Specification of an Exception Handling System for a Replicated Agent Environment

Christophe Dony
LIRMM, UMR 5506
CNRS and Univ. Montpellier 2
Montpellier, France
Christophe.Dony@lirmm.fr

Chouki Tibermacine
LIRMM, UMR 5506
CNRS and Univ. Montpellier 2
Montpellier, France
Chouki.Tibermacine@lirmm.fr

Christelle Urtado
LGI2P, Ecole des Mines d'Alès
Nîmes, France
Christelle.Urtado@ema.fr

Sylvain Vauttier
LGI2P, Ecole des Mines d'Alès
Nîmes, France
Sylvain.Vauttier@ema.fr

ABSTRACT

Exception handling and replication are two mechanisms that increase the reliability of applications. Exception handling helps programmers to control situations in which the normal flow of a program execution cannot continue. Replication handles system failures. Exceptions handling and replication do not apply to the same situations and are two complementary mechanisms to increase the reliability of applications. Moreover, each technique can benefit from the other: exception handling capabilities can rely on replication mechanisms while replication can be further secured by using exceptions. The paper proposes a specification of an execution history oriented exception handling system for an agent language and middleware providing replication. The paper proposes an original signaling algorithms taking into account replicated agents and a rationale of how exception handling and replication mechanisms can combine to increase the capability of programmers to achieve reliable agent-based applications.

1. INTRODUCTION

Exception handling and replication are two mechanisms (algorithms and architectures) dedicated to reliability and fault-tolerance that we wish to associate. Replication handles failures whereas exceptions enable programmers to dynamically handle those situations that prevent software from running normally.

An agent replication system [14, 8] is able to replace an agent (provided he has been replicated) thats fails by one of its (active or passive) replicas. This is transparent to software users and does not require any additional code from programmers. A failure is generally detected when an agent does not answer any more to messages since a given amount of time, either because network connections are lost or because the machine on which the agent ran is switched off. Active replication systems include algorithms capable of identifying the most critical agents to automatically create replicas of them. All messages sent to an agent in active replication are transmitted to all of its replicas, which process the same message in parallel. In passive replication, there is only one replica (called *leader*) which processes messages and sends periodically state updates to the other replicas. In semi-active replication, the replica selected as a leader, by the replication system, receives the messages from other client agents and forwards them to the other replicas.

Exceptions are situations in which the standard control flow of a program execution cannot continue. An exception is not a failure because it is a kind of answer from the agent. It indicates what the agent is unable to continue its task the standard way but that he is still alive.

Exception handling and replication do not apply to the same situations and are different in nature : replication is preventive and exception handling is curative. However, both mechanisms are obviously very complementary. In the context of the FACOMA[1] project that studies adaptive reliability of large scaled multi-agent applications we are studying an exception handling system capable of working on top of a replication system. The goal of this global project is then threefold:

- The combination of replication mechanisms and exception handling in general is a new an interesting challenge for software reliability.
- Exception handling can improve replication. Firstly, the replication system implementation can be made more robust by internally using exceptions. Secondly it can improve replication strategies, for example, with passive replication, the signaling of a system exception by the leader can become a new case where a replica can be activated in place of it.

[1] http://www-src.lip6.fr/homepages/facoma.officiel/

- Replication can also improve in various ways exception handling by providing active copies of the computation state.

The objective of this paper is to present our study on the first of the three above points: how an exception handling system can be combined with a replication mechanism to increase the reliability of agent-based applications. The bases of the study are our SaGE exception handling system dedicated to agents [21], components [22] and active objects [6] and the DIMAX replicated agent system [7, 8, 14] abstracted in section 2.

Our study lists and discusses issues, related to the adaptation of an computation histoty oriented exception handling system on top of such replicated agent systems. Here is a panel of the main ones:

- How to transparently exploit replication for exception handling.
- What should happen when an exception is raised by an agent that has one or more, active or passive, replicas?
- When and where should the replication system be able to take control when an agent signals an exception.
- Which decisions can be taken within a handler defined at the replication system level?
- How to distinguish an exception that can be interpreted as a failure from the point of view of the replication system and that cant thus entail the election of a new replica, from an exception which can be interpreted as a correct answer for the service caller? For example, signaling the "division-by-zero" exception is the normal answer from the "divide" function if its second argument is zero. In this case it is useless to activate another replica, because it will compute the same answer.
- Do standard resolution mechanisms used in distributed exception handling to concert exceptions apply to synthesize the results computed by different replicas of the same agent?

The remainder of the paper is structured as follows. Section 2 sets the context of this work, describing the agent model and the targeted replicated agent system. Section 3 abstracts the requirements for exception handling in a multi-agent world and provides the agent programmer-directed API of our X-SaGE EHS. Section 4 describes how exception handlers are searched for in a replicated agent system while Sect. 5 discusses how system-defined handlers can be defined in the replication system to integrate exception handling in the replication manager. Section 6 concludes with a short discussion on the benefits of our approach against state of the art exception handling systems and open perspectives to this work.

2. CONTEXT OF THE STUDY : OVERVIEW OF A REPLICATED MULTI-AGENT SYSTEM

The context of this work is the programming of reactive, collaborating agents that are deployed over a middleware which handles agent replication. The concepts exposed in this section are derived from the DIMAX software that combines the DIMA multi-agent system [8] and the DARX fault-tolerant middleware [14]. Initial names and principles are generalized here and adapted to the agent interaction scheme we studied in our previous work [6], namely peer-to-peer service exchanges.

2.1 The agent model

The agent concept has many concretisations. This section abstracts the reactive and collaborative agent model we have worked with.

An agent is a computation entity that executes in its own thread. This provides the agent with the properties of being active and autonomous. The behavior of a reactive agent consists of two parts: a control behavior which defines how the agent makes decisions to act, depending on its internal state and the state of its environment; several elemental behaviors that represent the different actions the agent knows to do. Figure 1 shows an abstract of the *BasicCommunicatingAgent* class, the base class in our context, used to implement reactive agents. The control behavior of the agent is defined by the *live* method. It implements a loop that is executed while the agent is alive. Each iteration of this control loop calls the *step* method. This method implements the decision mechanism that enables the agent to choose, step by step, the action it executes. A control behavior represents the existence of the agent and is executed in a separate thread, provided by the execution platform as an instance of the *AgentEngine* class. This enables the execution of the agents on top of different platforms, which can adapt their specific execution model to the management of an agent as an *AgentEngine* subclass. The other behaviors of the agents are represented as methods of the agent class. Some of these behaviors are executed upon the reception of a request from another agent. These behaviors are called services.

Agents interact by exchanging asynchronous messages. Each agent holds a message box and a communication interface respectively to send and receive messages (*cf.* the *MessageBox* and *CommunicationComponent* classes on Fig. 1). The communication interface is provided by the execution platform and is responsible for the delivery of the messages. As described above, specific asynchronous messaging mechanisms can be adapted to the agent model as a *CommunicationComponent* subclass. Each agent bears a unique identifier. These identifiers are used as logical references to agents, for instance as senders or recipients of messages. The execution platform uses name directories to convert these abstract agent identifiers to effective references, in order to deliver the messages to the agents.

A specific semantics is associated to messages in order to set up a request / response interaction protocol between agents. This protocol describes peer-to-peer collaborations in which a client agent asks a server agent for a service thanks to a request message. Conforming to a contract-based approach of software development, whenever a server agent accepts a request, it commits to send back a result, either standard or exceptional, to the client agent in a response message. Response messages are correlated with request messages.

When no response is received within a time defined by the client agent, a timeout exception is signaled.

As an illustration (*cf.* Fig. 2), we use the canonical *Travel Agency* example in which a *Client* can send to a *Broker* a reservation message in order to request a bid for a travel. The contacted broker sends in turn a bid request to several travel providers and collects their responses. Then, the *Broker* selects the best offer and requests the *Client* and the selected *Provider* to contract.

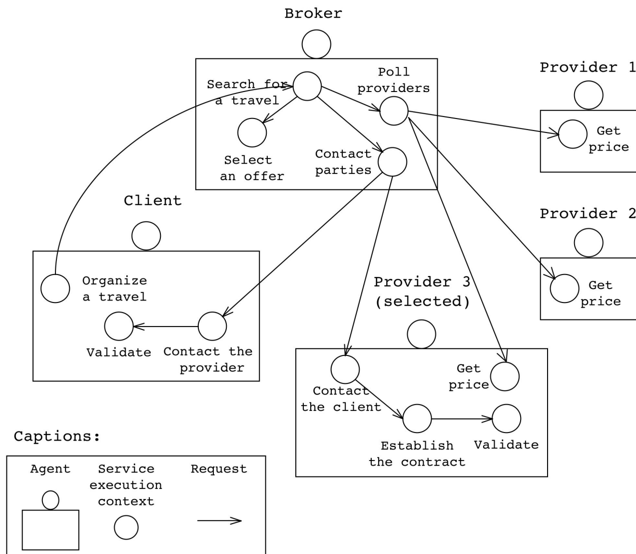

Figure 2: Execution resulting from a request to a travel agency

2.2 The Replication System

Agents are executed on a middleware which provides a fault-tolerant execution context thanks to a replication mechanism [14]. The execution context consists of a set of distributed replication servers which manage the execution of tasks (*ReplicatedTask* class of Fig. 1). Every task belongs to a replication group (*ReplicationGroup* class) that identifies the set of tasks which are replicas of a same logical task. Thus, all the tasks within a replication group have the same behavior (they actually are instances of the same task class). Moreover, the middleware maintains the consistency within the replication group so that all the tasks are in the same state after each computation.

Logical tasks are identified by logical names that are used to send them messages. The replication middleware is in charge of the location and the delivery of messages to the corresponding replicas. More precisely, the message is delivered first to the leader of the corresponding replication group. The leader is a replica which has the specific role to control the replication group. For this purpose, the leader holds a replication manager which monitors the messages sent to or by the replicas in the replication group. The replication manager maintains status information about the replicas and executes group management operations (creation, destruction of replicas, etc.). The replication manager distinguishes two kinds of replicas. Active replicas effectively execute treatments. The leader is necessarily an active replica. The leader forwards the messages sent to the task to the other active replicas so that they do the same computation and reach the same new state. Passive replicas only perform state updates. When the leader completes the computation, its new state is serialized and sent to all the passive replicas. After their update, the passive replicas are in the same state as the leader (and supposedly as the other active replicas).

Conversely, all the messages sent by the replicas are filtered by the replication middleware. Only the messages sent by the leader are actually delivered to other tasks. This way, replication is transparent. Whatever the number of replicas of a task, a unique message is sent to invoke a computation and a unique message is received as a response. In fact, active and passive replicas are to be considered as failover copies of the leader. Only the leader interacts with other tasks as long as it runs correctly.

The number and the type of replicas is determined by the replication policy, regarding the criticality of the task and the availability of resources (memory, CPU). In case of failures, new replicas can be dynamically created in order to maintain the redundancy required to provide an expected level of fault-tolerance. The type of replica (active, passive) can be changed to adapt resource consumption to criticality and risk. When the leader fails, its responsibility is transferred to another replica. When a passive replica is chosen to become the leader, its status is changed to active. The state of the task (meaning the state of all the replicas of the corresponding replication group) is thus rolled back to the state of the new leader (which represents the previous consistent state of the task, backed up in a passive replica). If no replica still exists, the task has been finally destroyed by the failure.

Every agent executes inside a task (*cf.* Fig. 1). As such, agents can be replicated by the middleware and benefit from this fault-tolerance mechanism. The following sections explain how exception handling is combined with replication to provide a more reliable and robust multi-agent system.

3. CONTROL STRUCTURES FOR EXCEPTION HANDLING WITH REPLICATED AGENTS

This section motivates and presents the first part of our proposal : the X-SaGE control structures for exception handling designed for agent programmers. These control structure only slightly differ from the SaGE system ones. Indeed, they are programmer-directed and replication mechanisms do not interfere in any programmer-directed capability as replication must be transparent to the agent programmer. Handling replication will intervene in the implementation of these control structures in sections 4 and 5.

3.1 Requirements for an Agent Programmer-directed EHS

The key requirements of the X-SaGE exception handling system, extended from [21, 6], are :

- to enforce agent encapsulation,

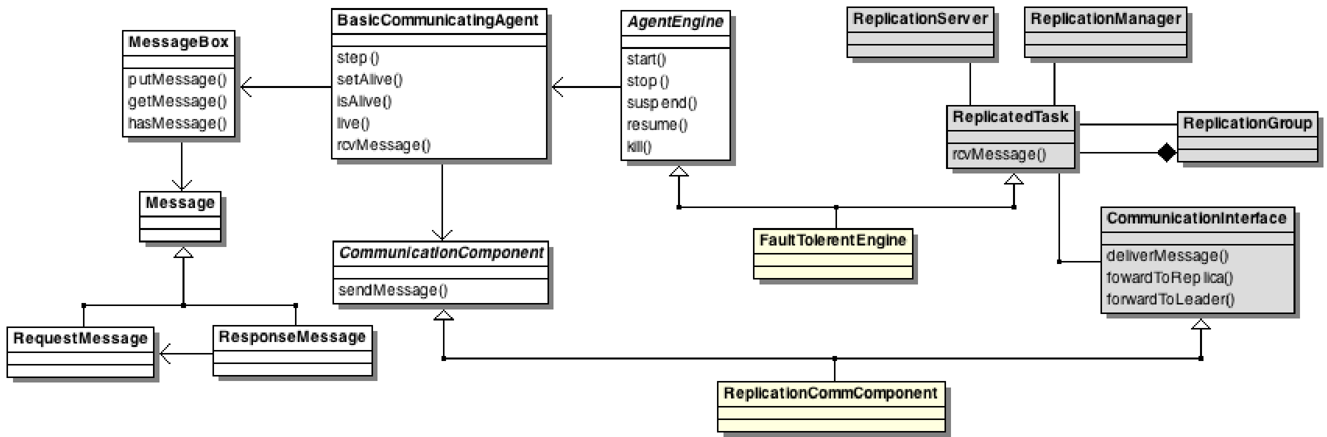

Figure 1: Excerpt from the agent model

- to provide a representation for collaborative concurrent activities [19] so that they can be coordinated and controlled [18],
- to look for handlers in the history of computation, instead of delegating exception handling to specialized agents, and to execute handlers in their lexical definition context; we call this caller contextualization [4] for handler definition and execution). When encapsulation and decoupling are enforced, non lexical handlers do not have access to the execution contexts where the exceptions are signaled (the agents which execute the faulty services). They can only use generic management operations (such as ervice or agent termination) to cope with the signaled exception.
- to handle concurrent exceptions with resolution functions [10],
- and, to support asynchronous signaling and handler search and thus maintain agent reactivity,

Our specification comes in four steps indicating: (1) to which program code units exception handlers can be attached, (2) how exceptions can be signaled, (3) what can be written within the code of exception handlers to put the system back into a coherent state and, (4) in which order handlers are searched for. The following two subsections are dedicated to items 1 to 3. Item number 4 involves interfacing with the agent replications mechanisms. It is discussed in Sect. 4 and 5.

3.2 Signaling Exceptions and Attaching Handlers

Figure 3 shows the java code of an X-SaGE agent that defines services and various exception handlers. X-SaGE takes advantage of the java annotations to make exception handling for agents as seamless as possible. It shows examples of service definitions (annotated by *@service*): lines 6–7 define the *pollProviders* service and lines 17–30 the *contactParties* service. It also illustrates (lines 20–28) how a message can be sent by (a service of) a client agent to request a server agent to provide him with some (sub-) service.

Signaling exceptions is done by the means of a classical *signal* primitive (*cf.* Fig. 3, line 11). Signaling is possible anywhere in the code. This includes the possibility of signaling an exception from within handlers.

Steps of the request / response interaction pattern highlight the role of three key entities: the request, the service and the active agent. They are the three program code units to which exception handlers can be attached:

- Exception handlers can be attached to **requests**. Such handlers can, for example, specify two distinct reactions to the occurrence of two identical exceptions raised by two invocations of the same service. Lines 23–27 of Fig. 3 show how a handler can be attached to a specific request.
- Exception handlers can be attached to **services**. Such handlers treat exceptions that are raised, directly or indirectly, by some service's execution. If the service is complex, the handler has to be able to deal with concurrent exceptions, to compose with partial results or to ignore partial failures. Lines 10–14 of Fig. 3 shows the code of two handlers attached to a same service (*@serviceHandler* annotation). Note that the *servicename* attribute of the annotation allows to identify the service the handler protects.
- Finally, exception handlers can be attached to **agents**. Such handlers act as if they were repeatedly attached to all of the agent's services. They can be used, for example, to uniformly maintain in the consistency of the agent's private data. Lines 3–4 of Fig. 3 shows how such handlers can be associated to agents using the*@agentHandler* annotation.

These capabilities are powerful enough to encompass most cases the agent programmer will be confronted to and simple enough to be easy to learn and use.

3.3 Defining Exception Handlers and Resolution Functions

Exception handlers are classically defined by the set of exception types they can catch and by their code body (as

```
( 1)  public class Broker extends X_SaGEAgent
( 2)    {
( 3)    // handler associated to the Broker agent
( 4)    @agentHandler public void handle (GlobalNetworkException exc)  { ... }
( 5)
( 6)    // service provided by the Broker agent
( 7)    @service public void pollProviders ()  { ... }
( 8)
( 9)    // handler associated to the PollProviders service
(10)    @serviceHandler(servicename=pollProviders) public void handle (BadParameterException exc)
(11)       { signal (new NoAirportInDestinationException ( ... ); }
(12)
(13)    // handler associated to the PollProviders service
(14)    @serviceHandler(servicename=pollProviders) public void handle (NoProviderException exc) { ... }
(15)
(16)    // service provided by the Broker agent
(17)    @service public void contactParties ()
(18)       {
(19)       ...
(20)       sendMessage (new RequestMessage (aServerAgent, "ContactSelectedProvider")
(21)          {
(22)          // handler associated to a request
(23)          @requestHandler public void handle (OffLineException exc)
(24)             {
(25)             wait(120);
(26)             retry();
(27)             }
(28)          });
(29)       ...
(30)       }
(31)
(32)    // resolution function associated to the pollProviders service
(33)    @serviceResolutionFunction(servicename=pollProviders) public TooManyProvidersException concert ()
(34)       {
(35)       int failed = 0;
(36)       for (int i=0; j<subServicesInfo.size(); i++)
(37)          if ((ServiceInfo) (subServicesInfo.elementAt(i)).getRaisedException() != null)  failed++;
(38)       if (failed > 0.3*subServicesInfo.size())  return new TooManyProvidersException(numberOfProviders);
(39)       return null;
(40)       }
(41)   }
```

Figure 3: Service, handler and resolution function definitions in X-SaGE using annotations

illustrated by Fig. 3, lines 23–27, for example). There are three main actions a handler can classically have:

- A handler can restore whatever should be, to put back data into a consistent state, and can **return** a value that becomes the value of the expression the handler is associated to. In case of a message sending expression (standard or broadcast), the value returned by the handler is the value of the expression. In case of a handler attached to a service, the value becomes the result of the service execution. In case of a handler attached to an agent, the value becomes the result of the execution of the service that raised the exception.
- A handler can **signal** a new exception (generally of a higher conceptual level) or **re-signal** the original one. This behavior is illustrated on Fig. 3 line 11. Of course, handlers cannot protect themselves from the exceptions they signal.
- A handler can **retry** the execution of the program unit it is attached to. Retry amounts to entirely re-execute the program unit it is attached to, generally after having modified the local environment, but in the same historical context. This possibility is illustrated on Fig. 3, line 27. In case of handlers attached to agents, retrying means re-executing the service that signaled the exception.

X-SaGE provides exception resolution support integrated to the handler search. It enables resolution functions to be defined at places where concurrent activities are launched and have to be co-ordinated (*i.e.,* at the service level). There is no need for a resolution function either at the request level, because requests are atomic, or at the agent level because all semantically sound activities of agents, that need to be co-ordinated, are accessible via services. The default behavior of the resolution function associated to a service is, once all recipients have replied, to aggregate all the exceptions that occurred into a concerted one. Another possible behavior is to transmit one reply as soon as it arrives without waiting for others. Such a use of resolution for concerted exception slightly differs from the original work of [10], a resolution function is executed each time an exception handler is searched for at the service level, this makes our system reactive, because our resolution function evaluates the situation each time an exception is signaled. Of course, a programmer can define his own exception resolution function using the *@serviceResolutionFunction* annotation as shown in the

example of Fig. 3, lines 33-40.

4. HANDER SEARCH IN A REPLICATED AGENT SYSTEM

Handler search requires that a tree of **service execution contexts** be monitored. Each node represents a service execution context and records the identities of the service being executed and the agent that owns the current service (*cf.* Fig. 4). Each node can optionally have a parent node that links to the calling context of the current service. In this parent node, the request that triggered the current service is recorded. Links between nodes (callee to caller links) are used to look for handlers. Figure 2 shows the service execution context tree that results from the services executed in the travel agency example.

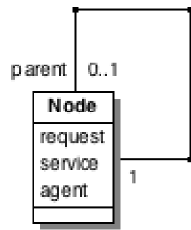

Figure 4: Node of a service execution context tree

The Figure 5 then shows the organigram that synthesizes the different steps of the handler search process.

If an exception is raised within an agent service, then the execution of the service is suspended and handler search is launched. The handler search process decomposes into four steps. First, a handler for the exception is searched in the list of handlers and resolution functions associated to the service. Concerning the resolution function, three cases are possible:

- the exception is critical for the service. The resolution function returns the exception object and the handler search process carries on.

- the resolution function evaluates that the exception is under-critical and that nothing more should be done yet. The exception is logged, the resolution function returns null and the handler search process stops. The collective activity is not affected. The only service that is terminated is the defective sub-service.

- the resolution function evaluates that the exception is under-critical but that there is a need to signal something, for example because too many under-critical exceptions have been logged. The resolution function returns a special exception that reflects the situation and the handler search carries on.

If no handler has been found at the service level, one is searched at the agent owner of the service. If a suitable handler is found, it is executed and its execution terminates the execution of the service. The agent is of course still alive. Along with the execution of the handler, all pending services called by the current one, if any, are terminated.

If no handler has been found at the agent level, and if the agent is replicated, the control is given to its replication manager. Each replication manager has a handler that traps all exceptions and which acts in fact as a resolution function the goal of which is to coordinate the answers given by the various replicas of the agent. The behavior of this replication manager handler is described in section 5).

If the replication manager does not want to handle the exception or if it propagates it, the search proceeds in the calling context. First, the caller service is suspended and the search for a handler is initiated in the calling service's context. The list of handlers associated to the request which initiated the called service is searched first. If a handler is found, it is executed and the search stops. Then, the search proceeds by starting again at step 1, searching the list of handlers associated to the current service, then, those associated to the owner agent of the current service, etc.

The same four steps are repeated until an adequate handler is found and executed, following callee to caller links in the service execution context graph. If no handler has been found when the root of the service execution context tree is reached, a default top-level handler is executed.

5. HANDLING EXCEPTIONS AT THE REPLICATION MANAGER LEVEL

With replication, we face the following global issues: (1) how to trap an exception raised by the leading replica of an agent before the exception is propagated to the caller? (2) What to do when it has been trapped? Solving issue 1 is done by invoking the replication manager during handler search as explained in the preceding section. We have added in each replication managers a resolution function and an associated handler that trap all exceptions. Concerning Issue 2, the replication manager handler will either, as described in the following section, put the system back into a coherent state, signal a new exception to the request caller of propagate one of those it has trapped. In this latter case, the handler search will continue as explained in section 4.

5.1 Typology of exceptions

The first global question for the replication manager handler of an agent when one of its replicas raises an exception is to know whether the same exception will also be raised by the others. Which exception is replica-specific (examples of this include exceptions raised when some resources specific to a given replica are unavailable) and which ones are replica-independent (an example is bad parameter in the request sent to the agent (and thus to all its replicas), leading for example to a division by zero)? In the worst case, it could be considered that all exceptions are replica-specific. It would mean that when one replica signals an exception, we could systematically elect a new one to retry the same computation. This would significantly slow down program executions.

We thus have conceived our algorithms on the base of a classification of exceptions. Goodenough's seminal paper has proposed a classification in *domain*, *range* and *monitoring* exceptions that highlights the reason why an exception is raised. It however appears that we have no way to know

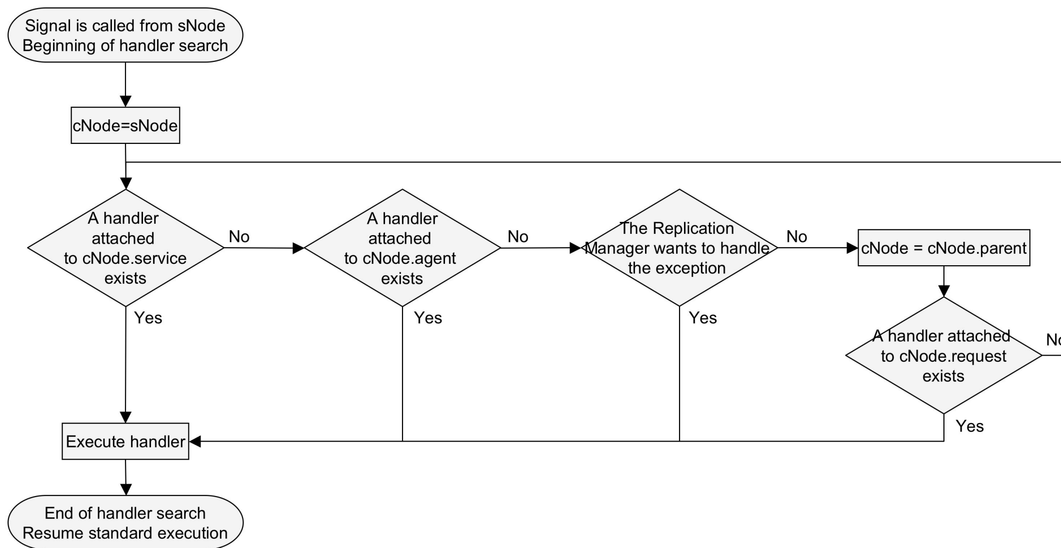

Figure 5: Organigram for handler search

whether a *range* exception (for example) is replica specific or independent. A classification in terms of *Error* (serious problem, should not be handled) and *Exception* (business problem, can be handled) as in Java, inherited from the *Flavors* system, highlights the exception gravity but cannot again be applied to our problem.

A more appropriate classification in our case is a classical semantic one in term of *business* (also called *domain* or *applicative* exceptions) and *resource* or *system* exceptions. Resource or system exceptions are raised by the computing environment and are likely to reflect a specific communication or resource lack problem. System exception can be considered as replica-specific. Business exceptions are direct consequences of a programmer's code. Under the hypothesis that all the replicas of an agent have the same deterministic behavior, exceptions identified as business exception can be considered as replica-independent : they will be raised by all replicas of a given agent. The question of knowing how to detect at run-time whether an exception is a system or business is left open at this point of the study.

Beyond this classification, the strategies of the exception handler of the replication manager also takes into account the composition of the replication group. Three strategies are described in the following sections.

5.2 Controlling one active replica with multiple passive ones

So-called *passive-replication strategy* uses one active replica (the leader) and a set of passive ones.

When the active replica (leader) raises a business exception, it is immediately propagated to the client agent. The leader is considered to be in a coherent state as a business exception is part of the behavior designed by the programmer of the agent. Moreover, this exception should be raised by any replica executing the same request message. It is useless to discard the current leader and to activate another replica to retry and execute the request.

When the active replica (leader) raises a system exception, it is handled as a failure of the leader. The leader, which is left in an undefined, potentially inconsistent and harmful state, is destroyed. One of the passive replicas is activated and it is asked to to retry the interpretation of the requested message. If this new leader raises a system exception too, another passive replica is used until their number runs out. If the last replica fails, the exception (or a m is finally signaled to the client agent to warn about the service failure.

An optimization can be introduced to manage more efficiently system exceptions. If the same kind of system exception is raised by all replicas, a global problem is suspected (for example, the unavailability of a shared resource or a faulty programming). After a given number of attempts, the replication manager stops retrying the execution of the request to prevents a useless consumption of replicas. This threshold is a configuration parameter of the replication policy and its value is defined by the administrator of the replicated agent environment.

5.3 Controlling a set of active replicas

When a business-related exception is raised by a replica, it is immediately propagated to the client agent since all the others one are intended to raise the same exception. The replication manager handler does not stop the execution of the request in the other replicas but filters the exceptions they raise (in order not to send the same exception to the client agent several times). This also enables to determine when the execution of the request is achieved for all the replicas, to verify that they have raised the same exception and are thus in the same new consistent state.

When a system-related exception is raised by a replica, it is recorded by the resolution function of the replication manager, until all the active replicas have achieved the execution of the request and have sent a response. Meanwhile, if a nor-

mal response is computed by one replica, it is immediately forwarded to the client agent. The other subsequent normal responses are discarded by the replication manager. When all the replicas have sent a response, the replication manager destroys all the faulty replicas. If the leader is destroyed, a new leader is chosen among the remaining replicas. If all the replicas are destroyed, an exception is then signaled to the client agent to warn it about the service failure.

5.4 Controlling a mix of active and passive replicas

In the case where active and passive replicas are mixed (the most general case), the handler first behaves as if there were only active replicas (the second case above). If system exceptions are successively signaled by all active replicas and active replicas are destroyed, it is possible to activate some of the passive replicas, whether to augment the number of active replicas for the next request or to retry and execute the current request. It is to be noticed that the creation of new replicas is not part of the behavior of replication managers (which are specific to each replication group) because it must be arbitrated between the different replication groups, according to the criticality of the tasks and the availability of computing resources. Replica creation is thus managed by another replication middleware mechanism based on the observation of task termination.

6. CONCLUSION AND RELATED WORKS

In this paper, we have proposed the specification of the X-SaGe exception handling system able to work with replicated agents. X-SaGe firstly offers to agent programmers an exception handling system that works transparently with replicated agents. It secondly offers to replication implementors the capacity to implement new replication strategies based on the signaling of programmers-code related exception by replicas. We have described such possible strategies for passive and active replicas. It can finally offer to replication implementors the capacity to internally control internal exceptions raised by the replication algotithms, as proposed in [13]. This last point concerns replication implementors and is not developed in this paper. We have proposed an original and light programming API, using java annotations to define handlers and resolution functions. We propose a handler search and a handler invocation algorithms that take into account the service execution history and, when possible, work asynchronously to improve agent reactivity. The implementation of our specification in the context of the Dimax[7] software (the DIMA agent language on top of the Darx replication system) in in progress.

Concerning related works, there are few studies on mixing exception handling and replication and as far as we know, no other in the agent context. [13] has proposed internal strategies to enhance a majority voting algorithm for replicated processes thanks to the handling of sequencing exceptions or hardware failures via an exception handling system. What is done for hardware failure has partially influenced our strategies for replication in presence of exceptions. The system is also able to report exceptions to callers. [11] has proposed an initial study to combine distributed object-oriented programming and N-version programming and [20] proposes an ADA framework for the same purpose. Our handler at the replication manager level globally plays the same role for exceptions than the "exception adjudicator" of [20]. One main difference with our proposal is that the control of the coherence of versions is more complex with N-version since versions are programmed by different programmers whereas replication simply duplicate agents that run the same code in the same environment on different processors. For this reason we have been able to propose different strategies to return responses to clients as soon as possible without waiting all responses from replicas.
Concerning our resulting exception handling system, as far as replication is transparent to programmers, it can be compared with existing ones. Various proposals address the issues related to exception handling for active object integrating asynchronous communication [2, 12, 3, 9, 15, 16, 17, 1]. Our solution is original in that it combines the following features : handling of request / response interactions between agents, handling of agent replicas, encapsulation and reactivity, ability to write context-dependent dynamic scope handlers (caller-contextualization), ability to coordinate and control group of active agents collaborating to a common task, ability to configure the exception propagation policy by defining *exception resolution functions* at the service level.

This specification and implementation are first steps and we wish to develop many points in future works. In a first step, the interactions between the replication mechanism and the exception handling system have to be further analyzed for system-level and application-level exceptions to rafine replication strategies. We also plan to enhance Dimax capabilities using the exception handling system as a "last chance" mechanism to signal failures when the Darx replication system has failed. This could be used in two distinct situations: (1) to signal that the last remaining replica of an agent died (failed) in order to allow to trigger less efficient modes the programmer might have coded at the agent level, and (2) to signal the death (failure) of an agent that was not considered to be critical. This would allow to provide a recover strategy if the estimated criticality was wrong. This would imply that the Darx component dedicated to failure detection also detect those specific situations and raise an exception. This could also lead to enhance SaGe capabilities using the meta-information on agents computed by the Dimax platform. For example, an agent's computed criticality could be used to tune concerted exception resolution. We could also propose a vulnerability measure that would use information on agents such as its reliability (using an exception history) its criticality and its current number of replicas. We also look forward to use replication as a support to give the core implementation of an exception handling system that supports a resumption policy. Indeed, even if handler search is stack destructive, as in most systems, a replica of an agent could restart the computation where it has been stopped in the original one.

Acknowledgments

Authors wish to thank the French Research Agency[2] (Anr) that supported part of this work through the Facoma project of the SetIn 2006 programme. They also want to thank Alexander Romanovsky for his help on the bibliography and

[2] http://www.agence-nationale-recherche.fr/.

all colleagues from the FACOMA project[3], Jean-Pierre Briot, Zahia Guessoum, Olivier Marin and Jean-François Perrot for fruitful and inspiring discussions.

7. REFERENCES

[1] N. Cacho, K. Damasceno, A. F. Garcia, A. Romanovsky, and C. J. P. de Lucena. Exception handling in context-aware agent systems: A case study. In R. Choren, A. F. Garcia, H. Giese, H. fung Leung, C. J. P. de Lucena, and A. B. Romanovsky, editors, *SELMAS*, volume 4408 of *Lecture Notes in Computer Science*, pages 57–76. Springer, 2006.

[2] R. Campbell and B. Randell. Error recovery in asynchronous systems. *IEEE Transactions on Software Engineering*, SE-12 number 8(8):811–826, August 1986.

[3] R. Carlsson, B. Gustavsson, and P. Nyblom. Erlang: Exception handling revisited. In *Proceedings of the Third ACM SIGPLAN Erlang Workshop*, September 2004.

[4] C. Dony. An object-oriented exception handling system for an object-oriented language. In *Proceedings ECOOP '88 (European Conference on Object-Oriented Programming)*, pages 146–161. Springer-Verlag, 1988.

[5] C. Dony, J. L. Knudsen, A. B. Romanovsky, and A. Tripathi, editors. *Advanced Topics in Exception Handling Techniques*, volume 4119 of *Lecture Notes in Computer Science*. Springer, september 2006.

[6] C. Dony, C. Urtado, and S. Vauttier. Exception handling and asynchronous active objects: Issues and proposal. In Dony et al. [5], chapter 5, pages 81–101.

[7] N. Faci, Z. Guessoum, and O. Marin. Dimax: A fault-tolerant multi-agent platform. In B. Dunin-Keplicz, A. Omicini, and J. A. Padget, editors, *EUMAS*, volume 223 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2006.

[8] Z. Guessoum, N. Faci, and J.-P. Briot. Adaptive replication of large-scale multi-agent systems: towards a fault-tolerant multi-agent platform. In *Proceedings of the 2006 international workshop on Software engineering for large-scale multi-agent systems (SELMAS'06). Alessandro Garcia, Ricardo Choren, Carlos Lucena, Paolo Giorgini, Tom Holvoet, and Alexander Romanovsky, editors.* Number 3914 Lecture Notes in Computer Science, Springer Verlag, 2006.

[9] A. Iliasov and A. Romanovsky. Exception handling in coordination-based mobile environments. In *Proceedings of 29th IEEE International Computer Software and Applications Conference (COMPSAC 2005)*, pages 341–350, Edinburgh, Scotland, UK, July 2005.

[10] V. Issarny. An exception handling model for parallel programming and its verification. In *Proceedings of the ACM SIGSOFT'91 Conference on Software for Critical Systems*, pages 92–100, New Orleans, Louisianna, USA, December 1991.

[11] V. Issarny. An exception handling mechanism for parallel object-oriented programming: Towards the design of reusable, and robust distributed software. *Journal of Object-Oriented Programming 6(6)*, pages 29–39, October 1993.

[12] A. W. Keen and R. A. Olsson. Exception handling during asynchronous method invocation. In B. Monien and R. Feldmann, editors, *Proceedings of Euro-Par 2002 Parallel Processing*, Lecture Notes in Computer Science, pages 656–660. Springer-Verlag, August 2002.

[13] L. Mancini and S. Shrivastava. Exception handling in replicated systems with voting. In *Digestofpapers, Fault Tol. Comp. Symp-16. Vienna*, pages 384–389, 1986.

[14] O. Marin, M. Bertier, and P. Sens. Darx—a framework for the fault-tolerant support of agent software. In *ISSRE '03: Proceedings of the 14th International Symposium on Software Reliability Engineering*, page 406. IEEE Computer Society, 2003.

[15] R. Miller and A. Tripathi. The guardian model and primitives for exception handling in distributed systems. *IEEE Trans. Software Eng.*, 30(12):1008–1022, 2004.

[16] S. Mostinckx, J. Dedecker, E. G. Boix, T. V. Cutsem, and W. D. Meuter. Ambient-oriented exception handling. In Dony et al. [5], pages 141–160.

[17] E. Platon, N. Sabouret, and S. Honiden. A definition of exceptions in agent-oriented computing. In G. M. P. O'Hare, A. Ricci, M. J. O'Grady, and O. Dikenelli, editors, *ESAW*, volume 4457 of *Lecture Notes in Computer Science*, pages 161–174. Springer, 2006.

[18] B. Randell, A. Romanovsky, C. Rubira-Calsavara, R. Stroud, Z. Wu, and J. Xu. From recovery blocks to concurrent atomic actions. In *Predictably Dependable Computing Systems*, ESPRIT Basic Research Series, pages 87–101, 1995.

[19] A. Romanovksy and J. Kienzle. *Advances in Exception Handling Techniques:*, volume 2022 of *Lecture Notes in Computer Science*, chapter Action-Oriented Exception Handling in Cooperative and Competitive Concurrent Object-Oriented Systems, pages 147–164. Springer, 2001.

[20] A. B. Romanovsky. An exception handling framework for n-version programming in object-oriented systems. In *3rd International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2000), 15-17 March 2000, Newport Beach, CA, USA*, pages 226–233. IEEE Computer Society, 2000.

[21] F. Souchon, C. Dony, C. Urtado, and S. Vauttier. Improving exception handling in multi-agent systems. In C. J. P. de Lucena, A. F. Garcia, A. B. Romanovsky, J. Castro, and P. S. C. Alencar, editors, *Software engineering for multi-agent systems II, Research issues and practical applications*, number 2940 in Lecture Notes in Computer Science, pages 167–188. Springer, February 2004.

[22] F. Souchon, C. Urtado, S. Vauttier, and C. Dony. Exception handling in component-based systems: a first study. In *Exception Handling in Object Oriented Systems: towards Emerging Application Areas and New Programming Paradigms Workshop (at ECOOP'03 international conference) proceedings*, pages 84–91, 2003.

[3] `http://www-src.lip6.fr/homepages/facoma.officiel/`