

The Inez Mathematical Programming Modulo Theories Framework

Panagiotis Manolios, Jorge Pais, and Vasilis Papavasileiou

Northeastern University *
{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 instances, where some variable symbols have meaning in background first-order theories [19]. 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 [15, 19]. 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 [7].

We have formalized BC(T) as a highly non-deterministic transition system [19]. 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).

*Inez*¹ extends the SCIP [7] 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

* 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.

¹ <https://github.com/vasilisp/inez>

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 [20]. `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 Section 4 discusses extensions on top of this setup. Section 5 provides an overview of related work. We conclude with Section 6.

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 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 `lnez`, data retrieval, data processing, and constraint solving can all happen side by side, in the same environment. We present concrete `lnez` 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 [9]. 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 `lnez` scripts. The module `Script` contains useful functions for interacting with `lnez`, while `Core.Std` refers to Jane Street Core, which is a featureful alternative to INRIA's OCaml base library. `lnez` uses Jane Street Core internally, and we recommend that `lnez` 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 $0 \leq i < n_cities$ and $0 \leq j < n_locations$, `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 $0 \leq j < n_locations$, `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.

```

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

```

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

```

19 |   Array.init n_locations ~f in
20 |   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 14 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 18 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.

```

21 | let cost (- : Int) = (~free : Int) ;;
22 |
23 | for i = 0 to n_locations - 1 do
24 |   let i = toi i in
25 |   constrain
26 |     (~logic (cost i >= hist_lb i && cost i <= hist_ub i))
27 | 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 21) 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 24) converts an OCaml integer to an integer term. We use `constrain` to bound the return values of `cost` (Line 26). 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 23-27).

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:

$$\forall x. \forall y. [x \leq y \Rightarrow \text{cost}(x) \leq \text{cost}(y)] \quad (1)$$

Such a constrain can be expressed on `lnez` by means of an axiom as follows:

```
28 assert_axiom
29 (~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 Section 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

$$\forall l \in \text{locations}. \left[\left(\sum_{c \in \text{cities}} \text{production}[c][l] \leq \text{build}[l] * \text{capacity}[l] \right) \right]. \quad (2)$$

The `lnez` encoding is

```
30 for l = 0 to n_locations - 1 do
31   constrain
32     (let cities = List.init n_cities ~f:Fn.id
33      and f c = production.(c).(l) in
34      ~logic
35      (sum cities ~f <= iite build.(l) (toi capacity.(l)) 0))
36 done ;;
```

We notably use the `lnez`-provided function `sum` (Line 35), 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 35) 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:

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

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:

$$\sum_{c \in \text{cities}, l \in \text{locations}} (\text{revenue}[c][l] \cdot \text{production}[c][l]) - \sum_{l \in \text{locations}} (\text{build}[l] \cdot \text{cost}(l))$$

The corresponding `lnez` code specifies the optimization criterion by means of the `maximize` function (and re-uses constructs that we have described already):

```

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

```

Line 48 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)* [16] allow us to represent terms

and formulas in a type-safe way; (b) the extensibility of `lnez` 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) `camlidl` 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 describe the operation of a BC(T)-based solver. We thus document the architecture empowering the backend of `lnez`.

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 `lnez`, 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 linear constraints, the input to the solver contains UF constraints. We assume that the input has been purified [18], resulting in ILP 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`.

3.1 High-Level Functions

The top-level B&C procedure, `bc`, accepts as its argument a set of linear constraints, p . p corresponds to the root node of the B&C search tree. `bc` keeps

```

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
         $\lfloor$  noop () // do nothing
      case Solved( $\beta$ )
         $\lfloor$   $\alpha \leftarrow \text{Some}(\beta)$ 
      case Branched( $Q$ )
         $\lfloor$   $P \leftarrow P \cup Q$ 
    match  $\alpha$  with
      case None
         $\lfloor$  return Unsat
      case Some( $\beta$ )
         $\lfloor$  return Optimal( $\beta$ )

```

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 $\text{Some}(\beta)$ 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. We expect a bias towards the children of the node that was more recently branched upon (*i.e.*, depth-first search).

`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 special constants $+\infty, -\infty$.) `solve_node` performs three processing stages: (a) propagation (Section 3.2); (b) solving a real relaxation of the linear constraints; and (c) enforcing constraints against a relaxation-obtained solution (Section 3.3). Enforcing may result in branching (Section 3.4). The aforementioned stages operate on one node at a time, always called p in the respective functions. (Their output may be multiple nodes, as a result of branching.)

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


```

function
solve_node(p : node, l : int) : (Unsat| Solved(assignment)| Branched({node}))
  match propagate(p) with
    case Unsat
    | return Unsat
    case Unchanged
    | noop ()
    case Modified(p')
    | p ← p'

  match solve_relaxation(p) with
    case Unsat
    | return Unsat
    case Modified(p')
    | return solve_node(p')
    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')
    | | | case Branched(P)
    | | | | return Branched(P)

```

(Modified(*p'*)); Unchanged 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. 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.

`propagate_cc` is described in a declarative way. Our concrete implementation is similar to the CC procedures in SMT. `propagate_cc` takes offsets into account [25]. `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

```

function propagate( $p$  : node) : (Unsat| Modified(node)|Unchanged)
   $m$   $\leftarrow$  false
  while * do
    if * then
      match propagate_ilp( $p$ ) with
        case Unsat
           $\perp$  return Unsat
        case Modified( $p'$ )
           $p$   $\leftarrow$   $p'$ 
           $m$   $\leftarrow$  true
        case Unchanged
           $\perp$  noop ()
    if * then
      match propagate_cc( $p$ ) with
        case Unsat
           $\perp$  return Unsat
        case Modified( $p'$ )
           $p$   $\leftarrow$   $p'$ 
           $m$   $\leftarrow$  true
        case Unchanged
           $\perp$  noop ()
  return  $m$  ? Modified( $p$ ) : Unchanged

```

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 Real Relaxations

Propagation is followed by solving a real relaxation (`solve_relaxation`). A response `Unsat` for the relaxation implies that the (strictly harder) integer constraints of the node are also unsatisfiable. If `solve_relaxation` returns an assignment α (case `Solved(α)`), `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 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` needs to process the node again.

```

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

```

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 [13, 21]. Exposition of cut generation is 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 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.

```

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

```

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 .

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. This means that the conflict provides a witness for the gap between 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. It is always

```

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

```

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 $v > v'$.

The branching strategy we outlined is in alignment with the Nelson-Oppen (NO) scheme for combining decision procedures [24, 18]. 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 conjunctions of integer linear constraints that appear under arbitrary propositional structure. `lnez` provides such support by utilizing *indicator constraints*. Such constraints have the form $l \Rightarrow \sum_{0 \leq i < n} [c_i \cdot x_i] \leq c$, where l is a possibly negated Boolean variable, c_i and c are constants, and x_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 [19], 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 $BC(T)$ just like their CC counterparts (Section 3).

```

function branch_cc( $p$  : node, ( $f, v, v', l, l'$ ) : conflict) : {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

```

4.2 Local Theory Extensions

We demonstrate how to support user-provided axioms within $\text{BC}(T)$ and Inez . Such axioms constrain newly defined function symbols (beyond the ones in the signature $\Sigma_{\mathcal{Z}}$ 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 in Equation 1. Equation 1 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. (Every universally quantified axiom over the union of Σ and the signature of arithmetic can be encoded as a set of universally quantified clauses.)

Our implementation of axioms in Inez builds upon results on *local theory extensions* [30] that allow us to replace axioms like the one in Equation 1 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 (Section 4.1), our implementation applies a more specialized procedure that retains the clausal structure. The literals are inequalities, *e.g.*, for our example of Equation 1 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 ³ for constraints of the form $\bigvee_i v_i \leq c_i$ that notably employs SAT-like techniques for clauses.

4.3 Databases

`lnez` provides an extension aimed at database analysis [20]. 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, `lnez` provides higher-level database-inspired modeling constructs.

Table membership fits in $BC(T)$ 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 [17, 31, 6] differ from `lnez` both with respect to the language constructs that they provide and the underlying constraint technology.

Backend: `lnez` seeks to combine the strengths of Mathematical Programming solvers [2, 3, 7] and solvers for first-order theories [23, 29, 24], *e.g.*, as implemented within `Lazy SMT` [28]. Previous work on combining ideas from these two areas has mostly focused on improving the arithmetic capabilities [11, 14] of SMT solvers by integrating MILP engines [12, 22, 8], and on implementing optimization within SMT [26, 10]. `MPMT` differs by having as its core an MILP solver, as opposed to a SAT solver. In contrast to SMT, `MPMT` targets applications where a MILP-based approach is appropriate (*e.g.*, the example of Section 2), but involve requirements that are hard or impossible to encode as linear inequalities.

6 Conclusions

We provided an overview of the techniques that empower the `lnez` constraint solver. `lnez` 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 `lnez` to efficiently implement $BC(T)$. `lnez` is an extension of OCaml that allows users to write programs that orchestrate arbitrary interleaving between general computation and `MPMT` constraint solving.

³ http://scip.zib.de/doc-3.1.0/html/cons_bounddisjunction_8h.php

References

1. Camlp4. See <https://github.com/ocaml/camlp4/wiki>.
2. CPLEX. See <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
3. Gurobi. See <http://www.gurobi.com>.
4. MySQL. See <https://www.mysql.com/>.
5. mysql-protocol. See https://github.com/slegrand45/mysql_protocol.
6. Z3py. See <http://rise4fun.com/z3py/tutorial>.
7. T. Achterberg. SCIP: Solving Constraint Integer Programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.
8. F. Besson. On using an inexact floating-point LP solver for deciding linear arithmetic in an SMT solver. In *SMT*, 2010.
9. E. Castillo, A. Conejo, P. Pedregal, R. Garca, and N. Alguacil. *Building and Solving Mathematical Programming Models in Engineering and Science*. Wiley, 2002.
10. A. Cimatti, A. Franzen, A. Griggio, R. Sebastiani, and C. Stenico. Satisfiability Modulo the Theory of Costs: Foundations and Applications. In *TACAS*, 2010.
11. B. Dutertre and L. de Moura. A Fast Linear-Arithmetic Solver for DPLL(T). In *CAV*, 2006.
12. G. Faure, R. Nieuwenhuis, A. Oliveras, and E. Rodriguez-Carbonell. SAT Modulo the Theory of Linear Arithmetic: Exact, Inexact and Commercial Solvers. In *SAT*, 2008.
13. R. E. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the AMS*, 64:275–278, 1958.
14. A. Griggio. A Practical Approach to Satisfiability Modulo Linear Integer Arithmetic. *JSAT*, 8:1–27, 2012.
15. C. Hang, P. Manolios, and V. Papavasileiou. Synthesizing Cyber-Physical Architectural Models with Real-Time Constraints. In *CAV*, 2011.
16. S. P. Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple Unification-Based Type Inference for GADTs. In *ICFP*, 2006.
17. A. S. Koksals, V. Kuncak, and P. Suter. Constraints as Control. In *POPL*, 2012.
18. Z. Manna and C. Zarba. Combining Decision Procedures. In *10th Anniversary Colloquium of UNU/IIST*, 2002.
19. P. Manolios and V. Papavasileiou. ILP Modulo Theories. In *CAV*, 2013.
20. P. Manolios, V. Papavasileiou, and M. Riedewald. ILP Modulo Data. In *FMCAD*, 2014.
21. J. E. Mitchell. Branch-and-Cut Algorithms for Combinatorial Optimization Problems. In *Handbook of Applied Optimization*, pages 223–233. Oxford University Press, 2000.
22. D. Monniaux. On using floating-point computations to help an exact linear arithmetic decision procedure. In *CAV*, 2009.
23. G. Nelson and D. Oppen. Fast Decision Algorithms Based on Union and Find. 1977.
24. G. Nelson and D. C. Oppen. Simplification by Cooperating Decision Procedures. *TOPLAS*, 1:245–257, 1979.
25. R. Nieuwenhuis and A. Oliveras. Congruence Closure with Integer Offsets. In *LPAR*, 2003.
26. R. Nieuwenhuis and A. Oliveras. On SAT Modulo Theories and Optimization Problems. In *SAT*, 2006.

27. R. Nieuwenhuis and A. Oliveras. Fast congruence closure and extensions. *Information and Computation*, 205:557–580, 2007.
28. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *JACM*, 53(6):937–977, 2006.
29. R. Shostak. A practical decision procedure for arithmetic with function symbols. *JACM*, 26(2):351–360, 1979.
30. V. Sofronie-Stokkermans. Hierarchic Reasoning in Local Theory Extensions. In *CADE*, 2005.
31. E. Torlak and R. Bodik. A Lightweight Symbolic Virtual Machine for Solver-Aided Host Languages. In *PLDI*, 2014.