

MiniSearch: a solver-independent meta-search language for MiniZinc

Andrea Rendl¹, Tias Guns², Peter J. Stuckey³, and Guido Tack¹

¹ National ICT Australia (NICTA) and Faculty of IT, Monash University, Australia
andrea.rendl@nicta.com.au, guido.tack@monash.edu

² KU Leuven, Belgium

tias.guns@cs.kuleuven.be

³ National ICT Australia (NICTA) and University of Melbourne, Victoria, Australia
pstuckey@unimelb.edu.au

Abstract. Much of the power of CP comes from the ability to create complex hybrid search algorithms specific to an application. Unfortunately there is no widely accepted standard for specifying search, and each solver typically requires detailed knowledge in order to build complex searches. This makes the barrier to entry for exploring different search methods quite high. Furthermore, search is a core part of the solver and usually highly optimised. Any imposition on the solver writer to change this part of their system is significant.

In this paper we investigate how powerful we can make a uniform language for meta-search *without placing any burden on the solver writer*. The key to this is to only interact with the solver when a solution is found. We present MINISEARCH, a meta-search language that can directly use any FLATZINC solver. Optionally, it can interact with solvers through an efficient C++ API. We illustrate the expressiveness of the language and performance using different solvers on a number of examples.

1 Introduction

When using constraint programming (CP) technology, one often needs to exert some control over the meta-search mechanism. Meta-search, such as Branch-and-Bound search (BaB) or Large Neighbourhood Search (LNS), happens on top of CP tree search, and aids finding good solutions, often by encoding meta-information a modeller has about the problem.

Unfortunately, there is no widely accepted standard for controlling search or meta-search. A wide range of high-level languages have been proposed that are quite similar in how constraints are specified. However, they differ significantly in the way search can be specified. This ranges from built-in minimisation and maximisation only to fully programmable search. The main trade-off in developing search specification languages is the expressivity of the language versus the required integration with the underlying solver. Fully programmable search, including meta-search, is most expressive but requires deep knowledge of and tight integration with a specific solver. Languages like OPL [24] and COMET [13] provide convenient abstractions for programmable search for the solvers bundled

with these languages. On the other hand, some solver-independent languages such as Esra [3], Essence [5] and Essence' [6] do not support search specifications. Zinc [12] and MINIZINC [16] have support for specifying variable and value ordering heuristics to CP solvers, but also no real control over the search.

Search combinators [19] was recently proposed as a generic meta-search language for CP solvers. It interacts with the solver *at every node* in the search tree. While very expressive, it requires significant engineering effort to implement for existing solvers, since typically the solvers' search engines are highly optimised and tightly integrated with other components, such as the propagation engine and state space maintenance.

In this paper, we introduce MINISEARCH, a new combinator-like meta-search language that has three design objectives: a *minimal solver interface* to facilitate solver support, *expressiveness*, and, most importantly, *solver-independence*.

The objective to obtain a minimal solver interface stems from lessons learnt in the design of search combinators that interact with the solver at every *node*. In contrast, MINISEARCH interacts with the underlying solving system only *at every solution*, which is a minimal interface. At every solution, constraints can be added or constraints in a well-defined *scope* can be removed, before asking for the next solution. If the underlying solver does not support dynamic adding and removing of constraints, MINISEARCH can emulate this behaviour, for little overhead.

Despite the lightweight solver interface, MINISEARCH is surprisingly expressive and supports many meta-search strategies such as BaB search, lexicographic BaB, Large Neighbourhood Search variants, AND/OR search, diverse solution search, and more. MINISEARCH can also be used to create interactive optimisation applications. Moreover, since MINISEARCH builds upon MINIZINC, all MINIZINC language features and built-ins can be used, for instance to formulate custom neighbourhoods.

Solver-independence is the most important contribution of MINISEARCH. All solvers that can read and solve FLATZINC, which the majority of CP solvers do [15,22], can be used with MINISEARCH. Moreover, solvers that provide native meta-search variants, such as branch-and-bound, can declare so and avoid executing the MINISEARCH decomposition instead. At the language level, this is similar to the handling of global constraints in MINIZINC. Thus, solvers can apply their strengths during meta-search, despite the minimal interface.

2 The MINISEARCH language

MINISEARCH is a high level meta-search language based on MINIZINC 2.0 [23]. MINIZINC is used for formulating the model and the constraints posted during search, with language extensions for specifying the search. A MINISEARCH specification is provided, together with the `search` keyword, as an argument to MINIZINC's `solve` item. It is executed by the MINISEARCH kernel (see Sec. 4). With the built-in language extensions summarised in Tab. 1, users can define functions such as the following branch-and-bound (BaB) minimisation:

Table 1. MINISEARCH built-ins

MINISEARCH built-ins	Description
<code>next()</code>	find the next solution
<code>post(c)</code>	post the MINIZINC constraint c in the current scope
<code>scope(s)</code>	open a local scope containing search s
$s_1 \wedge s_2$	run s_1 and iff successful, run s_2
$s_1 \setminus / s_2$	run s_1 and iff it fails, run s_2
<code>if s then s_1 else s_2</code>	if s is successful, run s_1 , otherwise s_2
<code>repeat(s)</code>	repeat search s until <code>break</code> is executed
<code>repeat(i in 1..N)(s)</code>	repeat search s N times or until <code>break</code> is executed
<code>break()</code>	break within a <code>repeat</code>
<code>fail()</code>	return ‘failure’
<code>skip()</code>	return ‘success’
<code>time_limit(ms,s)</code>	run s until timelimit ms is reached
<code>print(S)</code>	print MINIZINC output string S
<code>print()</code>	print solution according to model output specification
<code>c := v</code>	assign parameter c the value v
<code>commit()</code>	commit to last found solution in function scope
<code>sol(v)</code>	return solution value of variable v
<code>hasSol()</code>	returns true if a solution has been found

```

1 include "minisearch.mzn";
2 var int: obj; % other variables and constraints not shown
3 solve search minimize_bab(obj);
4 output ["Objective: " ++ show(obj)];
5
6 function ann: minimize_bab(var int: obj) =
7     repeat( if next() then commit() /\ print() /\ post(obj < sol(obj))
8             else break endif );

```

The `include` item on line 1 includes the built-in MINISEARCH function declarations. This is necessary for any MINIZINC model that uses MINISEARCH. Line 3 contains the MINIZINC `solve` item followed by the new `search` keyword and a user-defined MINISEARCH function that takes a variable representing the objective as argument. Line 4 is the MINIZINC `output` item, specifying how solutions should be printed. Lines 7–8 contain the actual MINISEARCH specification. We will explain the different built-ins in more detail below, but the specification can be read as follows: repeatedly try to find the next solution; and if that is successful, commit to the solution, print it and add the constraint that the objective must have a lower value than the current solution. If unsuccessful, break out of the `repeat`.

All MINISEARCH built-ins are typed as functions returning *annotations*. Semantically, however, every MINISEARCH built-in returns a value that represents either ‘success’ or ‘failure’, with respect to finding a solution. The handling of these implicit return values is done by the MINISEARCH interpreter (Sec. 4.1).

2.1 MINISEARCH built-ins involving the solver

Communication with the solver is restricted to three forms: invoking the solver, adding constraints/variables to the model, and scopes for temporary variables and constraints.

Invoking the solver. The MINISEARCH instruction for finding the next solution is `next()`. It is successful if a solution has been found, and fails otherwise. The variable/value labelling strategy (such as first-fail on the smallest domain value) can be set in two ways: either by annotating the `solve` item (as in standard MINIZINC), which sets the labelling globally, for every call to `next()`. Otherwise, by annotating any MINISEARCH function call, such as `minimize_bab`, with a labelling strategy. Note, however, that as in MINIZINC, solvers may ignore these annotations, for example if the labelling is not supported.

Solvers may declare support for native meta-search strategies, as with global constraints, in which case these MINISEARCH functions are treated as built-ins.

Adding constraints and variables. A constraint is added by calling the `post()` built-in with a constraint as argument. Constraints can be formulated using the same MINIZINC constructs as in the model, including global constraints, user-defined functions and predicates. Variables can be dynamically added during search too, using the MINIZINC `let` construct (see AND/OR search in Sec. 3.3).

Search Scopes. Search scopes define the lifespan of constraints and variables in the model. MINISEARCH has an implicit *global* search scope that contains all variables and constraints of the model. A new search scope can be created by using the `scope(s)` built-in that takes a MINISEARCH specification `s` as an argument. When entering a scope, search is assumed to start from the root again. Whenever the execution leaves a scope, all constraints and variables that were added in the scope are removed from the model and the solver. Execution in the enclosing scope resumes from the point where it left off.

2.2 MINISEARCH control built-ins

All MINISEARCH built-ins have an implicit return value that represents either ‘success’ (*true*) or ‘failure’ (*false*). Using this concept, we introduce MINISEARCH control built-ins. All built-ins execute their arguments in order.

And, Or, Repeat. The `/\`-built-in runs its arguments in order and stops to return *false* as soon as one of its arguments fails. Similarly, the `\/`-built-in stops and returns *success* as soon as one of its arguments succeeds. Existing control mechanisms of MINIZINC such as `if then else endif` expressions can be used as well. The `repeat(s)` built-in takes a MINISEARCH specification `s` and repeats it until a `break` built-in is executed; returns *false* if a break happened, otherwise returns what `s` returned. The delimited variant `repeat(i in 1..N)(s)` will execute `s` for `N` iterations (or until `break` is executed).

Time-Limits. The built-in `time_limit(ms, s)` imposes a time limit `ms` (in milliseconds) on any MINISEARCH specification `s`. This way, `s` stops whenever

the time limit is reached, returning its current status. Time-limits are handled transparently by the MINISEARCH kernel as an exception.

Assigning values to constants. In standard MINIZINC constant parameters such as `int: N=10;` cannot change their value. However, in MINISEARCH we often want to change constants across different iterations. For this purpose, we added the assignment operator `:=` which may only be used inside a MINISEARCH specification. It overwrites that constant’s current value by the value supplied.

2.3 Solution management

The strength of any meta-search language lies in using intermediate solutions to guide the remaining search. For instance, branch-and-bound needs to access the objective to post further constraints, and a Large Neighbourhood Search thaws some of the variables in a solution to continue in that neighbourhood.

To facilitate working with solutions, the most recently found solution is always accessible in MINISEARCH using the `sol` built-in, where `sol(x)` returns the value of `x` in the last solution. MINISEARCH also provides a `hasSol()` built-in to test whether a solution exists.

User-defined functions. When a MINISEARCH strategy is defined as a MINIZINC function, a local solution scope is created. This means that any solution found by a call to `next()` inside the function is visible for the code in the function body, but not for the caller of the function when the function returns. This architecture allows for calls to `next()` to be *encapsulated*, i.e., a function can make “tentative” calls to `next` in a nested search scope and only commit if these succeed. Sect. 3.3 shows how AND/OR search can be implemented based on this principle. In order to make the current solution accessible to the caller, the function must call the `commit` built-in. A function returns ‘success’ if it called `commit()` at least once, and ‘failure’ otherwise, and the last solution committed by the function will then become the current solution of the caller.

Printing Solutions & Debugging. The `print()` function without any arguments prints the last found solution in the format specified in the model’s `output` item. Alternatively, `print(s)` provides more fine-grained control over the output. It prints the string `s`, which can be constructed dynamically from values in the solution using calls to `sol`. MINISEARCH can be debugged using `print()` and MINIZINC’s `trace()` function to display values of parameters, variables, and arbitrary expressions during search. Furthermore, the MINISEARCH interpreter uses the C++ stack, so C++ debuggers can be used to follow the meta-search.

2.4 A library of search strategies

Using the MINISEARCH built-ins, we have defined and collected the most common meta-search approaches in the standard library `minisearch.mzn`.⁴ These meta-search approaches can be used within any MINIZINC model that includes the library. In the next section we present some of these meta-searches in detail.

⁴ The MINISEARCH library is part of the open-source MINISEARCH implementation [14]

3 MINISEARCH examples

Despite MINISEARCH’s limited communication with the solver, it provides enough power to implement many useful complex searches that we illustrate here.

3.1 Lexicographic BaB

In multi-objective optimisation, lexicographic optimisation can be used if the objectives can be ranked according to their importance. The idea is to minimise (or maximise) an array of objectives lexicographically. Lexicographic optimisation can be more efficient than the commonly used approach of obtaining a single objective term by multiplying the components of the lexicographic objective with different constants, as the latter approach leads to large objective values, and potentially overflow.

Analogous to the implementation of branch-and-bound we post the global constraint `lex_less` so that the next solution is lexicographically smaller than the previous one. Below we show the respective MINISEARCH specification.

```
1 function ann: minimize_lex(array[int] of var int: objs) =
2   next() /\ commit() /\ print() /\
3   repeat( scope(
4     post(lex_less(objs, [sol(objs[i]) | i in index_set(objs)])) /\
5     if next() then commit() /\ print() else break endif ) );
```

In line 2 we search for an initial solution and, if successful, repeatedly open a new scope (line 3). Then, we post the lexicographic (lex) constraint (line 4) and search for another solution in line 5. This way, in each iteration of `repeat`, we add one lex constraint, and all previously added lex constraints are removed due to the scope. This is not required but beneficial, since posting several lex-constraints can cause overhead if many intermediate solutions are found.

3.2 Large Neighbourhood Search (LNS)

Large area neighbourhood search (LNS) [20] is an essential method in the toolkit of CP practitioners. It allows CP solvers to find very good solutions to very large problems by iteratively searching large neighbourhoods (close) to optimality.

Randomised LNS explores a random neighbourhood of a given size, which can be surprisingly effective in practice as long as the neighbourhood size is chosen correctly. Cumulative scheduling is an example of a successful application [7].

The following MINISEARCH specification of randomised LNS takes the objective variable, an array of decision variables that will be searched on, the number of iterations, the destruction rate (the size of the neighbourhood) and a time limit for exploring each neighbourhood. We have two scopes: in the global scope, we post BaB style constraints (line 10); in the sub-scope (line 5), we search the neighbourhoods. The predicate `uniformNeighbourhood` defines the neighbourhood: for each search variable we decide randomly whether to set it to its solution value of the previous solution (line 15).

```

1 function ann: lns(var int: obj, array[int] of var int: vars,
2                 int: iterations, float: destrRate, int: exploreTime) =
3   repeat (i in 1..iterations) (
4     print("Iteration "++show(i)+"\n") /\
5     scope (
6       post(uniformNeighbourhood(vars,destrRate)) /\
7       time_limit(exploreTime, minimize_bab(obj)) /\
8       commit() /\ print()
9     ) /\
10    post(obj < sol(obj))
11  );
12 predicate uniformNeighbourhood(array[int] of var int: x, float: destrRate) =
13   if hasSol() then
14     forall(i in index_set(x)) (
15       if uniform(0.0,1.0) > destrRate then x[i] = sol(x[i]) else true endif )
16   else true endif;

```

Adaptive LNS modifies the neighbourhood size over the course of the iterations, depending on the success of previous iterations. Below is a simple variant, where the neighbourhood size parameter `nSize` (line 3) is step-wise enlarged each time no solution is found (line 9). The `fail` command fails the current conjunction and will hence avoid that the `post` command on line 10 is executed.

```

1 function ann: adaptive_lns(var int: obj, array[int] of var int: vars,
2                           int: iterations, int: initRate, int: exploreTime) =
3   let { int: nSize = initRate, int: step = 1; } in
4   repeat (i in 1..iterations) (
5     print("Iteration "++show(i)+" , rate="++show(nSize)+"\n") /\
6     scope ( ( post(uniformNeighbourhood(vars,nSize/100.0)) /\
7              time_limit(exploreTime, minimize_bab(obj)) /\
8              commit() /\ print()
9            ) \\/ (nSize := nSize + step /\ fail) )
10    /\ post(obj < sol(obj)) );

```

Custom Neighbourhoods can sometimes be effective if they capture some insight into the problem structure. For instance, in a Vehicle Routing Problem (VRP), we might want to keep certain vehicle tours or vehicle-customer assignments. Below we show such a custom neighbourhood that is easily specified in MINISEARCH. The predicate `keepTour` (line 1) posts the tour constraints of a given vehicle number `vNum`. If a solution exists, the neighbourhood predicate (line 4) determines the number of customers of each vehicle (line 7), and then randomly chooses to keep the vehicle's tour (line 8) where a high customer usage results in a higher chance of keeping the vehicle. This predicate can be used instead of the uniform neighbourhood in the LNS specifications above.

```

1 predicate keepTour(int: vNum) =
2   forall (i in 1..nbCustomers where sol(vehicle[i]) == vNum)
3     ( successor[i] = sol(successor[i]) );
4 predicate vehicleNeighbourhood() =
5   if hasSol() then
6     forall (v in 1..nbVehicles) (
7       let {int: usage = sum(c in 1..nbCustomers) (sol(vehicle[c]) == v) } in
8       if usage > uniform(0,nbCustomers) then
9         keepTour(v) % higher usage -> higher chance of keeping the vehicle
10        else true endif
11    )
12   else true endif;

```

3.3 AND/OR Search

Search in CP instantiates variables according to a systematic variable labelling, corresponding to an OR tree. AND/OR search decomposes problems into a master and several conjunctive slave sub-problems. An AND/OR search tree consists of an OR tree (master problem), an AND node with one branch for each sub-problem, and OR trees underneath for the sub-problems. A prominent example is stochastic two-stage optimisation [18], where the objective is to find optimal first-stage variable assignments (master problem) such that all second-stage variable assignments (sub-problems) are optimal for each scenario. AND/OR search for stochastic optimisation is called *policy based search* [27]. AND/OR search is also applied in other applications, such as graphical models [11].

Below is a MINISEARCH example of AND/OR search for stochastic two-stage optimisation. The variables and constraints of each scenario (sub-problem) are added incrementally during search.

```

1 function ann: policy_based_search_min(int:sc) =
2   let {
3     array[1..sc] of int: sc_obj = [ 0 | i in 1..sc ];
4     int: expectedCosts = infinity;
5   } in (
6     repeat (
7       if next() then % solution for master
8         repeat (s in 1..sc) (
9           scope( % a local scope for each subproblem
10            let {
11              array[int] of var int: recourse; % subproblem variables
12              var 0..maxCosts: scenarioCosts;
13            } in (
14              post(setFirstStageVariables() /\ % assign master variables
15                 secondStageCts(s,recourse)) /\ % subproblem constraints
16              if minimize_bab(scenarioCosts) then
17                sc_obj[s] := sol(scenarioCosts)
18              else print("No solution for scenario "++show(s)+"\n") /\
19                break endif
20            ) )
21          ) /\ % end repeat
22        if expectedCosts > expectedValue(sc_obj) then
23          expectedCosts := expectedValue(sc_obj) /\
24            commit() % we found a better AND/OR solution
25        else skip() endif
26        else break endif % no master solution
27      ) );
28 % the following predicates are defined in the model according to the problem class
29 predicate setFirstStageVariables();
30 predicate secondStageCts(int: scenario, array[int] of var: y);
31 function int: expectedValue(var int: sc_obj);

```

Lines 3-4 initialise parameters that represent the costs for each scenario/sub-problem and the combined, expected cost of the master and subproblems. Line 7 searches for a solution to the master problem (OR tree), and if this succeeds, we continue with the AND search by finding the optimal solution for each sub-problem, based on the master solution (line 9): we create each subproblem in a local scope, and add the respective variables (line 11- 12) and constraints (line 15). Furthermore, we set the master variables to the values in the master solution (line 14). Then we search for the minimal solution for the scenario (line 16), and, if successful, store it in `sc_obj` (line 17). If we find a solution for each scenario,

then we compute the combined objective (`expectedValue`) and compare it to the incumbent solution (line 22). If we found a better solution, we store its value (line 23) and commit to it (line 24) and continue in the master scope (line 7). Otherwise, if we find no solution for one of the scenarios (line 18), the master solution is invalid. We therefore break (line 19) and continue in the master problem scope, searching for the next master solution (line 7).

3.4 Diverse solutions

Sometimes we don't just require a satisfying or optimal solution, but a diverse set of solutions [9]. The MINISEARCH specification below implements the greedy approximation method that iteratively constructs a set of K diverse solutions:

```

1 function ann: greedy_maxDiverseKset(array[int] of var int: Vars, int: K) =
2   let { array[int,int] of int: Store = array2d(1..K, index_set(Vars),
3                                             [0 | x in 1..K*length(Vars)]),
4       int: L = 0 % current length
5   } in
6   next() /\ commit() /\ print() /\
7   repeat (
8     L := L+1 /\ repeat(i in 1..length(Vars)) (Store[L,i] := sol(Vars[i])) /\
9     if L <= K then
10      scope (
11        let { var int: obj; } in
12        post(obj = sum(j in 1..L, i in index_set(Vars)) (Store[j,i] != Vars[i])) /\
13        maximize_bab(obj) /\ commit() /\ print() )
14      else print(show(Store)+"\n") /\ break endif
15  );

```

The first few lines initialise the `Store`, which will contain the K diverse solutions, as well as the current *length* of the store up to which it already contains solutions. On line 8 the length is increased and the previously found solution is saved. If the length does not exceed K , we construct a new objective on line 12. This objective expresses how different a solution is from the previously found solutions using the Hamming distance. This objective is then maximised, resulting in the next most diverse solution, and the process repeats.

3.5 Interactive Optimisation

Interactive optimisation lets users participate in the solving process by inspecting solutions, adding constraints, and then re-solving. This has been shown to improve the trust of end-users into decision support systems, and a number of successful applications have been implemented [2].

MINISEARCH supports interactive optimisation by calling MINIZINC built-ins that ask for user input. The MINIZINC library contains functions such as `read_int()` which accept keyboard input. More advanced input facilities can be added through user-defined MINIZINC built-ins that execute arbitrary C++ code, such as consulting a user through a graphical user interface or other means.

The following is an example of an interactive Vehicle Routing Problem solver implemented using MINISEARCH. We use a Large Neighbourhood search where in every N -th iteration, the user can choose a vehicle route that should be kept. This code can be used with the LNS implementations described earlier.

```

1 predicate interactiveNeighbourhood(int: iteration, int: N) =
2   if iteration mod N = 0 then
3     let { string: msg = "Enter a vehicle tour to keep (0 for none):\n";
4         int: n = read_int(msg)
5     } in if n > 0 then keepTour(n) else true endif
6   else true endif;

```

4 The MINISEARCH kernel

This section describes the architecture of the MINISEARCH *kernel*, the engine that executes MINISEARCH specifications, interacting with both the MINIZINC compiler and the backend solver.

First, let us briefly review how MINIZINC models are solved. The MINIZINC compiler (usually invoked as `mzn2fzn`) takes text files containing the model and instance data and compiles them into FLATZINC, a low-level language supported by a wide range of solvers. The compiler generates FLATZINC that is specialised for the capabilities of the particular target solver using a solver-specific *library* of predicate declarations. Solvers read FLATZINC text files and produce text-based output for the solutions. Version 2.0 of MINIZINC is based on the `libminizinc` C++ library, which provides programmatic APIs, eliminating the need to communicate through text files. The library also defines an API for invoking solvers directly in C++. MINISEARCH is built on top of these APIs.

4.1 The MINISEARCH Interpreter

The MINISEARCH kernel implements an *interpreter* that processes MINISEARCH specifications and handles the communication between the MINIZINC compiler and the solver. The interpreter assumes that the solver interface is *incremental*, i.e. variables and constraints can be added dynamically during search. We provide an *emulation layer* (see Sec. 4.2) for solvers that do not support incremental operations, including solvers that cannot be accessed through the C++ interface.

The MINISEARCH interpreter is a simple stack-based interpreter that maintains the following state: A stack of *solutions*, one for each function scope; a stack of *time-outs* and *breaks*; and a stack of *search scopes*, containing the solver state of each scope. The interpreter starts by compiling the MINIZINC model into FLATZINC, and then interprets each MINISEARCH built-in as follows:

- `next()` invokes the solver for a new solution; if successful, it replaces the top-most solution on solution stack. If a time-out has been set, the call to the solver is only allowed to run up to the time-out. Returns `true` iff a new solution was found.
- `commit()` replaces the parent solution (2nd on the stack) with the current solution. This commits the current solution into the parent function scope. Returns `true`. Aborts if not in a function call.
- *function calls* duplicate the top of the solution stack before executing the function body. Return `true` if the function committed a new solution.

- `time_limit(l, s)` adds a new time-out `now+1` to the stack, executes `s`, and then pops the time-out. During the execution of `s`, calls to `next` and `repeat` check whether any time-outs `t` have expired (`t > now`), and if so they immediately break. Returns whatever `s` returned.
- `repeat(s)` pushes a break scope with a Boolean flag on the stack, then repeats the execution of `s` as long as the break flag is false and pops the break scope when `s` is finished. The break construct sets the break flag in the current break scope (similar to a time limit).
- `post(c)` compiles the MINIZINC expression `c` into FLATZINC. The compilation is incremental, i.e., the result of compiling `c` is added to the existing FLATZINC. The interpreter then adds the newly generated variables and constraints to the current solver instance.
- `scope(s)` creates a new local scope. The current implementation copies the flat model and creates a new solver instance based on this flat model.
- Other operations (`/\,\/,print`) are interpreted with respect to their semantics.

4.2 Emulating advanced solver behaviour

A key goal of this work is to make MINISEARCH available for any solver that supports FLATZINC. Current FLATZINC solvers, however, neither support incrementally adding constraints and variables, nor do they implement the `libminizinc` C++ API. We therefore need to *emulate* the incremental API.

In order to emulate incrementality, we can implement the dynamic addition of variables and constraints using a restart-based approach, re-running the solver on the entire updated FLATZINC. To avoid re-visiting solutions, after each call to `next` the emulation layer adds a no-good to the model that excludes the solution that was just found. This emulation reduces the requirements on the solver to simply being able to solve a given FLATZINC model.

In order to emulate the C++ API, we generate a textual representation of the FLATZINC and call an external solver process. The emulator then converts the textual output of the process back into `libminizinc` data structures. Using this emulation of the incremental C++ solver API, any current FLATZINC solver can be used with the MINISEARCH kernel.

4.3 Built-in primitives

Solvers can declare native support for a MINISEARCH search specification. For every MINISEARCH call such as `f(x, y, z)`, the kernel will check whether the declaration of `f` has a function body. If it does not have a function body, the function is considered to be a solver built-in, and it is executed by passing the call to a generic `solve` function in the C++ solver API. This is similar to the handling of global constraints in MINIZINC, where solvers can either use their own primitives for a given global constraint, or use the respective MINIZINC decomposition of the global constraint. This way solvers can easily apply their own primitives, but can use alternatives if they do not support a certain feature.

5 Experiments

In our experimental evaluation, we analyse MINISEARCH on practical examples to study the efficiency and compare the overhead of MINISEARCH for different interfaces. Furthermore, we study the benefits of specialised, heuristic MINISEARCH approaches to standard branch-and-bound optimisation; the only available option in solver-independent CP modelling languages.

5.1 Experimental Setup

The problems and instances are taken from the MINIZINC benchmarks repository⁵. The source code of MINISEARCH and the respective solver interfaces will be released in September 2015 [14].

Experiments are run on Ubuntu 14.04 machines with eight i7 cores and 16GB of RAM. The MINISEARCH kernel is based on the MINIZINC 2.0.1 toolchain. The native Gecode interface and incremental C++ interface were implemented using the latest Gecode source code (version 4.3.3+, 20 April 2015). The FZN solvers used are: Gecode (20 April 2015), Choco 3.3.0, Or-tools source (17 February 2015), and Opturion CPX 1.0.2. All solvers use the same variable/value ordering heuristic that is part of the MINIZINC model.

5.2 Overhead of different MINISEARCH interfaces

First, we study the comparative overhead of MINISEARCH with respect to the different interface levels. We do this by analysing the performance of a meta-search approach on the same solver, Gecode. We compare the following four approaches: the solver’s *native* approach (without MINISEARCH) through FLATZINC (Nat-F), the *incremental API* interface to MINISEARCH (MS-Inc) and the *non-incremental text-based* (MS-F) interface to MINISEARCH. We use Rectangle Packing [21] as problem benchmark, and standard BaB as meta-search, since it is available natively in FLATZINC for Gecode.

The results are summarised in Table 2, columns 2–4. It shows the runtimes (in seconds) taken to find the optimal solution for the different approaches using Gecode. The native approach through the FlatZinc interface is fastest, though MINISEARCH through the incremental API interface performs quite similarly. Moreover, MINISEARCH through the FZN interface is, as expected, slower, but only by a small factor. These results are very promising, since they show that the overhead of the MINISEARCH interfaces is low.

In addition, we analyse the overhead of MINISEARCH for other FLATZINC solvers in Table 2, columns 5–10. We ran the Rectangle Packing problem for the FLATZINC solvers or-tools, Choco, and Opturion CPX on a MINIZINC model using both native branch-and-bound (Nat-F), and MINISEARCH branch-and-bound (MS-F), both through the FLATZINC interface. Overall the MINISEARCH FLATZINC interface, while slower, still gives decent performance. Choco seems to incur most overhead of the repeated restarts.

⁵ <https://github.com/MiniZinc/minizinc-benchmarks>

Table 2. Rectangle Packing. Times (sec) averaged over 10 runs. Comparing FLATZINC-solvers with MINISEARCH BaB through FLATZINC (MS-F), MINISEARCH BaB through incremental API (MS-Inc), and native BaB with FLATZINC (Nat-F).

Rectangle Size (N)	Gecode			or-tools		choco		Opturion CPX	
	MS-F	MS-Inc	Nat-F	MS-F	Nat-F	MS-F	Nat-F	MS-F	Nat-F
14	0.7	0.1	0.2	0.5	0.2	6.4	1.0	0.3	0.1
15	0.8	0.3	0.4	0.6	0.3	6.0	1.3	0.4	0.2
16	1.3	0.5	0.5	1.1	0.5	9.1	1.6	0.6	0.2
17	4.5	2.9	2.9	5.2	3.8	16.6	5.7	0.9	0.4
18	37.4	37.4	36.8	41.2	41.5	38.0	29.8	6.7	6.2
19	77.5	43.2	42.0	59.0	32.8	59.2	28.8	12.7	7.7
20	96.4	97.9	94.1	99.2	98.9	93.9	83.0	6.4	5.6
21	472.4	469.7	462.5	416.3	410.0	250.1	239.2	83.8	82.1

5.3 Heuristic search

MINISEARCH strategies are not restricted to complete search. For example, different variants of Large Neighbourhood Search can be tried. To demonstrate the benefits of having heuristic search at hand, we compare the Randomised LNS approach from Section 3.2 with the standard out-of-the-box branch-and-bound approach for a selection of FLATZINC solvers on the Capacitated Vehicle Routing problem (CVRP).

We run each solver on (1) the standard MINIZINC model that uses branch-and-bound, and (2) the MINIZINC model with the MINISEARCH specification of Randomised LNS. We use the Augerat et al [1] CVRP instance sets A,B and P, where the 27 Set-A instances contain random customer locations and demands, the 22 Set-B instances have clustered locations and the 23 Set-P instances are modified versions of instances from the literature. The LNS approach uses a destroy rate of 0.3 (the neighbourhood size is 30%) and an exploration timeout (for each neighbourhood) of 5 seconds. Both the LNS and BaB approach have an overall 2 minute timeout, and we report the quality of the best solution found on time-out.

We can see that LNS provides better solutions than BaB, and Table 3 reports the *average* improvement of the objective of LNS over BaB on the three instance sets. We observe an improvement between 8-15%, which is considerable within the given time limit. Furthermore, we see that the improvements are similar for all solvers across the instances sets, in particular for the Set-B instances. This demonstrates the added value of MINISEARCH search strategies to the existing out-of-the-box optimisation approaches of standard modelling languages.

6 Related work

The starting point for MINISEARCH is search combinators [19], and MINISEARCH is indeed a combinator language. The MINISEARCH language is more restricted than search combinators. In particular, search combinators can interact with the

Table 3. Average relative improvement of MINISEARCH LNS over standard BaB for the achieved objective within a time limit for different solvers. The figures are averages over the Augerat et al [1] CVRP instance sets A,B and P.

CVRP Instance-Set	Gecode	or-tools	choco	Opt. CPX
Set-A	11.60%	11.76%	11.17%	12.11%
Set-B	13.38%	11.82%	12.62%	14.92%
Set-P	9.78%	10.53%	7.98%	11.35%

search at every node of the tree, essentially *replacing* the solver’s built-in search, while MINISEARCH only interacts with the solver’s search at the granularity of solutions. Examples not possible in MINISEARCH are the use of average depth of failure or custom limits based on how many search nodes with a particular property have been encountered, and heuristics such as limited discrepancy search [8] and iterative deepening. Other search combinator expressions are however expressible with MINISEARCH, as witnessed by the library of strategies presented in this paper.

Neither search combinators nor MINISEARCH support fully programmable search. Especially not variable and value ordering heuristics, which interact with the internal data structures of the solver. These kinds of search strategies are important, but cannot be implemented without deep hooks into the solver.

Objective CP [26] can express complex meta-search at the model level, which can then be executed by any of the Objective CP solvers. For example a CP-style search can be implemented using a MIP solver. So it also provides solver-independent search, but through a much more fine grained interface, supporting interaction at each search node. The fine grained interface however means that adding a new solver to Objective CP is complex, in particular supporting the reverse mapping of information from solver objects to model objects.

Possibly the closest system to MINISEARCH is OPL Script [25], a scripting language for combining models. It essentially also communicates to solvers via models and solutions. Distinct from MINISEARCH is that it is object-oriented and provides models as first-class objects, and hence supports the construction and modification of multiple models, as opposed to our simpler mechanisms. It can communicate more information than just variables and constraints, including dual values, bases, and various statistics. However, the language is tightly linked to the CP/MIP solver ILOG CPLEX CP Optimizer. MINISEARCH on the other hand can use any FlatZinc solver.

AMPL Script [4] is similar in many respects to OPL script, though a more basic scripting language, not supporting function definitions for example. It does not have a concept of scopes, but instead supports adding, dropping and restoring variables and constraints from the model dynamically. It interacts with the solver through requests to solve, which continue until some termination condition arises, like optimality proven or a timeout. Hence its natural interaction level is more like `minimize_bab()` than `next()`.

Any meta-search that can be implemented in MINISEARCH could also be implemented directly in a solver such as Gecode, Choco, Comet, or CP Optimizer, although the implementation burden for doing so could be significant. Similarly, APIs for CP embedded in scripting languages such as Numberjack [10] (through Python) could offer a MINISEARCH-like interface. The benefits of MINISEARCH are that it is programming language- and solver-independent, that it abstracts away solution and timeout handling. Furthermore, meta-search languages are easier to use when they extend the language for expressing constraints, in this case MiniZinc, since this is often a vital component of the search.

7 Conclusion and Future Work

In this paper we present MINISEARCH, a solver-independent lightweight language to specify meta-search in MINIZINC models. The first contribution of MINISEARCH is its *expressiveness*: we showed how to formulate different complex search approaches in MINISEARCH, such as Large Neighbourhood Search or AND/OR search. It reuses the MINIZINC modelling language for specifying constraints and variables to add during search. It even enables the implementation of interactive optimisation approaches where the user guides the search process.

The second contribution is the *minimal interface* of MINISEARCH. Since the kernel interprets all statements, the communication between the solver and MINISEARCH is only necessary at every *solution* (instead of at every search *node* as in search combinators). This makes it more realistic to implement for solver developers. Furthermore, in our experiments we observed that both the incremental and the text-based interface provide satisfactory performance.

The third and most important contribution is *solver-independence*. In contrast to existing search languages, MINISEARCH can already use any solver that can process FLATZINC, which is the majority of CP solvers. This allows users to test different solvers with complex meta-searches without having to commit to one single solver.

For future work, we plan to extend MINISEARCH with parallelism, so that independent searches can be executed in parallel and advanced parallel meta-search approaches such as Embarrassingly Parallel Search [17] can be specified. Furthermore, we plan to integrate portfolio search into MINISEARCH, so that different search functions can be executed by different solvers.

Acknowledgements NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. T.G. is supported by a Postdoc grant from the Research Foundation – Flanders.

References

1. Augerat, P., Belenguer, J., Benavent, E., Corberan, A., Naddef, D., Rinaldi, G.: Computational results with a branch and cut code for the capacitated vehicle

- routing problem. Technical Report 949-M, Universite Joseph Fourier, Grenoble (1995)
2. Belin, B., Christie, M., Truchet, C.: Interactive design of sustainable cities with a distributed local search solver. In: Integration of AI and OR Techniques in Constraint Programming - 11th International Conference, CPAIOR 2014, Cork, Ireland, May 19-23, 2014. Proceedings. pp. 104–119 (2014)
 3. Flener, P., Pearson, J., Ågren, M.: Introducing ESRA, a relational language for modelling combinatorial problems. In: Principles and Practice of Constraint Programming - CP 2003, 9th International Conference, CP 2003, Kinsale, Ireland, September 29 - October 3, 2003, Proceedings. p. 971 (2003)
 4. Fourer, R., Gay, D.M., Kernighan, B.W.: AMPL: A Modeling Language for Mathematical Programming. Cengage Learning (2002)
 5. Frisch, A.M., Harvey, W., Jefferson, C., Hernández, B.M., Miguel, I.: Essence : A constraint language for specifying combinatorial problems. *Constraints* 13(3), 268–306 (2008)
 6. Gent, I.P., Miguel, I., Rendl, A.: Tailoring solver-independent constraint models: A case study with Essence’ and Minion. In: Abstraction, Reformulation, and Approximation, 7th International Symposium, SARA 2007, Whistler, Canada, July 18-21, 2007, Proceedings. pp. 184–199 (2007)
 7. Godard, D., Laborie, P., Nuijten, W.: Randomized large neighborhood search for cumulative scheduling. In: Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS 2005), June 5-10 2005, Monterey, California, USA. pp. 81–89 (2005)
 8. Harvey, W.D., Ginsberg, M.L.: Limited discrepancy search. In: Proceedings of the 14th IJCAI. pp. 607–613 (1995)
 9. Hebrard, E., Hnich, B., O’Sullivan, B., Walsh, T.: Finding diverse and similar solutions in constraint programming. In: Veloso, M.M., Kambhampati, S. (eds.) AAAI. pp. 372–377. AAAI Press / The MIT Press (2005)
 10. Hebrard, E., O’Mahony, E., O’Sullivan, B.: Constraint programming and combinatorial optimisation in Numberjack. In: Lodi, A., Milano, M., Toth, P. (eds.) Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, Lecture Notes in Computer Science, vol. 6140, pp. 181–185. Springer Berlin Heidelberg (2010)
 11. Marinescu, R., Dechter, R.: AND/OR branch-and-bound search for combinatorial optimization in graphical models. *Artif. Intell.* 173(16-17), 1457–1491 (2009)
 12. Marriott, K., Nethercote, N., Rafeh, R., Stuckey, P.J., de la Banda, M.G., Wallace, M.: The design of the Zinc modelling language. *Constraints* 13(3), 229–267 (2008)
 13. Michel, L., Van Hentenryck, P.: The Comet programming language and system. In: Principles and Practice of Constraint Programming - CP 2005, 11th International Conference, CP 2005, Sitges, Spain, October 1-5, 2005, Proceedings. pp. 881–881 (2005)
 14. MiniSearch release. <http://www.minizinc.org/minisearch>
 15. MiniZinc challenge. <http://www.minizinc.org/challenge.html>
 16. Nethercote, N., Stuckey, P., Becket, R., Brand, S., Duck, G., Tack, G.: MiniZinc: Towards a standard CP modelling language. In: Bessiere, C. (ed.) Proceedings of the 13th International Conference on Principles and Practice of Constraint Programming. LNCS, vol. 4741, pp. 529–543. Springer-Verlag (2007)
 17. Régim, J., Rezgüi, M., Malapert, A.: Embarrassingly parallel search. In: Principles and Practice of Constraint Programming - 19th International Conference, CP 2013, Uppsala, Sweden, September 16-20, 2013. Proceedings. pp. 596–610 (2013)

18. Ruszczyński, A., Shapiro, A.: Stochastic Programming. Handbooks in operations research and management science, Elsevier (2003)
19. Schrijvers, T., Tack, G., Wuille, P., Samulowitz, H., Stuckey, P.J.: Search combinatorics. *Constraints* 18(2), 269–305 (2013)
20. Shaw, P.: Using constraint programming and local search methods to solve vehicle routing problems. In: *Principles and Practice of Constraint Programming-CP98*, pp. 417–431. Springer (1998)
21. Simonis, H., O’Sullivan, B.: Search strategies for rectangle packing. In: Stuckey, P. (ed.) *Principles and Practice of Constraint Programming, Lecture Notes in Computer Science*, vol. 5202, pp. 52–66. Springer Berlin Heidelberg (2008)
22. Stuckey, P.J., Feydy, T., Schutt, A., Tack, G., Fischer, J.: The MiniZinc challenge 2008-2013. *AI Magazine* 35(2), 55–60 (2014)
23. Stuckey, P.J., Tack, G.: MiniZinc with functions. In: *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 10th International Conference, CPAIOR 2013, Yorktown Heights, NY, USA, May 18-22, 2013. Proceedings.* pp. 268–283 (2013)
24. Van Hentenryck, P.: *The OPL Optimization Programming Language*. MIT Press, Cambridge, MA, USA (1999)
25. Van Hentenryck, P., Michel, L.: OPL Script: Composing and controlling models. In: Apt, K., Monfroy, E., Kakas, A., Rossi, F. (eds.) *New Trends in Constraints, Lecture Notes in Computer Science*, vol. 1865, pp. 75–90. Springer Berlin Heidelberg (2000)
26. Van Hentenryck, P., Michel, L.: The Objective-CP optimization system. In: Schulte, C. (ed.) *Principles and Practice of Constraint Programming, Lecture Notes in Computer Science*, vol. 8124, pp. 8–29. Springer Berlin Heidelberg (2013)
27. Walsh, T.: Stochastic Constraint Programming. In: van Harmelen, F. (ed.) *ECAI*. pp. 111–115. IOS Press (2002)