Control Abstractions for Local Search

Pascal Van Hentenryck and Laurent Michel

Department of Computer Science
Brown University
Providence, Rhode Island 02912

CS-03-06
March 2003

Control Abstractions for Local Search

Pascal Van Hentenryck[1] and Laurent Michel[2]

[1] Brown University, Box 1910, Providence, RI 02912
[2] University of Connecticut, Storrs, CT 06269-3155

Abstract. COMET is an object-oriented language supporting a constraint-based architecture for local search through declarative and search components. This paper proposes three novel and lightweight control abstractions for the search component, significantly enhancing the compositionality, modularity, and reuse of COMET programs. These abstractions, which includes events and checkpoints, rely on first-class closures as the enabling technology. They are especially useful for expressing, in a modular way, heuristic and meta-heuristics, unions of heterogeneous neighborhoods, and sequential composition of neighborhoods.

1 Introduction

Historically, most research on modeling and programming tools for combinatorial optimization has focused on systematic search, which is at the core of branch & bound and constraint satisfaction algorithm. It is only recently that more attention has been devoted to programming tools for local search and its variations (e.g., [6, 26, 23, 11, 14, 25]).

COMET [13] is a novel, object-oriented, programming language specifically designed to simplify the implementation of local search algorithms. Comet supports a constraint-based architecture for local search organized around two main components: a declarative component which models the application in terms of constraints and functions, and a search component which specifies the search heuristic and meta-heuristic. Constraints, which are a natural vehicle to express combinatorial optimization problems, are *differentiable objects* in COMET: They maintain a number of properties incrementally and they provide algorithms to evaluate the effect of various operations on these properties. The search component then uses these functionalities to guide the local search using multidimensional, possibly randomized, selectors and other high-level control structures. The architecture enables local search algorithms to be high-level, compositional, and modular. It is possible to add new constraints and to modify or remove existing ones, without having to worry about the global effect of these changes. COMET also separates the modeling and search components, allowing programmers to experiment with different search heuristics and meta-heuristics without affecting the problem modeling. This separation of concerns give COMET some flavor of aspect-oriented programming [9] and feature engineering [24], since constraints represent and maintain properties across a wide range of objects. COMET has been applied to many applications and can be implemented to be competitive with tailored algorithms, primarily because of its fast incremental algorithms [13].

This paper focuses on the search component and aims at fostering the compositionality, modularity, and genericity of COMET. It introduces three novel control abstractions whose main benefit is to separate, in the source code, components which are usually presented independently in scientific papers. Indeed, most local search descriptions cover the neighborhood, the search heuristic, and the meta-heuristic separately. Yet typical implementations of these algorithms exhibit complex interleavings of these independent aspects and/or require many intermediary classes and/or interfaces. The resulting code is opaque, less extensible, and less reusable. The new control abstractions address these limitations and reduce the distance between high-level descriptions and their implementations.

The first abstraction, *events*, enables programmers to isolate the search heuristic from the meta-heuristic, as well as animation code from the modeling and search components. The second abstraction, *neighbors*, aims at expressing naturally unions of heterogeneous neighborhoods, which often arise in complex routing and scheduling applications. It allows to separate the neighborhood definition from its exploration, while keeping move evaluation and execution textually close. The third abstraction, *checkpoints*, simplifies the sequential composition of neighborhoods, which is often present in large-scale neighborhood search.

These three control abstractions, not only share the same conceptual motivation, but are also based on a common enabling technology: *first-class closures*. Closures make it possible to separate the definition of a dynamic behaviour from its use, providing a simple and uniform implementation technology for the three control abstractions. Once closures are available, the control abstractions really become lightweight extensions, which is part of their appeal.

The rest of this paper is organized as follows. Section 2 briefly reviews the local search architecture and its implementation in COMET. Section 3 gives a brief overview of closures. Sections 4, 5, and 6 present the new control abstractions and sketches their implementation. Section 7 presents some experimental results showing the viability of the approach. Section 8 concludes the paper.

2 The Constraint-Based Architecture for Local Search

This section is a brief overview of the constraint-based architecture for local search and its implementation in COMET. See [13] for more detail. The architecture consists of a declarative and a search component organized in three layers. The kernel of the architecture is the concept of *invariants* over algebraic and set expressions [14]. Invariants are expressed in terms of incremental variables and specify a relation which must be maintained under modifications to its variables. Once invariants are available, it becomes natural to support the concept of *differentiable objects*, a fundamental abstraction for local search programming. *Differentiable objects maintain a number of properties (using invariants) and can be queried to evaluate the effect of local moves on these properties.* They are fundamental because many local search algorithms evaluate the effect of various

```
range Size = 1..1024;
LocalSolver ls();
UniformDistribution distr(Size);
inc{int} queen[i in Size](ls,Size) := distr.get();
int neg[i in Size] = -i;
int pos[i in Size] = i;

ConstraintSystem S(ls);
S.post(new AllDifferent(queen));
S.post(new AllDifferent(queen,neg));
S.post(new AllDifferent(queen,pos));
inc{set{int}} conflicts(ls) <- argMax(q in Size) S.violations(queen[q]);
m.close();

Counter it(ls);
while (!S.isTrue()) {
  select(q in conflicts)
    selectMin(v in Size)(S.getAssignDelta(queen[q],v))
      queen[q] := v;
  it++;
}
```

Fig. 1. The Queens Problem in Comet.

moves before selecting the neighbor to visit. Two important classes of differentiable objects are constraints and functions. A differentiable constraint maintains properties such as its satisfiability, its violation degree, and how much each of its underlying variables contribute to the violations. It can be queried to evaluate the effect of local moves (e.g., assignments and swaps) on these properties. *Differentiable objects also capture combinatorial substructures arising in many applications* and are appealing for two main reasons. On the one hand, they are high-level modeling tools which can be composed naturally to build complex local search algorithms. As such, they bring into local search some of the nice properties of modern constraint satisfaction systems. On the other hand, they are amenable to efficient incremental algorithms that exploit their combinatorial properties. The use of combinatorial constraints is also advocated in [3, 7, 17, 26].

These first two layers, invariants and differentiable objects, constitute the declarative component of the architecture. The third layer of the architecture is the search component which aims at simplifying the implementation of heuristics and meta-heuristics, another critical aspect of local search algorithms. It does not prescribe any specific heuristic or meta-heuristic. Rather, it features high-level constructs and abstractions to simplify the neighborhood exploration and the implementation of meta-heuristics. These includes several multidimensional selectors, abstractions to manipulate solutions, and advanced simulation techniques.

Figure 1 illustrates the architecture, and its implementation in COMET, on the queens problem. The COMET algorithm is based on the min-conflict heuristic [16]. The algorithm starts with an initial random configuration. Then, at each iteration, it chooses the queen violating the largest number of constraints and moves it to a position minimizing its violations. This step is iterated until a solution is found. Since a queen must be placed on every column, the algorithm uses an array `queen` of variables and `queen[i]` denotes the row of the queen placed on column `i`. Lines 1-6 declare a range, a local solver, a uniform distribution, an array of incremental variables for representing the row of each queen, as well as two arrays of constants. *The modeling component* is given in Lines 7-12. Line 7 declares a constraint system. Lines 8-10 add the three traditional `AllDifferent` constraints, showing how COMET supports "global" combinatorial constraints for local search. Line 11 expresses an invariant which maintains the set of queens with the most violations. Operator `argMax(v in S) E` simply returns the set of values `v` in `S` which maximizes `E`. *The search component* is given in lines 13-19. It iterates lines 15-17 until the constraint system is true, i.e., no constraint is violated. Line 15 selects a most violated queen, while line 16 selects a new value `v` for the selected queen. The value is selected to minimize the number of violations of the selected queen. To implement this min-conflict heuristic, COMET queries the constraint system, a differential object, to find out the effect of assigning queen `q` to each row. Line 17 simply executes the move, automatically updating all invariants and constraints. The use of the counter `it` will become clear later in the paper.

Observe that the search and declarative components are clearly separated in the program. It is thus easy to modify one of them (e.g., adding a constraint and/or changing the search heuristic) without affecting the other. Although the two components are physically separated in the program code, they closely collaborate during execution. The declarative component is used to guide the search, while the assignment `queen[q] := v` starts a propagation phase which updates all invariants and constraints. This compositionality and clear separation of concerns are some of the appealing features of the architecture. *This is precisely such properties which this paper tries to foster further.* Note also that the declarative component only specifies the properties of the solutions, as well as the data structures to maintain. It does not specify how to update them, which is the role of the incremental algorithms in the COMET runtime system.

3 Closures in COMET

Closures are the common enabling technology behind all three control abstractions introduced in this paper. A closure is a piece of code together with its environment. Closures are ubiquitous in functional programming languages, where they are first-class citizens. They are rarely supported in object-oriented languages however. To illustrate the use of closures in COMET, consider the following class

```
class DemoClosure {
  DemoClosure() {}
  Closure print(int i) {
    return new closure
      {cout << i << endl;}
  }
}
DemoClosure demo();
Closure c1 = demo.print(9);
Closure c2 = demo.print(5);
call(c2);
call(c1);
call(c2);
```

Method `print` receives an integer `i` and returns a closure which, when executed, prints `i` on the standard output. The following snippet shows how to use closures in COMET: the snippet displays 5, 9, and 5 on the standard output. Observe that closures are first-class citizens: They can be stored in data structures, passed as parameters, and returned as results. The two closures created in the example above share the same code (i.e., `cout << i << endl`), but their environments differ. Both contain only one entry (variable `i`), but they associate the value 9 (closure `c1`) and the value 5 (closure `c2`) to this entry. When a closure is created, its environment is saved and, when a closure is executed, the environment is restored before, and popped after, execution of its code. Closures can be rather complex and have environments containing many parameters and local variables, as will become clear later on.

4 Events for Modularity, Compositionality, and Reuse

One of the fundamental benefits of COMET is its ability to separate problem modeling from search. This separation of concerns is made possible by incremental variables, invariants, and differential objects. However, practical applications typically involve other components which would also benefit from such modularity. One such component is algorithm animation, which is valuable early in the development process to visualize the local search behavior. Another component is the meta-heuristic which is often orthogonal and independent from the search heuristic. This section introduces the concept of *publish/subscribe events* in COMET, which make this separation of concerns possible. Informally speaking, classes can publish events, which can be subscribed by event-handlers elsewhere in the code. Methods in the classes can then notify these events, which triggers the event-handler behaviour. We first focus on how to use events for animation and meta-heuristic. We then show how to publish and notify events.

Events for Animation Consider a graphical animation for the n-queens problem and assume the existence of an `Animation` class handling the graphics and providing a method `updatePosition(int q,int p)` to display the queen on column `q` on row `r`. Such an animation is obtained by inserting the snippet

```
forall(q in Size)
    whenever queen[q]@changes(int or,int nr)
       animation.updateQueen(q,nr);
```

just before the search component (between lines 12 and 13). The core of the snippet is an event-handler that specifies that, whenever the value of `queen[q]` changes from `or` to `nr`, the code `animation.updateQueen(q,nr)` must be executed. This event-handler is installed for all queens.

There are a few important points to highlight here. First, the animation code is completely separated from both the modeling and the search components. The glue between the components is the event `changes` on incremental variables which is *notified* whenever a variable is assigned a new value. The snippet achieves the same effect as calling `animation.updateQueen(q,nr)` after the assignment of queens, while clearly separating the two aspects and avoiding to clutter the heuristic with animation code. This makes the code more readable and easier to modify and extend. Second, observe that the event-handler behaviour `animation.updateQueen(q,nr)` is a closure which depends on the value of `q` in the environment and is created when the event is subscribed to. Closures make the animation code more natural, avoid the definition of intermediary classes, and feature a textual proximity between the event-handler condition (e.g., the queen is assigned a new value) and its behavior (e.g., update the display of the queen). In traditional object-oriented languages, event conditions and behaviors are separated, which complicates reading and requires new class definitions to store the information necessary to execute the behavior. Finally, observe that events are statically and strongly typed: they enable information to be transmitted from the notifier (e.g., the incremental variable) to the event-handler in a safe fashion with no downcasting.

Events are also compositional. Consider, for instance, adding the functionality of coloring the queens differently according to their number of violations. It is sufficient to add the instructions

```
inc{int} violation[q in Size](m) <- S.violations(queen[q]);
forall(q in Size)
    whenever violation[q]@changes(int ov,int nv)
       animation.updateColor(q,nv);
```

This snippet declares an array of incremental variables maintaining the number of violations of each queen, and updates the color of a queen each time its number of violations is updated. Note that the number of violations of a queen may change even when the queen is not moved. Hence, it is not possible to insert the behaviour elsewhere in the program, while remaining incremental, i.e., only considering the queens whose number of violations was modified. This example shows the strengths of events in COMET: they enable elegant animation codes, which would require complex control flows, the creation of intermediary classes, and/or less incrementality in other languages.

Events for Meta-Heuristics Events are also beneficial to separate the search heuristic and the meta-heuristic (e.g., tabu-search). They make it possible to divide the statement into modeling, search, and meta-heuristic components. For illustration purposes, consider upgrading the queen algorithm with a tabu-search strategy, which would make a queen tabu for a number of iterations, each time a queen is moved. The tabu-list management can be almost entirely separated from the search heuristic. For instance, the snippet

```
1. set{int} tabu();
2. forall(q in Size)
3.  whenever queen[q]@changes(int o,int n) {
4.     tabu.insert(q);
5.     when it@reaches[it+tLen]()
6.        tabu.remove(q);
7.  }
```

shows a simple management of the tabu list, which we now explain in detail. The code declares a set `tabu` to store the tabu queens and features two nested event-handlers. The outermost event-handler is notified each time a queen is moved. It inserts the queen in the tabu set and install the second event-handler (lines 5-6) whose goal is to remove `q` from the tabu set after `tLen` iterations, where `tLen` is the length of the tabu list. This second handler is interesting in several ways. First, it features a *key-event*, i.e., an event which is parametrized by a specific key which is in between brackets in the code. In this example, the key is an iteration number, i.e., the handler is notified when the counter `it` reaches the value `it+tLen`, i.e., the value of the counter when the handler is installed (subscription time) plus the length of the tabu-list. Second, the handler uses the `when` construct, which means that it will be notified only once.

Once this code is in place, the only modification in the search heuristic consists in selecting the queen with the largest number of violations among the non-tabu queens (instead of among all queens). As a consequence, the "glue" between the components (i.e., the counter and the tabu-set) is minimal and the proper behavior is achieved without interleaving the heuristic and the meta heuristic in the source code. Note that, in complex applications, this glue can be anticipated in the first place by assuming that moves are always selected from a restricted set specified by the modeling and/or meta-heuristic components.

Event Specification and Notification The examples above focused on the event-handler (the *subscription* part) and showed how the `when` and `whenever` are used to register a behaviour. Since they only used primitive objects, no explicit specification and notification of events (the *publish* part) was necessary. Of course, COMET makes it possible to define new events. Each class may publish some events or key-events by declaring them. Its methods are then responsible to notify these events appropriately. To illustrate event specification and notification, consider a possible implementation of the class `Counter` in COMET:

```
class Counter {
  inc{int} _cnt;
  Event changes(int ov,int nv);
  KeyEvent reaches();
  Counter();
  int ++();
}
```

```
Counter::Counter(){_cnt=new inc{int}(0);}
int Counter::++() {
  int old = _cnt++;
  notify changes(old,_cnt);
  notify reaches[_cnt]();
  return _cnt;
}
```

The class declares an incremental variable `_cnt`, an event `changes` with two parameters, a key-event `reaches` with no parameter, the constructor and the operator. The implementation of the operation (on the right part of the snippet) notifies the `changes` events, passing the old and new values of the incremental variable. It also notifies all the key-events `reaches`, whose keys are smaller or equal to the value of `_cnt`. These notifications triggers all the event-handlers associated with these events, i.e., it executes the closures which were registered at subscription time by the `when` and `whenever` instructions. In aspect-oriented terms, the `notify` instructions are joint-points and `when` and `whenever` statements are *dynamic* aspects, i.e., aspects associated with instances, not with classes as is typical in aspect-oriented languages.

Implementation of Events Conceptually, the implementation of events is close to the OBSERVER design pattern. An event is compiled into virtual machine instructions which explicitly use closures as shown below:

```
when x@changes(int o,int n)              aload x
    cout << n << endl;          ==>      newClosure "cout << n << endl;"
                                         subscribeEvent changes,<o,n>
```

The virtual machine is a JVM-like stack machine and `x` and the closure are retrieved from the stack in `subscribeEvent`. At the instance level, each event corresponds to a data structure which collects all the subscribers. Upon notification, the appropriate subscribers are executed, i.e., their parameters are properly initialized and their closures are executed.

5 Union of Heterogeneous Neighborhoods

Many complex applications in areas such as scheduling and routing use complex neighborhoods consisting of several heterogeneous moves. For instance, the elegant tabu-search of Dell'Amico and Trubian [5] consists of the union of the subneighborhoods, each of which consisting of several types of moves. Similarly, many advanced vehicle routing algorithms [10, 4, 2] use a variety of moves (e.g., swapping visit orders and relocating customers on other routes), each of which may involve a different number of customers and trucks.

The difficulty in expressing these algorithms come from the temporal disconnection between the move selection and execution. In general, a tabu-search or a greedy local search algorithm first scans the neighborhood to determine the best move, before executing the selected move. However, in these complex applications, the exploration cannot be expressed using a (multidimensional) selector, since the moves are heterogeneous and obtained by iterating over different sets. As a consequence, an implementation would typically create classes to store the information necessary to characterize the different types of moves. Each of these classes would inherit from a common abstract class (or would implement the same interface). During the scanning phase, the algorithm creates instances of these classes to represent selected moves and stores them in a selector whenever appropriate. During the execution phase, the algorithm extracts the selected move and applies its `execute` operation. The drawbacks of this approach are twofold. On the one hand, it requires the definition of a several classes to represent the moves. On the other hand, it fragments the code, separating the *evaluation* of a move from its *execution* in the program source. As a result, the program is less readable and more verbose.

The `neighbor` *Construct* COMET supports a `neighbor` construct, which relies heavily on closures and eliminates these drawbacks. It makes it possible to specify the move evaluation and execution in one place and avoids unnecessary class definitions. More important, it significantly enhances compositionality and reuse, since the various subneighborhoods do not have to agree on a common interface or abstract class. `neighbor` constructs are of the form

```
neighbor(δ,N) M
```

where `M` is a move, δ is its evaluation, and `N` is a neighbor selector, i.e., a container object to store one or several moves and their evaluations. COMET supports a variety of such selectors and users can define their own, since they all have to implement a common interface. For instance, a typical neighbor selector for tabu-search maintains the best move and its evaluation. The execution of the `neighbor` instruction queries selector `N` to find out whether it accepts a move of quality δ, in which case the closure of `M` is submitted to `N`.

Jobshop Scheduling We now illustrate how the `neighbor` construct significantly simplifies the implementation of the tabu-search algorithm of Dell'Amico and Trubian (DT) for jobshop scheduling. We first review the basic ideas behind the DT algorithm and then sketch how the neighborhood exploration is expressed in COMET. Algorithm DT uses neighborhood $NC = RNA \cup NB$, where RNA is a neighborhood swapping vertices on a critical path (critical vertices) and NB is a neighborhood where a critical vertex is moved toward the beginning or the end of its critical block. More precisely, RNA considers sequences of the form $\langle p, v, s\rangle$, where v is a critical vertex and p, v, s represent successive tasks on the same machine, and explores all permutations of these three vertices. Neighborhood NB considers a maximal sequence $\langle v_1, \ldots, v_i, \ldots, v_n\rangle$ of critical vertices on the same machine. For each such subsequence and each vertex v_i, it explores the schedule obtained by placing v_i at the beginning or at the end of the block, i.e.,

$$\langle v_i, v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n\rangle \vee \langle v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n, v_i\rangle$$

Since these schedules are not necessarily feasible, NB actually considers the leftmost and rightmost feasible positions for v_i (instead of the first and last position). NB is connected which is an important theoretical property of neighborhoods.

We now show excerpts of the neighborhood implementation in COMET. The top-level methods are as follows:

```
void executeMove() {
  MinNeighborSelector N();
  exploreN(N);
  if (N.hasMove())call(N.getMove());
}
```

```
void exploreN(NeighborSelector N)
{
   exploreRNA(N);
   exploreNB(N);
}
```

Method `executeMove` creates a selector, explores the neighborhood, and executes the best move (if any). Method `exploreN` explores the neighborhood and illustrates the compositionality of the approach: It is easy to add new neighborhoods without modifying existing code, since the subneighborhoods do not have to agree on a common interface or abstract class. The implementation of `exploreRNA` and `exploreNB` is of course where the `neighbor` construct is used.

```
1. void exploreNB(NeighborSelector N) {
2.   forall(v in _jobshop.getCriticalVertices()) {
3.      int lm = _jobshop.leftMostFeasible(v);
4.      if (lm > 0) {
5.         int delta = _jobshop.moveBackwardDelta(v,lm);
6.         if (acceptNBLeft(delta,v))
7.            neighbor(delta,N) _jobshop.moveBackward(v,lm);
8.      }
9.      int rm = _jobshop.rightMostFeasible(v);
10.     if (rm > 0) {
11.        int delta = _jobshop.moveForwardDelta(v,rm);
12.        if (acceptNBRight(delta,v))
13.           neighbor(delta,N) _jobshop.moveForward(v,rm);
14.     }
15.  }
16.}
```

Fig. 2. Exploration of Neighborhood *NB* in Comet.

Figure 2 gives the implementation of `exploreNB`: method `exploreRNA` is similar in spirit, but somewhat more complex, since it involves 5 different moves, as well as additional conditions to ensure feasibility. Method `exploreNB` uses the instance variable `_jobshop`, which is a differentiable object representing the disjunctive graph, a fundamental concept in jobshop scheduling [20]. This differential object maintains the release and tail dates of all vertices, as well as the critical paths, under various operations on the disjunctive graph. The `exploreNB` method iterates over all critical vertices. For each of them it finds the leftmost feasible insertion point in its critical block (line 3). If such a feasible insertion point exists, it evaluates the move (line 5) and then tests if the move is acceptable (line 6). In the DT algorithm, this involves testing the tabu status, a cycling condition, and the aspiration criterion. If the move is acceptable, the `neighbor` instruction is executed. The move itself consists of moving vertex `v` by `lm` positions backwards. Note that, although the move is specified in the `neighbor` instruction, it is not executed. Only the best move is executed and this takes place in method `executeMove` once the entire neighborhood has been explored. The remaining of method `exploreNB` handles the symmetric forward move.

The neighborhood exploration is particularly elegant (in our opinion). Although a move evaluation and its execution take place at different execution times, the `neighbor` construct makes it possible to specify them together, significantly enhancing clarity and programming ease. The move evaluation and execution are textually adjacent and the logic underlying the neighborhood is not made obscure by introducing intermediary classes and methods. Compositionality is another fundamental advantage of the code organization. As mentioned earlier, new moves can be added easily, without affecting existing code. Equally or more important perhaps, the approach separates the neighborhood definition (method `exploreN`) from its use (method `executeMove` in the DT algorithm). This makes it possible to use the neighborhood exploration in many different ways without any modification to its code. For instance, a semi-greedy strategy, which selects one of the k-best moves, only requires to use a semi-greedy selector. Similarly, method `exploreN` can be used to collect all neighbors which is useful in intensification strategies based on elite solutions [18].

Implementation of Neighbor The `neighbor` construct is only syntactic sugar once closures are available. Indeed, the syntactic form is rewritten as shown below:

```
forall(v in Size)                        forall(v in Size)
   neighbor(Δ(v),N)         ⟹              δ ← Δ(v)
      M(v);                                 if (N.accept(δ))
                                               N.insert(δ,new closure {M(v);});
```

The rewriting uses method `accept` on the selector to determine whether to accept a move. It also ensures that closures are constructed lazily.

6 Sequential Composition of Neighborhoods

This section discusses the use of checkpoint to express the sequential composition concisely. Sequential composition is often fundamental in very large neighborhood search, which explores sequences or trees of (possibly heterogeneous) moves and selects the best encountered neighbor (e.g., [8, 1]). This section illustrates these concepts using variable-depth neighborhood search (VDNS) [8], a very large neighborhood search shown very effective on graph-partitioning and traveling salesman problems.

Variable-Depth Neighborhood Search VDNS consists of exploring a sequences of k distinct moves and moving to the state with best evaluation in the sequence. By exploring sequences which include degrading moves, VDNS may avoid being trapped in local optima of poor quality.

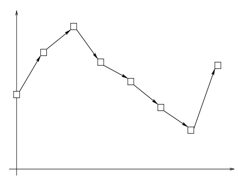

Fig. 3. A Sequence of Moves

Consider Figure 3 which plots the quality of a sequence of moves. Each node in the graph corresponds to a computation state and two successive nodes are neighbors in the transition graph of the local search. VDNS explores the whole sequence and then returns to the best computation state, i.e., the before-last node.

```
function boolean selectBest(LocalSolver ls,int l,inc{int} f,Closure Move) {
    boolean found = false;
    with checkpoint(ls) {
       Checkpoint chp(ls); int best = f;
       forall(i in 1..l) {
          call(Move);
          if (f < best) {
             found = true; best = f; chp = new Checkpoint(ls);
          }
       }
       chp.restore();
    }
    return found;
}
```

Fig. 4. The implementation of VDNS in COMET.

Checkpoints *Checkpoints* are a simple conceptual abstraction to express VLNS algorithms. A checkpoint is simply a data structure that implicitly represents the computation state of a local solver, i.e., the state of all incremental variables and data structures of the solver. Whenever a local solver is in checkpointing mode, checkpoints can be saved and, later, restored in order to reset all incremental variables, constraints, and data structures to their earlier states. Checkpoints are first-class citizens in COMET. They also encapsulate incremental algorithms to avoid saving entire computation states.

Variable-Depth Neighborhood Search in COMET We now illustrate how to express VDNS in COMET for graph partitioning [8], where moves consists of swapping two vertices, one from each set in the partition. The snippet

```
selectBest(ls,nb/2,cost)
    select(s in BestSwaps) {
       x[s.o] :=: x[s.d]; mark[s.o] := true; mark[s.d] := true;
    }
```

shows the core of the search procedure in COMET. In the snippet, `ls` is the local solver, `nb` is the number of vertices, `cost` is the cost of the partition, `bestSwaps` is an incremental set of tuples which maintains the best swaps, `x` is an array of incremental variables specifying which set of the partition a vertex belongs to, and `mark` is an array of incremental Boolean variables, indicating whether a vertex have been selected in the VDNS sequence already. The `selectBest` function is the cornerstone of the VDNS implementation. It receives four arguments: the local solver, the length of the sequence, the function to minimize (an incremental variable), and a closure representing the move. Here the move consists of selecting a tuple `s` in `BestSwaps` and to swap the vertices `s.o` and `s.d`. Both vertices are then marked in order to avoid selecting them again in the sequence.

Figure 4 depicts the implementation of function `selectBest`. It uses the `with checkpoint(ls)` statement to indicate the use of checkpointing inside the enclosed block. It saves the current state in variable `chp` using instruction `Checkpoint chp(ls)`. The `forall` loop explores a sequence of `l` moves, storing the best computation state in variable `chp`. After this exploration, instruction `chp.restore()` restores the best computation state encountered (possibly the initial state). Note that COMET supports the syntactic rewriting from `f(`$\mathtt{a}_1$`,...,`$\mathtt{a}_n$`) S` to `f(`$\mathtt{a}_1$`,...,`$\mathtt{a}_n$`,new closure { S })` when the last argument of function `f` is a closure. The VDNS implementation has a number of interesting features. First, it is entirely generic and reusable: It can be applied to an arbitrary move and separates search heuristic and the meta-heuristic. Second, checkpoints specify *what* to maintain, i.e., the "best" computation states, but not *how* to save or restore it. The implementation uses incremental algorithms to do so, but this is abstracted from programmers. Finally, observe the role of closures for the genericity of the VLNS implementation.

Implementation of Checkpoints We now discuss the checkpoint implementation. The key to an incremental implementation lies in a representation of computation states as sequences of primitive moves from an initial state (i.e., the state when the checkpoint statement is executed). In other words, a state s is a sequence $\langle m_0, \ldots, m_k\rangle$ where m_i is a primitive move. A primitive move in COMET is a function $f : State \rightarrow State$ from computation states to computation states which is invertible, i.e., there exists a function f^{-1} such that $f(f^{-1}(s)) = s$. For instance, a move `x[i]:=j` corresponds to a function $f(s) = s\{x[i]/j\}$ where $s\{y/v\}$ represents the state s where y is assigned the value v. The inverse move is of course $f^{-1}(s) = s\{x[i]/lookup(s^0, x[i])\}$ where s^0 is the computation state before executing the move, and *lookup* reads the value of a variable in a computation state. Consider now how to restore a state s_r from a state s_c where

$$s_c = \langle m_0, \ldots, m_n, m'_{n+1}, \ldots, m'_k\rangle$$
$$s_r = \langle m_0, \ldots, m_n, m''_{n+1}, \ldots, m''_l\rangle.$$

The COMET implementation exploits the common prefix of the two states. It undoes the suffix $\langle m'_{n+1}, \ldots, m'_k\rangle$ by using the inverse moves, and then executes the moves $\langle m''_{n+1}, \ldots, m''_l\rangle$. This implementation has several properties. First, its memory requirements are independent of the size of the computation states. Only moves are memorized and the size of a checkpoint c only depends on the length of the sequence from the initial state to c. Second, the runtime requirements are also minimal, since they either reexecute a subsequence executed before or they execute the inverse of such a subsequence. For VDNS, for instance, restoring the best state does not change the asymptotic complexity: in the worst case, restoring the checkpoint involves as much work as exploring the sequence.

The checkpoint implementation is related to techniques underlying generic search strategies (e.g., [19, 15, 22]). However, it does not use backtracking and/or trailing. Rather, it makes heavy use of inverse moves, which is efficient because the invariant propagation algorithm never updates the same incremental variable twice [14] (which is not the case in constraint satisfaction algorithms in general). Our implementation thus combines low memory requirements with incrementality, which is critical for many local search applications.

7 Experimental Results

This section describes some preliminary experimental results to demonstrate the practical viability of the abstractions and of closures. It compares various implementations of the tabu-search algorithm DT (the goal, of course, is not to compare various scheduling algorithms). In particular, it compares the original results [5], a C++ implementation [21], and the COMET implementation. Table 1 presents the results corresponding to Table 3 in [5]. Since DT is actually faster on the LA benchmarks (Table 4 in [5]), these results are representative. In the table, DT is the original implementation on a 33mhz PC, DT* is the scaled times on a 745mhz PC, KS is the C++ implementation on a 440 MHz Sun Ultra, KS* are the scaled times on a 745mhz PC, and CO are the COMET times on a 745mhz PC. Scaling was based on the clock frequency, which is favorable to slower machines (especially for the Sun). The times corresponds to the average over multiple runs (5 for DT, 20 for KS, and 50 for CO). Results for COMET are for the JIT compiler but include garbage collection. The results clearly indicate that COMET can be implemented to be competitive with specialized programs. Note also that the C++ implementation is more than 4,000 lines long, while the COMET program has about 400 lines.

	ABZ5	ABZ6	ABZ7	ABZ8	ABZ9	MT10	MT20	ORB1	ORB2	ORB3	ORB4	ORB5
DT	139.5	86.8	320.1	336.1	320.8	155.8	160.1	157.6	136.4	157.3	156.8	140.1
DT*	6.2	3.8	14.2	15.1	14.2	6.9	7.1	7.0	6.0	7.0	6.9	6.2
KS	7.8	8.2	20.7	23.1	20.3	8.7	16.4	9.2	7.8	9.3	8.5	8.1
KS*	4.6	4.8	12.2	13.6	11.9	5.1	9.6	5.4	4.6	5.5	5.0	4.8
CO	5.9	5.7	11.7	9.9	9.0	6.7	9.8	5.6	4.8	5.6	6.3	6.5

Table 1. Computational Results on the Tabu-Search Algorithm (DT)

8 Conclusion

This paper presented three novel control abstractions for COMET, which significantly enhance the compositionality, modularity, and reuse of COMET. These abstractions may significantly improve conciseness, extensibility, and clarity of the local search implementations. They all rely on first-class closures as the enabling technology and can be implemented efficiently.

One of the most appealing features of COMET is its small number of fundamental concepts, as well as their generality. First-class closures simplify many applications beyond local search (e.g., [12]) and are ubiquitous in functional programming. Events are related to many constructs in the logic and functional communities (e.g., delay mechanisms and reactive functional programming). Invariants (one-way constraints) and constraints are widely recognized as natural vehicles for many applications. These concepts provide significant support for local search, and may significantly reduce the distance between high-level descriptions of the algorithms and their actual implementations. Yet they are non-intrusive and impose minimal “constraints” on programmers, who keeps control of their algorithms and their code organization. An interesting topic for future research is to study how to unify the COMET architecture with the tree-search models proposed in [23, 11], since both approaches have orthogonal strengths.

References

1. E. Balas and A. Vazacopoulos. Guided local search with shifting bottleneck for job-shop scheduling. Management Science, 44(2), 1998.
2. R. Bent and P. Van Hentenryck. A Two-Stage Hybrid Local Search for the Vehicle Routing Problem with Time Windows. *Transportation Science*, 2001. (To Appear).
3. C. Codognet and D. Diaz. Yet Another Local Search Method for Constraint Solving. In *AAAI Fall Symposium on Using Uncertainty within Computation*, Cape Cod, MA., 2001.
4. B. De Backer et al. Solving Vehicle Routing Problems Using Constraint Programming and Metaheuristics. *Journal of Heuristics*, 6:501–523, 2000.
5. M. Dell'Amico and M. Trubian. Applying Tabu Search to the Job-Shop Scheduling Problem. *Annals of Operations Research*, 41:231–252, 1993.
6. L. Di Gaspero and A. Schaerf. Writing Local Search Algorithms Using EasyLocal++. in *Optimization Software Class Libraries*, Kluwer, 2002.
7. P. Galinier and J.-K. Hao. A General Approach for Constraint Solving by Local Search. In *CP-AI-OR'00*, Paderborn, Germany, March 2000.
8. B. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, 49:291–307, 1970.
9. G. Kiczales, J. Lamping, A. Menhdhekar, C. Maeda, C. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-oriented programming. In *ECOOP '97*.
10. Kindervater, G. and Savelsbergh, M.W. Vehicle routing: Handling edge exchanges. In *Local Search in Combinatorial Optimization*, Wiley, 1997.
11. F. Laburthe and Y. Caseau. SALSA: A Language for Search Algorithms. In *CP'98*.
12. D. Manolescu. Workflow Enactment with Continuation and Future Objects. In *OOPLSA'02*, Seattle, WA, 2002.
13. L. Michel and P. Van Hentenryck. A Constraint-Based Architecture for Local Search. In *OOPLSA'02*, Seattle, WA, 2002.
14. L. Michel and P. Van Hentenryck. Localizer. *Constraints*, 5:41–82, 2000.
15. L. Michel and P. Van Hentenryck. A Decomposition-Based Implementation of Search Strategies. *ACM Transactions on Computational Logic*, 2002. (To Appear).
16. S. Minton, M. Johnston, and A. Philips. Solving Large-Scale Constraint Satisfaction and Scheduling Problems using a Heuristic Repair Method. In *AAAI-90*.
17. A. Nareyek. *Constraint-Based Agents*. Springer Verlag, 1998.
18. E. Nowicki and C. Smutnicki. A fast taboo search algorithm for the job shop problem. *Management Science*, 42(6):797–813, 1996.
19. L. Perron. Search Procedures and Parallelism in Constraint Programming. In *CP'99*, Alexandra, VA, 1999.
20. B. Roy and B. Sussmann. Les problèmes d'ordonnancement avec contraintes disjonctives. Note DS No. 9 bis, SEMA, Paris, France, 1964.
21. K. Schmidt. Using Tabu-search to Solve the Job-Shop Scheduling Problem with Sequence Dependent Setup Times. ScM Thesis, Brown University, 2001.
22. C. Schulte. Comparing trailing and copying for constraint programming. In *ICLP'99*.
23. P. Shaw, B. De Backer, and V. Furnon. Improved local search for CP toolkits. *Annals of Operations Research*, 115:31–50, 2002.
24. C. Turner, A. Fuggetta, L. Lavazza, and A. Wolf. A conceptual basis for feature engineering. *Journal of Systems and Software*, 49(1):3–15, 1999.
25. S. Voss and D. Woodruff. *Optimization Software Class Libraries*. Kluwer, 2002.
26. J. Walser. *Integer Optimization by Local Search*. Springer Verlag, 1998.