
Analyses of Simple Hybrid Algorithms for the Vertex Cover Problem*

Tobias Friedrich tobias.friedrich@mpi-inf.mpg.de
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Jun He jqh@aber.ac.uk
Department of Computer Science, University of Wales, Aberystwyth, United Kingdom

Nils Hebbinghaus nils.hebbinghaus@mpi-inf.mpg.de
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Frank Neumann frank.neumann@mpi-inf.mpg.de
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Carsten Witt carsten.witt@cs.uni-dortmund.de
Fakultät für Informatik, LS 2, Technische Universität Dortmund, Dortmund, Germany

Abstract

Hybrid methods are very popular for solving problems from combinatorial optimization. In contrast, the theoretical understanding of the interplay of different optimization methods is rare. In this paper, we make a first step into the rigorous analysis of such combinations for combinatorial optimization problems. The subject of our analyses is the vertex cover problem for which several approximation algorithms have been proposed. We point out specific instances where solutions can (or cannot) be improved by the search process of a simple evolutionary algorithm in expected polynomial time.

Keywords

Hybrid algorithms, evolutionary computation, combinatorial optimization, running time analysis.

1 Introduction

Evolutionary Algorithms (EAs) have been widely applied to optimization problems (Goldberg, 1989; Holland, 1992; Michalewicz, 1996). Especially for combinatorial optimization they have been shown to produce good solutions. To be competitive with other methods such as approximation algorithms (Vazirani, 2001), Tabu Search (Glover and Laguna, 1997), Branch and Bound (Dakin, 1965), or Linear Programming (Schrijver, 1998), often a “pure” EA is not used but some knowledge is incorporated into the algorithm.

In contrast to the assumption that using hybridization in evolutionary computation often yields better results than a more general approach with problem-specific knowledge, there are up to now no theoretical investigations with respect to the runtime

*A preliminary version of this article appeared in *Proceedings of the 14th IEEE Congress on Evolutionary Computation (CEC 2007)*, pp. 2614–2621, Piscataway, NJ: IEEE Press, 2007.

behavior of such algorithms. A lot of theoretical results giving bounds on the runtime of simple evolutionary algorithms for combinatorial optimization problems have been obtained in recent years. Among these problems, there are some of the most popular polynomially solvable problems such as sorting and shortest path (Scharnow et al., 2004), the computation of a maximum matching (Giel and Wegener, 2003, 2006), the Eulerian cycle problem (Doerr et al., 2006; Neumann, 2004), and the computation of minimum spanning trees (Neumann and Wegener, 2004, 2005). In the case of NP-hard problems, results have been obtained for a scheduling problem on two machines (Witt, 2005) and the multi-objective minimum spanning tree problem (Neumann, 2007). Especially in the case of well-known NP-hard combinatorial optimization problems, often the initial solution is computed by an approximation algorithm and later on improved by a heuristic method. None of previous works on the runtime analysis of combinatorial optimization problems has considered this scenario. All investigations considered the case where the initial solution is drawn uniformly at random from the considered search space. Our aim is to make a first step to understand in which situations evolutionary algorithms are able to improve solutions obtained by specialized approximation algorithms. We do not aim to show general results on improving the approximation ratio of such a problem as this would be a statement that has to be made about all possible instances. We instead focus on describing different situations where the random search process of an evolutionary might (or might not) be useful.

The subject of our investigations is the vertex cover problem. Various EAs have been applied to this problem (Evans, 1998; Khuri and Bäck, 1994). Kehden and Neumann (2006) showed in an experimental study that in the case of random graphs, solutions obtained by a well-known approximation algorithm are often far from optimal and may be improved in a small number of iterations by evolutionary algorithms so that nearly optimal solutions are obtained. The first rigorous runtime analysis on classes of instances for this problem is given in He et al. (2005). Recently, it has been shown that the well-known (1+1) EA is not able to produce a good approximation for the vertex cover problem (Friedrich et al., 2007). In this work, it has been proven that even in the case of bipartite graphs the approximation ratio achievable by this algorithm in expected polynomial time is almost as bad as the trivial cover made up of all vertices of the given graph.

Such a bad approximation quality can be prevented by starting with an initial solution that has been obtained by running a good approximation algorithm. In the case of the vertex cover problem several approximation algorithms are known. The first idea for such an algorithm is to start with the empty vertex set and in each iteration add a vertex that covers the largest number of uncovered edges. It is well known that this approach achieves an approximation ratio of $\Theta(\log n)$, where n is the number of vertices in the graph. Some simple approximation algorithms obtain an approximation ratio of 2 which is asymptotically the best known up to now. Such an approximation quality can, for example, be obtained by computing a maximal matching of the given graph and including for each edge of this matching both endpoints into the cover.

We investigate for which cases the solutions obtained by the two described approximation algorithms can be improved by an evolutionary algorithm. We investigate the (1+1) EA starting with such a solution and point out situations where such a search process does (or does not) have the ability to improve a solution constructed by the approximation algorithms.

The outline of the paper is as follows. In Section 2 we introduce the vertex cover problem and the EA that is subject to our investigations. Section 3 considers

improvements achievable by the EA for the greedy approximation algorithm and Section 4 investigates the combination with the maximal matching approach. Finally, we finish with some conclusions.

2 The Vertex Cover Problem and the (1+1) EA

The vertex cover problem is one of the most studied NP-hard combinatorial optimization problems. Given an undirected graph $G = (V, E)$ where $|V| = n$ and $|E| = m$, the aim is to find a subset $V' \subseteq V$ of minimum cardinality such that for each $e \in E$, $e \cap V' \neq \emptyset$ holds.

We consider the well-known (1+1) EA (see Algorithm 1) for the vertex cover problem. The search space is $\{0, 1\}^n$ and each bit x_i of a solution x corresponds to a vertex $v_i \in V$. The vertex v_i is chosen in the current solution x if $x_i = 1$ and otherwise it is unchosen.

Algorithm 1: (1+1) EA

1. Choose an initial solution $x \in \{0, 1\}^n$.
2. Repeat
 - Create x' by flipping each bit of x with probability $1/n$.
 - If $f(x') \leq f(x)$, set $x := x'$.

Denote by $|x|_1$ and $|x|_0$ the number of ones respectively of zeros in a bitstring x . The fitness of a search point x is given by $f(x) = (u(x), |x|_1)$ where $u(x)$ denotes the number of uncovered edges of the solution x . In the case of the (1+1) EA, the function should be minimized with respect to the lexicographic order. This setting has already been examined in Friedrich et al. (2007) for randomly chosen initial solutions. We examine the effect of using an initial solution that has been computed by some approximation algorithm. Here in any case all edges are covered and the (1+1) EA does not accept solutions that do not constitute a vertex cover.

Our aim is to analyze the (1+1) EA by a rigorous runtime analysis until it has produced good solutions for the vertex cover problem. The measure of interest is the number of constructed solutions until certain goals have been achieved. Often the expectation of this value is considered and called the expected time to achieve such a goal.

Most of our investigations consider the approximation ability of the proposed algorithms. The worst-case approximation ratio of an algorithm A for a given minimization problem R is defined as $\max_{I \in R} \frac{A(I)}{\text{OPT}(I)}$ where $A(I)$ denotes the value obtained by A when applied to an instance I of R and $\text{OPT}(I)$ denotes the value of an optimal solution for the given instance. We are mainly interested in upper and lower bounds for the number of constructed solutions until a certain approximation ratio has been achieved by the introduced algorithms.

As the (1+1) EA does not accept worsenings, the approximation ratio achieved is at least as good as the approximation ratio of the algorithm to compute the initial solution.

3 Analysis of Hybrid EAs with the Greedy Method

In this section we examine how an initial solution produced by a greedy algorithm can be improved by the (1+1) EA. We show that the success of such a hybrid EA depends on




Figure 1: An illustration of Graph 1 for $k = 4$. The gray straps should indicate complete bipartite subgraphs. For better clarity, the edges incident to W_2 are omitted.

the specifics of the examined graphs. We introduce two graph classes and show that the hybrid EA may fail on the first one to find a solution with approximation ratio $o(\log n)$ in polynomial time and finds the optimum quickly for the second class of graphs.

The greedy method is based on the following idea. In the vertex covering problem, the aim is to find a vertex cover that uses the minimum number of vertices to cover all edges. Therefore, a vertex with a larger degree is more likely to appear in the optimal vertex cover. However, this is only a heuristic knowledge and it will not lead to an optimal solution in general.

We consider the greedy method described in Papadimitriou and Steiglitz (1998) to compute the initial solution x . In each step of the algorithm, a vertex with the largest degree is added to the solution and the vertex and the edges incident to it are removed from the graph. The process is iterated until the graph becomes empty. We can state the algorithm as follows (see Figure 1).

Algorithm 2: *Greedy Vertex Cover*

1. Set $x = 0^n$
2. Repeat
 - Choose a vertex v_k having the largest degree in G .
 - Set $x_k = 1$, $V := V \setminus \{v_k\}$, and $E := E \setminus \{e \mid e \cap v_k \neq \emptyset\}$
3. Until G is empty

The greedy approach achieves an approximation ratio of $O(\log n)$. We obtain a simple hybrid algorithm by computing the initial solution of the (1+1) EA using

Algorithm 2. The resulting algorithm we will denote by *Greedy Vertex Cover (1+1) EA* ((1+1) EA_G). It achieves the same approximation ratio $O(\log n)$ as the greedy approach as it does not accept solutions that are worse than the initial one.

Friedrich et al. (2007) have shown that the (1+1) EA starting with a solution that is chosen uniformly at random from the search space is not able to obtain a good approximation for a specific class of bipartite graphs in expected polynomial time. They have also shown that a greedy approach in the form of a multi-objective EA is able to produce an optimal solution for such problems quickly. The same holds for Algorithm 2. Hence, using the greedy procedure for the initial solution can make the difference between obtaining an optimal solution or achieving only a bad approximation of such a solution.

Compared with other approximation algorithms that achieve an approximation ratio of 2 for the vertex cover problem, the greedy approach behaves badly in the worst case. Therefore, the question arises whether a solution that is far from optimal can always be improved by the (1+1) EA. The following example shows that this is not always the case. The solution for Graph 1 computed by the greedy approach might be far from optimal and the (1+1) EA_G is likely to achieve not even a single improvement.

Graph 1: Let the vertex set be $W_1 \cup W_2 \cup \bigcup_{i=\lceil k/2 \rceil}^k V_i$ with

$$\begin{aligned} W_\ell &:= \{w_{\ell,j} \mid j = 1, \dots, 2^k\} \quad \text{for all } \ell \in \{1, 2\} \\ V_i &:= \{v_{i,j} \mid j = 1, \dots, 2^k\} \quad \text{for all } \lceil k/2 \rceil \leq i \leq k \end{aligned}$$

and the edge set be

$$\begin{aligned} E &:= \{\{v_{i,j_1}, w_{\ell,j_2}\} \mid \lceil k/2 \rceil \leq i \leq k, \ell \in \{1, 2\}, \\ &\quad 1 \leq j_1, j_2 \leq 2^k, \lceil j_1/2^i \rceil = \lceil j_2/2^i \rceil\}. \end{aligned}$$

The number of vertices is $n = (\lfloor k/2 \rfloor + 3)2^k$ and the vertices in set W_1 and W_2 have degree $\sum_{i=\lceil k/2 \rceil}^k 2^i = 2^{k+1} - 2^{\lceil k/2 \rceil}$ while the vertices in V_i have degree 2^{i+1} . The optimal cover is

$$C^* := W_1 \cup W_2.$$

In the following, we show that Algorithm 2 produces a factor $\Omega(\log n)$ approximation of an optimal solution and that such a solution is hard to improve by the search procedure of the (1+1) EA.

THEOREM 1: On Graph 1 the expected time for the (1+1) EA_G to obtain a solution whose approximation ratio is $o(\log n)$ is $2^{\Omega(\sqrt{n \log n})}$.

PROOF: We first show that the solution produced by Algorithm 2 is only a factor $\Omega(\log n)$ approximation and lower bound the expected time to improve such a solution afterward.

As the vertices in sets V_k have the largest degree, the greedy Algorithm 2 first chooses all vertices from V_k . After removing these nodes, the vertex degrees in W_1 and W_2 decrease to $\sum_{i=\lceil k/2 \rceil}^{k-1} 2^i = 2^k - 2^{\lceil k/2 \rceil}$ while the vertices in V_i still have degree 2^i for




Figure 2: An illustration of Graph 2 for $k = 6$.

$i < k$. Hence, the greedy method chooses the vertices of V_{k-1} to be added next to the cover. Iterating this procedure, the greedy method obtains the cover

$$C = V_k \cup V_{k-1} \cup \dots \cup V_{\lfloor k/2 \rfloor}$$

and the approximation ratio

$$\frac{|C|}{|C^*|} = \frac{(\lfloor k/2 \rfloor + 1)2^k}{2^k} = \Theta(k) = \Theta(\log n).$$

Hence, Graph 1 is, up to a constant factor, a worst case example for Algorithm 2 with respect to approximability.

To gain an improvement if C is the current solution, a certain number of vertices of $W_1 \cup W_2$ have to be added while a larger number of vertices of $\bigcup_i V_i$ have to be removed. As the minimum vertex degree in Graph 1 is $\delta = 2^{\lfloor k/2 \rfloor + 1}$, at least 2δ arbitrary bits have to be flipped to gain an improvement. It remains to bound the probability that the $(1+1)EA_G$ achieves such an improvement in one mutation step. For this, we use Stirling's formula, the identity $n = 2^{\log n}$, and (for sufficiently large n) $\delta \geq \sqrt{\frac{n}{\log n}}$ and $(\frac{2}{e}\sqrt{\frac{n}{\log n}})^3 \geq n$. Hence, the probability of an improvement can be upper bounded by

$$\frac{\binom{n}{2\delta}}{n^{2\delta}} \leq \frac{1}{(2\delta)!} \leq \left(\frac{2\delta}{e}\right)^{-2\delta} \leq n^{-\frac{2}{3}\sqrt{\frac{n}{\log n}}} = 2^{-\Omega(\sqrt{n \log n})}.$$

which proves the theorem. ■

Theorem 1 shows that the $(1+1)EA_G$ may be unable to improve a bad greedy solution. We will now show that there are also graph classes on which the above hybrid EA can play a positive role. On the following (bipartite) Graph 2, the $(1+1)EA$ finds the optimal vertex cover in expected polynomial time.

Consider the following well-known graph given in Papadimitriou and Steiglitz (1998); see Figure 2.

Graph 2: Let $k \in \mathbb{N}$. Let the vertex set be $B \cup \bigcup_{i=1}^k A_i$ with

$$B := \{b_j \mid j = 1, \dots, k\},$$

$$A_i := \{a_{i,j} \mid j = 1, \dots, \lfloor k/i \rfloor\} \text{ for all } 1 \leq i \leq k$$

and the edge set be

$$E := \{\{b_j, a_{i, \lceil j/i \rceil}\} \mid 1 \leq i \leq k, 1 \leq j \leq i \lfloor k/i \rfloor\}.$$

Graph 2 has the following properties. Since every vertex in A_i is connected with exactly i vertices in B and no two vertices in B have a common neighbor in any A_i , all vertices in B have degree at most k and all vertices in A_i have degree i . Let us denote the total number of vertices by n . Then $n = \Theta(k \log k)$. Due to Papadimitriou and Steiglitz (1998), the greedy algorithm (Algorithm 2) can determine the set $\bigcup_{i=1}^k A_i$ as a vertex cover. This set is a factor of $\Omega(\log n)$ away from the optimal vertex cover B . In contrast to the behavior of the (1+1) EA_G on Graph 1, the (1+1) EA_G determines the optimal vertex cover in expected time $O(n^3)$. We formulate this result in a more general way in the following theorem.

THEOREM 2: *The expected optimization time of the (1+1) on Graph 2 is $O(n^3)$ for each initial search point.*

PROOF: In Friedrich et al. (2007) it is shown that the (1+1) EA produces a vertex cover in expected time $O(n \log n)$ regardless of the chosen initial solution. Therefore, we now assume that a vertex cover has been obtained and upper bound the expected time to produce an optimal solution that consists of the vertex set B . If the current cover is not a minimal one at least one vertex can be removed to achieve an improvement. The corresponding probability for such an event is $\Omega(1/n)$ and its expected waiting time is $O(n)$.

If the current solution is minimal but not a minimum vertex cover at least one vertex v_1 in A_2 has been chosen. Otherwise, all vertices of B must be selected (to cover all edges between A_2 and B) and no other vertex can be chosen due to the minimality of the vertex cover. But this is the minimum vertex cover. The chosen vertex v_1 of A_2 has two neighbors v_2 and v_3 in B . In addition v_2 is adjacent to a vertex v_4 in A_1 and v_3 is adjacent to a vertex v_5 in A_1 . To reduce the number of vertices in the case that the current solution is minimal but not a minimum vertex cover, we consider special two-bit flips concerning the vertices of $\{v_2, v_3, v_4, v_5\}$. As the current solution is a cover that is minimal exactly two vertices of $\{v_2, v_3, v_4, v_5\}$ have been chosen and at least one of the chosen vertices belongs to A_1 . Deleting a chosen vertex of A_1 and including its neighbor of B is always accepted and can be done by a specific two-bit flip. Similarly, deleting a vertex of B and including its neighbor of A_1 can be done by a specific two-bit flip. However, the acceptance of such a step is determined by the choice of the other vertices of $V \setminus (B \cup A_1)$, for example, removing the vertex of B is not possible if at least one of its neighbors in $V \setminus (B \cup A_1)$ is unchosen. This shows that the probability of switching from a chosen vertex of A_1 to its neighbors in B is at least as high as the probability of switching from a chosen vertex of B to its neighbor of A_1 .

Only steps affecting the vertices in $\{v_2, v_3, v_4, v_5\}$ are relevant for us. If both vertices of B have not been chosen, the probability of increasing the number of vertices in B in such a step is at least $1/2$ the probability of the case where at least one of the two vertices of B is missing. Two such steps consecutively are sufficient to have both vertices of B chosen. Each such step chooses a missing B -vertex in favor of an A_1 -vertex independently of earlier steps. Therefore, the probability that two consecutive such steps both increase the number of B -vertices is at least $1/4$. The probability of carrying out a specific two-bit flip mutation step is at least $1/(en^2)$. Hence, after an expected

number of $O(n^2)$ steps, both vertices of B are chosen. Now the solution is not minimal as the vertex v_1 can be removed. The probability that the vertex v_1 is removed before v_4 (respectively v_5) are replaced by v_2 (respectively v_3) is $1 - O(\frac{1}{n})$.

Thus, the expected waiting time for such an improvement is $O(n^2)$ in any case and the number of vertices is $O(n)$. This implies that a minimum vertex cover is obtained after an expected number of $O(n^3)$ steps. ■

4 Analysis of Hybrid EAs with Two-Approximation Heuristics

It has been known for a long time (see e. g., Papadimitriou and Steiglitz, 1998) that the vertex cover problem admits a two-approximation using so-called *maximal matchings*. A matching is a subset of pairwise disjoint edges of a given undirected graph. Thus, the empty set is always a matching. A matching is called maximal if there is no edge left that can be added to the subset without violating the matching property. Therefore the maximal matching problem can be solved in linear time w. r. t. the number of edges by greedily adding edges to the current matching until the matching is maximal.

Note that a maximal matching is not necessarily a matching of maximum cardinality. From an algorithmic point of view, the latter problem—called the *maximum matching* problem—is more complicated yet can also be solved in polynomial time (Papadimitriou and Steiglitz, 1998). It is interesting that almost-maximum matchings can be found in polynomial time using pure evolutionary algorithms (Giel and Wegener, 2003).

The proposed two-approximation algorithm works as follows:

Algorithm 3: *Maximal Matching Vertex Cover*

1. Compute a maximal matching.
2. Output all endpoints of the edges of the matching.

The set of vertices we obtain is a vertex cover since otherwise, the matching would not have been maximal. Moreover, the set forms a two-approximation since at least one endpoint of each edge in the matching must be chosen for a minimum vertex cover. There are examples, for example, the bipartite graph consisting of $n/2$ vertex-disjoint edges, where the approximation ratio is tight since the matching chooses twice as many vertices as necessary. It is therefore interesting to study how such solutions can be improved by other heuristics.

We obtain a simple hybrid algorithm by computing the initial solution of the (1+1) EA using Algorithm 3. This is called the *Maximal Matching (1+1) EA* ((1+1) EA_M). Note that in fact, a family of algorithms is considered since the component computing the maximal matching has not been specified yet.

In the following, we study the interplay of the maximal matching component and the (1+1) EA in our hybrid algorithm. It is easy to see that for the above-mentioned graph on $n/2$ vertex-disjoint edges, the (1+1) EA_M finds a minimum vertex cover in expected polynomial time. Due to the NP-hardness of the problem, we however cannot expect the (1+1) EA_M to always find optimal solutions in expected polynomial time. In fact, we do believe that there are instances where it fails to find in polynomial time solutions that are significantly better than two-approximate. The best known polynomial-time approximation algorithm for vertex cover has an approximation ratio of $2 - \Theta(1/\sqrt{\log n})$ (Karakostas, 2005), so we cannot hope to show better results for our




Figure 3: An illustration of Graph 3 for $k = 4$.

simple hybrid algorithm. However, a simple observation proves that the evolutionary component helps to improve solutions that are exactly two-approximate.

PROPOSITION 1: *With probability $\Omega(1)$, the (1+1) EA_M produces in expected polynomial time a vertex cover with an approximation ratio of $2 - 2/n$.*

PROOF: Suppose that the optimal vertex cover has size k . If the solution of the maximal matching component is not $(2 - 2/n)$ -approximate, $k \leq n/2$ must hold and the solution must be of cardinality greater than $2k - 2k/n$, hence of cardinality $2k$. This means that an optimal cover consists of exactly one endpoint of each edge in the matching. With probability $\Omega(1)$, the first step that flips one of the $2k$ bits of the solution flips only one of the bits that do not belong to the optimal vertex cover. The new solution of $2k - 1$ vertices is a $(2 - 2/n)$ -approximation. ■

The following considerations show that the choice of the maximal matching can be crucial for the performance of the hybrid algorithm and that the random search of the (1+1) EA can still result in very different solutions for a given initial matching. We describe two examples where specific matchings can lead the (1+1) EA into local optima of bad quality. The first one is composed of bipartite subgraphs on vertex sets of size $k = n/5$. These bipartite subgraphs are connected in a chain-like manner such that the whole graph still is bipartite. Figure 3 depicts the graph for $n = 20$.

Graph 3: *Let the vertex set be the union of*

$$V_i = \{v_{i,1}, \dots, v_{i,k}\}$$

for $1 \leq i \leq 5$ and the edge set be the union of the sets

$$\{\{v_{i,r}, v_{i+1,s}\} \mid 1 \leq r, s \leq k\}$$

for $1 \leq i \leq 4$.

An optimal vertex cover for Graph 3 is obtained by choosing the sets V_2 and V_4 . A suboptimal vertex cover of approximation ratio $3/2$ is given by $V_1 \cup V_3 \cup V_4$. Both vertex covers are likely to be reached if the initial solution of the (1+1) EA is $V_1 \cup V_2 \cup V_3 \cup V_4$. This initial solution is created by, for example, the maximal matching

$$M^* := \bigcup_{r=1}^k \{v_{1,r}, v_{2,r}\} \cup \{v_{3,r}, v_{4,r}\}. \quad (1)$$

THEOREM 3: *Suppose the maximal matching component of the (1+1) EA_M creates the cover $V_1 \cup V_2 \cup V_3 \cup V_4$ for Graph 3. Then with probability $\Omega(1)$ each,*

- *The (1+1) EA creates the globally optimal solution $V_2 \cup V_4$ in polynomial time,*
- *Stays at the locally optimal solution $V_1 \cup V_3 \cup V_4$ for a superpolynomial number of steps.*

PROOF: We first observe that no vertex from V_4 can be removed from the cover unless a step adds all missing vertices from V_5 simultaneously to the cover. If the fitness function did not count the number of chosen vertices from V_5 , this number would follow a random walk on the $2^{n/5}$ possible assignments to the bits corresponding to V_5 , and the time to reach the all-one assignment would be $2^{\Omega(n)}$ with probability at least $1 - 2^{-\Omega(n)}$ using the arguments on needle-in-a-haystack functions in Wegener and Witt (2005). Since the fitness function is linear in the number of chosen vertices in V_5 and has to be minimized, the actual time to reach the all-one assignment is even larger. We therefore assume that no vertex from V_4 is removed from the cover in any phase of length 2^{cn} for small enough constant c , and hence do not have to take into account the sets V_4 in V_5 in the sequel. As a mutation step of the (1+1) EA_M with probability at least $(1 - 1/n)^{2n/5} = \Omega(1)$ does not flip a bit corresponding to V_4 or V_5 and since these bits are treated independently of the bits corresponding to $V_1 \cup V_2 \cup V_3$, the assumptions will not affect our following asymptotic statements.

We are now faced with a situation like in the study of the (1+1) EA on the bipartite graph B in the paper by Friedrich et al. (2007, Theorem 5). The set $V' := V_1 \cup V_3$ of Graph 3 plays the role of V_2 of graph B and $V'' := V_2$ plays the role of V_1 . With probability $\Omega(1)$, a vertex from V' of Graph 3 is removed before a vertex from V'' is removed. Then we apply the argumentation concerning the second phase in the proof of Theorem 5 in Friedrich et al. (2007), which will show us that all vertices from V' are removed in polynomial time with constant probability. This will imply that the global optimum is reached with constant probability in polynomial time since the expected time to remove possible nodes from V_5 is also polynomial. More precisely, we consider a phase of $n^{3/2}$ mutation steps and show that all vertices of V' are removed with probability $\Omega(1)$. In this phase (all vertices of V'' and some vertices of V' chosen), the only mutation steps accepted by the (1+1) EA_M are the following. Either all missing vertices of V' are chosen and at least as many vertices of V'' are removed, or all vertices of V'' are kept and the number of vertices in V' is decreased (or stays the same by adding and removing some vertices). The former mutation step has a probability of at most n^{-k} , where k denotes the current number of missing vertices in V' . For the latter kind of mutation steps we restrict ourselves to one-bit flips reducing the number of vertices in V' . The probability for such a mutation step is at least $(1/n)(1 - 1/n)^{n-1} \geq 1/(en)$. For our calculations we take only those two kinds of mutation steps into account, the “good event” with probability at least $1/(en)$ and the “bad event” with probability at most n^{-k} since all other accepted mutation steps reduce or preserve the number of vertices in V' . The probability that the “good event” occurs before the “bad event” is at least $\frac{1}{en} / (\frac{1}{en} + n^{-k}) = 1 - \frac{e}{n^{k-1} + e}$. Thus, the probability that the vertices of V' were all removed by the (1+1) EA_M before the “bad event” occurs is at least

$$\prod_{k=1}^{2n/5} \left(1 - \frac{e}{n^{k-1} + e}\right) \geq \frac{1}{1 + e} \left(1 - \frac{e}{n}\right)^{n-e} = \Omega(1).$$




Figure 4: An illustration of Graph 4 for $k = 4$.

The expected waiting time for removing all vertices of V' by the $(1+1) EA_M$ is $O(n \log n)$ and, therefore, all vertices of V' are removed within $n^{3/2}$ steps with probability $1 - o(1)$ using Markov's inequality (always assuming that the "bad event" does not occur during this phase). Hence, the probability that the $(1+1) EA_M$ determines the global optimum as vertex cover in polynomial time is at least $\Omega(1)$.

The second claim follows by the same arguments. With probability $\Omega(1)$, a vertex from V_2 is removed first and the local optimum is reached with constant probability in polynomial time. To leave the local optimum, a step that flips k bits simultaneously has to occur. This has probability $n^{-\Omega(n)}$ and with probability $1 - n^{-\Omega(n)}$, does not happen in a superpolynomially long phase of 2^{cn} steps for a small enough constant $c > 0$. ■

It seems tempting to decrease the size of V_2 in Graph 3 to trick the $(1+1) EA$ into local optima with an even worse approximation ratio than $3/2$. This, however, does not work since the maximal matching component would choose only a subset of V_1 and V_3 then.

Using a less obvious, still bipartite graph, we can show that the $(1+1) EA$ is very likely to get stuck at $(2 - o(1))$ -optimal solutions if the initial maximal matching is chosen badly. The idea is to start from a big complete bipartite graph on vertex sets of size $k = \Omega(n)$ and to add small extra components of only $\log k$ vertices in a circular manner. Depending on the choice of the vertices from the small components, it can be necessary to add all vertices from the big bipartite graph to obtain a vertex cover. An optimal vertex cover, however, chooses only one side of the big bipartite graph.

Graph 4: Let $n = 2k + 4 \log k$. Let the vertex set be the union of

$$V_i = \{v_{i,1}, \dots, v_{i,k}\}$$

for $1 \leq i \leq 2$ and

$$V_i = \{v_{i,1}, \dots, v_{i,\log k}\}$$

for $3 \leq i \leq 6$. Define the edge set as the union of the sets

$$\{\{v, w\} \mid v \in V_i, w \in V_{(i \bmod 6)+1}\}$$

for $1 \leq i \leq 6$.

Figure 4 shows Graph 4 for $k = 4$.

To obtain a minimum vertex cover, we can choose the sets $V_1 \cup V_3 \cup V_5$ or $V_2 \cup V_4 \cup V_6$. An optimal vertex cover, therefore, is of size $k + 2 \log k$. A bad initial solution for the (1+1) EA is produced if all vertices $V_1 \cup V_2$ are chosen but the vertices from $V_3 \cup V_6$ are missing. This happens, for example if the maximal matching chooses the edges between V_1 and V_2 and those between V_4 and V_5 . The following theorem shows that it is really hard to obtain improvements from this initial solution.

THEOREM 4: *Suppose the maximal matching component of the (1+1) EA_M creates the cover $V_1 \cup V_2 \cup V_4 \cup V_5$ for Graph 4. Then with probability $1 - 2^{-\Omega(\log^2 n)}$, the (1+1) EA needs $2^{\Omega(\log^2 n)}$ steps to obtain a solution whose approximation ratio is better than $2 - O((\log n)/n)$.*

PROOF: Note that the initial cover is of the approximation ratio $(2k + 2 \log k)/(k + 2 \log k) = 2 - O((\log n)/n)$. To remove one of the vertices in V_1 or V_2 from the cover, all vertices from the neighboring sets V_6 and V_3 , respectively, must be included in the cover. The (1+1) EA can include vertices in the cover only if the number of vertices in a different set is decreased at the same time. However, for the supposed cover, the situation is similar to the proof of Theorem 3. Vertices from a subset (e. g., V_6) can only be added by swapping all vertices from this subset with all vertices from the neighboring set (e. g., V_5) in a single step. The latter has probability at most $(1/n)^{\log k} = 2^{-\Omega(\log^2 n)}$. Hence, the probability of an improvement in a phase of $2^c \log^2 n$ steps is still $2^{-\Omega(\log^2 n)}$ if $c > 0$ is chosen small enough but constant. ■

In the examples above, the search of the (1+1) EA_M is likely to get stuck at local optima since a worst-case initial maximal matching is assumed. We now turn to a more general view. In the following, we use the above-mentioned greedy algorithm to compute the maximal matching in the (1+1) EA_M. This means that we choose uniformly free edges until there are no such edges left. Let the obtained algorithm be called *greedy (1+1) EA_M* ((1+1) EA_{GM}).

With respect to the previous examples, Graphs 3 and 4, the probability that the search of the (1+1) EA leads to an optimal vertex cover seems to become higher when using the greedy maximal matching algorithm. (We do not go into the details here.) As stated above, we however cannot expect the (1+1) EA_{GM} to find efficiently optimal solutions on arbitrary instances. The following example shows when it is indeed likely to run into local optima of bad quality.

We define the bipartite graph $B(k, \ell)$, $\ell < k$, on $2k + \ell$ vertices as follows.

Graph 5: *Let the vertex set be*

$$\underbrace{\{v_1, \dots, v_k\}}_{=:L} \cup \underbrace{\{v_{k+1}, \dots, v_{2k}\}}_{=:R} \cup \underbrace{\{v_{2k+1}, \dots, v_{2k+\ell}\}}_{=:C}$$

and the edge set be

$$\begin{aligned} & \{ \{v_i, v_{k+i}\} \mid i = 1, \dots, k \} \\ & \cup \{ \{v_i, v_{2k+j}\} \mid i = 1, \dots, k, j = 1, \dots, \ell \}. \end{aligned}$$

Hence, we obtain the whole edge set from the induced subgraph on $L \cup R$, which is a perfect matching, and the induced subgraph on $L \cup C$, which is a complete bipartite

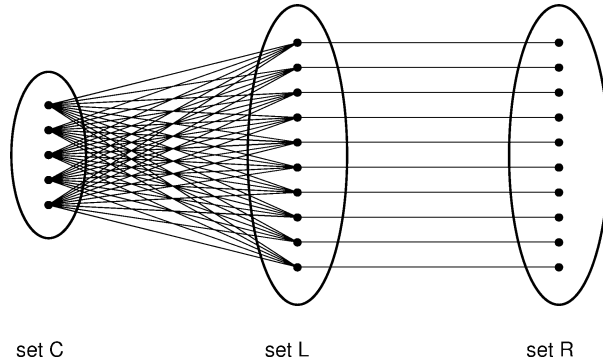


Figure 5: An illustration of Graph 5 for $k = 10$ and $\ell = 5$.

graph. A minimum vertex cover chooses all L -vertices while any vertex cover that leaves out an L -vertex in favor of its adjacent R -vertex must choose all C -vertices. The following theorem (note that $k = \Omega(n)$) shows that the latter event can mislead the $(1+1)$ EA_{GM} into a local optimum.

THEOREM 5: *Let $\ell \leq k - 2 \log n$. Then with probability $\Omega(1 - \ell/k)$, the $(1+1)$ EA_{GM} on $B(k, \ell)$ needs at least $2^{\Omega(k-\ell)}$ steps to create a solution that is better than $(1 + \ell/k)$ -approximate.*

PROOF: The proof outline is as follows. We first show that the maximal matching routine will choose all C -vertices with the claimed probability. Afterwards, it removes with this probability at least $\Omega(k - \ell)$ of the L -vertices from the cover before a C -vertex flips. From this situation, it takes with probability $\Omega(1)$ at least $2^{\Omega(k-\ell)}$ steps to obtain a cover that does not choose all C -vertices. Up to then, all covers contain at least $k + \ell$ vertices and are, therefore, no better than $(1 + \ell/k)$ -approximate.

We call the edges that are incident on C -vertices C -edges and the remaining edges, which are incident on R -vertices, R -edges. We are interested in matchings that consist of ℓ C -edges and $k - \ell$ R -edges. If i C -edges and j R -edges have already been chosen during the construction of the maximal matching, there are $\ell - i$ free C -vertices and $k - i - j$ free L -vertices left. Hence, the number of free C -edges equals $(\ell - i)(k - i - j)$ while there are only $k - i - j$ free R -edges left. The probability of choosing another R -edge in this situation is

$$\frac{(\ell - i)(k - i - j)}{(\ell - i)(k - i - j) + (k - i - j)} = \frac{\ell - i}{\ell - i + 1}$$

Therefore, for $i < \ell$, the expected number of chosen R -edges between the i -th and the $(i + 1)$ -st choice of a C -edge is at most

$$\frac{\ell - i + 1}{\ell - i} - 1 = \frac{1}{\ell - i}$$

since the random number of steps between the two choices follows a geometric distribution. Hence, the expected number of chosen R -edges until ℓ C -edges have been chosen is at most $\sum_{i=0}^{\ell-1} 1/(\ell - i) \leq (\ln \ell) + 1$. By Markov's inequality, the number is at

most $k - \ell$, that is, ℓ C -edges are chosen, with probability at least $1 - (\ln \ell + 1)/(k - \ell)$. Due to our assumption on ℓ , the last expression is positive and bounded from below by $\Omega(1 - \ell/k)$.

For the second part of the proof, we assume that a matching with ℓ C -edges and $k - \ell$ R -edges has been created and consider the subsets $L^* \subseteq L$ and $R^* \subseteq R$ consisting of the $k - \ell$ vertices in L resp. R on which the R -edges of the matching are incident. Starting from the vertex cover corresponding to the matching, the (1+1) EA accepts each step that flips only an L^* -bit. We consider the probability that an L^* -vertex flips before a C - or R^* -vertex flips. By simple calculations, this probability is bounded from below by $\Omega(1 - k/\ell)$. We assume that such a flip occurs. Let v^* be the flipping L^* -vertex. As long as v^* is not added to the cover again, no C -vertices can be removed.

In the following, we are interested in the event of reducing the number of chosen L^* -vertices even further. Consider a phase consisting of the $n/4$ steps after v^* has been removed. With probability $\Omega(1)$, v^* is not chosen again during the phase. We assume this to happen. The expected number of flipping R^* -bits in the phase is bounded from above by $(k - \ell)/4$. By Chernoff bounds, the number is at most $(k - \ell)/2$ with probability $1 - 2^{-\Omega(k-\ell)}$. If this event holds, there are throughout the phase still $(k - \ell)/2 - 1$ L^* -vertices left that can be (or already have been) removed without violating the cover. We apply Chernoff-bound arguments again. Altogether, in the end of the phase, we have with probability $\Omega(1) - 2^{-\Omega(k-\ell)}$ arrived at a cover that contains all C -vertices but misses at least $\Omega(k - \ell)$ of the L -vertices. We assume to be in such a situation.

To obtain a cover without C -vertices from this situation, all L -vertices have to be chosen at at least one point of time. We consider the $\Omega(k - \ell)$ R -edges whose L -vertex is unchosen. The random number of unchosen L -vertices can be increased or decreased. This process can be identified with the random walk of the (1+1) EA on a needle-in-a-haystack function where the R -edges correspond to bits and an unchosen L -vertex of an R -edge corresponds to a zero-bit and an unchosen R -vertex to a one-bit. With $\Omega(k - \ell)$ bits and starting from $\Omega(k - \ell)$ zero-bits, the time until the all-ones string is reached is bounded from below by $2^{\Omega(k-\ell)}$ with probability $1 - 2^{-\Omega(k-\ell)} = \Omega(1 - k/\ell)$ (using the results in Wegener and Witt, 2005). ■

Theorem 5 provides only a lower bound $\Omega(1 - \ell/k)$ on the probability of reaching the local optimum which seems to be too pessimistic for small ℓ since the bound is only $\Omega(1)$ then. If ℓ/k converges to 0, we however conjecture that the probability converges to 1. This can be made precise for the special case $\ell = 1$.

THEOREM 6: *Let $\ell = 1$. Then with probability $1 - O(1/\sqrt{n})$, the (1+1) EA_{GM} on $B(k, \ell)$ needs at least $2^{\Omega(n)}$ steps to create a solution that is better than $(1 + 1/n)$ -approximate.*

PROOF: The proof follows the same lines as the one of Theorem 5. Using the arguments from the first part of the proof, it is easy to see that the greedy maximal matching algorithm chooses with probability at least $1 - 1/n$ one C -edge and $k - 1$ R -edges. This case means that the initial solution of the (1+1) EA chooses the single C -vertex and $2(k - 1) = n - 3$ L - and R -vertices. We assume this to happen.

Next consider the phase of the \sqrt{n} steps after initialization. With probability $1 - O(1/\sqrt{n})$, the C -vertex does not flip in the phase. Using the Chernoff-bound arguments from the proof of Theorem 5, we prove that with probability $1 - 2^{-\Omega(\sqrt{n})}$, at least $\Omega(\sqrt{n})$ L -vertices are removed from the cover. Note that we need not assume that an L -vertex flips before the first R -vertex flips.

Assuming that $\Omega(\sqrt{n})$ L -vertices have been removed by the end of the phase, we apply the ideas from the last paragraph of the proof of Theorem 5. Hence, the probability that all L -vertices are in the cover again—which is necessary for the C -vertex to flip—before $2^{\Omega(n)}$ steps have elapsed is $2^{-\Omega(n)}$. ■

We are left with several interesting open problems. It would be desirable to have an example where the $(1+1)$ EA_{GM} gets stuck at $(2 - o(1))$ -optimal solutions with probability $1 - o(n^{-c})$ for any constant c , that is, the success probability should be only superpolynomially small. Such examples should exist if the vertex cover problem does not admit approximations with a constant factor less than 2. Moreover, one could work on generalizations of Proposition 1. Since the $(1+1)$ EA is able to flip any given subset of $\ln n / \ln \ln n$ bits in a single step with probability $n^{-O(1)}$, it might happen that it is able to explore in polynomial time the Hamming ball of radius $\ln n / \ln \ln n$ centered around the initial solution. This would imply that solutions with an approximation ratio $2 - O(\ln n / (n \ln \ln n))$ could be found in polynomial time with at least polynomially small probability.

5 Conclusions

In the case of combinatorial optimization problems hybrid methods are often used to obtain good solutions for a certain problem. The theoretical understanding of combining evolutionary algorithms with other methods is rather weak. We have made a first step into the rigorous analysis of such methods by considering the combination of well-known approximation algorithms with a simple evolutionary algorithm. In our analyses we have pointed out situations where approximate solutions can (or cannot) be improved by the search procedure of an EA.

There are several open questions concerning the topic of analyzing the combination of evolutionary algorithms with other methods. Some regarding the vertex cover problem and starting with solutions computed by the considered approximation algorithms have been pointed out in Sections 3 and 4. We also want to state a more general question. The EAs considered in this paper only use an initial solution that has been computed by another method. In general a different optimization procedure is applied more often during the run of an EA. These so-called memetic approaches have already been analyzed w. r. t. their runtime on toy problems (Sudholt, 2006a,b). It would be nice to have some results on such methods by rigorous analyses on a well-known combinatorial optimization problem.

Acknowledgments

J. He was supported by the UK Engineering and Physical Research Council under Grant No. EP/C520696/1. C. Witt was supported by the Deutsche Forschungsgemeinschaft (DFG) in terms of the Collaborative Research Center “Computational Intelligence” (SFB 531).

References

Dakin, R. J. (1965). A tree-search algorithm for mixed integer programming problems. *The Computer Journal*, 8:250–255.

- Doerr, B., Hebbinghaus, N., and Neumann, F. (2006). Speeding up evolutionary algorithms through restricted mutation operators. *Proceedings of PPSN '06*, vol. 4193 of LNCS, pp. 978–987.
- Evans, I. K. (1998). Evolutionary algorithms for vertex cover. In V. W. Porto, N. Saravanan, D. E. Waagen, and A. E. Eiben (Eds.), *Proceedings of Evolutionary Programming VII*, vol. 1447 of LNCS, pp. 377–386. Berlin: Springer.
- Friedrich, T., He, J., Hebbinghaus, N., Neumann, F., and Witt, C. (2007). Approximating covering problems by randomized search heuristics using multi-objective models. *Proceedings of GECCO '07*, pp. 797–804. New York: ACM Press.
- Giel, O., and Wegener, I. (2003). Evolutionary algorithms and the maximum matching problem. *Proceedings of STACS '03*, vol. 2607 of LNCS, pp. 415–426.
- Giel, O., and Wegener, I. (2006). Maximum cardinality matchings on trees by randomized local search. *Proceedings of GECCO '06*, pp. 539–546. New York: ACM Press.
- Glover, F., and Laguna, M. (1997). *Tabu search*. Norwell, MA: Kluwer.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley.
- He, J., Yao, X., and Li, J. (2005). A comparative study of three evolutionary algorithms incorporating different amount of domain knowledge for node covering problems. *IEEE Transactions on Systems, Man and Cybernetics*, 35(2):266–271.
- Holland, J. H. (1992). *Adaptation in natural and artificial system*, 2nd ed. Cambridge, MA: MIT Press.
- Karakostas, G. (2005). A better approximation ratio for the vertex cover problem. *Proceedings of ICALP '05*, vol. 3580 of LNCS, pp. 1043–1050. Berlin: Springer.
- Kehden, B., and Neumann, F. (2006). A relation-algebraic view on evolutionary algorithms for some graph problems. *Proceedings of EvoCop '06*, vol. 3906 of LNCS, pp. 147–158.
- Khuri, S., and Bäck, T. (1994). An evolutionary heuristic for the minimum vertex cover problem. In J. Hopf (Ed.), *Genetic algorithms within the framework of evolutionary computation, Proceedings of the KI-94 Workshop*, pp. 86–90.
- Michalewicz, Z. (1996). *Genetic algorithms + data structure = Evolution program*, 3rd ed. Berlin: Springer Verlag.
- Neumann, F. (2004). Expected runtimes of evolutionary algorithms for the Eulerian cycle problem. *Proceedings of CEC '04*, vol. 1, pp. 904–910. Piscataway, NJ: IEEE Press.
- Neumann, F. (2007). Expected runtimes of a simple evolutionary algorithm for the multi-objective minimum spanning tree problem. *European Journal of Operational Research*, 181(3):1620–1629.
- Neumann, F., and Wegener, I. (2004). Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Proceedings of GECCO '04*, vol. 3102 of LNCS, pp. 713–724.
- Neumann, F., and Wegener, I. (2005). Minimum spanning trees made easier via multi-objective optimization. *Proceedings of GECCO '05*, pp. 763–770. New York: ACM Press.
- Papadimitriou, C. H., and Steiglitz, K. (1998). *Combinatorial optimization: Algorithms and complexity*. Mineola, NY: Dover.
- Scharnow, J., Tinnefeld, K., and Wegener, I. (2004). The analysis of evolutionary algorithms on sorting and shortest paths problems. *Journal of Mathematical Modelling and Algorithms*, pp. 349–366.

- Schrijver, A. (1998). *Theory of linear and integer programming*. New York: John Wiley & Sons.
- Sudholt, D. (2006a). Local search in evolutionary algorithms: The impact of the local search frequency. *Proceedings of ISAAC '06*, vol. 4288 of LNCS, pp. 359–368.
- Sudholt, D. (2006b). On the analysis of the (1+1) memetic algorithm. *Proceedings of GECCO '06*, pp. 493–500. New York: ACM Press.
- Vazirani, V. (2001). *Approximation algorithms*. Berlin: Springer.
- Wegener, I., and Witt, C. (2005). On the optimization of monotone polynomials by simple randomized search heuristics. *Combinatorics, Probability and Computing*, 14:225–247.
- Witt, C. (2005). Worst-case and average-case approximations by simple randomized search heuristics. *Proceedings of STACS '05*, vol. 3404 of LNCS, pp. 44–56.