state changes are recorded for logging purposes.
- *Responses to client requests* - an OIS must also respond to explicit client requests, such as pulling up information regarding a particular passenger, and it may generate additional updates for events like changes in flights, crews or passengers.

In this paper, we model the information acquisition, manipulation, and dissemination done by such an OIS as an information flow graph (a sample flow-graph is shown in Figure 1). We then present techniques, based on this flow-graph formalization, to proactively manage OIS availability such that the net-utility achieved by the system is maximized. This is done by assigning per information flow availability guarantees which are aligned with the benefit that is derived from the information flow, and by proactively responding to perceived changes in system stability. We also present additional techniques, based on real-world case studies, that can help a system recover from non-transient failures.

2 System Overview

This section describes a model of the information flows under consideration, and it elaborates the fault model used for the proactive availability-management techniques explained later.

2.1 Information Flow Model

An information flow is represented as a directed acyclic graph $G(V_g, E_g, U_{net})$ with each vertex in V_g representing an information-source, an information-sink or a flow-operator that processes the information i.e. $V_g = V_{sources} \cup V_{sinks} \cup V_{operators}$. Edges E_g in the graph represent the flow of information, and may span multiple intermediate edges and nodes in the underlying network. The utility-function U_{net} is defined as:

$$U_{net} = Benefit - Cost \quad (1)$$

Both benefit and cost are expressed in terms of some unit of value delivered per unit time (e.g., dollars/second). Benefit is a user-supplied function that maps the delay, availability, etc. of the information flow to its corresponding value to the enterprise. Cost is also a user-supplied function; it maps resources such as CPU usage and bandwidth consumed to the expense incurred by the enterprise. We will expand the terms of this seemingly simple equation in upcoming sections.

2.2 Fault Model

We are concerned with failures that occur after the information flow has been deployed. In particular, we consider fail-stop failures of operators that process events. Such failures could result from problems in the operator code or in the underlying physical node. Other factors might also cause failures, but are not considered here, including problems with sources, problems with the sink, or link failures between nodes. While such issues can cause user-perceived failures, they must be addressed with other techniques. For example, link failures could be managed by retransmission or re-routing at the network level.

For the purpose of failure recovery, we assume that each flow-operator consists of a *static-state* S_{static} that contains the information about the edges connected to the operator and the enterprise logic embedded in the operator; in contrast, the *dynamic-state* $S_{dynamic}$ is the information that is a result of all the updates that have been processed by this operator (shown in Figure 1). Recovery therefore, is dependent upon the correct retrieval of the states S_{static} and $S_{dynamic}$, which jointly contain the information necessary for re-instantiation of flow-operator and information flow edges. However, as described next, simply recovering these states may not prevent the recurrence of a failure.

Transient Faults. A fault can be caused by a condition that is transient in nature (e.g., a memory overload due to a mis-behaving process). Such faults will not typically recur after system recovery. In our formulation, a transient fault would cause the failure of an operator, and correct retrieval of the two states associated with the operator would ensure permanent recovery from this fault. The techniques proposed in this paper are capable of effectively handling faults of this nature.

Non-transient Faults. Non-transient faults may be caused by some bugs in the code or some unhandled conditions. For information flows, this may mean recurrence of the fault even after recovery, particularly when recovery entails repeating the same sequence of messages that caused the fault. To deal with faults of this nature, we note that the output produced by a flow-operator in response to an input event E depends on the existing dynamic-state $S_{dynamic}$, the operator logic encoded as S_{static} , and the event E itself. Therefore, the failure of an operator on arrival of an event E is a result of the 3-tuple $\langle S_{dynamic}, S_{static}, E \rangle$. Thus, any technique that aims to deal with non-transient failures must have application-level methods for retrieving and appropriately modifying this 3-tuple. Our prior work presents examples of such methods [2], and we generalize such techniques here.

3 Utility-Driven Proactive Availability-Management

Traditional techniques for availability-management typically rely on undo-redo logs, active-replicas, or active-passive pairs. A new set of problems is presented by information flows that form the backbone of an enterprise. For instance, using traditional on-disk undo-redo logs for information flows would lead to unacceptable recovery times for the enterprise domain in face of machine or disk failures. The other end of the availability-management spectrum, which uses active replicas, would impose large additional communication and processing overheads due to the high arrival rate of updates, typically making it economically infeasible for the enterprise to use this option. In response, we take the active-passive pair algorithm [4], and customize it for enterprise-scale information flows. To do this, we will incorporate our previous work on soft-checkpoints [6], and add the ability to dynamically choose checkpointing intervals to reduce communication and processing overheads. For completeness, we first describe the existing active-passive pair and soft-checkpoint techniques, and then describe our enhancements.

3.1 Basic Active-Passive Pair Algorithm

To ensure high-availability for the flow-operator, in its simplest form, the active-passive pair replication requires: a *passive node* containing the static-state S_{static} of the flow-operator hosted on the active node, an *event log* at the flow-graph vertices directly upstream to the flow-operator in question, a mechanism to *detect duplicates* at the vertices directly downstream to the flow-operator, and a *failure detection* mechanism for the active node hosting the primary flow-operator.

In case of a failure, recovery proceeds as follows: the failure detection mechanism detects the failure and reports it to the passive node. On receipt of the failure message, the passive node instantiates the flow-operator, making use of the static-state, S_{static} , already available at the node. The instantiated operator then contacts the upstream vertices for retransmission of the events in their event log. The newly instantiated operator node processes these re-transmitted events in a normal fashion, generating output events, and leaving it to the downstream nodes to detect the resulting duplicates. Once the retransmission of the event log has been completed, the resulting dynamic-state,

$S_{dynamic}$, will be recovered to the state of the failed operator, and normal operations can resume. Unfortunately, this simple algorithm can lead to long recovery times, large event logs at the upstream nodes, and large associated retransmission costs. The remedy to these problems is the ‘soft-checkpoint’ technique, described next.

The event logs at the upstream nodes and their retransmission to the recovered operator are required for reconstructing the dynamic-state $S_{dynamic}$, of the failed operator. However, in practice, it is advantageous to retain additional stable state at the passive node in order to avoid the need to re-transmit the entire event log. Such state saving is called soft-checkpointing, because it is not needed for correctness. Soft checkpoints can be updated on an intermittent basis in the background. Once taken, the component receiving the checkpoint no longer requires the events on which the state depends for reconstructing $S_{dynamic}$. This in turn permits upstream nodes to discard the event logs for which the soft-checkpoint has been taken. Soft-checkpointing, therefore, is an optimization that reduces worst-case recovery time and permits the reclamation of logs.

The introduction of soft-checkpoints requires small modifications to the recovery mechanism described earlier in this section. The flow-operator at the active node in the duration prior to failure would intermittently send messages to the passive node that contain information about the incremental change to its dynamic-state since the last message. The passive node, after the receipt of complete state update message from the active node, applies the incremental modifications to the state it holds and then sends a message to the flow-operator’s upstream neighbors about the most recent event contained in the message from the active node. The upstream nodes can use such information to purge their event logs. In case of a failure, the algorithm proceeds exactly as described earlier, but only a small fraction of the events needs to be re-transmitted and processed.

3.2 Availability-Utility Formulation

In this section, we use a basic availability formulation to better describe the effects and trade-offs in soft-checkpoint-based active-passive replication. Availability $\mathcal{A}_T$ is described in terms of Mean-Time Between Failure, $MTBF$ and Mean-Time To Repair, $MTTR$.

$$\mathcal{A}_T = \frac{MTBF}{MTBF + MTTR} \quad (2)$$

As stated earlier, our approach contributes to a reduction in recovery time and also reduces the processing and communication overhead imposed as a result of ensuring a certain level of availability. The reduction in recovery time results in lower MTTR and a reduction in associated overheads diminishes cost. Jointly, both result in higher net-utility U_{net} , which is the actual utility provided by the system.

With our approach, MTTR depends on two factors: (1) the time to detect a failure, and (2) the time to reconstruct the dynamic-state of the operator. Failure detection mechanisms generally rely on time-outs to detect failures and therefore, depend on the coarseness of the timer used for this purpose. Some research in the domain of fault-tolerance has focused on multi-resolution timeouts [7], but to simplify analysis, henceforth, we assume that the time to detect a failure is a constant. The second factor contributing to MTTR depends on the soft-checkpoint algorithm. Specifically, a higher

frequency f_{cp} , expressed in per unit time, of such checkpoints would lead to a smaller number of events required to reconstruct $S_{dynamic}$ in case of a failure. Therefore:

$$MTTR \propto \frac{1}{f_{cp}} \quad (3)$$

For simplicity, we next derive the availability-utility formulation for a single information flow (self-configuration across multiple information flows is addressed in Section 3.3), and we assume that the *Benefit* and *Cost* depend only on availability. In this case, in general, the benefit derived from a system is directly proportional to its availability. Thus:

$$Benefit \propto \frac{MTBF}{MTBF + k_1/f_{cp}} \quad (4)$$

The above formulation may lead one to believe that a higher f_{cp} is good for the system. Unfortunately, a higher f_{cp} also means more cost to propagate checkpoints from the active node to the passive node. Therefore:

$$Cost \propto f_{cp} \quad (5)$$

Note that a higher f_{cp} also results in fewer events retransmitted per soft-checkpoint; however, for large values of MTBF this effect is minor compared to the effects described above (increase in benefit due to better availability, and compared to the increase in cost due to a higher frequency of checkpoints). Experiments reported in Section 5.2 study the effects of soft-checkpoint frequency on the cost and availability of information flows.

Combining equations 1, 4, 5, and replacing proportionality using constants, we arrive at:

$$U_{net} = \frac{k_2 \times MTBF}{MTBF + k_1/f_{cp}} - k_3 \times f_{cp}, \quad (6)$$

which represents the business-utility calculation model and the constants are determined by business level objectives [8,5], or using more detailed formulation described later. This equation expresses the key insight that net-utility depends not only on MTBF, but also on the soft-checkpoint frequency used in a system, the latter both positively contributing to net-utility (by reducing the denominator) and directly reducing net-utility (by increasing the term being subtracted). Intuitively, this means that frequent checkpointing can improve utility by reducing MTBF, but that it can also reduce utility by using resources that would otherwise directly benefit the information flow.

3.3 Availability-Aware Self-configuration

Ideally, we would like to maximize the availability of an information flow, but given that there is an associated cost, our actual goal is to choose a value of availability that maximizes its net-utility. In our algorithm and its mathematical formulation, f_{cp} is the factor that governs availability. By setting the derivative of equation 6 equal to zero, we find that the value of f_{cp} that maximizes net-utility is:

$$f_{cp} = \sqrt{\frac{k_1 \times k_2}{k_3 \times MTBF}} - \frac{k_1}{MTBF} \quad (7)$$

Fig. 2. Enterprise error-log showing predictable behavior of failures

In the presence of multiple information flows, each with a different benefit-function, the resource assignment for availability is driven by the need to maximize net-utility across all deployed information flows. Total net-utility of the entire system, then, is the sum of individual net-utilities of information flows. For a system with n information flows, we will need to calculate $\{f_{cp}^1, f_{cp}^2, \dots, f_{cp}^n\}$, which will automatically determine resource assignments. The value of f_{cp} for each information flow can be calculated using partial differentials, and the involved calculations are omitted due to space constraints.

3.4 Proactive Availability-Management

We have established that net-utility depends on checkpoint frequency and MTBF. However, the MTBF in a real system is not a constant. Instead, the rate of failures fluctuates, with more failures occurring when the system is in an unstable state. For example, during periods of extreme overload, the system is likely to experience many component failures. If we can better approximate the current MTBF, and in particular predict when there will be many failures, we can make better decisions about checkpointing, increasing the checkpoint rate when the current MTBF is low (and failures are imminent.)

Failure Prediction. An effective way to estimate the current MTBF is to use failure prediction techniques to generate 'early alarms' when a failure seems to be imminent. By using failure prediction, our approach can be 'better prepared' for an imminent failure, by taking more frequent soft-checkpoints. Analysis logs provided to us by one of our industry partners strengthens our belief in the usefulness of dynamic failure prediction. These logs contain error messages and warnings that were recorded at a middle-ware broker over a period of 7 days, along with their time-stamps. Figure 2 shows the distribution and severity of errors recorded at the broker node. One interesting observation of these logs is that errors recur at almost the same time (around 9:00am as read from the log time-stamp) beginning from the 2nd day. Another interesting observation about the same set of logs is that 128 errors of severity level 1 occurred from 7:30pm in the first day before a series of level 4 errors occurred from 8pm. Based on such logs, it would be reasonable, therefore, to assume lower MTBF (i.e., predict imminent failures) for the 9am time period and the period when a large number of less severe errors occur,

than for other time periods in which this application executes. We note that similar time- or load-dependent behaviors have been observed for other distributed applications[9].

We implemented the Sequential Probability Ratio Test (SPRT) used in MSET [10,11] failure prediction method, to predict failures injected by the FIMD [12] failure injection tool, including timing delay, omission, message corruption datatype, message corruption length, message corruption destination, message corruption tag, message corruption data, memory leak, and invalid memory access. The SPRT method is a run-time statistical hypothesis test that can detect statistical changes in noisy process signals at the earliest possible time, e.g., before the process crashes or when severe service degradation occurs. SPRT has been applied successfully to monitor nuclear power plants, and it has recently been used for software aging problems, e.g., for the database latch contention problem, memory leaks, unreleased file locks, data corruption, etc. For example, an early warning may be raised about 30 seconds (the 'early warning capability') before a memory leak fault causes the service to degrade dramatically or the process crashes. For database shared-memory-pool latch contention failures, early warning capabilities of 5 minutes to 2 hours have been observed. For additional information about SPRT and associated MSET method, please refer to [10] and to an extended version of this text in a technical report [13].

Modulating Checkpoint Frequency. The idea behind proactive availability-management is to use failure prediction to modulate f_{cp} . We first provide the important yet simple guideline regarding checkpoint frequency modulation, we then develop a detailed formulation for enterprise-scale information flows, and finally, present a formulation and method to meet some specific availability requirement while also maximizing net-utility.

General guideline. Intuitively, if a failure prediction turns out to be correct, the system 'benefits' because of reduced $MTTR$; if a prediction turns out to be a false-positive, the system still operates correctly, but it pays the extra 'cost' due to increased f_{cp} . Stated more formally, let:

$$\begin{aligned}\alpha &= \text{prediction false-positive rate} \\ \beta &= \text{prediction false-negative rate} \\ f'_{cp} &= \text{modulated checkpoint frequency after a failure is predicted} \\ T_{proactive} &= \text{duration of increased checkpoint frequency} \\ k &= \text{timeout after which an operator is concluded to have failed}\end{aligned}$$

Earlier, $Cost$ was shown to be proportional to soft-checkpoint frequency. The new cost, $Cost'$, due to modulated f'_{cp} , is:

$$Cost' = Cost \times f'_{cp} / f_{cp} \quad (8)$$

This increased cost is incurred for a duration equal to $T_{proactive}$, and it is incurred each time a prediction is made. Therefore, the additional cost incurred per prediction is:

$$\delta Cost = (f'_{cp} / f_{cp} - 1) \times Cost \times T_{proactive} \quad (9)$$

Table 1. Four types of cost

$Cost^{cp} = [1 - P(1 - \beta + \alpha)]f_{cp}C_1,$
$Cost^{fp} = \alpha Pf'_{cp}t_oC_1,$
$Cost^{fn} = \beta PC_2/(2f_{cp}) + \beta P(k+1)/(2f_{cp})C_3,$
$Cost^{ps} = (1 - \beta)P[C_2/(2f'_{cp}) + (k+1)/(2f'_{cp})C_3 + f'_{cp}t_oC_1].$

The increase in f_{cp} also affects the availability of the system and therefore, the benefit, $Benefit'$, derived from the system. Using equation 4, we have:

$$Benefit' = \frac{MTBF + k_1/f'_{cp}}{MTBF + k_1/f_{cp}} \times Benefit \quad (10)$$

Therefore, the increase in benefit due to a correct prediction that affects a period equal to $MTBF$ is:

$$\delta Benefit = (Benefit' - Benefit) \times MTBF \quad (11)$$

Since λ is the fraction of false-positives and because there is no increase in benefit due to a false positive, the following condition expresses when proactive availability-management based on failure prediction is beneficial for an entire system:

$$\delta Cost < (1 - \alpha) \times \delta Benefit \quad (12)$$

Proactive availability-management. Different systems could have different types and formulations of benefit and cost, and the above analysis provides the general guideline regarding proactive availability-management. For the enterprise information flow system targeted by this paper, the proactive availability-management problem can be formulated in more details as follows. Proactive availability-management regulates checkpoint frequency based on stability predictions to maximize net business utility. By considering ‘total cost’, including the cost of checkpointing and the utility loss because of a failure (i.e. the extra utility the system could offer if there had been no failure), the problem of maximizing net-utility can be converted to the problem of minimizing total cost. This total cost consists of the cost of normal checkpointing (at frequency f_{cp}), $Cost^{cp}$, the cost due to false-positive failure prediction (i.e., the failure predictor raises a false alarm), $Cost^{fp}$, the cost due to false-negative failure prediction (i.e., a failure is not predicted successfully), $Cost^{fn}$, and finally, the cost associated with failure recovery when a failure is successfully predicted, $Cost^{ps}$.

These four types of cost are summarized in Table 1. For the cost of normal checkpoints, $Cost^{cp}$, C_1 is the cost for each checkpoint update (e.g., the communication cost), and P is the possibility an operator could fail from any time t to $t+1$ (seconds). Here, $P(1 - \beta + \alpha)$ is the fraction of time when the checkpoint frequency is f'_{cp} , due to correct failure predictions and false alarms. For the cost of false-positive failure prediction, t_o is the average time a predictor raises an early alarm for a severe failure. In the equation for the cost due to false-negative prediction, $Cost^{fn}$, the first term is the total state recovery cost, i.e., the cost for the passive node to recover from the latest

checkpointed state to the state when the failure occurred, including retransmission cost and re-computation cost. C_2 is the average recovery cost per unit time($\$/sec$). The second term is the total utility loss from the time when failure occurs to the time when the system recovers to normal operational status. In other words, this term represents the utility the system could provide if there had been no such failure. C_3 is the utility the system provides per second($\$/sec$) if there is no failure. The cost associated with failure recovery when a failure is successfully predicted, $Cost^{ps}$, is determined in a similar manner as $Cost^{fn}$.

To regulate checkpoint frequency, proactive fault tolerance finds the best checkpoint frequency, f_{cp} , when there is no failure predicted, and the best checkpoint frequency, f'_{cp} , after the time a failure is predicted. This is done by minimizing the total cost.

Meet specific availability requirement. Often, enterprises have specific requirements for system availability. For example, a 365 x 24 system with maximum allowed average downtime of 8.76 hours (i.e., 525 minutes) per year requires 99.9 percent availability, while a system with only 3 minutes of service outage must have at least a 99.999 percent availability. To achieve such availability is difficult due to the high cost of fault tolerance services and equipments. Proactive availability-management is able to strike a balance between these two factors by jointly considering availability and utility when regulating checkpoint frequency. Notice that MTTR can be expressed as:

$$MTTR = (1/2f_{cp} + k)\beta + (1/2f'_{cp} + k)(1 - \beta), \quad (13)$$

where k is the timeout after which we conclude that a module actually failed, the availability is given by:

$$\begin{aligned} A_I &= \frac{MTBF}{MTBF + MTTR} = \frac{1 - P \cdot MTTR}{1} \\ &= 1 - p[(1/2f_{cp} + k)\beta + (1/2f'_{cp} + k)(1 - \beta)] \end{aligned} \quad (14)$$

Proactive fault tolerance meets the minimum availability requirement and also maximizes net utility by solving the following equation:

$$\text{Minimize}\{Cost = Cost^{cp} + Cost^{fp} + Cost^{fn} + Cost^{ps}\}, \text{ subject to:}$$

$$1 - p[(1/2f_{cp} + k)\beta + (1/2f'_{cp} + k)(1 - \beta)] \geq A_I^{required} \quad (15)$$

This optimization problem is of small size with two variables and one constraint, and is solved using standard Quasi-Newton method with inverse barriers.

3.5 Handling Non-transient Faults

Non-transient failures are a result of bugs or unhandled conditions in operator code. Traditional techniques for ensuring high-availability that use undo/redo logs [3,6] are useful for transient failures, but for non-transient failures, they may result in recurrence of faults during recovery. The same applies to replication-based approaches [14], for which all replicas would fail simultaneously for non-transient faults.

As described in Section 2.2, a non-transient failure of the information flow in our model is a result of the 3-tuple $\langle S_{static}, S_{dynamic}, E \rangle$. The active-passive pair approach for ensuring high-availability has sufficient information during recovery to change this 3-tuple. The passive-node during recovery has access to S_{static} , a stale state $S'_{dynamic}$, and a set of updates T from the upstream nodes that when applied to $S'_{dynamic}$, would lead to $S_{dynamic}$. The rationale behind our approach to avoid non-transient failures is simple: avoid the 3-tuple that caused the failure. This can be done in a number of ways, and the retransmitted updates T along with application-level knowledge holds the key:

- *Dropping Updates*: the simplest solution to avoid recurrence of a fault is to avoid processing the update that caused the failure. Our earlier work on ‘poison messages’ used this technique [2].
- *Update Reordering*: changing the order in which updates are applied to $S'_{dynamic}$ during recovery can avoid $S_{dynamic}$. This makes use of application-level knowledge to ensure correctness.
- *Update Fusion*: combining updates to avoid an intermediate state could be an option. A simple example of this approach could be the use of this technique to avoid ‘division by zero’ error.
- *Update Decomposition*: decomposing an update into a number of equivalent updates can be an option with several applications, and this can potentially avoid the fault.

While seemingly simple, the techniques described above are often successful in realistic settings. For example, one of our collaborators, reported an occasional surge in the usage of resources connected to their Operational Information System (OIS) [15] that traced back to a particular uncommon message type. The resulting performance hit caused other subsystem’s requests to build up, including those from the front ends used by clients, ultimately threatening operational failure (e.g., inappropriately long response times) or revenue loss (e.g., clients going to alternate sites). Such uncommon request/message, termed ‘Poison Messages’, were later found to be identifiable by certain characteristics. The solution then adopted was to either drop or re-route the poison message in order to maintain operational integrity.

4 Middleware Implementation

IFLOW [5] is an information flow middleware developed at Georgia Tech. IFLOW implements the information flow abstraction of Section 2.1 and provides methods to deploy and then optimize (by migrating operators) the information flow. For more details please refer to [16].

We now briefly describe the features that enable proactive availability-management in the IFLOW middleware. These features are implemented both at the *control plane* and the *data plane* of this middleware infrastructure.

4.1 Control Plane

The control plane in IFLOW is the basis for managing information flows. Self-management methods involve running a self-configuration and a self-optimization

algorithm, carried out by exchanging control messages between physical nodes that are external to the data fast paths used to transport IFLOW data. Control actions involve operations like flow-control, operator re-instantiation, etc. The main new features of the IFLOW control plane that are used for proactive fault tolerance are described below:

- *Availability-aware self-configuration module*: the benefit-formulation in IFLOW allows for availability goals to be specified, and determines the best value of f_{cp} by using the formulation described in Section 3.2.
- *Failure detection & prediction*: IFLOW attempts to use the regular traffic from a node to determine its liveness, but it switches to specific detection messages if there is no regular traffic from the node to the monitoring node. We also have a provision for multi-resolution timeouts to reduce the load imposed by the failure detection algorithm. Finally, state can be maintained to use failure history for predicting failures, but we have not yet implemented any specific technique into IFLOW.
- *Control messages*: SOAP calls are used to notify active-node failure, to communicate log purge points to upstream vertices, etc.
- *Update re-direction in case of failure*: a simple control mechanism exists at the upstream vertices to re-direct updates to the passive node in case of failure. The connection between upstream vertices and the passive node is created at the time of flow deployment.

4.2 Data Plane

A *fast data-path* is one of the key design philosophies of the IFLOW middleware. We have taken care that the features required for proactive availability-management have minimal impact on the data-path. In order to ensure proactive availability-management, the state of an operator on the data plane needs to be soft-checkpointed and the changes need to be periodically communicated to the passive-node. The fact that a soft-checkpoint is not necessary for correctness of proactive availability-management ensures minimal impact on the data-path. Specifically, the active-node can transfer the soft-checkpoint to the passive node asynchronously, and this will not compromise the correctness of our algorithm. The specific features required for proactive availability-management are described below:

- *Logging at upstream vertices*: any update that is sent out from the source vertex is logged to enable retransmission in case of failure. Additional logs can be established at intermediate nodes (an operator vertex is a source for downstream vertices) to enable faster recovery. The log module also implements a mechanism to purge the log when a message is received from the downstream node after a soft-checkpoint is completed.
- *Soft-checkpoint module at operator vertices*: the soft-checkpoint module tracks the changes in $S_{dynamic}$ since the last soft-checkpoint. It is also responsible for sending soft-checkpoints to the passive node.
- *Duplicate detection at the downstream node*: the duplicate detection mechanism is based on the monotonic update system proposed in our earlier work [6]. When the updates cannot be ordered using the contained attributes, a monotonically increasing

Fig. 3. Sample testbed. The testbed topology is generated using GT-ITM and is configured at emulab facility.

attribute (e.g., the real-time clock) is appended to the out-going update that uniquely identifies this update.

- *Additional edge between active-passive pair:* a supplementary data-flow between the active-passive pair delivers the soft-checkpoints to the passive vertex.
- *Maintaining checkpoints at passive-node:* the passive vertex contains the logic that applies an incoming soft-checkpoint to the recorded active node state.

5 Experiments

Experiments are designed to evaluate the performance our proactive availability-management techniques. First, simulations are used to better understand the behavior of the self-configuration module that determines the availability requirement based on the user-supplied benefit function. Next, an end-to-end setup is created on Emulab [17], representing an enterprise-scale information flow to compare our approach against the traditional approaches and to study the effect of different soft-checkpoint intervals and proactivity on aspects like MTTR, recovery cost, and net-utility. Results show that proactive availability-management is effective at providing low-cost failure resilience for information flow applications, while also maximizing the application's net-utility.

5.1 Simulation Study

A simulation study is used to compare utility-based availability management to simple approaches that are not availability-aware. The study uses a 128 node topology generated with the GT-ITM internetwork topology generator [18]. The formulation of net-utility U_{net} determines benefit as: $benefit = k_1 \times (k_2 - delay)^2 \times availability \times$

Table 2. Self-Determining Availability based on Benefit

Optimization Criterion	Utility	Cost	Delay
Net-Utility (dollars/sec)	431991	52670	2160
Cost (dollars/sec)	79875	14771	80315
Delay (msec)	222	444	191
f_{cp} (sec^{-1})	0.050	0.018	0.020
Availability (percent)	99.88	99.66	99.70

$availableBandwidth/requiredBandwidth$, and cost is calculated as: $cost = dataRate \times bandwidthCostPerByte$. Random costs are assigned to the network links, expressed in dollars per byte. We substitute ($k_1 = 1.0$, $k_2 = 150.0$) in the benefit formulation for this specific simulation [5]. The MTBF is assumed to be 86400sec. and the MTTR is assumed to be 864sec. for a f_{cp} value of 0.01Hz. (Many values are possible for these variables. However, we must choose some values when conducting our simulations, and the ones we chose are reasonable for the enterprise environment.) We first deploy the flow-graph using the net-utility specification from equation 1 as the optimization criteria, and the results are shown in Table 2 under the column labeled ‘Utility’. The results show a high achieved net-utility with acceptable values for delay, f_{cp} and availability. The second deployment (under ‘Cost’) focuses instead on minimizing the cost, and it uses $1/cost$ as the optimization criteria. The effect of choosing a different criteria is evident in the reduced cost, achieved by allowing a higher delay and a lower availability (resulting from lower f_{cp}). The final experiment uses $1/delay$ to drive the deployment. This results in a reduction of delay achieved for the flow-graph, but at the expense of net-utility and availability.

5.2 Testbed Experiments Using IFLOW

This set of experiments is conducted on Emulab [17], and the network topology is again generated using the GT-ITM internetwork topology generator. In many cases, enterprises would hand tune their topology for availability and performance, instead of using an arbitrary topology. For example, an enterprise may explicitly designate a primary and secondary data center. An arbitrary topology is used in our experiments in order to understand how our techniques perform without the benefit of additional hand tuning. Figure 3 shows the testbed used for experimental evaluations. Background traffic is generated using cmu-scen-gen [19], injected into the testbed using rate-controlled udp connections. For the testbed depicted in Figure 3, background traffic is composed of 900 CBR connections. We use the utility formulation in Equation 15 to better study the net-utility and the costs associated with checkpointing and failures. Required availability is 99.9% if not stated otherwise.

Variation of Net-Utility for Different Approaches. The first experiment studies the variation of net-utility with different availability-management approaches in the presence of failures. For simplicity, only one failure is injected into the system. We conduct experiments with the active replication approach, the passive replication approach with varying soft-checkpoint intervals, and our proactive replication approach. Figure 4

Fig. 4. Net utility rate variations using active, passive or proactive fault tolerance approaches. A failure is injected into one operator node at the time $t = 40s$.

clearly demonstrates that the active replication approach provides lowest net-utility. This is because of the high amount of replicated communication traffic when using this approach. After a failure, net-utility of the active approach increases slightly; there is less replication traffic, because the failed node no longer sends replicated output updates. The experiment also corroborates the analysis in Section 3.2: a lower soft-checkpoint interval for the passive approach imposes higher communication cost on the system and therefore, results in lower net-utility. Note that if availability were a predominant factor in the net-utility formulation, then a lower soft-checkpoint interval could have resulted in higher net-utility. The cost of soft-checkpoints is almost negligible when the interval is greater than 5 seconds, but its effect is evident for an interval of 2 seconds.

Our proactive approach provides the highest net-utility overall, as it modulates the soft-checkpoint interval and takes into account the perceived system to offer preventive fault tolerance. For instance, it switches to a smaller soft-checkpoint interval just before the failure and is therefore able to recover as fast as the passive approach with a 2 seconds update interval, while performing as well as the passive approach with a 30 seconds update interval at other times. We note that evaluation of failure prediction techniques is not the focus of this paper (such kind of evaluations appear in [13]). To investigate how prediction accuracy affects the system, these experiments simulate a predictor for the proactive approach, with failure prediction statistically generated at various levels of accuracy. In particular, we notify the soft-checkpoint mechanism that a failure is imminent, no matter whether the prediction is correct or a false positive.

Variation of MTTR for Different Approaches. The variation of MTTR and its standard deviation with different approaches are shown in Figure 5. For each approach, nine experiments are used to obtain the mean and standard deviation. The active replication approach (not shown in the graph) has no explicit recovery time. This is because the node downstream of the replicated operator continues to receive processed updates even after the failure of one active replica. On the other hand, the passive replication approach which attempts to avoid the high cost of active replication incurs recovery

Fig. 5. MTTR and standard deviation of recovery time under three replication strategies. Standard deviation is represented by vertical error bars.

Fig. 6. Utility before failure and during recovery, and the total cost to recover from one failure

times that increase with the soft-checkpoint interval. The reason for this increase is the time taken for reconstructing the operator state: the higher the soft-checkpoint interval, the larger the number of updates required to rebuild the state. Recovery time for the passive replication approach depends on the soft-checkpoint interval. It ranges from 3.7 seconds (for a 2 second interval) to 14.8 seconds (for a 30 second interval). Our proactive approach, as expected, performs well as compared to other passive replication approaches, since it is able to change over to a very small soft-checkpoint interval just before the failure, and hence, has low MTTR. The experiment demonstrates the importance of choosing the right soft-checkpoint interval automatically to maximize availability at low cost and thereby maximize the net-utility of information flows.

Cost & Net-Utility During Recovery. Our proactive availability-management approach increases soft-checkpoint activity when a failure is predicted in the near future, but it maintains a low soft-checkpoint activity at other times. The analysis of net-utility value before failure, during failure recovery, and the total cost to recover from failure are summarized in Figure 6. Net-utility using proactive availability-management is higher than any other approach, because it contains a very recent soft-checkpoint for the operator state and therefore, incurs the least cost during recovery. Note that passive replication with an interval of 2 seconds also incurs a low cost during recovery, but this is achieved by losing non-negligible net-utility at normal operation time.

Effects of Checkpoint Frequency and Prediction Accuracy on Cost and Availability. The next experiment closely examines the effect of checkpoint frequency on the system, both in terms of system availability and the cost imposed to gain a unit amount of utility. As mentioned in Section 3.2, a higher f_{cp} leads to a higher number of soft-checkpoint messages from the active to the passive node, but it also leads to a smaller number of updates being required to reconstruct the operator state during recovery. The conflicting behavior of incurred cost due to f_{cp} is represented in Figure 7 by the two parabolic curves. Ideally, we would like to spend the minimum cost to achieve a unit amount of utility and would therefore, like to choose a value of f_{cp} that is located

Fig. 7. Effect of checkpoint frequency on cost and availability. Checkpoint frequency affects the cost (left y-axis) in a non-linear way, and it is important to optimize it. Note that there is also a sweet spot in the graph, where cost is minimized and availability (right y-axis) is also high. Our proactive approach can achieve the same level of availability with significantly less cost compared to the passive approach.

Fig. 8. Effect of prediction accuracy on cost of ensuring availability. Better prediction accuracy helps reduce the cost incurred for ensuring high-availability, especially when the checkpoint frequency is low (the curve with the deepest slope). When checkpoint frequency is sufficiently large (the four curves with $f_{cp} \geq 1Hz$), higher accuracy has less effect on the cost due to the fact that less time is required to recover from an unpredicted failure.

at the dip of the parabolic curve. Note that the cost/utility ratio is consistently higher for the passive vs. the proactive approach. We also show the effect of f_{cp} on the availability of the system: the change is in line with the formulation described in Equation 4. However, the interesting insight from this experiment is the direct correspondence between the lowest achievable cost/utility and the flattening of the availability curve.

Our final experiment studies the effect of prediction accuracy λ , on the achieved cost/utility ratio. It is intuitive that better prediction accuracy would lead to lower cost/utility for proactive availability-management, and this is clearly depicted in Figure 8. It is interesting to note the behavior of proactive availability-management with a lower f_{cp} value. When prediction accuracy is low, a small f_{cp} leads to very high recovery times with low net-utility during that period. However, if f_{cp} is modulated properly to handle failures, recovery time decreases and a far lower cost/utility is achieved. Meanwhile, the effect of prediction accuracy is less prominent when a higher value of f_{cp} is used, as the recovery times don't improve much, even with a correct prediction.

6 Related Work

Traditional Fault-Tolerance. Redundancy is probably the earliest form of fault-tolerance; the approach popularly known as the active replication approach is well-studied, and a thorough description appears in [4]. Log-based recovery is well-known in the database domain. Here, a failure is handled with an undo-redo log [3]. Fault-tolerance has also been studied in the context of transactions [20] and distributed systems [21]. Dynamically trading consistency for availability is proposed in [22] using a continuous consistency model. A number of factors distinguish our approach from

these traditional mechanism, the first and the foremost being its utility-awareness. Another distinction is our ability to use failure prediction to reduce the overhead of ensuring high-availability.

Failure Detection & Prediction. [7] focuses on the implementation of fault detection, and proposed a scalable fault detection/collection framework. More recently, researchers in the autonomic domain have used statistical monitoring techniques to detect failures in component-based Internet services [11,23]. MSET or multi-variate state estimation techniques [11] constitute an early warning system that enables failure prediction with low false alarm probability and has been successfully applied to the thermal control domain, and more recently, to software aging problems, including predicting memory leaks, data corruption, shared memory pool latching, etc. In [9], instrumentation data is correlated to system states using statistical induction techniques to identify system-level metrics that correlate with high-level performance states. In addition, these techniques are used to forecast service level objective violations, with prediction accuracy reported to be around 90%. Our system provides a framework in which several such failure detection and prediction techniques can be implemented to provide high-availability while imposing a low-overhead.

Fault-Tolerant Distributed Information Systems. Stars [21] presents a fault-tolerance manager for distributed application, using a distributed file manager which performs actions like message backups and checkpoints storage for user files. Its reliance on causal and atomic group multi-cast, however, demands additional solutions in the context of today's widely geographically distributed enterprise systems [24].

MTTR may be improved with solutions like Microreboot [25], which proposes a fast recovery technique for large systems. It is based on the observation that a significant fraction of software failures in large-scale Internet systems can be cured by rebooting. While rebooting can be expensive and cause nontrivial service disruption, microrebooting is a fine-grain technique for surgically recovering faulty application components, without disturbing the rest of the components of the application. Our work could benefit from such techniques.

GSpace [26] and replica management in Grids [27] studied dynamic data replication policy and modeling in distributed component-based systems when multiple replicas of data are desired, e.g., for global configuration data, or in a highly dynamic environment, to improve availability. For this kind of data replication management, efficient read-one write-all protocol [28] can be used when updates of the replicated data occur frequently.

IFLOW's techniques may be directly compared to the fault-tolerance offered in systems like Fault-Tolerant CORBA [29,30], Arjuna [31] and REL [32], which replicate selected application/service objects. Multiple replicas allow an object to continue to provide service even when one of its replicas fails. Passive replication is also provided. Here, the system records both the state of the currently executing member (primary member) and the entire sequence of method invocations. While CORBA focuses on the client-server model of communication, recent systems like Borealis [14] and SMILE [6] have focused on fault-tolerance for applications that process data streams. The former uses replication-based fault-recovery, and the authors propose to trade consistency for recovery time. The latter proposes the soft-checkpointing mechanism that can be used to implement a low-overhead passive replication scheme for fault tolerance. We

differ from such earlier work because of our explicit consideration of system utility for managing system availability, and because our system also provides a framework for incorporating failure prediction techniques.

Utility-Functions. The specific notions of utility used in this paper mirror the work presented in [8], which uses utility functions for autonomic data-centers. Autonomic self-optimization according to business objectives is studied in [33], and self-management of information flow applications in accordance with utility functions is studied in [5]. A preliminary discussion about availability-aware self-configuration in autonomic systems appears in [34]. Our middleware carefully integrates the ideas from the above systems and other domains to build a comprehensive framework for fault-tolerant information flows.

7 Conclusion

We have proposed techniques for managing the tradeoff between availability and cost in information flow middleware. First, a net-utility-based formulation of the benefits an enterprise derives from its information flows combines both performance and reliability attributes of such flows. The goal is not simply to attain high utility, but to reliably provide high utility to large-scale information flow applications. Second, since reliability techniques incur costs, thereby reducing utility, proactive methods for availability-management regulates resources used to guarantee availability and take into account the fact that system and application behaviors change over time. A specific example is a higher likelihood of failure in high load vs. low load conditions. Reliability costs, therefore, are reduced by exploiting knowledge about the current ‘perceived’ system stability. Additional cost savings result from the use of failure prediction methods. Third, the implementation presented in this paper can deal with both transient and non-transient failures, the latter relying on application-specific techniques for fault avoidance. Finally, utility-driven proactive availability-management techniques has been integrated into our infrastructure for large-scale information flows, where it is shown to impose low additional communication and processing overheads on information flows. Experimental results with IFLOW attained on Emulab [17] demonstrate the effectiveness of proactive fault tolerance in recovering from failures.

Future work will experiment with richer failure prediction techniques, and investigate specific enterprise environments. For instance, we will model the redundant data-centers mandated by government rules, and will consider the attainment of high availability and net-utility in information flows that cross multiple organizational boundaries.

References

1. IBM: IBM global services: Improving systems availability. (<http://www.cs.cmu.edu/~priya/hawht.pdf>)
2. Gavrilovska, A., Schwan, K., Oleson, V.: A practical approach for zero’ downtime in an operational information system. In: Proc. of ICDCS. (2002)

3. Gray, J., McJones, P.R., Blasgen, M.W., Lindsay, B.G., Lorie, R.A., Price, T.G., Putzolu, G.R., Traiger, I.L.: The recovery manager of the system R database manager. *ACM Comput. Surv.* **13**(2) (1981)
4. Randell, B., Lee, P., Treleaven, P.C.: Reliability issues in computing system design. *ACM Comput. Surv.* **10**(2) (1978)
5. Kumar, V., Cai, Z., Cooper, B.F., Eisenhauer, G., Schwan, K., Mansour, M., Seshasayee, B., Widener, P.: Implementing diverse messaging models with self-managing properties using iflow. In: *Proc. of ICAC*. (2006)
6. Strom, R.E.: Fault-tolerance in the smile stateful publish-subscribe system. In: *Proc. of the Int'l Workshop on Distributed Event-Based Systems*. (2004)
7. Stelling, P., Foster, I., Kesselman, C., Lee, C., Laszewski, G.V.: A fault detection service for wide area distributed computations. In: *Proc. of HPDC*. (1998)
8. Walsh, W.E., Tesauro, G., Kephart, J.O., Das, R.: Utility functions in autonomic systems. In: *Proc. of ICAC*. (2004)
9. Cohen, I., Goldszmidt, M., Kelly, T., Symons, J., Chase, J.: Correlating instrumentation data to system states: A building block for automated diagnosis and control. In: *Proc. of OSDI*. (2004)
10. Gross, K.C., Lu, W.: Early detection of signal and process anomalies in enterprise computing systems. In: *Proc. of IEEE International Conference on Machine Learning and Applications*. (2002)
11. Zavaljevski, N., Gross, K.C.: Uncertainty analysis for multivariate state estimation in mission-critical and safety-critical applications. In: *Proc. MARCON*. (2000)
12. Blough, D., Liu, P.: FIMD-MPI: A tool for injecting faults into mpi applications. In: *Proc. of IPDPS*. (2000)
13. Cai, Z., Kumar, V., Cooper, B.F., Eisenhauer, G., Schwan, K., Strom, R.: Utility-driven availability-management in enterprise-scale information flows. <http://www.cercs.gatech.edu/tech-reports/> (2006) Technical report.
14. Balazinska, M., Balakrishnan, H., Madden, S., Stonebraker, M.: Fault-tolerance in the borealis distributed stream processing system. In: *Proc. of the ACM SIGMOD international conference on Management of data*. (2005)
15. Mansour, M., Schwan, K.: I-rmi: Performance isolation in information flow applications. In: *Proc. of ACM/IFIP/IEEE Middleware*. (2005)
16. Schwan, K., et al.: Autoflow: Autonomic information flows for critical information systems. In: *Autonomic Computing: Concepts, Infrastructure, and Applications*, ed. Manish Parashar and Salim Hariri, CRC Press. (2006)
17. Lepreau, J., et. al.: The Utah network testbed. (<http://www.emulab.net/>)
18. Zegura, E.W., Calvert, K., Bhattacharjee, S.: How to model an internetwork. In: *Proc. of IEEE INFOCOM*. (1996)
19. VINT Project: The network simulator - ns-2. (<http://www.isi.edu/nsnam/ns/>)
20. Wu, H., Kemme, B.: Fault-tolerance for stateful application servers in the presence of advanced transactions patterns. In: *Proc. of SRDS*. (2005)
21. Sens, P., Folliot, B.: STAR A fault-tolerant system for distributed applications. *Software - Practice and Experience* **28**(10) (1998)
22. Yu, H., Vahdat, A.: The costs and limits of availability for replicated services. In: *Proc. of SOSP*. (2001)
23. Fox, A., Kiciman, E., Patterson, D.: Combining statistical monitoring and predictable recovery for self-management. In: *Proc. of SIGSOFT workshop on Self-managed systems*. (2004)
24. Cheriton, D., Skeen, D.: Understanding the limitations of causally and totally ordered communication. In: *Proc. of SOSP*. (1993)

25. Candea, G., Kawamoto, S., Fujiki, Y., Friedman, G., Fox, A.: Microreboot - a technique for cheap recovery. In: Proc. of OSDI. (2004)
26. Russello, G., Chaudron, M., van Steen, M.: Dynamically adapting tuple replication for high availability in a shared data space. In: Proc. Int'l Conf. on Coordination Models and Languages. (2005)
27. Schintke, F., Reinefeld, A.: Modeling replica availability in large data grids. *Journal of Grid Computing* (2003)
28. Rabinovich, M., Lazowska, E.: An efficient and highly available read-one write-all protocol for replicated data management. In: Proc. of the Int'l Conf. on Parallel and Distributed Information Systems. (1993)
29. Group, O.M.: Final adopted specification for Fault Tolerant CORBA. In: OMG Technical Committee Document ptc/00-04-04. (2000)
30. Moorsel, A., Yajnik, S.: Design of a resource manager for fault-tolerant corba. In: Proc. of the Workshop on Reliable Middleware. (1999)
31. Parrington, G.D., Shrivastava, S.K., Wheeler, S.M., Little, M.C.: The design and implementation of Arjuna. *USENIX Computing Systems* (1995)
32. Friese, T., Muller, J., Freisleben, B.: Self-healing execution of business processes based on a peer-to-peer service architecture. In: Proc. of ICAC. (2005)
33. Aiber, S., Gilat, D., Landau, A., Razinkov, N., Sela, A., Wasserkrug, S.: Autonomic self-optimization according to business objectives. In: Proc. of ICAC. (2004)
34. Chess, D.M., Kumar, V., Segal, A., Whalley, I.: Availability-aware self-configuration in autonomic systems. In: *Distributed Systems Operations and Management*. (2003)