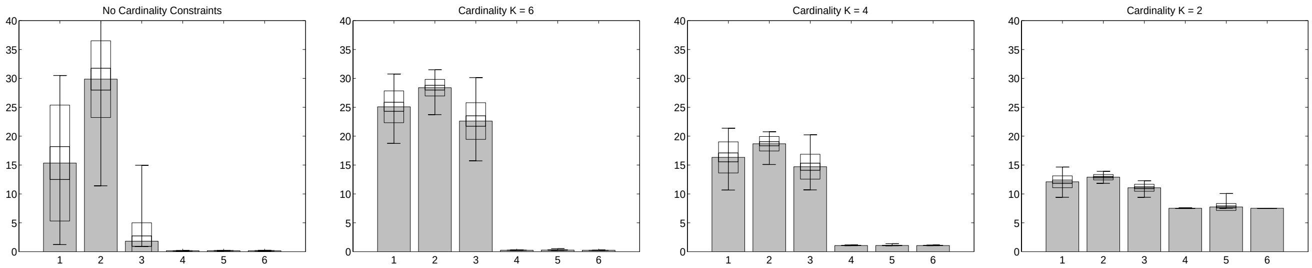

Fig. 4. Δ_{area} for the experiments on the Hang Seng data set without additional constraints l_i and c_i (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover)

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.

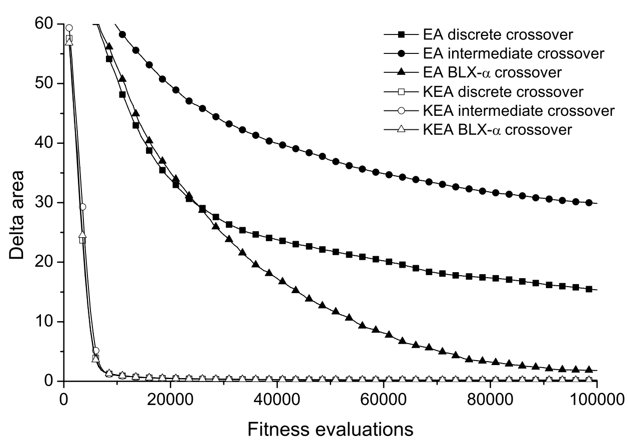

Fig. 5. Convergence behavior of Δ_{area} on the Hang Seng data set without cardinality constraints and without l_i and c_i constraints

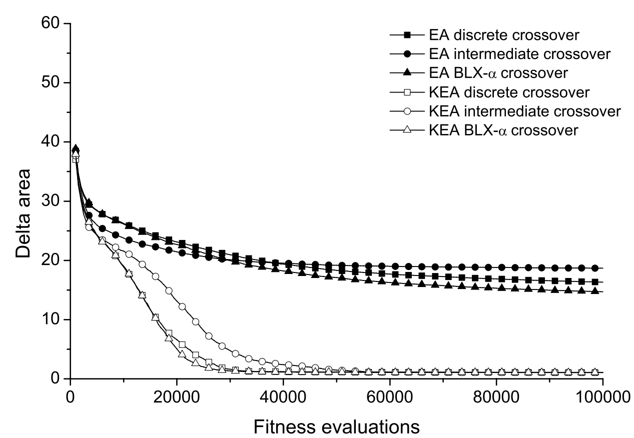

Fig. 6. Convergence behavior of Δ_{area} on the Hang Seng data set with $K = 4$ and without l_i and c_i constraints

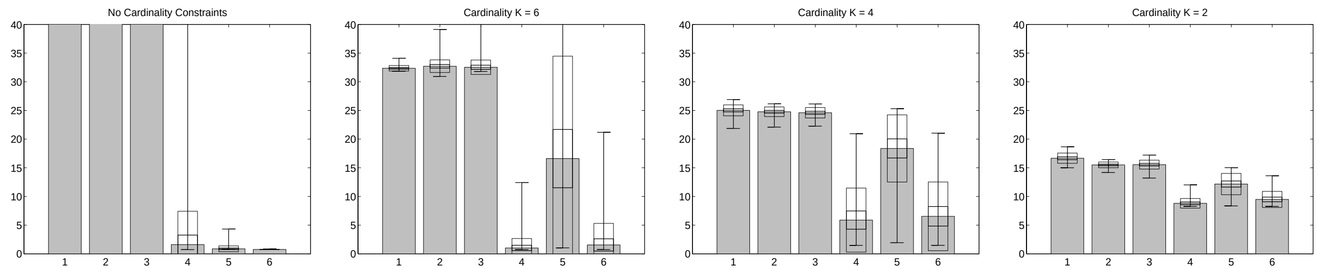

Fig. 7. Δ_{area} for the experiments on the Hang Seng data set with $l_i = 0.1$ and $c_i = 0.02$ (1: EA with discrete crossover; 2: EA with intermediate crossover; 3: EA with BLX-α crossover 4: KEA with discrete crossover; 5: KEA with intermediate crossover; 6: KEA with BLX-α crossover)

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.

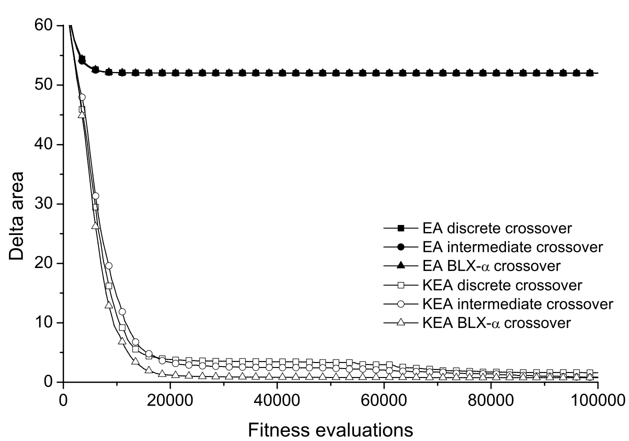

Fig. 8. Convergence behavior of Δ_{area} on the Hang Seng data set without cardinality constraints, $l_i = 0.1$ and $c_i = 0.02$ constraints

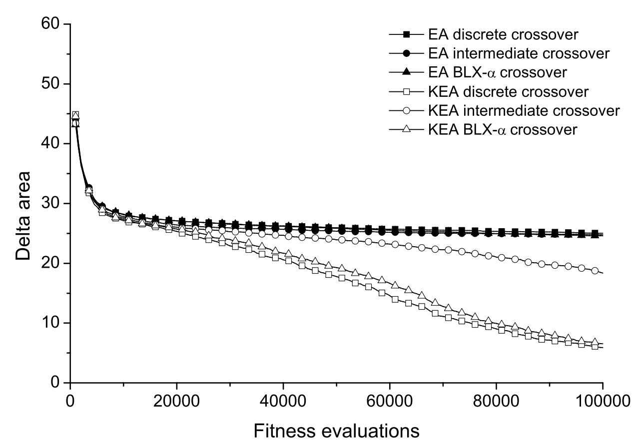

Fig. 9. Convergence behavior of Δ_{area} on the Hang Seng data set with $K = 4$, $l_i = 0.1$ and $c_i = 0.02$ constraints