
General Upper Bounds on the Runtime of Parallel Evolutionary Algorithms*

Jörg Lässig

Department of Computer Science, University of Applied Sciences
Zittau/Görlitz, Germany

jlæssig@hszg.de

Dirk Sudholt

Department of Computer Science, University of Sheffield, UK

d.sudholt@sheffield.ac.uk

doi:10.1162/EVCO_a_00114

Abstract

We present a general method for analyzing the runtime of parallel evolutionary algorithms with spatially structured populations. Based on the fitness-level method, it yields upper bounds on the expected parallel runtime. This allows for a rigorous estimate of the speedup gained by parallelization. Tailored results are given for common migration topologies: ring graphs, torus graphs, hypercubes, and the complete graph. Example applications for pseudo-Boolean optimization show that our method is easy to apply and that it gives powerful results. In our examples the performance guarantees improve with the density of the topology. Surprisingly, even sparse topologies such as ring graphs lead to a significant speedup for many functions while not increasing the total number of function evaluations by more than a constant factor. We also identify which number of processors lead to the best guaranteed speedups, thus giving hints on how to parameterize parallel evolutionary algorithms.

Keywords

Parallel evolutionary algorithms, runtime analysis, island model, spatial structures.

1 Introduction

Due to the increasing number of CPU cores, exploiting possible speedups by parallel computations is currently more important than ever. Parallel evolutionary algorithms (EAs) form a popular class of heuristics with many applications to computationally expensive problems (Luque and Alba, 2011; Nedjah et al., 2006; Tomassini, 2005). This includes island models, also called distributed EAs, multi-deme EAs, or coarse-grained EAs. Evolution is parallelized by evolving subpopulations, called islands, on different processors. Individuals are periodically exchanged in a process called migration, where selected individuals, or copies of these, are sent to other islands, according to a migration topology that determines which islands are neighboring. More fine-grained models are also known, where neighboring subpopulations communicate in every generation, first and foremost in cellular EAs (Tomassini, 2005).

By restricting the flow of information through spatial structures and/or infrequent communication, diversity in the whole system is increased. Researchers and

*A preliminary version of this paper with parts of the results was published at PPSN 2010 (Lässig and Sudholt, 2010b).

practitioners frequently report that parallel EAs speed up the computation time, and at the same time lead to a better solution quality (Luque and Alba, 2011).

Despite these successes, a long history (Cantú Paz, 1997) and very active research in this area (Alba, 2005; Luque and Alba, 2011; Rudolph, 2006), the theoretical foundation of parallel EAs is still in its infancy. The impact of even the most basic parameters on performance is not well understood (Skolicki and De Jong, 2005). Past and present research is mostly empirical, and a solid theoretical foundation is missing. Theoretical studies are mostly limited to artificial settings. In the study of takeover times, an important question is how long it takes for a single optimum to spread throughout the whole parallel EA, if the EA uses selection and migration but neither mutation nor crossover (Rudolph, 2000, 2006). This gives a useful indicator for the speed at which communication is spread, but it does not give any formal results about the runtime of evolutionary algorithms with mutation and/or crossover.

One way of gaining insight into the capabilities and limitations of parallel EAs is by means of rigorous runtime analysis (He and Yao, 2003; Wegener, 2002). By asymptotic bounds on the runtime, it is possible to compare different implementations of parallel EAs and assess the speedup gained by parallelization in a rigorous manner. Many runtime analyses have been presented (Auger and Doerr, 2011; Jansen, 2013; Neumann and Witt, 2010; Oliveto, He, and Yao, 2007), from simple pseudo-Boolean test functions (Droste et al., 2002) to NP-hard problems from combinatorial optimization (Friedrich et al., 2010; Horoba, 2010; Witt, 2005; Yu et al., 2012).

Lässig and Sudholt (2010a) presented the first runtime analysis of a parallel evolutionary algorithm with a nontrivial migration topology. It was demonstrated for a constructed problem that migration is essential in the following way. A suitably parameterized island model with migration has a polynomial runtime; whereas the same model without migration as well as comparable panmictic populations¹ need exponential time, with overwhelming probability. Neumann, Oliveto, Rudolph, and Sudholt (2011) presented a similar result for island models using crossover. If islands perform crossover with immigrants during migration, this can drastically speed up optimization. This was demonstrated for a pseudo-Boolean example as well as for instances of the VERTEXCOVER problem (Neumann et al., 2011).

In this work, we take a broader view and consider the speedup gained by parallelization in terms of the number of generations, for various common pseudo-Boolean functions and function classes of varying difficulty. A general method is presented for proving upper bounds on the parallel runtime of parallel EAs. The latter is defined as the number of generations of the parallel EA until a global optimum is found for the first time. This allows us to estimate the speedup gained by parallelization, defined as the ratio of the expected parallel runtime of a single island and the expected runtime of an island model with multiple islands (see Section 2 for formal definitions). It also can be used to determine how to choose the number of islands such that the best possible upper bounds on the parallel runtime are obtained, while still maintaining an asymptotically optimal speedup.

Our method is based on the fitness-level method or method of f -based partitions, a simple and well-known tool for the analysis of evolutionary algorithms (Droste et al., 2002; Wegener, 2002). The main idea of this method is to divide the search space into sets A_1, \dots, A_m , strictly ordered according to fitness values of elements therein. Elitist EAs, that is, EAs where the best fitness value in the population can never decrease, can

¹A panmictic population is an unstructured population with no mating restrictions.

only increase their current best fitness. If, for each set A_i we know a lower bound s_i on the probability that an elitist EA finds an improvement, that is, for finding a new search point in a new best fitness-level set $A_{i+1} \cup \dots \cup A_m$, this gives rise to an upper bound $\sum_{i=1}^{m-1} 1/s_i$ on the expected runtime. The method is described in more detail in Section 2.

In Section 3 we first derive a general upper bound for parallel EAs, based on fitness levels. Our general method is then tailored toward different spatial structures often used in fine-grained or cellular evolutionary algorithms and parallel architectures in general: ring graphs (Theorem 8 in Section 4), torus graphs (Theorem 10 in Section 5), hypercubes (Theorem 12 in Section 6), and complete graphs (Theorems 14 and 17 in Section 7).

The only assumption made is that islands run elitist algorithms, and that in each generation each island has a chance of transmitting individuals from its best current fitness level to each neighboring island, independent with probability at least p . We call the latter the transmission probability. It can be used to model various stochastic effects such as disruptive variation operators, the impact of selection operators, probabilistic migration, probabilistic emigration and immigration policies, and transient faults in the network. This renders our method widely applicable to a broad range of settings.

1.1 Main Results

Our estimates of parallel runtimes from Theorems 8, 10, 12, 14, and 17 are summarized in the following theorem, hence characterizing our main results. Throughout this work μ always denotes the number of islands.

THEOREM 1: *Consider an island model with μ islands where each island runs an elitist EA. For each island let there be a fitness-based partition A_1, \dots, A_m such that for all $1 \leq i < m$ all points in A_i have a strictly worse fitness than all points in A_{i+1} , and A_m contains all global optima. We say that an island is in A_i if the best search point on the island is in A_i . Let s_i be a lower bound for the probability that in one generation a fixed island in A_i finds a search point in $A_{i+1} \cup \dots \cup A_m$.*

Further, assume that for each edge in the migration topology in every iteration there is a probability of at least p that the following holds, independently from other edges and for all $1 \leq i < m$. If the source island is in A_i then after the generation the target island is in $A_i \cup \dots \cup A_m$. Then the expected parallel runtime of the island model is bounded from above by

1. $O(\frac{1}{p^{1/2}} \sum_{i=1}^{m-1} \frac{1}{s_i^{1/2}}) + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$ for every ring graph or any other strongly connected² topology (Theorem 8),
2. $O(\frac{1}{p^{2/3}} \sum_{i=1}^{m-1} \frac{1}{s_i^{1/3}}) + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$ for every undirected grid or torus graph whose side lengths are at least $\sqrt{\mu}$ in both directions (Theorem 10),
3. $O(\frac{m + \sum_{i=1}^{m-1} \log(1/s_i)}{p}) + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$ for the $(\log \mu)$ -dimensional hypercube graph (Theorem 12),
4. $O(m/p) + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$ for the complete topology K_μ , as well as $O(m + \frac{m \log \mu}{\min\{p, 1\}}) + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$ (Theorems 14 and 17).

²A directed graph is strongly connected if for each pair of vertices u, v there is a directed path from u to v and vice versa.

A remarkable feature of our method is that it can automatically transfer upper bounds for panmictic EAs to parallel versions thereof. The only requirement is that bounds on panmictic EAs must be derived using the fitness-level method, and that the partition A_1, \dots, A_m and the probabilities for improvements s_1, \dots, s_{m-1} used therein are known. Then the expected parallel time of the corresponding island model can be estimated for all mentioned topologies simply by plugging the s_i into Theorem 1. Fortunately, many published runtime analyses use the fitness-level method—either explicitly or implicitly—and the mentioned details are often stated or easy to derive. Hence even researchers with limited expertise in runtime analysis can easily reuse previous analyses to study parallel EAs.

Further, note that we can easily determine which choice of μ , the number of islands, will give an upper bound of order $1/\mu \cdot \sum_{i=1}^{m-1} 1/s_i$ —the best upper bound we can hope for, using the fitness-level method. In all bounds from Theorem 1 we have a first term that varies with the topology and p , and a second term that is always $1/\mu \cdot \sum_{i=1}^{m-1} 1/s_i$. The first term reflects how quickly information about good fitness levels is spread throughout the island model. Choosing μ such that the second term becomes asymptotically as large as the first one, or larger, we get an upper bound of $O(1/\mu \cdot \sum_{i=1}^{m-1} 1/s_i)$. For settings where $\sum_{i=1}^{m-1} 1/s_i$ is an asymptotically tight upper bound for a single island, this corresponds to an asymptotic linear speedup. The maximum feasible value for μ depends on the problem, the topology, and the transmission probability p .

We give simple examples that demonstrate how our method can be applied. Our examples are from pseudo-Boolean optimization, but the method works in any setting where the fitness-level method is applicable. The simple (1+1) EA is used on each island (see Section 2 for details). Table 1 summarizes the resulting runtime bounds for the considered algorithms and problem classes. For simplicity we assume $p = 1$; a more detailed table for general transmission probabilities is presented in the appendix, see Table 3. The number of islands μ was chosen as explained above: to give the smallest possible parallel runtime, while not increasing the sequential time, asymptotically. The table also shows the expected communication effort, defined as the total number of individuals migrated throughout the run. Details are given in Theorems 8, 10, 12, 14, and 17. Bounds on the expected communication effort follow easily from bounds on the parallel runtime using Theorem 2. The functions used in this table are explained in Section 2. Table 3 in the appendix shows all our results for a variable number of islands μ and variable transmission probabilities p .

The method has already found a number of applications and it spawned a number of follow-up papers. After the preliminary version of this work (Lässig and Sudholt, 2010b) was presented, the authors applied it for various problems from combinatorial optimization: the sorting problem (as maximizing sortedness), finding shortest paths in graphs, and Eulerian cycles (Lässig and Sudholt, 2011b). Very recently, Mambrini, Sudholt, and Yao (2012) also used it for studying how quickly island models find good approximations for the NP-hard SETCOVER problem. This work has also led to the discovery of simple adaptive schemes for changing the number of islands dynamically throughout the run, see Lässig and Sudholt (2011a). These schemes lead to near-optimal parallel runtimes, while asymptotically not increasing the sequential runtime on many examples (Lässig and Sudholt, 2011a). These schemes are tailored towards island models with complete topologies, which includes offspring populations as special case. The study of offspring populations in comma strategies is another recent development that was inspired by this work (Rowe and Sudholt, 2012).

Table 1: Asymptotic bounds on expected parallel (T^{par} , number of generations) and sequential (T^{seq} , number of function evaluations) runtimes and expected communication efforts (T^{com} , total number of migrated individuals) for various n -bit functions and island models with μ islands running the (1+1) EA and using migration probability $p = 1$. The number of islands μ was always chosen to give the best possible upper bound on the parallel runtime, while not increasing the upper bound on the sequential runtime by more than a constant factor. For unimodal functions $d + 1$ denotes the number of function values. See Droste et al. (2002) for bounds for the (1+1) EA. Results for Jump_k were restricted to $3 \leq k = O(n/\log n)$ for simplicity. All upper bounds for OneMax and LO stated here are asymptotically tight, as follows from general results in Sudholt (2013).

	(1+1) EA	Ring	Grid/torus	Hypercube	Complete
OneMax					
best μ		$\mu = \Theta(\log n)$	$\mu = \Theta(\log n)$	$\mu = \Theta(\log n)$	$\mu = \Theta(\log n)$
$E(T^{\text{par}})$	$\Theta(n \log n)$	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$
$E(T^{\text{seq}})$	$\Theta(n \log n)$	$\Theta(n \log n)$	$\Theta(n \log n)$	$\Theta(n \log n)$	$\Theta(n \log n)$
$E(T^{\text{com}})$	0	$\Theta(n \log n)$	$\Theta(n \log n)$	$\Theta(n(\log n) \log \log n)$	$\Theta(n \log^2 n)$
LO					
best μ		$\mu = \Theta(n^{1/2})$	$\mu = \Theta(n^{2/3})$	$\mu = \Theta(\frac{n}{\log n})$	$\mu = \Theta(n)$
$E(T^{\text{par}})$	$\Theta(n^2)$	$\Theta(n^{3/2})$	$\Theta(n^{4/3})$	$\Theta(n \log n)$	$\Theta(n)$
$E(T^{\text{seq}})$	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n^2)$
$E(T^{\text{com}})$	0	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n^2 \log^2 n)$	$\Theta(n^3)$
Unimodal					
best μ		$\mu = \Theta(n^{1/2})$	$\mu = \Theta(n^{2/3})$	$\mu = \Theta(\frac{n}{\log n})$	$\mu = \Theta(n)$
$E(T^{\text{par}})$	$O(dn)$	$O(dn^{1/2})$	$O(dn^{1/3})$	$O(d \log n)$	$O(d)$
$E(T^{\text{seq}})$	$O(dn)$	$O(dn)$	$O(dn)$	$O(dn)$	$O(dn)$
$E(T^{\text{com}})$	0	$O(dn)$	$O(dn)$	$O(dn \log n)$	$O(dn^2)$
Jump_k					
best μ		$\mu = \Theta(n^{k/2})$	$\mu = \Theta(n^{2k/3})$	$\mu = \Theta(n^{k-1})$	$\mu = \Theta(n^{k-1})$
$E(T^{\text{par}})$	$\Theta(n^k)$	$O(n^{k/2})$	$O(n^{k/3})$	$O(n)$	$O(n)$
$E(T^{\text{seq}})$	$\Theta(n^k)$	$O(n^k)$	$O(n^k)$	$O(n^k)$	$O(n^k)$
$E(T^{\text{com}})$	0	$O(n^k)$	$O(n^k)$	$O(kn^k \log n)$	$O(n^{2k-1})$

2 Preliminaries

We are interested in the following performance measures, where the number of islands, μ , in the considered island model is obvious from the context. First we define the parallel runtime T^{par} as the number of generations until the first global optimum is evaluated. Let T_μ be the number of generations before the island model finds a global optimum for the first time. Then

$$T^{\text{par}} := T_\mu.$$

The sequential runtime T^{seq} is defined as the number of function evaluations until the first global optimum is evaluated. It thus captures the overall effort across all processors. It is formally defined as

$$T^{\text{seq}} := \mu \cdot T_\mu.$$

In both measures, T^{par} and T^{seq} , the cost of the initialization is neglected, as this only adds a fixed term to the runtimes.

The speedup is defined as the ratio of expected runtimes of a single island and an island model with μ islands:

$$E(T_1) / E(T_\mu).$$

The definition of speedup used here is called weak orthodox speedup in Alba's (2002) taxonomy. If the speedup is at least of order μ , that is, if it is $\Omega(\mu)$, we speak of a linear speedup. In this work it is generally understood in an asymptotic sense, unless we call it a perfect linear speedup.

Note that our notion of speedup is defined with regard to the number of generations. In terms of wall-clock time, when the algorithm is stopped when reaching a global optimum, the execution time of an island model with μ islands can be estimated as

$$\text{Exec}_\mu = T_\mu \cdot (\text{Exec}_\mu^{\text{gen}} + \text{Exec}_\mu^{\text{migr}}) \tag{1}$$

where Exec^{gen} is the execution time for one generation (excluding migration) and $\text{Exec}^{\text{migr}}$ is the execution time for migration. The latter depends on the number of islands, the communication topology, and the number of individuals migrated. For sequential algorithms ($\mu = 1$) we have $\text{Exec}_1^{\text{migr}} = 0$. Contrarily, for homogeneous parallel systems we may assume that $\text{Exec}_\mu^{\text{gen}}$ is fixed, that is, independent of μ .

Many common definitions of speedup consider the wall-clock time, where the differences in $\text{Exec}^{\text{migr}}$ play a role. Here speedups are considered with regard to the number of generations only, ignoring the differences in $\text{Exec}^{\text{migr}}$. This makes sense in setting where $\text{Exec}^{\text{gen}} \gg \text{Exec}^{\text{migr}}$; for instance, when fitness evaluations are so expensive that they dominate the execution time. Otherwise, the speedups stated here may be optimistic as the overhead induced by migration is ignored. Note, however, that with additional information about Exec^{gen} and $\text{Exec}^{\text{migr}}$, and Equation (1), the results of the current paper easily extend to more sophisticated notions of speedups.

To get a more complete picture of the resources used in a parallel system and to take into account the overhead by communication, this paper also considers the communication effort T^{com} . It is defined as the total number of individuals migrated to other islands during the course of a run. The communication effort therefore captures the total bandwidth used during a run of an island model. It represents an important factor for determining the performance of a parallel EA, alongside the parallel runtime.

The expected communication effort is a multiple of the parallel expected runtime, with the factor depending on the number of (directed) edges in the topology, the transmission probability p and the number of individuals migrated in each migration event.

The following theorem lists various topologies: a unidirectional ring is a graph consisting of a single directed cycle, whereas a directed ring has undirected edges. Note that an undirected edge can be regarded as two directed edges. In a torus graph, all vertices are arranged on a two-dimensional grid, with undirected edges wrapping around (vertices in the top row are neighbors to the ones in the bottom row and vice versa, similarly for the leftmost and rightmost columns). Each vertex in a torus thus has four distinct neighbors, provided that the torus has at least three rows and at least three columns. Hypercubes are formally defined in Section 6, and the complete graph K_μ contains undirected edges between all pairs of nodes.

THEOREM 2: *Consider an island model with a directed graph $T = (V, E)$ as topology, such that in each generation along each directed edge migration takes place independently with probability*

p. Assume that ν individuals are migrated in each migration event. Let T^{par} be the parallel runtime of the island model, then the expected communication effort $E(T^{com})$ is

$$E(T^{com}) = \nu \cdot |E| \cdot E(T^{par}). \tag{2}$$

Thereby, if $|V| = \mu$, we have

- $|E| = \mu$ for a unidirectional ring and $|E| = 2\mu$ for a bidirectional ring,
- $|E| = 4\mu$ for any torus graph where both sides have length at least three,
- $|E| = \mu \log \mu$ for the $(\log \mu)$ -dimensional hypercube, and
- $|E| = \mu(\mu - 1)$ for the complete graph K_μ .

PROOF: Fix a single edge e , then the expected number of migration events across e equals

$$\sum_{t=0}^{\infty} \text{Prob}(T^{par} = t) \cdot t \cdot p = p \cdot E(T^{par})$$

by definition of $E(T^{par})$. By linearity of expectations, we can add these values for all edges to get the expected number of migration events across the whole topology. This yields a factor of $|E|$. Additionally multiplying by ν gives the expected communication effort. \square

Hence to estimate the expected communication effort it suffices to analyze the expected parallel runtime.

In the example applications, this work considers the maximization of a pseudo-Boolean function $f : \{0, 1\}^n \rightarrow \mathbb{R}$. It is easy to adapt the method for minimization. The number of bits is always denoted by n . The following well known example functions have been chosen because they exhibit different probabilities for finding improvements in a typical run of an EA. For a search point $x \in \{0, 1\}^n$ write $x = x_1 \dots x_n$; then $\text{OneMax}(x) := \sum_{i=1}^n x_i$ counts the number of ones in x and $\text{LO}(x) := \sum_{i=1}^n \prod_{j=1}^i x_j$ counts the number of leading ones in x , that is, the length of the longest prefix containing only 1-bits. A function is called unimodal if every nonoptimal search point has a Hamming neighbor (that is, a point with Hamming distance 1 to it) with strictly larger fitness. Observe that LO is unimodal as flipping the first 0-bit results in a fitness increase. For LO, every nonoptimal point has exactly one Hamming neighbor with a better fitness. For $1 \leq k \leq n$ we also consider

$$\text{Jump}_k := \begin{cases} k + \sum_{i=1}^n x_i, & \text{if } \sum_{i=1}^n x_i \leq n - k \text{ or } x = 1^n, \\ \sum_{i=1}^n (1 - x_i) & \text{otherwise.} \end{cases}$$

This function was introduced by Droste, Jansen, and Wegener (2002) as a function with tunable difficulty as evolutionary algorithms typically have to perform a jump to overcome a gap by flipping k specific bits. It is also interesting because it is one of very few examples where crossover has been proven to be essential (Jansen and Wegener, 2002; Kötzing et al., 2011).

Our method for proving upper bounds is based on the fitness-level method (Droste et al., 2002; Wegener, 2002). The idea is to partition the search space into sets A_1, \dots, A_m called fitness levels that are ordered with respect to fitness values. We say that an

algorithm is in A_i or on level i if the current best individual in the population is in A_i . An evolutionary algorithm where the best fitness value in the population can never decrease (called an elitist EA) can only improve the current fitness level. If one can derive lower bounds on the probability of leaving a specific fitness level toward higher levels, this yields an upper bound on the expected runtime.

THEOREM 3 (FITNESS-LEVEL METHOD): *For two sets $A, B \subseteq \{0, 1\}^n$ and a fitness function f let $A <_f B$ if $f(a) < f(b)$ for all $a \in A$ and all $b \in B$. Partition the search space into nonempty sets A_1, A_2, \dots, A_m such that $A_1 <_f A_2 <_f \dots <_f A_m$ and A_m only contains global optima. For an elitist EA let s_i be a lower bound on the probability of creating a new offspring in $A_{i+1} \cup \dots \cup A_m$, provided the population contains a search point in A_i . Then the expected number of iterations of the algorithm to find the optimum is bounded from above by*

$$\sum_{i=1}^{m-1} \frac{1}{s_i}.$$

In contrast to other methods such as drift analysis (He and Yao, 2004; Johannsen, 2010), the fitness-level method is applicable in cases where “easy” and “hard” fitness levels are mixed up, so that the progress toward the optimum cannot reasonably be bounded by a closed formula.

The fitness-level method has also been applied to other elitist optimization methods, including elitist ant colony optimizers (Gutjahr and Sebastiani, 2008; Neumann et al., 2009) and a binary particle swarm optimizer (Sudholt and Witt, 2010). It gives rise to powerful tail inequalities (Zhou et al., 2012) and it can be used to prove lower bounds as well, when combined with additional knowledge on transition probabilities (Sudholt, 2013). Finally, Lehre (2011) recently showed that the fitness-level method can be extended toward nonelitist EAs with additional mild conditions on transition probabilities and the population size.

Note that the method only requires a finite number of fitness-level sets, not a finite set of fitness values. In principle, the method can be applied to continuous fitness functions as well, provided a suitable discretization is made and the goal is to find the best fitness-level set. However, this might not be the most practical approach.

In the following, the fitness-level method is applied to parallel EAs. For the considered EAs the assumption of this work is that there is a migration topology, given by a directed graph. Islands represent vertices of the topology and directed edges indicate neighborhoods between the islands. We often describe undirected graphs for use as migration topology, understanding that for an undirected edge $\{u, v\}$ we have two directed edges (u, v) and (v, u) . In other words, though formally the migration topology is a directed graph, we often use the language of undirected graphs to describe it.

Our methods for proving upper bounds require that the islands run elitist evolutionary algorithms. All islands create new offspring independently by mutation and/or recombination among individuals in the island. In every generation there is a chance that migration will send an individual on the current best fitness level to some target island, and that this individual will be included on the target island. This would effectively increase the fitness level of the target island to the current best level (or an even better one). For every pair of connected islands, we call this probability *transmission probability* and denote it p . Note that for any pair of islands, the mentioned transmission events are independent.

Algorithm 1: Parallel (1+1) EA with μ islands and migration probability p

For all $1 \leq i \leq \mu$ choose $x^i \in \{0, 1\}^n$ uniformly at random.

repeat

For all $1 \leq i \leq \mu$ **do in parallel**

Create y^i by flipping each bit in x^i with probability $1/n$.

if $f(y^i) \geq f(x^i)$ **then** $x^i := y^i$.

Send a copy of x^i to each neighboring island, independently with probability p .

Choose z^i with maximum fitness among all incoming migrants.

if $f(z^i) \geq f(x^i)$ **then** $x^i := z^i$.

The transmission probability can model various settings, where randomness and stochasticity may be involved:

- Migrations do not take place in every generation, but only probabilistically with probability p ,
- Islands do not automatically select individuals on the best fitness level for emigration, but there is a probability of at least p that this happens,
- Similarly, islands do not automatically include immigrants on higher fitness levels, but only with probability of at least p ,
- During migration crossover is performed, and p is a lower bound on the probability that crossover does not disrupt the fitness of an individual on a current best fitness level (if a crossover probability p_c is used, then clearly $p \geq 1 - p_c$), and
- The physical architecture suffers from transient faults and p is a lower bound on the probability that migration is executed correctly.

Of course, the transmission probability can also model any combination of the above, in which case the product of all above probabilities gives a lower bound on the transmission probability.

Most of our results also apply when instead of probabilistic migration a fixed migration interval τ is used. This is similar to a migration probability $p = 1/\tau$; in fact, it can be regarded as a derandomized or quasirandom version of probabilistic migration. With a fixed migration interval the variance in the information propagation is reduced, and all islands operate in synchronicity. Probabilistic migrations are asynchronous; this simplifies the analysis as we do not need to keep track on how much time has passed since the last migration. We expect our results for probabilistic migration to transfer to the study of migration intervals. The only notable exception is the case of a complete topology, when the migration probability is rather small (Theorem 17), as there, synchronous and asynchronous migrations lead to different effects.

As elaborated above, our method is robust and it applies in various settings, and for various types of EAs simulated on the islands. In our applications for illustrating concrete speedups for test problems, we use a simple (1+1) EA for all islands. The (1+1) EA maintains a single current search point, and in each generation it creates an offspring by mutation. The offspring replaces its parent if its fitness is not worse. The resulting island model is shown in Algorithm 1.

3 Proving Upper Bounds for Parallel EAs

3.1 A General Upper Bound

This section covers how to prove upper bounds on the runtime of parallel EAs. In contrast to panmictic EAs, in an island model several islands might participate in the search for improvements from the current-best fitness level. The number of islands may vary over time according to the spread of information.

The following theorem transfers upper bounds for panmictic EAs derived by the fitness-level method into upper bounds for parallel EAs in a systematic way.

THEOREM 4 (FITNESS-LEVEL METHOD FOR PARALLEL EAs): *Consider a partition of the search space into fitness levels $A_1 <_f A_2 <_f \dots <_f A_m$ such that A_m only contains global optima. Let s_i be (a lower bound on) the probability that a fixed island running an elitist EA creates a new offspring in $A_{i+1} \cup \dots \cup A_m$, provided the island contains a search point in A_i . Let μ_t for $t \in \mathbb{N}$ denote (a lower bound on) the number of islands that have discovered an individual in $A_i \cup \dots \cup A_m$ in the t generation after the first island has found such an individual. Then the expected parallel runtime of the parallel EA on f is bounded from above by*

$$E(T^{par}) \leq \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} (1 - s_i)^{\sum_{j=1}^t \mu_j}.$$

PROOF: Let T_i denote the random time until the first island finds an individual on a fitness level $i + 1, \dots, m$, starting with at least one individual on fitness level i in the whole population. The expected parallel runtime can be written as

$$\begin{aligned} E(T^{par}) &\leq \sum_{i=1}^{m-1} E(T_i) = \sum_{i=1}^{m-1} \sum_{t=1}^{\infty} \text{Prob}(T_i \geq t) \\ &= \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} \text{Prob}(T_i \geq t + 1). \end{aligned}$$

A necessary condition for $T_i \geq t + 1$ is that during all t generations after the first individual has reached fitness level i , all islands are unsuccessful in finding an improvement. In the j th of these generations, there are at least μ_j islands with individuals in $A_i \cup \dots \cup A_m$. Each island is successful with probability at least s_i . Using that the islands create new offspring independently, the probability of all islands being unsuccessful is at most $(1 - s_i)^{\mu_j}$. Thus,

$$\begin{aligned} \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} \text{Prob}(T_i \geq t + 1) &\leq \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} \prod_{j=1}^t (1 - s_i)^{\mu_j} \\ &= \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} (1 - s_i)^{\sum_{j=1}^t \mu_j}. \quad \square \end{aligned}$$

The upper bound from Theorem 4 is very general as it does not restrict the communication among the islands in any way. These aspects are hidden in the definition of the variables μ_t . When looking at one particular fitness level, say level i , we also speak of islands being informed if and only if they contain an individual on level i . The variable μ_t then gives the number of informed islands t generations after the first island has become informed by reaching level i .

The spread of information obviously depends on the migration topology, the migration interval, and the selection strategies used to choose migrants that are sent and how migrants are included in the population. The basic method works for all choices of these design aspects. We elaborate on these aspects and then move on to more specific scenarios where we can obtain more concrete results.

3.2 How to Deal with Migration Intervals

With a migration interval of $\tau > 1$, the μ_t value remains fixed for periods of τ generations, unless further islands are raised toward the current best fitness level by variation. If we pessimistically ignore this effect, then we have for appropriate t that $\mu_t = \mu_{t+1} = \dots = \mu_{t+\tau-1}$. In any case, we have $\mu_t \leq \mu_{t+1} \leq \dots \leq \mu_{t+\tau-1}$, hence the sum of μ values is at least $\sum_{j=1}^t \mu_j \geq \tau \sum_{j=1}^{t/\tau} \mu_{(j-1)\tau+1}$. This implies the following simplified upper bound.

COROLLARY 5: *For a parallel EA with migration interval τ the bound from Theorem 4 simplifies to*

$$E(T^{par}) \leq \sum_{i=1}^{m-1} \sum_{t=0}^{\infty} (1 - s_i)^\tau \sum_{j=1}^{t/\tau} \mu_{(j-1)\tau+1} .$$

The values $\mu_{(j-1)\tau}$ can be estimated like the values μ_j in a setting with $\tau = 1$. In order to keep the presentation simple, in the following applications we only consider the case that $\tau = 1$, that is, migration happens in every generation. This reflects common principles used in fine-grained or cellular evolutionary algorithms. The following considerations can always be combined with the above arguments to handle migration intervals larger than 1.

3.3 Stochastic Communication and Finding Improvements

In order to arrive at more concrete bounds on the parallel runtime for common migration topologies, it is important to understand how the number of informed islands grows on each fitness level, that is, the growth curves underlying the μ_j variables. Note that these variables are random variables in all settings where we have a transmission probability less than 1. This means that getting a closed formula for the expected parallel runtime is not easy. In Theorem 4 we cannot simply replace the μ_j variables by their expectations as by Jensen’s (1906) inequality this would yield an estimation in the wrong direction (i.e., it would give a lower bound where an upper bound is needed). More work is required in order to arrive at closed formulas for common topologies.

Instead of arguing with the random number of informed islands, it is easier to argue with expected hitting times for the time until a specified number of islands is informed. If we know such expected hitting times, or upper bounds thereof, the time until the parallel EA finds a better fitness level can be estimated.

LEMMA 6: *Consider an island model running elitists EAs and fix some fitness level i with success probability s_i for each island. Let $\xi(k)$ denote the random number of generations until at least k islands are informed. Then for every $k \leq \mu$, the expected time until this fitness level is left toward a better one is at most*

$$E(\xi(k)) + 1 + \frac{1}{k} \cdot \frac{1}{s_i} .$$

PROOF: The following argument is similar to previous work by Witt (2006) on parent populations. After $E(\xi(k))$ expected generations, there are at least k informed islands.

Then the probability of leaving the fitness level is at least $1 - (1 - s_i)^k$, and the expected time is bounded from above by

$$\frac{1}{1 - (1 - s_i)^k} \leq 1 + \frac{1}{k} \cdot \frac{1}{s_i}, \tag{3}$$

where the inequality was published in Rowe and Sudholt (2012, Lemma 3) and additionally stated as Lemma 19 in the appendix. Together, this proves the claim. \square

A good choice for k is one where $E(\xi(k)) \approx \frac{1}{k} \cdot \frac{1}{s_i}$ as this is likely to minimize the bound from Lemma 6, at least asymptotically.

The lemma ignores the fact that during the first $\xi(k)$ generations islands can already find improvements. It also ignores that the number of informed islands might grow beyond k after this time. However, we will see that for appropriate choices of k , the lemma can still give near-optimal results. In the first generations, the number of informed islands is likely to be too small to yield a significant benefit. In addition, after k islands have been informed, this number is large enough to guarantee that improvements are found quickly, for appropriate k .

3.4 Information Propagation in Networks

The task remains to estimate the first hitting time for informing a certain number of vertices. Note that this is similar to studying growth curves and takeover times. In fact, $\xi(\mu)$ is the expected time until the whole island model is informed. Growth curves and takeover times have been studied in artificial settings where no variation takes place (see Alba and Luque, 2004; Giacobini, Alba, et al., 2005; Giacobini, Alba, et al., 2003; Giacobini, Tomassini, et al., 2003; Giacobini, Tomassini, et al., 2005; Rudolph, 2000, 2006; Sarma and De Jong, 1997) or recent surveys (Luque and Alba, 2011, Chapter 4; Sudholt, 2013).

In the following, we refer to our model of transmission probabilities as it is a general model that captures many stochastic components in the dynamic behavior of island models. But at the same time, it is simple enough to allow for a theoretical analysis.

Transmission probabilities give rise to a stochastic information propagation process in networks. Each informed vertex in the network independently tries to inform all its neighbors in every iteration, and information is successfully transmitted across any of these edges with probability p . This process was studied by Rowe et al. (2008), who considered the propagation time as the time until all vertices in the network are informed. They presented bounds for interesting graph classes as well as a general upper bound of

$$\frac{8 \text{diam}(G) + 8 \log n}{p(1 - e^{-1})}$$

for the propagation time on an undirected graph G . Thereby, $\text{diam}(G)$ denotes the diameter of G , defined as the maximum number of edges on any shortest path between two vertices in the graph.

Interestingly, the same probabilistic process also underlies the way randomized search heuristics find shortest paths in weighted undirected graphs. Doerr et al. (2007, 2011) showed that the (1+1) EA can find shortest paths in graphs by simulating the Bellman-Ford algorithm. The task is to find shortest paths from a source v^* to all other vertices. For vertices whose shortest paths have few edges, shortest paths are found quickly. In our terminology, these vertices would be called informed. If u is informed and

the graph contains an edge $\{u, v\}$, then v can become informed with a fixed probability during a lucky mutation, if the shortest path from v^* to v contains u . This way, shortest paths propagate through the graph in the same fashion as information does. The same can be observed for ant colony optimizers (Sudholt and Thyssen, 2012).

Doerr et al. (2007) independently used a different argument for bounding the expected propagation time. Fix a shortest path in the graph, leading from v^* to some fixed vertex v , then in every generation there is a chance of informing the first uninformed vertex on the path, until eventually the information reaches v . If the path has at least $\log n$ edges, the time until v is informed is highly concentrated. Using tail bounds, the probability of significantly exceeding the expectation is very small. This allows us to apply a union bound for all considered vertices v .

Following the proof of Doerr et al. (2007, Lemma 3), we get the following lemma. An advantage over the general bound from Rowe et al. (2008) is that it not only bounds the propagation time for the whole network; it also bounds expected hitting times for informing smaller numbers of vertices.

LEMMA 7: Consider propagation with transmission probability p on any undirected graph where initially a single vertex v^* is informed. For $i \in \mathbb{N}_0$ let V_i contain all vertices v whose shortest path from v^* to v contains i edges. Let $n_k := \sum_{i=1}^k |V_i|$. The probability of not having informed n_k vertices in time $\lambda k/p$, $\lambda \geq 2$, is at most

$$n_k \cdot \exp\left(-\frac{(\lambda - 1)^2}{2\lambda} \cdot k\right) \leq n_k \cdot \exp\left(-\frac{\lambda k}{8}\right).$$

The expected time until n_k vertices are informed is at most

$$\frac{\frac{c}{c-1} \cdot \max\{2k, 8 \ln(cn_k)\}}{p}$$

for every $c > 1$.

PROOF: The first claim follows from the proof of Lemma 3 in Doerr et al. (2007) and the fact that $(\lambda - 1)^2/\lambda \geq \lambda/4$ for $\lambda \geq 2$.

If $(8/k) \cdot \ln(cn_k) \geq 2$ we use $\lambda := (8/k) \cdot \ln(cn_k)$ and have that after $\lambda k/p$ iterations the probability of not having informed all vertices is at most

$$n_k \cdot \exp(-\ln(cn_k)) = \frac{1}{c}.$$

If not, we repeat the argument with another phase of $\lambda k/p$ iterations. As each phase is successful with probability at least $1 - 1/c$, the expected propagation time is at most

$$\frac{1}{1 - 1/c} \cdot \frac{\lambda k}{p} = \frac{\frac{c}{c-1} \cdot 8 \ln(cn_k)}{p}.$$

If $(8/k) \cdot \ln(cn_k) < 2$ then $k/4 > \ln(cn_k)$. The first statement with $\lambda := 2$ then gives a probability bound of

$$n_k \cdot \exp\left(-\frac{k}{4}\right) \leq n_k \cdot \exp(-\ln(cn_k)) \leq \frac{1}{c}$$

and using the same arguments as before we get a time bound of

$$\frac{1}{1 - 1/c} \cdot \frac{\lambda k}{p} = \frac{\frac{c}{c-1} \cdot 2k}{p}. \quad \square$$

Note that putting $k := \text{diam}(G)$ and $c = 2$, we get a bound of

$$\max \left\{ \frac{4 \text{diam}(G)}{p}, \frac{16 \ln(2n)}{p} \right\} \leq \frac{4 \text{diam}(G) + 11.2 \log(n) + 11.2}{p}.$$

For all nonempty graphs, this is better than the general upper bound

$$\frac{8 \text{diam}(G) + 8 \log n}{p(1 - e^{-1})} \approx \frac{12.7 \text{diam}(G) + 12.7 \log n}{p}$$

from Rowe et al. (2008). However, the asymptotic behavior of both bounds is the same (since $\max\{x, y\} = \Theta(x + y)$ for all $x, y \in \mathbb{R}_0^+$).

The next section will analyze parallel EAs with concrete topologies.

4 Parallel EAs with Ring Structures

Ring graphs are the starting point as they are often used as topologies (Tomassini, 2005). Rings can either be unidirectional, in which case there is exactly one directed cycle, or bidirectional, when all edges are undirected. The following theorem holds for both kinds of graphs, and in fact for all strongly connected graphs. Recall that a directed graph is called strongly connected if for every two vertices u, v there is a directed path from u to v (implying that there is also a path from v to u).

THEOREM 8: *Consider an island model running elitists EAs on a function f with a fitness-level partition $A_1 <_f \dots <_f A_m$ and success probabilities s_1, \dots, s_{m-1} . Let p be (a lower bound on) the probability that a specific island on fitness level i informs a specific neighbor in the topology in one generation. The expected parallel runtime on f with a unidirectional or bidirectional ring and μ islands—or in fact any strongly connected topology—is bounded from above by*

$$\frac{2}{p^{1/2}} \sum_{i=1}^{m-1} \frac{1}{s_i^{1/2}} + \frac{1}{\mu} \cdot \sum_{i=1}^{m-1} \frac{1}{s_i}.$$

The expected communication effort for ring graphs is by a factor of at most $2p\mu$ larger than the expected parallel time.

The shape of this formula deserves some explanation. The second term $\frac{1}{\mu} \cdot \sum_{i=1}^{m-1} \frac{1}{s_i}$ is by a factor of μ smaller than the upper bound for a single island by Theorem 3. If the latter is asymptotically tight, the second term in Theorem 8, regarded in isolation, would give a perfect linear speedup. The first term is related to the speed at which information is propagated; it reflects the time needed to bring a reasonably large number of islands to the current best fitness level. Unlike for the second term, it is independent of μ , but it depends on the transmission probability p . We do have a linear speedup if the first term $\frac{2}{p^{1/2}} \sum_{i=1}^{m-1} \frac{1}{s_i^{1/2}}$ asymptotically does not grow faster than the second term, again assuming that the bound for a single island is tight.

As μ grows, the second term becomes smaller, while the first term remains fixed. Thus, if we have a linear speedup for small μ , there is a point where with growing μ the linear speedup disappears. This threshold can be easily computed by checking which value of μ gives rise to the first and second terms being of equal asymptotic order. As will be seen in the next sections, the same also holds for other migration topologies.

PROOF OF THEOREM 8: For the unidirectional ring we have $E(\xi(k)) \leq (k - 1)/p$ since a new island is informed with probability at least p . As this happens independently in each generation, the expected wait time until this happens is at most $1/p$. In fact,

this argument holds for all strongly connected topologies and in particular for the bidirectional ring.

Now, if $1 \leq k := p^{1/2}/s_i^{1/2} \leq \mu$ (ignoring rounding issues), by Lemma 6 the expected number of generations on fitness level i is bounded from above by

$$\frac{k-1}{p} + 1 + \frac{1}{k} \cdot \frac{1}{s_i} \leq \frac{1}{p^{1/2}s_i^{1/2}} + \frac{1}{p^{1/2}s_i^{1/2}} = \frac{2}{p^{1/2}s_i^{1/2}}.$$

In the case where $p^{1/2}/s_i^{1/2} < 1$, we trivially get an upper bound of

$$\frac{1}{s_i} \leq \frac{1}{p^{1/2}s_i^{1/2}}.$$

If $p^{1/2}/s_i^{1/2} > \mu$, Lemma 6 for $k := \mu$ gives an upper bound of

$$\frac{\mu-1}{p} + 1 + \frac{1}{\mu} \cdot \frac{1}{s_i} < \frac{1}{p^{1/2}s_i^{1/2}} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Taking the maximum of the above upper bounds gives

$$\max\left(\frac{2}{p^{1/2}s_i^{1/2}}, \frac{1}{p^{1/2}s_i^{1/2}} + \frac{1}{\mu} \cdot \frac{1}{s_i}\right) \leq \frac{2}{p^{1/2}s_i^{1/2}} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Summing over all fitness levels proves the claim.

The claim on the expected communication effort follows from Theorem 2. □

As remarked in the proof, the bound from Theorem 8 holds for arbitrary strongly connected topologies as the unidirectional ring is a worst case for the μ_i values. Along with Theorem 2, this also gives a general upper bound on the expected communication effort for any strongly connected topology.

For bidirectional rings, we have $E(\xi(k)) \leq \frac{k}{2p}$. This can be seen from applying Johannsen’s drift theorem, stated in the appendix as Theorem 20, applied to the difference between k and the current number of informed vertices. If there is more than one uninformed vertex, there are always at least two vertices neighboring to informed ones. The number of informed vertices then increases by $2p$ in expectation. This means that we can use $h(1) = p$ and $h(x) = 2p$ for $x > 1$ as the drift function. This decreases the constant 2 in the first term toward $\sqrt{2}$, at the expense of an additional term $m - 1$. In some settings this upper bound may be better than the upper bound from Theorem 8 for unidirectional rings; where it is not, we may still use the latter for bidirectional rings as Theorem 8 applies to all strongly connected topologies.

Also note that if $p < s_i$ then the trivial bound $1/s_i$ gives a better estimate for the time that this fitness level remains. If this holds for all fitness levels, information is propagated too slowly and our method does not give any provable speedups for the parallel model.

Contrarily, if, say, $p = \Omega(1)$, compared to a single island in a ring, the expected wait time for every fitness level can be replaced by its square root. This can yield significant speedups. We make this precise for concrete functions in the following theorem. For comparing these times with runtime bounds for the (1+1) EA, we refer to Table 1.

THEOREM 9: *The following holds for the parallel (1+1) EA with transmission probability at least p on a unidirectional or bidirectional ring (or any other strongly connected topology):*

- $E(T^{par}) = O\left(\frac{n}{p^{1/2}} + \frac{n \log n}{\mu}\right)$ for OneMax,

- $E(T^{par}) = O\left(\frac{dn^{1/2}}{p^{1/2}} + \frac{dn}{\mu}\right)$ for every unimodal function with $d + 1$ function values,
- $E(T^{par}) = O\left(\frac{n^{k/2}}{p^{1/2}} + \frac{n^k}{\mu}\right)$ for Jump_k with $k \geq 2$.

PROOF: For OneMax, we choose the canonical partition $A_i := \{x \mid \text{OneMax}(x) = i\}$. The probability of increasing the current fitness from fitness level i is at least $s_i \geq (n - i) \cdot 1/(en)$, since there are $n - i$ Hamming neighbors of larger fitness and a specific Hamming neighbor is created with probability at least $1/n \cdot (1 - 1/n)^{n-1} \geq 1/(en)$. The second sum in Theorem 8 is

$$\frac{1}{\mu} \cdot \sum_{i=0}^{n-1} \frac{en}{n-i} = \frac{en}{\mu} \sum_{i=1}^n \frac{1}{i} = O\left(\frac{n \log n}{\mu}\right).$$

The first sum in Theorem 8 is

$$\begin{aligned} 2 \sum_{i=0}^{n-1} \left(\frac{en}{n-i} \cdot \frac{1}{p}\right)^{1/2} &= 2 \left(\frac{en}{p}\right)^{1/2} \sum_{i=1}^n \frac{1}{\sqrt{i}} \\ &\leq 2 \left(\frac{en}{p}\right)^{1/2} \int_0^n \frac{1}{\sqrt{i}} \, di \leq 2 \left(\frac{en}{p}\right)^{1/2} \cdot \sqrt{n} = O(n/p^{1/2}). \end{aligned}$$

For unimodal functions we choose a partition A_1, \dots, A_{d+1} where A_i contains all search points with the i th smallest function value. The probability of improving the fitness from level i is at least $s_i \geq 1/(en)$ because there is at least one search point in the next fitness level which is at Hamming distance one. Theorem 8 gives an upper bound of

$$2 \sum_{i=1}^d \left(\frac{en}{p}\right)^{1/2} + \sum_{i=1}^d \frac{en}{\mu} \leq 2d \cdot \left(\frac{en}{p}\right)^{1/2} + \frac{den}{\mu} = O\left(\frac{dn^{1/2}}{p^{1/2}} + \frac{dn}{\mu}\right).$$

For Jump_k we again choose A_i to contain all search points with fitness i , leading to sets A_1, \dots, A_n , and A_{n+k} . The levels A_1, \dots, A_{k-1} resemble those for OneMax (modulo swapping the meaning of zeros and ones), and the same estimations for the s_i apply. The same holds for levels A_k, \dots, A_{n-1} , where A_i corresponds to a success probability of s_{i-k} in the context of OneMax. Thus, the expected time for getting to either of the two best fitness levels, the set $A_n \cup A_{n+k}$, is bounded by twice the above upper bound for OneMax. In case we have reached A_n , to reach the highest level A_{n+k} , a specific bit string with Hamming distance k has to be created. This has a probability of at least

$$s_n \geq \left(\frac{1}{n}\right)^k \cdot \left(1 - \frac{1}{n}\right)^{n-k} \geq \left(\frac{1}{n}\right)^k \cdot \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{en^k}.$$

Theorem 8 and the above bound for OneMax give

$$O\left(\frac{n}{p^{1/2}} + \frac{n \log n}{\mu}\right) + 2 \left(\frac{en^k}{p}\right)^{1/2} + \frac{en^k}{\mu} = O\left(\frac{n^{k/2}}{p^{1/2}} + \frac{n^k}{\mu}\right). \quad \square$$

The speedups obtained are indeed significant, particularly for those functions where improvements are hard to find.

The proof of Theorem 9 uses well-known fitness-level partitions (Droste et al., 2002; Wegener, 2002), and hence it simply consists of plugging in known values s_i and simplifying. This shows how easy it is to obtain results for parallel EAs based on analyses of panmictic EAs.

5 Parallel EAs with Two-Dimensional Grids and Tori

For two-dimensional grids and tori, Theorem 4 is adapted in a similar manner, in order to get the best possible leading constant in the first term of the runtime bound. We also consider applications of the resulting theorem similar to the applications for ring graphs.

THEOREM 10: *Consider an island model running elitists EAs on a function f with a fitness-level partition $A_1 <_f \dots <_f A_m$ and success probabilities s_1, \dots, s_{m-1} . Let p be (a lower bound on) the probability that a specific island on fitness level i informs a specific neighbor in the topology in one generation. The expected parallel runtime of the island model on a grid or torus topology whose side lengths are at least $\sqrt{\mu}$ in both directions is bounded from above by*

$$\frac{3^{5/3}}{p^{2/3}} \sum_{i=1}^{m-1} \frac{1}{s_i^{1/3}} + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}.$$

The expected communication effort for μ vertex topologies is by a factor of at most $4p\mu$ larger than the expected parallel time.

PROOF: Note that within a square area of $\sqrt{k} \times \sqrt{k}$ vertices in the graph, all shortest paths between any two vertices have at most $2\sqrt{k} - 2$ edges. Thus, for every vertex in this area, the number of vertices that can be reached via up to $k' := 2\sqrt{k} - 2$ edges is $n'_k \geq k$. The first, pessimistic, assumption is that $n'_k = k$; if not, a slower propagation is considered where the edges of the graph are removed to ensure $n'_k = k$. Applying Lemma 7 (with respect to the primed variables $k' := 2\sqrt{k} - 2$, $n'_k = k$ and $c = 4$) gives, for every $k \leq \mu$, the expected time until k islands are informed is bounded from above by

$$\max \left\{ \frac{16/3 \cdot \sqrt{k} - 16/3}{p}, \frac{32/3 \cdot \ln(4k)}{p} \right\}.$$

We also get an upper bound of $k/(2p)$ using Johannsen’s (2010) variable drift theorem, shown in Theorem 20 in the appendix, as before. If there is more than one uninformed vertex, there are always at least two vertices neighboring to informed ones. Thus, the expected number of informed vertices increases by $2p$ in expectation. Applying Johannsen’s drift theorem as for the bidirectional ring gives an upper bound of $k/(2p)$. It is easy to check that the best upper bound is as follows: for all $k \in \mathbb{N}$

$$\min \left\{ \frac{k}{2p}, \max \left\{ \frac{16/3 \cdot \sqrt{k} - 16/3}{p}, \frac{32/3 \cdot \ln(4k)}{p} \right\} \right\} \leq \frac{6\sqrt{k} - 1}{p}.$$

Now, if $1 \leq k := 3^{-2/3} \cdot (p/s_i)^{2/3} \leq \mu$ (ignoring rounding issues) by Lemma 6 the expected number of generations on fitness level i is bounded from above by

$$\frac{6\sqrt{k} - 1}{p} + 1 + \frac{1}{k} \cdot \frac{1}{s_i} \leq \frac{6 \cdot 3^{-1/3}}{p^{2/3}s_i^{1/3}} + \frac{3^{2/3}}{p^{2/3}s_i^{1/3}} = \frac{3^{5/3}}{p^{2/3}s_i^{1/3}}.$$

If $3^{-2/3} \cdot (p/s_i)^{2/3} < 1$, then the upper bound is, trivially,

$$\frac{1}{s_i} \leq \frac{3^{2/3}}{p^{2/3}s_i^{1/3}}.$$

If $3^{-2/3} \cdot (p/s_i)^{2/3} > \mu$, we get for $k := \mu$ an upper bound of

$$\frac{6\sqrt{\mu} - 1}{p} + 1 + \frac{1}{\mu} \cdot \frac{1}{s_i} < \frac{2 \cdot 3^{2/3}}{p^{2/3}s_i^{1/3}} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Taking the maximum of the above upper bounds gives

$$\max \left(\frac{3^{5/3}}{p^{2/3}s_i^{1/3}}, \frac{2 \cdot 3^{2/3}}{p^{2/3}s_i^{1/3}} + \frac{1}{\mu} \cdot \frac{1}{s_i} \right) \leq \frac{3^{5/3}}{p^{2/3}s_i^{1/3}} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Summing over all fitness levels yields the claim.

The claim on the expected communication effort follows from Theorem 2. □

Note that the communication effort in one generation is asymptotically as large as for ring graphs, but for large p the upper bound on the parallel runtime is generally smaller (or asymptotically equal, in the case where upper bounds are dominated by the term $1/\mu \cdot \sum_{i=1}^{m-1} 1/s_i$). If $p < 3s_i$, then again the trivial upper bound $1/s_i$ is better, because then the spread of information is too slow.

Compared to a single island, in a torus, the expected wait time for every fitness level can be replaced by its third root. This leads to improved upper bounds for unimodal functions and Jump_k .

THEOREM 11: *The following holds for the parallel (1+1) EA with transmission probability p on a grid or torus topology whose side lengths are at least $\sqrt{\mu}$ in both directions:*

- $E(T^{par}) = O\left(\frac{n}{p^{2/3}} + \frac{n \log n}{\mu}\right)$ for *OneMax*,
- $E(T^{par}) = O\left(\frac{dn^{1/3}}{p^{2/3}} + \frac{dn}{\mu}\right)$ for every unimodal function with $d + 1$ function values,
- $E(T^{par}) = O\left(\frac{n+n^{k/3}}{p^{2/3}} + \frac{n^k}{\mu}\right)$ for Jump_k with $k \geq 2$.

PROOF: We choose the same partitions as in the proof of Theorem 9. Note that the second terms in Theorem 8 and 10 are identical, so we only estimate the first terms and refer to Theorem 9 for the second terms.

For *OneMax*, the first sum in Theorem 10 is

$$\begin{aligned} \frac{3^{5/3}}{p^{2/3}} \sum_{i=0}^{n-1} \left(\frac{en}{n-i}\right)^{1/3} &= \frac{3^{5/3}e^{1/3}n^{1/3}}{p^{2/3}} \sum_{i=1}^n \left(\frac{1}{i}\right)^{1/3} \\ &\leq \frac{3^{5/3}e^{1/3}n^{1/3}}{p^{2/3}} \int_{i=0}^n \left(\frac{1}{i}\right)^{1/3} di \\ &= \frac{3^{5/3}e^{1/3}n^{1/3}}{p^{2/3}} \cdot \frac{3}{2} \cdot n^{2/3} = \frac{3^{8/3}/2 \cdot e^{1/3}n}{p^{2/3}}. \end{aligned}$$

This gives an upper bound of

$$O\left(\frac{n}{p^{2/3}} + \frac{n \log n}{\mu}\right).$$

For unimodal functions, Theorem 10 gives

$$\frac{3^{5/3}}{p^{2/3}} \sum_{i=1}^d (en)^{1/3} + \sum_{i=1}^d \frac{en}{\mu} \leq \frac{3^{5/3}d \cdot e^{1/3}n^{1/3}}{p^{2/3}} + \frac{den}{\mu} = O\left(\frac{dn^{1/3}}{p^{2/3}} + \frac{dn}{\mu}\right).$$

For Jump_k , the result is

$$O\left(\frac{n}{p^{2/3}} + \frac{n \log n}{\mu}\right) + 3^{5/3} \cdot \frac{(en^k)^{1/3}}{p^{2/3}} + \frac{en^k}{\mu} = O\left(\frac{n+n^{k/3}}{p^{2/3}} + \frac{n^k}{\mu}\right). \quad \square$$

6 Parallel EAs with Hypercube Graphs

Hypercube graphs are popular topologies in parallel computation. In a d -dimensional hypercube, each vertex has a label of d bits. Two vertices are neighboring if and only if their labels differ in exactly one bit. The number of vertices is then 2^d , and each vertex has d neighbors. The diameter of a d -dimensional hypercube is d ; hence, only logarithmic in the size of the graph. The small diameter implies that in many communication models information is spread rapidly, even though the degree of vertices is quite small. With regard to the propagation process investigated here, we get a small first term in the following runtime bound, and still have a very moderate communication effort.

THEOREM 12: *Consider an island model running elitists EAs on a function f with a fitness-level partition $A_1 <_f \dots <_f A_m$ and success probabilities s_1, \dots, s_{m-1} . Let p be (a lower bound on) the probability that a specific island on fitness level i informs a specific neighbor in the topology in one generation. The expected parallel runtime of the island model on a $(\log \mu)$ -dimensional hypercube graph with μ islands is bounded from above by*

$$\frac{25m + 12 \sum_{i=1}^{m-1} \log(\frac{1}{s_i})}{p} + \frac{1}{\mu} \cdot \sum_{i=1}^{m-1} \frac{1}{s_i}.$$

The expected communication effort is by a factor of at most $p\mu \log \mu$ larger than the expected parallel time.

PROOF: In the notation of Lemma 7, given the hypercube and $1 \leq k \leq \log \mu$, the number of vertices reachable with at most k edges is bounded from below by

$$n_k = \sum_{i=1}^k \binom{\mu}{i} \geq 2^k.$$

As in the proof of Theorem 8, $n'_k = 2^k$ is considered instead of n_k , justified by the fact that we can remove edges in the graph to slow down propagation. Invoking Lemma 7 with $k, n'_k = 2^k$, and $c = 2$, the expected time until 2^k vertices are informed is therefore at most

$$\frac{16 \ln(2 \cdot 2^k)}{p} = \frac{16 \cdot (\ln 2) \cdot (k + 1)}{p} < \frac{12(k + 1)}{p}.$$

By Lemma 6, the expected time on fitness level i is hence bounded, for any integer $0 \leq k \leq \log \mu$, by

$$\frac{12(k + 1)}{p} + 1 + \frac{1}{2^k} \cdot \frac{1}{s_i} \leq \frac{13}{p} + \frac{12k}{p} + \frac{1}{2^k} \cdot \frac{1}{s_i}. \tag{4}$$

If $p/(12s_i) < 1$, there is a trivial upper bound of $1/s_i \leq 12/p$. If $p/(12s_i) > \mu$, which implies $d < \log(p/(12s_i)) \leq \log(1/s_i)$, there is an upper bound of

$$\frac{13 + 12d}{p} + \frac{1}{\mu} \cdot \frac{1}{s_i} < \frac{13 + 12 \log(\frac{1}{s_i})}{p} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Otherwise, Equation (4) is minimized for $2^k = p/(12s_i)$, leading to

$$\frac{13}{p} + \frac{12 \log(\frac{p}{12s_i})}{p} + \frac{12}{p} \leq \frac{25}{p} + \frac{12 \log(\frac{1}{s_i})}{p}.$$

The maximum over all these bounds is at most

$$\frac{25 + 12 \log(\frac{1}{s_i})}{p} + \frac{1}{\mu} \cdot \frac{1}{s_i}.$$

Summing over all fitness levels yields the claim.

The claim on the expected communication effort follows from Theorem 2. □

The results for the example applications are as follows.

THEOREM 13: *The following holds for the parallel (1+1) EA with transmission probability p on a $(\log \mu)$ -dimensional hypercube:*

- $E(T^{par}) = O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right)$ for OneMax,
- $E(T^{par}) = O\left(\frac{d \log n}{p} + \frac{dn}{\mu}\right)$ for every unimodal function with $d + 1$ function values,
- $E(T^{par}) = O\left(\frac{n+k \log n}{p} + \frac{n^k}{\mu}\right)$ for Jump_k with $k \geq 2$.

PROOF: For OneMax,

$$\begin{aligned} \sum_{i=0}^{n-1} \log\left(\frac{1}{s_i}\right) &= \sum_{i=1}^n \log\left(\frac{en}{i}\right) = \log\left(\prod_{i=1}^n \frac{en}{i}\right) \\ &= \log\left(\frac{e^n n^n}{n!}\right) \leq \log\left(\frac{e^n n^n}{(n/e)^n}\right) = \log(e^{2n}) = 2n \log(e). \end{aligned}$$

Theorem 12 gives an upper bound of

$$\frac{25n + 24n \log e}{p} + O\left(\frac{n \log n}{\mu}\right) = O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right).$$

For unimodal functions, Theorem 12 gives

$$\frac{25d + 12d \log(en)}{p} + O\left(\frac{dn}{\mu}\right) = O\left(\frac{d \log n}{p} + \frac{dn}{\mu}\right).$$

For Jump_k , using $\log(en^k) \leq \log(e^k n^k) = O(k \log n)$ yields

$$\begin{aligned} O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right) + \frac{25 + 12 \log(en^k)}{p} + O\left(\frac{n^k}{\mu}\right) \\ = O\left(\frac{n + k \log n}{p} + \frac{n^k}{\mu}\right). \end{aligned} \quad \square$$

If $p = \Omega(1)$, linear speedups arise for OneMax if $\mu = O(\log n)$, and linear speedups for unimodal functions where the bound $O(dn)$ for a single island is tight, if $\mu = O(n/\log n)$. For Jump_k , if $k = O(n/\log n)$, $\mu = O(n^{k-1})$ can be chosen to get a linear speedup. Note from Table 1 that the upper bounds on the expected parallel times for LO and Jump_k are much better for the hypercube than for rings and torus graphs, if p is large.

7 Parallel EAs with Complete Topologies

Finally, the densest topology is considered, the complete graph K_μ , where every island is a neighbor to every other island. The complete graph is interesting because it represents an extreme case: the largest possible communication effort with regard to one generation, but also the fastest possible spread of information.

For the special case of $p = 1$, a parallel (1+1) EA is basically equivalent to a $(1+\mu)$ EA, which creates μ offspring independently and then compares a best offspring against the current search point. The only difference is that the parallel (1+1) EA can store different individuals of the same fitness. However, this issue is irrelevant when using the fitness-level method. Hence, the results for a parallel (1+1) EA with a complete topology and $p = 1$ also apply to the $(1 + \mu)$ EA. For $p < 1$, the two models are generally different.

To state a simple argument, clearly, if there is at least one informed island, each other island will become informed with a probability of at least p .

THEOREM 14: *Consider an island model running elitists EAs on a function f with a fitness-level partition $A_1 <_f \dots <_f A_m$ and success probabilities s_1, \dots, s_{m-1} . Let p be (a lower bound on) the probability that a specific island on fitness level i informs a specific neighbor in the topology in one generation. The expected parallel runtime of the island model on a complete topology is*

$$E(T^{par}) \leq m + \frac{2m}{p} + \frac{2}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}.$$

The expected communication effort is by a factor of at most $p\mu(\mu - 1) < p\mu^2$ larger than the expected parallel time.

PROOF: The expected time is estimated until at least $\mu/2$ islands are informed after an improvement. If more than $\mu/2$ islands are uninformed, the expected number of islands that become informed in one generation is at least $p\mu/2$. By standard drift analysis arguments (He and Yao, 2004) the desired expectation is bounded from above by $2/p$.

By Lemma 6, the expected time on fitness level i is at most

$$1 + \frac{2}{p} + \frac{2}{\mu} \cdot \frac{1}{s_i}.$$

Adding these times for all fitness levels proves the claim.

The claim on the expected communication effort follows from Theorem 2. □

As mentioned previously, the complete graph leads to a maximal spread of information. In addition, the best upper bounds for the considered function classes are found. However, a maximum amount of migration takes place in each generation, so the expected total communication effort is also highest (cf. Tables 1 and 3; 3 is shown in the appendix).

THEOREM 15: *Let $\mu \in \mathbb{N}$. The following holds for the expected parallel runtime of the parallel (1+1) EA with topology K_μ . In the case where $p = 1$, the same holds for the $(1+\mu)$ EA:*

- $E(T^{par}) = O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right)$ for *OneMax*,
- $E(T^{par}) = O\left(\frac{d}{p} + \frac{dn}{\mu}\right)$ for every unimodal function with $d + 1$ function values, and
- $E(T^{par}) = O\left(\frac{n}{p} + \frac{n^k}{\mu}\right)$ for *Jump_k* with $k \geq 2$.

The proof is obvious by now.

The term $2/p$ for the time until at least $\mu/2$ islands are informed is a reasonable estimate if p is large (e.g., $p = \Omega(1)$). However, for small p , this estimation is quite loose

as we have completely neglected that all informed vertices have a chance to inform other islands.

This work, therefore, also presents a more detailed analysis for small p . The motivation for studying complete graphs and small p is that it captures random migration policies. Assume that each island decides randomly with probability p for each other island whether to migrate individuals to that island. Then this can be regarded as a complete topology with transmission probability p .

Values around $p = 1/\mu$ seem particularly interesting as then in each generation one migration takes place for each island in expectation. In fact, there are different results for $p > 1/\mu$ and $p < 1/\mu$.

LEMMA 16: Consider propagation with transmission probability p on the complete topology with μ vertices. Let $\xi(k)$ be as in Lemma 6, then

$$\xi(\mu) \leq \frac{8 \log(\mu)}{\min(p\mu, 1)}.$$

PROOF: The claim is obvious for $\mu \leq 2$, so $\mu \geq 3$ is assumed in the following. Let X_t denote the random number of informed vertices after t iterations. First, the expected time until at least $\mu/2$ vertices become informed is estimated, and then an estimate is made of how long it takes to get from $\mu/2$ informed vertices to μ informed ones.

If $X_t = i$, each presently uninformed vertex is being informed in one iteration with probability (using Lemma 19 in the appendix)

$$1 - (1 - p)^i \geq 1 - \frac{1}{1 + ip} = \frac{ip}{1 + ip}.$$

This holds independently from other presently uninformed vertices. In fact, the number of newly informed vertices follows a binomial distribution with parameters $\mu - i$ and $\frac{ip}{1+ip}$. The median of this binomial distribution is $i(\mu - i) \cdot \frac{p}{1+ip}$ (assuming that this is an integer); hence, with probability at least $1/2$, we have at least $i(\mu - i) \cdot \frac{p}{1+ip}$ newly informed vertices in one iteration. Hence, it takes an expected number of at most two iterations to increase the number of informed vertices by $i(\mu - i) \cdot \frac{p}{1+ip}$, which for $i \leq \mu/2$ is at least $i \cdot \frac{p\mu}{2+p\mu}$.

For every $0 \leq j \leq \log(\mu) - 2$, the following holds. If $i \geq 2^j$ and $i \leq \mu/2$, then in an expected number of two generations, at least $2^j \cdot \frac{p\mu}{2+p\mu}$ new vertices are informed. The expected number of iterations for informing a total of 2^j new vertices is therefore at most $2 \cdot \frac{2+p\mu}{p\mu}$. Therefore, the number of informed vertices has gone from at least 2^j to at least 2^{j+1} informed vertices. Summing up all times across all j , the expected time until at least $2^{\log(\mu)-1} = \mu/2$ vertices are informed is at most

$$2(\log(\mu) - 1) \cdot \frac{2 + p\mu}{p\mu}.$$

For $p\mu \leq 1$ we have $\frac{2+p\mu}{p\mu} \leq 3/(p\mu)$, yielding an upper time bound of $6(\log(\mu) - 1)/(p\mu)$. Otherwise, we use $\frac{2+p\mu}{p\mu} \leq 3$ to get a bound of $6(\log(\mu) - 1)$.

For the time to get from $\mu/2$ to μ informed vertices, observe that the expected number of newly informed vertices is still $\frac{i(\mu-i)p}{1+ip}$, if currently i vertices are informed. Equivalently, if i vertices are uninformed, the expected decrease of the number of uninformed vertices is $\frac{i(\mu-i)p}{1+(\mu-i)p}$. This function is monotonically increasing if $i \leq \mu/2$.

Applying Johannsen’s drift theorem (Theorem 20 in the appendix), for the number of uninformed nodes, using $h(i) := \frac{i(\mu-i)p}{1+(\mu-i)p}$ as the drift function, gives an upper bound of

$$\begin{aligned} & \frac{1 + (\mu - 1)p}{(\mu - 1)p} + \int_1^{\mu/2} \frac{1 + (\mu - i)p}{i(\mu - i)p} \, di \\ & \leq \frac{1 + (\mu - 1)p}{(\mu - 1)p} + \frac{\ln(\mu - 1)(1 + p\mu)}{p\mu} \\ & \leq \frac{1 + p\mu}{p\mu} + \frac{1}{\mu(\mu - 1)p} + \frac{\ln(\mu - 1)(1 + p\mu)}{p\mu} \\ & \leq \frac{(\ln(\mu) + 1)(1 + p\mu) + \frac{1}{\mu-1}}{p\mu}. \end{aligned}$$

For $p\mu \leq 1$, this is at most

$$\frac{2 \ln(\mu) + 2 + \frac{1}{\mu-1}}{p\mu} \leq \frac{2 \ln(\mu) + 5/2}{p\mu}.$$

Otherwise, this is at most

$$\frac{(\ln(\mu) + 1) \cdot 2p\mu + \frac{p\mu}{\mu-1}}{p\mu} \leq 2 \ln(\mu) + 5/2.$$

Together, along with $\ln(\mu) \leq \log(\mu)$, this proves the claim. □

Combining Lemma 16 with Lemma 6 gives the following. Apart from an additive term m , the case of $p \leq 1/\mu$ yields a bound where the first term is smaller by a factor on the order of $\log(\mu)/\mu$. For fairly large transmission probabilities, $p \geq 1/\mu$, in the first term we have replaced the factor $1/p$ by $\log(\mu)$. These improvements reflect that the complete graph can spread information much more quickly than previously estimated in the proof of Theorem 14.

THEOREM 17: *Consider an island model running elitist EAs on a function f with a fitness-level partition $A_1 <_f \dots <_f A_m$ and success probabilities s_1, \dots, s_{m-1} . Let p be (a lower bound on) the probability that a specific island on fitness level i informs a specific neighbor in the topology in one generation. The expected parallel runtime of the island model on a complete topology is bounded as follows. If $p \geq 1/\mu$,*

$$E(T^{par}) \leq m + 8m \log \mu + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}$$

and if $p \leq 1/\mu$,

$$E(T^{par}) \leq m + \frac{8m \log \mu}{p\mu} + \frac{1}{\mu} \sum_{i=1}^{m-1} \frac{1}{s_i}.$$

For the example applications, the refinements in Theorem 17 result in the following refined bounds. As there are only improved upper bounds for $p = O(1/\mu)$, it is not necessary to mention the special case of the $(1+\mu)$ EA with $p = 1$.

THEOREM 18: *Let $\mu \in \mathbb{N}$. The following holds for the expected parallel runtime of the parallel $(1+1)$ EA with topology K_μ :*

- $E(T^{par}) = O\left(n \log(\mu) + \frac{n \log n}{\mu}\right)$ for *OneMax* if $p \geq 1/\mu$ and $E(T^{par}) = O\left(n + \frac{n \log \mu}{p\mu} + \frac{n \log n}{\mu}\right)$ otherwise,
- $E(T^{par}) = O\left(d \log(\mu) + \frac{dn}{\mu}\right)$ for unimodal functions with $d + 1$ values, if $p \geq 1/\mu$, and $E(T^{par}) = O\left(d + \frac{d \log(\mu)}{p\mu} + \frac{dn}{\mu}\right)$ otherwise, and
- $E(T^{par}) = O\left(n \log(\mu) + \frac{n^k}{\mu}\right)$ for *Jump_k* with $k \geq 2$, if $p \geq 1/\mu$ and $E(T^{par}) = O\left(n + \frac{n \log(\mu)}{p\mu} + \frac{n^k}{\mu}\right)$ otherwise.

8 Experiments

In order to complement the analytical results above, this paper now gives experimental results on the behavior of island models for different topologies. As a detailed experimental evaluation is beyond the scope of this paper, only illustrative results for the two functions *OneMax* and *LO* are shown.

The first investigation is of the parallel runtime T^{par} and the communication effort for different transmission probabilities. The experiments were repeated 100 times per data point for the parallel (1+1) EA with $\mu = 64$ islands and an instance size of $n = 256$ for both example functions, varying the transmission probability p in steps of 0.01. Figure 1 shows the behavior for the topologies K_{64} , a bidirectional ring graph, an 8×8 torus graph, and a 6-dimensional hypercube.

Looking at the influence of the transmission probability on the runtime, a higher transmission probability improves the runtime behavior of the algorithm, also according to the expectations from the theoretical analysis. In particular, all values p that are not too small lead to much smaller runtimes compared to the pathological setting $p = 0$, where there is no communication, but μ independent runs of the (1+1) EA. This demonstrates for our functions that parallelization and migration can lead to drastic speedups. For larger or intermediate values for p , the parallel runtime does not vary much, as then for all topologies the runtime is dominated by the second terms from our bounds: $1/\mu \cdot O(n \log n)$ and $1/\mu \cdot O(n^2)$ for *OneMax* and *LO*, respectively.

Comparing the behavior of those topologies, we see that the parallel runtime indeed depends on the density of the topology, that is, more dense topologies spread information more efficiently, which results in a faster convergence. As expected, the topology K_μ performs best, and the ring graph performs worst.

Two-sided Mann-Whitney U tests were used on the data from Figure 1(a) and 1(b), along with a comparison of mean ranks, to make pairwise comparisons between the topologies concerning the parallel runtime. Separate tests were performed for each individual data point (e.g., each tested transition probability), as this illuminates in which settings one topology is better than another.³ For both *OneMax* and *LO* and all transmission probabilities of at least 0.01, the outcome is that $K_\mu < \text{hypercube} < \text{torus} < \text{ring}$ on a significance level of 0.001.

³This is particularly relevant as for the upcoming experiments in Figure 2 some of the curves cross, and then it cannot be said that one topology is always better than another. The large number of tests means that the possibility cannot be excluded that some of the results may be false positives. To minimize this issue, results are reported for a very low significance level of 0.001 wherever possible.

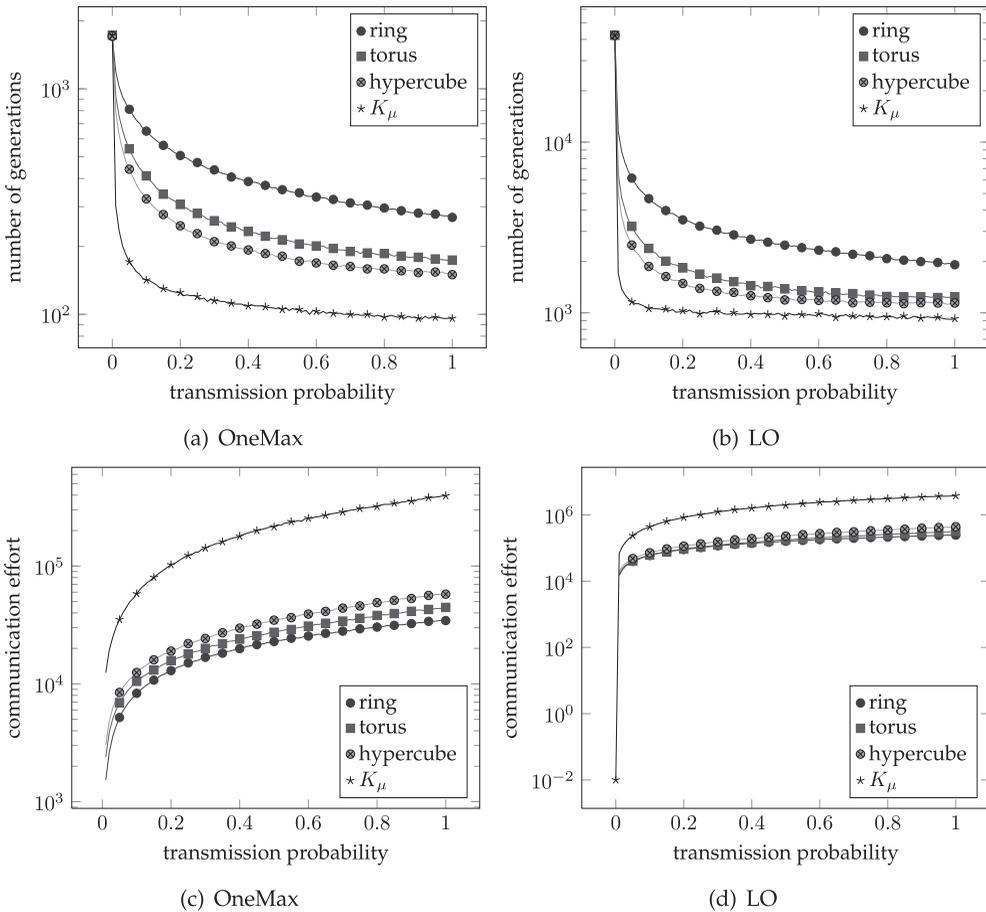


Figure 1: Average parallel runtime and communication effort for the parallel (1+1) EA using $\mu = 64$ islands and different transmission probabilities, both for OneMax and LO on $n = 256$ bits. The lines connect all data points, that is, varying transmission probabilities in increments of 0.01; plot marks are shown in increments of 0.05.

Looking at the communication effort, Figure 1(c) and 1(d), it seems that it is larger for more dense topologies, as expected. Hence, although the topology K_μ shows the best runtime behavior, the communication effort is highest for all transmission probabilities. Interestingly, the communication effort is about the same for the other three topologies. This is in particular the case for LO, that is, although, for example, the ring graph is more sparse, its parallel runtime is higher, so that the communication effort remains similar to the hypercube and the torus graph.

Applying the Mann-Whitney U test for the communication effort and OneMax yields the result that $\text{ring} < \text{torus} < \text{hypercube} < K_\mu$ for all transmission probabilities at least 0.01 on a significance level of 0.001. Looking at LO, the level of significance is slightly less when comparing ring and torus. For the communication effort and LO, the relation $\text{ring} < \text{torus}$ for a transmission probability of at least 0.01 only holds on a significance level of 0.05. Additionally, there are three exceptions: for transmission probabilities 0.13, 0.14, and 0.17, the results were not significant.

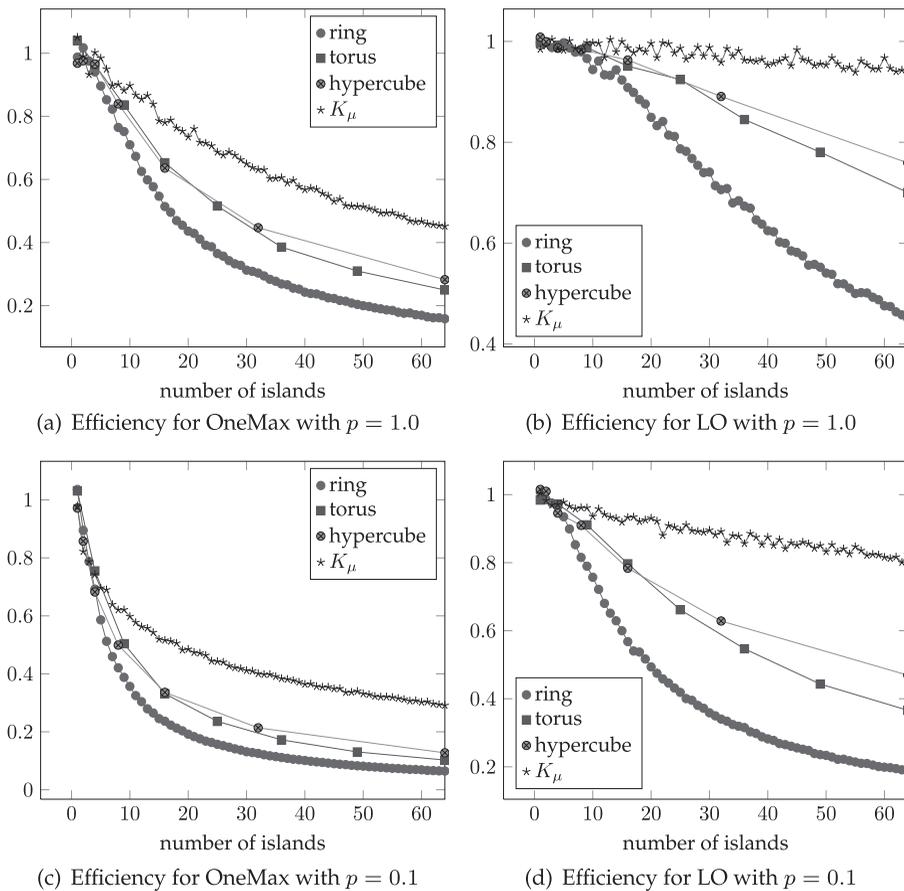


Figure 2: Efficiency for the parallel (1+1) EA with transmission probabilities $p \in \{0.1, 1\}$ for $\mu \in \{1, \dots, 64\}$ numbers of islands.

Next, the impact of the number of islands on the performance is investigated, with regard to different topologies and transmission probabilities, see Figures 2(a) and 2(b) for a transmission probability $p = 1.0$ and Figures 2(c) and 2(d) for a transmission probability $p = 0.1$. As the parallel runtime shows a steep decrease, the efficiency is plotted instead, defined as

$$\frac{T^{\text{seq}}}{T^{\text{par}} \cdot \mu}.$$

It can be regarded as a normalized version of speedup, normalized by the number of islands. Small efficiencies indicate small speedups, whereas large efficiencies indicate good speedups. An efficiency of 1 corresponds to a perfect linear speedup.

Again, the instance size of the benchmark functions was set to $n = 256$ and the number of islands μ was chosen from 1 to 64. Only square torus graphs were used. Therefore, our torus graphs and hypercubes are only defined for square numbers and powers of 2, respectively, leading to fewer data points.

For lower numbers of islands, the efficiency of the algorithm is better than for larger numbers of islands. This is expected, as a single (1+1) EA, that is, the setting with $\mu = 1$

Table 2: A summary of test results regarding pairwise comparisons between the different migration topologies with Mann-Whitney U tests, according to the efficiency shown in Figure 2. Here the lines of the table are for different fitness functions and migration probabilities (0.1 or 1.0) and a line-column combination describes the number of islands of the model, where the Mann-Whitney U test was significant on a significance level of 0.001.

		Torus	Hypercube	Complete
OneMax, 0.1	Ring	$\mu > 1$	$\mu > 4$	$\mu > 4$
OneMax, 0.1	Torus	—	$\mu > 16$	$\mu > 4$
OneMax, 0.1	Hypercube	—	—	$\mu > 1$
LO, 0.1	Ring	$\mu > 4$	$\mu > 4$	$\mu > 5$
LO, 0.1	Torus	—	$\mu > 16$	$\mu > 4$
LO, 0.1	Hypercube	—	—	$\mu > 1$
OneMax, 1.0	Ring	$\mu > 4$	$\mu > 4$	$\mu > 7$
OneMax, 1.0	Torus	—	$\mu > 16$	$\mu > 9$
OneMax, 1.0	Hypercube	—	—	$\mu > 8$
LO, 1.0	Ring	$\mu > 16$	$\mu > 16$	$\mu > 8$
LO, 1.0	Torus	—	$\mu > 16$	$\mu > 16$
LO, 1.0	Hypercube	—	—	$\mu > 16$

minimizes the number of function evaluations for both OneMax and LO (Sudholt, 2013), among all EAs that only use standard bit mutation. This excludes superlinear speedups on OneMax and LO, for such EAs, from a theoretical perspective.

It can be seen that more dense topologies are more efficient than sparse topologies. In accordance with the theoretical analyses, the efficiency decreases more rapidly for OneMax. For OneMax, and $p = \Omega(1)$, only values $\mu = O(\log n)$ were guaranteed to give a linear speedup. And indeed, the efficiency in Figure 2(c) degrades quite quickly for OneMax and $p = 1.0$.

Higher numbers of islands are still efficient for LO. For the ring, the range of good μ values is up to $\mu = O(\sqrt{n})$. This is reflected in Figure 2(d) as the efficiency degrades as μ increases beyond $\sqrt{n} = 16$. For denser topologies, the efficiency only degrades for large μ . The complete graph remains effective throughout the whole scale—and even stronger, for values up to $\mu = 256$ (not shown in Figure 2) the efficiency was always above 0.75. This was also expected as $\mu = \Theta(n)$ still guarantees a linear speedup for LO.

Comparing the runtime behavior for different transmission probabilities, the plots confirm again that in our examples a higher transmission probability for individuals allows for a better overall performance.

Also, for the results in Figure 2, pairwise two-sided Mann-Whitney U tests were performed as before. For the efficiency, the mean ranks in each setup indicate the ordering of ring < torus < hypercube < K_{μ} , but with different levels of significance. In Table 2 we list those numbers of islands for which the pairwise comparisons are statistically significant at a significance level of 0.001. Note that torus and hypercube by definition only share very few data points (squares of powers of two): $\mu \in \{1, 4, 16, 64\}$. For very small values of μ , the results are not significant, as the topologies are too similar and hence show indistinguishable performance. But for larger topologies, that is, $\mu > 16$, all comparisons are indeed significant at a very low level of 0.001.

Downloaded from http://direct.mit.edu/evco/article-pdf/22/3/405/1513871/evco_a_00114.pdf by guest on 29 September 2021

9 Conclusions

We have provided a general method for the runtime analysis of parallel evolutionary algorithms, including applications to a set of well-known and illustrative example functions. Our method provides a way of automatically transforming runtime bounds obtained for panmictic EAs to parallel EAs with spatial structures. In addition to a general result, we have provided methods tailored toward specific topologies: ring graphs, torus graphs, hypercubes, and complete graphs. The latter also covers offspring populations and random migration topologies as special cases. Our results can estimate the expected parallel runtime from above, thus lower-bounding the speedup obtained through parallelization with regard to the number of generations. They also bound the expected total communication effort in terms of the total number of individuals migrated as an indicator of the bandwidth used.

The example applications revealed insights that are remarkable in their own right (see Table 1, and a more general version in Table 3 in the appendix). Compared to upper bounds obtained for a single panmictic island by the fitness-level method, for ring graphs the expected wait time for an improvement can be replaced by its square root in the parallel runtime, provided that the number of islands is large enough and improvements are transmitted efficiently, that is, $p = \Omega(1)$. This leads to a speedup on the order of $\log n$ for OneMax and on the order of \sqrt{n} for some unimodal functions, such as LO. On $\text{Jump}_{k'}$, the speedup is even on the order of at least $n^{k/2}$. A similar effect is observed for torus graphs where the expected wait time can be replaced by its cube root. The hypercube reduces the (upper bound on the) expected wait time on each level to its logarithm, and on the complete graph it is reduced to a constant, again provided there are sufficiently many islands. This way, even on functions such as LO and Jump_k [$3 \leq k = O(n/\log n)$] the expected parallel time can be reduced to $O(n)$. In all these results, the population size can be chosen in such a way that the total number of function evaluations does not increase, in an asymptotic sense. The population sizes leading to the best possible upper bounds on the parallel runtime have been stated explicitly (cf. Table 1 and also Table 3, which is in the appendix), therefore giving hints on how to parameterize parallel EAs.

The tables also reveal that in certain situations there is a tradeoff between the expected parallel time and the communication effort—at least with regard to the upper bounds shown here. For instance, on LO the torus graph has the smallest communication effort of $O(n^2)$ at the expense of a higher parallel time bound of $O(n^{4/3})$. The complete graph has the smallest bound for the parallel time, $O(n)$, but the largest bound for the communication effort: $O(n^3)$. The hypercube provides a good compromise, combining the smallest bounds up to polylogarithmic factors. A similar observation can be made for $\text{Jump}_{k'}$, but there the hypercube is a better choice than the complete graph (strictly better in terms of communication effort and equally good in the parallel time bound). In all the examples, the ring never gave better upper bounds than torus graphs in both objectives.

We also gave experimental results for the parallel (1+1) EA on OneMax and LO with the four topologies: ring, torus, hypercube, and complete graph (in order of increasing density). These results and statistical tests confirm that, also in a non-asymptotic sense, in almost all parameters settings, the following holds. Going from sparse to dense topologies, the parallel runtime of the parallel (1+1) EA decreases, but the communication effort increases.

Future work should deal with lower bounds on the runtime of parallel evolutionary algorithms to establish in which cases the upper bounds are asymptotically

tight. It might be possible to extend drift analysis methods (He and Yao, 2004; Johannsen, 2010) in such a way that both the current best fitness and the number of islands on the current best fitness level are considered; a similar strategy was used for the analysis of population-based evolutionary algorithms by Lehre and Yao (2012). Also, in the example functions, no diversity was needed. Further studies are needed in order to better understand how the topology and the parameters of migration affect diversity, and how diversity helps for optimizing more difficult, multimodal problems.

Acknowledgments

The authors were supported by postdoctoral fellowships from the German Academic Exchange Service while visiting the International Computer Science Institute, Berkeley, CA, USA. Dirk Sudholt was also supported by EPSRC grant EP/D052785/1. We would like to thank the anonymous reviewers for their comments that helped to improve the paper.

References

- Alba, E. (2002). Parallel evolutionary algorithms can achieve super-linear performance. *Information Processing Letters*, 82(1):7–13.
- Alba, E. (2005). *Parallel metaheuristics: A new class of algorithms*. New York: Wiley-Interscience.
- Alba, E., and Luque, G. (2004). Growth curves and takeover time in distributed evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, Lecture notes in computer science*, vol. 3102, pp. 864–876. Berlin: Springer-Verlag.
- Auger, A., and Doerr, B. (Eds.). (2011). *Theory of randomized search heuristics; Foundations and recent developments*. Singapore: World Scientific.
- Cantú Paz, E. (1997). A survey of parallel genetic algorithms. Technical report, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana Champaign, Urbana, IL.
- Doerr, B., Happ, E., and Klein, C. (2007). A tight analysis of the (1+1)-EA for the single source shortest path problem. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC '07*, pp. 1890–1895.
- Doerr, B., Happ, E., and Klein, C. (2011). Tight analysis of the (1+1)-EA for the single source shortest path problem. *Evolutionary Computation*, 19(4):673–691.
- Droste, S., Jansen, T., and Wegener, I. (2002). On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81.
- Friedrich, T., He, J., Hebbinghaus, N., Neumann, F., and Witt, C. (2010). Approximating covering problems by randomized search heuristics using multi-objective models. *Evolutionary Computation*, 18(4):617–633.
- Giacobini, M., Alba, E., Tettamanzi, A., and Tomassini, M. (2005). Selection intensity in cellular evolutionary algorithms for regular lattices. *IEEE Transactions on Evolutionary Computation*, 9:489–505.
- Giacobini, M., Alba, E., and Tomassini, M. (2003). Selection intensity in asynchronous cellular evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '03*, pp. 955–966.
- Giacobini, M., Tomassini, M., and Tettamanzi, A. (2003). Modelling selection intensity for linear cellular evolutionary algorithms. In *Proceedings of the Sixth International Conference on Artificial Evolution, Evolution Artificielle*, pp. 345–356.

- Giacobini, M., Tomassini, M., and Tettamanzi, A. (2005). Takeover time curves in random and small-world structured populations. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '05*, pp. 1333–1340.
- Gutjahr, W. J., and Sebastiani, G. (2008). Runtime analysis of ant colony optimization with best-so-far reinforcement. *Methodology and Computing in Applied Probability*, 10:409–433.
- He, J., and Yao, X. (2003). Towards an analytic framework for analyzing the computation time of evolutionary algorithms. *Artificial Intelligence*, 145(1–2):59–97.
- He, J., and Yao, X. (2004). A study of drift analysis for estimating computation time of evolutionary algorithms. *Natural Computing*, 3(1):21–35.
- Horoba, C. (2010). Exploring the runtime of an evolutionary algorithm for the multi-objective shortest path problem. *Evolutionary Computation*, 18(3):357–381.
- Jansen, T. (2013). *Analyzing evolutionary algorithms; The computer science perspective*. Berlin: Springer.
- Jansen, T., and Wegener, I. (2002). On the analysis of evolutionary algorithms; A proof that crossover really can help. *Algorithmica*, 34(1):47–66.
- Jensen, J. (1906). Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30:175–193.
- Johannsen, D. (2010). Random combinatorial structures and randomized search heuristics. Ph.D. thesis, Universität des Saarlandes, Saarbrücken, Germany; and the Max-Planck-Institut für Informatik.
- Kötzing, T., Sudholt, D., and Theile, M. (2011). How crossover helps in pseudo-Boolean optimization. In *Proceedings of the 13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011*, pp. 989–996.
- Lässig, J., and Sudholt, D. (2010a). The benefit of migration in parallel evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2010*, pp. 1105–1112.
- Lässig, J., and Sudholt, D. (2010b). General scheme for analyzing running times of parallel evolutionary algorithms. In *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature, PPSN 2010. Lecture notes in computer science*, vol. 6238, pp. 234–243. Berlin: Springer-Verlag.
- Lässig, J., and Sudholt, D. (2011a). Adaptive population models for offspring populations and parallel evolutionary algorithms. In *Proceedings of the 11th Workshop on Foundations of Genetic Algorithms, FOGA 2011*, pp. 181–192.
- Lässig, J., and Sudholt, D. (2011b). Analysis of speedups in parallel evolutionary algorithms for combinatorial optimization. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC 2011). Lecture notes in computer science*, vol. 7074, pp. 405–414. Berlin: Springer-Verlag.
- Lehre, P. K. (2011). Fitness-levels for non-elitist populations. In *Proceedings of the 13th Annual Genetic and Evolutionary Computation Conference, GECCO '11*, pp. 2075–2082.
- Lehre, P. K., and Yao, X. (2012). On the impact of mutation-selection balance on the runtime of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 16(2):225–241.
- Luque, G., and Alba, E. (2011). *Parallel genetic algorithms; Theory and real world applications. Studies in computational intelligence*, vol. 367, Berlin: Springer-Verlag.
- Mambrini, A., Sudholt, D., and Yao, X. (2012). Homogeneous and heterogeneous island models for the set cover problem. In *Parallel Problem Solving from Nature, PPSN 2012. Lecture notes in computer science*, vol. 7491, pp. 11–20. Berlin: Springer-Verlag.

- Nedjah, N., de Macedo Mourelle, L., and Alba, E. (2006). *Parallel Evolutionary Computations*. Berlin: Springer-Verlag.
- Neumann, F., Oliveto, P. S., Rudolph, G., and Sudholt, D. (2011). On the effectiveness of crossover for migration in parallel evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2011*, pp. 1587–1594.
- Neumann, F., Sudholt, D., and Witt, C. (2009). Analysis of different MMAS ACO algorithms on unimodal functions and plateaus. *Swarm Intelligence*, 3(1):35–68.
- Neumann, F., and Witt, C. (2010). *Bioinspired computation in combinatorial optimization: Algorithms and their computational complexity*. Berlin: Springer-Verlag.
- Oliveto, P. S., He, J., and Yao, X. (2007). Time complexity of evolutionary algorithms for combinatorial optimization: A decade of results. *International Journal of Automation and Computing*, 4(3):281–293.
- Rowe, J., Mitavskiy, B., and Cannings, C. (2008). Propagation time in stochastic communication networks. In *Proceedings of the Second IEEE International Conference on Digital Ecosystems and Technologies*, pp. 426–431.
- Rowe, J., and Sudholt, D. (2012). The choice of the offspring population size in the $(1, \lambda)$ EA. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2012*, pp. 1349–1356.
- Rudolph, G. (2000). On takeover times in spatially structured populations: Array and ring. In *Proceedings of the 2nd Asia-Pacific Conference on Genetic Algorithms and Applications*, pp. 144–151.
- Rudolph, G. (2006). Takeover time in parallel populations with migration. In *Proceedings of BIOMA 2006*, pp. 63–72.
- Sarma, J., and De Jong, K. (1997). An analysis of local selection algorithms in a spatially structured evolutionary algorithm. In *Proceedings of the 7th International Conference on Genetic Algorithms*, pp. 181–186.
- Skolicki, Z., and De Jong, K. (2005). The influence of migration sizes and intervals on island models. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2005*, pp. 1295–1302. ACM.
- Sudholt, D. (to appear). Parallel evolutionary algorithms. In *Handbook of computational intelligence*. Berlin: Springer-Verlag.
- Sudholt, D. (2013). A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 17(3):418–435.
- Sudholt, D., and Thyssen, C. (2012). Running time analysis of ant colony optimization for shortest path problems. *Journal of Discrete Algorithms*, 10:165–180.
- Sudholt, D., and Witt, C. (2010). Runtime analysis of a binary particle swarm optimizer. *Theoretical Computer Science*, 411(21):2084–2100.
- Tomassini, M. (2005). *Spatially structured evolutionary algorithms: Artificial evolution in space and time*. Berlin: Springer-Verlag.
- Wegener, I. (2002). Methods for the analysis of evolutionary algorithms on pseudo-Boolean functions. In R. Sarker, X. Yao, and M. Mohammadian (Eds.), *Evolutionary optimization*, pp. 349–369. Dordrecht, The Netherlands: Kluwer.
- Witt, C. (2005). Worst-case and average-case approximations by simple randomized search heuristics. In *Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science, STACS '05. Lecture notes in computer science*, vol. 3404, pp. 44–56. Berlin: Springer-Verlag.
- Witt, C. (2006). Runtime analysis of the $(\mu+1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation*, 14(1):65–86.

Yu, Y., Yao, X., and Zhou, Z.-H. (2012). On the approximation ability of evolutionary optimization with application to minimum set cover. *Artificial Intelligence*, 180–181:20–33.

Zhou, D., Luo, D., Lu, R., and Han, Z. (2012). The use of tail inequalities on the probable computational time of randomized search heuristics. *Theoretical Computer Science*, 436:106–117.

Appendix

The following inequality was brought to our attention by Jon Rowe. A proof is found in Rowe and Sudholt (2012, Lemma 3).

LEMMA 19: For any $0 \leq x \leq 1$, and any $n > 0$

$$(1 - x)^n \leq \frac{1}{1 + nx}.$$

We also state Johanssen’s variable drift theorem (Johanssen, 2010), in a version with slightly improved conditions (Rowe and Sudholt, 2012).

THEOREM 20 (JOHANNSEN’S VARIABLE DRIFT THEOREM) (Johanssen, 2010; Rowe and Sudholt, 2012): Consider a stochastic process $\{X_t\}_{t \geq 0}$ on $\{0, 1, \dots, m\}$, with $m \in \mathbb{N}$. Suppose there is a monotonic increasing function $h : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that the function $1/h(x)$ is integrable on $\{1, \dots, m\}$, and with

$$E(X_t - X_{t+1} \mid X_t = k) \geq h(k)$$

for all $k \in \{1, \dots, m\}$. Then the expected first hitting time of state 0 is at most

$$\frac{1}{h(1)} + \int_1^m \frac{1}{h(x)} dx.$$

Table 3: Asymptotic bounds on expected parallel (T^{par} , number of generations) and sequential (T^{seq} , number of function evaluations) runtimes and expected communication efforts (T^{com} , total number of migrated individuals) for various n -bit functions and island models with μ islands running the (1+1) EA and using migration probability p . The number of islands μ was always chosen to give the best possible upper bound on the parallel running time, while not increasing the upper bound on the sequential running time by more than a constant factor. For unimodal functions $d + 1$ denotes the number of function values. See Droste et al. (2002) for bounds for the (1+1) EA. Results for Jump_k were restricted to $3 \leq k = O(n/\log n)$ for simplicity.

	(1+1) EA	Ring	Grid/Torus	Hypercube	Complete/ K_μ	K_μ with $p = O(1/\mu)$
OneMax						
$E(T^{\text{par}})$	$\Theta(n \log n)$	$O\left(\frac{n}{p^{1/2}} + \frac{n \log n}{\mu}\right)$	$O\left(\frac{n}{p^{2/3}} + \frac{n \log n}{\mu}\right)$	$O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right)$	$O\left(\frac{n}{p} + \frac{n \log n}{\mu}\right)$	$O\left(\frac{n \log \mu}{p\mu} + \frac{n \log n}{\mu}\right)$
$E(T^{\text{seq}})$	$\Theta(n \log n)$	$O\left(\frac{\mu n}{p^{1/2}} + n \log n\right)$	$O\left(\frac{\mu n}{p^{2/3}} + n \log n\right)$	$O\left(\frac{\mu n}{p} + n \log n\right)$	$O\left(\frac{\mu n}{p} + n \log n\right)$	$O\left(\frac{\mu(\log \mu)^n}{p\mu} + n \log n\right)$
$E(T^{\text{com}})$	0	$O(p^{1/2} \mu n + pn \log n)$	$O(p^{1/3} \mu n + pn \log n)$	$O(\mu n(\log \mu) + p(\log \mu)n \log n)$	$O(\mu^2 n + p\mu n \log n)$	$O(\mu(\log \mu)n + p\mu n \log n)$
LO						
$E(T^{\text{par}})$	$\Theta(n^2)$	$O\left(\frac{n^{3/2}}{p^{1/2}} + \frac{n^2}{\mu}\right)$	$O\left(\frac{n^{4/3}}{p^{2/3}} + \frac{n^2}{\mu}\right)$	$O\left(\frac{n \log n}{p} + \frac{n^2}{\mu}\right)$	$O\left(\frac{n}{p} + \frac{n^2}{\mu}\right)$	$O\left(\frac{n \log \mu}{p\mu} + \frac{n^2}{\mu}\right)$
$E(T^{\text{seq}})$	$\Theta(n^2)$	$O\left(\frac{\mu n^{3/2}}{p^{1/2}} + n^2\right)$	$O\left(\frac{\mu n^{4/3}}{p^{2/3}} + n^2\right)$	$O\left(\frac{\mu n \log n}{p} + n^2\right)$	$O\left(\frac{\mu n}{p} + n^2\right)$	$O\left(\frac{\mu(\log \mu)^n}{p\mu} + n^2\right)$
$E(T^{\text{com}})$	0	$O(p^{1/2} \mu n^{3/2} + pn^2)$	$O(p^{1/3} \mu n^{4/3} + pn^2)$	$O(\mu n(\log n)(\log \mu) + p(\log \mu)n^2)$	$O(\mu^2 n + p\mu n^2)$	$O(\mu(\log \mu)n + p\mu n^2)$
unimodal						
$E(T^{\text{par}})$	$O(dn)$	$O\left(\frac{dn^{1/2}}{p^{1/2}} + \frac{dn}{\mu}\right)$	$O\left(\frac{dn^{1/3}}{p^{2/3}} + \frac{dn}{\mu}\right)$	$O\left(\frac{d \log n}{p} + \frac{dn}{\mu}\right)$	$O\left(\frac{d}{p} + \frac{dn}{\mu}\right)$	$O\left(\frac{d \log \mu}{p\mu} + \frac{dn}{\mu}\right)$
$E(T^{\text{seq}})$	$O(dn)$	$O\left(\frac{d \mu n^{1/2}}{p^{1/2}} + dn\right)$	$O\left(\frac{d \mu n^{1/3}}{p^{2/3}} + dn\right)$	$O\left(\frac{d \mu \log n}{p} + dn\right)$	$O\left(\frac{d \mu}{p} + dn\right)$	$O\left(\frac{d \log \mu}{p\mu} + dn\right)$
$E(T^{\text{com}})$	0	$O(p^{1/2} d \mu n^{1/2} + pdn)$	$O(p^{1/3} d \mu n^{1/3} + pdn)$	$O(d \mu(\log n)(\log \mu) + p(\log \mu)dn)$	$O(d \mu^2 + pd \mu n)$	$O(d \mu \log \mu + p \mu dn)$
Jump_k						
$E(T^{\text{par}})$	$\Theta(n^k)$	$O\left(\frac{n^{k/2}}{p^{1/2}} + \frac{n^k}{\mu}\right)$	$O\left(\frac{n^{k+1/3}}{p^{2/3}} + \frac{n^k}{\mu}\right)$	$O\left(\frac{n}{p} + \frac{n^k}{\mu}\right)$	$O\left(\frac{n}{p} + \frac{n^k}{\mu}\right)$	$O\left(\frac{n \log \mu}{p\mu} + \frac{n^k}{\mu}\right)$
$E(T^{\text{seq}})$	$\Theta(n^k)$	$O\left(\frac{\mu n^{k/2}}{p^{1/2}} + n^k\right)$	$O\left(\frac{\mu(n+n^{k/3})}{p^{2/3}} + n^k\right)$	$O\left(\frac{\mu n}{p} + n^k\right)$	$O\left(\frac{\mu n}{p} + n^k\right)$	$O\left(\frac{n \log \mu}{p} + n^k\right)$
$E(T^{\text{com}})$	0	$O(p^{1/2} \mu n^{k/2} + pn^k)$	$O(p^{1/3} \mu n^{k/3} + pn^k)$	$O(\mu n(\log \mu) + p(\log \mu)n^k)$	$O(\mu^2 n + p \mu n^k)$	$O(\mu(\log \mu)n + p \mu n^k)$

