

RELEASE: A High-level Paradigm for Reliable Large-scale Server Software

Olivier Boudeville¹, Francesco Cesarini², Natalia Chechina³,
Kenneth Lundin⁴, Nikolaos Papaspyrou⁵, Konstantinos Sagonas⁶,
Simon Thompson⁷, Phil Trinder³, and Ulf Wiger²

¹ EDF R&D, 92140 Clamart, France

² Erlang Solutions AB, 113 59 Stockholm, Sweden

³ Heriot-Watt University, Edinburgh, EH14 4AS, UK

⁴ Ericsson AB, 164 83 Stockholm, Sweden

⁵ Institute of Communication and Computer Systems (ICCS), Athens, Greece

⁶ Uppsala University, 751 05 Uppsala, Sweden

⁷ University of Kent, Canterbury, CT2 7NF, UK

<http://www.release-project.eu>

Abstract. Erlang is a functional language with a much-emulated model for building reliable distributed systems. This paper outlines the RELEASE project, and describes the progress in the first six months. The project aim is to scale the Erlang’s radical concurrency-oriented programming paradigm to build reliable general-purpose software, such as server-based systems, on massively parallel machines.

Currently Erlang has inherently scalable computation and reliability models, but in practice scalability is constrained by aspects of the language and virtual machine. We are working at three levels to address these challenges: evolving the Erlang virtual machine so that it can work effectively on large scale multicore systems; evolving the language to Scalable Distributed (SD) Erlang; developing a scalable Erlang infrastructure to integrate multiple, heterogeneous clusters. We are also developing state of the art tools that allow programmers to understand the behaviour of massively parallel SD Erlang programs. We will demonstrate the effectiveness of the RELEASE approach using demonstrators and two large case studies on a Blue Gene.

Keywords: Erlang, scalability, multicore systems, massive parallelism

1 Introduction

There is a widening gap between state of the art in hardware and software. Architectures are inexorably becoming manycore, with the numbers of cores per chip following Moore’s Law. Software written using conventional programming languages, on the other hand, is still essentially sequential: with a substantial effort, some degree of concurrency may be possible, but this approach just doesn’t scale to 100s or 1000s of cores. However, manycore programming is not only about concurrency. We expect 100,000 core platforms to become commonplace,

and the predictions are that core failures on such an architecture will become relatively common, perhaps one hour mean time between core failures [1]. So manycore systems need to be both scalable and robust.

The RELEASE project aim is to scale the radical concurrency-oriented programming paradigm to build reliable general-purpose software, such as server-based systems, on massively parallel machines. The trend-setting concurrency-oriented programming model we will use is Erlang/OTP (Open Telecom Platform) – designed in the telecoms sector, is based on highly-scalable lightweight processes *which share nothing*, and used in strategic Ericsson products such as the AXD301 telecoms switch [2]. Erlang/OTP provides high-level coordination with concurrency and robustness built-in: it can readily support 10,000 processes per core, and transparent distribution of processes across multiple machines, using message passing for communication. The robustness of the Erlang distribution model is provided by hierarchies of supervision processes which manage recovery from software or hardware errors. Erlang is an Actor-based programming language where actions are performed by concurrent processes named actors [3]. Actors communicate with each other via asynchronous message passing, and each actor has an address and a mailbox. An actor also has a behaviour that may change in a response to a received message. Actors may create and kill other actors [4].

The Erlang/OTP has inherently scalable computation and reliability models, but in practice scalability is constrained by the transitive sharing of connections between all nodes and by explicit process placement. Moreover programmers need support to engineer applications at this scale and existing profiling and debugging tools do not scale, primarily due to the volumes of trace data generated. In the RELEASE we tackle these challenges working at three levels (Figure 1):

1. *evolving the Erlang Virtual Machine (VM)* so that it can work effectively on large scale multicore systems (Section 3.1);
2. *evolving the language to Scalable Distributed (SD) Erlang*, and adapting the OTP framework to provide both constructs like locality control, and reusable coordination patterns to allow SD Erlang to effectively describe computations on large platforms, while preserving performance portability (Section 3.2);
3. *developing a scalable Erlang infrastructure* to integrate multiple, heterogeneous clusters (Section 3.3).

These developments will be supported by state of the art tools to allow programmers to understand the behaviour of large scale SD Erlang programs, and to refactor standard Erlang programs into SD Erlang (Section 3.4). We will demonstrate the effectiveness of the RELEASE approach through building two significant demonstrators: a simulation based on a port of SD Erlang to the Blue Gene architecture [5] and a large-scale continuous integration service – and by investigating how to apply the model to an Actor framework [6] for a mainstream language (Section 4).

The Erlang community is growing exponentially, moreover it has defined a widely accepted concurrency-oriented programming and become a beacon lan-




Fig. 1. RELEASE Project Research Areas

guage for reliable distributed computing. As such it influences the design and implementation of numerous Actor programming languages, libraries and frameworks like Cloud Haskell [7], Scala and its Akka framework [8], F# [9], and Kilim [10]. Hence, we expect that RELEASE will have a similar impact on the design and implementation of a range of languages, libraries and frameworks, and thus deliver significant results beyond the Erlang community.

2 Progress Beyond the State-of-the-art

The RELEASE project targets reliable scalable general purpose computing on heterogeneous platforms. Our application area is that of general server-side computation, e.g. a web or messaging server. This form of computation is ubiquitous, in contrast to more specialised forms such as traditional High Performance Computing (HPC). Moreover, this is computation on stock platforms, with standard hardware, operating systems and middleware. We aim for 10^5 cores, for example, the Blue Gene/Q that will be exploited during the project has 65,000 cores. Our focus on commodity hardware implies that we do not aim to exploit the experimental many-core architectures like the Intel Tera-scale [11].

The project makes advances in a number of areas, primarily in Virtual Machine (VM) support for high-level concurrency (Section 2.1), in scalable high-level distributed languages (Section 2.2), in tools for concurrent software engineering (Section 2.3), in cloud hosting infrastructure (Section 2.4), and in massive scale simulation (Section 2.5).

2.1 VM Support for High-level Concurrency

Due to the development of multicore architectures and the distributed nature of modern computer systems, the quantity and variance of processors on which software is executed has increased by orders of magnitude in recent years and is expected to increase even further. In this setting the role of high-level concurrent programming models is very important. Such models must abstract from variations introduced by differences between multicore architectures and uniformly

treat different hardware architectures, number of cores and memory access characteristics. Currently, the implementation of such models in the form of high-level languages and libraries is not sufficient; these two must be complemented with efficient and reliable VMs that provide inherent support for high-level concurrency [12].

Historically the efficient implementation of VMs for single core processor systems has presented a number of challenges, largely unrelated to concurrency. For example, in order to optimally use the hardware, a VM has to exploit deep hierarchies of cache memory by reorganizing data layouts, and to support e.g. out-of-order execution, the hardware's prefetching heuristics, branch prediction. With multicore processors, concurrency and the inherent shared memory model introduce new challenges for VMs, as it is not only arbitrary thread interleavings but parallel execution on limited shared resources which has to be taken into account. On top of the implementation considerations mentioned before, cache coherence becomes a critical issue, memory barriers become a necessity and compiler optimisations, e.g. instruction reordering, must often be more conservative to be semantics-preserving. Furthermore, multicore machines currently rely on Non-Uniform Memory Access (NUMA) architectures, where the cost of accessing a specific memory location can be different from core to core and data locality plays a crucial role for achieving good performance.

In the past few years, there has been sustained research on the development of VMs for software distributed shared memory. Some recent research aims to effectively employ powerful dedicated and specialized co-processors like graphic cards. Notable VM designs in this direction are the CellVM [13] for the Cell Broadband Engine, a VM with a distributed Java heap on a homogeneous TILE-64 system [14]; an extension of the JikesVM to detect and offload loops on CUDA devices [15]; and VMs for Intel's Larrabee GPGPU architecture [16]. Most of these designs are specific to VMs for Java-like languages which are based on a shared-memory concurrent programming model. Despite much research, the implementation of shared memory concurrency still requires extensive synchronisation and for this reason is inherently non-scalable. For example, in languages with automatic memory management, the presence of a garbage collector for a heap which is shared among all processes/threads imposes a point of synchronisation between processes and thus becomes a major bottleneck.

Although a language based on the Actor concurrency model is in principle better in this respect, its VM implementation on top of on shared memory hardware in an efficient and scalable way presents many challenges [17]. In the case of the implementation of Erlang, various runtime system architectures have been explored by the High Performance Erlang (HiPE) group in Uppsala, based either on process-local memory areas, on a communal heap which is shared among all threads, or following some hybrid scheme. However, the performance evaluation [18] was conducted on relatively small multiprocessor machines (up to 4 CPUs) in 2006 and cannot be considered conclusive as far as scalability on the machines that the RELEASE project is aiming at. More generally, the technology required to reach the scalability target of the current project requires

significant extensions to the state of the art in the design of VMs for message passing concurrency and will stretch the limits of data structures and algorithms for distributed memory management.

2.2 Scalable Reliable Programming Models

Shared memory concurrent programming models like OpenMP [19] or Java Threads are generally simple and high level, but do not scale well beyond 10^2 cores. Moreover reliability mechanisms are greatly hampered by the shared state, for example, a lock becomes permanently unavailable if the thread holding it fails. In the HPC environment the distributed memory model provided by the MPI communication library [20] dominates. Unfortunately MPI is not suitable for producing general purpose concurrent software as it is too low level with explicit, synchronous message passing. Moreover the most widely used MPI implementations offer no fault recovery⁸: if any part of the computation fails, the entire computation fails.

For scalable high-level general purpose concurrent programming a more flexible model is required. Actors [6] are a widely used model and are built into languages like Scala [22], Ptolemy [23], and Erlang [3]. There are also Actor frameworks and libraries for many languages, for example, Termite Scheme [24], PARLEY for Python [25], and Kilim [10] for Java. The Erlang style concurrency has the following key aspects [26]: fast process creation/destruction; scalability to support more than 10^4 concurrent processes; fast asynchronous message passing; copying message-passing semantics, i.e. share-nothing concurrency; process monitoring; selective message reception.

Figure 2 illustrates Erlang’s support for concurrency, multicores and distribution. A rectangle represents a host with an IP address, and an arc represents a connection between nodes. Multiple Erlang processes may execute in a node, and a node can exploit multiple processors, each having multiple cores.

⁸ Some fault tolerance is provided in less widely used MPI implementations like [21].




Fig. 2. Conceptual View of Erlang’s Concurrency, Multicore Support and Distribution

supports single core concurrency where a core may support as many as 10^8 lightweight processes [2]. In the Erlang distribution model a node may be on a remote host, and this is almost entirely transparent to the processes. Hosts need not be identical, nor do they need to run the same operating system.

Erlang currently delivers reliable medium scale distribution, supporting up to 10^2 cores, or 10^2 distributed memory processors. However, the scalability of a distributed system in Erlang is constrained by the transitive sharing of connections between all nodes and by explicit process placement. The transitive sharing of connections between nodes means that the underlying implementation needs to maintain data structures that are quadratic in the number of processes, rather than considering the communication locality of the processes. While it is possible to explicitly place large numbers of processes in a regular, static way, explicitly placing the irregular or dynamic processes required by many servers and other applications is far more challenging. The project addresses these limitations.

2.3 Tools for Concurrency and Erlang

From the inception of parallel programming, the availability of tools that ease the effective deployment of programs on parallel platforms has been a crucial criterion for success. The Intel Trace Analyzer and Collector [27] is a typical modern tool, which supports the analysis, optimisation and deployment of applications on Intel-based clusters with MPI communication.

The Erlang VM is equipped with a comprehensive, low-level tracing infrastructure provided by the trace built-in functions [3]. There is a number of higher-level facilities built upon this, including the debugger and the Trace-Tool Builder (TTB) [28]. The later provides facilities for managing tracing across a set of distributed Erlang nodes. TTB is designed to be extensible to provide different types of tracing tuned to different applications and environments. While these tools support monitoring Erlang programs across distributed nodes, there is a problem in dealing with the volume of data generated; on a multicore chip, or highly parallel system, it will be impossible to ship all the data off the chip without swamping communication entirely [29]. So, we will ensure that local monitoring and analysis can be performed within the system, leveraging the locality supported by SD Erlang to support a hierarchical 'tracing architecture'.

In building analyses we will be able to leverage work coming from the FP7 ProTest project [30] including the on-line monitoring tool Inviso, which has recently been integrated into TTB, and the off-line monitoring tool Exago. Higher-level analysis can be provided independently [31], and building on this it is also possible to analyse the process structure of Erlang applications, particularly those structured to use the OTP standard behaviours [32, Chapter 6]. These existing systems need to be extended to integrate monitoring with process discovery.

Building on the Erlang `syntax_tools` package and the standard Erlang compiler tool chain, Wrangler is a tool for refactoring Erlang programs [33]. In the RELEASE project we will develop and implement refactorings from Erlang to SD

Erlang within Wrangler. To give users guidance about which refactorings could be applied we will develop a refactoring assistant, that will suggest refactorings of a given system on the basis of the system behaviour.

Traditional breakpoint debuggers are of little use for Reliable Massively Parallel (RMP) software; there are too many processes and breakpoints conflict with the frequent timeouts in Erlang's communication. The RELEASE project will build a debugger for massively parallel systems that retains a recent history of the current computation state, as well as saving fault-related information, so that this can be retraced and explored for debugging purposes.

2.4 Cloud Hosting

The cloud hosting area is moving rapidly, and progress is driven mainly by entrepreneurs and competing cloud infrastructure providers. While users can already choose between many hosting providers offering price/performance alternatives, provisioning is typically manual and time consuming, making it unnecessarily difficult to switch providers.

With increased competition comes specialisation, which in its turn introduces an integration challenge for users. This calls for a broker layer between users and cloud providers [34], but creating such a broker layer, especially a dynamic one, is no easy task, not least because Cloud Provision APIs currently do not support ad-hoc, capability-based provisioning and coordination of cloud resources. A problem is that basic cloud APIs must primarily serve mainstream languages, not least REST (stateless) clients, where management of dynamic and complex state is considered extremely difficult [35].

We intend to advance the concept of a capability-based dynamic cloud broker layer, building on the cloud broker component of Erlang Solutions' recently launched Hosted Continuous Integration tool SWARM. SWARM is capable of allocating a mix of cloud and physical resources on demand and running complex tests in parallel, and draws on the ease with which Erlang can spawn and coordinate parallel activities, as well as its considerable strengths as a test automation and load testing environment. We intend to use SWARM itself as a testbed for the concepts in the RELEASE project, aiming to increase both the capabilities of SWARM and its cost-effectiveness.

At a basic level, this broker would be useful for any kind of cloud application (even non-Erlang) making it easier for users to switch provider, and create on-demand, capability-driven mashups of specialised clusters. For more dynamic capabilities, we will lead by providing a Cloud Broker API customised for the Erlang programming model.

2.5 Scalable Simulation

Many business and scientific fields would benefit from larger-scale simulations as the need for intrinsically risky extrapolation can be removed if we are able to simulate a complex system as a whole. Moreover, some important behaviours

only appear in detailed simulations. Distributed memory clusters are cost effective scalable platforms. However, simulations are typically initially developed using sequential, imperative technologies and these are ill suited for distributed execution. Additionally parallelism poses a number of challenges for simulation. We believe that declarative programming is an appropriate tool to harness parallel platforms.

Sim-Diasca is a distributed engine for large scale discrete simulations implemented in Erlang. It is among the most scalable discrete simulation engines and currently is able to handle more than one million relatively complex model instances using only hundreds of cores. However, more accurate models are required – ideally we would like 50-200 million model instances. We would also like to introduce reliability. The achievement of these goals requires scalable reliability improvements across the software stack, i.e. VM, language, and engine.

3 Developing a Scalable Programming Model

This section gives an overview of the technical and research components of the project covering the following areas: scaling the Erlang VM (Section 3.1), scaling the Erlang programming model (Section 3.2), scalable virtualisation infrastructure (Section 3.3), and tool support for the development of RMP software (Section 3.4).

3.1 Scaling the Erlang VM

The Erlang VM was initially designed as a machine to support cooperative concurrent execution of processes (“green threads”) on physical machines with a single CPU. Since 2006, the Erlang VM has been extended to support Symmetric MultiProcessing (SMP) in the form that is commonly found these days in multicore machines. The goal has been to focus on stability while giving incremental performance improvements in each release. This approach has worked well for the important class of high-availability server software running on machines with up to 16 cores, but needs significant extensions to achieve scalability on bigger multicore machines.

In the RELEASE project we aim to re-design and improve some core aspects of Erlang’s VM so that it becomes possible for applications to achieve highly scalable performance on high-end multicore machines of the future with minimal refactoring of existing applications. Our goal is to push a big part of the responsibility for achieving scalability from the application programmer to the VM. In particular, we will investigate architectural changes and alternative implementations of key components of Erlang’s VM that currently hinder scalability of applications on large multicore machines. A key consideration is to design scalable components without sacrificing crucial aspects of the existing architecture such as process isolation, high-reliability, robustness and soft real-time properties.

To achieve these goals, we will begin our work by a detailed study of the performance and scalability characteristics of a representative set of existing Erlang applications running on the current VM. This will help us both to identify the major scalability bottlenecks and to prioritize changes and extensions of the runtime system and of key VM components. Independently of the results of this study however, there are parts of the VM which could definitely benefit from extensions or redesign. One of them is the Erlang Term Storage (ETS) mechanism.

Although Erlang processes do not share any memory at the language level, at the implementation level, the Erlang VM provides built-ins that allow processes to store terms in shared global data structures, called ETS tables, and to destructively modify their contents. Currently, many Erlang applications make extensive use of this mechanism, either directly in their code or indirectly via using in-memory databases such as Mnesia [3]. With the current implementation of ETS, when the number of processes gets big, an access to these tables becomes serialisation points for applications and hinder their scalability. Moreover, due to the nature of the garbage collector currently employed by the Erlang VM, an access to terms in ETS tables requires a physical copy of the term from the table to the heap of the process. We will investigate scalable designs of ETS tables that avoid these performance bottlenecks. In addition we will experiment with runtime system organisations and design language built-ins that avoid the need for copying data from ETS tables to the process-local memory areas when it is provably safe to do so.

In the current Erlang VM processes allocate the majority of their data in process-local memory areas. Whenever processes need to communicate, they must explicitly copy their data from the heap of the sender to that of the receiver. The processes also need to wait to get rescheduled when execution reaches a receive statement with no appropriate messages in the process mailbox. We will design and investigate scalable runtime system architectures that allow groups of processes to communicate without the need for explicit copying. We will also develop the runtime support for processes to temporarily yield to their appropriate senders when reaching a receive statement that would otherwise block. One possible such architecture is a clustered shared heap aided by the presence of language constructs such as fibers available in languages like C++.

To identify “frequently-communicating processes” and to guide the VM schedulers a scalable design of the Erlang VM needs to be supported both by language extensions and by static and dynamic analysis. A significant effort in this task is not only to design and implement the analysis, but also to integrate it in the development environment in a smooth and seamless way. No matter how scalable the underlying VM will get, large scale applications will need tool support to identify bottlenecks (Section 3.4). Finally, to test the scalability of our implementation and to enable a case study we will port the Erlang VM to a massively parallel platform. Our current plan is to use a Blue Gene/Q machine. This plan may change, or be extended to include more platforms, if we gain access to more powerful such machines during the duration of the project.

3.2 Scaling the Erlang Programming Model

We will extend both Erlang, to produce Scalable Distributed (SD) Erlang, and the associated OTP library. The SD Erlang name is used only as a convenient means of identifying the language and VM extensions as we expect them to become standard Erlang in future Erlang/OTP releases.

Controlling Connections. The scalability of a distributed Erlang system is constrained by the transitive sharing of connections between all nodes and by explicit process placement. SD Erlang will regain scalability using layering, and by controlling connection locality by grouping nodes and by controlling process placement affinity.

Process Placement. Currently the Erlang distribution model permits explicit process placement: a process is spawned on a named node. Such a static, directive mechanism is hard for programmers to manage for anything other than small scale, or very regular process networks. We propose to add an abstraction layer that maintains a tree of node groups, abstractly modelling the underlying architecture. We will provide mechanisms for controlling *affinity*, i.e. how close process must be located, e.g. two rapidly communicating processes may need to be located in the same node. We will also provide mechanisms for controlling *distribution*, i.e. how far the process must be from the spawning process. For example, two large computations, such as simulation components, may need to be placed on separate clusters.

Scaling Reliability. Erlang/OTP has world leading language level reliability. The challenge is to maintain this reliability at massive scale. For example, any node with a massive number of connections should be placed in a different node group from its supervisor. A new OTP principle could be to structure systems with the supervision tree preserving this property.

Performance Portability. The abstract computational control mechanisms are not strongly related to a specific architecture. As far as possible we intend to construct performance portable programs, and computational patterns, by computing distance metrics for the affinity and distribution metrics. Moreover, locality control enables us to use layering as a principle to structure systems: for example, the control processes for a layer appearing in a different group from the controlled processes. This facilitates performance portability as the top layers can be refactored for the new architecture while preserving the lower layers unchanged.

Scalable and Portable OTP. Some OTP principles and behaviours will need to be extended with new scalable principles, and perhaps to some extent redesigned and refactored to control locality and support layering. The supervisor group discussed above and the control layering are examples of new scalable principles.

3.3 Scalable Virtualisation Infrastructure

Given the aim of the RELEASE project to develop a model of participating clusters, it is logical to also explore the possibility of creating super-clusters of on-demand clusters provisioned from competing Cloud providers. This would make

it possible to cost-optimize large clusters by matching capability profiles against the requested work, and combining groups of instance types, possibly from different providers, into a larger grid. The complexities of running a computing task across such a cluster generally fall into the categories addressed within the RELEASE project: providing a layer of distribution transparency across cooperating clusters; monitoring neighbouring clusters and recovering from partial failures; and tolerance to latency variations across different network links.

A possible small-scale use for a Cloud cluster broker could be to act as a “Pricerunner” for on-demand computing resources. For Erlang Solutions, it is a natural extension of their Hosted Continuous Integration and Testing as a Service offerings (SWARM), where the system can match the computing needs of different build-and-test batches against availability and pricing of virtual images from different providers. In the context of Hosted Continuous Integration, this capability can also be used to simulate both server- and client-side, using different capabilities, and possibly different providers, for each. It needs to be easy to configure and automate.

The infrastructure that we construct will be novel as few cloud computing providers today offer a powerful enough API for this task. For this reason, we will build our own virtualisation environment, e.g. based on the Eucalyptus Cloud software, which is API-compatible with Amazon EC2, but available as open source.

3.4 Tool Support for Developing Reliable Parallel Software

The Erlang programming model provides abundant concurrency, with no theoretical limit to the number of concurrent processes existing – and communicating – at any particular time. In practice, however, there can be problems in the execution of highly concurrent systems on a heterogeneous multicore framework. Two particular difficulties are, first, *balancing of computational load between cores*. Each Erlang process will run on a single core, and a desirable property of the system is that each core is utilised to a similar extent. Secondly, there may be *bottlenecks due to communication*. Each process executes on a single core, but a typical process will communicate with other processes which are placed on other cores, and in a large system this can itself become a bottleneck. Using more cores can ease load balancing, but communication bottlenecks may be alleviated by keeping related processes close together, potentially on the same core, but these two are in tension.

We will supply tools that can measure and visualise performance in a number of different ways. The Erlang VM is equipped with a comprehensive, low-level tracing infrastructure on which other tools can be built. DTrace also provides dynamic tracing support at the operating system level on a number of Unix platforms, complementing the built-in facilities. Because of the volume of data generated, we will need to ensure that local monitoring and analysis are performed, leveraging the locality properties of the running system. The tools will be built in sequence. First we will develop textual reports and graphical visualisations for off-line monitoring, secondly, these will be generated as snapshots of

a running system, and finally we will build tools to support interactive, real-time on-line system monitoring.

For SD Erlang to be taken up in practice, it will be essential to provide users with a migration path from Erlang: using the Wrangler refactoring platform we will provide the infrastructure for users to migrate their programs to SD Erlang, and moreover provide decision-support tools suggesting migration routes for particular code. Finally, we will supply tools to debug concurrency bugs (or Heisenbugs) in SD Erlang programs, and to develop an *intelligent debugger* based on saving partial histories of computations.

4 Case Studies

The case studies will demonstrate and validate the new reliable massively parallel tools and methodologies in practice. The major studies using SD Erlang and the scalable infrastructure are large-scale discrete simulations for EDF, i.e. 10^7 model instances, and dynamically scalable continuous integration service for Erlang Solutions. Key issues include performance portability, scalability and reliability for large-scale software. We will also investigate the feasibility of impacting dominant programming models by adding our scalable reliability technologies to an Actor framework for a mainstream programming language, such as Java.

EDF Simulation. The goals of this first case study are to enhance the reliability and scalability of an existing open-source simulation engine, Sim-Diasca that stands for “Simulation of Discrete Systems of All Scales”. Sim-Diasca is a general-purpose engine dedicated to all kinds of distributed discrete simulations, and currently used as the corner stone for a few simulators of smart metering systems. The increased scalability will enable Sim-Diasca to execute simulations at an unprecedented scale by exploiting large-scale HPC resources, clusters or Blue Gene supercomputer that in total provide 10^5 cores.

Currently, Sim-Diasca is designed to halt if any node fails during a simulation. The goal is to make the engine able to resist to the crash of up to a pre-determined number of computing nodes, e.g. 3 nodes. This is a twofold task:

- The main part of the task is to add application-level fault tolerance to the engine, so that the engine is able to store its state based on triggers, e.g. wall clock durations elapsed, simulation intermediate milestone met, in such a way that state is spread over all the computed nodes involved. Thus, a distributed snapshot is to be created, and the state of a given node must be duplicated to k other nodes depending on the targeted k -reliability class, either in RAM or on a non-volatile memory, e.g. local files, a distributed file system, or a replicated database. This check-pointing is meant to allow for a later restart on a possibly different resource configuration.
- The second part of the task is to integrate the lower-level mechanisms for reliability provided by SD Erlang. Some mechanisms will be transparent, whereas others will be used as building blocks for the first part of the task.

The distributed snapshot/restart feature will be implemented using a trigger, e.g. every N simulation ticks all the attributes of all the simulation instances are replicated on k other nodes. This requires that simulation agents like the time managers and data-loggers have means of serialising their state for future re-use. The main difficulty is to do the action for all model instances.

Tolerating the common case of losing a “slave” node at the leaf of the simulation tree is the basic requirement. More advanced challenges to be investigated include the following:

- Tolerating the loss of intermediate nodes, or even the root node. For the latter case consensus and leader election are required, and may be generic reliability mechanisms for Sim-Diasca.
- The dynamic addition of nodes that naturally follows from tolerating the loss of nodes.
- Supporting instance migration, for example, for planned down-time or load balancing.

Target platforms include both large clusters such as the Blue Genes and multicore cards like Tiler. The advantages of the latter are ready access and local administration.

Continuous Integration Service. The second case study will be to integrate the support for on-demand, capability-driven heterogeneous super clouds into Erlang Solutions’ Continuous Integration framework, SWARM. We envisage integrating up to 4 Amazon EC2 clusters ranging from small, e.g. four 1GHz cores, to large, e.g. hundreds of 3GHz cores. To meet client requirements we will also include in-house clusters, such as bespoke embedded device clusters, e.g. 20 ARM cores, or clusters of dedicated machines currently not available in virtualised environments.

Erlang Solutions has a dynamic stream of customer projects and the exact configurations will depend on the clients’ use cases, but already, prospective clients of SWARM include massively scalable NoSQL databases, multisite instant messaging systems and middleware for nation-wide trading infrastructures. Erlang Solutions is also currently developing ad-hoc networking solutions for mobile devices. We plan to explore the potential of Hosted Continuous Integration in all these areas, and derive from that experience what a common provisioning framework for heterogeneous clusters should look like.

Scalable Reliability for a Mainstream Actor Framework. The goal of this study is to evaluate how the scalable reliable concurrency oriented paradigm we will develop can be effectively applied to mainstream software development. We will investigate the feasibility and limitations of adding SD Erlang scalability constructs to a popular Actor framework for a dominant language. The investigation will focus on an open source framework with an active user base. One such framework that meets the requirements is Kilim for Java [10].

5 Progress So Far

In the first six months of the RELEASE project we have made the following progress [36].

We have started to design SD Erlang by making an overview of architecture trends, possible failures, and Erlang modules that might impinge on Erlang scalability [37]. We have formulated the main design principles, and analysed in-memory and persistent data storage mechanisms. We also have developed an initial SD Erlang design that includes implicit process placement and scalable groups to reduce node connections.

We have started to benchmark and trace multicore systems using both the built-in Erlang tracing and DTrace for Unix systems. We are also compiling a set of programs for Erlang and in particular Distributed Erlang, so that we can benchmark the performance of multicore SD Erlang on a set of practical problems [38].

We have produced a survey of the state-of-the-art cloud providers and management infrastructures. The survey is a foundation for the initial work on providing access to cloud infrastructure for different tasks, such as system building, testing and deployment. We have added a few useful features to Sim-Diasca to ease troubleshooting, e.g. a distributed instance tracker. We are also working on a scalable simulation case that can be used for benchmarking purposes.

6 Conclusion

The RELEASE project intends to scale Erlang's concurrency-oriented programming paradigm to build reliable general-purpose software, such as server-based systems, on massively parallel commodity hardware. We outline the state of the art (Section 2), and then explain how we plan to tackle the challenge at the following three levels. We are evolving the Erlang VM for large scale multicore systems (Section 3.1). We are evolving the language to Scalable Distributed (SD) Erlang, and adapting the OTP framework to provide both constructs like locality control, and reusable coordination patterns (Section 3.2). We are developing a scalable Erlang infrastructure to integrate multiple, heterogeneous clusters (Section 3.3). To support these activities we are developing tools to enable programmers to understand the behaviour of large scale Erlang systems, and to refactor standard Erlang programs into SD Erlang (Section 3.4). We have outlined how case studies will be used to investigate and validate the new reliable massively parallel tools and methodologies in practice (Section 4). Finally, we have briefly outlined our initial progress (Section 5). We look forward to reporting on the progress of the project in the years to come.

7 Acknowledgements

We would like to thank all our colleagues who work on the RELEASE project. This work has been supported by the European Union grant RII3-CT-2005-

026133 'SCIENCE: Symbolic Computing Infrastructure in Europe', IST-2011-287510 'RELEASE: A High-Level Paradigm for Reliable Large-scale Server Software', and by the UK's Engineering and Physical Sciences Research Council grant EP/G055181/1 'HPC-GAP: High Performance Computational Algebra and Discrete Mathematics'.

References

1. Trew, A.: Parallelism and the exascale challenge. Distinguished Lecture. St Andrews University, St Andrews, UK (2010)
2. Ericsson AB: Erlang/OTP Efficiency Guide, System Limits. (2011) erlang.org/doc/efficiency_guide/advanced.html#id67011.
3. Cesarini, F., Thompson, S.: Erlang Programming. O'Reilly (2009)
4. Agha, G.: An overview of Actor languages. SIGPLAN Not. **21**(10) (1986) 58–67
5. IBM Research: IBM Research Blue Gene Project (2012) www.research.ibm.com/bluegene/index.html.
6. Hewitt, C., Bishop, P., Steiger, R.: A universal modular ACTOR formalism for artificial intelligence. In: IJCAI'73, San Francisco, CA, USA, Morgan Kaufmann (1973) 235–245
7. Epstein, J., Black, A.P., Peyton-Jones, S.: Towards Haskell in the Cloud. In: Haskell '11, ACM (2011) 118–129
8. Odersky, M., et al.: The Scala programming language (2012) www.scala-lang.org.
9. Syme, D., Granicz, A., Cisternino, A.: Expert F#. Springer (2007)
10. Srinivasan, S., Mycroft, A.: Kilim: Isolation-typed actors for Java. In: ECOOP'08, Berlin, Heidelberg, Springer-Verlag (2008) 104–128
11. Intel Corporation: Tera-scale computing architecture overview (2012) www.intel.com/content/www/us/en/research/intel-labs-terascale-computing-demo.html.
12. Marr, S., Haupt, M., Timbermont, S., Adams, B., D'Hondt, T., Costanza, P., Meuter, W.D.: Virtual machine support for many-core architectures: Decoupling abstract from concrete concurrency models. In: PLACES. (2009) 63–77
13. Noll, A., Gal, A., Franz, M.: CellVM: A homogeneous virtual machine runtime system for a heterogeneous single-chip multiprocessor. In: Workshop on Cell Systems and Applications. (2009)
14. Ungar, D., Adams, S.S.: Hosting an object heap on manycore hardware: An exploration. SIGPLAN Not. **44**(12) (2009) 99–110
15. Leung, A., Lhoták, O., Lashari, G.: Automatic parallelization for graphics processing units. In: PPPJ'09, New York, NY, USA, ACM (2009) 91–100
16. Seiler, L., Carmean, D., Sprangle, E., Forsyth, T., Abrash, M., Dubey, P., Junkins, S., Lake, A., Sugeran, J., Cavin, R., Espasa, R., Grochowski, E., Juan, T., Hanrahan, P.: Larrabee: A many-core x86 architecture for visual computing. ACM Trans. Graph. **27**(3) (2008) 18:1–18:15
17. Karmani, R.K., Shali, A., Agha, G.: Actor frameworks for the JVM platform: A comparative analysis. In: PPPJ '09, New York, NY, USA, ACM (2009) 11–20
18. Sagonas, K., Wilhelmsson, J.: Efficient memory management for concurrent programs that use message passing. Sci. Comput. Program. **62**(2) (2006) 98–121
19. Chandra, R., Dagum, L., Kohr, D., Maydan, D., McDonald, J., Menon, R.: Parallel Programming in OpenMP. Morgan Kaufmann, San Francisco, CA, USA (2001)

20. Snir, M., Otto, S.W., Walker, D.W., Dongarra, J., Huss-Lederman, S.: MPI: The Complete Reference. MIT Press, Cambridge, MA, USA (1995)
21. Dewolfs, D., Broeckhove, J., Sunderam, V., Fagg, G.E.: FT-MPI, fault-tolerant metacomputing and generic name services: A case study. In: EuroPVM/MPI'06, Berlin, Heidelberg, Springer-Verlag (2006) 133–140
22. Odersky, M., Altherr, P., Cremet, V.: The Scala language specification. Technical report, EFPL, Lausanne, Switzerland (April 2004)
23. Rajan, H., Leavens, G.T.: Ptolemy: A language with quantified, typed events. In: ECOOP '08, Berlin, Heidelberg, Springer-Verlag (2008) 155–179
24. Germain, G.: Concurrency oriented programming in termite scheme. In: ER-LANG'06, New York, NY, USA, ACM (2006) 20–20
25. Lee, J., et al.: Python actor runtime library (2010) osl.cs.uiui.edu/parley/.
26. Wiger, U.: What is Erlang-style concurrency? (2010) www.ulf.wiger.net/weblog/2008/02/06/what-is-erlang-style-concurrency/.
27. Intel Software Network: Intel trace analyzer and collector (2011) www.software.intel.com/en-us/articles/intel-trace-analyzer/.
28. Erlang Online Documentation: Trace tool builder (2011) www.erlang.org/doc/apps/observer/ttb_ug.html.
29. Zhao, J., Madduri, S., Vadlamani, R., Bureson, W., Tessier, R.: A dedicated monitoring infrastructure for multicore processors. Very Large Scale Integration (VLSI) Systems, IEEE Transactions on **19**(6) (2011) 1011–1022
30. The ProTest Project: Framework 7 project 215868 (2008-11) www.protest-project.eu.
31. Arts, T., Fredlund, L.A.: Trace analysis of Erlang programs. SIGPLAN Not. **37**(12) (2002) 18–24
32. Nystrom, J.: Analysing Fault Tolerance for Erlang Application. PhD thesis, Division of Computer Systems, Uppsala University (2009)
33. Thompson, S., et al.: Wrangler Tool (2011) www.cs.kent.ac.uk/projects/wrangler/.
34. Gartner Group: Gartner says Cloud consumers need brokerages to unlock the potential of Cloud services (2009) www.gartner.com/it/page.jsp?id=1064712.
35. Eucalyptus Systems: Cloud APIs (2010) www.eucalyptus.com/blog/2010/03/11/cloud-apis.
36. The RELEASE Project: Framework 7 project 287510 (2011-14) www.release-project.eu.
37. Chechina, N., Trinder, P., Ghaffari, A., Green, R., Lundin, K., Virding, R.: The design of scalable distributed (SD) Erlang. In: Draft Proceedings of IFL '12, Oxford University (2012) 461–476
38. Aronis, S., et al.: BenchErl: A Scalability Benchmark Suite for Erlang/OTP (2012) release.softlab.ntua.gr/bencherl/.