

Distributed Exception Handling : ideas, lessons and issues with recent exception handling systems

Aurélien Campéas*, Christophe Dony*, Christelle Urtado**, Sylvain Vauttier**

5th October 2004

* LIRMM, 161 rue ADA, 34392 Montpellier

** LGI2P, Ecole des Mines d'Alès, site EERIE, Parc scientifique G. Besse, 30035 Nîmes

Abstract

Exception handling is an important feature of the tool-set that enables the building of fault-tolerant concurrent and distributed programs. While transactional distributed systems have been studied for a long time, distributed exception handling techniques are only now evolving towards maturity, especially within asynchronous multi agents or component based systems. In this paper, we review two recent proposals for distributed exceptions handling systems (DEHS), namely SaGE and Guardian, in the light of what the Erlang programming language brings to the table : native constructs for concurrency, distributedness and exception handling across processes. We expand on the merits and possible drawbacks or limits of these systems. We advocate the need to introduce the notion of *resumption*, an often downplayed feature of EHSs, to modern day DEHS, in order to address the problem of multi agent systems facing the «real world».

Keywords : distributed exception handling, distribution, concurrency, fault tolerance, resumption, restart, SaGE, Guardian, Erlang.

1 Introduction

Exception handling mechanism have been a feature of programming languages for about thirty years now. The complete feature set of an exception handling system (EHS) is well known and stable (chap.1.1) in the context of non-concurrent, non-distributed programs [3, 21, 29, 17]. Yet within concurrent and distributed systems, it remains an active research topic [23, 14, 25, 7, 8, 15, 16].

Distributed computing has been addressed in length, with relation to programming languages. On one hand, distributed systems have been built with languages without any support for it except system call level primitives for process instantiation and message passing (like C or C++). On the other hand, actor's model programming language have been devised to make distribution and concurrency high-level and first-class concepts available to the programmer (recent examples : *Erlang* [1], *JoCaml* [27], *Oz* [20]). Thus, it is no surprise that in those recent languages we find building blocks, and even first stage implementations of distributed EHS. We will begin to review the most mature of them, Erlang.

There seems to be a middle ground between those poles, around Java. Java has already its own primitives to help building multi-threaded and distributed (RMI) programs. But in Java, there is no language support to manage exception handling across thread groups and process boundaries.

Thus, in this paper, we also review two proposals around this limitation of the Java platform : the *SaGE* [5] and *Guardian* [19, 28] systems, which specify and implement a DEHS for asynchronous component based and multi agents [4] systems written in Java. We believe that Erlang, SaGE and Guardian present important insights for the future makers of a DEHS, while they expand on earlier work

As a programming language feature, an EHS is mostly orthogonal to other language features, but can be disastrous - essentially useless, albeit tempting to the programmer - in the presence of low-level memory (de)allocation primitives, like in C++, for instance. It has been shown that an EHS pairs well with an automatic memory management subsystem. It remains to be seen upon which assumptions it may behave well in distributed and concurrent settings, for some properties of such systems may. Through this study, we uncover some of them.

We will see overall that the choice of concurrency and distributed constructs have an impact on the structure of a distributed EHS, and that

until now experimental outcomes have promoted a fail-fast philosophy hindering the development of a full-fledged distributed EHS. Finally, we advocate a DEHS supporting the notion of resumption of a computation after an exception has been handled, where «fail-fast» shows its limits.

1.1 The EHS feature set

Before we delve in DEHS, let us remember the commonly accepted features of an EHS [6]. Firstly, EHSs have been devised in the context of functional and imperative programming languages, which all carry the notion of a program execution as a sequential series of jumps from function (or procedure, or method) to function ; we could say from current continuation to next normal continuation.

An EHS introduces the notion of an exceptional set of continuations to be executed in the presence of an exceptional situation arising during a computation.

The act of *signaling* is what occurs when a program encounters an exceptional situation. It can be triggered automatically, by perusal of some built-in exception-throwing primitive, or explicitly via some signalment primitive. Signaling implies the building of an "exception" object carrying run-time information about the exceptional situation, and finding the innermost exception handler matching its type.

Exception handlers (EH) allow to define, at any arbitrary point in a program, a new set of exceptional continuations, that adds up to, or overrides part of, the current set ; each exceptional continuation is usually selected using the type of the exception object that is built at signalment time ; an EH has typically dynamic extent and global scope¹.

Most programming language since the past twenty five years, from Ada [10] to Java [26] or Haskell [13], have had this kind of exception handling capabilities, which is also referred to as the «termination» model. Indeed, the research of handlers is done «destructively», that is the call stack is unwound up to the handler to be selected.

The possibility to resume activity after signalment exists in the «*resumption*» model, where the signalment is merely a function call. The ability to resume depends on the type of the exception being signaled (an *error* implying termination, and a *continuable error* allowing resumption).

¹Some systems offer statically (class) or lexically scoped handlers.

Thus, a handler may have the choice between termination (then the stack will be unwound up to the handler) or a list of possible resuming points, also described in the Common Lisp terminology as *restarts* [22]. Let us define the notion of restart.

Restarts allow to define a set of continuations whose goal is to provide an EH with reparation and restarting capabilities. Restarts, being closed over some lexical environment between the signaling place and the innermost last EH, can be told to modify their environment and do some non-local jump to this environment in order to resume the computation, if abortion is not wanted. Restarts have also dynamic extent and global scope.

Resumption is supported in a few programming languages (like Smalltalk [3, 2] and various Lisps). To our knowledge, only Common Lisp and its poster child Dylan have restarts ; also Scheme may have them potentially through the usage of first class continuations and some macrology. We will argue in the last chapter in favor of the resumption model with restarts in the context of distributed and concurrent programs.

1.2 Implications of this feature set

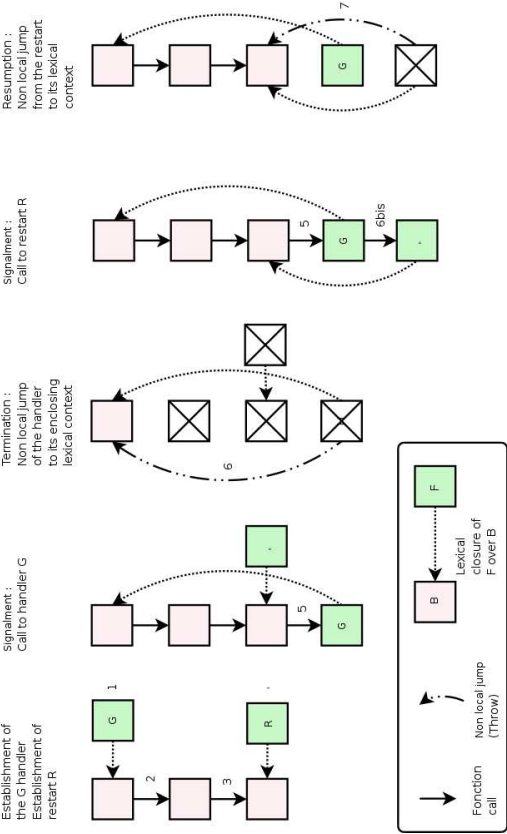
The main assumptions under which EHS have been built are :

- functional or imperative programming languages (where the notion of continuation makes sense),
- automatic memory management (garbage-collected languages),
- one single control flow,
- synchronous calls (or jumps) between parts of a program,
- determinism of execution and total-orderliness of any program's trace².

Distribution and concurrency raise new specific and difficult issues for fault-tolerance that have up to now remained incompletely studied.

²This one has to be taken with a grain of salt. Haskell has a simple exception handling mechanism that works in spite of the lazy evaluation strategy ; however it needed a pinch of cleverness to get it working.

Figure 1: Resumption model with restarts



1.3 Properties of distributed, concurrent systems

Distributed and concurrent systems are build from small sequential and process-like blocks ; those blocks may be executed independently and concurrently, as long as they don't interact ; they may be distributed on different places, physically speaking ; they may have to cooperate in various ways to achieve a desired outcome : then, communication happens through asynchronous message passing.

[24] identifies three concurrency patterns :

1. disjoint concurrency, where concurrent activities share nothing and do not synchronize (in the manner of independent UNIX processes),
2. competitive concurrency, where different running activities compete for some shared resources (like in the case of transactional systems),
3. cooperative concurrency, where different active entities act collectively to reach a common outcome.

Sometimes it seems like the distinction between the problem of one system's consistency and the handling of distributed, concurrent exceptions, is blurred. We defend the view that exception handling is a low-level construct which ought to remain orthogonal, albeit compatible, with transactional systems. Both are important with relation to software reliability and fault-tolerance, but they really are independent things. From now on, we will mainly focus on cooperative concurrency, without after-thoughts about transactions. The questions we face are :

- how to deal with asynchronous message-passing,
- how to design signalment path and boundaries, and exception handlers scope.

Older proposals, such as Guide or Argus [18], helped raise the importance of the notion of *concertation* amongst a set of cooperating entities. As we will show, SaGE and Guardian have retained this notion.

2 An overview of three systems

Erlang is a high-level programming language with built-in features to provide fault-tolerance to massively concurrent and distributed programs. It is a complete, integrated system. It has been proven in the field [1].

On the other hand, SaGE and Guardian are independent of any language. But it must be noted that they fill an important gap in the Java world and their current implementations are purely Java based.

2.1 Erlang

Erlang has been built out of the need of an efficient, robust and massively parallel programming language for the programming of world-class telecoms switches ; it has proven successful with relation to these objectives.

2.1.1 Language constructs

At its core, Erlang is a pure *functional* programming language (i.e closures and function calls are everything) with a simple EHS providing exception signalment and handling but *no restarts* (i.e. no resumption).

Erlang has a notion of *process*, which is an independent, distributable sequential piece of code. Erlang processes communicate through asynchronous message passing : no shared state is ever allowed between concurrent entities.

There are two types of processes : worker and supervisor processes. A typical Erlang program is a tree of supervisor processes whose leaf nodes are worker processes. At run time, the program is unfolded from the root supervisor process (which can be seen as the «main» process) to the leaves made of worker processes. The act of creating a new process, and message passing between processes, are always explicit.

2.1.2 Exception handling

Whenever an exception is signaled (typically in a worker process), it can be handled within the process, with the classical set of operators ; if it isn't handled there, then the process automatically sends a message carrying failure notification to its direct supervisor - the last step

in signalment - and suicides. Depending of the signaling process's expected durability (permanent, transient or temporary), it is restarted *from scratch* by its supervisor or nothing happens.

Thus, a well-written Erlang program has its functionality distributed into a tree structured hoard of independent worker processes, such as to limit the consequences of a process crash (an unhandled exception) to a subtree.

It must be noted that since Erlang is not object oriented, exception «objects» are built with tuples ; there is at least one entry carrying type information in order to enable the selection of the correct handler according to the type of the error.

2.1.3 Conclusions

This has proven to be a very efficient, allegedly because simple, strategy for Erlang's initial target, namely the software running massive telecoms switches. However, it is not a complete transposition of the EHS feature set.

2.2 SaGE

SaGE is a specialized DEHS, in that it first defines protocol, in other words a fixed set of interactions between groups of concurrent and (possibly) distributed blocks, and then defines a proper EHS on this protocol, which is named "service". SaGE has been implemented for the Mad-kit [11] multi agents system and the Jonas J2EE [12] component-based framework.

2.2.1 Services

A service is basically modelled after method invocation in object oriented languages ; it associates an agent, a name, a set of formal parameters and a piece of functionality. A service is said to be provided by an agent A, invoked by an agent B, and then executed by agent A.

There are however important differences with method invocation. In a non-concurrent object-oriented program, there is a call chain of method invocations ; with services, there is a tree of service invocations.

This is because :

- one agent can execute in parallel many instances of any service (mapping to an even number of service invocations by possibly different agents),
- one service execution which invokes another service execution (it is then said to be a "complex" service) is not synchronously stuck to its normal or exceptional outcome, but can err on its own and indeed invoke many other sub-services, for the sake of redundancy or merely for it is profitable to exploit the distributedness of the agents, making the worst-case execution time the maximum of the execution of the distributed sub-services instead of the sum.

Such a protocol is said to be "semi-synchronous", that is an invoking process is free to do whatever it pleases until some later time when it decides to collect the answers, which it is, contractually, bound to do. This is much like Multilisp's "futures" [7].

It is important to note that the service protocol presents itself to the programmer as an API which hides whenever possible the low level messaging layer, and tries to match closely the semantics of object orientedness in the manner of Smalltalk or Java, notwithstanding the extensions to the exception handling capabilities.

2.2.2 Exception handling

On a call-tree degenerated into a chain (or stack), the EHS behaves quite like in a monolithic Java program. First, there is a new type of exception : *SageException*. Then, for each service, one can associate a set of handlers with any set of subtypes of *SageException*. The exception handler is quite limited in its capability for it cannot re-enter the service for which it handles an exception ; however this is bound to implementation issues, that could be alleviated without too much heroism. It is declared and bound to a service and a set of exceptions through a *handle (...)* {...} block. Finally, there is a special signaling primitive (*signal ()*).

In the general case, there is an important piece of functionality available that does not belong to the EHS of non concurrent systems. We have said that one service can issue many sub-service invocations for the sake of redundancy ; when doing this, one does not want to have our service execution killed because just one (or a few amongst many) redundant sub-service execution failed and signaled an exception ; instead it is

practical to craft a special function that filters the exceptions and can decide not to handle or propagate upwards, that is to merely do nothing special and go on with the computation. This function is named the "*concertation*" function, after V. Issarny's work [9].

2.2.3 Limits and conclusion

There are limits to SaGE, by definition : we can model a problem with SaGE as long as it fits well into the notion of service. That is, it works well for any problem for which a functional decomposition comes in handily like, for instance, information retrieval systems.

Indeed, SaGE tries hard to put the least possible amount of ordering constraints on the execution of services, in order to benefit as much as possible of concurrency of the executions. The only constraint is that a service execution starts before any of its direct and indirect sub-services, and terminates only after all of them have terminated. Thus, when service execution side-effects an agent's mental state or the outer world, then suddenly every kind of interlocking and race conditions raise its head ; at this point, the programmer has to cope with the ordering and serialisation of access to the shared entities through low level constructs like locks and mutexes, which are hard to get right and may have scalability issues, or alternatively he could rely on a (costly but reliable) transactional system.

While restricted in various ways, SaGE is nevertheless an efficient and expressive framework that provides a simple DEHS on top of an oo-like interaction model.

2.3 Guardian

While SaGE defines a DEHS upon a set of well defined interactions, Guardian strives to be a complete DEHS, in the sense that it can work with any protocol for which an exception handling strategy can make sense. This mighty goal seems to be attained with Guardian, but at a price ; we will see how and why soon.

2.3.1 Components and primitives

A multi-agent system extended with Guardian has two main new components : a special agent named "leader", and sets of agents engaged

in a collective activity, whose members are said to be "participants". It adds to the participants a set of primitives to manage the exception handling whose scope is the collective activity. Exception handling at the level of a set of participants is said to be global ; accordingly, Guardian introduces a "GlobalException" sub hierarchy to the built-in exception taxonomy, like with SaGE and its "SageException".

Those primitives mimic the Java EHS model, with one addition : methods to enable and disable local *contexts*. Contexts are a generalisation of the lexical context, from the viewpoint of exception handling of the monolithic EHSs ; in those, it is easy to think about the places where signaling and handling happen : all EHSs provide syntax to help the programmer map an exception handler or a signaling instruction to their contexts. Guardian contexts represent phases of the program, those phases being now decoupled from the block structure of the programming language. Moreover, Guardian, being built on Java, cannot benefit from any syntactic sugar ; therefore it has to make context management available to the programmer through the "{enable|disable}Context" pair of primitives.

So, when writing an application exploiting the capabilities of Guardian, one has to think about the specification of the participants, which are bound to use the Guardian primitives to stack up contexts, to install exception handlers on those contexts and to signal global exceptions ; and about the leader, which computes the set of exceptions to be handled in the participants in response to a set of concurrent global exceptions.

According to the authors, the gist of exception handling is the choice of the «semantically correct» global exception to be signaled by a participant, so as to allow the leader to compute a meaningful concerted exception set in return. Off course, the choice depends heavily on the type of the collective activity.

It helps remind us that whatever the EHS we are in, the two most important questions an exception handler can ask an exception object are : who are you and where do you come from ?

2.3.2 The signalment process

Signaling is a two phases mechanism, after a first global exception is signaled :

1. the leader is alerted, tells every other participant to suspend any work and wait synchronously for instructions, and then it collects all pending global exceptions,
2. when all the participants are ready, the leader computes a set of concerted exceptions to be raised in the participants in response to the set of global exception it has received, then it sends the computed exceptions back to the participants, which can then enter their handler and go on with their (concurrent) activities.

We should note that the leader is akin to the concertation function from SaGE : concertation is the central mechanism in Guardian whereas it is an auxiliary mechanism in SaGE.

2.3.3 Limits and conclusion

Guardian has been built to address many, if not all, of the situations a DEHS has to cope with. For instance, coordinated atomic actions, conversations, even some level of transaction processing can be built using the Guardian framework.

But it does not scale well : the need to freeze all participants' activity whenever a global exception is signaled makes it not usable when the participant set becomes big or if some participants have to perform some uninterruptible, high-priority or real-time service. Moreover, the system would be sensible to crashes or unavailability of the leader ; uniqueness of the leader seems like a weak point.

Furthermore, one important problem remains the complexity of the system, for it is quite hard to figure out how to use it properly.

Those points relate to the fact that Guardian is more than a DEHS framework and seems to encourage the programmer to build transactional-like systems using exception handling capabilities, since those come with automatic agent synchronisation for free. In fact, there is probably a wrong coupling between EHS semantics and execution serialisation : both should probably not come in the same package, and under the same name.

3 Delving further : the case for resumption with restarts

3.1 Qualities and possible shortcomings of the three systems

Both SaGE and Guardian define a scope for the definition of handlers and a signalment process. SaGE has an efficient and scalable DEHS system, but it is restricted to the service protocol, which may not be suitable to any agent situation (for instance, conversations or transactions are not easily modelled in terms of services). Guardian offers building blocks to help building a general purpose DEHS, but fails to address all distributed problems, due to the centralised and synchronised handling of global exceptions.

Also, both focus on providing an EHS in its simplest terms : there are no provisions made for repairing and resuming some, or all of the entities engaged in a collective activity.

Finally, the Erlang programming language has excellent capabilities with relation to fault tolerance, based on an extremely simple DEHS. Erlang does not try to do any concertation when an exception propagates from one process up to its supervisor : the process, and possibly a whole group of processes involved in a coordinated operation, will be restarted anew.

All three systems lack the ability to resume an activity around the signaling point. While there are many reasons for not wanting to use an EHS with resumption, even in the context of monolithic programs, we have yet to see a good one to dismiss it as unworthy. We want to close this paper in making the case for resumption using restarts in a DEHS.

Even if not heavily used within programs, restarts offer surprising opportunities and unmatched ease for those that need to build a program's test suit ; in other words, it can help a lot, even if indirectly, in the debugging phase. It is just one example, albeit an important one. One could collect many other use cases that exhibit the usefulness of restarts in monolithic programs. However we are not going to do this, for it would be going backwards. We want to show there exists a world in which DEHS plus restarts are a good thing.

3.2 Entropy and the EHS

A terminating DEHS can only help dealing with hard, uncorrectable errors. In practice, there are a lot of non-error conditions that are to be dealt with, or possibly just ignored. The more a system is open to «real world» entities, the more it has to deal with those conditions : transient network failures, laptop batteries going low, one's cup of tea getting cold, in other world the manifestations of entropy. Most of the time, those conditions can be safely ignored by some program. And in this case, we need a way to say «I don't care» in no more words.

Sometimes, these conditions have to be taken seriously : someone ought to plug in this network cable again, so that we can go on with our current work set. We need a way to be notified that something is going wrong, yet ask for a reparation and just resume work afterwards.

Very often, smallish conditions will grow, if unhandled, into proper errors. That is because there is a link between the two ; after all, a low batteries notification shortly precedes a power outage, if nothing is done in time.

So we have to correct the small glitches in order to prevent the bigger breakages. This is even more needed in distributed and concurrent systems facing the real world, that is, a lot of unexpected albeit not yet deadly events.

3.3 No «fail early, fail often» philosophy ?

The designers of Erlang had the good taste to provide an explanation for the lack of resumption in the language ; an explanation, that is not arguments such as «resumption is too hard to implement in our language», «users don't want it» or even worse «our users do not and will not ever need it, by decree of the marketing department».

They had to devise a simple but expressive language (thus functional), with strong grips on massive concurrency and distributedness (hence the notion of processes as part of the language, not out of what an operating system can offer) and advanced serviceability and reliability features (thus, amongst many other features, a *simple, terminal* exception handling system). They had to built software that run continuously for years without interruption, and allowing incremental live upgrades.

The «fail-fast» strategy embodied by Erlang has its roots in the early work on high-reliability systems (such as Tandems systems), where the following properties were spotted :

- the less complicated the code, the less bugs there is,
- there are really two kinds of bugs : «Bohr bugs» which show up deterministically, and «Heisenbugs», that occur under a set of statistically unreproducible conditions,
- deterministic bugs can be caught early and hammered down, but heisenbugs will remain uncovered for undetermined extents of time,
- process pairing is a very efficient solution to catch the heisenbugs : it is likely that out of two process doing the same work, an heisenbug will raise its head only in one of them !

We could go on like this but we have enough material to get to our point : in Erlang, resumption is deemed too complicated (it complicates code path indeed) to solve even transient or minor problem ; in the context of hundred of thousands of concurrently running processes, it is wise to keep it simple, stupid, and be just brutal : when a process or a group of process is in fault, we had better «reboot» it, in the hope that it will work. And it just works.

But it works in the context of telecoms switches, which exposes the real world to the software in a (we believe) limited and controlled manner. The cost of rebooting a subtree of the running application may be, in those circumstances, lower than the cost of more complicated, also possibly buggy since less tested, code paths.

We argue, with no more tools than common sense, that it cannot apply to all the distributed, concurrent systems we may have to build : whenever a robot, in a factory, experiences an exceptional and locally unsolvable situation, we don't think that destroying the robot and replacing it is a good way to handle the problem. We need a more subtle solution, that is some yet-to-come full-blown DEHS.

3.4 Conclusion

Restarts are a useful construct that prevent, where needed, a lot of ad-hockery. We need them in the future DEHS. Are those going to be the

same as in Common Lisp ? We don't know yet. A system supporting them has to be built first, and then experiments have to be carried out. We are far from it. We doubt however that something can be done with programming languages that lack support for them, like Java. The Java world will forever and sorely miss them. It would be however an interesting experiment to bring them to Erlang, or to extend Common Lisp with primitives to deal, *à la* Erlang, with issues of concurrency and distributedness.

References

- [1] ARMSTRONG, J. *Making Reliable Distributed Systems in the presence of Software Errors*. PhD thesis, 2003.
- [2] DONY, C. Exception handling and object-oriented programming : towards a synthesis. *ACM SIGPLAN Notices* 25, 10 (Oct. 1990), 322–330. *OOPSLA/ECOOP '90 Proceedings*, N. Meyrowitz (editor).
- [3] DONY, C. A fully object-oriented exception handling system : Rationale and smalltalk implementation. In *Advances in Exception Handling Techniques* (2001), A. Romanovsky, C. Dony, J. L. Knudsen, and A. Tripathi, Eds., LNCS (Lecture Notes in Computer Science) 2022, Springer-Verlag.
- [4] FERBER, J. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Addison-Wesley Pub Co; 1st edition (February 25, 1999).
- [5] FRÉDÉRIC SOUCHON, SYLVAIN VAUTIER, C. U. C. D. Fiabilité des applications multi-agents : le système de gestion d'exception sage.
- [6] GOODENOUGH, J. B. Exception handling: Issues and a proposed notation. *Communications of the ACM* 18, 12 (Dec. 1975), 683–696.
- [7] HALSTEAD, R., AND LOAIZA, J. Exception handling in multilisp. In *1985 Int'l. Conf. on Parallel Processing* (Aug. 1985), pp. 822–830.

- [8] ICHISUGI, Y., AND YONEZAWA, A. Exception handling and real-time features in an object oriented concurrent language. In *Concurrency: Theory, Language and Architecture* (1991), Springer-Verlag, Ed., pp. 92–109.
- [9] ISSARNY, V. Concurrent exception handling. In *Advances in Exception Handling Techniques* (2001).
- [10] J. ICHBIAH, J.G.P. BARNES, J. H. B. K.-B. O. R. B. W. Rationale for the design of the ada programming language. In *ACM Sigplan Notices* (1979), vol. 14(6B).
- [11] JACQUES FERBER, O. G. Madkit: Organizing heterogeneity with groups in a platform for multiple multi-agent systems.
- [12] Java open application server (JOnAS) 4.1 : a J2EE platform, June 2004. <http://www.objectweb.org/jonas/current/doc/JOnASWP.html>.
- [13] JONES, S. P. Tackling the awkward squad : monadic input/output, concurrency, exceptions and foreign-language calls in haskell.
- [14] J.P. BANÂTRE, V. I. Exception handling in communicating sequential processes. In *Technical Rreport 660, IRISA, INRIA Rennes* (1992).
- [15] KLEIN, M., AND DELLAROCAS, C. Exception handling in agent systems. In *Proceedings of the Third Annual Conference on Autonomous Agents (AGENTS-99)* (New York, May 1–5 1999), O. Etzioni, J. P. Müller, and J. M. Bradshaw, Eds., ACM Press, pp. 62–68.
- [16] KNUDSEN, J. L. Fault tolerance and exception handling in beta. In *Advances in Exception Handling Techniques* (2001), A. Romanovsky, C. Dony, J. L. Knudsen, and A. Tripathi, Eds., LNCS (Lecture Notes in Computer Science) 2022, Springer-Verlag.
- [17] KOENIG, A. R., AND STROUSTRUP, B. Exception handling for C++. In *Proceedings «C++ at Work» Conference* (nov 1989).
- [18] LISKOV, B. Distributed programming in argus. In *Communications of the ACM, vol. 31, n° 3* (1988), pp. 300–312.
- [19] MILLER, R. *The Guardian Model for Exception Handling in Distributed Systems*. PhD thesis, 2003.

- [20] PETER VAN ROY, S. H. Mozart, a programming system for agents applications.
- [21] PITMAN, K. Exceptional situations in lisp. In *EUROPAL'90* (1990).
- [22] PITMAN, K. Condition handling in the lisp language family. In *Advances in Exception Handling Techniques* (2001).
- [23] ROMANOVSKY, A., DONY, C., KNUDSEN, J. L., AND A. TRIPATHI. *Advances in Exception Handling Techniques*. LNCS (Lecture Notes in Computer Science) 2022. Springer-Verlag, 2001.
- [24] ROMANOVSKY, A. B., AND KIENZLE, J. Action-oriented exception handling in cooperative and competitive concurrent object-oriented systems. In *Advances in Exception Handling Techniques* (2001), pp. 147–164.
- [25] SOUCHON, F., URTADO, C., VAUTIER, S., AND DONY, C. 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* (2003), pp. 84–91.
- [26] SUN MICROSYSTEMS. *Java 2 Platform, Standard Edition (J2SE)*. Mountain View, Calif., June 2004. <http://java.sun.com/j2se>.
- [27] SYLVAIN CONCHON, F. L. F. Jocaml : mobile agents for objective caml.
- [28] TRIPATHI, A., AND MILLER, R. Exception handling in agent-oriented systems. In *Advances in Exception Handling Techniques* (2001), A. Romanovsky, C. Dony, J. L. Knudsen, and A. Tripathi, Eds., LNCS (Lecture Notes in Computer Science) 2022, Springer-Verlag.
- [29] WEINREB, D. L. Signalling and handling conditions. Technical report, Symbolics, Inc., Cambridge, MA, Jan. 1983.