The Inez Mathematical Programming Modulo Theories Framework

Panagiotis Manolios^(✉), Jorge Pais, and Vasilis Papavasileiou

Northeastern University, Boston, USA
{pete,jpais,vpap}@ccs.neu.edu

Abstract. Our Mathematical Programming Modulo Theories (MPMT) constraint solving framework extends Mathematical Programming technology with techniques from the field of Automated Reasoning, *e.g.*, solvers for first-order theories. In previous work, we used MPMT to synthesize system architectures for Boeing’s Dreamliner and we studied the theoretical aspects of MPMT by means of the Branch and Cut Modulo T (BC(T)) transition system. BC(T) can be thought of as a blueprint for MPMT solvers. This paper provides a more practical and algorithmic view of BC(T). We elaborate on the design and features of *Inez*, our BC(T) constraint solver. *Inez* is an open-source, freely available superset of the OCaml programming language that uses the SCIP Branch and Cut framework to extend OCaml with MPMT capability. *Inez* allows users to write programs that arbitrarily interweave general computation with MPMT constraint solving.

1 Introduction

The ILP (or, more generally, Mathematical Programming) Modulo Theories (IMT or MPMT) framework accommodates Mathematical Programming (MP) instances, where some variable symbols have meaning in background first-order theories [21]. In previous work, we used this approach to solve systems architectural synthesis problems with hard real-time constraints for Boeing and we introduced the Branch and Cut Modulo T (BC(T)) architecture for solving MPMT [17, 21]. BC(T) combines Branch and Cut (B&C) with theory reasoning. B&C is the most established family of algorithms for solving ILP instances, empowering such powerful solvers as CPLEX [2], Gurobi [3], and SCIP [8].

We have formalized BC(T) as a highly non-deterministic transition system [21]. By abstracting away solver implementation details, the BC(T) transition system captures a wide range of possible implementations, and facilitates theoretical analysis. BC(T) can be thought of a design space for MPMT solvers. Implementing an MPMT solver involves zooming in on a region of this design space, with assorted performance trade-offs. To inform efficient solver design, this paper provides an algorithmic (and more deterministic) view of BC(T).

This research was supported in part by DARPA under AFRL Cooperative Agreement No. FA8750-10-2-0233 and by NSF grants CCF-1117184 and CCF-1319580.

Inez¹ extends the SCIP [8] solver and we show how to implement MPMT on top of a B&C-based solver. We explain as much of the operation of the B&C core as needed to demonstrate where theory solvers fit, with an emphasis on the interface between theory solvers and B&C. We do not cover purely internal operations of either side. For example, we treat Simplex (which handles real relaxations within B&C) purely as a black box. We use congruence closure (CC) as an example of a background procedure. Given that the core operations of CC are well-known [27], our discussion only covers the BC(T)-specific aspects. Our choice of CC is motivated by its wide applicability and by the relatively simple (but not trivial) constraints and algorithms involved.

We provide an overview of the features of **Inez**. Notably, **Inez** provides database techniques for reasoning in the presence of data [22]. **Inez** additionally supports user-provided axioms through *local theory extensions* [30]. **Inez** is implemented in OCaml, and makes extensive use of OCaml language constructs and technologies. In fact, the standard way of interacting with the solver is via scripts in a superset of OCaml. Programming with **Inez** is qualitatively different from programming in a standard programming language because **Inez** allows us to write programs that arbitrarily interleave general computation with MPMT constraint solving. To our knowledge, **Inez** is the first system that allows expressing constraints over uninterpreted functions within a programming language, with minimal syntactic overhead, while providing type-safety.

The rest of the paper is organized as follows. Section 2 introduces our superset of the OCaml language through a worked example, and explains how OCaml facilitates our implementation efforts. Section 3 describes the core BC(T) setup as a set of algorithms, while Sect. 4 discusses extensions on top of this setup. Section 5 provides an overview of related work. We conclude with Sect. 6. In the interest of space, we do not provide experimental results. We refer the interested readers to our previous work for a comparison of **Inez** against SMT solvers on instances from the SMT-LIB [21], and also on database problems [22].

2 The Inez Language

In this section, we introduce some of the most notable features of **Inez** by means of a worked example. We focus on the user-facing aspects of **Inez**, *i.e.*, on its input language, which is a superset of OCaml. Our extensions over OCaml are language constructs (and supporting APIs) for easily expressing logical constraints and seamlessly integrating with the underlying constraint solver. Building on top of OCaml allows us to provide a mixed functional and constraint language that users can utilize to express their models in a compact and self-contained way.

Inez utilizes the Camlp4 framework [1] to extend OCaml by assigning meaning to programs that are *syntactically valid* (*i.e.*, recognized by the unmodified OCaml grammar) but *semantically invalid*. The semantics of programs accepted by unmodified OCaml do not change under **Inez**. This design decision has multiple benefits. First, the syntax of **Inez** programs is natural, given that these

¹ <https://github.com/vasilisp/inez>.

programs are syntactically valid OCaml anyway. Also, there are no additional syntactic constructs to cause trouble for editors and other tools. Finally, the implementation is cleaner, because all that it does is transform Abstract Syntax Trees (ASTs) produced by the Camlp4 parser.

Our integration of OCaml and constraints has great impact from a user perspective. For instance, consider a problem that depends on raw data defined and stored in a different location than the problem code, *e.g.*, in a plain text file, spreadsheet, database or web service. With Inez, data retrieval, data processing, and constraint solving can all happen side by side, in the same environment. We present concrete Inez code that obtains data from a database, defines data structures to store and manipulate this data, and finally produces and solves an MPMT instance. In the interest of succinctness, we omit the data retrieval code. A complete implementation (based on MySQL [4] and the `mysql_protocol` [5] library) can be found online.²

The example is based on a facility location problem [10]. We are given a finite set of locations and a finite set of cities. Each city requires a certain number of units of some product. We have to decide where to place facilities in order to satisfy the needs of the cities, while maximizing our earnings.

```

1 open Script ;;
2 open Core.Std ;;
3
4 let n_cities = db_get_n_cities() ;;
5 let locations = db_get_locations() ;;
6 let revenue = db_get_revenue() ;;
7 let capacity = db_get_capacity() ;;
8 let demand = db_get_demand() ;;
9 let n_locations = Array.length locations ;;

```

Lines 1 and 2 are a typical preamble for Inez scripts. The module `Script` contains useful functions for interacting with Inez, while `Core.Std` refers to Jane Street Core, which is a featureful alternative to INRIA's OCaml base library. Inez uses Jane Street Core internally, and we recommend that Inez scripts also use this library. Lines 4 to 8 perform queries to a database instance to obtain data relevant to the problem: (4) an integer `n_cities` with the number of cities we plan to serve; (5) an array `locations` where each position corresponds to the ZIP code of a potential location; (6) a matrix `revenue`, such that for $0 \leq i < \text{n_cities}$ and $0 \leq j < \text{n_locations}$, `revenue.(i).(j)` represents the revenue of selling to city i a unit of product fabricated at location j ; (7) an array `capacity`, where for $0 \leq j < \text{n_locations}$, `capacity.(j)` is the production capacity for a factory in location j ; and (8) an array `demand` that represents the demand, in units of product, from each city. Finally, we define `n_locations` as the size of the array `locations`, *i.e.*, the number of potential locations.

² <http://www.ccs.neu.edu/home/pete/2015/cav-example.zip>.

```

10 let build =
11   let f _ = fresh_bool_var () in Array.init n_locations ~f ;;
12
13 let production =
14   let f _ =
15     let f _ =
16       let v = fresh_int_var () in
17       constrain (~logic (v >= 0)); v in
18     Array.init n_locations ~f in
19   Array.init n_cities ~f ;;

```

Each city is identified by an integer $c \in [0 \dots n_cities - 1]$. Each location is identified by the corresponding ZIP code in the `locations` array. Line 10 defines an array of size `n_locations`, where each element is an `Inez` Boolean variable, created by the function `fresh_bool_var`. The library function `Array.init` initializes each element of the array, by calling its `f` argument (a function) with the corresponding index as the argument. (In our case, the argument to `f` is ignored, hence the underscore.) Each Boolean variable corresponds to a location and represents whether a facility is built there or not. Similarly, line 13 defines a two-dimensional matrix of `Inez` integer variables. The two dimensions correspond to cities and locations: for each possible pair of city c and location l , `production.(c).(l)` represents the planned production (in units of product) of a factory to-be-built in location l destined to city c . The `Inez`-provided function `constrain` adds a formula to the solver context. In line 17 we constrain each integer variable so that it only takes positive values.

For expressing these constraints, we utilize the `~logic` keyword. `~logic` allows expressing terms and formulas with minimal syntactic overhead. We utilize Camlp4 infrastructure to preprocess applications of `~logic` to ensure that the intended meaning over terms and formulas applies. Specifically, (a) integer literals become `Inez` integer terms; (b) the literals `true` and `false` become formulas; and (c) operators like `+` and `&&` obtain meaning over terms and formulas (as opposed to their standard meaning over OCaml integers and Booleans, respectively). For instance, given integer variables `x` and `y`, `~logic (x + 1)` is an `Inez` term, while `~logic (1 <= y && x <= 0)` is an `Inez` formula. `Inez` integer terms and formulas are regular OCaml values that can be passed around.

```

20 let cost (_ : Int) = (~free : Int) ;;
21
22 for i = 0 to n_locations - 1 do
23   let i = toi i in
24   constrain (~logic (cost i >= hist_lb i &&
25                     cost i <= hist_ub i))
26 done ;;

```

Now consider the following situation. During an early planning phase, the exact cost of building a facility on a given location may be unknown. However, experience from similar previous developments could provide bounds for these

costs. We use an uninterpreted function (UF) to express this. Given an integer representing the ID of a location, the UF `cost` (Line 20) returns an integer that corresponds to the cost of building a facility on that location. The syntax for UFs follows closely the standard syntax for defining OCaml functions. Inez recognizes the declaration of `cost` as its own responsibility because of the keyword `~free` in the function body. The declaration produces an actual OCaml function `cost` from integer terms to integer terms. (*Integer terms* belong to an OCaml datatype that describes symbolic integer expressions; integer terms differ from OCaml integers.) Functions can also operate over Booleans. (The `Int` annotations could have been omitted, because integer is the default.)

The function `toi` (Line 23) converts an OCaml integer to an integer term. We use `constrain` to bound the return values of `cost` (Line 25). The upper (respectively lower) bound for each location is computed by the OCaml function `hist_ub` (respectively `hist_lb`), which retrieves historical construction data from a database and analyzes the current situation in order to provide estimate bounds for the cost of building. We impose this constraint across all locations by means of a standard OCaml `for` loop (Lines 22-26).

Also, suppose that we have some knowledge about the global building costs for each location and how they compare to one another. That is, given two ZIP codes, we can determine where it is cheaper to build a factory. This knowledge allows us to define an ordering among ZIP codes, and thus assign to each a unique identifier in the range $[0 \dots n_locations - 1]$ such that the ZIP code with *id* 0 is the cheapest location and the one with *id* $n_locations - 1$ is the most expensive. Given such ids, `cost` is monotonically increasing. Monotonicity can be expressed in Inez by means of an axiom, as follows:

```

27 assert_axiom
28   (~forall x (~forall y ([x <= y], cost x <= cost y))) ;;

```

The function `assert_axiom` is used to introduce an axiom. The keyword `~forall` defines two universally quantified variables `x` and `y`. We subsequently provide a list of assumptions (in this case just `x <= y`), followed by a conclusion (`cost x <= cost y`). We provide details on our implementation of axioms in Sect. 4.2.

We subsequently add constraints to ensure that the units produced by each factory that is built will not exceed its capacity, *i.e.*, that

$$\forall l \in locations. \left[\left(\sum_{c \in cities} production[c][l] \leq build[l] * capacity[l] \right) \right].$$

The Inez encoding is

```

29 for l = 0 to n.locations - 1 do
30   constrain
31     (let cities = List.init n.cities ~f:Fn.id
32      and f c = production.(c).(l) in
33      ~logic

```

```

34 | (sum cities ~f <= iite build.(1) (toi capacity.(1)) 0))
35 | done ;;

```

We notably use the **lnez**-provided function **sum** (Line 34), to express the sum of **lnez** terms resulting from the application of the function **f** on each element of the list **cities**. Additionally, we use the function **iite** (Line 34) that encodes an if-then-else condition. The first argument to **iite** is an **lnez** formula, while the second and third are **lnez** integer expressions for each possible case. Our application of **iite** ensures that, if a factory is built on location **l**, then we obtain the capacity from the corresponding array, otherwise the capacity is zero.

Our concrete example additionally enforces that the demands of each city are satisfied, which can be expressed mathematically as:

$$\forall c \in \text{cities}. \left[\left(\sum_{l \in \text{locations}} \text{production}[c][l] \right) = \text{demand}[c] \right]$$

In the interest of brevity, we omit the corresponding **lnez** code.

Finally, we define an objective function, which is to maximize the earnings, *i.e.*, the total revenue minus the cost of building the factories. The corresponding **lnez** code specifies the optimization criterion by means of the **maximize** function (and re-uses constructs that we have described already):

```

36 | maximize
37 | (let cities = List.init n_cities ~f:Fn.id
38 |   and locations = List.init n_locations ~f:Fn.id in
39 |   let s1 =
40 |     ~logic (sum cities ~f:(fun c -> sum locations ~f:(fun l ->
41 |       revenue.(c).(1) * production.(c).(l)))
42 |   and s2 =
43 |     ~logic (sum locations ~f:(fun l ->
44 |       iite build.(1) (cost (toi l)) 0)) in
45 |   ~logic (s1 - s2)) ;;
46 |
47 | solve_print_result () ;;

```

Line 47 starts the solving process and prints the result (which can be one of **opt**, **sat**, **unsat**, **unbounded**, or **unknown**) to the standard output. Note that our example builds a single set of constraints, and calls the underlying solver once. In general, **lnez** provides an incremental *push/pop* interface that allows the user to add and remove constraints, and perform multiple queries. As an example, consider that for the presented problem we had two different optimization criteria: first maximize the earnings and second minimize the number of factories. One could achieve this by *push*-ing the first maximization criterion, solving the problem, registering the maximum value obtained, and finally *pop*-ing the criterion. One could then add a constrain that restricts the first criterion to be equal to the registered value and minimize the second criterion. The full power of OCaml is available to determine future steps by examining intermediate results. **lnez** thus provides a framework for constraint-based algorithms.

We conclude this section with an overview of the OCaml features that we have utilized to provide the functionality described in this section. Interestingly, (a) *Generalized Algebraic Data Types (GADTs)* [18] allow us to represent terms and formulas in a type-safe way; (b) the extensibility of Inez is reflected on the *module system*, i.e., extending the backend amounts to instantiating a *functor* that given a theory solver (wrapped up as a module) produces a solver for the resulting logic (i.e., another module); (c) the *toplevel system* allows us to build custom read-evaluate-print loops that interactively interpret OCaml plus our logic fragments; finally, (d) `camlid1` enables relatively seamless interaction with C/C++ code (like SCIP and our implementation of CC).

3 An Algorithmic View of $BC(T)$

This section provides a set of interconnected algorithms that together implement $BC(T)$. We thus document the architecture empowering the backend of Inez. The algorithms primarily operate upon *nodes* and sets thereof. Each node is described by a set of *integer linear constraints*, i.e., constraints of the form $c_1 \cdot v_1 + \dots + c_n \cdot v_n \{< | \leq | = | \geq | >\} c$, where c_i are integer constants, v_i are variable symbols, and the right-hand side c is an integer constant. While we provide support for *mixed integer linear constraints* (i.e., integer and real variables side-by-side) through an experimental version of Inez, our discussion focuses on the integer case for simplicity. A node characterizes an open subproblem that needs to be explored. Nodes also carry metadata, like known variable bounds.

In addition to the integer linear constraints, the input to the solver contains uninterpreted function (UF) constraints. We assume that variable abstraction [20] has happened as a preprocessing step, resulting in linear constraints that do not involve UF terms. The `definitions` symbol that is used in the pseudocode stands for a collection of atomic formulas of the form $v = f(l)$, where v is a variable symbol, f is a UF symbol, and l is a list of arguments of the form $w + k$, where w is a variable symbol and k is an integer constant. Entirely concrete terms are a special case that can be encoded with a single integer variable fixed to zero. UF terms thus involve limited arithmetic, as is common practice [27]. `definitions` is an immutable global constant.

Our pseudocode uses *sum types* (also known as *tagged unions*) for some of the variables. Sum types have multiple *constructors* that correspond to different cases for the values carried. The constructor of a particular element serves as a tag denoting which case the element belongs in. Furthermore, the magic constant `*` stands for non-deterministic Boolean choice. `*` is used in conditionals where heuristics apply. $\langle e \rangle$ denotes that standard operators within e are to be interpreted over syntactic objects, e.g., $\langle v - w \rangle$ is not a concrete integer or real, but a term representing the subtraction of w from v . We follow a generally applicative style, e.g., operations that modify a node (by producing new linear constraints and bounds) produce a new node. Our presentation is top-down.

Our CC solver is implemented by the functions with suffix `_cc` (`check_cc`, `enforce_cc`, `propagate_cc`, and `branch_cc`). These are the functions that need to be replaced (or enhanced) for supporting a theory other than $\emptyset$.

```

function bc( $p$  : node) : (Unsat|Optimal(assignment))
   $P \leftarrow \{p\}$ 
   $\alpha \leftarrow \text{None}$ 
  while  $P \neq \emptyset$  do
     $q \leftarrow \text{pick}(P)$ 
     $P \leftarrow P \setminus \{q\}$ 
    match solve_node( $q$ , obj( $\alpha$ )) with
      case Unsat
      | noop() // do nothing
      case Solved( $\beta$ )
      |  $\alpha \leftarrow \text{Some}(\beta)$ 
      case Branched( $Q$ )
      |  $P \leftarrow P \cup Q$ 
    match  $\alpha$  with
      case None
      | return Unsat
      case Some( $\beta$ )
      | return Optimal( $\beta$ )

```

3.1 High-Level Functions

The top-level B&C procedure, **bc**, accepts as its argument a set of integer linear constraints, p . p corresponds to the root node of the B&C search tree. **bc** keeps track of a set P of nodes to be examined (initialized with $\{p\}$). α carries a candidate satisfying (integer) assignment. α , belonging in a sum type, is of the form **Some**(β) if an assignment β is known; α is **None** otherwise. The loop body in **bc** picks one of the remaining nodes in P and processes it by calling **solve_node**. The implementation of **pick** (not provided) may involve sophisticated heuristics for the choice of next node to be examined. A bias towards the children of the node that was more recently branched upon (*i.e.*, depth-first search) is common.

solve_node receives as arguments the node p to be processed, in addition to an upper bound l for the objective values of the solutions of interest. l corresponds to an already-known solution. (We assume that the function **obj** that computes l and our comparisons with l take care of the possibility of no known solution or unbounded solution, by supporting the special constants $+\infty, -\infty$.) **solve_node** performs three processing stages: (a) propagation (Sect. 3.2); (b) solving a continuous relaxation of the linear constraints; and (c) enforcing constraints against a relaxation-obtained solution (Sect. 3.3). Enforcing may result in branching (Sect. 3.4). The aforementioned stages operate on one node at a time, always called p in the respective functions, while their output may be multiple nodes, as a result of branching.


```

function
solve_node( $p : \text{node}, l : \text{int}$ ) : (Unsat|Solved(assignment)|Branched({node}))
  match propagate( $p$ ) with
    case Unsat
      | return Unsat
    case Unmodified
      | noop()
    case Modified( $p'$ )
      |  $p \leftarrow p'$ 
  match solve_relaxation( $p$ ) with
    case Unsat
      | return Unsat
    case Modified( $p'$ )
      | return solve_node( $p', l$ )
    case Solved( $\alpha$ )
      | if obj( $\alpha$ )  $\geq l$  then
          | return Unsat
      | else
          match enforce( $p, \alpha$ ) with
            case Sat
              | return Solved( $\alpha$ )
            case Unsat
              | return Unsat
            case Modified( $p'$ )
              | return solve_node( $p', l$ )
            case Branched( $P$ )
              | return Branched( $P$ )

```

3.2 Propagation

The function `propagate` attempts to reduce the domain of variables. In the process of doing so, it may detect infeasibility (response `Unsat`); if it succeeds, `propagate` returns a version p' of the original node p modified with new bounds (`Modified( $p'$ )`); `Unmodified` means that no propagation was possible, neither was the function able to detect infeasibility. The implementation we provide combines ILP (`propagate_ilp`) and CC (`propagate_cc`) propagation techniques. We do not explain `propagate_ilp`, which is internal to the ILP solver, and as such orthogonal to $\text{BC}(T)$. Either kind of propagation can be skipped. We repeatedly perform propagation, until a fixpoint is reached, or until a heuristic for termination returns `true`, *e.g.*, after a fixed number of rounds. In practice, SCIP employs various *constraint handlers* that provide propagation procedures of different *priority*. The top-level propagation procedure takes into account priorities to combine the sub-procedures.

```

function propagate (p : node) : (Unsat|Modified(node)|Unmodified)
  m ← false
  while * do
    if * then
      match propagate_ilp(p) with
        case Unsat
          ⊥ return Unsat
        case Modified(p')
          | p ← p'
          | m ← true
        case Unmodified
          ⊥ noop()
    if * then
      match propagate_cc(p) with
        case Unsat
          ⊥ return Unsat
        case Modified(p')
          | p ← p'
          | m ← true
        case Unmodified
          ⊥ noop()
  return m ? Modified(p) : Unmodified

```

`propagate_cc` is described in a declarative way. Our concrete implementation is similar to the CC procedures in SMT, and therefore takes offsets into account [27]. `equalities(p)` stands for known equalities of the form $v = w + k$, where v and w are integer variables and k is an integer constant. We implement this by defining an auxiliary variable $d_{v,w} = v - w$ for every interesting pair of variables v and w . We can subsequently query whether $d_{v,w}$ is fixed. For any equality $v = w + k$ implied by the already known equalities (conjoined with definitions), we try to fix the upper and lower bound of $v - w$ to k (via the functions `set_lb` and `set_ub` that provide an interface to the ILP solver), and report unsatisfiability if this is impossible. The outer forall statement should be read as a declarative specification (*i.e.*, we range over all relevant v, w), not as a suggestion for efficient implementation.

3.3 Enforcing Continuous Relaxations

In `solve_node`, propagation is followed by solving a continuous relaxation. A response `Unsat` for the relaxation implies that the integer constraints of the node (which are strictly harder than the real constraints of the relaxation) are also unsatisfiable. If `solve_relaxation` returns an assignment α (case `Solved( $\alpha$ )`), `solve_node` first checks whether α is better than the already known solution ($\text{obj}(\alpha) < l$), and does not further process the node if not; integer solutions can be

```

function propagate_cc( $p$  : node) : (Unsat|Modified(node)|Unmodified)
   $m \leftarrow \text{false}$ 
  forall  $v, w \in \text{variables}(p)$  do
    if  $\text{equalities}(p) \wedge \text{definitions} \models_{\mathbb{Z}} v = w + k$  for some  $k \in \mathbb{Z}$  then
      match  $\text{set\_lb}(p, \langle v - w \rangle, k)$  with
        case Unsat
           $\perp$  return Unsat
        case Modified( $p'$ )
           $m \leftarrow \text{true}$ 
           $p \leftarrow p'$ 
        case Unmodified
           $\perp$  noop()
      match  $\text{set\_ub}(p, \langle v - w \rangle, k)$  with
        case Unsat
           $\perp$  return Unsat
        case Modified( $p'$ )
           $m \leftarrow \text{true}$ 
           $p \leftarrow p'$ 
        case Unmodified
           $\perp$  noop()
     $\perp$ 
  return  $m ? \text{Modified}(p) : \text{Unmodified}$ 

```

at most as good as the solution to the relaxation. Otherwise, **enforce** is executed. If α is not integer, or if it is theory-inconsistent, **enforce** is responsible for explaining why, *e.g.*, by introducing implied linear constraints violated by α . **enforce** may determine that α satisfies all constraints (response **Sat**), or that the node (and not just α) is infeasible (response **Unsat**). In either of these cases, **solve_node** has solved p . Enforcing may result in learning new linear constraints or bounds (case **Modified**(p')), in which case **solve_node** re-processes the node.

enforce combines different kinds of enforcement, in much the same way that **propagate** combines different kinds of propagation. The part of enforcement that is related to integrality (**enforce_ilp**) may branch around a real solution, or apply cut generation techniques [12, 15]. ILP cut generation techniques are beyond the scope of this paper. Conversely, the implementation of **enforce_ilp** is not shown. We proceed to describe CC enforcement (**enforce_cc**). First, **enforce_cc** calls **check_cc** to check whether α is theory-consistent. **check_cc** reports that α does not satisfy the UF constraints if there exist calls $v = f(l)$ and $v' = f(l')$ of some function f , such that all arguments in the respective positions of the lists of arguments l and l' have the same value under α , but $\alpha(v) \neq \alpha(v')$. **check_cc** then returns the *conflict* (f, v, v', l, l') to explain what is wrong with α . If no conflict is found, **bc** receives α , and α becomes the new candidate solution. In case **check_cc** returns a conflict, **enforce_cc** ensures that

```

function
enforce( $p$  : node,  $\alpha$  : assignment) : (Sat|Unsat|Modified(node)|Branched({node}))
  match enforce_ilp( $p$ ,  $\alpha$ ) with
    case Sat
       $\sqsubseteq$  return enforce_cc( $p$ ,  $\alpha$ )
    case Unsat
       $\sqsubseteq$  return Unsat
    case Modified( $p'$ )
       $\sqsubseteq$  return Modified( $p'$ )
    case Branched( $P$ )
       $\sqsubseteq$  return Branched( $P$ )

```

```

function enforce_cc( $p$  : node,  $\alpha$  : assignment) :
  (Sat|Unsat|Modified(node)|Branched({node}))
  match check_cc( $p$ ,  $\alpha$ ) with
    case Conflict( $c$ )
      match propagate_cc( $p$ ) with
        case Unsat
           $\sqsubseteq$  return Unsat
        case Modified( $p'$ )
           $\sqsubseteq$  return Modified( $p'$ )
        case Unmodified
           $\sqsubseteq$  return Branched(branch_cc( $p$ ,  $c$ ))
    case Sat
       $\sqsubseteq$  return Sat

```

propagation has happened by calling **propagate_cc** again. The latter function may have been skipped during the propagation stage. **enforce_cc** only needs to act further if propagation can neither detect unsatisfiability, nor produce new information. In this case, **enforce_cc** proceeds by branching.

Note that CC enforcement happens after the corresponding method for the ILP constraints. CC enforcement thus only ever deals with integer assignments, which yields cleaner implementation. Additionally, this design prioritizes ILP-related over theory-related operations, thus emphasizing ILP-heavy problems.

3.4 Branching

Branching is what our CC implementation performs when all else fails. Concretely, the following invariant holds when we get to **branch_cc**. There exists an integer solution for the non-theory constraints of p (given that integrality enforcement has succeeded), but the integer bounds that hold for p do not allow any information to be propagated, neither can we deduce unsatisfiability of p .

```

function check_cc( $p$  : node,  $\alpha$  : assignment) : (Sat|Conflict(conflict))
   $m \leftarrow \{\}$  //  $m$  is a map
  foreach  $\langle v = f(l) \rangle \in \text{definitions}$  do
     $c \leftarrow [\alpha(w) + k \mid \langle w + k \rangle \text{ in } l]$  // list comprehension over  $l$ 
    if  $(f, c) \in \text{keys}(m)$  then
       $(v', l') \leftarrow m[(f, c)]$ 
      if  $\alpha(v) \neq \alpha(v')$  then
        return Conflict( $f, v, v', l, l'$ )
    else
       $m[(f, c)] \leftarrow (v, l)$ 
  return Sat

```

When we call `branch_cc` from `enforce_cc`, we have access to a conflict (f, v, v', l, l') . Note that there must be some position $i \in [0, \text{arity}(f) - 1]$ such that the equality $l[i] = l'[i]$ is not implied by the bounds visible to `propagate_cc`. Otherwise, all arguments would have been equal, and `propagate_cc` would have produced the equality $v = v'$, which is violated. It is always possible to branch on whether $l[i] < l'[i]$, $l[i] = l'[i]$, or $l[i] > l'[i]$. If, according to the bounds on $v - v'$, $v = v'$ is a possibility, then we may instead choose to branch on whether $v < v'$, $v = v'$, or $v > v'$. Our branching involves very little guesswork. A conflict (f, v, v', l, l') provides a witness for the gap between the assignment under examination α (which is not feasible with respect to UF) and the more limited information that is available as bounds in p (which do not entail infeasibility). In order to steer the ILP solver away from the problematic assignment α (and other assignments similar to it), we have to examine the aforementioned gap. We do so by branching driven by the conflict.

The branching strategy we outlined is in alignment with the Nelson-Oppen (NO) scheme for combining decision procedures [20, 25]. We branch on pairs of variables that are shared between ILP and UF, *i.e.*, make progress towards an *arrangement* of the shared variables. Such branching will eventually produce subproblems for which CC has all the information on the shared variables that it needs to determine (in)feasibility of the UF constraints (in definitions); similarly, for the UF-feasible subproblems, the ILP solver (with no more input possible from CC) has all the information it needs to apply complete techniques and determine feasibility. We thus guarantee termination of the combination.

4 Extensions

4.1 Propositional Structure

We have so far not discussed integer linear constraints that appear under arbitrary propositional structure. Inez provides such support by utilizing *indicator constraints*. Such constraints have the form $l \Rightarrow \Sigma_{0 \leq i < n} [c_i \cdot x_i] \leq c$, where l is a

```

function branch.cc( $p : \text{node}, (f, v, v', l, l') : \text{conflict} ) : \{\text{node}\}$ 
  if  $* \wedge \text{lb}(p, \langle v - v' \rangle) \leq 0 \wedge \text{ub}(p, \langle v - v' \rangle) \geq 0$  then
     $P \leftarrow \{ \langle p \wedge v = v' \rangle \}$ 
    if  $\text{lb}(p, \langle v - v' \rangle) < 0$  then
       $P \leftarrow P \cup \{ \langle p \wedge v < v' \rangle \}$ 
    if  $\text{ub}(p, \langle v - v' \rangle) > 0$  then
       $P \leftarrow P \cup \{ \langle p \wedge v > v' \rangle \}$ 
    return  $P$ 
  for  $i \in [0, \text{arity}(f) - 1]$  do
    if  $\alpha(l[i]) = \alpha(l'[i])$  then
      if  $\text{lb}(p, \langle l[i] - l'[i] \rangle) < 0 \vee \text{ub}(p, \langle l[i] - l'[i] \rangle) > 0$  then
         $P \leftarrow \{ \langle p \wedge l[i] = l'[i] \rangle \}$ 
        if  $\text{lb}(p, \langle l[i] - l'[i] \rangle) < 0$  then
           $P \leftarrow P \cup \{ \langle p \wedge l[i] < l'[i] \rangle \}$ 
        if  $\text{ub}(p, \langle l[i] - l'[i] \rangle) > 0$  then
           $P \leftarrow P \cup \{ \langle p \wedge l[i] > l'[i] \rangle \}$ 
      return  $P$ 
  assert(false) // unreachable

```

possibly negated Boolean variable, c_i and c are constants, and v_i are variables. Indicator constraints can establish equivalence between a Boolean variable b and an inequality $\sum_{0 \leq i < n} [c_i \cdot x_i] \leq c$ via the constraints $b \Rightarrow \sum_{0 \leq i < n} [c_i \cdot x_i] \leq c$ and $\neg b \Rightarrow -\sum_{0 \leq i < n} [c_i \cdot x_i] \leq -c - 1$. Once we have Boolean variables like b , encoding propositional structure can be done via clauses (which are a special case of integer linear inequalities) in a Tseitin-like fashion.

Indicator constraints can be encoded in terms of integer linear constraints [21], based on a technique that is known as Big-M. SCIP deals with indicator constraints via a specialized *constraint handler* (rather than via Big-M). This handler implements indicator constraints through propagation, enforcing, and branching functions that fit in $\text{BC}(T)$ just like their CC counterparts (Sect. 3).

4.2 Local Theory Extensions

We demonstrate how to support user-provided axioms within $\text{BC}(T)$ and *lnez*. Such axioms constrain newly defined function symbols (beyond the ones in the signature Σ_Z of Linear Integer Arithmetic, *e.g.*, $+$). We thus *extend* QFLIA by axiomatizing new functions. Throughout this section, we assume a first-order signature Σ , comprised of the axiomatized function symbols.

An example of the kinds of axioms we support was given via the axiom in Line 28 of our introductory example. Our axiom is only meaningful as an *extension* of (Integer Linear) Arithmetic. The intended meaning (monotonicity of **cost**) is only achieved because $\leq$ is already constrained by Arithmetic. More generally, we support axioms that are universally quantified disjunctions

of inequalities that may contain function symbols. Our focus on clauses is not a restriction; collections of clausal axioms can be used to encode axioms with more complex structure.

Our implementation of axioms in *Inez* builds upon results on *local theory extensions* [30] that allow us to replace axioms with a finite set of instances thereof (computed based on the set of terms that appear in the formula). In our case, the instantiation procedure produces clauses, where the literals involve arithmetic and the Σ -function symbols.

While in principle we can simply encode the axiom instances of interest as part of the input formula (Sect. 4.1), our implementation applies a specialized procedure that retains the clausal structure. The literals are inequalities, *e.g.*, for our monotonicity example we have inequalities of the form $x \leq y$ and $\text{cost}(x) \leq \text{cost}(y)$ over x and y that appear in the input as arguments to **cost**. By introducing fresh variables, we simplify these literals by rewriting them to the form $v \leq c$, where v is a variable and c is a constant. We then employ a SCIP handler for constraints of the form $\bigvee_i v_i \leq c_i$, that notably employs SAT-like techniques for clauses [6].

4.3 Databases

Inez provides an extension aimed at database analysis [22]. The workhorse of this extension is what we call *table membership* constraints, which have the form

$$(x_1 + c_1, \dots, x_k + c_k) \in \{(y_{1,1} + d_{1,1}, \dots, y_{1,k} + d_{1,k}), \dots, (y_{l,1} + d_{l,1}, \dots, y_{l,k} + d_{l,k})\},$$

where $x_i, y_{i,j}$ are variables and $c_i, d_{i,j}$ are (integer) constants. On top of table membership, *Inez* provides higher-level database-inspired modeling constructs.

Table membership fits in $\text{BC}(T)$ just like CC. Functions **propagate_db** and **enforce_db** replace (or enhance) the corresponding CC functions, while everything else remains unchanged. Design decisions in **enforce_db** resemble the ones in **enforce_cc**, *e.g.*, branching (driven by the data) happens only as a last resort.

5 Related Work

Frontend: Existing projects that enhance programming languages with constraints [7, 19, 31] differ from *Inez* both with respect to the language constructs that they provide and the underlying constraint technology.

Backend: *Inez* seeks to combine the strengths of MP solvers [2, 3, 8] and solvers for first-order theories [24, 25, 29], *e.g.*, as implemented within Lazy SMT [28]. Previous work on similar combinations has focused on improving the arithmetic capabilities [13, 16] of SMT solvers by integrating MP engines [9, 14, 23], and on extending SMT for optimization [11, 26]. MPMT differs by having as its core an MP solver, as opposed to a SAT solver.

6 Conclusions

We provided an overview of the techniques that empower the *Inez* constraint solver. *Inez* is an open-source, freely available system that instantiates the $BC(T)$ architecture for Mathematical Programming Modulo Theories. We described the concrete algorithms used to in *Inez* to efficiently implement $BC(T)$. *Inez* is an extension of OCaml that allows users to write programs that orchestrate arbitrary interleaving between general computation and MPMT constraint solving.

References

1. Camlp4. <https://github.com/ocaml/camlp4/wiki>
2. CPLEX. <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>
3. Gurobi. <http://www.gurobi.com>
4. MySQL. <https://www.mysql.com/>
5. mysql_protocol. https://github.com/slegrand45/mysql_protocol
6. SCIP Constraint Handler for Bound Disjunction Constraints. http://scip.zib.de/doc-3.1.1/html/cons__bounddisjunction_8h.php
7. Z3py. <http://rise4fun.com/z3py/tutorial>
8. Achterberg, T.: SCIP: solving constraint integer programs. *Math. Program. Comput.* **1**(1), 1–41 (2009)
9. Besson, F.: On using an inexact floating-point LP solver for deciding linear arithmetic in an SMT solver. In: *SMT* (2010)
10. Castillo, E., Conejo, A., Pedregal, P., Garca, R., Alguacil, N.: *Building and Solving Mathematical Programming Models in Engineering and Science*. Wiley, Hoboken (2002)
11. Cimatti, A., Franzén, A., Griggio, A., Sebastiani, R., Stenico, C.: Satisfiability modulo the theory of costs: foundations and applications. In: Esparza, J., Majumdar, R. (eds.) *TACAS 2010*. LNCS, vol. 6015, pp. 99–113. Springer, Heidelberg (2010)
12. Cornuejols, G.: Valid inequalities for mixed integer linear programs. *Math. Program.* **112**(1), 3–44 (2008)
13. Dutertre, B., de Moura, L.: A fast linear-arithmetic solver for DPLL(T). In: Ball, T., Jones, R.B. (eds.) *CAV 2006*. LNCS, vol. 4144, pp. 81–94. Springer, Heidelberg (2006)
14. Faure, G., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E.: SAT modulo the theory of linear arithmetic: exact, inexact and commercial solvers. In: Kleine Büning, H., Zhao, X. (eds.) *SAT 2008*. LNCS, vol. 4996, pp. 77–90. Springer, Heidelberg (2008)
15. Gomory, R.E.: Outline of an algorithm for integer solutions to linear programs. *Bull. AMS* **64**, 275–278 (1958)
16. Griggio, A.: A practical approach to satisfiability modulo linear integer arithmetic. *JSAT* **8**, 1–27 (2012)
17. Hang, C., Manolios, P., Papavasileiou, V.: Synthesizing cyber-physical architectural models with real-time constraints. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 441–456. Springer, Heidelberg (2011)
18. Jones, S.P., Vytiniotis, D., Weirich, S., Washburn, G.: Simple unification-based type inference for GADTs. In: *ICFP* (2006)

19. Koksai, A.S., Kuncak, V., Suter, P.: Constraints as Control. In: POPL (2012)
20. Manna, Z., Zarba, C.: Combining Decision Procedures. In: 10th Anniversary Colloquium of UNU/IIST (2002)
21. Manolios, P., Papavasileiou, V.: ILP modulo theories. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 662–677. Springer, Heidelberg (2013)
22. Manolios, P., Papavasileiou, V., Riedewald, M.: ILP Modulo Data. In: FMCAD (2014)
23. Monniaux, D.: On using floating-point computations to help an exact linear arithmetic decision procedure. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 570–583. Springer, Heidelberg (2009)
24. Nelson, G., Oppen, D.: Fast Decision Algorithms Based on Union and Find (1977)
25. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. *TOPLAS* **1**, 245–257 (1979)
26. Nieuwenhuis, R., Oliveras, A.: On SAT modulo theories and optimization problems. In: Biere, A., Gomes, C.P. (eds.) SAT 2006. LNCS, vol. 4121, pp. 156–169. Springer, Heidelberg (2006)
27. Nieuwenhuis, R., Oliveras, A.: Fast congruence closure and extensions. *Inf. Comput.* **205**, 557–580 (2007)
28. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: from an abstract davis-putnam-logemann-loveland procedure to DPLL(T). *JACM* **53**(6), 937–977 (2006)
29. Shostak, R.: A practical decision procedure for arithmetic with function symbols. *JACM* **26**(2), 351–360 (1979)
30. Sofronie-Stokkermans, V.: Hierarchic reasoning in local theory extensions. In: Nieuwenhuis, R. (ed.) CADE 2005. LNCS (LNAI), vol. 3632, pp. 219–234. Springer, Heidelberg (2005)
31. Torlak, E., Bodik, R.: A lightweight symbolic virtual machine for solver-aided host languages. In: PLDI (2014)