Evaluating a Hybrid Encoding and Three Crossover Operators on the Constrained Portfolio Selection Problem

Felix Streichert, Holger Ulmer, and Andreas Zell
Centre for Bioinformatics Tübingen (ZBIT)
University of Tübingen, Sand 1, 72076 Tübingen, Germany
Email: streiche@informatik.uni-tuebingen.de

***Abstract*— In this paper we investigate the impact of different crossover operators for a real-valued Evolutionary Algorithm on the constrained portfolio selection problem based on the Markowitz mean-variance model. We also introduce an extension of a real-valued genotype, which increases the performance of the Evolutionary Algorithm significantly, independent of the crossover operator used. This extension is based on the effect that most efficient portfolios only consist of a selection of few assets. Therefore, the portfolio selection problem is actually a combination of a knapsack and continuous parameter problem. We also introduce a repair mechanism and examine the impact of Lamarckism on the performance of the Evolutionary Algorithm.**

I. Introduction

There are numerous optimization problems in the area of financial engineering like index tracking, credit scoring, identifying default rules, time series prediction, trading rules, etc. But one of the most prominent is the portfolio selection problem, which is given by the task of how to distribute a limited amount of money between multiple assets available for a profitable investment strategy.

Markowitz made an early approach to give the portfolio selection problem a mathematical background, the Markowitz mean-variance model [12], [13]. This model assumes that an investor would always try to maximize the return of his investments while at the same time securing his investments from a possible loss. Therefore, the portfolio problem gives a multi-objective optimization problem (MOOP), maximizing the expected return on the one hand and on the other hand minimizing the risk (variance) of the portfolio.

While the unconstrained portfolio selection problem based on the Markowitz mean-variance model can be solved through quadratic programming, this is not the case for the constrained portfolio selection problem. Therefore, there have been several alternative approaches to the portfolio selection problem.

One of the first groups to apply Evolutionary Algorithms (EA) on the portfolio selection problem were Tettamanzi et al. [1], [11], [10]. Here, the MOOP was transformed into a single-objective problem by using a trade-off function. They used multiple EA populations with individual trade-off coefficients and found that a parallel implementation performed much better than a sequential one. More recently Crama et al. applied Simulated Annealing (SA) to the portfolio selection problem [5]. They especially pointed out that SA and similar heuristics have the major advantage that they can be easily applied to any kind of model with arbitrary constraints without much modification. For the same reason Beasley et al. compared Tabu Search, SA and EA on the portfolio selection to evaluate their performance [4]. They solved the MOOP by using one objective as constraint, which was increased iteratively to obtain the complete Pareto front. But they found that no heuristic performed significantly better and concluded that only a combination of all three heuristics was satisfying.

Unfortunately, all these papers did not utilize the advantages of multi-objective EAs (MOEA) to the portfolio selection problem, although MOEA have shown to be very useful on MOOPs [9], [6], [20]. In this paper we apply such a MOEA and further suggest a new hybrid encoding of the portfolio selection that proves to be more efficient than a standard encoding. We also examine the impact of several real-valued crossover operators and the effect of an additional repair mechanism to search for feasible solutions with and without Lamarckism.

In the next section we give a short introduction to the Markowitz mean-variance model and the constraints we applied to the portfolio selection problem. In sec. III we give details of the EA, the multi-objective optimization strategy we applied, the crossover operators and the hybrid encoding we suggest. Experimental results are presented in sec. IV and conclusions and an outlook on future work are given in sec. V and sec. VI.

II. The Portfolio Optimization Problem

In this paper we use the standard Markowitz mean-variance model for one time step. The optimization problem is to find a portfolio p consisting of N assets with specific volumes given as weights w_i by

- minimizing the variance σ_p of the portfolio

$$\sigma_p = \sum_{i=1}^{N} \sum_{j=1}^{N} w_i \cdot w_j \cdot \sigma_{ij} \quad (1)$$

- while maximizing the return μ_p of the portfolio

$$\mu_p = \sum_{i=1}^{N} w_i \cdot \mu_i \quad (2)$$

subject to

$$\sum_{i=1}^{N} w_i = 1 \quad , \tag{3}$$
$$0 \leq w_i \leq 1 \quad ; \quad i = 1, .., N \tag{4}$$

where N is the number of assets available, μ_i the expected return of asset i, σ_{ij} the covariance between asset i and j. Usually μ_i and σ_{ij} are to be estimated from historic data.

Eq. 1 and 2 give the two competing objectives, which are to be optimized. Eq. 3 and 4 give the constraints of a feasible portfolio: all the money available is to be invested, and all investments should be positive, i.e. no short sales are allowed. As said before, this basic form of the mean-variance model is a quadratic optimization problem, for which computationally effective algorithms exist, unfortunately this is not the case when we add real-world constraints:

Cardinality Constraints restrict the maximum number of assets used in the portfolio:

$$\sum_{i=1}^{N} \text{sign}(w_i) = K \tag{5}$$

Buy-in Thresholds give the minimum amount to be purchased, in case the asset should be in the portfolio:

$$w_i \geq l_i \quad \forall \quad w_i > 0; \quad i = 1, .., N \tag{6}$$

Roundlots give the smallest volumes c_i that can be purchased for each asset:

$$w_i = y_i \cdot c_i; \quad i = 1, .., N \text{ and } y_i \in \mathbb{Z} \tag{7}$$

These constraints are often hard constraints, i.e. they cannot be violated. Other real-world constraints like sector/industry constraints, immunization/duration matching and taxation constraints can be considered soft constraints. Soft constraints may be violated, because a violation in such constraints may lead to significantly higher performance in the other objectives or because valid solutions do not even exist. Therefore, it is more reasonable to implement violations of soft constraints as additional objectives, which are to be minimized. For this reason we currently do not consider such soft constraints since they would just increase the output dimension of the MOOP.

III. The Evolutionary Algorithm

In our experiments we apply a generational EA population strategy with a population size of 500 individuals. We use tournament selection with a tournament group size of 8 together with objective space based fitness sharing with a sharing distance of $\sigma_{share} = 0.01$. The selection mechanism prefers individuals that are better than other individuals in at least one objective value, i.e. which are not dominated by another individual. To maintain the currently known Pareto front we use an archive of 250 individuals and use the archive as elite to achieve a faster speed of convergence. Details of this MOEA strategy can be found in [16].

In this paper an EA with a real-valued genotype is applied. We use local mutation with one strategy parameter σ_i for each decision variable on the EA genotype, which mutates each decision variable w_i by adding a random gaussian number with the deviation σ_i [15].

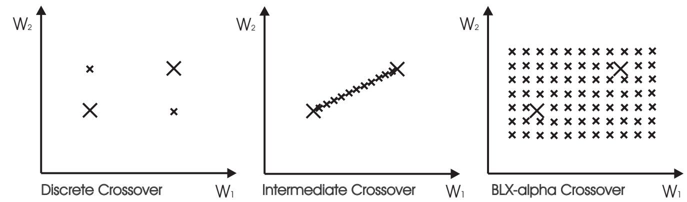

Fig. 1. Effects of the crossover operators. A big cross (X) indicates a parent, while a small cross (x) gives a possible offspring.

We compare three crossover operators: a discrete 3-point-crossover [19], the BLX-α crossover [7] and the intermediate crossover [14]. An example for the general effects of each crossover operator is given in Fig. 1.

The **discrete N-point-crossover** equals the mechanism used in bit-string crossover, N points ($\in \{1, 2, ..., n-1\}$) are selected where the chromosomes of the parents are swapped to produce the offsprings.

An example for 1-point crossover:

$$\begin{aligned} \mathbf{W}'^1 = & \quad (w_1^1, w_2^1, ..., w_i^1, w_{i+1}^2, ..., w_n^2) \\ \mathbf{W}'^2 = & \quad (w_1^2, w_2^2, ..., w_i^2, w_{i+1}^1, ..., w_n^1) \end{aligned} \tag{8}$$

The **intermediate crossover** operator uses a linear combination of w_i of all k parents to set the decision variables of the offsprings: $w\text{'}_i^j = \sum_{j=0}^{k} \alpha_j \cdot w_i^j$. The linear factor α_i is a random variable with $\sum_{j=0}^{k} \alpha_j = 1$ and $0 \leq \alpha_j \leq 1 \; \forall \; j$.

The **BLX-α crossover** operator reinitializes the values of W' with values from a extended range given by the parents, $w\text{'}_i^j = \textit{UniformRandomDouble}\ (w_{i,min} - I \cdot \alpha, w_{i,max} + I \cdot \alpha)$ with $w_{i,min} = Min(w_i^1, w_i^2)$, $w_{i,max} = Max(w_i^1, w_i^2)$ and $I = w_{i,max} - w_{i,min}$. We use $\alpha = 0.5$ as suggested in [8].

Our EA uses a mutation probability of $p_m = 1.0$, a crossover probability of $p_c = 0.5$ and uses $k = 2$ parents for crossover. The general EA parameters were selected from preliminary experiments.

In the following subsections details are given on the repair mechanism, which creates feasible solutions from arbitrary decision variables w_i, and the hybrid encoding.

A. The Repair Mechanism

Without additional real-world constraints the repair algorithm is rather simple. Since our implementation of the EA encodes each decision variable in the desired range, $w_i \in \{0, 1\}$, only Eq. 3 must be met. This can be easily done by normalization, $w_i' = w_i / \sum w_i$.

If cardinality constraints are added, the repair mechanism sets all surplus decision variables w_i to zero and keeps only the K largest values of w_i before applying the normalization.

With buy-in constraints the algorithm sets all w_i below their given buy-in threshold to zero after applying the cardinality repair mechanism and the following normalization step is only allowed to redistribute the amount of the portfolio less the already assigned amount needed to meet the buy-in threshold.

To meet roundlot constraints the algorithm rounds w_i to the next roundlot level, $w_i' = w_i - (w_i \ \ mod \ c_i)$, after cardinality repair, buy-in repair and normalization was applied. The remainder of the rounding process, $\sum_i (w_i \ \ mod \ c_i)$, is spent in quantities of c_i on those w_i', which had the biggest values for $w_i \ \ mod \ c_i$ until all of the remainder is spent.
Fortunately, the repair algorithm is deterministic. Therefore, an individual is always assigned to the same solution after repair if the genotype did not change. Depending on the original w_i and the constraints, the repair algorithm may fail to find feasible solutions, in that case the fitness of the individual will be set to the worst possible value instead of assigning a random feasible solution. The repair algorithm only affects the phenotype of an individual, while the genotype remains unaltered.

B. The Hybrid Encoding

Preliminary experiments indicated that pareto-optimal solutions for the portfolio selection problem are rarely composed of all available assets, but only a limited selection of the available assets, especially in case of cardinality constraints, see Fig. 2. For $K = 2$ there are several distinct regimes of two assets combinations that form the Pareto front. The same holds true for larger values of K. But the less restrictive the cardinality constraints are, the less distinct the regimes.
The problem to find the best combinations of assets in the portfolio resembles a one-dimensional binary knapsack problem. This kind of problem has already been addressed by means of EA using a binary genotype. We suggest to use the very same genotype in addition to the vector of decision variables **W**, see Fig. 3. Each bit of the bit-string **B** determines whether the associated asset will be an element of the portfolio or not, so that the actual value of the decision variable is $w_i' = b_i \cdot w_i$. This is the value that will be processed by the repair mechanism. With this hybrid encoding it is much easier for the EA to add or remove the associated asset simply by mutating the bit-string **B**.

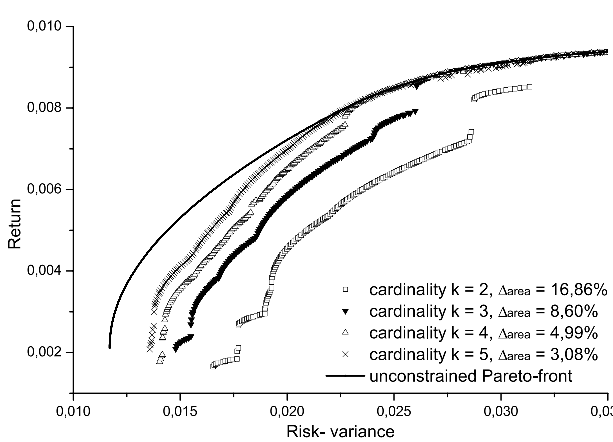

Fig. 2. Solutions generated by EA with the hybrid encoding on the *DAX* data set with 81 assets as given in [2].

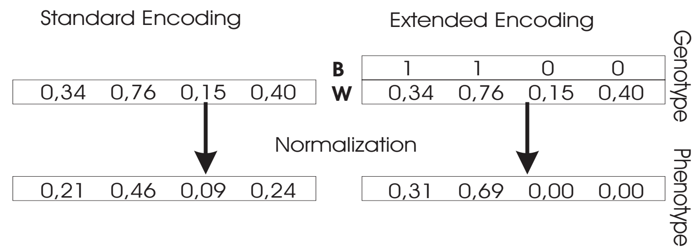

Fig. 3. Comparing the standard encoding to the hybrid encoding.

The hybrid encoding is altered by mutating/crossing each genotype **B** and **W** separately from each other. Binary one-point mutation ($p_m = 0.1$) and 3-point-crossover ($p_c = 1.0$) is used on the bit-string **B** and the real-valued operators mentioned before are used on the decision variables **W**.
The extended EA is abbreviated KEA (Knapsack-EA). An general comparison between GA and ES against a KGA and KES has been performed in [17].

IV. EXPERIMENTAL RESULTS

The comparison of the different EA implementations was performed on benchmark data sets given by Beasley [2] available at *http://mscmga.ms.ic.ac.uk/info.html*. The numerical results presented here were performed on the *Hang Seng* data set with 31 assets. On this data set we use several combinations of real-world constraints to compare the performance of the different EA encodings and crossover operators. First, we compare the portfolio selection problem without cardinality constraints and with cardinality constraints $K = 6$, $K = 4$ and $K = 2$. In a second set of experiments we also add buy-in thresholds ($l_i = 0.1$) and roundlots constraints ($c_i = 0.02$) to the portfolio selection problem.
We measure the performance of the algorithms by calculating the S-metric [21], i.e. the area under the currently achieved Pareto front bounded by μ_{max} of the maximum return asset and $\sigma = 0$. We compare this area to the area under the Pareto front of the unconstrained portfolio selection problem calculated through quadratic programming also given in the benchmark data set. The percentage difference (Δ_{area}) of the EA calculated solution and the reference solution is to be minimized and gives the measure of quality. But only without any real-world constraints can this measure drop to zero, otherwise the Δ_{area} is limited by the constraints, compare Fig. 2. Additionally the Δ_{area} is limited due to the limited size of the archive population, which gives the Pareto front identified by the EA.
To obtain reliable results we repeat each EA experiment for 50 times for each parameter setting and problem instance. A single EA run is terminated after 100,000 fitness evaluations. We calculate the mean value, the standard deviation, the maximum and minimum values and the 90 % confidence intervals of the Δ_{area} value to compare the performance of each EA setting.

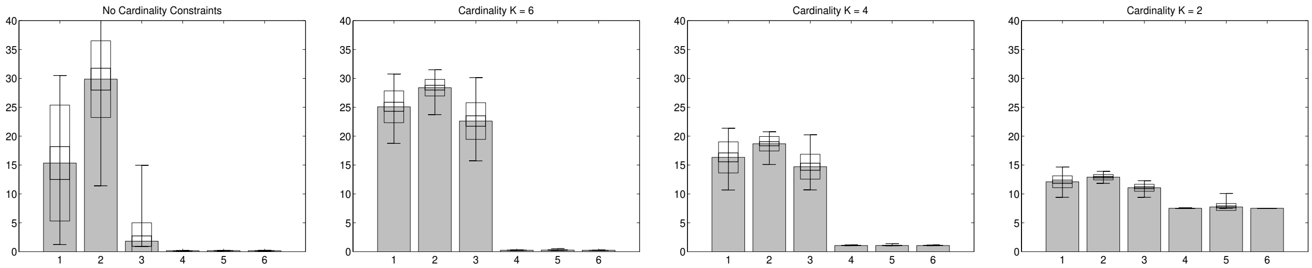

Fig. 4. Δ_{area} for the experiments on the Hang Seng data set without additional constraints l_i and c_i (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover)

A. Results without Lamarckism

In the experiments without Lamarckism only the phenotype of an individual is altered by the repair mechanism while the genotype remains unaltered, see Fig. 3. If the repair mechanism is interpreted as local search mechanism, then the optimization process is guided by the Baldwin effect. In this case the search space becomes neutral to some extent, i.e. a mutation of the genotype does not necessarily change the phenotype. Neutrality caused by the Baldwin effect is said to support the optimization, since it enables the EA to escape from local optima by chance due to genetic drift on a plateau of equal fitness [18].

1) Without Additional Constraints: On all four problem instances without additional constraints l_i and c_i the hybrid encoding performs significantly better than the standard encoding regardless of the crossover operator used, see Fig. 4. The confidence intervals indicate that the hybrid encoding is also very reliable compared to the standard encoding.
Further, the KEA not only outperforms the standard EA regarding the quality of the solution found, but also the speed of convergence, see Fig. 5 and Fig. 6. Without cardinality constraints the KEA even has a better start since the additional bit-string removes about half of the possibly unnecessary assets from the portfolio, which have to be removed in case of the standard encoding by other means. With $K = 4$ the initial quality of the EA solutions for the standard encoding equals that of the hybrid encoding, see Fig. 6. This is because of the repair mechanism, which comes into action as a result of the cardinality constraints. The repair mechanism removes the surplus assets from the portfolio for both the KEA and the standard EA. This way the standard EA starts with the same sparse portfolios as the KEA. But the speed of convergence for the EA is significantly slower than that of the KEA, see Fig. 6. This is because the KEA is more efficient to create portfolios with smaller cardinalities than given by the constraints. The standard EA even starts to stagnate far from the global optimum.
When comparing the different crossover operators for the standard EA implementation the intermediate crossover performs worst, while the BLX-α crossover performs significantly better than the other two crossover operators. Especially without cardinality constraints the result of the BLX-α crossover on the standard EA even comes close to the result of the KEA. But with increasing cardinality the difference becomes less significant. When compared on the KEA the crossover operators do not really differ from each other. Only the variance for the intermediate crossover is slightly higher and it converges slower than discrete and BLX-α crossover, see Fig. 6.

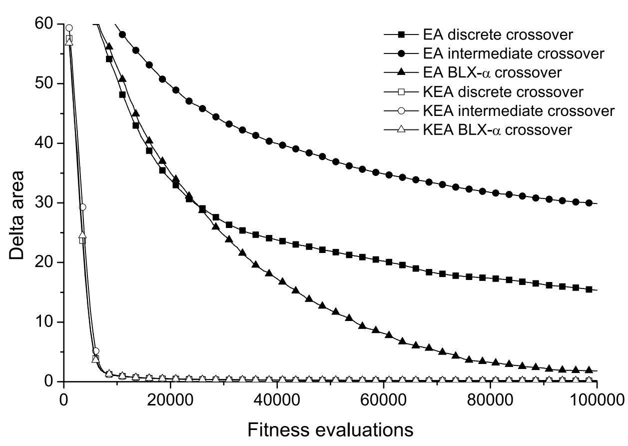

Fig. 5. Convergence behavior of Δ_{area} on the Hang Seng data set without cardinality constraints and without l_i and c_i constraints

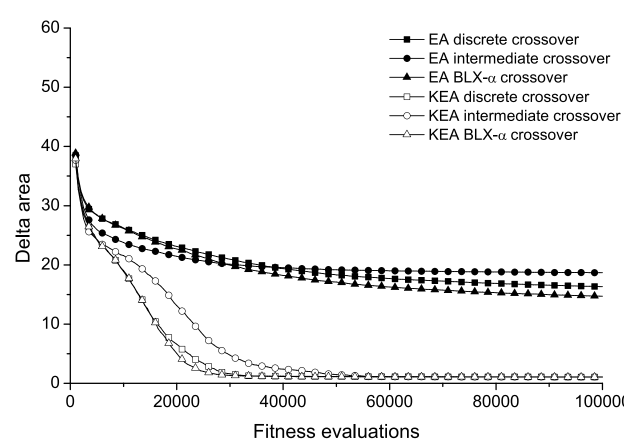

Fig. 6. Convergence behavior of Δ_{area} on the Hang Seng data set with $K = 4$ and without l_i and c_i constraints

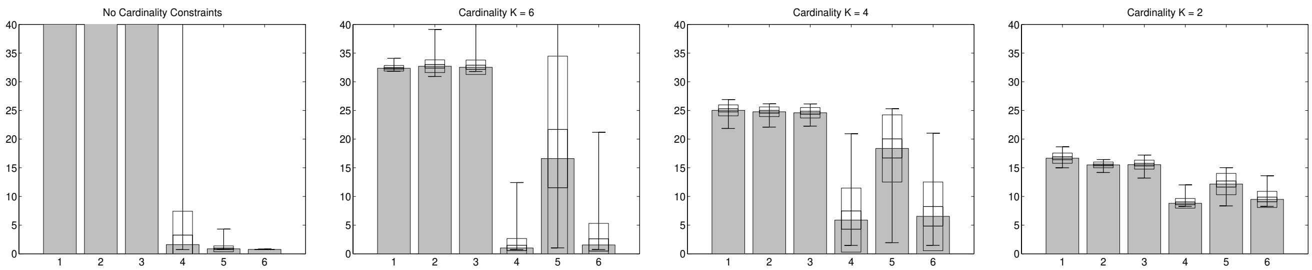

Fig. 7. Δ_{area} for the experiments on the Hang Seng data set with $l_i = 0.1$ and $c_i = 0.02$ (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover)

2) With Additional Constraints: With additional buy-in thresholds and roundlot constraints the portfolio selection problem becomes much more complicated and the performance of both EA approaches drops considerably. But the additional constraints cause the EA to have the same initial quality as the KEA, because $l_i = 0.1$ behaves like a cardinality constraint of $K = 10$, see Fig. 8 and Fig. 9.

Again the KEA performs much better than the standard EA, see Fig. 7. But the KEA converges slower to the Pareto front and is not as reliable as it was the case without additional constraints, but the KEA is basically able to find the Pareto front. This is demonstrated by the best results of the KEA runs, which are very close to the true Pareto front, see Fig. 7. The standard EA on the other hand starts to stagnate very fast and converges to a local optimum, see Fig. 8 and Fig. 9. In absence of cardinality constraints this can be easily explained. The buy-in threshold $l_i = 0.1$ acts like a cardinality constraint of $K = 10$. A randomly initialized decision vector will have mean values of $1/2$ for each asset. The repair mechanism will remove any surplus assets from the portfolio and keep only the $K = 10$ assets with the biggest values w_i. Unfortunately, the standard EA will not be able to create portfolios of lower cardinality, since every time an asset is removed from the portfolio through mutation or crossover it will be replaced by the next biggest w_i of the $N - K$ assets previously not element of the portfolio. Since the other $N - K$ asset weights w_i will have random values due to genetic drift in the neutral search space caused by the repair mechanism. Therefore, the standard EA will only search the subspace where portfolios are of cardinality K and the assets in the portfolio are assigned weights of $w_i \approx 1/K$. This is the reason why the standard EA converges to suboptimal Pareto fronts. The same effect can also be observed on problem instances with cardinality constraints, see Fig. 9. This problem will be discussed more detailed in sec. IV-B.2.

Comparing the crossover operators for the EA and the KEA on the portfolio problem with additional constraints the situation of sec. IV-A.1 is reversed. For the EA no significant difference can be observed neither regarding the resulting quality nor the convergence behavior. For the KEA this is only the case, if no cardinality constraints are present. Except for one extreme outlier occurring during the discrete crossover runs, the intermediate crossover is only insignificantly worse than discrete and BLX-α crossover. This difference grows, if cardinality constraints are added. The intermediate crossover performs worst and has also a slower speed of convergence, see Fig. 9. But in this case the discrete crossover performs a little better than the BLX-α crossover.

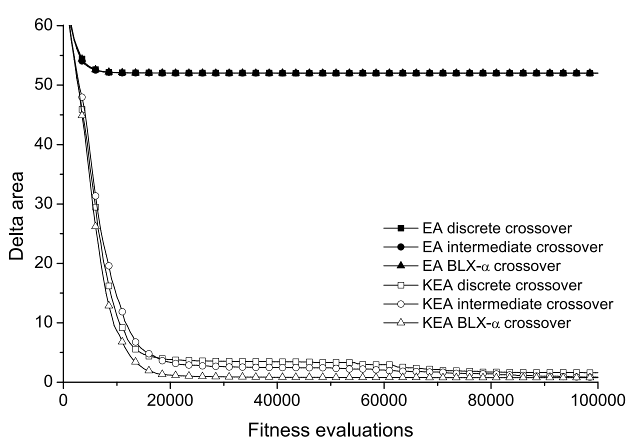

Fig. 8. Convergence behavior of Δ_{area} on the Hang Seng data set without cardinality constraints, $l_i = 0.1$ and $c_i = 0.02$ constraints

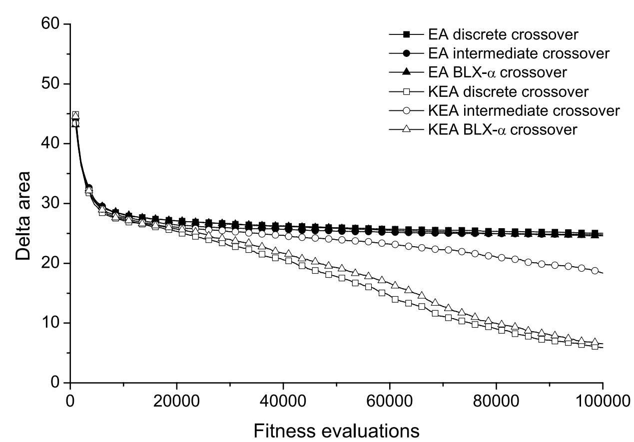

Fig. 9. Convergence behavior of Δ_{area} on the Hang Seng data set with $K = 4$, $l_i = 0.1$ and $c_i = 0.02$ constraints

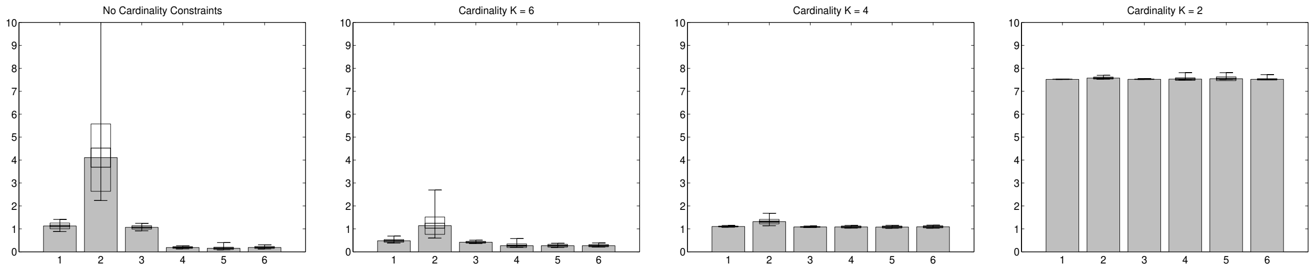

Fig. 10. Δ_{area} for the experiments on the Hang Seng data set with Lamarckism, without l_i and c_i (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover))

B. Results with Lamarckism

With Lamarckism the repair algorithm alters the genotype of an individual according to the phenotype. This way Lamarckism removes the neutrality of the search space, since any neutral mutation will be reversed by Lamarckism.

One major effect of Lamarckism in case of cardinality constraints is that the vector of decision variables will become sparse, since surplus assets are removed from the genotype. With such sparse decision vectors the search of the standard EA will become as efficient as the KEA, since it can add and remove assets from the portfolio as easily.

1) Without Additional Constraints: As expected the advantage of the hybrid encoding becomes counterbalanced through the application of Lamarckism on all problems with cardinality constraints, see Fig. 10. Especially with increasing cardinality constraints the difference between the standard and the hybrid encoding becomes more and more negligible and vanishes completely for $K = 2$, see Fig. 10. Even the convergence behavior of the standard EA equals the KEA and both are accelerated significantly, compare Fig. 12 to Fig. 6. Although the KEA still has the edge on the standard EA.

But even without cardinality constraints Lamarckism has an advantageous effect on the standard EA, see Fig. 10 and Fig. 11. This can be explained by the reduction of the search space through Lamarckism, since mutation and crossover only act on valid solutions. This also favors easy removal of assets from the portfolio. While a randomly initialized individual has a mean value for w_i of 0.5, an individual after application of Lamarckism and without cardinality constraints has a mean value of $1/N$. With such a low average value for w_i it is much easier for mutation to remove surplus assets from the portfolio, which leads to sparse vectors of good assets, and allows crossover to search for combinations of effective assets instead of permuting between all assets available.

The KEA also benefits from the use of Lamarckism. Due to Lamarckism the speed of convergence is increased significantly on the cardinality constrained problem instances, compare Fig. 12 to Fig. 6, and even in absence of cardinality constraints, compare Fig. 11 to Fig. 5.

With Lamarckism the difference between the crossover operators become less distinct, but corresponds again to the results presented in sec. IV-A.1. The intermediate crossover performs worse than discrete and BLX-α crossover for the standard encoding, but it catches up with increasing cardinality constraints, see Fig. 10. And there is still a very slight advantage for the BLX-α crossover compared to the discrete crossover. For the KEA on the other hand again no significant

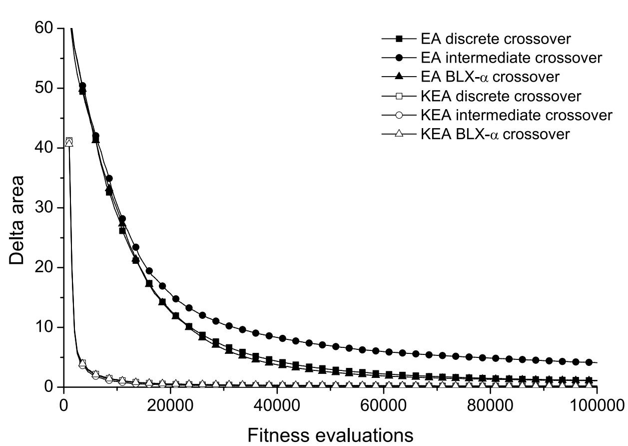

Fig. 11. Convergence behavior of Δ_{area} on the Hang Seng data set with Lamarckism, without cardinality constraints and without l_i and c_i constraints

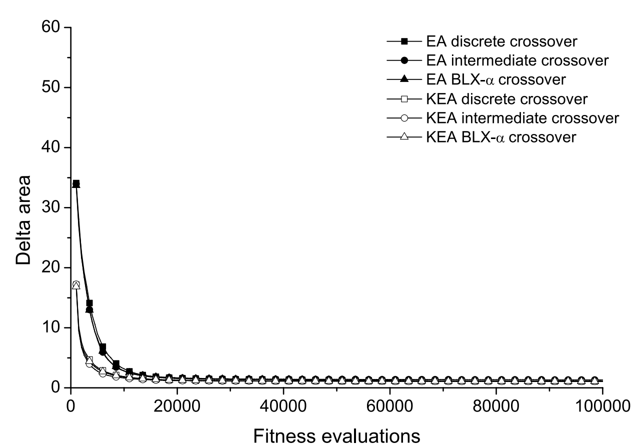

Fig. 12. Convergence behavior of Δ_{area} on the Hang Seng data set with Lamarckism, $K = 4$ and without l_i and c_i constraints

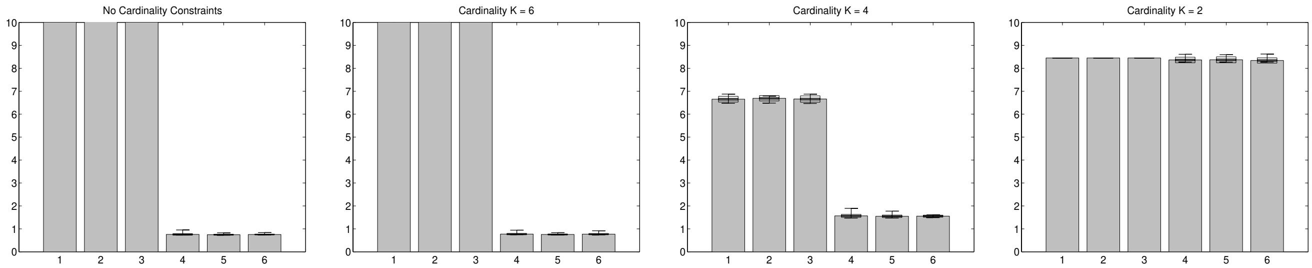

Fig. 13. Δ_{area} for the experiments on the Hang Seng data set with Lamarckism, $l_i = 0.1$ and $c_i = 0.02$ (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover)

distinctions can be made. Even the slightly higher variance of the intermediate crossover disappeared.

2) With Additional Constraints: The effect of Lamarckism leads to such a uniform convergence behavior of all three crossover operators that a reasonable comparison is no longer possible, see Fig. 13. But still there are some differences between the standard EA and the KEA, which are to be explained.

First, we want to discuss the behavior of the standard EA. Again as in sec. IV-A.2 the standard EA converges very reliably to suboptimal solutions depending on the cardinality. Without cardinality constraints and with or without Lamarckism the standard EA converges to the very same suboptimal Pareto front, compare Fig. 14 to Fig. 8. Again this is due to the implicit cardinality of $K = 10$ imposed by the buy-in threshold constraint $l_i = 0.1$.

With more restrictive cardinality constraints ($K < 10$) it becomes more complicated. Here the standard EA with Lamarckism also suffers from premature convergence, but finds much better results than without Lamarckism, compare Fig. 15 to Fig. 9. This is because the standard EA without Lamarckism searches not only the subspace of portfolio of cardinality K but also the weights of the assets in the portfolio are limited to values of $w_i \approx 1/K$, because the repair mechanism always chooses the K biggest values from $\mathbf{W}$ to be in the resulting portfolio, and normalizes them to $w_i \approx 1/K$. With Lamarckism on the other hand the cardinality of the portfolio is still limited to K but the weights for those assets selected can be effectively explored through mutation and crossover, since surplus the weights of surplus assets are set to zero.

Finally, the KEA is again not limited to such subspaces and explores the whole search space very effectively and also reliably, see Fig. 13. But with Lamarckism the KEA performs much better than the KEA without Lamarckism. While without cardinality constraint the difference is not as big, compare Fig. 14 to Fig. 8, the convergence speed is significantly increased in case of additional cardinality constraints, compare Fig. 15 to Fig. 9.

V. CONCLUSION

In this paper we were able to show that the new proposed hybrid encoding is able to solve the portfolio optimization problem more efficiently than the standard encoding based on a single real-valued vector of decision variables. This was shown on multiple problem instances and for several crossover operators. We were able to verify the positive effect of the hybrid encoding, by creating a similar effect for the standard EA through the application of Lamarckism on problems with

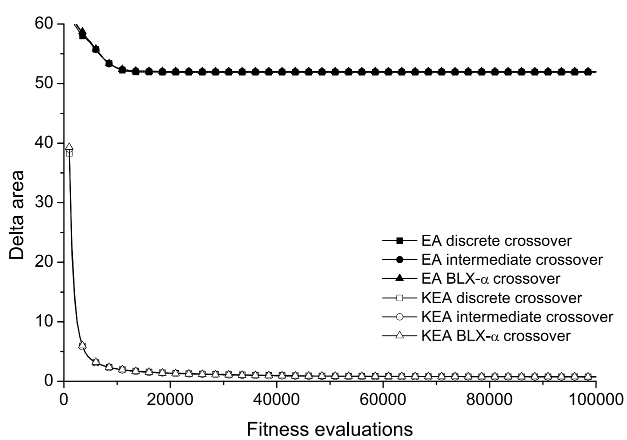

Fig. 14. Convergence behavior of Δ_{area} on the Hang Seng data set with Lamarckism, without cardinality constraints, $l_i = 0.1$ and $c_i = 0.02$

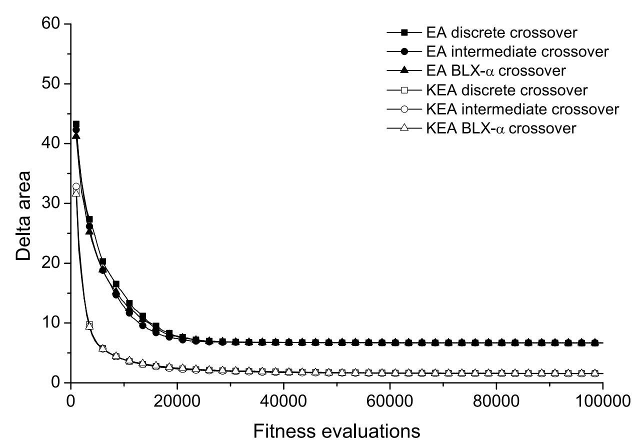

Fig. 15. Convergence behavior of Δ_{area} on the Hang Seng data set with Lamarckism, $K = 4$, $l_i = 0.1$ and $c_i = 0.02$ constraints

cardinality constraints and without additional real-world constraints.

We also examined the impact of the repair mechanism on the performance of the EA implementations. We showed that without Lamarckism and with additional real-world constraints like buy-in thresholds and roundlot constraints the standard EA fails due to the neutrality of the search space. The neutrality introduced through the repair mechanism misled the standard EA to search only a subspace of the true search space where the portfolios are of cardinality K and the assets used have weights of $w_i \approx 1/K$. While the standard EA was not able to explore the search space beyond this subspace, the more efficient hybrid encoding enabled the KEA to explore the full search space.

With Lamarckism on the other hand both EA implementations performed much better. But finally, on the problem instances with additional real-world constraints the KEA with Lamarckism outperformed the standard EA significantly, since the standard EA is still limited to a subspace of the true search space where the portfolios are of cardinality K.

Regarding the different crossover operators we showed that the intermediate crossover, which could be considered as an intermediate step between the discrete crossover and the BLX-α crossover, performed worst. And although the difference between discrete crossover and BLX-α crossover is less distinct, the BLX-α crossover performed better on some problem instances. The common element between the discrete crossover and the BLX-α crossover seems to be the diversity created through the crossover operators, by placing the offsprings on the edges of a hyper cube given by the parents. Unfortunately, the differences between the crossover operators became leveled out either due to the reduced subspace searched by the standard EA in case with additional real-world constraints or due to the high speed of convergence in case of the KEA with Lamarckism.

VI. Future Work

Our future work will concentrate on evaluating the performance of alternative MOEA implementations on the portfolio selection problem. We believe that the choice of the MOEA strategy will become crucial, if more real-world constraints are added like sector/industry constraints, immunization/duration matching and taxation constraints, which may increase the output dimension of the portfolio selection problem.

Another area of improvement could be the application of more sophisticated local search heuristics. There are numerous alternatives to the simple search for feasible solutions, but they have to be carefully evaluated regarding their ability to handle real-world constraints. Further we plan to extend our experiments to other portfolio models like for example the Black-Litterman model [3].

References

[1] S. Arnone, A. Loraschi, and A. Tettamanzi. A genetic approach to portfolio selection. *Neural Network World, International Journal on Neural and Mass-Parallel Computing and Information Systems*, 3:597–604, 1993.

[2] J. B. Beasley. OR-Library: distributing test problems by electronic mail. *Journal of the Operational Research*, 8:429–433, 1996.

[3] F. Black and R. Litterman. Global portfolio optimization. *Financial Analysts Journal*, pages 28–43, September-October 1992.

[4] T.-J. Chang, N. Meade, J. B. Beasley, and Y. Sharaiha. Heuristics for cardinality constrained portfolio optimization. *Computers and Operations Research*, 27:1271–1302, 2000.

[5] Y. Crama and M. Schyns. Simulated annealing for complex portfolio selection problems. Working paper GEMME 9911, Universit de Lige, 1999.

[6] K. Deb, S. Agrawal, A. Pratab, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, and H.-P. Schwefel, editors, *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 849–858, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

[7] L. J. Eshelman and J. D. Schaffer. Real-coded genetic algorithms and interval-schemata. In L. D. Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 187–202. Morgan Kaufmann, San Mateo, CA, 1993.

[8] F. Herrera, M. Lozano, and J. L. Verdegay. Tackling real-coded genetic algorithms: Operators and tools for behavioural analysis. *Artificial Intelligence Review*, 12(4):265–319, 1998.

[9] J. Knowles and D. Corne. The pareto archived evolution strategy: A new baseline algorithm for pareto multiobjective optimisation. In P. J. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzala, editors, *Proceedings of the Congress on Evolutionary Computation*, volume 1, pages 98–105, Mayflower Hotel, Washington D.C., USA, 1999. IEEE Press.

[10] A. Loraschi and A. Tettamanzi. An evolutionary algorithm for portfolio selection in a downside risk framework. *Working Papers in Financial Economics*, 6:8–12, June 1995.

[11] A. Loraschi, A. Tettamanzi, M. Tomassini, and P. Verda. Distributed genetic algorithms with an application to portfolio selection problems. In D. W. Pearson, N. C. Steele, and R. F. Albrecht, editors, *Artificial Neural Networks and Genetic Algorithms*, pages 384–387, Wien, 1995. Sringer.

[12] H. M. Markowitz. Portfolio selection. *Journal of Finance*, 1(7):77–91, 1952.

[13] H. M. Markowitz. *Portfolio Selection: efficient diversification of investments*. John Wiley & Sons, 1959.

[14] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolutions Programs*. Springer-Verlag, New York, 1992.

[15] H.-P. Schwefel. *Numerical Optimization of Computer Models*. John Wiley & Sons, Chichester, U.K., 1977.

[16] N. Srinivas and K. Deb. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation*, 2(3):221–248, 1994.

[17] F. Streichert, H. Ulmer, and A. Zell. Evolutionary algorithms and the cardinality constrained portfolio selection problem. In D. Ahr, R. Fahrion, M. Oswald, and G. Reinelt, editors, *Operations Research Proceedings 2003, Selected Papers of the International Conference on Operations Research (OR 2003), Heidelberg, September 3-5, 2003*. Springer, 2003.

[18] D. L. Whitley, V. S. Gordon, and K. E. Mathias. Lamarckian evolution, the baldwin effect and function optimization. In Y. Davidor, H.-P. Schwefel, and R. Männer, editors, *Parallel Problem Solving from Nature – PPSN III*, pages 6–15, Berlin, 1994. Springer.

[19] A. H. Wright. Genetic algorithms for real parameter optimization. In G. J. Rawlins, editor, *Foundations of genetic algorithms*, pages 205–218. Morgan Kaufmann, San Mateo, CA, 1991.

[20] E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. Technical Report 103, Gloriastrasse 35, CH-8092 Zurich, Switzerland, 2001.

[21] E. Zitzler and L. Thiele. Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999.