Egida: An Extensible Toolkit For Low-overhead Fault-Tolerance

Sriram Rao* Lorenzo Alvisi* Harrick M. Vin†
Department of Computer Sciences
The University of Texas at Austin
Taylor Hall 2.124, Austin, Texas 78712-1188, USA.

Abstract

We discuss the design and implementation of Egida, an object-oriented toolkit designed to support transparent rollback-recovery. Egida exports a simple specification language that can be used to express arbitrary rollback recovery protocols. From this specification, Egida automatically synthesizes an implementation of the specified protocol by gluing together the appropriate objects from an available library of "building blocks". Egida is extensible and facilitates rapid implementation of rollback recovery protocols with minimal programming effort. We have integrated Egida with the MPICH implementation of the MPI standard. Existing MPI applications can take advantage of Egida without any modifications: fault-tolerance is achieved transparently—all that is needed is a simple re-link of the MPI application with Egida.

1 Introduction

Building reliable distributed application is complex. Application developers not only have to worry about the intricacies of their applications, but also have to understand the subtleties of distributed fault-tolerance. Over the past decade, toolkits such as Horus [22], and Transis [8] have addressed this issue by providing higher level primitives for implementing fault-tolerant broadcasting and group communication. These primitives are well-suited for implementing fault-tolerance protocols based on active replication, desirable for mission-critical applications that require the highest degree of availability and the capability to tolerate arbitrary failures.

As distributed computing becomes commonplace, however, a growing number of non mission-critical applications are emerging. For instance, parallel scientific applications that used to be run on supercomputers can be executed on a distributed cluster of servers. For these applications, fault-tolerance is highly desirable, but only if it can be provided with low-overhead in terms of dedicated resources and execution time. Log-based rollback recovery protocols—such as checkpointing and message logging—provide an attractive low-overhead solution for building non-critical distributed applications that can tolerate crash failures.

In this paper, we discuss the design and implementation of *Egida*, an object-oriented toolkit designed to support transparent rollback recovery for low-overhead fault-tolerance. Rather than providing monolithic implementations of a set of protocols, we build Egida around a library of objects that implement a set of functionalities that are at the core of all log-based rollback recovery protocols and define a grammar for configuring protocols from the library of objects. Our approach has several advantages. First, it promotes extensibility and flexibility by allowing multiple implementations of each of the core functionalities. Second, it facilitates rapid implementation of rollback recovery protocols with minimal programming effort by gluing together objects from the available library of building blocks. This enables application developers to experiment with different protocols and then use the one that closely matches the requirements of their application. Finally, Egida enables designers of fault-tolerance protocols to develop new rollback recovery protocols by combining different implementations of the core functionalities in novel ways.

Egida shares some of the design goals of other frameworks—such as OTEC [20] and RENEW [18]—for implementing rollback recovery protocols. Egida is unique in that it allows rollback recovery protocols to be configured from basic building blocks. This approach is similar to the one used in other systems—for instance, in Adaptive [23] and xkernel [13] for networking protocols, and in Horus [22] and Cactus [11] for distributed computing and group membership protocols.

We have integrated Egida with the MPICH implementation of the Message Passing Interface (MPI) standard [25]. This enables existing MPI applications to take advantage of Egida without any modifications. Conversely, the performance of rollback recovery protocols can now be evaluated using a large set of demanding applications.

This paper is organized as follows. In Section 2, we provide a brief overview of log-based rollback recovery protocols. In Section 3, we identify the set functionalities that are at the core of all log-based rollback recovery protocols, and present a grammar for synthesizing these protocols from its core components. Section 4 describes the architecture of Egida. Section 5 describes the implementation of Egida and its integration with the MPICH library [10]. Section 6 presents the related work, and finally, Section 7 summarizes our contributions.

*These authors supported in part by the National Science Foundation (CAREER award CCR-9734185, Research Infrastructure Award CDA-9734185, and DARPA/SPAWAR grant N66001-98-8911).

†This author supported in part by an AT&T Foundation Award, IBM Faculty Development Award, Intel, the National Science Foundation (CAREER award CCR-9624757, and Research Infrastructure Award CDA-9624082), Lucent Bell Laboratories, NASA, Mitsubishi Electric Research Laboratories (MERL), and Sun Microsystems Inc.

2 Flavors of Log-based Rollback Recovery

Log-based rollback recovery protocols—such as checkpointing and message logging—provide a low-overhead solution for building distributed applications that tolerate crash failures. These protocols come in several flavors. Checkpoints can be independent, coordinated, or induced by specific patterns of communication. Logging can be pessimistic, optimistic, or causal. Pessimistic protocols [3, 28] allow processes to communicate only from recoverable states. These protocols enforce this condition by synchronously logging to stable storage any information critical for recovery before letting processes communicate. Optimistic protocols (for instance, [7, 14, 27]) allow processes to communicate with other processes even from states that are not yet recoverable. These protocols guarantee that these states will eventually become recoverable, but only if no failures occur. Failures may indeed render some states permanently unrecoverable, forcing the rollback of any process that depends on such states. Causal protocols [2, 9] weaken the condition imposed by pessimistic protocols and allow the possibility that a state from which a process communicates may become unrecoverable because of a failure, but only if no correct process depends on that state. This is enforced by appending to all communication the information necessary to recover the state from which the communication originates. This information is replicated in the volatile memory of the processes that causally depend [15] on the originating state. Hence, causal protocols never roll back correct processes, but do not require synchronous writes to stable storage.

3 Deconstructing Log-Based Rollback-Recovery Protocols

The diversity of rollback-recovery protocols reflects the heterogeneity in the requirements of applications. For instance, applications that don't interact frequently with the external environment and can tolerate rollback of correct processes may use checkpointing or optimistic protocols. Pessimistic and causal protocols are better suited for applications in which communication with the environment is frequent and rollbacks are unacceptable; pessimistic protocols are preferred when fast and simple recovery is required. Applications that instead demand minimal overhead during failure-free executions would use causal or optimistic protocols. A closer look at these protocols, however, reveals that behind this diversity lies a simple event-driven structure that all these protocols share and that all protocols are interested in the same set of "relevant" events. In this section we identify such events and present a language that can be used to specify a protocol in terms of which actions it takes in response to each of these events.

3.1 Log-based Rollback-Recovery Protocols as Event-Driven Programs

There are five types of events that are relevant to all rollback recovery protocols. These are: (1) non-deterministic events, (2) dependency-generating events, (3) output-commit events, (4) checkpointing events, and (5) failure-detection events.

1. **Non-deterministic events**: A non-deterministic event is an event whose outcome may change for different executions of the same program. For example, a message deliver event is often non-deterministic since its outcome depends on many factors, including process scheduling, routing, and flow control. Thus, a recovering process may not produce the same run upon recovery even if the same set of messages are sent to it.

   If a non-deterministic event cannot be replayed during the recovery of a process, all correct processes whose state depends on that unrecoverable event must be rolled back. These processes are called *orphans*. The information necessary to reproduce the results of the non-deterministic events is called the event's *determinant*. Restoring the system to a consistent state while limiting the extent of rollbacks depends on the ability of the rollback-recovery protocol to replay the determinants of the non-deterministic events executed by the failed processes before crashing. Hence, determinants must be saved to stable storage so that they are available during recovery. In practice, different protocols choose different ways to log determinants and rely on different implementations of stable storage.

2. **Dependency-generating events**: These events can increase the number of processes that depend on the non-deterministic events executed by a process. For example, in message-passing applications, a message-send event is a dependency-generating event since it can create dependencies between non-deterministic events executed by the sender and the state of the receiver. If the sender fails and any of those non-deterministic events is unrecoverable, then the receiver becomes an orphan.

   To enable orphan detection and appropriate rollback, rollback recovery protocols typically append to application messages information that tracks these dependencies—logging protocols piggyback logical clocks [15] or vector clocks [17] and checkpointing protocols piggyback an index that records the number of checkpoints taken by a process. In addition, to speedup the replay of non-deterministic events, it may be desirable to log some information on executing a dependency-generating event. For instance, in sender-based message-logging protocols processes log the content of each message they send.

3. **Output-commit events**: These events can make the external environment depend on the non-deterministic events executed by a process.

   Since it is not reasonable to expect the external environment to roll back in response to a process crash, log-based rollback recovery protocols invoke *output commit* procedures that write synchronously to stable storage the determinants of all non-deterministic events that precede the communication with the external environment, thus making them recoverable. Depending on how determinants are logged, this procedure may require coordination among processes.

4. **Checkpointing events**: These events instruct the protocols to write to stable storage the state of one or more processes. These event can either be transparent to the application (e.g. a signal from a timer indicating that some pre-determined time has elapsed since the last checkpoint) or be explicitly driven by the application.

   Checkpoints can be coordinated, independent, or communication induced. Processes may checkpoint their state incrementally, and checkpoints can be saved to stable storage either synchronously or asynchronously.

5. **Failure-detection events**: These events are generated on detecting the failure of one or more processes.

In response to these events, faulty processes must be restarted, restored to a previously checkpointed state and rolled forward. In turn, correct orphan processes may need to be detected and rolled back. Different protocols implement different strategies to collect the determinants to be replayed during recovery and to detect and rollback orphans.

Implementing a specific protocol therefore amounts to selecting the set of actions performed in response to each relevant event. We now show how a simple language can be used to specify these choices. In Section 4 we show how protocols can be automatically synthesized starting from these specifications.

3.2 Specifying Design Choices in Rollback-Recovery Protocols

Figure 1 shows a simple language that can be used to specify rollback-recovery protocols. The language reflects the discussion of the previous section. A protocol is defined in terms the actions it takes in response to non-deterministic events, dependency generating events, output commit events, checkpointing events and failure-detection events. To define a protocol completely, it is necessary to instantiate a set of variables which specify, for instance, the set of non-deterministic events, the form of their determinant, the implementation of stable storage, etc. Figure 2 shows a set of representative instantiations of these variables used in several existing protocols.

Figure 3 illustrates how the language can be used to specify Damani and Garg's optimistic protocol [7], a hybrid protocol we have recently developed that combines causal logging with asynchronous receiver-based logging to stable storage in order to speed up crash recovery [21]. Our language, however, can be used for more than just specifying existing recovery protocols. Once deconstructing rollback-recovery protocols becomes possible, it is then easy to reassemble the fundamental components in novel ways to implement new protocols. An example is given in Figure 4(b), which shows a specification of a sender-based pessimistic message-logging protocol that tolerates f concurrent failures and implements stable storage by replicating data in the volatile memory of processes. In this protocol, which to our knowledge has not appeared in the literature, a process executing a non-deterministic event broadcasts the corresponding determinant to the other processes and does not send any application message until it has received at least f acknowledgments to its broadcast. This protocol is a variant of Strom, Bacon, and Yemini's sender-based pessimistic protocol [28], whose specification is shown in Figure 4(a).

4 The Architecture of Egida

Figure 1 defines the structure of log-based rollback recovery protocols, while the variables in Figure 2 identify the building blocks which when incorporated into the protocol structure yield different rollback recovery protocols. Egida instantiates these building blocks in a modular, object-oriented architecture. We begin this section by describing the functionality of each module and the interfaces it exports. We then explain how these modules are invoked in response to events and how they can be composed to synthesize rollback-recovery protocols.

4.1 Module Definitions and Interfaces

Egida defines the following modules:

EventHandler: Handles the five types of relevant events defined in Section 3. The methods exported by this module are event-dependent.

Determinant: Creates the determinants for each non-deterministic event. It exports a single method, *CreateDeterminant*.

HowToOutputCommit: Determines how to save the information necessary to recover the system to the state in which an output commit event is executed. It exports a single method, *OutputCommit*.

LogEventDeterminant: Saves the determinants of non-deterministic events to stable storage. It exports four methods: (1) *Log*, (2) *Retrieve*, (3) *Flush*, and (4) *GarbageCollect*.

LogEventInfo: Logs data associated with an event. It exports the same interfaces as the LogEventDeterminant module.

HowToLog: Determines how logged information is accessed. It exports two methods, *Read* and *Write*.

WhereToLog: Directs requests for accessing the logs to the appropriate form of storage. It exports the same methods as HowToLog.

StableStorage: Implements stable storage for determinants and data associated with an event. It also exports two methods, *Read* and *Write*.

VolatileStorage: Implements volatile storage for determinants and data associated with an event. It exports the same methods as StableStorage.

PiggybackLogging: Piggybacks onto application messages the information necessary to track dependencies among process states and, in the case of causal logging, appends the non-stable determinants logged by the message sender. It also processes, and if appropriate logs the piggybacked information. This module exports two methods: *GetPiggyback* and *ProcessPiggyback*.

PiggybackCheckpointing: Piggybacks the information necessary to track the dependencies among checkpoints. It also processes piggybacked information, and if appropriate triggers a checkpoint. It exports the same methods as Piggyback-Logging.

Checkpoint: Implements the checkpointing protocol. It exports two methods, *TakeCheckpoint* and *RestoreFromCheckpoint*.

HowToCheckpoint: Determines how the checkpoint information is accessed. It exports two methods, *Read* and *Write*.

CollectDeterminants: Collects determinants during recovery. It exports a single method, *RetrieveDeterminants*.

OrphanDetection: Implements the protocol for orphan detection. It exports a single method, *DetectOrphans*.

4.2 Module Invocation

Figure 5 shows the function call graph for the modules introduced in Section 4.1. In addition to the modules discussed in Section 4.1, it shows four other modules:

1. An API module that intercepts application-level calls to relevant events and invokes the appropriate handler exported by the EventHandler.

Protocol	:=	⟨non-det-event-stmt⟩* ⟨output-commit-event-stmt⟩*
		⟨dep-gen-event-stmt⟩* ⟨ckpt-stmt⟩$_{opt}$ ⟨recovery-stmt⟩$_{opt}$
⟨non-det-event-stmt⟩	:=	⟨event⟩ : *determinant* : ⟨determinant-structure⟩
		⟨*Log* ⟨event-info-list⟩ ⟨how-to-log⟩ *on* ⟨stable-storage⟩⟩$_{opt}$
⟨output-commit-event-stmt⟩	:=	⟨output-commit-proto⟩ *output commit on* ⟨event-list⟩
⟨event-list⟩	:=	⟨event⟩ | ⟨event⟩, ⟨event-list⟩
⟨dep-gen-event-stmt⟩	:=	⟨event⟩ : ⟨piggyback-stmt⟩ | ⟨logging-stmt⟩ |
		⟨piggyback-stmt⟩⟨logging-stmt⟩
⟨logging-stmt⟩	:=	*Log* ⟨event-info-list⟩ ⟨how-to-log⟩ *on* ⟨where-to-log⟩
⟨event-info-list⟩	:=	⟨event-info⟩ | ⟨event-info⟩, ⟨event-info-list⟩
⟨piggyback-stmt⟩	:=	*Piggyback* ⟨pb-value-list⟩
⟨pb-value-list⟩	:=	⟨pb-value⟩ | ⟨pb-value⟩, ⟨pb-value-list⟩
⟨ckpt-stmt⟩	:=	⟨ckpt-proto⟩ *checkpoint* ⟨how-to-ckpt⟩ *on* ⟨stable-storage⟩
		Implementation : ⟨ckpt-impl⟩
⟨recovery-stmt⟩	:=	⟨recovery-get-det-stmt⟩$_{opt}$
		⟨orphan-detection-stmt⟩$_{opt}$ ⟨rollback-stmt⟩$_{opt}$
⟨recovery-get-det-stmt⟩	:=	*Protocol to retrieve determinants* : ⟨recovery-get-det-proto⟩
⟨orphan-detection-stmt⟩	:=	*Protocol to detect orphans* : ⟨orphan-detection-proto⟩
⟨rollback-stmt⟩	:=	*On rollback* : ⟨rb-logging-stmt⟩ | ⟨rb-ckpt-stmt⟩
⟨rb-logging-stmt⟩	:=	*Log* ⟨event-info-list⟩ *on* ⟨stable-storage⟩
⟨rb-ckpt-stmt⟩	:=	*Take independent checkpoint on* ⟨stable-storage⟩

Figure 1 : A grammar for specifying rollback-recovery protocols. In the grammar, we use the following notational convention: ⟨x⟩ means that x is a production rule, ⟨y⟩ means that y is a variable that has to be mapped to a specific value, ⟨z⟩$_{opt}$ means that ⟨z⟩ is an optional statement, and *a* means that *a* is a keyword

2. A Timer module that enables Egida to specify the intervals at which timer interrupts need to be generated to perform periodic actions, such as checkpointing and flushing of volatile logs to stable storage.
3. A FailureDetector module that implements a protocol for detecting process crashes and triggers the appropriate recovery mechanisms through the EventHandler.
4. A Network module that invokes a method of EventHandler to process incoming messages.

To illustrate the interactions among Egida's module, we present two examples that show how this dependency graph is traversed in response to specific events.

Example 1 *Handling a Send event for a causal logging protocol (*see [2]*):*
The API module intercepts the Send event and invokes the corresponding handler in EventHandler. The handler in this example performs three tasks: (1) it saves in a volatile log kept by sender the content of the message being sent; (2) it retrieves the set of non-stable determinants from the log kept by the server; (3) it piggybacks these determinants on the application message and sends the message to the destination.

1. To save the content of the message in the sender's volatile log, the handler invokes the *Log* method of LogEventInfo. *Log* in turn invokes the *Write* method of HowToLog, which invokes the *Write* of WhereToLog, which finally invokes the *Write* method exported by VolatileStorage.
2. To obtain the determinants to piggyback, the handler invokes the *GetPiggyback* method of Piggyback-Logging. This method identifies the non-stable determinants and invokes the *Retrieve* method of LogEvent-Determinant. *Retrieve* in turn invokes the *Read* method of HowToLog, which invokes the *Read* of WhereToLog, which finally invokes the *Read* of VolatileStorage and, if necessary, of StableStorage. These methods return the requested determinants which are then passed up through the call chain back to *GetPiggyback*.
3. The event handler then piggybacks the determinants to the application message and invokes the send function in the transport layer. ■

Example 2 *Handling a Send event to the external environment for the optimistic protocol of Damani and Garg (*see [7]*):*
The API module intercepts the send event and invokes the corresponding handler in EventHandler. The handler in this example invokes the output commit procedure prior to sending the message. The output commit procedure in turn performs two tasks: (1) it flushes to stable storage the determinants of all the non-deterministic events executed by the process prior to the send to the external environment, thereby making them recoverable; and (2) it initiates a coordinated output commit operation to ensure that all the

⟨event⟩	:=	send | receive | read | write
⟨determinant-structure⟩	:=	{source, sesn, dest, desn}
⟨output-commit-proto⟩	:=	independent | co-ordinated
⟨event-info⟩	:=	determinant | message
⟨how-to-log⟩	:=	synchronously | asynchronously
⟨where-to-log⟩	:=	⟨stable-storage⟩ | ⟨volatile-storage⟩
⟨volatile-storage⟩	:=	local disk | volatile memory of self
⟨stable-storage⟩	:=	local disk | NFS disk | volatile memory of processes
⟨pb-value⟩	:=	determinants⟨, ⟨dependency matrix | stability matrix | stability vector⟩⟩$_{opt}$ | vector clock | logical clock | checkpoint counter
⟨ckpt-proto⟩	:=	independent | co-ordinated | communication-induced
⟨how-to-ckpt⟩	:=	synchronously | asynchronously
⟨ckpt-impl⟩	:=	full | incremental
⟨recovery-get-det-proto⟩	:=	centralized | distributed
⟨orphan-detection-proto⟩	:=	broadcast logical clock | exchange vector clock

Figure 2 : Representative instantiations for the variables defined in the grammar shown in Figure 1.

/* comment: ⟨non-det-event-stmt⟩ */
receive:
determinant : {source, sesn, dest, desn}
Log determinant, data asynchronously on NFS disk

{comment: ⟨output-commmit-event-stmt⟩}
Co-ordinated output commit on write

/* comment: ⟨dep-gen-event-stmt⟩ */
send:
Piggyback vector clock

/* comment: ⟨recovery-stmt⟩ */
Protocol to detect orphans : broadcast logical clock
On rollback: Log determinant, message on NFS disk

(a)

/* comment: ⟨non-det-event-stmt⟩ */
receive:
determinant : {source, sesn, dest, desn}
Log determinant, data asynchronously on NFS disk

/* comment: ⟨dep-gen-event-stmt⟩ */
send:
Piggyback determinants
Log message synchronously on volatile memory of self

/* comment: ⟨recovery-stmt⟩ */
Protocol to retrieve determinants : centralized

(b)

Figure 3 : Example specifications of existing protocols: (a) Damani and Garg's optimistic message logging protocol and (b) hybrid causal message logging protocol

/* comment: ⟨non-det-event-stmt⟩ */
receive:
determinant : {source, sesn, dest, desn}
Log determinant synchronously on NFS disk

/*comment: ⟨dep-gen-event-stmt⟩} */
send:
Log message synchronously on volatile memory of self

/* comment: ⟨ckpt-stmt⟩} */
Independent checkpoint asynchronously on NFS disk
Implementation : incremental

(a)

/* comment: ⟨non-det-event-stmt⟩ */
receive:
determinant : {source, sesn, dest, desn}
Log determinant synchronously on volatile memory of processes

/* comment: ⟨dep-gen-event-stmt⟩ */
send:
Log message synchronously on volatile memory of self

/* comment: ⟨ckpt-stmt⟩ */
Independent checkpoint asynchronously on NFS disk
Implementation : incremental

(b)

Figure 4 : Specification of (a) Strom, Bacon, and Yemini's sender-based pessimistic message logging protocol and (b) a novel sender-based pessimistic protocol

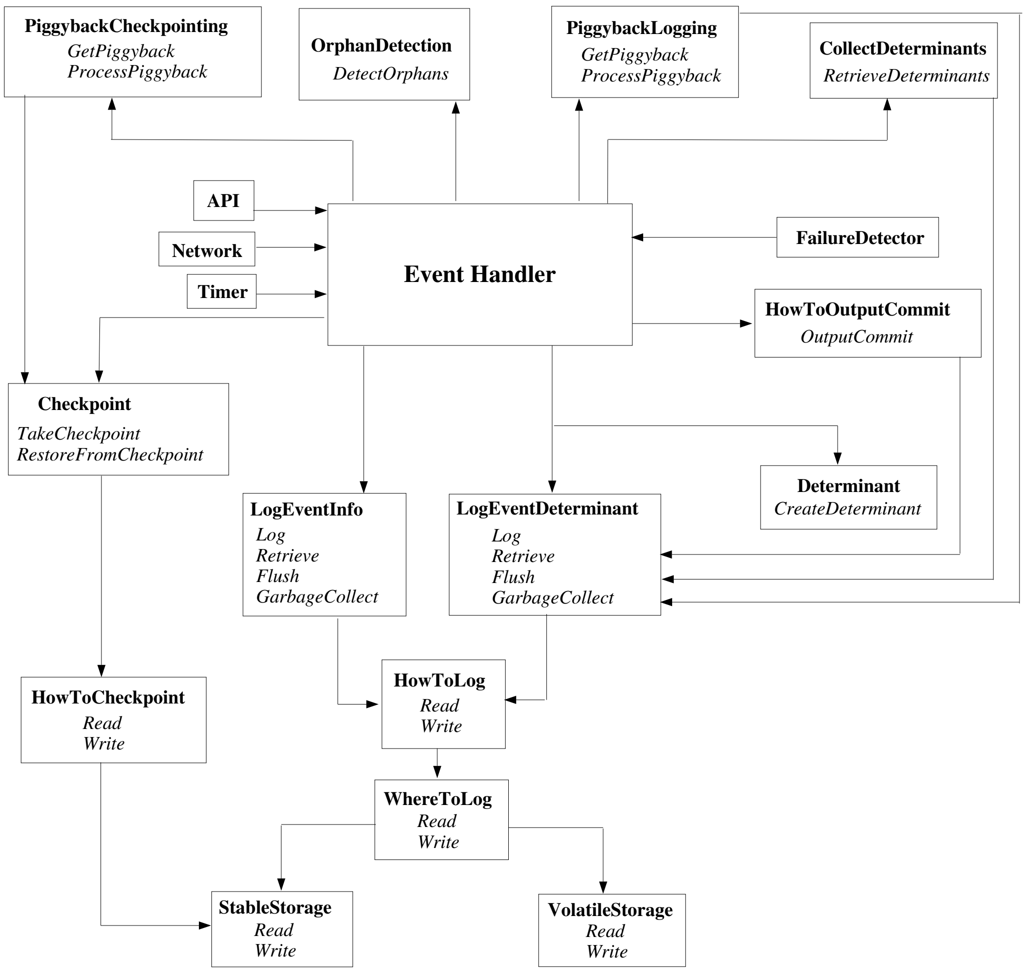

Figure 5 : The architecture of Egida

non-deterministic events that causally precede the send are recoverable.

1. To perform output commit, the handler invokes the *OutputCommit* method of HowToOutputCommit. The *OutputCommit* methods performs the following two tasks:
   (a) To perform a local output commit, *OutputCommit* invokes the *Flush* method of LogEventDeterminant. *Flush* then invokes the *Write* method of HowToLog, which invokes the *Write* of WhereToLog, which finally invokes the *Write* method exported by StableStorage.
   (b) To perform a co-ordinated output commit, *OutputCommit* broadcasts a message requesting an output commit. On receiving this message, the Network module invokes the appropriate handler in EventHandler, which in turn invokes the *OutputCommit* method of HowToOutputCommit[1]. The *OutputCommit* method then invokes the same sequence of methods as in (a).

[1]Note that the *OutputCommit* method exported by the instance of the HowToOutputCommit module invoked by the handler for processing network messages performs only a local output commit.

2. The handler sends the application message to the external environment. ■

4.3 Synthesizing Protocols through Module Composition

Egida allows the co-existence of multiple implementations for each of the modules introduced in Section 4.1. Hence, to synthesize a protocol, one must select a specific implementation of each module. Egida maintains a binding between the values for the variables on the left hand side of Figure 2 and their corresponding implementations. Therefore, synthesizing a protocol requires processing the specification along with this binding information to initialize the modules to their appropriate implementations. It is easy to extend Egida to recognize new implementations of modules; as more implementations of a module functionality become available, the specification language can be extended by simply (1) adding a new value to the right hand side of the corresponding variable in Figure 2, and (2) registering with Egida the binding between the new value and its corresponding implementation.

5 Implementation and Experience

Egida contains multiple implementations for each module introduced in Section 4.1. Each module is defined as a C++ class. The classes we have implemented so far correspond to the values on the right hand side of Figure 2.

To make Egida available to a large number of applications, we have integrated Egida with MPICH [10], a freely available, portable implementation of the MPI (Message Passing Interface) standard 1.1. MPICH has a two-layer architecture: the upper layer exports MPI's application programming interface (API), while the lower layer contains platform-specific message-passing libraries. For our testbed environment consisting of a network of workstations connected by an Ethernet, MPICH's lower layer is implemented using the p4 library [4].

To integrate Egida with MPICH, we replaced in the MPICH's upper layer all the send and receive calls to p4 with corresponding calls to Egida's API; modules in Egida in turn invoke the message passing operations of p4. We had to make two modifications to p4. First, we enhanced p4 to enable transfer of data from non-contiguous buffers. This was necessary to avoid a potential data copy that may occur when the information piggybacked by Egida has to be prepended to an MPICH message. Second, we modified p4 to handle socket errors that occur whenever processes fail and to allow a recovering process to establish socket connections with the surviving processes. For detecting failures and automatically re-starting crashed processes, in Egida, we have implemented a watchdog [12] for each application process. Our implementation provides transparent fault-tolerance to MPI applications—only a simple re-link of the application with Egida and our modified MPICH library is necessary.

We have used Egida to study the performance of rollback recovery protocols with a set of NAS [6] benchmarks. The results of our experimental evaluation are described elsewhere [1, 21].

6 Related Work

The literature contains several systems that have been designed to provide transparent fault-tolerance using rollback-recovery. They include libckpt [19], Fail-safe PVM [16], MIST [5], Co-check [26], Manetho [9] and libft [12]. As opposed to Egida, however, these systems are designed to support only a specific checkpointing or message-logging protocol.

Egida's emphasis on code reusability and extensibility is also present in the recently-introduced OTEC simulator [20]. OTEC's object-oriented design provides some degree of code reusability and can be used to compare the performance of different checkpointing protocols, but only through simulations (for example, varying message rates and checkpoint sizes). The run-time system for rollback-recovery that is closest to Egida in terms of extensibility is RENEW [18], a system that provides a framework for implementing and evaluating different checkpointing protocols. Egida is unique in allowing implementations of protocols to be automatically synthesized from reusable components. The specification language exported by Egida identifies key properties of rollback recovery protocols, simplifying their understanding, while hiding the tedium of combining the various properties to implement a complete protocol.

Egida's architecture, is similar to the one used in other systems—for instance, in Adaptive [23] and xkernel [13] for networking protocols, in Horus [22] and Cactus [11] for distributed computing and group membership protocols, and in COMERA [29] and Quarterware [24] for communications middleware. However, to our knowledge, Egida is the first application of this approach to rollback recovery protocols.

7 Concluding Remarks

We have presented the design and implementation of *Egida*, an object-oriented toolkit designed to support transparent rollback recovery. Egida exports a simple specification language that can be used to express arbitrary rollback recovery protocols. From this specification, Egida authomatically synthesizes an implementation of the specified protocol by glueing together the appropriate objects from an available library of "building blocks".

We have implemented a prototype of Egida and integrated it with the MPICH implementation of the Message Passing Interface (MPI) standard. Through Egida, MPI applications can be made fault-tolerant without any modifications. Conversely, Egida allows to study the robustness and performance of different rollback-recovery protocols with respect to a large set of demanding applications.

Our early experience with the prototype suggests that Egida is extensible, flexible, and facilitates rapid implementations of rollback-recovery protocols with minimal programming effort. We expect that the design of Egida will evolve as we gain more experience specifying and implementing different rollback recovery protocols.

References

[1] L. Alvisi, E. N. Elnozahy, S. Rao, S. A. Husain, and A. De Mel. An Analysis of Communication-induced Checkpointing. In *Proceedings of the IEEE Fault-Tolerant Computing Symposium (FTCS-29)*, Madison, WI, June 1999.

[2] L. Alvisi, B. Hoppe, and K. Marzullo. Nonblocking and Orphan-Free Message Logging Protocols. In *Proceedings of the 23rd Fault-Tolerant Computing Symposium*, pages 145–154, June 1993.

[3] A. Borg, J. Baumbach, and S. Glazer. A message system supporting fault tolerance. In *Proceedings of the Symposium on Operating Systems Principles*, pages 90–99. ACM SIGOPS, October 1983.

[4] R. Butler and E. Lusk. Monitors, Message, and Clusters: the p4 Parallel Programming System.

[5] J. Casas, D. Clark, P. Galbiati, R. Konuru, S. Otto, R. Prouty, and J. Walpole. MIST: PVM with transparent migration and checkpointing. In *3rd Annual PVM Users' Group Meeting*, Pittsburgh, PA, May 1995.

[6] NASA Ames Research Center. NAS Parallel Benchmarks. http://science.nas.nasa.gov/Software/NPB/, 1997.

[7] O. P. Damani and V. K. Garg. How to Recover Efficiently and Asynchronously when Optimism Fails. In *Proceedings of the 16th International Conference on Distributed Computing Systems*, pages 108–115, 1996.

[8] D. Dolev and D. Malki. The Transis Approach to High Availability Cluster Communication. *Communications of the ACM*, 39(4), April 1996.

[9] E. N. Elnozahy and W. Zwaenepoel. Manetho: Transparent rollback-recovery with low overhead, limited rollback and fast output commit. *IEEE Transactions on Computers*, 41(5):526–531, May 1992.

[10] William D. Gropp and Ewing Lusk. *User's Guide for* `mpich`*, a Portable Implementation of MPI*. Mathematics and Computer Science Division, Argonne National Laboratory, 1996. ANL-96/6.

[11] M. A. Hiltunen and R. D. Schlichting. A Configurable Membership Service. *IEEE Transactions on Computers*, 47(5):573—586, May 1998.

[12] Y. Huang and C. Kintala. Software Implemented Fault Tolerance: Technologies and Experience. In *Proceedings of the IEEE Fault-Tolerant Computing Symposium*, pages 2–9, 1993.

[13] N. C. Hutchinson and L. L. Peterson. The x-Kernel: An architecture for implementing network protocols. *IEEE Transactions on Software Engineering*, 17(1):64–76, January 1991.

[14] D. B. Johnson and W. Zwaenepoel. Recovery in Distributed Systems Using Optimistic Message Logging and Checkpointing. *Journal of Algorithms*, 11:462–491, 1990.

[15] L. Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *Communications of the ACM*, 21(7):558–565, July 1978.

[16] J. Leon, A. L. Ficher, and P. Steenkiste. Fail-safe PVM: A portable package for distributed programming with transparent recovery. Technical Report CMU-CS-93-124, School of Computer Science, Carnegie Mellon University, February 1993.

[17] F. Mattern. Virtual Time and Global States of Distributed Systems. In M. Cosnard et. al., editor, *Parallel and Distributed Algorithms*, pages 215–226. Elsevir Science Publishers B. V., 1989.

[18] N. Neves and W. K. Fuchs. RENEW: A Tool for Fast and Efficient Implementation of Checkpoint Protocols. In *Proceedings of the 28th IEEE Fault-Tolerant Computing Symposium (FTCS)*, Munich, Germany, June 1998.

[19] J. S. Plank, M. Beck, G. Kingsley, and K. Li. Libckpt:Transparent checkpointing under Unix. In *Proceedings of the USENIX Technical Conference*, pages 213–224, January 1995.

[20] B. Ramamurthy, S. J. Upadhyaya, and R. K. Iyer. An Object-Oriented Testbed for the Evaluation of Checkpointing and Recovery Systems. In *Proceedings of the 27th IEEE Fault-Tolerant Computing Symposium*, pages 194–203, June 1997.

[21] S. Rao, L. Alvisi, and H. Vin. The Cost of Recovery in Message Logging Protocols. Technical Report 99-12, University of Texas at Austin, March 1999.

[22] R. Van Renesse, T. Hickey, and K. Birman. Design and Performance of Horus: A Lightweight Group Communications System. Technical Report TR94-1442, Cornell University Computer Science Department, August 1994.

[23] D. Schmidt, D. Box, and T. Suda. Adaptive: A Dynamically Assembled Protocol Transformation, Integration, and Evaluation Environment. *Concurrency: Practice and Experience*, 5(4):269—286, June 1993.

[24] A. Singhai, A. Sane, and R. Campbell. Quarterware for Middleware. In *Proceedings of the International Conference on Distributed Computing and Systems (ICDCS'98)*, pages 192—201, Amsterdam, Holland, May 1998.

[25] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI: The Complete Reference*. Scientific and Engineering Computation Series. The MIT Press, Cambridge, MA, 1996.

[26] G. Stellner. CoCheck: Checkpointing and process migration for MPI. In *Proceedings of the International Parallel Processing Symposium*, pages 526–531, April 1996.

[27] R. B. Strom and S. Yemeni. Optimistic recovery in distributed systems. *ACM Transactions on Computer Systems*, 3(3):204–226, April 1985.

[28] R. E. Strom, D. F. Bacon, and S. A. Yemini. Volatile Logging in n-Fault-Tolerant Distributed Systems. In *Proceedings of the Eighteenth Annual International Symposium on Fault-Tolerant Computing*, pages 44–49, 1988.

[29] Y. M. Wang and W. J. Lee. COMERA: COM Extensible Remoting Architecture. In *Proceedings of the 4th Conference on Object Oriented Technologies and Systems (COOTS)*, April 1998.