Noname manuscript No.
(will be inserted by the editor)

The Missing Link! A new Skeleton for Evolutionary Multi-Agent Systems in Erlang

Jan Stypka · Wojciech Turek · Aleksander Byrski · Marek Kisiel-Dorohinicki · Adam D. Barwell · Chris Brown · Kevin Hammond · Vladimir Janjic

Received: date / Accepted: date

Abstract Evolutionary multi-agent systems (EMAS) play a critical role in many artificial intelligence applications that are in use today. In this paper, we present a new generic skeleton for parallel EMAS computations, written in Erlang. The skeleton enables us to capture a wide variety of concrete evolutionary computations that can exploit the same underlying parallel implementation. We demonstrate the use of our skeleton on two different evolutionary computing applications: i) computing the minimum of the Rastrigin function; and ii) solving an urban traffic optimization problem. We show that we can obtain very good speedups (up to 142.44× the sequential performance) on a variety of different parallel hardware, while requiring very little parallelisation effort.

Keywords Multi-core programming · Erlang · Agent-based computing · Metaheuristics · Many-core programming · Algorithmic Skeletons

1 Introduction

Evolutionary Multi-Agent Systems [9] (EMAS) are a critical part of many modern artificial intelligence systems. The computations that they perform are usually very expensive and, since they involve individual autonomous entities (agents) with no central authority involved in their operation, they should be highly amenable to parallelisation. Desipite this, there has been relatively

J. Stypka, W. Turek, A. Byrski, M. Kisiel-Dorohinicki
AGH University of Science and Technology
Al. Mickiewicza 30, 30-059 Krakow, Poland
E-mail: janstypka@gmail.com, {wojciech.turek,olekb,doroh}@agh.edu.pl

A.D. Barwell, C. Brown, K. Hammond, V. Janjic
School of Computer Science
The University of St Andrews
St Andrews, U.K. E-mail: {adb23,cmb21,kh8,vj32}@st-andrews.ac.uk

little effort spent to date on developing parallel solutions for EMASs and the solutions that are available are usually tailored to a specific instance of a more general problem. At the same time, we are witnessing the emergence of many-core systems across various levels of the computing spectrum, from low-power devices targetting data-centres on a chip/cyber-physical systems all the way to high-performance supercomputers. These systems offer significant processing potential and are ideally suited to improve the performance of EMAS. However, harnessing that potential is still a huge issue, both for EMAS and for many other applications. This is mainly due to the widespread use of low-level native parallel programming models that are commonly used to program these systems (e.g. pThreads and OpenMP).

In this paper, we present a new implementation of an algorithmic skeleton for evolutionary multi-agent systems. Algorithmic skeletons are implementations of common parallel patterns, parameterised over worker functions and other implementation-specific information. A skeleton can be specialised to a specific problem by providing concrete instantiations of worker functions and other required information. Our EMAS skeleton is implemented in the functional programming language Erlang. Erlang is widely used in the telecommunications industry, but is also beginning to be used more widely for high-reliability/ highly-concurrent systems, e.g. databases [30], AOL's *Marketplace by Adtech* [34] and WhatsApp [31]. It has excellent support for concurrency and distribution, including built-in fault tolerance. The latter is critical for long-running computations such as evolutionary computations. Functional programming approaches naturally provide high-level abstractions through e.g. higher-order functions and Erlang was therefore an ideal choice for developing a new generic parallel skeleton.

The specific research contributions of this paper are

1. We design and implement three versions of a new *domain-specific parallel skeleton* for Evolutionary Multi-Agent Systems (EMAS) using Erlang;
2. We implement two realistic evolutionary computing use cases that use our skeleton;
3. We evaluate these use cases on a number of different parallel architectures, ranging from small-scale low-power systems to large-scale high-performance multicore machines, demonstrating that we can obtain very good performance improvements of up to a factor of 150 over the sequential implementations.

2 Evolutionary Multi Agent Systems and Parallel Skeletons

In Agent-oriented programming (AOP), the construction of the software is based on the concept of specialised, independent, autonomous software *agents*. There are many different agent-base programming frameworks, ranging from general-purpose ones, like RePast [28], MASON [26], MadKit [22], JADE [6] to more specialised ones, like ParadisEO [11] or AgE [29]. Some of these frameworks leverage existing multi-core architectures (e.g. ParadisEO). However,

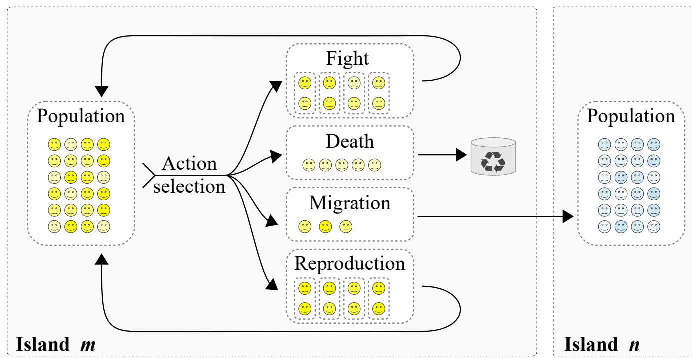

Fig. 1 The General Structure of an Evolutionary Multi-Agent System.

their scalability to larger systems has not been explored. In this paper, we focus on applications where agents use *evolutionary algorithms* in their operation, most frequently for learning/reasoning or coordination of some group activity.

Evolutionary computations [5, 19] use the idea of mimicking the mechanisms underlying biological evolution to solve complex adaptation and optimisation problems. They work on a population of individuals, each of which represents a point in the search space of potential solutions to a given problem. The population undergoes subsequent modification steps by means of randomised genetic operations that model recombination, mutation and selection. Following initialisation, the algorithms loop through these operators until some termination criterion is satisfied. Each of these cycles is called a *generation*. *Evolutionary Multi-Agent Systems* (EMASs) are a hybrid meta-heuristic that combines multi-agent systems with evolutionary algorithms [8, 9] (see Fig. 1). In a multi-agent system, no global knowledge is available to individual agents; agents remain autonomous and no central authority is involved in making their decisions. Therefore, in contrast to traditional evolutionary algorithms, in an evolutionary computing system, *selective pressure* (essentially, a method for selecting which members of the population survive and move to the next generation) needs to be decentralised. Using agent terminology, we can say that selective pressure is required to emerge from peer-to-peer interactions between agents instead of being globally-driven. In EMAS, emergent selective pressure is achieved by giving agents a single non-renewable resource called energy. Agents with high energy are more likely to reproduce whereas agents with low energy are more likely to die. The algorithm is designed to transfer energy from better to worse agents without central control.

In a basic EMAS implementation, every agent is provided with a real-valued vector representing a potential solution to the optimization problem, along with the corresponding fitness metric. Agents start with an initial amount of energy and meet randomly. If their energy is below a death threshold, they die. If it is above some reproduction threshold, they reproduce and yield new

agents – the genotype of the children is derived from their parents using variation operators and some amount of energy is also inherited. If neither of these two conditions is met, agents fight in tournaments by comparing their fitness values. The better (winning) agents then sap energy from the worse (losing) ones [9]. The system as a whole is *stable*, since the total energy remains constant. However, the number of agents may vary in order to adapt to the difficulty of the problem – small numbers of agents with high energy or large numbers of agents with low energy, for example. The number of agents can also be altered dynamically by varying the total energy of the system. As in other evolutionary algorithms, agents can be split into separate populations, or *islands*. Such islands help preserve diversity by introducing allopatric speciation and can also execute in parallel. Information is exchanged between islands through agent migrations. Evolutionary multi-agent systems have been formally shown to be a general optimization tool by constructing a detailed Markov-chain based model and proving its ergodicity [8].

2.1 Parallel Programming in Erlang

Erlang [3] is a strict, impure, functional programming language with built-in support for concurrency. It supports a *lightweight* threading model, where *processes* model small units of computation. The scheduling of processes is handled automatically by the Erlang Virtual Machine, which also provides basic load balancing mechanisms. A key aspect of the Erlang design is that it adopts a *shared-nothing* approach, in which processes do not implicitly share state – all data is communicated via channels, using explicit *send* and *receive* primitives. It has proven to be an efficient tool for building large-scale systems for multi-core processors.

Most of the industrial applications written in Erlang are deployed on systems with up to 24 CPU cores. Scalability of solutions is usually provided by using clusters of multicores and running Erlang in distributed configuration. There are few reports on using a single Erlang Virtual Machine on architectures exceeding 32 physical cores. In [4] the authors present a test suite for measuring different aspects of Erlang applications performance. The exemplary test running on a 64-core machine shows that in most cases the speedup is non-linear and it degrades for high number of cores and schedulers. The problem of Erlang Term Storage scalability on a computer with 32 physical cores have been considered in [32]. Promising results of using Erlang on a Intel Xeon Phi coprocessor have been shown in [35]. Basic benchmarks show good scalability up to 60 cores, which is the number of physical cores of the coprocessor. However, there are hardly any reports on scaling complex, computationally intensive Erlang applications on many-core architectures.

Our main motivation for using Erlang for evolutionary computations is to achieve scalability. While the performance on single-core machines of Erlang is still behind more mature and established languages, like C and C++, it is designed to be easily and transparently scalable across distributed systems. In

the era of many-core architectures, the ability of scaling the performance with the growing number of cores/nodes will become far more important than the effectiveness of single core utilisation.

2.2 Algorithmic Skeletons

Algorithmic skeletons abstract commonly-used patterns of parallel computation, communication, and interaction [15] into parameterised templates. For example, we might define a *parallel map* skeleton, whose functionality is identical to a standard *map* function, but which creates a number of Erlang processes (*worker processes*) to execute each element of the map in parallel. Using a skeleton approach allows the programmer to adopt a top-down *structured* approach to parallel programming, where skeletons are composed to give the overall parallel structure of the program. Details such as communication, task creation, task or data migration, scheduling, etc. are embedded within the skeleton implementation, which may be tailored to a specific architecture or class of architectures. This offers an improved level of portability over typical low-level approaches. A recent survey of algorithmic skeleton approaches can be found in [20].

2.3 The *Skel* Library for Erlang

In our previous work, we have developed the *Skel* [7, 23] library, which defines a small set of classical skeletons for Erlang. Each skeleton operates over a stream of input values, producing a corresponding stream of results. Skel also allows simple composition and nesting of skeletons. Skel supports the folowing skeletons:

- `func` is a simple wrapper skeleton that encapsulates a sequential function as a streaming skeleton.
- `pipe` models a composition of skeletons $s_1, s_2, \ldots, s_n$ over a stream of inputs. Within the pipeline, each of the s_i is applied in parallel to the result of the s_{i-1}.
- `farm` models application of the same operation over a stream of inputs. Each of the n farm *workers* is a skeleton that operates in parallel over independent values of the input stream.
- `cluster` is a data parallel skeleton, where each independent input, x_i can be partitioned into a number of sub parts, $x_1, x_2, \ldots, x_n$, that can be worked upon in parallel. A skeleton, s, is then applied to each element of the sub-stream in parallel.
- `feedback` wraps a skeleton s, feeding the results of applying s back as new inputs to s, provided they match a filter function, f.

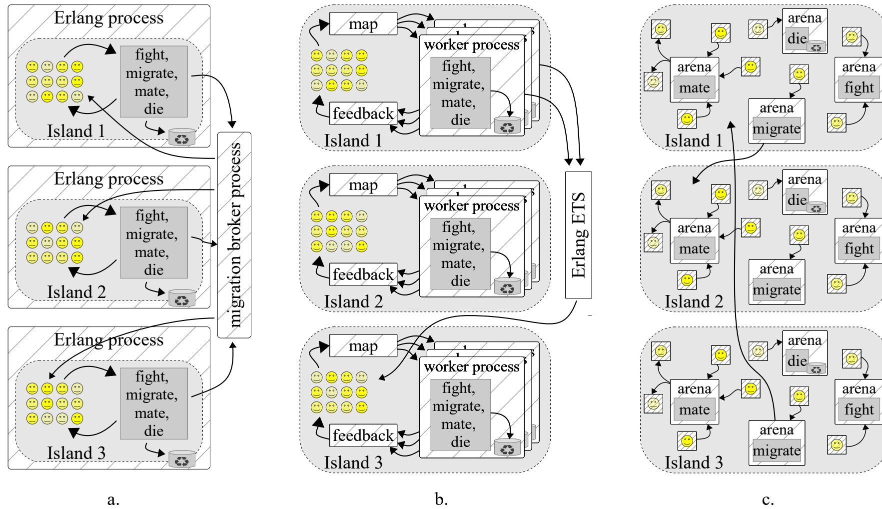

Fig. 2 Three versions of EMAS in Erlang: a. the Hybrid version with sequential islands computed concurrently; b. the Skel version with configurable number of concurrent workers; c. the Concurrent version with each solution represented by autonomous process.

3 A New Skeleton for Evolutionary Multi-Agent Systems

Our EMAS skeleton accepts two types of input:

- a set of parameters for the computation, including the number of agents, initial energy and details about reproduction (e.g. the maximum energy given by a parent to a child), mutation (e.g. the chance of mutation during reproduction) and recombination (e.g. the probability of recombining parent solutions during reproduction);
- problem-specific `solution`, `evaluation`, `mutation` and `recombination` functions.

The skeleton implementation then iterates over these functions for each agent, creating Erlang processes that allow agents to perform their work in parallel. We have implemented three different versions of the skeleton, which differ in the precise way in which agents are captured by Erlang processes:

a. *The Hybrid version* (Figure 2.a.), which represents an optimised version of the *Skel* version (described below). The population of agents is split into islands and every island is contained in a separate Erlang process. Within the process, the `fight`, `migrate`, `mate` and `die` functions are done sequentially, and a separate dedicated Erlang process handles migration of agents. This version represents coarse-grained implementation of the EMAS computation, with a small number of potentially expensive Erlang processes than in the previous two versions;
b. *The Skel version* (Figure 2.b.), which uses the *feedback* parallel pattern from the Skel pattern library (described in Section 2.3). Here, we set a

	pi	*titanic*	*zeus*	*power*	*phi*
Arch	Arm	AMD	AMD	IBM	Intel
Proc	Cortex-A7	Opteron 6176	Opteron 6276	Power8	Xeon Phi 7120
Cores	4	24	64	20	61
Threads	4	24	64	160	244
Freq.	900MHz	2.3GHz	2.3Ghz	3.69GHz	1.2GHz
L2 Cache	256 KB	24 x 512 KB	32 x 2 MB	20 x 512 KB	61 x 512KB
L3 Cache	-	4 x 6 MB	4 x 8 MB	20 x 8 MB	-
RAM	1 GB	32 GB	256 GB	256 GB	16 GB

Fig. 3 Experimental Platforms.

fixed number of Erlang processes that will be created (usually matching the number of OS threads that we want to use) and then distribute agents from a population to these worker processes. When one generation finishes computing, the resulting new population is merged by the feedback process. Separate Erlang processes are created to handle the mutation and recombination parts of the computation (or, in our case, fighting, migrating, mating and dying).

c. *The Concurrent version* (Figure 2.c.), where each agent is represented by a separate Erlang process. The agents use dedicated arena processes to trigger the behavior, mutation and recombination functions. This is a fine-grained implementation of the EMAS computation, where we potentially have a very large number of very small Erlang processes;

Since it is most coarse-grained and has the least synchronisation between processes, we expect *Hybrid* to give the best performance of these three versions.

4 Evaluation

We have evaluated the performance of our EMAS skeleton on two different optimisation examples that typify agent-based evolutionary computations:

1. Finding the minimum of the Rastrigin function, a popular global optimisation benchmark [17];
2. A complex problem of urban traffic optimization, a multi-variant simulation [24] that is applied to the mobile robotics domain. Our model anticipates possible situations on the road, preparing plans to deal with certain situations and apply these plans when certain traffic conditions arise.

Since the algorithms are bounded by time, we measure the average rate of reproductions that the algorithm achieves per second, representing its throughput. We also measure the increase in throughput as the number of cores increases. This corresponds to an increase in the application's performance (speedup). We ran each experiment a total of 10 times, with each experiment set to take 5 minutes. The number of islands that the algorithm is able to use is set to be the number of cores. We prioritised the evaluation of the hybrid

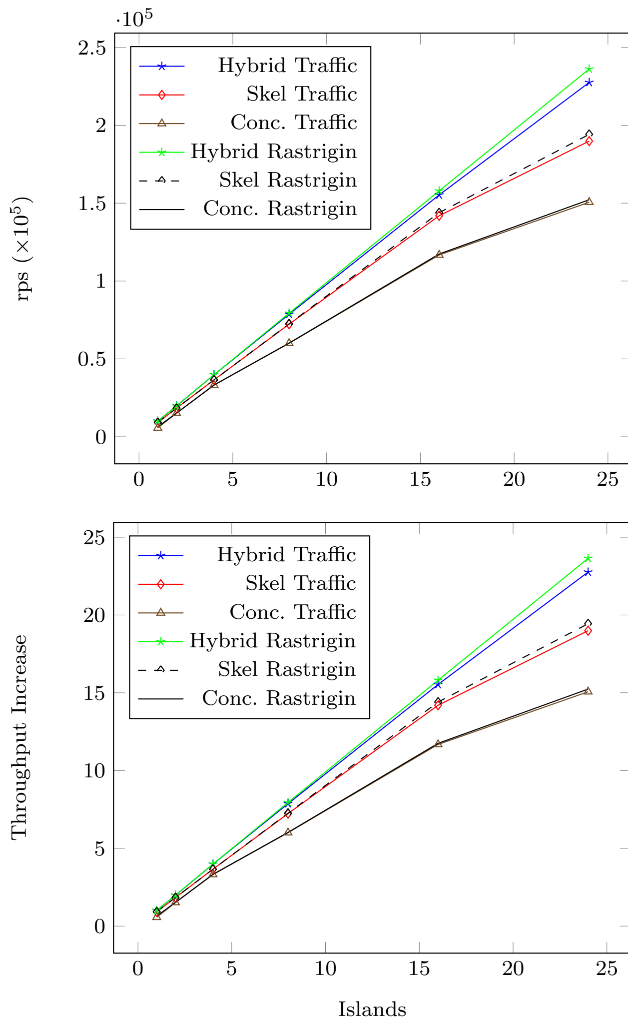

Fig. 4 Reproductions per Second (*rps*, above) and Throughput Increase (below) for Traffic and Rastrigin on *titanic*.

configuration, since it has proven itself to be the best of the three parallel configurations, but have included other configurations in our results where available. For all experiments, we have used version 18 of the Erlang system, except where otherwise highlighted. In order to evaluate the EMAS skeleton on a number of different architectures and in different settings, our experiments have been conducted on five different systems, ranging from low-power small-scale ARM systems (*pi*) through more powerful x86 (*titanic*, *zeus*) and (*power*) OpenPower multicores to a standalone Intel Xeon Phi many-core accelerator. Full details of the systems that we have used are given in Figure 3.

4.1 *titanic* results

The *titanic* system is an example of a medium-scale parallel server, with 24 cores and 32GB of RAM. Figure 4 shows the number of reproductions per second (*rps*) and the increases in throughput for Traffic and Rastrigin on this machine. For both applications, we observe that the Hybrid version achieves almost linear increases in throughput (and, therefore, performance). For Rastrigin, the Hybrid version achieves 236,140 *rps* on 24 cores versus 10,038 *rps* on one core, yielding 23.52× the throughput (or 23.64× the base sequential throughput). For Traffic, the Hybrid version on *titanic* achieves 227,509 *rps* on 24 cores versus 10,059 *rps* on one core, yielding 22.62× the throughput (or 22.76× the base sequential throughput of 9,996 *rps*). The Skel and Concurrent versions initally also perform well, but start to tail off after about 16 cores (islands) for both applications. This is consistent with other applications we have run on this machine, and is probably due to cache contention. Although the Hybrid version is always best, the untuned Skel version always performs better than the Concurrent version. This shows the benefit of reducing the number of Erlang processes by grouping the computations that multiple agents perform into a single process.

4.2 *zeus* results

The *zeus* system is an example of a larger-scale multicore system, with 64 cores and 256 GB of RAM. Figure 5 shows the *rps* and corresponding throughput increase for the Hybrid version only. Here, Rastrigin achieves 349,226 *rps* on 64 cores versus 5,752 *rps* on one core, for a throughput increase of 60.71× the sequential version. Traffic achieves 94,950 *rps* on 64 cores versus 1,473 *rps* on one core, for a throughput increase of 64× the sequential version. There is clearly a tail-off in throughput increase for Traffic at 32 and 48 cores, but further experiments are required to explain the reasons for this, given the clear improvement with 64 cores. A slight dip can also be observed for Rastrigin at 32 cores (27.8× throughput increase), but this has recovered at 48 cores (46.13× throughput increase). In comparison with *titanic*, the raw performance per core is lower for both applications (despite a nominally similar processor architecture), but *zeus* delivers higher total throughput.

4.3 *phi* results

Our motivation for considering the *phi* system was to evaluate the performance of our skeletons (and, generally, of the Erlang runtime system) on a many-core accelerator. We have therefore focused on using just the accelerator, without also using the multi-core host system. Figure 6 shows the reproductions per second and throughput increase for the Hybrid version of the two use cases. Note that the system has 61 physical cores, but can run 244 simultaneous

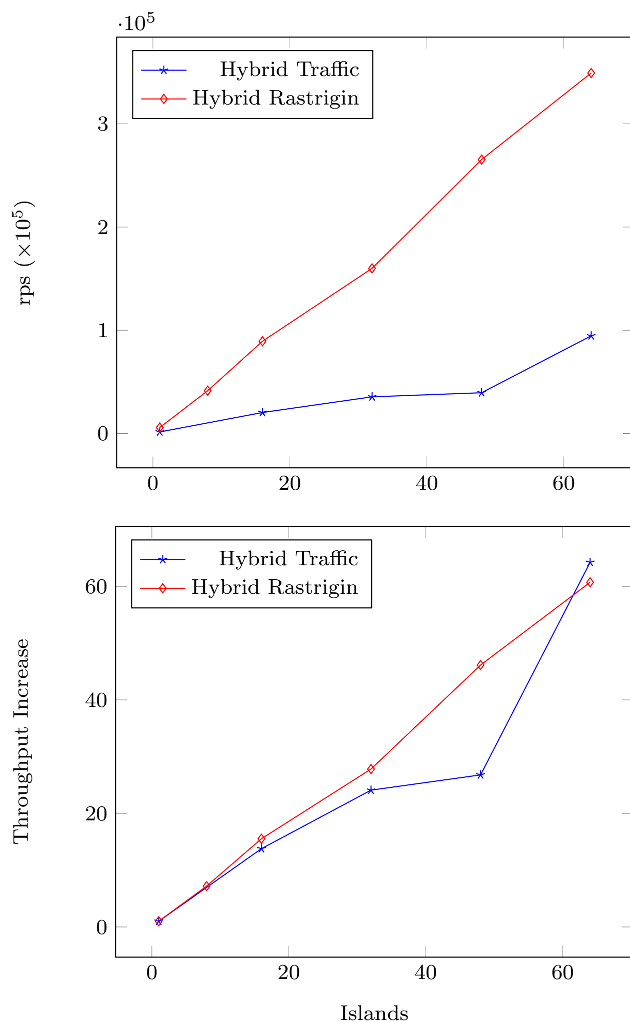

Fig. 5 Reproductions per Second (*rps*, above) and Throughput Increase (below) for Traffic and Rastrigin on *zeus*.

threads via 4 way hyper-threading. Due to the sharing of resources between multiple threads, we do not expect to see 244 times increase in throughput on this architecture. Indeed, we can clearly see that the performance improvement tails off when more than 120 threads are used. Even so, we were able to achieve very good results, improving performance more than 100 times (compared to the sequential version) when 244 threads were used for both applications, and achieving excellent improvements up to 61 cores for both applications. Here, Rastrigin achieves 188,240 *rps* on 244 cores versus 1,681 *rps* on one core, for a throughput increase of 111.96× the sequential version. Traffic achieves 42,902

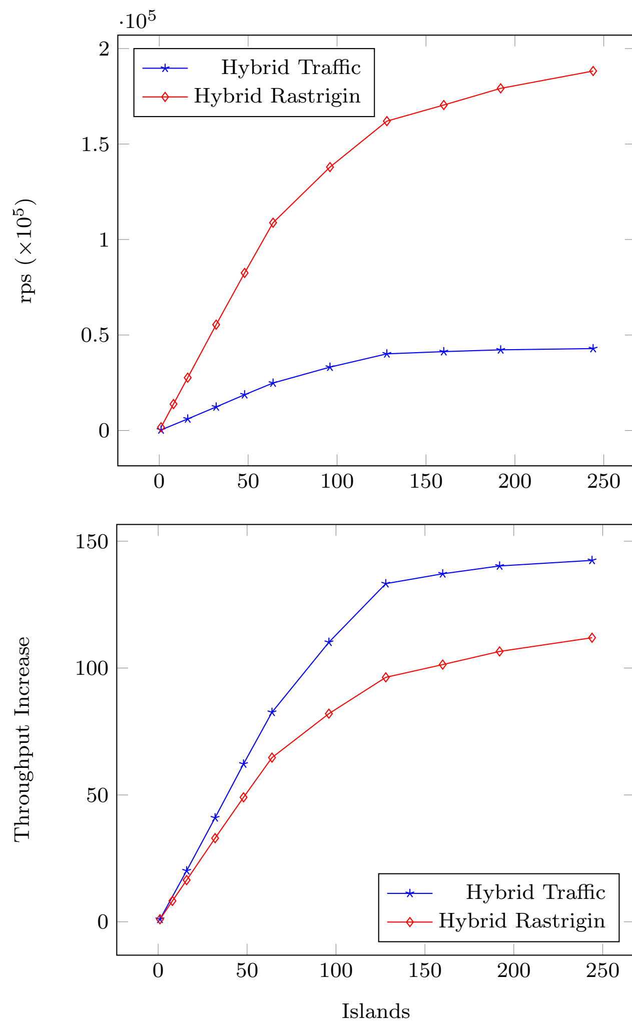

Fig. 6 Reproductions per Second (*rps*, above) and Throughput Increase (below) for Traffic and Rastrigin on *phi*.

rps on 244 cores versus 302 *rps* on one core, for a throughput increase of 142.44× the sequential version. Further experiments are needed to determine the cause of the super-linearity up to 120 cores.

4.4 *power* results

The *power* system represents the IBM Power8 architecture, which is designed to allow a highly multi-threaded chip implementation. Each core is capable of handling 8 hardware threads, which gives a total of 160 threads on a system

with 20 physical cores. At 3.69Ghz, the clock frequency is also noticably higher than for the other systems that we have considered. However, as with the Xeon Phi, resources are shared between multiple threads and we therefore do not see a linear increase in throughput with the number of cores (islands). Due to problems with running *Rastrigin*, we only present the results for the *Traffic* use case (Figure 7). Since the latest version of Erlang (Erlang 18) is currently unsupported on the OpenPower architecture, our experiments use an older version, Erlang 16. Our tests on other systems suggest that this dramatically lowers the absolute throughput of the system (by about a factor of 3). Overall, we achieve 172,078 *rps* on 160 cores versus 1,681 *rps* on one core, for a throughput increase of 39.53× the sequential version. We observe very good scaling up to about 40 threads (2 per physical core), achieving throughput improvement of 25.75 on 32 cores and 34.76 on 64 cores. Beyond that, as with the *phi* system, performamce improvements tail off rapidly, but smoothly.

4.5 *pi* results

The *pi* system is a quad-core Raspberry Pi 2, model B. We have chosen it as an example of a low-power parallel architecture, that is intended to be used in e.g. high-end embedded systems. While we expect the absolute performance of this system to be significantly lower than the other, heavyweight, parallel systems we have studied, it is still interesting to in evaluating how well our skeletons perform on such a system. Figures 8 shows the results for *pi*. We have considered all three versions of the EMAS skeleton. As for *titanic*, the Hybrid version gives the best performance, and the *Skel* version outperforms the *Concurrent* version. For Traffic, the Hybrid version on *titanic* achieves 5,962 *rps* on 4 cores versus 1,568 *rps* on one core, yielding 3.92× the throughput. For Rastrigin, the Hybrid version achieves almost identical results of 5,954 *rps* on 4 cores versus 1,556 *rps* on one core, also yielding 3.92× the throughput. We observe similar results for the other two versions: with throughput improvements of 3.54/3.54 for the Skel version on Traffic/Rastrigin (5,387/5,401 *rps*), and 2.93 for the Concurrent version (4,453/4,447 *rps*). Although, as expected, the absolute peak performance (about 6000 reproductions per second) was orders of magnitute weaker than on other systems; however, it is worth noting that Raspberry Pi power consumption is just a fraction of these of heavyweight servers, so it gives the best pefrormance-per-watt ratio of all systems we tested.

5 Related Work

Algorithmic skeletons [14] have been the focus of much research since the 1980s, with a number of skeleton libraries being produced in a range of languages [20, 2]. These include our own Skel library [7], which is currently the only available skeleton library for Erlang. A recent alternative approach is to define an

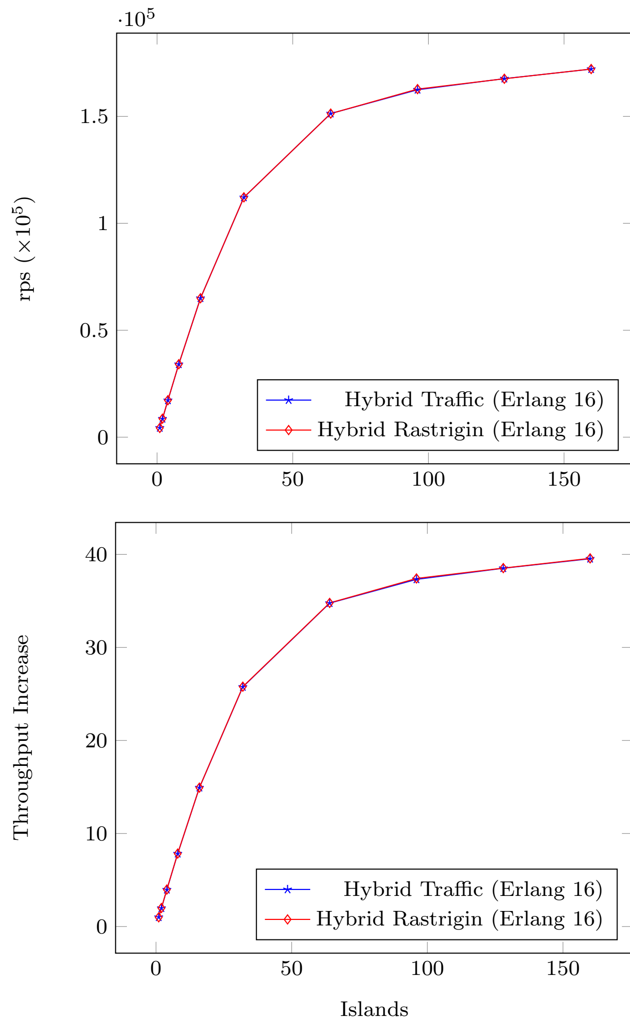

Fig. 7 Reproductions per Second (*rps*, above) and Throughput Increase (below) for Traffic and Rastrigin on *power* (Erlang 16).

embedded domain-specific language (EDSL) [25, 13, 21]. This can provide a structured approach to parallelism, similar to our approach using skeletons. Skeletons, however, have the advantage of representing *language-independent* patterns of parallelism, that can more easily be transferred to other language settings. Other high-level approaches that hide the low-level parallel mechanics from the programmer include: futures [18], evaluation strategies [33], and parallel monads [27].

There have been a few attempts to parallelise agent systems. For example Al-Jaroodi *et al.* have presented an implementation of a Java-based agent-

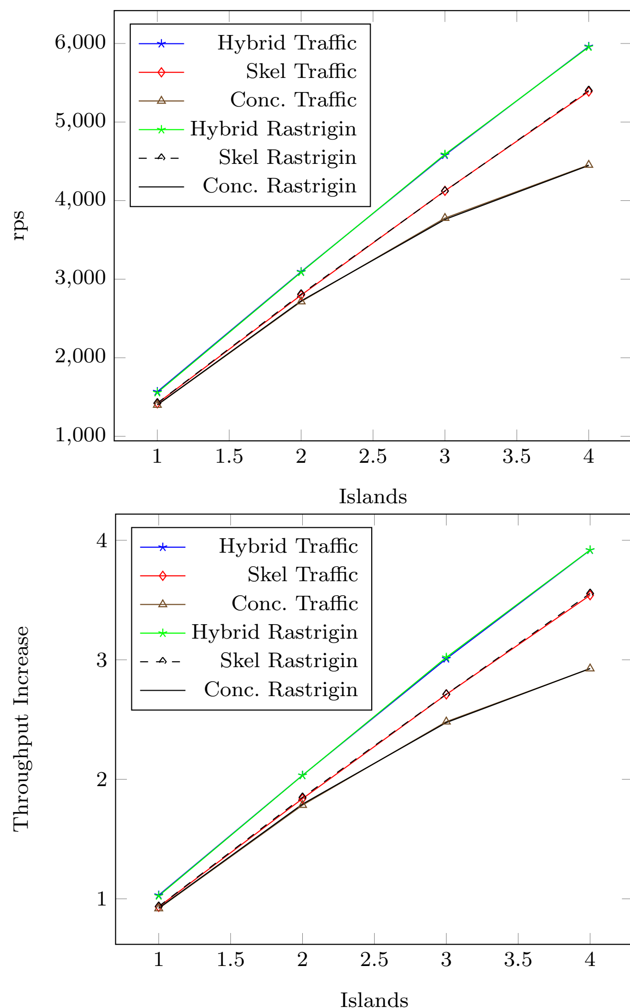

Fig. 8 Reproductions per Second (*rps*, above) and Throughput Increase (below) for Traffic and Rastrigin on *pi*

oriented system, achieving linear scalability up to 8 cores [1], and Cicirelli and Nigri have achieved similar scale-up in a system that is aimed at simulation [12]. Cosenza *et al.* describe an implementation of a C++/MPI scalable simulation platform for spatial simulations of particle movement, achieving linear scalability up to 64 cores [16]. Finally, Cahon *et al.* have achieved linear speedup on up to 10 cores for a C++ implementation of ParadisEO [10]. All of these systems were tested on either small scale parallel systems, or are tailored to a specific evolutionary computing application. In contrast, our skeleton is sufficiently general to cover a wide class of evolutionary computing applications, and we have also demonstrated its scalability on systems with up to

244 hardware threads, achieving maximum throughput improvements of up to 142.44× the sequential version.

6 Conclusions and Future Work

Evolutionary multi-agent systems (EMAS) are a very important component of many artificial intelligence systems. In this paper, we have described the design and implementation of a new domain-specific skeleton for evolutionary multi-agent systems (EMAS) in Erlang. By implementing a skeleton for a widely used computational pattern, we enable easy parallelisation of a large class of applications, where programmer is required to supply only problem-specific sequential components, and all parallelism is handled by the skeleton implementation. We also enable easy cross-platform portability. We have developed three different versions of the skeleton - the *Concurrent* version that is based on Erlang processes, the *Skel* version that is based on the Skel library of parallel skeletons, and the *Hybrid* tuned version of the *Skel* version. We evaluated these different versions of the skeleton on two different evolutionary computing applications: i) finding the minimum of the Rastrigin function; and ii) an urban traffic optimisation. These applications come from different domains and were adapted to use the same skeleton. We have achieved very good improvements in performance on a number of different architectures, ranging from small-scale low-power multi-core systems (a quad-core Raspberry Pi) to larger multi-core servers (a 64-core AMD Opteron) and a many-core accelerator (Intel Xeon Phi) with very little coding effort. This showed the scalability and adaptability of our skeleton implementation in a number of different deployment settings. As expected, the best implementation was consistently the Hybrid version. Our results show that we can achieve throughput improvements for this version of up to a factor of 142.44 compared to the sequential version (for the Xeon Phi), with near-linear improvements on the servers and Raspberry Pi. Overall, the 64-core multi-core system (*zeus*) gave the best result for Rastrigin (349,226 *rps*). However, rather surprisingly, the 24-core server (*titanic*) achieves the best result for Traffic (227,509 *rps vs* 94,950 *rps*). Given their pricing relative to the other architectures, the results for the Xeon Phi and Raspberry Pi are very creditable. Although the OpenPower results are a little disappointing in absolute terms, this is largely because an older version of Erlang was used. If the newest version of Erlang was available, we would expect this architecture to deliver the best absolute results. Finally, we had not expected the architectures to have such widely varying performance. The fact that our skeletons (notably the Hybrid version) dealt well and, on the whole, predictably with all of these architectures is a positive result.

Our main line of future work is to adapt our skeleton to distributed-memory systems, so allowing further scaling to supercomputers. This would, however, require significant reengineering of the skeleton. We also plan to consider more complex real-life applications that conform to the skeleton, to investigate the

anomalies that we have observed on e.g. the *zeus* and *phi* systems, and to retry the results for the *power* system once Erlang 18 is available.

References

1. (2001) Agent-based parallel computing in java proof of concept. Tech. Rep. TR-UNL-CSE-2001-1004, University of Nebraska—Lincoln
2. Aldinucci M, Campa S, Danelutto M, Kilpatrick P, Torquati M (2016) Pool evolution: A parallel pattern for evolutionary and symbolic computing. International Journal of Parallel Programming 44(3):531–551, DOI 10.1007/s10766-015-0358-5, URL `http://dx.doi.org/10.1007/s10766-015-0358-5`
3. Armstrong J, Virding S, Williams M (1993) Concurrent Programming in Erlang. Prentice-Hall
4. Aronis S, Papaspyrou N, Roukounaki K, Sagonas K, Tsiouris Y, Venetis IE (2012) A Scalability Benchmark Suite for Erlang/OTP. In: Proc. Erlang Workshop, ACM, Erlang '12, pp 33–42
5. Back T (1994) Selective Pressure in Evolutionary Algorithms: A Characterization of Selection Mechanisms. In: Proc. WCCI, IEEE, pp 57–62 vol.1
6. Bellifemine F, Poggi A, Rimassa G (2001) Jade: A fipa2000 compliant agent development environment. In: Proceedings of the Fifth International Conference on Autonomous Agents, ACM, New York, NY, USA, AGENTS '01, pp 216–217, DOI 10.1145/375735.376120, URL `http://doi.acm.org/10.1145/375735.376120`
7. Brown C, et al (2013) Cost-Directed Refactoring for Parallel Erlang Programs. Intnl Journal of Parallel Programming pp 1–19
8. Byrski A, Schaefer R, Smołka M (2013) Asymptotic guarantee of success for multi-agent memetic systems. Bulletin of the Polish Academy of Sciences—Technical Sciences 61(1)
9. Byrski A, Drezewski R, Siwik L, Kisiel-Dorohinicki M (2015) Evolutionary multi-agent systems. The Knowledge Engineering Review 30:171–186
10. Cahon S, Talbi E, Melab N (2003) Paradiseo: a framework for parallel and distributed biologically inspired heuristics. In: Parallel and Distributed Processing Symposium, 2003. Proceedings, pp 9 pp.–
11. Cahon S, Melab N, Talbi EG (2004) Paradiseo: A framework for the reusable design of parallel and distributed metaheuristics. Journal of Heuristics 10(3):357–380, DOI 10.1023/B:HEUR.0000026900.92269.ec, URL `http://dx.doi.org/10.1023/B:HEUR.0000026900.92269.ec`
12. Cicirelli F, Nigro L (2013) AsiaSim 2013: 13th International Conference on Systems Simulation, Singapore, November 6-8, 2013. Proceedings, Springer Berlin Heidelberg, Berlin, Heidelberg, chap An Agent Framework for High Performance Simulations over Multi-core Clusters, pp 49–60

13. Claessen K, Sheeran M, Svensson BJ (2012) Expressive Array Constructs in an Embedded GPU Kernel Programming Language. In: Proc. of DAMP, pp 21–30
14. Cole M (2004) Bringing Skeletons out of the Closet: A Pragmatic Manifesto for Skeletal Parallel Programming. Elsevier Science Publishers B. V., Amsterdam, The Netherlands, The Netherlands, vol 30, pp 389–406
15. Cole MI (1988) Algorithmic skeletons: A structured approach to the management of parallel computation. PhD thesis, aAID-85022
16. Cosenza B, Cordasco G, De Chiara R, Scarano V (2011) Distributed load balancing for parallel agent-based simulations. In: Parallel, Distributed and Network-Based Processing (PDP), 2011 19th Euromicro International Conference on, pp 62–69
17. Digalakis J, Margaritis K (2002) An experimental study of benchmarking functions for evolutionary algorithms. International Journal of Computer Mathemathics 79(4):403–416
18. Fluet M, Rainey M, Reppy J, Shaw A, Xiao Y (2007) Manticore: A Heterogeneous Parallel Language. In: Proc. of the 2007 Workshop on Declarative Aspects of Multicore Programming, ACM, New York, NY, USA, DAMP '07, pp 37–44
19. Fogel DB (2000) What is Evolutionary Computation? IEEE Spectrum 37(2):26, 28–32, DOI 10.1109/6.819926
20. González-Vélez H, Leyton M (2010) A Survey of Algorithmic Skeleton Frameworks: High-level Structured Parallel Programming Enablers. Softw Pract Exper 40(12):1135–1160
21. Grossman M, Simion Sbîrlea A, Budimlić Z, Sarkar V (2011) CnC-CUDA: Declarative Programming for GPUs. In: Languages and Compilers for Parallel Computing, Lecture Notes in Computer Science, vol 6548, Springer Berlin Heidelberg, pp 230–245, DOI 10.1007/978-3-642-19595-2˙16
22. Gutknecht O, Ferber J (2000) Madkit: A generic multi-agent platform. In: Proc. AGENTS '00, ACM, pp 78–79, URL http://doi.acm.org/10.1145/336595.337048
23. Janjic V, Brown C, Hammond K (2015) Lapedo: Hybrid Skeletons for Programming Heterogeneous Multicore Machines in Erlang. In: Proc. Intl. Conf on Parallel Computing (ParCo 2015), IOS Press
24. Krzywicki D, Turek W, Byrski A, Kisiel-Dorohinicki M (2015) Massively concurrent agent-based evolutionary computing. Journal of Computational Science 11:153 – 162, DOI http://dx.doi.org/10.1016/j.jocs.2015.07.003, URL http://www.sciencedirect.com/science/article/pii/S1877750315300041
25. Lee HJ, et al (2011) Implementing Domain-Specific Languages for Heterogeneous Parallel Computing. Micro, IEEE 31(5):42–53
26. Luke S, Cioffi-Revilla C, Panait L, Sullivan K, Balan G (2005) Mason: A multiagent simulation environment. Simulation 81(7):517–527, DOI 10.1177/0037549705058073, URL http://dx.doi.org/10.1177/0037549705058073

27. Marlow S, Newton R, Peyton Jones S (2011) A Monad for Deterministic Parallelism. In: Proceedings of the 4th ACM Symposium on Haskell, ACM, New York, NY, USA, Haskell '11, pp 71–82, DOI 10.1145/2034675.2034685, URL http://doi.acm.org/10.1145/2034675.2034685
28. North M, Collier N, Ozik J, Tatara E, Macal C, Bragen M, Sydelko P (2013) Complex adaptive systems modeling with repast simphony. Complex Adaptive Systems Modeling 1(1):3, DOI 10.1186/2194-3206-1-3, URL http://www.casmodeling.com/content/1/1/3
29. Piętak K, Kisiel-Dorohinicki M (2013) Transactions on Computational Collective Intelligence X, Springer Berlin Heidelberg, Berlin, Heidelberg, chap Agent-Based Framework Facilitating Component-Based Implementation of Distributed Computational Intelligence Systems, pp 31–44. DOI 10.1007/978-3-642-38496-7˙3, URL http://dx.doi.org/10.1007/978-3-642-38496-7_3
30. Rashkovskii, Yurii (2013) Genomu: A Concurrency-Oriented Database. In: Erlang Factory SF
31. Reed, Rick (2012) Scaling to Milions of Simultaneous Connections. In: Erlang Factory SF
32. Sagonas K, Winblad K (2014) More scalable ordered set for ets using adaptation. In: Proceedings of the Thirteenth ACM SIGPLAN Workshop on Erlang, ACM, New York, NY, USA, Erlang '14, pp 3–11, DOI 10.1145/2633448.2633455
33. Trinder PW, Hammond K, Loidl HW, Peyton Jones SL (1998) Algorithm + Strategy = Parallelism. Cambridge University Press, New York, NY, USA, vol 8, pp 23–60, DOI 10.1017/S0956796897002967, URL http://dx.doi.org/10.1017/S0956796897002967
34. Wilson, Ken (2012) Migrating a C++ Team to Using Erlang to Deliver a Real-Time Bidding Ad System. In: Erlang Factory SF
35. Zheng S, Long X, Yang J (2013) Using Many-core Coprocessor to Boost Up Erlang VM. In: Proc. Erlang Workshop, ACM, pp 3–14