

Kent Academic Repository

Full text document (pdf)

Citation for published version

Thompson, Simon and Li, Huiqing and Baker, Robert and Adams, Stephen and Trinder, Phil and Chechina, Natalia and Papaspyrou, Nikolaos and Sagonas, Konstantinos and Aronis, Stavros and Bihari, Eva and Boudeville, Olivier and Cesarini, Francesco and de Stefano, Maurizio and Eriksson, Sverker and Fordos, Viktoria and Ghaffari, Amir and Green, Rickard and Hoch, csaba

DOI

<https://doi.org/10.1145/3107937>

Link to record in KAR

<http://kar.kent.ac.uk/63197/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Scaling Reliably: Improving the Scalability of the Erlang Distributed Actor Platform

PHIL TRINDER, NATALIA CHECHINA, NIKOLAOS PAPASPYROU,
KONSTANTINOS SAGONAS, SIMON THOMPSON, STEPHEN ADAMS, STAVROS ARONIS,
ROBERT BAKER, EVA BIHARI, OLIVIER BOUDEVILLE, FRANCESCO CESARINI,
MAURIZIO DI STEFANO, SVERKER ERIKSSON, VIKTORIA FORDOS, AMIR GHAFARI,
AGGELOS GIANTSIOS, ROCKARD GREEN, CSABA HOCH, DAVID KLAFTENEGGER,
HUIQING LI, KENNETH LUNDIN, KENNETH MACKENZIE, KATERINA ROUKOUNAKI,
YIANNIS TSIOURIS, KJELL WINBLAD

RELEASE EU FP7 STREP (287510) project
www.release-project.eu/

April 25, 2017

Abstract

Distributed actor languages are an effective means of constructing scalable reliable systems, and the Erlang programming language has a well-established and influential model. While the Erlang model conceptually provides reliable scalability, it has some inherent scalability limits and these force developers to depart from the model at scale. This article establishes the scalability limits of Erlang systems, and reports the work of the EU RELEASE project to improve the scalability and understandability of the Erlang reliable distributed actor model.

We systematically study the scalability limits of Erlang, and then address the issues at the virtual machine, language and tool levels. More specifically: (1) We have evolved the Erlang virtual machine so that it can work effectively in large scale single-host multicore and NUMA architectures. We have made important changes and architectural improvements to the widely used Erlang/OTP release. (2) We have designed and implemented Scalable Distributed (SD) Erlang libraries to address language-level scalability issues, and provided and validated a set of semantics for the new language constructs. (3) To make large Erlang systems easier to deploy, monitor, and debug we have developed and made open source releases of five complementary tools, some specific to SD Erlang.

Throughout the article we use two case studies to investigate the capabilities of our new technologies and tools: a distributed hash table based Orbit calculation and Ant Colony Optimisation (ACO). Chaos Monkey experiments show that two versions of ACO survive random process failure and hence that SD Erlang preserves the Erlang reliability model. While we report measurements on a range of NUMA and cluster architectures, the key scalability experiments are conducted on the Athos cluster with 256 hosts (6144 cores). Even for programs with no global recovery data to maintain, SD Erlang partitions the network to reduce network traffic and hence improves performance of the Orbit and ACO benchmarks above 80 hosts. ACO measurements show that maintaining global recovery data dramatically limits scalability; however scalability is recovered by partitioning the recovery data. We exceed the established scalability limits of distributed Erlang, and do not reach the limits of SD Erlang for these benchmarks at this scale (256 hosts, 6144 cores).

I. INTRODUCTION

Distributed programming languages and frameworks are central to engineering large scale systems, where key properties include scalability and reliability. By scalability we mean that performance increases as hosts and cores are added, and by large scale we mean architectures with hundreds of hosts and tens of thousands of cores. Experience with high performance and data centre computing shows that reliability is critical at these scales, e.g. host failures alone account for around one failure per hour on commodity servers with approximately 10^5 cores [11]. To be usable, programming languages employed on them must be supported by a suite of deployment, monitoring, refactoring and testing tools that work at scale.

Controlling shared state is the only way to build reliable scalable systems. State shared by multiple units of computation limits scalability due to high synchronisation and communication costs. Moreover shared state is a threat for reliability as failures corrupting or permanently locking shared state may poison the entire system.

Actor languages avoid shared state: actors or processes have entirely local state, and only interact with each other by sending messages [2]. Recovery is facilitated in this model, since actors, like operating system processes, can fail independently without affecting the state of other actors. Moreover an actor can *supervise* other actors, detecting failures and taking remedial action, e.g. restarting the failed actor.

Erlang [6, 19] is a beacon language for reliable scalable computing with a widely emulated distributed actor model. It has influenced the design of numerous programming languages like Clojure [43] and F# [74], and many languages have Erlang-inspired actor frameworks, e.g. Kilim for Java [73], Cloud Haskell [31], and Akka for C#, F# and Scala [64]. Erlang is widely used for build-

ing reliable scalable servers, e.g. Ericsson's AXD301 telephone exchange (switch) [79], the Facebook chat server, and the Whatsapp instant messaging server [77].

In Erlang, the actors are termed processes and are managed by a sophisticated Virtual Machine on a single multicore or NUMA host, while *distributed Erlang* provides relatively transparent distribution over networks of VMs on multiple hosts. Erlang is supported by the Open Telecom Platform (OTP) libraries that capture common patterns of reliable distributed computation, such as the client-server pattern and process supervision. Any large-scale system needs scalable persistent storage and, following the CAP theorem [39], Erlang uses and indeed implements Dynamo-style NoSQL DBMS like Riak [49] and Cassandra [50].

While the Erlang distributed actor model conceptually provides reliable scalability, it has some inherent scalability limits, and indeed large-scale distributed Erlang systems must depart from the distributed Erlang paradigm in order to scale, e.g. not maintaining a fully connected graph of hosts. The EU FP7 RELEASE project set out to establish and address the scalability limits of the Erlang reliable distributed actor model [67].

After outlining related work (Section II) and the benchmarks used throughout the article (Section III) we investigate the scalability limits of Erlang/OTP, seeking to identify specific issues at the virtual machine, language and persistent storage levels (Section IV).

We then report the RELEASE project work to address these issues, working at the following three levels.

1. We have designed and implemented a set of Scalable Distributed (SD) Erlang libraries to address *language-level* reliability and scalability issues. An operational semantics is provided for the key new `s_group` construct, and the implementa-

tion is validated against the semantics (Section V).

2. We have *evolved the Erlang virtual machine* so that it can work effectively in large-scale single-host multicore and NUMA architectures. We have improved the shared ETS tables, time management, and load balancing between schedulers. Most of these improvements are now included in the Erlang/OTP release, currently downloaded approximately 50K times each month (Section VI).
3. To facilitate the development of scalable Erlang systems, and to make them maintainable, we have developed three new tools: Devo, SDMon and WombatOAM, and enhanced two others: the visualisation tool Percept, and the refactorer Wrangler. The tools support refactoring programs to make them more scalable, easier to deploy at large scale (hundreds of hosts), easier to monitor and visualise their behaviour. Most of these tools are freely available under open source licences; the WombatOAM deployment and monitoring tool is a commercial product (Section VII).

Throughout the article we use two benchmarks to investigate the capabilities of our new technologies and tools. These are a computation in symbolic algebra, more specifically an algebraic ‘orbit’ calculation that exploits a non-replicated distributed hash table, and an Ant Colony Optimisation (ACO) parallel search program (Section III).

We report on the reliability and scalability implications of our new technologies using Orbit, ACO, and other benchmarks. We use a Chaos Monkey instance [14] that randomly kills processes in the running system to demonstrate the reliability of the benchmarks and to show that SD Erlang preserves the Erlang language-level reliability model. While we

report measurements on a range of NUMA and cluster architectures as specified in Appendix A, the key scalability experiments are conducted on the Athos cluster with 256 hosts and 6144 cores. Having established scientifically the folklore limitations of around 60 connected hosts/nodes for distributed Erlang systems in Section 4, a key result is to show that the SD Erlang benchmarks exceed this limit and do not reach their limits on the Athos cluster (Section VIII).

Contributions. This article is the first systematic presentation of the coherent set of technologies for engineering scalable reliable Erlang systems developed in the RELEASE project.

Section IV presents the first scalability study covering Erlang VM, language, and storage scalability. Indeed we believe it is the first comprehensive study of any distributed actor language at this scale (100s of hosts, and around 10K cores). Individual scalability studies, e.g. into Erlang VM scaling [8], or language and storage scaling have appeared before [38, 37].

At the language level the design, implementation and validation of the new libraries (Section V) have been reported piecemeal [21, 60], and are included here for completeness.

While some of the improvements made to the Erlang Virtual Machine (Section i) have been thoroughly reported in conference publications [66, 46, 47, 69, 70, 71], others are reported here for the first time (Sections iii, ii).

In Section VII, the WombatOAM and SDMon tools are described for the first time, as is the revised Devo system and visualisation. The other tools for profiling, debugging and refactoring developed in the project have previously been published piecemeal [52, 53, 55, 10], but this is their first unified presentation.

All of the performance results in Section VIII are entirely new, although a comprehensive study of SD Erlang performance is now available in a recent article by [22].

II. CONTEXT

i. Scalable Reliable Programming Models

There is a plethora of shared memory concurrent programming models like PThreads or Java threads, and some models, like OpenMP [20], are simple and high level. However synchronisation costs mean that these models generally do not scale well, often struggling to exploit even 100 cores. Moreover, reliability mechanisms are greatly hampered by the shared state: for example, a lock becomes permanently unavailable if the thread holding it fails.

The High Performance Computing (HPC) community build large-scale (10^6 core) distributed memory systems using the *de facto* standard MPI communication libraries [72]. Increasingly these are hybrid applications that combine MPI with OpenMP. Unfortunately, MPI is not suitable for producing general purpose concurrent software as it is too low level with explicit message passing. Moreover, the most widely used MPI implementations offer no fault recovery:¹ if any part of the computation fails, the entire computation fails. Currently the issue is addressed by using what is hoped to be highly reliable computational and networking hardware, but there is intense research interest in introducing reliability into HPC applications [33].

Server farms use commodity computational and networking hardware, and often scale to around 10^5 cores, where host failures are routine. They typically perform rather constrained computations, e.g. Big Data Analytics, using reliable frameworks like Google MapReduce [26] or Hadoop [78]. The idempotent nature of the analytical queries makes it relatively easy for the frameworks to provide implicit reliability: queries are monitored and failed queries are

¹Some fault tolerance is provided in less widely used MPI implementations like [27].

simply re-run. In contrast, actor languages like Erlang are used to engineer reliable general purpose computation, often recovering failed stateful computations.

ii. Actor Languages

The actor model of concurrency consists of independent processes communicating by means of messages sent asynchronously between processes. A process can send a message to any other process for which it has the address (in Erlang the “process identifier” or pid), and the remote process may reside on a different host. While the notion of actors originated in AI [42], it has been used widely as a general metaphor for concurrency, as well as being incorporated into a number of niche programming languages in the 1970s and 80s. More recently it has come back to prominence through the rise of not only multicore chips but also larger-scale distributed programming in data centres and the cloud.

With built-in concurrency and data isolation, actors are a natural paradigm for engineering reliable scalable general-purpose systems [1, 41]. The model has two main concepts: actors that are the unit of computation, and messages that are the unit of communication. Each actor has an *address-book* that contains the addresses of all the other actors it is aware of. These addresses can be either locations in memory, direct physical attachments, or network addresses. In a pure actor language, messages are the only way for actors to communicate.

After receiving a message an actor can do the following: (i) send messages to another actor from its *address-book*, (ii) create new actors, or (iii) designate a behaviour to handle the next message it receives. The model does not impose any restrictions in the order in which these actions must be taken. Similarly, two messages sent concurrently can be received in any order. These features enable actor based systems to support indeterminacy and quasi-

commutativity, while providing locality, modularity, reliability and scalability [41].

Actors are just one message-based paradigm, and other languages and libraries have related message passing paradigms. Recent example languages include Go [28] and Rust [62] that provide explicit channels, similar to actor mailboxes. Probably the most famous message passing library is MPI [72], with APIs for many languages and widely used on clusters and High Performance Computers. It is, however, arguable that the most important contribution of the actor model is the one-way asynchronous communication [41]. Messages are not coupled with the sender, and neither they are transferred synchronously to a temporary container where transmission takes place, e.g. a buffer, a queue, or a mailbox. Once a message is sent, the receiver is the only entity responsible for that message.

Erlang [6, 19] is the pre-eminent programming language based on the actor model, having a history of use in production systems, initially with its developer Ericsson and then more widely through open source adoption. There are now actor frameworks for many other languages; these include Akka for C#, F# and Scala [64], CAF² for C++, Pykka³, Cloud Haskell [31], PARLEY [51] for Python, and Termit Scheme [35], and each of these is currently under active use and development. Moreover, the recently defined Rust language [62] has a version of the actor model built in, albeit in an imperative context.

iii. Erlang’s Support for Concurrency

In Erlang, actors are termed *processes*, and virtual machines are termed *nodes*. The key elements of the actor model are: fast process creation and destruction; lightweight processes, e.g. enabling 10⁶ concurrent processes on a single host with 8GB RAM; fast asynchronous

message passing with copying semantics; process monitoring; strong dynamic typing, and selective message reception.

By default Erlang processes are addressed by their *process identifier* (*pid*), e.g.

```
Pong_PID = spawn(fun some_module:pong/0)
```

spawns a process to execute the anonymous function given as argument to the `spawn` primitive, and binds `Pong_PID` to the new process identifier. Here the new process will execute the `pong/0` function which is defined in `some_module`. A subsequent call

```
Pong_PID ! finish
```

sends the message `finish` to the process identified by `Pong_PID`. Alternatively, processes can be given names using a call of the form:

```
register(my_funky_name, Pong_PID)
```

which registers this process name in the node’s *process name table* if not already present. Subsequently, these names can be used to refer to or communicate with the corresponding processes (e.g. send them a message):

```
my_funky_process ! hello.
```

A *distributed Erlang* system executes on multiple nodes, and the nodes can be freely deployed across hosts, e.g. they can be located on the same or different hosts. To help make distribution transparent to the programmer, when any two nodes connect they do so transitively, sharing their sets of connections. Without considerable care this quickly leads to a fully connected graph of nodes. A process may be spawned on an explicitly identified node, e.g.

```
Remote_Pong_PID = spawn(some_node,
                        fun some_module:pong/0).
```

After this, the remote process can be addressed just as if it were local. It is a significant burden on the programmer to identify the remote nodes in large systems, and we will return to this in Sections ii, i.2.

²<http://actor-framework.org>

³<http://pykka.readthedocs.org/en/latest/>




Figure 1: *Conceptual view of Erlang's concurrency, multicore support and distribution.*

iv. Scalability and Reliability in Erlang Systems

Erlang was designed to solve a particular set of problems, namely those in building telecommunications' infrastructure, where systems need to be scalable to accommodate hundreds of thousands of calls concurrently, in soft real-time. These systems need to be highly-available and reliable: i.e. to be robust in the case of failure, which can come from software or hardware faults. Given the inevitability of the latter, Erlang adopts the "let it fail" philosophy for error handling. That is, encourage programmers to embrace the fact that a process may fail at any point, and have them rely on the supervision mechanism, discussed shortly, to handle the failures.

Figure 1 illustrates Erlang's support for concurrency, multicores and distribution. Each Erlang node is represented by a yellow shape, and each rectangle represents a *host* with an IP address. Each red arc represents a *connection* between Erlang nodes. Each node can run on multiple cores, and exploit the inherent concurrency provided. This is done automatically by the VM, with no user intervention needed. Typically each core has an associated *scheduler* that schedules processes; a new process will be spawned on the same core as the process that spawns it, but work can be moved to a different scheduler through a work-stealing allocation algorithm. Each scheduler allows a process

that is ready to compute at most a fixed number of computation steps before switching to another. Erlang *built-in functions* or BIFs are implemented in C, and at the start of the project were run to completion once scheduled, causing performance and responsiveness problems if the BIF had a long execution time.

Scaling in Erlang is provided in two different ways. It is possible to scale *within a single node* by means of the multicore virtual machine exploiting the concurrency provided by the multiple cores or NUMA nodes. It is also possible to scale *across multiple hosts* using multiple distributed Erlang nodes.

Reliability in Erlang is multi-faceted. As in all actor languages each process has private state, preventing a failed or failing process from corrupting the state of other processes. Messages enable stateful interaction, and contain a deep copy of the value to be shared, with no references (e.g. pointers) to the senders' internal state. Moreover Erlang avoids type errors by enforcing strong typing, albeit dynamically [7]. Connected nodes check liveness with heartbeats, and can be monitored from outside Erlang, e.g. by an operating system process.

However, the most important way to achieve reliability is *supervision*, which allows a process to monitor the status of a child process and react to any failure, for example by spawning a substitute process to replace a failed process. Supervised processes can in turn supervise other processes, leading to a *supervision tree*. The supervising and supervised processes may be in different nodes, and on different hosts, and hence the supervision tree may span multiple hosts or nodes.

To provide reliable distributed service registration, a global namespace is maintained on every node, which maps *process names* to pids. It is this that we mean when we talk about a 'reliable' system: it is one in which a named process in a distributed system can be restarted without requiring the client processes also to

be restarted (because the name can still be used for communication).

To see global registration in action, consider a pong server process

```
global:register_name(pong_server,
                    Remote_Pong_PID).
```

Clients of the server can send messages to the registered name, e.g.

```
global:whereis_name(pong_server)!finish.
```

If the server fails the supervisor can spawn a replacement server process with a new pid and register it with the same name (`pong_server`). Thereafter client messages to the `pong_server` will be delivered to the new server process. We return to discuss the scalability limitations of maintaining a global namespace in Section V.

v. ETS: Erlang Term Storage

Erlang is a pragmatic language and the actor model it supports is not pure. Erlang processes, besides communicating via asynchronous message passing, can also share data in public memory areas called *ETS tables*.

The Erlang Term Storage (ETS) mechanism is a central component of Erlang’s implementation. It is used internally by many libraries and underlies the in-memory databases. ETS tables are key-value stores: they store Erlang tuples where one of the positions in the tuple serves as the lookup key. An ETS table has a *type* that may be either `set`, `bag` or `duplicate_bag`, implemented as a hash table, or `ordered_set` which is currently implemented as an AVL tree. The main operations that ETS supports are table creation, insertion of individual entries and atomic bulk insertion of multiple entries in a table, deletion and lookup of an entry based on some key, and destructive update. The operations are implemented as C built-in functions in the Erlang VM.

The code snippet below creates a `set` ETS table keyed by the first element of the entry; atomically inserts two elements with keys

`some_key` and `42`; updates the value associated with the table entry with key `42`; and then looks up this entry.

```
Table = ets:new(my_table,
                [set, public, {keypos, 1}]),
ets:insert(Table, [
                {some_key, an_atom_value},
                {42, {a,tuple,value}}]),
ets:update_element(Table, 42,
                  [{2, {another,tuple,value}}]),
[Key, Value] = ets:lookup(Table, 42).
```

ETS tables are heavily used in many Erlang applications. This is partly due to the convenience of sharing data for some programming tasks, but also partly due to their fast implementation. As a shared resource, however, ETS tables induce contention and become a scalability bottleneck, as we shall see in Section i.

III. BENCHMARKS FOR SCALABILITY AND RELIABILITY

The two benchmarks that we use throughout this article are Orbit, that measures scalability without looking at reliability, and Ant Colony Optimisation (ACO) that allows us to measure the impact on scalability of adding global namespaces to ensure reliability. The source code for the benchmarks, together with more detailed documentation, is available at <https://github.com/release-project/benchmarks/>. The RELEASE project team also worked to improve the reliability and scalability of other Erlang programs including a substantial (approximately 150K lines of Erlang code) SimDiasca simulator [16] and an Instant Messenger that is more typical of Erlang applications [23] but we do not cover these systematically here.

i. Orbit

Orbit is a computation in symbolic algebra, which generalises a transitive closure compu-




Figure 2: *Distributed Erlang Orbit (D-Orbit) architecture: workers are mapped to nodes.*

tation [57]. To compute the *orbit* for a given space $[0..X]$, a list of generators g_1, g_2, \dots, g_n are applied on an initial vertex $x_0 \in [0..X]$. This creates new values $(x_1 \dots x_n) \in [0..X]$, where $x_i = g_i(x_0)$. The generator functions are applied on the new values until no new value is generated.

Orbit is a suitable benchmark because it has a number of aspects that characterise a class of real applications. The core data structure it maintains is a set and, in distributed environments is implemented as a *distributed hash table* (DHT), similar to the DHTs used in replicated form in NoSQL database management systems. Also, in distributed mode, it uses standard peer-to-peer (P2P) techniques like a credit-based termination algorithm [61]. By choosing the orbit size, the benchmark can be parameterised to specify smaller or larger computations that are suitable to run on a single machine (Section i) or on many nodes (Section ii). Moreover it is only a few hundred lines of code.

As shown in Figure 2, the computation is initiated by a master which creates a number of workers. In the single node scenario of the benchmark, workers correspond to processes but these workers can also spawn other processes to apply the generator functions on a subset of their input values, thus creating *intra-worker parallelism*. In the distributed version of the benchmark, processes are spawned by the master node to worker nodes, each maintaining a DHT fragment. A newly spawned

process gets a share of the parent’s credit, and returns this on completion. The computation is finished when the master node collects all credit, i.e. all workers have completed.

ii. Ant Colony Optimisation (ACO)

Ant Colony Optimisation [29] is a meta-heuristic which has been applied to a large number of combinatorial optimisation problems. For the purpose of this article, we have applied it to an NP-hard scheduling problem known as the Single Machine Total Weighted Tardiness Problem (SMTWTP) [63], where a number of jobs of given lengths have to be arranged in a single linear schedule. The goal is to minimise the *cost* of the schedule, as determined by certain constraints.

The ACO method is attractive from the point of view of distributed computing because it can benefit from having multiple cooperating colonies, each running on a separate compute node and consisting of multiple “ants”. Ants are simple computational agents which concurrently compute possible solutions to the input problem guided by shared information about good paths through the search space; there is also a certain amount of stochastic variation which allows the ants to explore new directions. Having multiple colonies increases the number of ants, thus increasing the probability of finding a good solution.

We implement four distributed coordination patterns for the same multi-colony ACO com-




Figure 3: *Distributed Erlang Two-level Ant Colony Optimisation (TL-ACO) architecture.*

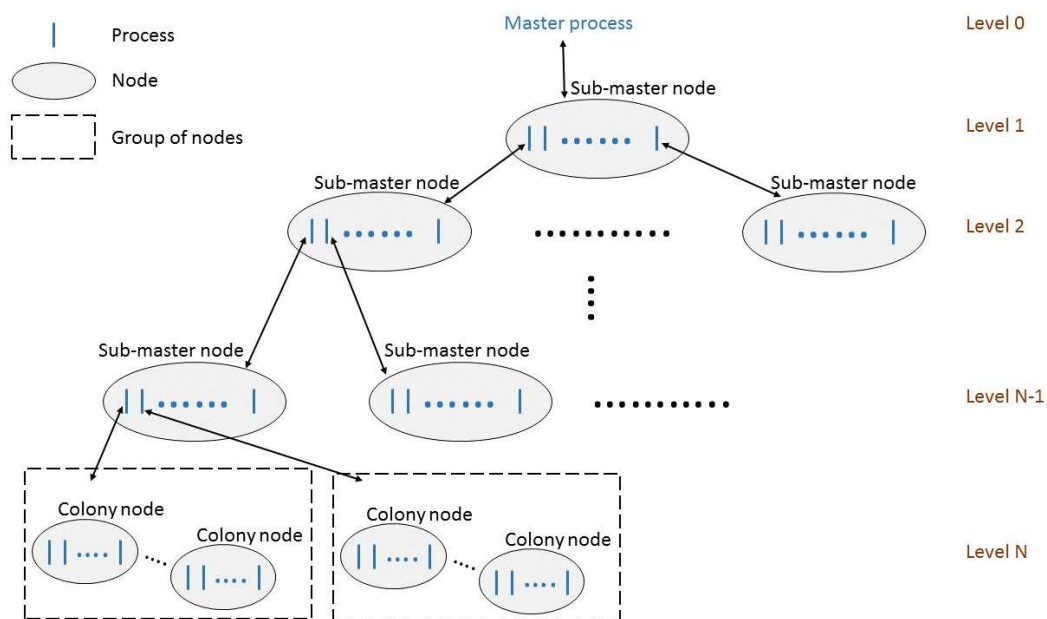


Figure 4: Distributed Erlang Multi-level Ant Colony Optimisation (ML-ACO) architecture.

putation as follows. In each implementation, the individual colonies perform some number of *local iterations* (i.e. generations of ants) and then report their best solutions; the globally-best solution is then selected and is reported to the colonies, which use it to update their pheromone matrices. This process is repeated for some number of *global iterations*.

Two-level ACO (TL-ACO) has a single master node that collects the colonies' best solutions and distributes the overall best solution back to the colonies. Figure 3 depicts the process and node placements of the TL-ACO in a cluster with N_C nodes. The master process spawns N_C colony processes on available nodes. In the next step, each colony process spawns N_A ant processes on the local node. Each ant iterates I_A times, returning its result to the colony master. Each colony iterates I_M times, reporting their best solution to, and receiving the globally-best solution from, the master process. We validated the implementation by applying TL-ACO to a number of standard SMTWTP

instances [25, 13, 34], obtaining good results in all cases, and confirmed that the number of perfect solutions increases as we increase the number of colonies.

Multi-level ACO (ML-ACO). In TL-ACO the master node receives messages from *all* of the colonies, and thus could become a bottleneck. ML-ACO addresses this by having a tree of submasters (Figure 4), with each node in the bottom level collecting results from a small number of colonies. These are then fed up through the tree, with nodes at higher levels selecting the best solutions from their children.

Globally Reliable ACO (GR-ACO). In ML-ACO if a single colony fails to report back the system will wait indefinitely. GR-ACO adds fault tolerance, supervising colonies so that a faulty colony can be detected and restarted, allowing the system to continue execution.

Scalable Reliable ACO (SR-ACO) also adds fault-tolerance, but using supervision within our new `s_groups` from Section i.1, and the architecture of SR-ACO is discussed in detail

there.

IV. ERLANG SCALABILITY LIMITS

This section investigates the scalability of Erlang at VM, language, and persistent storage levels. An aspect we choose not to explore is the security of large scale systems where, for example, one might imagine providing enhanced security for systems with multiple clusters or cloud instances connected by a Wide Area Network. We assume that existing security mechanisms are used, e.g. a Virtual Private Network.

i. Scaling Erlang on a Single Host

To investigate Erlang scalability we built BenchErl, an extensible open source benchmark suite with a web interface.⁴ BenchErl shows how an application’s performance changes when resources, like cores or schedulers, are added; or when options that control these resources change:

- the *number of nodes*, i.e. the number of Erlang VMs used, typically on multiple hosts;
- the *number of cores* per node;
- the *number of schedulers*, i.e. the OS threads that execute Erlang processes in parallel, and their *binding* to the topology of the cores of the underlying computer node;
- the *Erlang/OTP release and flavor*; and
- the *command-line arguments* used to start the Erlang nodes.

Using BenchErl, we investigated the scalability of an initial set of twelve benchmarks and two substantial Erlang applications using a single Erlang node (VM) on machines

⁴Information about BenchErl is available at <http://release.softlab.ntua.gr/bencher1/>.

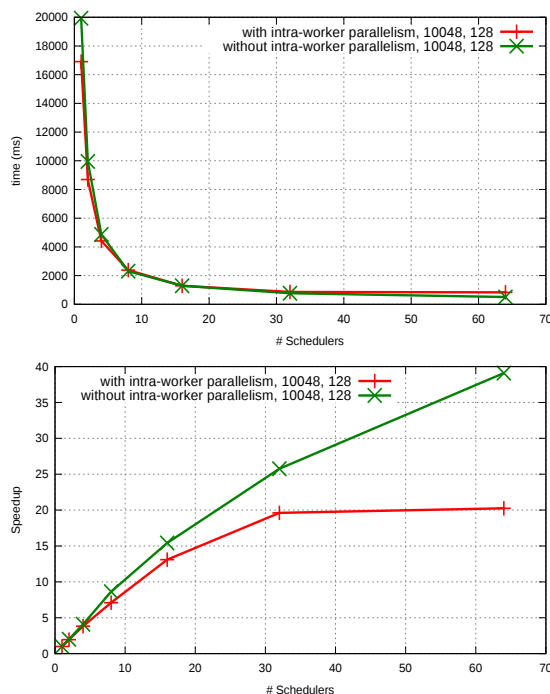


Figure 5: Runtime and speedup of two configurations of the Orbit benchmark using Erlang/OTP R15B01.

with up to 64 cores, including the Bulldozer machine specified in Appendix A. This set of experiments, reported by [8], confirmed that some programs scaled well in the most recent Erlang/OTP release of the time (R15B01) but also revealed VM and language level scalability bottlenecks.

Figure 5 shows runtime and speedup curves for the Orbit benchmark where master and workers run on a single Erlang node in configurations with and without intra-worker parallelism. In both configurations the program scales. Runtime continuously decreases as we add more schedulers to exploit more cores. The speedup of the benchmark without intra-worker parallelism, i.e. without spawning additional processes for the computation (green curve), is almost linear up to 32 cores but in-

creases less rapidly from that point on; we see a similar but more clearly visible pattern for the other configuration (red curve) where there is no performance improvement beyond 32 schedulers. This is due to the asymmetric characteristics of the machine’s micro-architecture, which consists of modules that couple two conventional x86 out-of-order cores that share the early pipeline stages, the floating point unit, and the L2 cache with the rest of the module [3].

Some other benchmarks, however, did not scale well or experienced significant slowdowns when run in many VM schedulers (threads). For example the `ets_test` benchmark has multiple processes accessing a shared ETS table. Figure 6 shows runtime and speedup curves for `ets_test` on a 16-core (eight cores with hyperthreading) Intel Xeon-based machine. It shows that runtime increases beyond two schedulers, and that the program exhibits a slowdown instead of a speedup.

For many benchmarks there are obvious reasons for poor scaling like limited parallelism in the application, or contention for shared resources. The reasons for poor scaling are less obvious for other benchmarks, and it is exactly these we have chosen to study in detail in subsequent work [8, 46].

A simple example is the parallel BenchErl benchmark, that spawns some n processes, each of which creates a list of m timestamps and, after it checks that each timestamp in the list is strictly greater than the previous one, sends the result to its parent. Figure 7 shows that up to eight cores each additional core leads to a slowdown, thereafter a small speedup is obtained up to 32 cores, and then again a slowdown. A small aspect of the benchmark, easily overlooked, explains the poor scalability. The benchmark creates timestamps using the `erlang:now/0` built-in function, whose implementation acquires a *global lock* in order to return a *unique timestamp*. That is, two calls to `erlang:now/0`, even from different processes




Figure 6: Runtime and speedup of the `ets_test` BenchErl benchmark using Erlang/OTP R15B01.

are guaranteed to produce monotonically increasing values. This lock is precisely the bottleneck in the VM that limits the scalability of this benchmark. We describe our work to address VM timing issues in Figure iii.

Discussion Our investigations identified contention for shared ETS tables, and for commonly-used shared resources like timers, as the key VM-level scalability issues. Section VI outlines how we addressed these issues in recent Erlang/OTP releases.

ii. Distributed Erlang Scalability

Network Connectivity Costs When any normal distributed Erlang nodes communicate, they share their connection sets and this typically leads to a fully connected graph of nodes.




Figure 7: Runtime and speedup of the BenchErl benchmark called parallel using Erlang/OTP R15B01.

So a system with n nodes will maintain $O(n^2)$ connections, and these are relatively expensive TCP connections with continual maintenance traffic. This design aids transparent distribution as there is no need to discover nodes, and the design works well for small numbers of nodes. However at emergent server architecture scales, i.e. hundreds of nodes, this design becomes very expensive and system architects must switch from the default Erlang model, e.g. they need to start using hidden nodes that do not share connection sets.

We have investigated the scalability limits imposed by network connectivity costs using several Orbit calculations on two large clusters: Kalkyl and Athos as specified in Appendix A. The Kalkyl results are discussed by [21], and Figure 29 in Section i shows representative

results for distributed Erlang computing orbits with 2M and 5M elements on Athos. In all cases performance degrades beyond some scale (40 nodes for the 5M orbit, and 140 nodes for the 2M orbit). Figure 32 illustrates the additional network traffic induced by the fully connected network. It allows a comparison between the number of packets sent in a fully connected network (ML-ACO) with those sent in a network partitioned using our new `s_groups` (SR-ACO).

Global Information Costs Maintaining global information is known to limit the scalability of distributed systems, and crucially the process namespaces used for reliability are global. To investigate the scalability limits imposed on distributed Erlang by such global information we have designed and implemented *DE-Bench*, an open source, parameterisable and scalable peer-to-peer benchmarking framework [37, 36]. DE-Bench measures the throughput and latency of distributed Erlang commands on a cluster of Erlang nodes, and the design is influenced by the Basho Bench benchmarking tool for Riak [12]. Each DE-Bench instance acts as a peer, providing scalability and reliability by eliminating central coordination and any single point of failure.

To evaluate the scalability of distributed Erlang, we measure how adding more hosts increases the throughput, i.e. the total number of successfully executed distributed Erlang commands per experiment. Figure 8 shows the parameterisable internal workflow of DE-Bench. There are three classes of commands in DE-Bench: (i) *Point-to-Point (P2P) commands*, where a function with tunable argument size and computation time is run on a remote node, include `spawn`, `rpc`, and synchronous calls to server processes, i.e. `gen_server` or `gen_fsm`. (ii) *Global commands*, which entail synchronisation across all connected nodes, such as `global:register_name` and

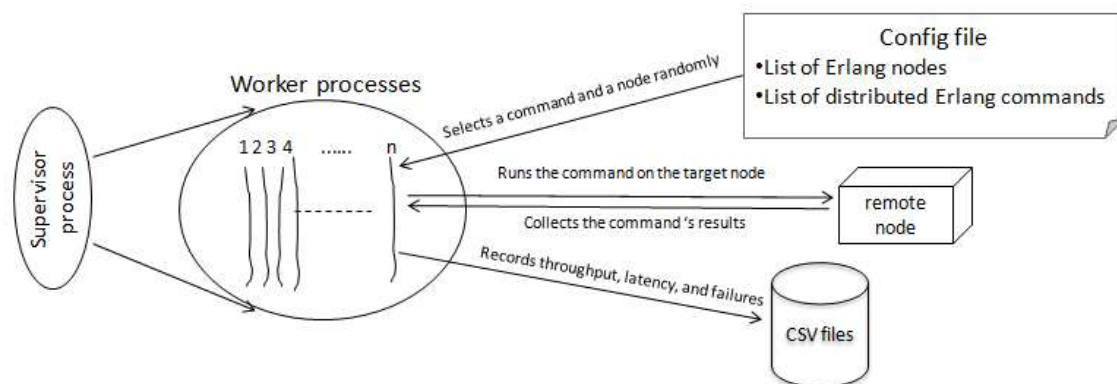


Figure 8: DE-Bench's internal workflow.

global:unregister_name. (iii) *Local commands*, which are executed independently by a single node, e.g. `whereis_name`, a look up in the local name table.

The benchmarking is conducted on 10 to 100 host configurations of the Kalkyl cluster (in steps of 10) and measures the throughput of successful commands per second over 5 minutes. There is one Erlang VM on each host and one DE-Bench instance on each VM. The full paper [37] investigates the impact of data size, and computation time in P2P calls both independently and in combination, and the scaling properties of the common Erlang/OTP generic server processes `gen_server` and `gen_fsm`.

Here we focus on the impact of different proportions of global commands, mixing global with P2P and local commands. Figure 9 shows that even a low proportion of global commands limits the scalability of distributed Erlang, e.g. just 0.01% global commands limits scalability to around 60 nodes. Figure 10 reports the latency of all commands and shows that, while the latencies for P2P and local commands are stable at scale, the latency of the global commands increases dramatically with scale. Both results illustrate that the impact of global operations on throughput and latency in a distributed Erlang system is severe.

Explicit Placement While network connectivity and global information impact *performance* at scale, our investigations also identified explicit process placement as a *programming* issue at scale. Recall from Section iii that distributed Erlang requires the programmer to identify an explicit Erlang node (VM) when spawning a process. Identifying an appropriate node becomes a significant burden for large and dynamic systems. The problem is exacerbated in large distributed systems where (1) the hosts may not be identical, having different hardware capabilities or different software installed; and (2) communication times may be non-uniform: it may be fast to send a message between VMs on the same host, and slow if the VMs are on different hosts in a large distributed system.

These factors make it difficult to deploy applications, especially in a scalable and *portable* manner. Moreover while the programmer may be able to use platform-specific knowledge to decide where to spawn processes to enable an application to run efficiently, if the application is then deployed on a different platform, or if the platform changes as hosts fail or are added, this becomes outdated.

Discussion Our investigations confirm three language-level scalability limitations of Erlang




Figure 9: Scalability vs. percentage of global commands in Distributed Erlang.

from developer folklore. (1) Maintaining a *fully connected network* of Erlang nodes limits scalability, for example Orbit is typically limited to just 40 nodes. (2) *Global operations*, and crucially the global operations required for reliability, i.e. to maintain a *global namespace*, seriously limit the scalability of distributed Erlang systems. (3) Explicit *process placement* makes it hard to build performance portable applications for large architectures. These issues cause designers of reliable large scale systems in Erlang to depart from the standard Erlang model, e.g. using techniques like hidden nodes and storing pids in data structures. In Section V we develop language technologies to address these issues.

iii. Persistent Storage

Any large scale system needs reliable scalable persistent storage, and we have studied the scalability limits of Erlang persistent storage alternatives [38]. We envisage a typical large server having around 10^5 cores on around 100 hosts. We have reviewed the requirements for scalable and available persistent stor-

age and evaluated four popular Erlang DBMS against these requirements. For a target scale of around 100 hosts, Mnesia and CouchDB are, unsurprisingly, not suitable. However, Dynamo-style NoSQL DBMS like Cassandra and Riak have the potential to be.

We have investigated the current scalability limits of the Riak NoSQL DBMS using the Basho Bench benchmarking framework on a cluster with up to 100 nodes and independent disks. We found that the scalability limit of Riak version 1.1.1 is 60 nodes on the Kalkyl cluster. The study placed into the public scientific domain what was previously well-evidenced, but anecdotal, developer experience.

We have also shown that resources like memory, disk, and network do not limit the scalability of Riak. By instrumenting the `global` and `gen_server` OTP libraries we identified a specific Riak remote procedure call that fails to scale. We outline how later releases of Riak are refactored to eliminate the scalability bottlenecks.




Figure 10: Latency of commands as the number of Erlang nodes increases.

Discussion We conclude that Dynamo-like NoSQL DBMSs have the potential to deliver reliable persistent storage for Erlang at our target scale of approximately 100 hosts. Specifically an Erlang Cassandra interface is available and Riak 1.1.1 already provides scalable and available persistent storage on 60 nodes. Moreover the scalability of Riak is much improved in subsequent versions.

V. IMPROVING LANGUAGE SCALABILITY

This section outlines the Scalable Distributed (SD) Erlang libraries [21] we have designed and implemented to address the distributed Erlang scalability issues identified in Section ii. SD Erlang introduces two concepts to improve scalability. *S_groups* partition the set of nodes in an Erlang system to reduce network connectivity and partition global data (Section i.1). *Semi-explicit placement* alleviates the issues of explicit process placement in large heterogeneous networks (Section i.2). The two features are independent and can be used separately or in combination. We overview SD Erlang in Section i, and outline *s_group* semantics and

validation in Sections ii, iii respectively.

i. SD Erlang Design

i.1 S_groups

reduce both the number of connections a node maintains, and the size of name spaces, i.e. they minimise global information. Specifically names are registered on, and synchronised between, only the nodes within the *s_group*. An *s_group* has the following parameters: a name, a list of nodes, and a list of registered names. A node can belong to many *s_group*s or to none. If a node belongs to no *s_group* it behaves as a usual distributed Erlang node.

The *s_group* library defines the functions shown in Table 1. Some of these functions manipulate *s_group*s and provide information about them, such as creating *s_group*s and providing a list of nodes from a given *s_group*. The remaining functions manipulate names registered in *s_group*s and provide information about these names. For example, to register a process, *Pid*, with name *Name* in *s_group* *SGroupName* we use the following function. The name will only be registered if the process is

Table 1: Summary of *s_group* functions.

Function	Description
<code>new_s_group(SGroupName, Nodes)</code>	Creates a new <i>s_group</i> consisting of some nodes.
<code>delete_s_group(SGroupName)</code>	Deletes an <i>s_group</i> .
<code>add_nodes(SGroupName, Nodes)</code>	Adds a list of nodes to an <i>s_group</i> .
<code>remove_nodes(SGroupName, Nodes)</code>	Removes a list of nodes from an <i>s_group</i> .
<code>s_groups()</code>	Returns a list of all <i>s_groups</i> known to the node.
<code>own_s_groups()</code>	Returns a list of <i>s_group</i> tuples the node belongs to.
<code>own_nodes()</code>	Returns a list of nodes from all <i>s_groups</i> the node belongs to.
<code>own_nodes(SGroupName)</code>	Returns a list of nodes from the given <i>s_group</i> .
<code>info()</code>	Returns <i>s_group</i> state information.
<code>register_name(SGroupName, Name, Pid)</code>	Registers a name in the given <i>s_group</i> .
<code>re_register_name(SGroupName, Name, Pid)</code>	Re-registers a name (changes a registration) in a given <i>s_group</i> .
<code>unregister_name(SGroupName, Name)</code>	Unregisters a name in the given <i>s_group</i> .
<code>registered_names({node, Node})</code>	Returns a list of all registered names on the given node.
<code>registered_names({s_group, SGroupName})</code>	Returns a list of all registered names in the given <i>s_group</i> .
<code>whereis_name(SGroupName, Name)</code>	Return the pid of a name registered in the given <i>s_group</i> .
<code>whereis_name(Node, SGroupName, Name)</code>	
<code>send(SGroupName, Name, Msg)</code>	Send a message to a name registered in the given <i>s_group</i> .
<code>send(Node, SGroupName, Name, Msg)</code>	

being executed on a node that belongs to the given *s_group*, and neither *Name* nor *Pid* are already registered in that group.

```
s_group:register_name(SGroupName, Name,
                    Pid) -> yes | no
```

To illustrate the impact of *s_groups* on scalability we repeat the global operations experiment from Section ii (Figure 9). In the SD Erlang experiment we partition the set of nodes into *s_groups* each containing ten nodes, and

hence the names are replicated and synchronised on just ten nodes, and not on all nodes as in distributed Erlang. The results in Figure 11 show that with 0.01% of global operations throughput of distributed Erlang stops growing at 40 nodes while throughput of SD Erlang continues to grow linearly.

The connection topology of *s_groups* is extremely flexible: they may be organised into a hierarchy of arbitrary depth or branching, e.g. there could be multiple levels in the tree




Figure 11: Global operations in Distributed Erlang vs. SD Erlang.

of `s_groups`; see Figure 12. Moreover it is not necessary to create an hierarchy of `s_groups`, for example, we have constructed an Orbit implementation using a ring of `s_groups`.

Given such a flexible way of organising distributed systems, key questions in the design of an SD Erlang system are the following. *How should `s_groups` be structured?* Depending on the reason the nodes are grouped – reducing the number of connections, or reducing the namespace, or both – `s_groups` can be freely structured as a tree, ring, or some other topology. *How large should the `s_groups` be?* Smaller `s_groups` mean more inter-group communication, but the synchronisation of the `s_group` state between the `s_group` nodes constrains the maximum size of `s_groups`. We have not found this constraint to be a serious restriction. For example many `s_groups` are either relatively small, e.g. 10-node, internal or terminal elements in some topology, e.g. the leaves and nodes of a tree. *How do nodes from different `s_groups` communicate?* While any two nodes can communicate in an SD Erlang system, to minimise the number of connections communication between nodes from different `s_groups`

is typically routed via gateway nodes that belong to both `s_groups`. *How do we avoid single points of failure?* For reliability, and to minimise communication load, multiple gateway nodes and processes may be required.

Information to make these design choices is provided by the tools in Section VII and by benchmarking. A further challenge is how to systematically refactor a distributed Erlang application into SD Erlang, and this is outlined in Section i. A detailed discussion of distributed system design and refactoring in SD Erlang provided in a recent article [22].

We illustrate typical SD Erlang system designs by showing refactorings of the Orbit and ACO benchmarks from Section III. In both distributed and SD Erlang the computation starts on the *Master* node and the actual computation is done on the *Worker* nodes. In the distributed Erlang version all nodes are interconnected, and messages are transferred directly from the sending node to the receiving node (Figure 2). In contrast, in the SD Erlang version nodes are grouped into `s_groups`, and messages are transferred between different `s_groups` via *Sub-master* nodes (Figure 13).

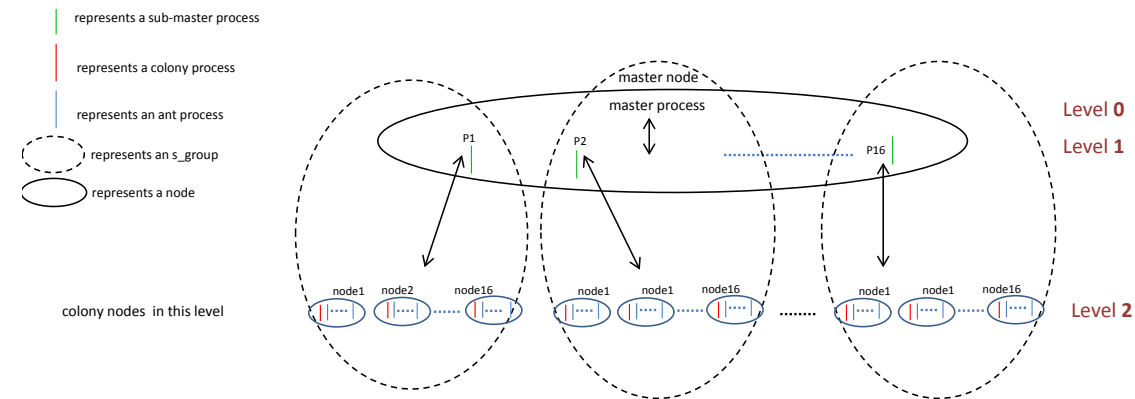


Figure 12: SD Erlang ACO (SR-ACO) architecture.

A fragment of code that creates an `s_group` on a *Sub-master* node is as follows:

```
create_s_group(Master,
               GroupName, Nodes0) ->
case s_group:new_s_group(GroupName,
                        Nodes0) of
{ok, GroupName, Nodes} ->
    Master ! {GroupName, Nodes};
_ -> io:format("exception: message")
end.
```

Similarly, we introduce `s_groups` in the GR-ACO benchmark from Section ii to create Scalable Reliable ACO (SR-ACO); see Figure 12. Here, apart from reducing the number of connections, `s_groups` also reduce the global namespace information. That is, instead of registering the name of a pid globally, i.e. with all nodes, the names is registered only on all nodes in the `s_group` with `s_group:register_name/3`.

A comparative performance evaluation of distributed Erlang and SD Erlang Orbit and ACO is presented in Section VIII.

i.2 Semi-Explicit Placement

Recall from Section iii that distributed Erlang spawns a process onto an explicitly named Erlang node, e.g.

```
spawn(some_node, fun some_module:pong/0).
```

and also recall the portability and programming effort issues associated with such explicit placement in large scale systems discussed in Section ii.

To address these issues we have developed a semi-explicit placement library that enables the programmer to select nodes on which to spawn processes based on run-time information about the properties of the nodes. For example, if a process performs a lot of computation one would like to spawn it on a node with considerable computation power, or if two processes are likely to communicate frequently then it would be desirable to spawn them on the same node, or nodes with a fast interconnect.

We have implemented two Erlang libraries to support semi-explicit placement [60]. The first deals with *node attributes*, and describes properties of individual Erlang VMs and associated hosts, such as total and currently available RAM, installed software, hardware configuration, etc. The second deals with a notion of *communication distances* which models the communication times between nodes in a distributed system. Therefore, instead of specifying a node we can use the `attr:choose_node/1` function to define the target node, i.e.

```
spawn(attr:choose_node(Params),
```




Figure 13: SD Erlang (SD-Orbit) architecture.

```
fun some_module:pong/0).
```

[60] report an investigation into the communication latencies on a range of NUMA and cluster architectures, and demonstrate the effectiveness of the placement libraries using the ML-ACO benchmark on the Athos cluster.

ii. S_group Semantics

For precise specification, and as a basis for validation, we provide a small-step operational semantics of the `s_group` operations [21]. Figure 14 defines the state of an SD Erlang system and associated abstract syntax variables. The abstract syntax variables on the left are defined as members of sets, denoted $\{\}$, and these in turn may contain tuples, denoted $()$, or further sets. In particular nm is a process name, p a pid, ni a node_id, and nis a set of node_ids. The state of a system is modelled as a four tuple comprising a set of s_groups , a set of $free_groups$, a set of $free_hidden_groups$, and a set of $nodes$. Each type of group is associated with nodes and has a namespace. An s_group additionally has a name, whereas a $free_hidden_group$ consists of only one node,

i.e. a hidden node simultaneously acts as a node and as a group, because as a group it has a namespace but does not share it with any other node. Free normal and hidden groups have no names, and are uniquely defined by the nodes associated with them. Therefore, group names, gr_names , are either *NoGroup* or a set of s_group_names . A namespace is a set of $name$ and process id, pid , pairs and is replicated on all nodes of the associated group.

A *node* has the following four parameters: $node_id$ identifier, $node_type$ that can be either hidden or normal, $connections$, and $group_names$, i.e. names of groups the node belongs to. The node can belong to either a list of s_groups or one of the free groups. The type of the free group is defined by the node type. Connections are a set of $node_ids$.

Transitions in the semantics have the form $(state, command, ni) \longrightarrow (state', value)$ meaning that executing $command$ on node ni in $state$ returns $value$ and transitions to $state'$.

The semantics is presented in more detail by [21], but we illustrate it here with the `s_group:registered_names/1` function from Section i.1. The function returns a list of names

$$\begin{aligned}
 (grs, fgs, fhs, nds) \in \{state\} &\equiv \{(\{s_group\}, \{free_group\}, \{free_hidden_group\}, \{node\})\} \\
 gr \in grs &\equiv \{s_group\} \equiv \{(s_group_name, \{node_id\}, namespace)\} \\
 fg \in fgs &\equiv \{free_group\} \equiv \{(\{node_id\}, namespace)\} \\
 fh \in fhs &\equiv \{free_hidden_group\} \equiv \{(node_id, namespace)\} \\
 nd \in nds &\equiv \{node\} \equiv \{(node_id, node_type, connections, gr_names)\} \\
 gs \in \{gr_names\} &\equiv \{NoGroup, \{s_group_name\}\} \\
 ns \in \{namespace\} &\equiv \{\{(name, pid)\}\} \\
 cs \in \{connections\} &\equiv \{\{node_id\}\} \\
 nt \in \{node_type\} &\equiv \{Normal, Hidden\} \\
 s \in \{NoGroup, s_group_name\} &
 \end{aligned}$$

Figure 14: SD Erlang state [21].

registered in `s_group` `s` if node `ni` belongs to the `s_group`, an empty list otherwise (Figure 15). Here \oplus denotes disjoint set union; `IsSGroupNode` returns true if node `ni` is a member of some `s_group` `s`, false otherwise; and `OutputNms` returns a set of process names registered in the `ns` namespace of `s_group` `s`.

iii. Semantics Validation

As the semantics is concrete it can readily be made *executable* in Erlang, with lists replacing

sets throughout. Having an executable semantics allows users to engage with it, and to understand how the semantics behaves *vis à vis* the library, giving them an opportunity to assess the correctness of the library against the semantics.

Better still, we can automatically assess how the system behaves in comparison with the (executable) semantics by executing them in lockstep, guided by the constraints of which operations are possible at each point. We do that by building an abstract state machine model of the library. We can then generate random sequences (or traces) through the model, with appropriate library data generated too. This random generation is supported by the QuickCheck property-based testing system [24, 9].

The architecture of the testing framework is shown in Figure 16. First an abstract state machine embedded as an “`eqc_statem`” module is derived from the executable semantic specification. The state machine defines the abstract state representation and the transition from one state to another when an operation is applied. Test case and data generators are then defined to control the test case generation;

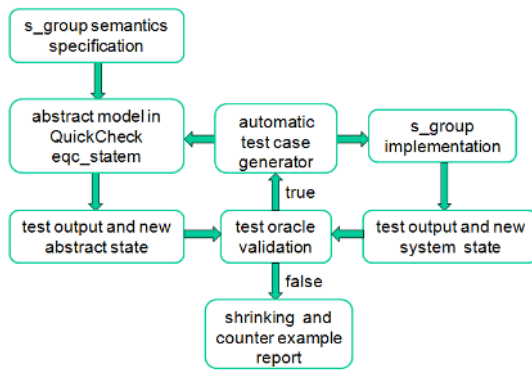


Figure 16: Testing `s_groups` using QuickCheck.

$$\begin{aligned}
 & ((grs, fgs, fhs, nds), s_group : registered_names(s), ni) \\
 & \longrightarrow ((grs, fgs, fhs, nds), nms) && \text{if } IsSGroupNode(ni, s, grs) \\
 & \longrightarrow ((grs, fgs, fhs, nds), \{\}) && \text{otherwise}
 \end{aligned}$$

where

$$\begin{aligned}
 & \{(s, \{ni\} \oplus nis, ns)\} \oplus grs' \equiv grs \\
 & nms \equiv OutputNms(s, ns)
 \end{aligned}$$

$$IsSGroupNode(ni, s, grs) = \exists nis, ns, grs' . \{(s, \{ni\} \oplus nis, ns)\} \oplus grs' \equiv grs$$

$$OutputNms(s, ns) = \{(s, nm) \mid (nm, p) \in ns\}$$

Figure 15: SD Erlang Semantics of `s_group:registered_names/1` Function

this includes the automatic generation of eligible `s_group` operations and the input data to those operations. Test oracles are encoded as the postcondition for `s_group` operations.

During testing, each test command is applied to both the abstract model and the `s_group` library. The application of the test command to the abstract model takes the abstract model from its current state to a new state as described by the transition functions; whereas the application of the test command to the library leads the system to a new actual state. The actual state information is collected from each node in the distributed system, then merged and normalised to the same format as the abstract state representation. For a test to be successful, after the execution of a test command, the test oracles specified for this command should be satisfied. Various test oracles can be defined for `s_group` operations; for instance one of the generic constraints that applies to all the `s_group` operations is that after each `s_group` operation, the normalised system state should be equivalent to the abstract state.

Thousands of tests were run, and three kinds of errors — which have subsequently been corrected — were found. Some errors in the

library implementation were found, including one error due to the synchronisation between nodes, and the other related to the `remove_nodes` operation, which erroneously raised an exception. We also found a couple of trivial errors in the semantic specification itself, which had been missed by manual examination. Finally, we found some situations where there were inconsistencies between the semantics and the library implementation, despite their states being equivalent: an example of this was in the particular values returned by functions on certain errors. Overall, the automation of testing boosted our confidence in the correctness of the library implementation and the semantic specification. This work is reported in more detail by [54].

VI. IMPROVING VM SCALABILITY

This section reports the primary VM and library improvements we have designed and implemented to address the scalability and reliability issues identified in Section i.

i. Improvements to Erlang Term Storage

Because ETS tables are so heavily used in Erlang systems, they are a focus for scalability improvements. We start by describing their redesign, including some improvements that pre-date our RELEASE project work, i.e. those prior to Erlang/OTP R15B03. These historical improvements are very relevant for a scalability study and form the basis for our subsequent changes and improvements. At the point when Erlang/OTP got support for multiple cores (in release R11B), there was a single reader-writer lock for each ETS table. Optional fine grained locking of hash-based ETS tables (i.e. `set`, `bag` or `duplicate_bag` tables) was introduced in Erlang/OTP R13B02-1, adding 16 reader-writer locks for the hash buckets. Reader groups to minimise read synchronisation overheads were introduced in Erlang/OTP R14B. The key observation is that a single count of the multiple readers must be synchronised across many cache lines, potentially far away in a NUMA system. Maintaining reader counts in multiple (local) caches makes reads fast, although writes must now check every reader count. In Erlang/OTP R16B the number of bucket locks, and the default number of reader groups, were both upgraded from 16 to 64.

We illustrate the scaling properties of the ETS concurrency options using the `ets_bench` BenchErl benchmark on the Intel NUMA machine with 32 hyperthreaded cores specified in Appendix A. The `ets_bench` benchmark inserts 1M items into the table, then records the time to perform 17M operations, where an operation is either a lookup, an insert, or a delete. The experiments are conducted on a hash-based (`set`) ETS table with different percentages of update operations, i.e. insertions or deletions.

Figure 17 shows the runtimes in seconds of 17M operations in different Erlang/OTP versions, varying the number of schedulers (x-axis), reflecting how the scalability of ETS ta-




Figure 17: Runtime of 17M ETS operations in Erlang/OTP releases: 10% updates (left) and 1% updates (right). The y-axis shows time in seconds (lower is better) and the x-axis is number of OS threads (VM schedulers).

bles has improved in more recent Erlang/OTP releases. Figure 18 shows the runtimes in seconds of 17M operations on an ETS table with different numbers of reader groups, again varying the number of schedulers. We see that one reader group is not sufficient with 10% updates, nor are two with 1% updates. Beyond that, different numbers of reader groups have little impact on the benchmark performance except that using 64 groups with 10% updates slightly degrades performance.

We have explored four other extensions or redesigns in the ETS implementation for better scalability. 1. Allowing more programmer control over the number of bucket locks in hash-based tables, so the programmer can reflect the number of schedulers and the expected




Figure 18: Runtime of 17M ETS operations with varying numbers of reader groups: 10% updates (left) and 1% updates (right). The y-axis shows runtime in seconds (lower is better) and the x-axis is number of schedulers.

access pattern. 2. Using contention-adapting trees to get better scalability for `ordered_set` ETS tables as described by [69]. 3. Using queue delegation locking to improve scalability [47]. 4. Adopting schemes for completely eliminating the locks in the meta table. A more complete discussion of our work on ETS can be found in the papers by [69] and [47].

Here we outline only our work on contention-adapting (CA) trees. A CA tree monitors contention in different parts of a tree-shaped data structure, introducing routing nodes with locks in response to high contention, and removing them in response to low contention. For experimental purposes two variants of the CA tree have been implemented to represent `ordered_sets` in the virtual ma-

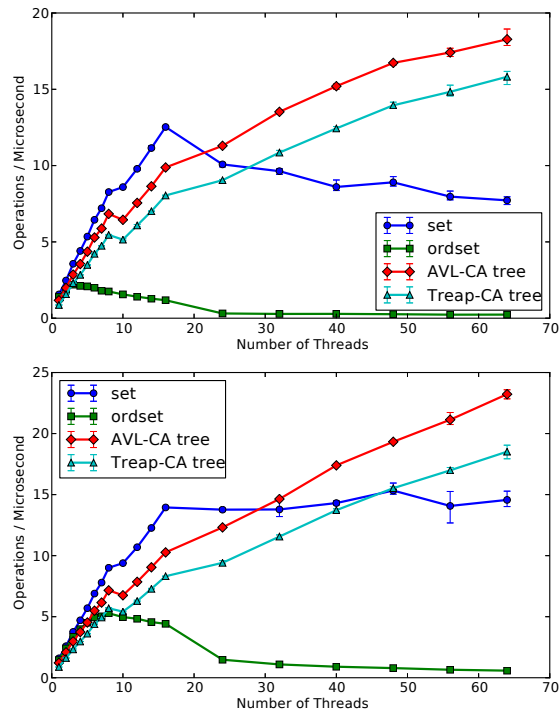


Figure 19: Throughput of CA tree variants: 10% updates (left) and 1% updates (right).

chine of Erlang/OTP 17.0. One extends the existing AVL trees in the Erlang VM, and the other uses a *Treap* data structure [5]. Figure 19 compares the throughput of the CA tree variants with that of `ordered_set` and `set` as the number of schedulers increases. It is unsurprising that the CA trees scale so much better than an `ordered_set` which is protected by a single readers-writer lock. It is more surprising that they also scale better than `set`. This is due to hash tables using fine-grained locking at a fixed granularity, while CA trees can adapt the number of locks to the current contention level, and also to the parts of the key range where contention is occurring.

ii. Improvements to Schedulers

In the Erlang VM, a scheduler is responsible for executing multiple processes concurrently, in a timely and fair fashion, making optimal use of hardware resources. The VM implements preemptive multitasking with soft real-time guarantees. Erlang processes are normally scheduled on a reduction count basis where one reduction is roughly equivalent to a function call. Each process is allowed to execute until it either blocks waiting for input, typically a message from some other process, or until it has executed its quota of reductions.

The Erlang VM is usually started with one scheduler per logical core (SMT-thread) available on the host machine, and schedulers are implemented as OS threads. When an Erlang process is spawned it is placed in the run queue of the scheduler of its parent, and it waits on that queue until the scheduler allocates it a slice of core time. Work stealing is used to balance load between cores, that is an idle scheduler may *migrate* a process from another run queue. Scheduler run queues are visualised in Figure iii.2.

The default load management mechanism is load compaction that aims to keep as many scheduler threads as possible fully loaded with work, i.e. it attempts to ensure that scheduler threads do not run out of work. We have developed a new optional *scheduler utilisation balancing* mechanism that is available from Erlang/OTP 17.0. The new mechanism aims to balance scheduler utilisation between schedulers; that is, it will strive for equal scheduler utilisation on all schedulers.

The scheduler utilisation balancing mechanism has no performance impact on the system when not enabled. On the other hand, when enabled, it results in changed timing in the system; normally there is a small overhead due to measuring of utilisation and calculating balancing information, which depends on the underlying primitives provided by the operating

system.

The new balancing mechanism results in a better distribution of processes to schedulers, reducing the probability of core contention. Together with other VM improvements, such as interruptable BIFs and garbage collection, it results in lower latency and improved responsiveness, and hence reliability, for soft real-time applications.

iii. Improvements to Time Management

Soon after the start of the RELEASE project, time management in the Erlang VM became a scalability bottleneck for many applications, as illustrated by the parallel benchmark in Figure i. The issue came to prominence as other, more severe, bottlenecks were eliminated. This subsection motivates and outlines the improvements to time management that we made; these were incorporated into Erlang/OTP 18.x as a new API for time and time warping. The old API is still supported at the time of writing, but its use is deprecated.

The original time API provides the `erlang:now/0` built-in that returns “Erlang system time” or time since Epoch with micro second resolution. This time is the basis for all time internally in the Erlang VM.

Many of the scalability problems of `erlang:now/0` stem from its specification, written at a time when the Erlang VM was not multi-threaded, i.e. SMT-enabled. The documentation promises that values returned by it are strictly increasing and many applications ended up relying on this. For example applications often employ `erlang:now/0` to generate unique integers.

Erlang system time should align with the operating system’s view of time since Epoch or “OS system time”. However, while OS system time can be freely changed both forwards and backwards, Erlang system time cannot, without invalidating the strictly increasing value

guarantee. The Erlang VM therefore contains a mechanism that slowly adjusts Erlang system time towards OS system time if they do not align.

One problem with time adjustment is that the VM deliberately presents time with an inaccurate frequency; this is required to align Erlang system time with OS system time smoothly when these two have deviated, e.g. in the case of clock shifts when leap seconds are inserted or deleted. Another problem is that Erlang system time and OS system time can differ for very long periods of time. In the new API, we resolve this using a common OS technique [59], i.e. a monotonic time that has its zero point at some unspecified point in time. Monotonic time is not allowed to make leaps forwards and backwards while system time is allowed to do this. Erlang system time is thus just a dynamically varying offset from Erlang monotonic time.

Time Retrieval Retrieval of Erlang system time was previously protected by a global mutex, which made the operation thread safe, but scaled poorly. Erlang system time and Erlang monotonic time need to run at the same frequency, otherwise the time offset between them would not be constant. In the common case, monotonic time delivered by the operating system is solely based on the machine's local clock and cannot be changed, while the system time is adjusted using the Network Time Protocol (NTP). That is, they will run with different frequencies. Linux is an exception with a monotonic clock that is NTP adjusted and runs with the same frequency as system time [76]. To align the frequencies of Erlang monotonic time and Erlang system time, we adjust the frequency of the Erlang monotonic clock. This is done by comparing monotonic time and system time delivered by the OS, and calculating an adjustment. To achieve this scalably, one VM thread calculates the time adjustment to use at least once a minute. If the adjustment

needs to be changed, new adjustment information is published and used to calculate Erlang monotonic time in the future.

When a thread needs to retrieve time, it reads the monotonic time delivered by the OS and the time adjustment information previously published and calculates Erlang monotonic time. To preserve monotonicity it is important that all threads that read the same OS monotonic time map this to exactly the same Erlang monotonic time. This requires synchronisation on updates to the adjustment information using a readers-writer (RW) lock. This RW lock is write-locked only when the adjustment information is changed. This means that in the vast majority of cases the RW lock will be read-locked, which allows multiple readers to run concurrently. To prevent bouncing the lock cache-line we use a bespoke reader optimised RW lock implementation where reader threads notify about their presence in counters on separate cache-lines. The concept is similar to the reader indicator algorithm described by [48, Fig. 11] and alternatives include the ingress-egress counter used by [18] and the SNZI algorithm of [30].

Timer Wheel and BIF Timer The timer wheel contains all timers set by Erlang processes. The original implementation was protected by a global mutex and scaled poorly. To increase concurrency, each scheduler thread has been assigned its own timer wheel that is used by processes executing on the scheduler.

The implementation of timers in Erlang/OTP uses a built in function (BIF), as most low-level operations do. Until Erlang/OTP 17.4, this BIF was also protected by a global mutex. Besides inserting timers into the timer wheel, the BIF timer implementation also maps timer references to a timer in the timer wheel. To improve concurrency, from Erlang/OTP 18 we provide scheduler-specific BIF timer servers as Erlang processes. These keep information about timers in private ETS

tables and only insert one timer at the time into the timer wheel.

Benchmarks We have measured several benchmarks on a 16-core Bulldozer machine with eight dual CPU AMD Opteron 4376 HEs.⁵ We present three of them here.

The first micro benchmark compares the execution time of an Erlang receive with that of a receive after that specifies a timeout and provides a default value. The receive after sets a timer when the process blocks in the receive, and cancels it when a message arrives. In Erlang/OTP 17.4 the total execution time with standard timers is 62% longer than without timers. Using the improved implementation in Erlang/OTP 18.0, total execution time with the optimised timers is only 5% longer than without timers.

The second micro benchmark repeatedly checks the system time, calling the built-in `erlang:now/0` in Erlang/OTP 17.4, and calling both `erlang:monotonic_time/0` and `erlang:time_offset/0` and adding the results in Erlang/OTP 18.0. In this machine, where the VM uses 16 schedulers by default, the 18.0 release is more than 69 times faster than the 17.4 release.

The third benchmark is the parallel BenchErl benchmark from Section i. Figure 20 shows the results of executing the original version of this benchmark, which uses `erlang:now/0` to create monotonically increasing unique values, using three Erlang/OTP releases: R15B01, 17.4, and 18.1. We also measure a version of the benchmark in Erlang/OTP 18.1 where the call to `erlang:now/0` has been substituted with a call to `erlang:monotonic_time/0`. The graph on its left shows that: 1. the performance of time management has remained roughly unchanged between Erlang/OTP releases prior to 18.0; 2. the improved time management in Erlang/OTP 18.x make time management less

⁵See §2.5.4 of the RELEASE project Deliverable 2.4 (<http://release-project.eu/documents/D2.4.pdf>).




Figure 20: *BenchErl* parallel benchmark using `erlang:monotonic_time/0` or `erlang:now/0` in different Erlang/OTP releases: runtimes (left) and speedup obtained using `erlang:monotonic_time/0` (right).

likely to be a scalability bottleneck even when using `erlang:now/0`, and 3. the new time API (using `erlang:monotonic_time/0` and friends) provides a scalable solution. The graph on the right side of Figure 20 shows the speedup that the modified version of the parallel benchmark achieves in Erlang/OTP 18.1.

VII. SCALABLE TOOLS

This section outlines five tools developed in the RELEASE project to support scalable Erlang systems. Some tools were developed from scratch, like Devo, SDMon and WombatOAM, while others extend existing tools, like Percept and Wrangler. These include tooling to transform programs to make them more scalable,

to deploy them for scalability, to monitor and visualise them. Most of the tools are freely available under open source licences (Devo, Percept2, SD-Mon, Wrangler); while Wombat-OAM is a commercial product. The tools have been used for profiling and refactoring the ACO and Orbit benchmarks from Section III.

The Erlang tool “ecosystem” consists of small stand-alone tools for tracing, profiling and debugging Erlang systems that can be used separately or together as appropriate for solving the problem at hand, rather than as a single, monolithic, super-tool. The tools presented here have been designed to be used as part of that ecosystem, and to complement already available functionality rather than to duplicate it. The Erlang runtime system has built-in support for tracing many types of events, and this infrastructure forms the basis of a number of tools for tracing and profiling. Typically the tools build on or specialise the services offered by the Erlang virtual machine, through a number of built-in functions. Most recently, and since the RELEASE project was planned, the Observer⁶ application gives a comprehensive overview of many of these data on a node-by-node basis.

As actor frameworks and languages (see Section ii) have only recently become widely adopted commercially, their tool support remains relatively immature and generic in nature. That is, the tools support the language itself, rather than its distinctively concurrent aspects. Given the widespread use of Erlang, tools developed for it point the way for tools in other actor languages and frameworks. For example, just as many Erlang tools use tracing support provided by the Erlang VM, so can other actor frameworks, e.g. Akka can use the Kamon⁷ JVM monitoring system. Similarly, tools for other actor languages or frameworks could use data derived through OS-level

tracing frameworks DTrace⁸ and SystemTap⁹ probes as we show in this section for Erlang, provided that the host language has tracing hooks into the appropriate infrastructure.

i. Refactoring for Scalability

Refactoring [65, 32, 75] is the process of changing how a program works without changing what it does. This can be done for readability, for testability, to prepare it for modification or extension, or — as is the case here — in order to improve its scalability. Because refactoring involves the transformation of source code, it is typically performed using machine support in a refactoring tool. There are a number of tools that support refactoring in Erlang: in the RELEASE project we have chosen to extend Wrangler¹⁰ [56]; other tools include Tidier [68] and RefactorErl [44].

Supporting API Migration The SD Erlang libraries modify Erlang’s `global_group` library, becoming the new `s_group` library; as a result, Erlang programs using `global_group` will have to be refactored to use `s_group`. This kind of API migration problem is not uncommon, as software evolves and this often changes the API of a library. Rather than simply extend Wrangler with a refactoring to perform this particular operation, we instead added a framework for the automatic generation of API migration refactorings from a user-defined adaptor module.

Our approach to automatic API migration works in this way: when an API function’s interface is changed, the author of this API function implements an *adaptor function*, defining calls to the old API in terms of the new. From this definition we automatically generate the refactoring that transforms the client code to use the new API; this refactoring can also be

⁶<http://www.erlang.org/doc/apps/observer/>

⁷<http://kamon.io>

⁸<http://dtrace.org/blogs/about/>

⁹<https://sourceware.org/systemtap/wiki>

¹⁰<http://www.cs.kent.ac.uk/projects/wrangler/>

supplied by the API writer to clients on library upgrade, allowing users to upgrade their code automatically. The refactoring works by generating a set of rules that “fold in” the adaptation to the client code, so that the resulting code works directly with the new API. More details of the design choices underlying the work and the technicalities of the implementation can be found in a paper by [52].

Support for Introducing Parallelism We have introduced support for parallelising explicit list operations (`map` and `foreach`), for process introduction to complete a computationally intensive task in parallel, for introducing a worker process to deal with call handling in an Erlang “generic server” and to parallelise a tail recursive function. We discuss these in turn now; more details and practical examples of the refactorings appear in a conference paper describing that work [55].

Uses of `map` and `foreach` in list processing are among of the most obvious places where parallelism can be introduced. We have added a small library to Wrangler, called `para_lib`, which provides parallel implementations of `map` and `foreach`. The transformation from an explicit use of sequential `map/foreach` to the use of their parallel counterparts is very straightforward, even manual refactoring would not be a problem. However a `map/foreach` operation could also be implemented differently using recursive functions, list comprehensions, etc.; identifying this kind of implicit `map/foreach` usage can be done using Wrangler’s code inspection facility, and a refactoring that turns an implicit `map/foreach` to an explicit `map/foreach` can also be specified using Wrangler’s rule-based transformation API.

If the computations of two non-trivial tasks do not depend on each other, then they can be executed in parallel. The *Introduce a New Process* refactoring implemented in Wrangler can be used to spawn a new process to execute

a task in parallel with its parent process. The result of the new process is sent back to the parent process, which will then consume it when needed. In order not to block other computations that do not depend on the result returned by the new process, the `receive` expression is placed immediately before the point where the result is needed.

While some tail-recursive list processing functions can be refactored to an explicit `map` operation, many cannot due to data dependencies. For instance, an example might perform a recursion over a list while accumulating results in an accumulator variable. In such a situation it is possible to “float out” some of the computations into parallel computations. This can only be done when certain dependency constraints are satisfied, and these are done by program slicing, which is discussed below.

Support for Program Slicing Program slicing is a general technique of program analysis for extracting the part of a program, also called the *slice*, that influences or is influenced by a given point of interest, i.e. the *slicing criterion*. Static program slicing is generally based on program dependency including both control dependency and data dependency. *Backward intra-function slicing* is used by some of the refactorings described above; it is also useful in general, and made available to end-users under Wrangler’s Inspector menu [55].

Our work can be compared with that in PaRTE¹¹ [17], a tool developed in another EU project that also re-uses the Wrangler front end. This work concentrates on skeleton introduction, as does some of our work, but we go further in using static analysis and slicing in transforming programs to make them suitable for introduction of parallel structures.

¹¹http://paraphrase-enlarged.elte.hu/downloads/D4-3_user_manual.pdf

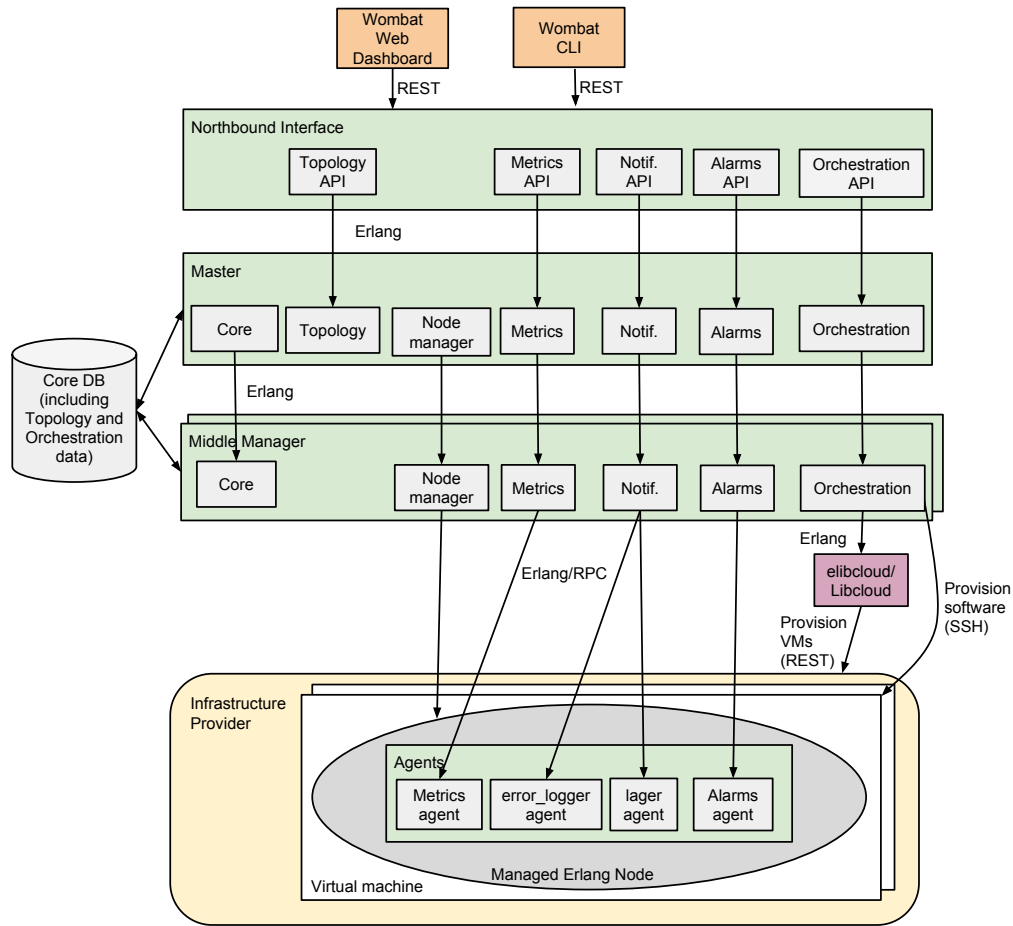


Figure 21: The architecture of WombatOAM.

ii. Scalable Deployment

We have developed the WombatOAM tool¹² to provide a deployment, operations and maintenance framework for large-scale Erlang distributed systems. These systems typically consist of a number of Erlang nodes executing on different hosts. These hosts may have different hardware or operating systems, be physical or virtual, or run different versions of Erlang/OTP.

¹²WombatOAM (<https://www.erlang-solutions.com/products/wombat-oam.html>) is a commercial tool available from Erlang Solutions Ltd.

Prior to the development of WombatOAM, deployment of systems would use scripting in Erlang and the shell, and this is the state of the art for other actor frameworks; it would be possible to adapt the WombatOAM approach to these frameworks in a straightforward way.

Architecture The architecture of WombatOAM is summarised in Figure 21. Originally the system had problems addressing full scalability because of the role played by the central *Master* node; in its current

version an additional layer of *Middle Managers* was introduced to allow the system to scale easily to thousands of deployed nodes. As the diagram shows, the “northbound” interfaces to the web dashboard and the command-line are provided through RESTful connections to the Master. The operations of the Master are delegated to the Middle Managers that engage directly with the managed nodes. Each managed node runs a collection of *services* that collect metrics, raise alarms and so forth; we describe those now.

Services WombatOAM is designed to collect, store and display various kinds of information and event from running Erlang systems, and these data are accessed and managed through an AJAX-based Web Dashboard; these include the following.

Metrics WombatOAM supports the collection of some hundred metrics — including, for instance, numbers of processes on a node and the message traffic in and out of the node — on a regular basis from the Erlang VMs running on each of the hosts. It can also collect metrics defined by users within other metrics collection frameworks, such as Folsom¹³, and interface with other tools such as graphite¹⁴ which can log and display such information. The metrics can be displayed as histograms covering different windows, such as the last fifteen minutes, hour, day, week or month.

Notifications As well as metrics, it can support *event-based* monitoring through the collection of notifications from running nodes. Notifications, which are *one time events* can be generated using the Erlang System Architecture Support Libraries, SASL, which is part of the standard distri-

bution, or the lager logging framework¹⁵, and will be displayed and logged as they occur.

Alarms Alarms are more complex entities. Alarms have a state: they can be *raised*, and once dealt with they can be *cleared*; they also have identities, so that the same alarm may be raised on the same node multiple times, and each instance will need to be dealt with separately. Alarms are generated by SASL or lager, just as for notifications.

Topology The Topology service handles adding, deleting and discovering nodes. It also monitors whether they are accessible, and if not, it notifies the other services, and periodically tries to reconnect. When the nodes are available again, it also notifies the other services. It doesn’t have a middle manager part, because it doesn’t talk to the nodes directly: instead it asks the Node manager service to do so.

Node manager This service maintains the connection to all managed nodes via the Erlang distribution protocol. If it loses the connection towards a node, it periodically tries to reconnect. It also maintains the states of the nodes in the database (e.g. if the connection towards a node is lost, the Node manager changes the node state to DOWN and raises an alarm). The Node manager doesn’t have a REST API, since the node states are provided via the Topology service’s REST API.

Orchestration This service can deploy new Erlang nodes on already running machines. It can also provision new virtual machine instances using several cloud providers, and deploy Erlang nodes on those instances. For communicating with the cloud providers, the Orchestration

¹³Folsom collects and publishes metrics through an Erlang API: <https://github.com/boundary/folsom>

¹⁴<https://graphiteapp.org>

¹⁵<https://github.com/basho/lager>

service uses an external library called Libcloud,¹⁶ for which Erlang Solutions has written an open source Erlang wrapper called `elibcloud`, to make Libcloud easier to use from WombatOAM. Note that WombatOAM Orchestration doesn't provide a platform for writing Erlang applications: it provides infrastructure for deploying them.

Deployment The mechanism consists of the following five steps.

1. *Registering a provider.* WombatOAM provides the same interface for different cloud providers which support the OpenStack standard or the Amazon EC2 API. WombatOAM also provides the same interface for using a fixed set of machines. In WombatOAM's backend, this has been implemented as two driver modules: the *elibcloud* driver module which uses the `elibcloud` and `Libcloud` libraries to communicate with the cloud providers, and the *SSH* driver module that keeps track of a fixed set of machines.
2. *Uploading a release.* The release can be either a proper Erlang release archive or a set of Erlang modules. The only important aspect from WombatOAM's point of view is that start and stop commands should be explicitly specified. This will enable WombatOAM start and stop nodes when needed.
3. *Defining a node family.* The next step is creating the node family, which is the entity that refers to a certain release, contains deployment domains that refer to certain providers, and contains other information necessary to deploy a node.

¹⁶The unified cloud API [4]: <https://libcloud.apache.org>

4. *Defining a deployment domain.* At this step a deployment domain is created that specifies (i) which providers should be used for provisioning machines; (ii) the username that will be used when WombatOAM connects to the hosts using *SSH*.
5. *Node deployment.* To deploy nodes a WombatOAM user needs only to specify the number of nodes and the node family these nodes belong to. Nodes can be dynamically added to, or removed from, the system depending on the needs of the application. The nodes are started, and WombatOAM is ready to initiate and run the application.

iii. Monitoring and Visualisation

A key aim in designing new monitoring and visualisation tools and adapting existing ones was to provide support for systems running on parallel and distributed hardware. Specifically, in order to run modern Erlang systems, and in particular SD Erlang systems, it is necessary to understand both their single host ("multi-core") and multi-host ("distributed") nature. That is, we need to be able to understand how systems run on the Erlang multicore virtual machine, where the scheduler associated with a core manages its own run queue, and processes migrate between the queues through a work stealing algorithm; at the same time, we have to understand the dynamics of a distributed Erlang program, where the user explicitly spawns processes to nodes.¹⁷

iii.1 Percept2

Percept2¹⁸ builds on the existing Percept tool to provide *post hoc* offline analysis and visuali-

¹⁷This is in contrast with the multicore VM, where programmers have no control over where processes are spawned; however, they still need to gain insight into the behaviour of their programs to tune performance.

¹⁸<https://github.com/RefactoringTools/percept2>

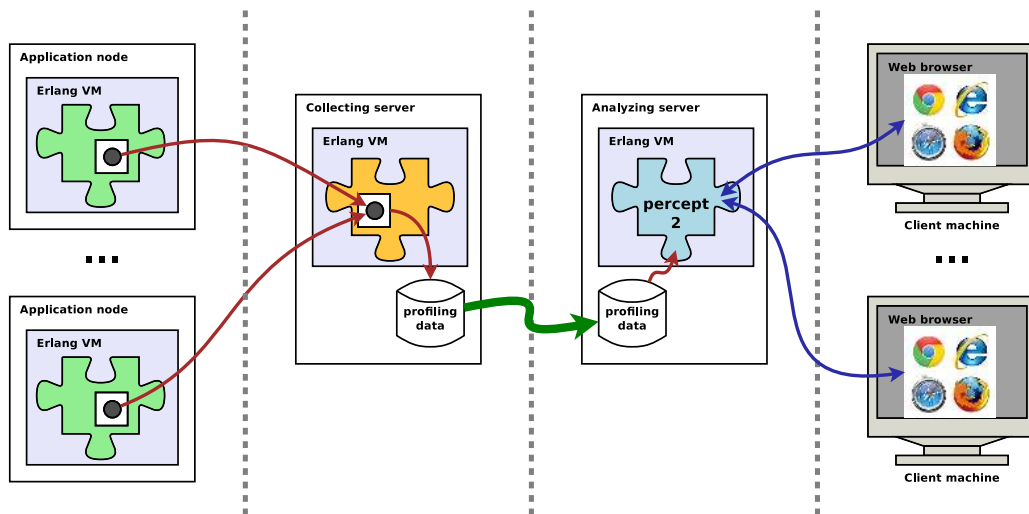


Figure 22: Offline profiling of a distributed Erlang application using Percept2.

sation of Erlang systems. Percept2 is designed to allow users to visualise and tune the parallel performance of Erlang systems on a single node on a single manycore host. It visualises Erlang application level concurrency and identifies concurrency bottlenecks. Percept2 uses Erlang built in tracing and profiling to monitor process states, i.e. waiting, running, runnable, free and exiting. A waiting or suspended process is considered an inactive and a running or runnable process is considered active. As a program runs with Percept, events are collected and stored to a file. The file is then analysed, with the results stored in a RAM database, and this data is viewed through a web-based interface. The process of offline profiling for a distributed Erlang application using Percept2 is shown in Figure 22.

Percept generates an application-level zoomable concurrency graph, showing the number of active processes at each point during profiling; dips in the graph represent low concurrency. A lifetime bar for each process is also included, showing the points during its lifetime when the process was active, as well as other per-process information.

Percept2 extends Percept in a number of ways — as detailed by [53] — including most importantly:

- Distinguishing between running and runnable time for each process: this is apparent in the process runnability comparison as shown in Figure 23, where orange represents runnable and green represents running. This shows very clearly where potential concurrency is not being exploited.
- Showing scheduler activity: the number of active schedulers at any time during the profiling.
- Recording more information during execution, including the migration history of a process between run queues; statistics about message passing between processes: the number of messages and the average message size sent/received by a process; the accumulated runtime per-process: the accumulated time when a process is in a running state.
- Presenting the process tree: the hierarchy




Figure 23: *Percept2*: showing processes running and runnable (with 4 schedulers/run queues).

structure indicating the parent-child relationships between processes.

- Recording dynamic function call graph/count/time: the hierarchy structure showing the calling relationships between functions during the program run, and the amount of time spent in a function.
- Tracing of `s_group` activities in a distributed system. Unlike global group, `s_group` allows dynamic changes to the `s_group` structure of a distributed Erlang system. In order to support SD Erlang, we have also extended *Percept2* to allow the profiling of `s_group` related activities, so that the dynamic changes to the `s_group` structure of a distributed Erlang system can be captured.

We have also improved on *Percept* as follows.

- Enabling finer-grained control of what is profiled. The profiling of port activities, schedulers activities, message passing, process migration, garbage collection and `s_group` activities can be enabled/disabled, while the profiling of process runnability (indicated by the “proc” flag) is always enabled.

- Selective function profiling of processes. In *Percept2*, we have built a version of `fprof`, which does not measure a function’s own execution time, but measures everything else that `fprof` measures. Eliminating measurement of a function’s own execution time gives users the freedom of not profiling all the function calls invoked during the program execution. For example, they can choose to profile only functions defined in their own applications’ code, and not those in libraries.
- Improved dynamic function callgraph. With the dynamic function callgraph, a user is able to understand the causes of certain events, such as heavy calls of a particular function, by examining the region around the node for the function, including the path to the root of the graph. Each edge in the callgraph is annotated with the number of times the target function is called by the source function as well as further information.

Finally, we have also improved the scalability of *Percept* in three ways. First we have parallelised the processing of trace files so that multiple data files can be processed at the same time. We have also compressed the representation of call graphs, and cached the history of

generated web pages. Together these make the system more responsive and more scalable.

iii.2 DTrace/SystemTap tracing

DTrace provides dynamic tracing support for various flavours of Unix, including BSD and Mac OS X, and SystemTap does the same for Linux; both allow the monitoring of live, running systems with minimal overhead. They can be used by administrators, language developers and application developers alike to examine the behaviour of applications, language implementations and the operating system during development or even on live production systems. In comparison to other similar tools and instrumentation frameworks, they are relatively lightweight, do not require special recompiled versions of the software to be examined, nor special post-processing tools to create meaningful information from the data gathered. Using these probes, it is possible to identify bottlenecks both in the VM itself and in applications.

VM bottlenecks are identified using a large number of probes inserted into Erlang/OTP's VM, for example to explore scheduler run-queue lengths. These probes can be used to measure the number of processes per scheduler, the number of processes moved during work stealing, the number of attempts to gain a run-queue lock, how many of these succeed immediately, etc. Figure 24 visualises the results of such monitoring; it shows how the size of run queues vary during execution of the bang BenchErl benchmark on a VM with 16 schedulers.

Application bottlenecks are identified by an alternative back-end for Percept2, based on DTrace/SystemTap instead of the Erlang built-in tracing mechanism. The implementation reuses the existing Percept2 infrastructure as far as possible. It uses a different mechanism for collecting information about Erlang programs, a different format for the trace files, but the




Figure 24: DTrace: runqueue visualisation of bang benchmark with 16 schedulers (time on X-axis).

same storage infrastructure and presentation facilities.

iii.3 Devo

Devo¹⁹ [10] is designed to provide real-time on-line visualisation of both the low-level (single node, multiple cores) and high-level (multiple nodes, grouped into s_groups) aspects of Erlang and SD Erlang systems. Visualisation is within a browser, with web sockets providing the connections between the JavaScript visualisations and the running Erlang systems, instrumented through the trace tool builder (ttb).

Figure 25 shows visualisations from devo in both modes. On the left-hand side a single compute node is shown. This consists of two physical chips (the upper and lower halves of the diagram) with six cores each; with hyperthreading this gives twelve virtual cores, and hence 24 run queues in total. The size of these run queues is shown by both the colour and height of each column, and process migrations are illustrated by (fading) arc between the queues within the circle. A green arc shows mi-

¹⁹<https://github.com/RefactoringTools/devo>




Figure 25: *Devoo: low and high-level visualisations of SD-Orbit.*

gration on the same physical core, a grey one on the same chip, and blue shows migrations between the two chips.

On the right-hand side of Figure 25 is a visualisation of SD-Orbit in action. Each node in the graph represents an Erlang node, and colours (red, green, blue and orange) are used to represent the *s_group*s to which a node belongs. As is evident, the three nodes in the central triangle belong to multiple groups, and act as routing nodes between the other nodes. The colour of the arc joining two nodes repre-

sents the current intensity of communication between the nodes (green quiescent; red busiest).

iii.4 SD-Mon

SD-Mon is a tool specifically designed for monitoring SD-Erlang systems. This purpose is accomplished by means of a *shadow network* of agents, that collect data from a running system. An example deployment is shown in Figure 26, where blue dots represent nodes in the target system and the other nodes make up the SD-Mon infrastructure. The network is deployed on the basis of a configuration file describing the network architecture in terms of hosts, Erlang nodes, global group and *s_group* partitions. Tracing to be performed on monitored nodes is also specified within the configuration file.

An agent is started by a master SD-Mon node for each *s_group* and for each free node. Configured tracing is applied on every monitored node, and traces are stored in binary format in the agent file system. The shadow network follows system changes so that agents are started and stopped at runtime as required, as shown in Figure 27. Such changes are persistently stored so that the last configuration can be reproduced after a restart. Of course, the shadow network can be always updated via the User Interface.

Each agent takes care of an *s_group* or of a free node. At start-up it tries to get in contact with its nodes and apply the tracing to them as stated by the master. Binary files are stored in the host file system. Tracing is internally used in order to track *s_group* operations happening at runtime. An asynchronous message is sent to the master whenever one of these changes occurs. Since each process can only be traced by a single process at a time, each node (included those belonging to more than one *s_group*) is controlled by only one agent. When a node is removed from a group or when the group is

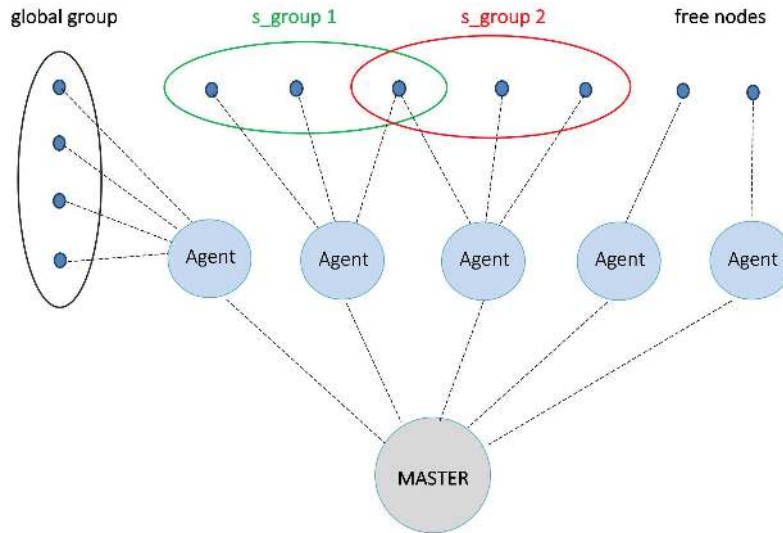


Figure 26: SD-Mon architecture.

deleted, another agent takes over, as shown in Figure 27. When an agent is stopped, all traces on the controlled nodes are switched off.

The monitoring network is also supervised, in order to take account of network fragility, and when an agent node goes down another node is deployed to play its role; there are also periodic consistency checks for the system as a whole, and when an inconsistency is detected then that part of the system can be restarted.

SD-Mon does more than monitor activities one node at a time. In particular inter-node and inter-group messages are displayed at runtime. As soon as an agent is stopped, the related tracing files are fetched across the network by the master and they are made available in a readable format in the master file system.

SD-Mon provides facilities for online visualisation of this data, as well as *post hoc* offline analysis. Figure 28 shows, in real time, messages that are sent *between* s_groups. This data can also be used as input to the animated Devo visualisation, as illustrated in the right-hand side of Figure 25.




Figure 27: SD-Mon evolution. Before (top) and after eliminating s_group 2 (bottom).

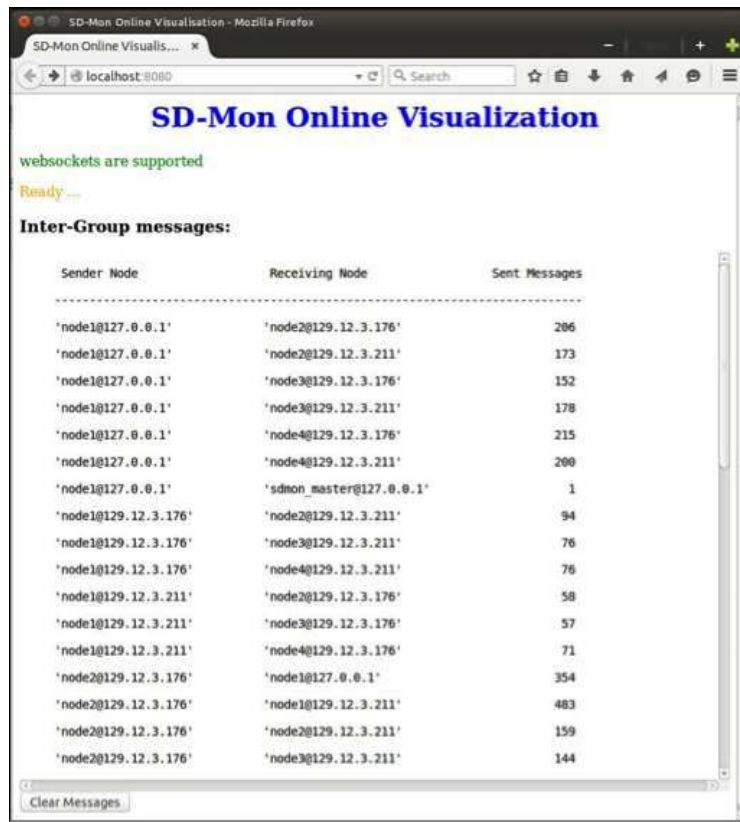


Figure 28: SD-Mon: online monitoring.

VIII. SYSTEMIC EVALUATION

Preceding sections have investigated the improvements of individual aspects of an Erlang system, e.g. ETS tables in Section i. This section analyses the impact of the new tools and technologies from Sections V, VI, VII in concert. We do so by investigating the deployment, reliability, and scalability of the Orbit and ACO benchmarks from Section III. The experiments reported here are representative. Similar experiments show consistent results for a range of micro-benchmarks, several benchmarks, and the very substantial (approximately 150K lines of Erlang) Sim-Diasca case study [16] on several state of the art NUMA architectures, and

the four clusters specified in Appendix A. A coherent presentation of many of these results is available in an article by [22] and in a RELEASE project deliverable²⁰. The bulk of the experiments reported here are conducted on the Athos cluster using Erlang/OTP 17.4 and the associated SD Erlang libraries.

The experiments cover two measures of scalability. As Orbit does a fixed size computation, the scaling measure is relative speedup (or strong scaling), i.e. speedup relative to execution time on a single core. As the work in ACO increases with compute resources, weak scaling is the appropriate measure. The benchmarks

²⁰See Deliverable 6.2, available online. <http://www.release-project.eu/documents/D6.2.pdf>




Figure 29: Speedup of distributed Erlang vs. SD Erlang Orbit (2M and 5M elements) [Strong Scaling].

also evaluate different aspects of `s_groups`: Orbit evaluates the scalability impacts of network connections, while ACO evaluates the impact of both network connections and the global namespace required for reliability.

i. Orbit

Figure 29 shows the speedup of the D-Orbit and SD-Orbit benchmarks from Section i.1. The measurements are repeated seven times, and we plot standard deviation. Results show that D-Orbit performs better on a small number of nodes as communication is direct, rather than via a gateway node. As the number of nodes grows, however, SD-Orbit delivers better speedups, i.e. beyond 80 nodes in case of 2M orbit elements, and beyond 100 nodes in case of 5M orbit elements. When we increase the size of Orbit beyond 5M, the D-Orbit version fails due to the fact that some VMs exceed the available RAM of 64GB. In contrast SD-Orbit experiments run successfully even for an orbit with 60M elements.

ii. Ant Colony Optimisation (ACO)

Deployment The deployment and monitoring of ACO, and of the large (150K lines of Erlang) Sim-Diasca simulation using WombatOAM (Section ii) on the Athos cluster is detailed in [22].

An example experiment deploys 10,000 Erlang nodes without enabling monitoring, and hence allocates three nodes per core (i.e. 72 nodes on each of 139 24-core Athos hosts). Figure 30 shows that WombatOAM deploys the nodes in 212s, which is approximately 47 nodes per second. It is more common to have at least one core per Erlang node, and in a related experiment 5000 nodes are deployed, one per core (i.e. 24 nodes on each of 209 24-core Athos hosts) in 101s, or 50 nodes per second. Crucially in both cases the deployment time is linear in the number of nodes. The deployment time could be reduced to be logarithmic in the number of nodes using standard divide-and-conquer techniques. However there hasn't been the demand to do so as most Erlang systems are long running servers.

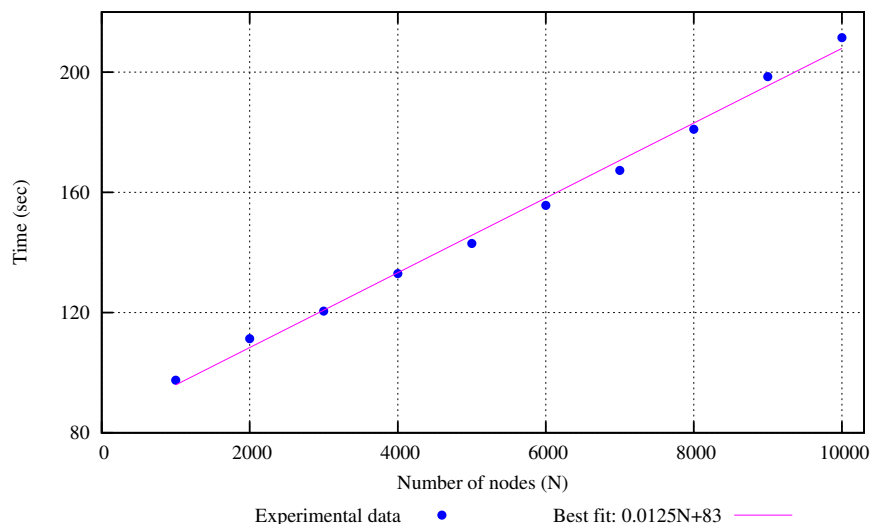


Figure 30: *Wombat Deployment Time*

The measurement data shows two important facts: it shows that WombatOAM scales well (up to a deployment base of 10,000 Erlang nodes), and that WombatOAM is non-intrusive because its overhead on a monitored node is typically less than 1.5% of effort on the node.

We conclude that WombatOAM is capable of deploying and monitoring substantial distributed Erlang and SD Erlang programs. The experiments in the remainder of this section use standard distributed Erlang configuration file deployment.

Reliability SD Erlang changes the organisation of processes and their recovery data at the language level, so we seek to show that these changes have not disrupted Erlang’s world-class reliability mechanisms at this level. As we haven’t changed them we don’t exercise Erlang’s other reliability mechanisms, e.g. those for managing node failures, network congestion, etc. A more detailed study of SD Erlang reliability, including the use of replicated databases for recovering Instant Messenger chat sessions, finds similar results [23].

We evaluate the reliability of two ACO versions using a Chaos Monkey service that kills processes in the running system at random [14]. Recall that GR-ACO provides reliability by registering the names of critical processes globally, and SR-ACO registers them only within an `s_group` (Section ii).

For both GR-ACO and SR-ACO a Chaos Monkey runs on every Erlang node, i.e. master, submasters, and colony nodes, killing a random Erlang process every second. Both ACO versions run to completion. Recovery, at this failure frequency, has no measurable impact on runtime. This is because processes are recovered within the Virtual machine using (globally synchronised) local recovery information. For example, on a common X86/Ubuntu platform typical Erlang process recovery times are around 0.3ms, so around 1000x less than the Unix process recovery time on the same platform [58]. We have conducted more detailed experiments on an Instant Messenger benchmark, and obtained similar results [23].

We conclude that both GR-ACO and SR-ACO are reliable, and that SD Erlang preserves




Figure 31: ACO execution times, Erlang/OTP 17.4 (RELEASE) [Weak Scaling].

distributed Erlang reliability model. The remainder of the section outlines the impact of maintaining the recovery information required for reliability on scalability.

Scalability Figure 31 compares the runtimes of the ML, GR, and SR versions of ACO (Section ii) on Erlang/OTP 17.4(RELEASE). As outlined in Section i.1, GR-ACO not only maintains a fully connected graph of nodes, it registers process names for reliability, and hence scales significantly worse than the unreliable ML-ACO. We conclude that providing reliability with standard distributed Erlang process registration dramatically limits scalability.

While ML-ACO does not provide reliability, and hence doesn't register process names, it maintains a fully connected graph of nodes which limits scalability. SR-ACO, that maintains connections and registers process names only within `s_groups` scales best of all. Figure 31 illustrates how maintaining the process namespace, and fully connected network, impacts performance. This reinforces the evidence from the Orbit benchmarks, and oth-

ers, that partitioning the network of Erlang nodes significantly improves performance at large scale.

To investigate the impact of SD Erlang on network traffic, we measure the number of sent and received packets on the GPG cluster for three versions of ACO: ML-ACO, GR-ACO, and SR-ACO. Figure 32 shows the total number of sent packets. The highest traffic (the red line) belongs to the GR-ACO and the lowest traffic belongs to the SR-ACO (dark blue line). This shows that SD Erlang significantly reduces the network traffic between Erlang nodes. Even with the `s_group` name registration SR-ACO has less network traffic than ML-ACO that has no global name registration.

iii. Evaluation Summary

We have shown that WombatOAM is capable of deploying and monitoring substantial distributed Erlang and SD Erlang programs like ACO and Sim-Diasca. The Chaos Monkey experiments with GR-ACO and SR-ACO show that both are reliable, and hence that SD Erlang




Figure 32: Number of sent packets in ML-ACO, GR-ACO, and SR-ACO.

preserves the distributed Erlang language-level reliability model.

As SD Orbit scales better than D-Orbit, SR-ACO scales better than ML-ACO, and SR-ACO has significantly less network traffic, we conclude that, even when global recovery data is not maintained, partitioning the fully-connected network into `s_groups` reduces network traffic and improves performance. While the distributed Orbit instances ($W=2M$) and ($W=5M$) reach scalability limits at around 40 and 60 nodes, Orbit scales to 150 nodes on SD Erlang (limited by input size), and SR-ACO is still scaling well on 256 nodes (6144 cores). Hence not only have we exceeded the 60 node scaling limits of distributed Erlang identified in Section ii, we have not reached the scaling limits of SD Erlang on this architecture.

Comparing GR-ACO and ML-ACO scalability curves shows that maintaining global recovery data, i.e. a process name space, dramatically limits scalability. Comparing GR-ACO and SR-ACO scalability curves shows that scalability can be recovered by partitioning the nodes into appropriately-sized `s_groups`,

and hence maintaining the recovery data only within a relatively small group of nodes. These results are consistent with other experiments.

IX. DISCUSSION

Distributed actor platforms like Erlang, or Scala with Akka, are a common choice for internet-scale system architects as the model, with its automatic, and VM-supported reliability mechanisms makes it *extremely easy to engineer scalable reliable systems*. Targeting emergent server architectures with hundreds of hosts and tens of thousands of cores, we report a systematic effort to improve the scalability of a leading distributed actor language, while preserving reliability. The work is a *vade mecum* for addressing scalability of reliable actor languages and frameworks. It is also high impact, with downloads of our improved Erlang/OTP running at 50K a month.

We have undertaken the first systematic study of scalability in a distributed actor language, covering VM, language and persistent storage levels. We have developed the

BenchErl and DE-Bench tools for this purpose. Key VM-level scalability issues we identify include contention for shared ETS tables and for commonly-used shared resources like timers. Key language scalability issues are the costs of maintaining a fully-connected network, maintaining global recovery information, and explicit process placement. Unsurprisingly the scaling issues for this distributed actor language are common to other distributed or parallel languages and frameworks with other paradigms like CHARM++ [45], Cilk [15], or Legion [40]. We establish scientifically the folklore limitations of around 60 connected nodes for distributed Erlang (Section IV).

The actor model is no panacea, and there can still be scalability problems in the algorithms that we write, either within a single actor or in the way that we structure communicating actors. A range of pragmatic issues also impact the performance and scalability of actor systems, including memory occupied by processes (even when quiescent), mailboxes filling up, etc. Identifying and resolving these problems is where tools like Percept2 and WombatOAM are needed. However, *many of the scalability issues arise where Erlang departs from the private state principle of the actor model, e.g. in maintaining shared state in ETS tables, or a shared global process namespace for recovery.*

We have designed and implemented a set of Scalable Distributed (SD) Erlang libraries to address *language-level* scalability issues. The key constructs are `s_groups` for partitioning the network and global process namespace, and semi-explicit process placement for deploying distributed Erlang applications on large heterogeneous architectures in a portable way. We have provided a state transition operational semantics for the new `s_groups`, and validated the library implementation against the semantics using QuickCheck (Section V).

To improve the scalability of the Erlang VM and libraries we have improved the implementation of shared ETS tables, time management

and load balancing between schedulers. Following a systematic analysis of ETS tables, the number of fine-grained (bucket) locks and of reader groups have been increased. We have developed and evaluated four new techniques for improving ETS scalability: (i) programmer control of number of bucket locks; (ii) a contention-adapting tree data structure for `ordered_sets`; (iii) queue delegation locking; and (iv) eliminating the locks in the meta table. We have introduced a new scheduler utilisation balancing mechanism to spread work to multiple schedulers (and hence cores), and new synchronisation mechanisms to reduce contention on the widely-used time management mechanisms. By June 2015, with Erlang/OTP 18.0, the majority of these changes had been included in the primary releases. In any scalable actor language implementation such thoughtful design and engineering will be required to schedule large numbers of actors on hosts with many cores, and to minimise contention on shared VM resources (Section VI).

To facilitate the development of large Erlang systems, and to make them understandable we have developed a range of tools. The proprietary WombatOAM tool deploys and monitors large distributed Erlang systems over multiple, and possibly heterogeneous, clusters or clouds. We have made open source releases of four concurrency tools: Percept2 now detects concurrency bad smells; Wrangler provides enhanced concurrency refactoring; the Devo tool is enhanced to provide interactive visualisation of SD Erlang systems; and the new SD-Mon tool monitors SD Erlang systems. We anticipate that these tools will guide the design of tools for other large scale distributed actor languages and frameworks (Section VII).

We report on the reliability and scalability implications of our new technologies using a range of benchmarks, and consistently use the Orbit and ACO benchmarks throughout the article. While we report measurements on a range of NUMA and cluster architectures, the

Table 2: Cluster Specifications (RAM per host in GB).

Name	Hosts	Cores		Processor	RAM	Inter-connection
		per host	total avail.			
GPG	20	16	320	320	Intel Xeon E5-2640v2 8C, 2GHz	64 10GB Ethernet
Kalkyl	348	8	2,784	1,408	Intel Xeon 5520v2 4C, 2.26GHz	24–72 InfiniBand 20 Gb/s
TinTin	160	16	2,560	2,240	AMD Opteron 6220v2 Bulldozer 8C, 3.0GHz	64–128 2:1 oversubscribed QDR Infiniband
Athos	776	24	18,624	6,144	Intel Xeon E5-2697v2 12C, 2.7GHz	64 Infiniband FDR14

key scalability experiments are conducted on the Athos cluster with 256 hosts (6144 cores). Even when global recovery data is not maintained, partitioning the network into `s_groups` reduces network traffic and improves the performance of the Orbit and ACO benchmarks above 80 hosts. Crucially we exceed the 60 node limit for distributed Erlang and do not reach the scalability limits of SD Erlang with 256 nodes/VMs and 6144 cores. Chaos Monkey experiments show that two versions of ACO are reliable, and hence that SD Erlang preserves the Erlang reliability model. However the ACO results show that maintaining global recovery data, i.e. a global process name space, dramatically limits scalability in distributed Erlang. Scalability can, however, be recovered by maintaining recovery data only within appropriately sized `s_groups`. These results are consistent with experiments with other benchmarks and on other architectures (Section VIII).

In future work we plan to incorporate RELEASE technologies, along with other technologies in a generic framework for building performant large scale servers. In addition, pre-

liminary investigations suggest that some SD Erlang ideas could improve the scalability of other actor languages.²¹ For example the Akka framework for Scala could benefit from semi-explicit placement, and Cloud Haskell from partitioning the network.

APPENDIX A: ARCHITECTURE SPECIFICATIONS

The specifications of the clusters used for measurement are summarised in Table 2. We also use the following NUMA machines. (1) An AMD Bulldozer with 16M L2/16M L3 cache, 128GB RAM, four AMD Opteron 6276s at 2.3 GHz, 16 “Bulldozer” cores each, giving a total of 64 cores. (2) An Intel NUMA with 128GB RAM, four Intel Xeon E5-4650s at 2.70GHz, each with eight hyperthreaded cores, giving a total of 64 cores.

²¹See Deliverable 6.7 (<http://www.release-project.eu/documents/D6.7.pdf>), available online.

ACKNOWLEDGEMENTS

We would like to thank the entire RELEASE project team for technical insights and administrative support. Roberto Aloï and Enrique Fernandez Casado contributed to the development and measurement of WombatOAM. 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] Gul Agha. *ACTORS: A Model of Concurrent Computation in Distributed Systems*. PhD thesis, MIT, 1985.
- [2] Gul Agha. An overview of Actor languages. *SIGPLAN Not.*, 21(10):58–67, 1986.
- [3] Bulldozer (microarchitecture), 2015. [https://en.wikipedia.org/wiki/Bulldozer_\(microarchitecture\)](https://en.wikipedia.org/wiki/Bulldozer_(microarchitecture)).
- [4] Apache SF. Libcloud, 2016. <https://libcloud.apache.org/>.
- [5] C. R. Aragon and R. G. Seidel. Randomized search trees. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pages 540–545, October 1989.
- [6] Joe Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf, 2007.
- [7] Joe Armstrong. Erlang. *Commun. ACM*, 53:68–75, 2010.
- [8] Stavros Aronis, Nikolaos Papaspyrou, Katerina Roukounaki, Konstantinos Sagonas, Yiannis Tsiouris, and Ioannis E. Venetis. A scalability benchmark suite for Erlang/OTP. In Torben Hoffman and John Hughes, editors, *Proceedings of the Eleventh ACM SIGPLAN Workshop on Erlang*, pages 33–42, New York, NY, USA, September 2012. ACM.
- [9] Thomas Arts, John Hughes, Joakim Johansson, and Ulf Wiger. Testing telecoms software with Quviq QuickCheck. In *Proceedings of the 2006 ACM SIGPLAN workshop on Erlang*, pages 2–10, New York, NY, USA, 2006. ACM.
- [10] Robert Baker, Peter Rodgers, Simon Thompson, and Huiqing Li. Multi-level Visualization of Concurrent and Distributed Computation in Erlang. In *Visual Languages and Computing (VLC): The 19th International Conference on Distributed Multimedia Systems (DMS 2013)*, 2013.
- [11] Luiz André Barroso, Jimmy Clidaras, and Urs Hölzle. *The Datacenter as a Computer*. Morgan and Claypool, 2nd edition, 2013.
- [12] Basho Technologies. Riakdocs. Basho Bench, 2014. <http://docs.basho.com/riak/latest/ops/building/benchmarking/>.
- [13] J. E. Beasley. OR-Library: Distributing test problems by electronic mail. *The Journal of the Operational Research Society*, 41(11):1069–1072, 1990. Datasets available at <http://people.brunel.ac.uk/~mastjjb/jeb/orlib/wtinfo.html>.
- [14] Cory Bennett and Ariel Tseitlin. Chaos monkey released into the wild. Netflix Blog, 2012.
- [15] Robert D Blumofe, Christopher F Joerg, Bradley C Kuszmaul, Charles E Leiserson, Keith H Randall, and Yuli Zhou. Cilk: An efficient multithreaded runtime system. In

- Proceedings of the fifth ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 207–216. ACM, 1995.
- [16] Olivier Boudeville. Technical manual of the sim-diasca simulation engine. *EDF R&D*, 2012.
- [17] István Bozó, Viktória Fördös, Dániel Horpácsi, Zoltán Horváth, Tamás Kozsik, Judit Kőszegi, and Melinda Tóth. Refactorings to enable parallelization. In Jurriaan Hage and Jay McCarthy, editors, *Trends in Functional Programming: 15th International Symposium, TFP 2014. Revised Selected Papers*, volume 8843 of *LNCS*, pages 104–121. Springer, 2015.
- [18] Irina Calciu, Dave Dice, Yossi Lev, Victor Luchangco, Virendra J. Marathe, and Nir Shavit. NUMA-aware reader-writer locks. In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 157–166, New York, NY, USA, 2013. ACM.
- [19] Francesco Cesarini and Simon Thompson. *Erlang Programming: A Concurrent Approach to Software Development*. O’Reilly Media, 1st edition, 2009.
- [20] Rohit Chandra, Leonardo Dagum, Dave Kohr, Dror Maydan, Jeff McDonald, and Ramesh Menon. *Parallel Programming in OpenMP*. Morgan Kaufmann, San Francisco, CA, USA, 2001.
- [21] Natalia Chechina, Huiqing Li, Amir Ghaffari, Simon Thompson, and Phil Trinder. Improving the network scalability of Erlang. *J. Parallel Distrib. Comput.*, 90(C):22–34, 2016.
- [22] Natalia Chechina, Kenneth MacKenzie, Simon Thompson, Phil Trinder, Olivier Boudeville, Viktória Fördös, Csaba Hoch, Amir Ghaffari, and Mario Moro Hernandez. Evaluating scalable distributed Erlang for scalability and reliability. *IEEE Transactions on Parallel and Distributed Systems*, 2017.
- [23] Natalia Chechina, Mario Moro Hernandez, and Phil Trinder. A scalable reliable instant messenger using the SD Erlang libraries. In *Erlang’16*, pages 33–41, New York, NY, USA, 2016. ACM.
- [24] Koen Claessen and John Hughes. QuickCheck: a lightweight tool for random testing of Haskell programs. In *Proceedings of the 5th ACM SIGPLAN International Conference on Functional Programming*, pages 268–279, New York, NY, USA, 2000. ACM.
- [25] H. A. J. Crauwels, C. N. Potts, and L. N. van Wassenhove. Local search heuristics for the single machine total weighted tardiness scheduling problem. *INFORMS Journal on Computing*, 10(3):341–350, 1998. The datasets from this paper are included in Beasley’s ORLIB.
- [26] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
- [27] David Dewolfs, Jan Broeckhove, Vaidy Sunderam, and Graham E. Fagg. FT-MPI, fault-tolerant metacomputing and generic name services: A case study. In *EuroPVM/MPI’06*, pages 133–140, Berlin, Heidelberg, 2006. Springer-Verlag.
- [28] Alan AA Donovan and Brian W Kernighan. *The Go programming language*. Addison-Wesley Professional, 2015.
- [29] Marco Dorigo and Thomas Stützle. *Ant Colony Optimization*. Bradford Company, Scituate, MA, USA, 2004.
- [30] Faith Ellen, Yossi Lev, Victor Luchangco, and Mark Moir. SNZI: scalable NonZero

- indicators. In *Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing*, pages 13–22, New York, NY, USA, 2007. ACM.
- [31] Jeff Epstein, Andrew P. Black, and Simon Peyton-Jones. Towards Haskell in the Cloud. In *Haskell '11*, pages 118–129. ACM, 2011.
- [32] Martin Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [33] Ana Gainaru and Franck Cappello. Errors and faults. In *Fault-Tolerance Techniques for High-Performance Computing*, pages 89–144. Springer International Publishing, 2015.
- [34] Martin Josef Geiger. New instances for the Single Machine Total Weighted Tardiness Problem. Technical Report Research Report 10-03-01, March 2010. Datasets available at <http://logistik.hsu-hh.de/SMTWTP>.
- [35] Guillaume Germain. Concurrency oriented programming in termite scheme. In *Proceedings of the 2006 ACM SIGPLAN workshop on Erlang*, pages 20–20, New York, NY, USA, 2006. ACM.
- [36] Amir Ghaffari. De-bench, a benchmark tool for distributed erlang, 2014. <https://github.com/amirghaffari/DEbench>.
- [37] Amir Ghaffari. Investigating the scalability limits of distributed Erlang. In *Proceedings of the 13th ACM SIGPLAN Workshop on Erlang*, pages 43–49. ACM, 2014.
- [38] Amir Ghaffari, Natalia Chechina, Phip Trinder, and Jon Meredith. Scalable persistent storage for Erlang: Theory and practice. In *Proceedings of the 12th ACM SIGPLAN Workshop on Erlang*, pages 73–74, New York, NY, USA, 2013. ACM.
- [39] Seth Gilbert and Nancy Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33:51–59, 2002.
- [40] Andrew S Grimshaw, Wm A Wulf, and the Legion team. The Legion vision of a worldwide virtual computer. *Communications of the ACM*, 40(1):39–45, 1997.
- [41] Carl Hewitt. Actor model for discretionary, adaptive concurrency. *CoRR*, abs/1008.1459, 2010.
- [42] Carl Hewitt, Peter Bishop, and Richard Steiger. A universal modular ACTOR formalism for artificial intelligence. In *IJCAI’73*, pages 235–245, San Francisco, CA, USA, 1973. Morgan Kaufmann.
- [43] Rich Hickey. The Clojure programming language. In *DLS’08*, pages 1:1–1:1, New York, NY, USA, 2008. ACM.
- [44] Zoltán Horváth, László Lövei, Tamás Kozsik, Róbert Kitlei, Anikó Nagyné Víg, Tamás Nagy, Melinda Tóth, and Roland Király. Building a refactoring tool for Erlang. In *Workshop on Advanced Software Development Tools and Techniques, WAS-DETT 2008*, 2008.
- [45] Laxmikant V Kale and Sanjeev Krishnan. Charm++: a portable concurrent object oriented system based on c++. In *ACM Sigplan Notices*, volume 28, pages 91–108. ACM, 1993.
- [46] David Klaftenegger, Konstantinos Sagonas, and Kjell Winblad. On the scalability of the Erlang term storage. In *Proceedings of the Twelfth ACM SIGPLAN Workshop on Erlang*, pages 15–26, New York, NY, USA, September 2013. ACM.
- [47] David Klaftenegger, Konstantinos Sagonas, and Kjell Winblad. Delegation locking libraries for improved performance of

- multithreaded programs. In *Euro-Par 2014 Parallel Processing.*, volume 8632 of *LNCS*, pages 572–583. Springer, 2014.
- [48] David Klaftenegger, Konstantinos Sagonas, and Kjell Winblad. Queue delegation locking. *IEEE Transactions on Parallel and Distributed Systems*, 2017. To appear.
- [49] Rusty Klophaus. Riak core: Building distributed applications without shared state. In *ACM SIGPLAN Commercial Users of Functional Programming, CUFPP ’10*, pages 14:1–14:1, New York, NY, USA, 2010. ACM.
- [50] Avinash Lakshman and Prashant Malik. Cassandra: A decentralized structured storage system. *SIGOPS Oper. Syst. Rev.*, 44(2):35–40, April 2010.
- [51] J. Lee et al. Python actor runtime library, 2010. osl.cs.uiui.edu/parley/.
- [52] Huiqing Li and Simon Thompson. Automated API Migration in a User-Extensible Refactoring Tool for Erlang Programs. In Tim Menzies and Motoshi Saeki, editors, *Automated Software Engineering, ASE’12*. IEEE Computer Society, 2012.
- [53] Huiqing Li and Simon Thompson. Multi-core profiling for Erlang programs using Percept2. In *Proceedings of the Twelfth ACM SIGPLAN workshop on Erlang*, 2013.
- [54] Huiqing Li and Simon Thompson. Improved Semantics and Implementation Through Property-Based Testing with QuickCheck. In *9th International Workshop on Automation of Software Test*, 2014.
- [55] Huiqing Li and Simon Thompson. Safe Concurrency Introduction through Slicing. In *PEPM ’15*. ACM SIGPLAN, January 2015.
- [56] Huiqing Li, Simon Thompson, György Orosz, and Melinda Tóth. Refactoring with Wrangler, updated. In *ACM SIGPLAN Erlang Workshop*, volume 2008, 2008.
- [57] Frank Lubeck and Max Neunhoffer. Enumerating large Orbits and direct condensation. *Experimental Mathematics*, 10(2):197–205, 2001.
- [58] Andreea Lutac, Natalia Chechina, Gerardo Aragon-Camarasa, and Phil Trinder. Towards reliable and scalable robot communication. In *Proceedings of the 15th International Workshop on Erlang*, Erlang 2016, pages 12–23, New York, NY, USA, 2016. ACM.
- [59] LWN.net. The high-resolution timer API, January 2006. <https://lwn.net/Articles/167897/>.
- [60] Kenneth MacKenzie, Natalia Chechina, and Phil Trinder. Performance portability through semi-explicit placement in distributed Erlang. In *Proceedings of the 14th ACM SIGPLAN Workshop on Erlang*, pages 27–38. ACM, 2015.
- [61] Jeff Matocha and Tracy Camp. A taxonomy of distributed termination detection algorithms. *Journal of Systems and Software*, 43(221):207–221, 1998.
- [62] Nicholas D Matsakis and Felix S Klock II. The Rust language. In *ACM SIGAda Ada Letters*, volume 34, pages 103–104. ACM, 2014.
- [63] Robert McNaughton. Scheduling with deadlines and loss functions. *Management Science*, 6(1):1–12, 1959.
- [64] Martin Odersky et al. The Scala programming language, 2012. www.scala-lang.org.
- [65] William F. Opdyke. *Refactoring Object-Oriented Frameworks*. PhD thesis, University of Illinois at Urbana-Champaign, 1992.

- [66] Nikolaos Papaspyrou and Konstantinos Sagonas. On preserving term sharing in the Erlang virtual machine. In Torben Hoffman and John Hughes, editors, *Proceedings of the 11th ACM SIGPLAN Erlang Workshop*, pages 11–20. ACM, 2012.
- [67] RELEASE Project Team. EU Framework 7 Project 287510 (2011–2015), 2015. <http://www.release-project.eu>.
- [68] Konstantinos Sagonas and Thanassis Avgerinos. Automatic refactoring of Erlang programs. In *Principles and Practice of Declarative Programming (PPDP'09)*, pages 13–24. ACM, 2009.
- [69] Konstantinos Sagonas and Kjell Winblad. More scalable ordered set for ETS using adaptation. In *Proc. ACM SIGPLAN Workshop on Erlang*, pages 3–11. ACM, September 2014.
- [70] Konstantinos Sagonas and Kjell Winblad. Contention adapting search trees. In *Proceedings of the 14th International Symposium on Parallel and Distributed Computing*, pages 215–224. IEEE Computing Society, 2015.
- [71] Konstantinos Sagonas and Kjell Winblad. Efficient support for range queries and range updates using contention adapting search trees. In *Languages and Compilers for Parallel Computing - 28th International Workshop*, volume 9519 of LNCS, pages 37–53. Springer, 2016.
- [72] Marc Snir, Steve W. Otto, D. W. Walker, Jack Dongarra, and Steven Huss-Lederman. *MPI: The Complete Reference*. MIT Press, Cambridge, MA, USA, 1995.
- [73] Sriram Srinivasan and Alan Mycroft. Kilim: Isolation-typed actors for Java. In *ECOOP'08*, pages 104–128, Berlin, Heidelberg, 2008. Springer-Verlag.
- [74] Don Syme, Adam Granicz, and Antonio Cisternino. *Expert F# 4.0*. Springer, 2015.
- [75] Simon Thompson and Huiqing Li. Refactoring tools for functional languages. *Journal of Functional Programming*, 23(3), 2013.
- [76] Marcus Völker. Linux timers, May 2014.
- [77] WhatsApp, 2015. <https://www.whatsapp.com/>.
- [78] Tom White. *Hadoop - The Definitive Guide: Storage and Analysis at Internet Scale (3. ed., revised and updated)*. O'Reilly, 2012.
- [79] Ulf Wiger. Industrial-Strength Functional programming: Experiences with the Ericsson AXD301 Project. In *Implementing Functional Languages (IFL'00)*, Aachen, Germany, September 2000. Springer-Verlag.