

# Chapter 12

## Memetic Algorithms and Complete Techniques

Carlos Cotta, Antonio J. Fernández Leiva, and José E. Gallardo

### 12.1 Introduction

As mentioned in previous chapters in this volume, metaheuristics (and specifically MAs) have a part of their *raison d'être* in practically solving problems whose resolution would be otherwise infeasible by means of other non-heuristic approaches. Such alternative non-heuristic approaches are complete methods that –unlike heuristics– do guarantee that the deviation from optimality of the solution they will provide is somehow bounded (and as a particular case, that the optimal solution will be found). These methods are eventually limited by the curse of dimensionality, yet they may still constitute a very interesting resource either from the application point of view, or from the lessons that can be learnt from them. Indeed, in some sense these approaches could be considered complementary to metaheuristics rather than mere “rivals”. Even more so in the case of MAs, whose philosophy has been since its inception much more flexible and integrative rather than dogmatic or exclusive.

This said, despite the eclosion of metaheuristics as powerful optimization techniques during the 80s and 90s, inter-breeding between the fields of provably problem-solving and heuristic problem-solving was relatively limited until the last decade (some seminal works dating back from the mid 90s – e.g., [165]). The last years however have witnessed a remarkable increase in the number of works trying to combine ideas from these two areas. Certainly, MAs have also played an important role in this cross-fertilization of search paradigms. Along this chapter we will review some of the lines of research that have emerged in this regard. To this end, we will begin by briefly revisiting complete techniques to highlight their strengths and weaknesses, and what they have to offer to metaheuristics. Subsequently, we will outline some of the efforts that have been made in the literature to classify

---

Carlos Cotta · Antonio J. Fernández Leiva · José E. Gallardo  
Dept. Lenguajes y Ciencias de la Computación, Universidad de Málaga,  
ETSI Informática, Campus de Teatinos, 29071 Málaga, Spain  
e-mail: {ccottap, afdez, pepeg}@lcc.uma.es

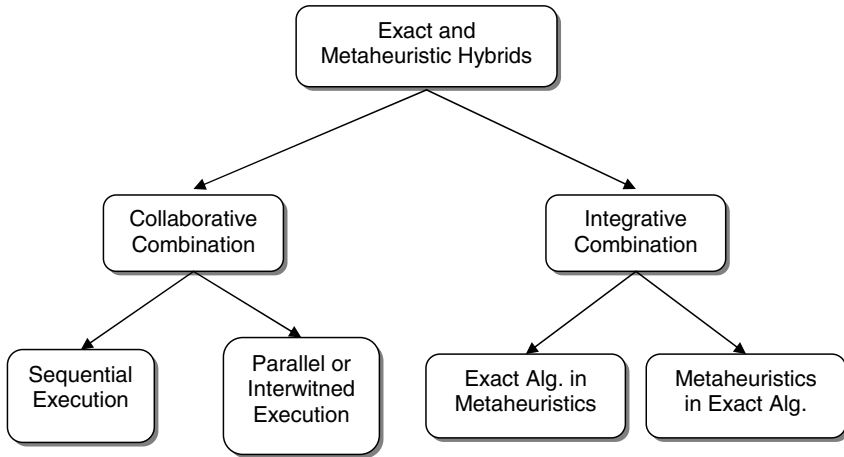
hybrid approaches. Although these classifications are usually general and intended to cover more than just complete-heuristic combinations, they will provide a framework within which actual combinations of MAs with exact techniques (or from the broad interpretation of memetic algorithms, MAs incorporating exact techniques) can be studied. This will be done in Sections 12.4 and 12.5.

## 12.2 Background

Complete techniques are those whose results can be proved to be at bounded distance from the optimum. From a very general point of view, these techniques can be further subdivided into techniques that guarantee finding the optimal solutions, i.e., exact techniques, and techniques that only provide a fixed or adjustable bound (that is, a bound that can be reduced by spending more computational effort), i.e., approximation techniques. Curiously, and this is something that may be worth some further analysis from a sociological and/or philosophical point of view, the community of researchers working on approximation theory has been traditionally more skeptical with respect to the value of metaheuristic optimization. Conversely, it is also true that the usefulness of approximation algorithms has not been always appreciated by the metaheuristic community, in part due to the inherent limitations of the former in many practical contexts – see for example [224] for a glimpse of the computational complexity of PTAS (a polynomial time approximation scheme, probably one of the jewels of the crown in approximation theory) for several common problems.

Focusing thus on exact techniques, such as for example branch and bound [507], dynamic programming [57], branch and cut [654], etc. these are characterized by the fact that they guarantee finding optimal solutions at the cost of a non-polynomial growth of computation time (and often memory consumption too). Their limitations are those emanating from the theory of computational complexity, such as the conspicuous P vs NP question. It must be noted however that such classical (unidimensional) hardness characterizations are not necessarily correlated with practical performance. A much more interesting characterization can be obtained from the field of parameterized complexity [221], in which hardness is approached from a multidimensional perspective, factoring out some *parameter(s)* from the input and trying to isolate the problem's difficulty in them. If this can be done –formally, if the complexity of the problem can be shown to be polynomially related to the input size (once the parameter is factored out), and the degree of the polynomial is unrelated to the value of the parameter– the problem is said to be fixed-parameter tractable (FPT). FPT problems can be solved for small values of the parameter using the arsenal developed by the parameterized-complexity community – e.g., [667]. Hard problems can be nevertheless detected from a parameterized perspective, and for such problems metaheuristics are fully in order.

There are many ways in which the hybridization of metaheuristics in general (and MAs in particular) with exact techniques can be fruitful: exact techniques can, for example, reduce their resource consumption if they obtain valuable input from metaheuristics (e.g., improved bounds); on the other hand, metaheuristics routinely



**Fig. 12.1.** Puchinger and Raidl’s classification of exact-heuristic hybrid algorithms.

use search mechanisms –recombination, mutation, etc.– in which exact techniques can play an important role to intensify the search. Furthermore, hybridization of MAs and exact techniques can be defined at several nested levels thus providing multiple ways of boosting each other’s performance. In the following we will survey some successful hybridization models reported in the literature along these lines just depicted. Subsequently, we will overview several approaches to classify these hybrid models.

### 12.3 Classification of Hybridization Approaches

Several taxonomical attempts have been proposed to classify hybrid optimization algorithms. For example, Talbi [872] proposed a mixed hierarchical-flat classification scheme. The hierarchical component captured the structure of the hybrid, whereas the flat component specified the features of the algorithms involved in the hybrid. More precisely, the hierarchical portion of the taxonomy firstly distinguished between low-level (a given function of a metaheuristic is replaced by another metaheuristic) and high-level (combined algorithms are self-contained) hybridization. Secondly, it was distinguished between relay hybridization (a set of metaheuristics is applied in a pipeline fashion) and teamwork hybridization (cooperative optimization models). Cotta [147] proposed another related taxonomy with the dichotomy *strong* vs. *weak* as its root. This distinction referred to whether problem-knowledge was placed in the core of the algorithm, affecting its internal components (e.g., representation and/or genotype-phenotype mapping, operators, etc.), or in the combination of different search algorithms that retained their identity. This terminology is consistent with the classification of problem-solving strategies in artificial intelligence as *strong* and *weak* methods [595].

A much more interesting classification for the purposes of this chapter is that proposed by Puchinger and Raidl [741]. This classification is specifically intended for exact-metaheuristic combinations, and establishes two main categories for such hybrid algorithms:

- *Collaborative combinations*, where an exact algorithm and a metaheuristic method exchange some information, but none of them are part of the other, and
- *Integrative combinations*, where one technique is a subordinate of the other, i.e., there is a *master* algorithm that uses the other one.

These two categories can be further refined depending on the particular of the combination as shown in Figure 12.1. Thus, a collaborative combination can be sequential or parallel/intertwined, depending on how the control flow passes from one algorithm to the other. Similarly, an integrative combination can be subdivided in models in which an exact technique plays the role of master (i.e., the metaheuristic is embedded in an exact technique), and models in which the opposite is true.

As mentioned before, this latter classification fits nicely into context of this chapter, so we will consider it in order to survey existing hybrid approaches combining MAs and exact techniques.

## 12.4 Integrative Combinations

One basic form of integrative collaboration consists of endowing a memetic algorithm with an exact technique (ET) so that this ET is a subordinate of the MA. The most common implementation consists of combining an EA with a procedure to perform a complete local search (which can consider the whole neighborhood and in this sense can be viewed as an exact technique). This is usually done after evaluation, although it must be noted however that the integration does not simply reduce itself to this particular scheme. In fact, the purpose of using an ET inside a MA is to provide specific knowledge that can help to a better optimization process. For instance, Algorithm 22 shows a general picture of where an ET can be incorporated inside an MA.

As it can be seen, during the initialization of the population some complete method may be used to generate high quality initial solutions. Of course, this complete method may only consider a subset of the search space, a relaxed version of the problem, or may perform just a truncated search, since otherwise the problem would just be solved at that stage (not to mention the computational cost). An example of relaxed initialization using complete techniques can be found in [148], where a backtracking algorithm is used to create feasible initial solutions for a protein structure prediction problem (thus relaxing optimality to mere feasibility). Another related approach will be discussed in next subsection in the context of collaborative models, and considers a variant of a B&B algorithm—namely beam search—to initialize the population of a MA with the aim of improving its performance. In addition to this an exhaustive LS could be applied to improve the individuals generated initially with the aim of providing a first population of better quality. This procedure

---

**Algorithm 22.** Pseudocode of a basic MA based on an integrative collaboration with an exact technique ET

---

```

1 for  $i \in \{1, \dots, \text{POPULATION SIZE}\}$  do
2    $pop[i] \leftarrow \text{RANDOM-SOLUTION}()$ ;
3   if  $\text{Rand}[0, 1] < p_{ET}$  then // ET is applied with probability  $p_{ET}$ 
4     EXACT-TECHNIQUE( $pop[i]$ ); // Usually ET = Local
      Improvement
5   endif
6 endfor
7  $i \leftarrow 0$ ;
8 while  $i < \text{MaxEvals}$  do
9   RANK-POPULATION( $pop$ ); // sort population according to
      fitness
10   $parent_1 \leftarrow \text{SELECT}(pop)$ ;
11  if  $\text{Rand}[0, 1] < p_X$  then // recombination is done
12     $parent_2 \leftarrow \text{SELECT}(pop)$ ;
13     $child \leftarrow \text{RECOMBINE}(parent_1, parent_2)$ ; // RECOMBINE might be an
      Exact Technique
14  else
15     $child \leftarrow parent_1$ ;
16  endif
17   $child \leftarrow \text{MUTATE}(child, p_M)$ ; //  $p_M$  is the mutation probability
      per gene
18  if  $\text{Rand}[0, 1] < p'_{ET}$  then // ET is applied
19    EXACT-TECHNIQUE( $child$ ); // Usually Local Improvement
      applied here
20  endif
21   $pop[\mu] \leftarrow child$ ; // replace worst
22 endw
23 return best solution in  $pop$ ;

```

---

is intimately related to the idea of local branching by Fischetti and Lodi [263], and to Congram's Dynasearch [138, 139].

Another proposal that can be devised from the general schema shown above is the use of an ET as a recombination operator. Recombination or mutation operators can be intelligently designed so that specific problem knowledge is used in order to improve the offspring. For instance, Cotta *et al.* [165] used a problem-specific B&B approach for the Travelling Salesman Problem based on 1-trees and the Lagrangean relaxation [910], to build a hybrid recombination operator. More precisely, the B&B was used in order to build the best possible tour within the (Hamiltonian) subgraph defined by the union of edges in the parents. This recombination procedure was costly, but provided better results than blind edge recombination. This model was later extended to a more general operator termed dynamically optimal recombination (DOR) [164]. The term refers to the *dynastic potential*, which in the framework of Forma Analysis [750] denotes the set of children attainable from a certain set of parents. DOR thus consists of finding the best children in this dynastic potential, i.e., that with the best

combination of parental features. This is done by “intelligently” exploring this set, using an adequate complete algorithm, check e.g. [163].

Related to the previous approach, [295] presented a memetic algorithm, embedded with tabu search, for weighted constraint satisfaction problems (see next section for a more detailed discussion of this kind of problems) in which bucket elimination (BE) [202] is used as a mechanism for recombining solutions, providing the best possible child from the parental set. BE is an exact technique related to dynamic programming which based on variable elimination and is commonly used for solving constraint satisfaction problems. This algorithm, with another collaborative proposals, was applied to the resolution of the maximum density still life problem, a hard constraint optimization problem based on Conway’s game of life.

Additionally, problem knowledge can be incorporated in the genotype to phenotype mapping present in many MAs, like when repairing an infeasible solution. This technique is used, for instance, in the MA designed by Chu and Beasley [127] for the multidimensional 0-1 knapsack problem. The use of complete techniques, again relaxed or truncated, can be here considered as well, check, e.g., [148] for an example of using backtracking to repair infeasible solutions.

Another place where an exact method can be particularly useful when used inside an MA is in the optimization of problems where different representations are considered. In these cases, an exact technique can be specifically useful in the codification-decodification phase. For instance, Puchinger and Raidl [742] represented another attempt to incorporate exact methods in metaheuristics. This work considered different heuristics algorithms for a real world glass cutting problem and a combined GA and B&B approach was proposed. The GA used an order-based representation that was decoded with a greedy heuristic. Incorporating B&B in the decoding for occasionally (with a certain probability) locally optimizing subpatterns turned out to increase the solution quality in a few cases.

Note also that applying always the ET in each generation of the MA (or initially on each individual in the initial population) is not always the best option (as shown in [858] for the application of LS on each generated new individual). For instance, if one considers LS as the technique to embed inside a MA, partial Lamarckianism [396], namely applying local search only to a fraction of individuals, can result in better performance. These individuals to which local search will be applied can be selected in many different ways [665]. Thus, LS can be applied to improve the individual with certain probability  $p_{LS}$ ; in case of application, the improvement uses up a number of  $LSevals$  evaluations (or in the case of specific local search such as HC, until it stagnates, whatever comes first). It is easy to extrapolate these results from the use of LS to any ET embedded in a MA.

In general, the underlying idea of this kind of integration is to combine the intensifying capabilities of the embedded ET method, with the diversifying features of MA, i.e., the population will spread over the search space providing starting points for a deeper (probably local) exploration. As generations go by, promising regions will start to be spotted, and the search will concentrate on them. Ideally, this combination should be synergistic, providing better results than either the MA or the ET by themselves. Regarding this issue, one can find in the literature a number of

proposals that explore the intensification/diversification balance within the memetic algorithm. Some works lean towards a more explorative combination, by using a blind recombination operator in the MA whereas other models incorporate an intense exploration of the dynamic potential of the solutions being recombined.

The other possibility for integrative combinations is to incorporate a metaheuristic into an exact algorithm. One example is the hybrid algorithm combining Genetic Algorithms and Integer Programming B&B approaches to solve MAX-SAT problems described in [287]. This hybrid algorithm gathered information during the run of a linear programming-based B&B algorithm, and used it to build the population of an EA population. The EA was eventually activated, and the best solution found was used to inject new nodes in the B&B search tree. The hybrid algorithm was run until the search tree was exhausted, and hence it is an exact approach. However, in some cases it expands more nodes than the B&B algorithm alone.

## 12.5 Collaborative Combinations

As mentioned in Section 12.3, the class of *collaborative combinations* includes hybrid algorithms which exchange information, but such that none of them is a subordinate of the other. Two subcases can be here considered in order to execute both algorithms:

- *Sequential execution*, in which one of the algorithms is completely executed before the other. Examples of this group are those in which one of the techniques can act as a kind of preprocessing for the other or those where the result of one algorithm can be used as data to initialize the other.
- *Parallel or intertwined execution*, where both techniques are executed simultaneously, either in parallel (i.e., running at the same time on different processors) or in an intertwined way by alternating between both algorithms.

As an example of sequential combinations we can cite the work of Klau et al. [469], in which a branch and cut algorithm is used analogously to the idea of dynamically optimal recombination mentioned in previous section, to combine the final population provided by a MA. This hybrid algorithm is applied to the prize-collecting Steiner tree problem. We will focus here on hybrid collaborative techniques in the second group. MAs and B&B techniques can be integrated by way of a *direct collaboration*, so that both techniques work alone in parallel (i.e., both processes perform independently) at the same level. Under this scheme, both processes will share the incumbent solution to the problem being solved. Whenever one of the algorithms finds a better approximation, it can update the solution. Two straightforward ways of obtaining a benefit of this parallel execution are [165]:

- The B&B algorithm can use the lower bound provided by the MA to purge its problem queue. Problems whose upper bound are smaller than the one obtained by the MA cannot improve the incumbent solution and can be safely removed.

- The B&B algorithm can provide information about more promising regions of the search space into the MA population. The aim of this process is to guide the MA search towards these promising regions of the search space.

Several implementations of these schemes are possible. For example, Cotta *et al.* [165] proposed a collaborative approach such as the one described above for the TSP. Puchinger and Raidl [740] consider the parallel combination of a MA and a branch and cut algorithm for the multidimensional knapsack problem. The MA provides improved bounds to the branch and cut algorithm, and the latter provides both new best-so-far solutions and the corresponding dual variable values, to be used for repairing and local search. More recently, Gallardo *et al.* [294] defined a hybrid algorithm that starts by running a MA (with a randomly initialized population) in isolation, so that a first approximation to the solution is obtained. This initial solution is later used by a B&B algorithm to purge its problem queue. As it can be seen, no information from the B&B algorithm was used in this first execution of the MA. In a subsequent phase, the B&B algorithm starts its execution. New solutions found by the B&B are incorporated into the MA population (by replacing the worst individual). Whenever a new solution is found, the B&B phase is paused and the MA is run to stabilization. In addition, pending nodes in the B&B queue are incorporated into the MA population periodically. The intention of this transfer is to direct the MA to these regions of the search space, that represent the subset of the search space still unexplored by the Branch and Bound. In this way, the MA is used for finding probably good solutions in those regions. Upon finding an improved lower bound (or upon stabilization of the MA if no improvement is found), B&B is resumed. This process is repeated until the search tree is exhausted, or a time limit is reached. One interesting property of this hybrid algorithm is that it acts as an anytime algorithm, providing both a quasi-optimal solution, and an indication of the maximum distance to the optimum. In [293, 294] this implementation schema is used to tackle large instances of the multidimensional knapsack problem. Experimental results showed that the hybrid approach can provide high quality results, better than those obtained by the MA and B&B on their own.

An alternative implementation of the previous model consist on using beam search (BS)[46] instead of B&B. This is an incomplete derivative of the later and acts thus as an heuristic method. In essence, BS extends every partial solution from a set  $\mathcal{B}$  (called the *beam*) in at most  $k_{ext}$  possible ways, generating a new beam. When all solutions in  $\mathcal{B}$  have been extended, the algorithm reduces the new beam by selecting the best up to  $k_{bw}$  (called the *beam width*) solutions and proceeds. A very interesting feature of this heuristic is that it extends in parallel a set of different partial solutions in several possible ways. For this reason, it can be used to provide periodically diverse promising partial solutions to a population based search method such as a MA. A general description of the resulting hybrid algorithm is given in Algorithm 23.

The beam search part of the algorithm can be iterated for each level of the search tree that corresponds to the problem at hand. The hybrid algorithm starts by executing this process for an initial number of levels (parameter  $l_0$  of the algorithm). Subsequently, both parts of the hybrid algorithm are alternatively executed until a



**Algorithm 23.** Beam Search + MA hybrid algorithm

---

```

1 for  $l_0$  levels do run BS;
2 repeat
3   select popsize nodes from problem queue;
4   initialize MA population with selected nodes;
5   run MA;
6   if MA solution better than BS solution then
7     | let BS solution  $\leftarrow$  MA solution;
8   endif
9   for  $l$  levels do run BS ;
10 until timeout or tree-exhausted ;
11 return BS solution;

```

---

termination condition is reached. Similar to the first implementation, for every execution of the MA, its population is initialized using the nodes in the BS queue. As the size of the BS queue is usually larger than the MA population size, a criteria, such as selecting the best nodes according to some measure of quality or selecting a subset that provides high diversity, has to be used in order to select a subset from the queue. Nodes in the BS queue represent partial solutions in which some genes are fixed but others are indeterminate, so they must first be converted to full solutions in a problem dependent way. This must be considered when instantiating the general template for different combinatorial problems. This kind of collaborative integration of Beam Search and MAs has been used to tackle different combinatorial optimization problems. In [297] the hybrid algorithm was experimentally evaluated on the multidimensional 0-1 knapsack problem and on the shortest common super-sequence problem, a NP-hard classical problem from the realm of string analysis. For both problems, it was shown the benefits of using the hybrid approach when compared to the constituents algorithms. Additionally, an analysis of the dynamics and sensitivity on different parameters of the algorithm was carried out. In [298], the hybrid algorithm was applied to the inference of phylogenetic trees, an important problem in Systematic Biology, that aims to represent the evolutionary history for a collection of organisms. That work focused in the ultrametric model for phylogenetic inference. A robust setting for the different parameters of the algorithm was determined, and the hybrid algorithm was experimentally shown to also be synergic for this problem.

A related hybridization model has been defined in [299] for weighted constraint satisfaction problems (WCSP) [795]. A WCSP is a constraint satisfaction problem (CSP) in which preferences among solutions can be expressed. Formally, a WCSP can be defined by a tuple  $(\mathcal{X}, \mathcal{D}, \mathcal{F})$ , where  $\mathcal{D} = \{D_1, \dots, D_n\}$  is a set of *finite domains*,  $\mathcal{X} = \{x_1, \dots, x_n\}$  is a set of variables taking values from their finite domains and  $\mathcal{F}$  is a set of *cost functions* (also called *soft constraints* or *weighted constraints*) used to declare preferences among possible solutions. Each  $f \in \mathcal{F}$  is defined over a subset of variables,  $var(f) \subseteq \mathcal{X}$ , called its *scope*. The objective function  $F$  –to be minimized– is defined as the sum of all functions in  $\mathcal{F}$ , i.e.,  $F = \sum_{f \in \mathcal{F}} f$ . WCSP were tackled using a algorithmic model based on the hybridization of MAs with




Fig. 12.2. Schematic description of the multilevel hybrid algorithm.

exact techniques at two levels: within the MA (as an embedded operator), and outside it (in a cooperative model). Figure 12.2 depicts the different components of the algorithm and their relationships. The first level of hybridization has already been described in Section 12.4, so we will describe here the second level, in which the MA cooperates with a beam search algorithm that further uses the technique of mini-buckets as a lower bound.

---

**Algorithm 24.** Hybrid algorithm for a WCSP

---

```

1  $sol \leftarrow \infty$ ;
2  $\mathcal{B} \leftarrow \{ () \}$ ;
3 for  $i \leftarrow 1$  to  $n$  do
4    $\mathcal{B}' \leftarrow \{ \}$ ;
5   for  $s \in \mathcal{B}$  do
6     for  $a \in D_i$  do
7        $\mathcal{B}' \leftarrow \mathcal{B}' \cup \{ s \cdot (x_i = a) \}$ ;
8     endfor
9   endfor
10   $\mathcal{B} \leftarrow$  select best  $k_{pw}$  nodes from  $\mathcal{B}'$ ;
11  if  $i \geq k_{MA}$  then
12    initialize MA population with best popsize nodes from  $\mathcal{B}'$ ;
13    run MA;
14     $sol \leftarrow \min(sol, \text{MA solution})$ ;
15  endif
16 endfor
17 return  $sol$ ;

```

---

The proposed hybrid algorithm, that executes BS and the MA in an interleaved way, is depicted in Algorithm 24. Here, a (possibly partial) solution for a WCSP instance is represented by a vector of variables  $s = (x_1, x_2, \dots, x_i)$ ,  $i \leq n$ , where  $s \cdot (x_i = a)$  stands for the extension of partial solution  $s$  by assigning value  $a$  to its  $i$ -th variable. The hybrid algorithm proceeds by constructing a search tree, so that its leaves are complete solutions to the problem and internal nodes at level  $i$  represent solutions that are partially specified up to the  $i$ -th variable. The algorithm traverses

this tree heuristically in a breadth first way using a BS algorithm that only maintains the best  $k_{bw}$  nodes at each level of the tree. During each iteration of BS (lines 5-16), a variable is assigned for every solution in the beam (line 8). The interleaved execution of the MA starts only when partial solutions in the beam have at least  $k_{MA}$  variables (line 12). For each iteration of BS, the best *popsize* solutions in the beam are selected with the purpose of initializing the population of the MA (line 13). The solution provided after the execution of the MA is used to update the incumbent solution (*sol*), and this process is iterated until the search tree is exhausted.

The performance of this algorithm will depend on the quality of the heuristic function used to estimate partial solutions (line 11). In order to compute tight, yet computationally inexpensive, lower bounds for the remaining part of the solutions, the technique of mini-buckets (MB) can be used. As described by Kask and Dechter[445], the intermediate functions created by applying the MB scheme can be used as a general mechanism to compute heuristic functions that estimate the best cost of yet unassigned variables in partial solutions. This can be achieved by running MB as a preprocessing stage. The set of augmented buckets computed during this process can be used as estimations of the best cost extension to partial solutions (check [445] for details).

In [296, 299], such a multilevel algorithm was used to tackle the Maximum Density Still Life Problem, a hard constrained problem defined in the context of John Conway's game of life. The resulting algorithm was able to find optimal solutions for currently solved instances of the problem in considerable less time than state-of-the-art approaches. Additionally, it was able to find new best known solutions for very large instances whose exact solutions are yet unknown.

## 12.6 Conclusions

Throughout this chapter we have surveyed existing work on MAs that incorporate at some level a complete technique. Several notes have to be done here. Notice firstly from a 'terminological' point of view that many evolutionary techniques hybridized with complete techniques can be considered memetic regardless of whether a classical trajectory-based local search algorithm is also used or not. For example, in an evolutionary algorithm that used an exact technique for recombination, the latter could be regarded as a generalized local-search operator working on set of solutions rather than on single solutions, and using a neighborhood composed of all solutions in the corresponding dynastic potential. This is also related to the so-called crossover hill-climbing idea defined in [536] for continuous optimization.

From a practical point of view, this kind of hybrid approaches must carefully control the computational complexity of the problems submitted to complete search. This draws again a connection to parameterized complexity by noting that this complexity is typically related to some structural parameter of the problem (e.g., a higher similarity of the parents during exact recombination reduces the size of the dynastic potential, thus making its exploration more amenable in principle). Even though no efficient (in the FPT sense) algorithmic resolution were available for the problem at

hand, the combinatorial explosion could be kept within acceptable levels by checking these parameters and resorting to other approaches (truncated exact search, fast heuristic search, or even a blind procedure) if the possibility of a prohibitive computational cost cannot be excluded prior to a certain invocation of the complete method.

The any-time nature of MAs has to be considered as well. Whether a hybrid approach including complete techniques is itself complete or not, it is very important that it provides better and better solutions for any increasing computational budget allowed. This is not always possible within the context of complete techniques, e.g., a B&B algorithm using a best-first policy may exhaust its allotted time and/or memory without producing a single feasible solution. On the other hand, the very same B&B algorithm using a LIFO policy may quickly provide a solution but take a long time to improve it. MAs are however ideal for anytime search, and this can be exploited in a synergistic combination. Note for example that a parallel collaborative model using a MA and an exact technique may end up providing both an upper and a lower bound for the optimal solution, and the higher the computational budget available, the tighter these bounds will be.

**Acknowledgements.** This work is supported by Spanish MICINN under project NEMESIS (TIN2008-05941) and Junta de Andalucía under project TIC-6083.