

An Introductory Practical Guide to Rust Bindings for the Z3 Solver

Mehrad Haghshenas*

June 2025

Acknowledgement

Some of the responses and content in this document were developed with the assistance of *LLMs*. Likewise, in some cases where clear explanations were provided in the official documentation, they were directly used. This document is not intended to introduce novel material. The primary goal was to centralize information and provide illustrative examples to create a beginner-friendly practical guide in using the Rust bindings for Z3.

*<https://mehrad31415.github.io/>

This page has been intentionally left blank.

1 Introduction

Rust has [high-level](#) and [low-level](#) bindings for the [Z3 solver](#). The [z3 crate](#) provides high-level bindings, and the [z3-sys crate](#) provides the low-level bindings to the Z3 solver. The `z3` crate is built on top of the `z3-sys` crate and encapsulates it. The `z3-sys` crate provides the raw, unsafe, low-level C API that Z3 exposes. More often than not, the `z3` crate should be used. The `z3-sys` crate is used only when some Z3 feature isn't wrapped into high-level bindings in the `z3` crate. Another time `z3-sys` is directly used is when writing your custom high-level API for Z3. There is no indication that the `z3-sys` crate offers better performance; thus, the `z3` crate will be used unless a limitation is confronted.

Both `z3` and `z3-sys` crates have dependencies on Z3. They can be added by writing the following in `Cargo.toml`:

```
1 [dependencies]
2 z3 = "0.12"
3 z3-sys = "0.8"
```

There are 3 ways for the crates to currently find Z3 on the local computer:

1. By default, it will look for a system-installed copy of Z3. On macOS, this will be via Homebrew (`brew install z3`). To reiterate, to use the crates on a macOS computer, run `brew install z3` in the terminal. Confirm that Z3 has been installed by running `which z3` in the terminal. After that, write the following lines in the `Cargo.toml` file of the workspace:¹

```
1 [dependencies]
2 z3 = { git="https://github.com/prove-rs/z3.rs.git" }
3 z3-sys = { git="https://github.com/prove-rs/z3.rs.git" }
```

2. Adding the *bundled* feature will use *CMake* to build a locally bundled copy of Z3. This copy is provided via a git submodule, nevertheless the time to build the project is considerably higher. This will look like `z3 = {version="0.12", features = ["bundled"]}` for example.²
3. Enabling the *vcpkg* feature will use *vcpkg* to build and install a copy of Z3. This will look like `z3 = {version="0.12", features = ["vcpkg"]}` for example.³

The source code related to the `z3` crate can be found in the `src` subdirectory. The following is a list of the rust modules:

1. `lib.rs`
2. `config.rs`
3. `context.rs`
4. `solver.rs`
5. `ast.rs`
6. `sort.rs`
7. `symbol.rs`

¹For some reason writing `z3 = "0.12"` or `z3-sys = "0.8"` in the `Cargo.toml` will give the error *fatal error: 'z3.h' file not found*. Note that the `z3` crate internally uses the `z3-sys` crate, which in turn requires a `z3.h` during the build time. You may need to ensure Z3 headers are discoverable via `PKG_CONFIG_PATH` or `Z3_SYS_INCLUDE_DIR`.

²Note that `static-link-z3` is a deprecated alias for `bundled`. Therefore, writing `z3 = { version = "0.12", features = ["static-link-z3"] }` is the same, but again it gets stuck at *Blocking waiting for file lock on build directory*. This is because `static-link-z3` - similar to `bundled` - links Z3 statically; i.e., the native `libz3` code is compiled directly into your binary and there is no need for Z3 to be installed on the system at runtime. Nevertheless, building the project takes a longer amount of time.

³*CMake* builds Z3 from source and gives full control over compilation options, making it ideal for customization. `vcpkg` installs prebuilt Z3 libraries with minimal setup, making it convenient for quick integration, especially on Windows which lacks a native package manager for C/C++ libraries.

8. `ops.rs`
9. `datatype_builder.rs`
10. `rec_func_decl.rs`
11. `func_decl.rs`
12. `func_entry.rs`
13. `func_interp.rs`
14. `model.rs`
15. `pattern.rs`
16. `goal.rs`
17. `tactic.rs`
18. `probe.rs`
19. `optimize.rs`
20. `params.rs`
21. `statistics.rs`
22. `version.rs`

By looking at the `lib.rs` module and the defined module tree system, the following functionalities are provided by the public API of the `z3` crate:

1. `AstKind`, `GoalPrec`, `SortKind`, `DeclKind` from the `z3-sys` crate are re-exported. See sections `AstKind`, `SortKind`, `DeclKind`, and `GoalPrec` for details.
2. `get_global_param`, `reset_all_global_params`, `set_global_param` from the `params` module are re-exported. See section `params` for explanations.
3. `StatisticsEntry`, `StatisticsValue` from the `statistics` module are re-exported. See section `statistics` for explanations.
4. `full_version`, `version`, `Version` from the `version` module are re-exported. See section `version` for explanations.
5. Everything marked `pub` inside the `lib.rs` file is accessible. This includes any struct or enum declared as `pub`. Although their full implementations (e.g., methods) are defined in other modules, any `pub` methods in those `impl` blocks are also accessible by external crates. See section `lib` for explanations.
6. Everything `public` inside the `ast` module is accessible.
7. Everything `public` inside the `datatype_builder` module is accessible.

We will begin by explaining points one to four. After that, we will discuss the public data structures defined in `lib.rs`, and then proceed to explain each module in turn, covering the implementations of the structs and enums introduced in `lib.rs`.

2 z3-sys re-exports

The `AstKind`, `GoalPrec`, `SortKind`, `DeclKind` from `z3-sys` are accessible:

1. **AstKind:** In Z3, everything is represented as an AST (Abstract Syntax Tree). `AstKind` tells you what kind of AST node you're working with. It corresponds to `Z3_ast_kind` in the C API,

defining the different kinds of Z3 AST (abstract syntax trees). Specifically, it is used to define *terms*, *formulas*, *types*. It has the following variants:

- **Numerical**: Used whenever the node is a literal number or bit-vector literal.
- **App**: Function or operator applications, including constant symbols. For example, every function application like `f(a,b)`, predicate like the constant `true`, or built-in operator like `x + y` is an App node.
- **Var**: Bound variables inside quantifiers (using *de Bruijn* indices). In `forall x. P(x)`, the occurrence of `x` is represented as a Var AST. Nevertheless, Z3 doesn't store the name `x`; it uses an index like `Var(0)`.
- **Quantifier**: `forall` and `exists` nodes. Wraps the whole quantifier; for example, `z3::ast::Quantifier::forall`.
- **Sort**: Type ASTs themselves, like Int, Bool, or user-declared sorts. When you inspect an AST that represents a type, you get `AstKind::Sort`.
- **FuncDecl**: Function-declaration ASTs (uninterpreted functions & constructors). Symbols when you declare a function or a datatype constructor.
- **Unknown**: Anything otherwise!

Here is an example of how the abstract syntax kind can be retrieved:

```

1 use z3::{
2     ast::{self, Ast},
3     Config, Context,
4 };
5
6 fn main() {
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9
10    // Build a small term: x + 1
11    let x = ast::Int::new_const(&ctx, "x");
12    let one = ast::Int::from_i64(&ctx, 1);
13    let term = x + one;
14
15    // Inspect its kind
16    match term.kind() {
17        z3::AstKind::App => println!("func app."),
18        z3::AstKind::Numeral => println!("num."),
19        other => println!("Unexpected {:?}", other),
20    }
21 }
```

2. **SortKind**: This corresponds to `Z3_sort_kind` in the C API and represents the different kind of Z3 *types*. This enum has the following variants:

- (a) **Uninterpreted**: This type is used to model abstract types with no built-in semantics (e.g., generics or opaque types). Do not use this type to model self-defined enums or structs. To reiterate, even unit structs (e.g., `struct MyUnitStruct;`) should be modeled as *Datatypes*, not uninterpreted sorts. **When to choose each sort**:

Use **Uninterpreted** when the type is truly opaque and you only need an identity:

- **Generic parameters**. For example in Rust if we have `fn f<T>(x: T) ...` the generic will translate to `(declare-sort T 0)` in Z3.

- **Opaque handles.** For example file descriptors, DB keys, capability tokens, ... which do not have constructors, translate well to uninterpreted types like (`declare-const fd File`).

– Rust (interface only)

```

1  /// An OS file descriptor; the function caller never sees its fields.
2  pub struct FileHandle { _priv_field: () }
3
4  /// Close the handle; returns true on success.
5  pub fn close(fd: FileHandle) -> bool { /* kernel call */ }
```

– SMT-LIB

```

(declare-sort FileHandle 0)
(declare-fun close (FileHandle) Bool)

; Concrete handle values the program might manipulate
(declare-const fd_stdin FileHandle)
(declare-const fd_temp FileHandle)
```

No constructors or pattern matching are necessary—the verifier only needs to reason about *identity* - (in)equality - and the effects of the operations such as `close`.

- **Trait objects or Implementations.** The caller sees only a behaviour contract, not a shape.

– Rust

```

1  trait Drawable { fn draw(&self); }
2
3  fn render(obj: &dyn Drawable) {           // caller sees only &dyn Drawable
4      obj.draw();
5  }
6
7  // returning an opaque iterator via `impl Trait`
8  fn counter() -> impl Iterator<Item = i32> {
9      0..10
10 }
```

– SMT-LIB

```

; A type with no visible structure
(declare-sort Drawable 0)

; Functions that use it
(declare-fun render (Drawable) Bool)

; The concrete iterator type returned by `counter` is hidden,
; so we model it as another uninterpreted sort.
(declare-sort CounterIter 0)
(declare-fun next (CounterIter) Int)
```

- **Phantom types.** Compile-time markers that never exist at runtime (e.g., `PhantomData<T>`).

– Rust

```

1  use std::marker::PhantomData;
2
3  // Distinguishes IDs that are all just u64 at runtime.
```

```

4  struct Id<Tag> {
5      raw: u64,
6      _p: PhantomData<Tag>,
7  }
8
9  struct UserTag;
10 struct OrderTag;
11
12 type UserId = Id<UserTag>;
13 type OrderId = Id<OrderTag>;

```

- SMT-LIB
 - ; Each phantom specialisation becomes an opaque sort


```
(declare-sort UserId 0)
(declare-sort OrderId 0)
```
 - ; The runtime payload can be modelled separately if needed


```
(declare-fun rawUserId (UserId) Int)
(declare-fun rawOrderId (OrderId) Int)
```

Use *Datatype* when the type has a fixed set of constructible values or internal structure:

- **Unit structs.** struct A; in Rust translates to (declare-datatypes () ((A A))).
- **Fielded structs.** struct Pt { x: Int, y: Int } in Rust translates to (declare-datatypes () ((Pt (mkPt (x Int) (y Int))))).
- **Enums/variants.** enum E { V1, V2(i32) } in Rust translates to one datatype with two constructors V1 and V2.
- **Option/Result-like sum types.** Needed for pattern matches.
- **Recursive data structures.** Lists, trees, etc. *Datatypes* give you built-in induction and selectors.

```

1 use z3::{
2     ast::{Ast, Dynamic},
3     Config, Context, Sort,
4 };
5 fn main() {
6     let cfg = Config::new();
7     let ctx = Context::new(&cfg);
8
9     let my_sort = Sort::uninterpreted(&ctx, "MySort".into());
10    let a = Dynamic::new_const(&ctx, "a", &my_sort);
11    let b = Dynamic::new_const(&ctx, "b", &my_sort);
12
13    let eq = a.eq(&b);
14    println!("Expression eq: {}", eq);
15
16    let _eq = a._eq(&b);
17    println!("Expression _eq: {}", _eq);
18 }
19

```

(b) **Bool:** This type is used for two-valued logic.

```

1 fn main() {
2     use z3::{ast::Bool, Config, Context, Solver};
3
4     let cfg = Config::new();

```

```

5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     // Create a boolean variable and assert it
9     let b = Bool::new_const(&ctx, "b");
10    solver.assert(&b);
11
12    // Check if the assertion is satisfiable
13    match solver.check() {
14        z3::SatResult::Sat => println!("The assertion is satisfiable."),
15        z3::SatResult::Unsat => println!("The assertion is unsatisfiable."),
16        z3::SatResult::Unknown => println!("The satisfiability is unknown."),
17    }
18}

```

(c) **Int:** This type is used to model integers.

```

1 fn main() {
2     use z3::{ast::Ast, ast::Int, Config, Context, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     // Create integer variables
9     let x = Int::new_const(&ctx, "x");
10    let y = Int::new_const(&ctx, "y");
11
12    solver.assert(&x.gt(&Int::from_i64(&ctx, 0))); // x > 0
13    solver.assert(&y.gt(&Int::from_i64(&ctx, 0))); // y > 0
14    solver.assert(&Int::add(&ctx, [&x, &y])._eq(&Int::from_i64(&ctx, 10))); // x + y
15    ↵      = 10
16    solver.assert(&x.lt(&Int::from_i64(&ctx, 7))); // x < 7
17
18    // Check if the constraints are satisfiable
19    match solver.check() {
20        z3::SatResult::Sat => {
21            let model = solver.get_model().unwrap();
22            println!(
23                "Solution found: x = {}, y = {}",
24                model.eval(&x, true).unwrap(),
25                model.eval(&y, true).unwrap()
26            );
27        }
28        _ => println!("No solution found."),
29    }
}

```

One extra note for the `model.eval(..., model_completion)`: When the *model_completion* flag is set to true, Z3 will always return a value for the given expression. This means that even if that value wasn't explicitly assigned in the model or there were no explicit constraints related to that variable, arbitrary values will be assigned to the expression which are consistent with the constraints. However, when set to false, Z3 will return `None` unless the variable or expression was explicitly assigned a value by the solver or used in an assertion constraint. This stricter mode is useful when you want to know exactly which parts of your formula the model truly depends on. For example, if you query an *unused* variable with `model_completion = false`, you'll get `None`, whereas setting it to true will give you some default value, even though it plays no role in the actual solution.

- (d) **Real:** Represents mathematical real numbers with exact algebraic precision, allowing you to write constraints that reason about fractions, square-roots, and polynomial roots without the rounding errors of floating-point arithmetic – i.e., Z3 internally treats *reals* as exact rational values (like $1/3$ stays as $1/3$, not 0.333). This is not suited for hardware or low-level binary math—use BV (bit-vectors) instead.

```

1 fn main() {
2     use z3::{ast::Ast, ast::Real, Config, Context, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     // Declare real variables
9     let x = Real::new_const(&ctx, "x");
10    let y = Real::new_const(&ctx, "y");
11
12    // Constraints: x + y = 3/2, x * y = 1/2
13    solver.assert(&(x.clone() + y.clone())._eq(&Real::from_real(&ctx, 3, 2))); // x +
14    ↪ y = 1.5
15    solver.assert(&(x.clone() * y.clone())._eq(&Real::from_real(&ctx, 1, 2))); // x *
16    ↪ y = 0.5
17
18    // Solve and retrieve values
19    match solver.check() {
20        z3::SatResult::Sat => {
21            let model = solver.get_model().unwrap();
22            let x_val = model.eval(&x, true).unwrap().as_real().unwrap();
23            let y_val = model.eval(&y, true).unwrap().as_real().unwrap();
24            println!(
25                "Solution: x = {} / {}, y = {} / {}",
26                x_val.0, x_val.1, y_val.0, y_val.1
27            );
28        }
29        _ => println!("No solution"),
30    }
31 }
```

- (e) **BV:** The theory of bit-vectors allows modeling the precise semantics of unsigned and of signed two-complements arithmetic. Bit-vectors model fixed-size binary numbers, like 32-bit or 64-bit integers. Use this for low-level operations such as bitwise logic, shifts, and overflow behavior. This is ideal for reasoning about hardware.

```

1 fn main() {
2     use z3::{ast::Ast, ast::BV, Config, Context, SatResult, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     // Create 8-bit variables: a = 255, b = 1
9     let a = BV::from_u64(&ctx, 255, 8);
10    let b = BV::from_u64(&ctx, 1, 8);
11
12    // Add them (wraps around due to 8-bit width)
13    let sum = a.bvadd(&b);
14
15    // Assert sum == 0 (255 + 1 wraps to 0 in 8-bit)
16    solver.assert(&sum._eq(&BV::from_u64(&ctx, 0, 8)));
17 }
```

```

17
18     // Verify and print model
19     match solver.check() {
20         SatResult::Sat => {
21             let model = solver.get_model().unwrap();
22             println!("Result: {}", model.eval(&sum, true).unwrap()); // result = #x00
23             // (hexadecimal with 8 bits)
24         }
25         _ => panic!("Unsatisfiable"),
26     }
27 }
```

- (f) **Array:** Used to stored key-value mappings. Each array term is an immutable function from an *index* sort (the key) to a *range* sort (the value). The underlying *select/store* axioms capture the behavior of read and write operations, while *extensionality* ensures that two arrays are equal iff they agree on all indices. Beyond simple maps, boolean-valued arrays encode *sets*.

```

1 fn main() {
2     use z3::{
3         ast::{Array, Ast, Int},
4         Config, Context, Solver,
5     };
6
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9     let solver = Solver::new(&ctx);
10
11    // array : Int + Int
12    let array = Array::new_const(&ctx, "mem", &z3::Sort::int(&ctx),
13                                &z3::Sort::int(&ctx));
14
15    // Write: mem[2] := 10
16    let mem1 = array.store(&Int::from_i64(&ctx, 2), &Int::from_i64(&ctx, 10));
17
18    // Read: r = mem1[2] ; r2 = mem1[3]
19    let r = mem1.select(&Int::from_i64(&ctx, 2));
20    let r2 = mem1.select(&Int::from_i64(&ctx, 3));
21
22    // Constraints: r=10, r2 != 10
23    solver.assert(&r._eq(&Int::from_i64(&ctx, 10).into()));
24    solver.assert(&r2._eq(&Int::from_i64(&ctx, 10).into()).not());
25
26    assert_eq!(solver.check(), z3::SatResult::Sat);
27    let model = solver.get_model().unwrap();
28    let mem1_in_model = model.eval(&mem1, true).unwrap();
29    println!("mem1_in_model = {}", mem1_in_model);
```

- (g) **Datatype:** Note that self-defined types are best represented via *Datatype* sort. Specifically, this sort is used to model enums and structs.

```

1 fn main() {
2     use z3::{ast::Ast, ast::Datatype, Config, Context, DatatypeBuilder, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     // Define an enum Color { Red, Green, Blue }
9     let color_enum = DatatypeBuilder::new(&ctx, "Color")
```

```

10      .variant("Red", vec![]) // No fields (unit variant)
11      .variant("Green", vec![]) // No fields
12      .variant("Blue", vec![]) // No fields
13      .finish();
14
15  // Create variables of Color type
16  let x = Datatype::new_const(&ctx, "x", &color_enum.sort);
17  let y = Datatype::new_const(&ctx, "y", &color_enum.sort);
18
19  // Assert x is Red
20  let red_value = color_enum.variants[0].constructor.apply(&[]);
21  solver.assert(&x._eq(&red_value.as_datatype().unwrap()));
22
23  // Assert y is not Red (must be Green or Blue)
24  let is_red = color_enum.variants[0].tester.apply(&[&y]);
25  solver.assert(&is_red.as_bool().unwrap().not());
26
27  match solver.check() {
28      z3::SatResult::Sat => {
29          println!("Satisfiable: x is Red, y is not Red");
30          let model = solver.get_model().unwrap();
31          let x_value = model.eval(&x, true).unwrap();
32          let y_value = model.eval(&y, true).unwrap();
33          println!("x: {:?}", x_value);
34          println!("y: {:?}", y_value);
35      }
36      z3::SatResult::Unsat => {
37          println!("Unsatisfiable");
38      }
39      z3::SatResult::Unknown => {
40          println!("Unknown result");
41      }
42  }
43 }

```

- (h) **Relation:** Represents finite relations for Datalog-style fixed-point engines. Relations in Z3 are predicate symbols that can be used in Datalog-style logic programming. They differ from regular functions because they're designed to work with Z3's fixedpoint solver rather than the standard SMT solver. They follow least-fixedpoint semantics - a relation is empty unless explicitly defined by facts or rules. **Note that Z3's fixedpoint engine (for Datalog/Horn clauses) is not fully supported in the high-level Rust binding of Z3.**
- (i) **FiniteDomain:** Used to model small, fixed-size enumerations (like an enum with a known number of opaque values). By contrast, *enums* in *Datatype* sorts require giving each element a distinct name. A FiniteDomain sort of size N provides exactly N distinct values (0 through N-1) with no further structure. **The high-level Rust binding doesn't currently expose *finite_domain_sort* directly.**
- (j) **FloatingPoint & RoundingMode:** IEEE-754 floats and their rounding modes. Z3 supports all five IEEE-754 rounding *modes*, each as a Float AST of the RoundingMode sort:
- `Float::round_nearest_ties_to_even(&ctx)` How this works is that it rounds to the nearest representable floating-point value. When halfway between two values, rounds to the even number. (RNE – default mode)
 - `Float::round_nearest_ties_to_away(&ctx)` How this works is that it rounds to the nearest representable floating-point value. When halfway between two values, rounds to the number away from zero. (RNA)

- `Float::round_toward_zero(&ctx)` Always rounds the number to the number closest to zero. (RTZ)
- `Float::round_toward_negative(&ctx)` Always rounds the number to the number closest to $-\infty$. (RTN)
- `Float::round_toward_positive(&ctx)` Always rounds the number to the number closest to $+\infty$. (RTP)

In the high-level Rust binding of Z3, `Float::round_nearest_ties_to_even(&ctx)` and `Float::round_nearest_ties_to_away(&ctx)` are unavailable. These modes can be accessed via `z3_sys`.

```

1 use z3::{ast::Ast, ast::Float, Config, Context, Solver};
2
3 use z3::{ast::Ast, ast::Float, Config, Context, Solver};
4
5 fn main() {
6     let cfg = Config::new();
7     let ctx = Context::new(&cfg);
8     let solver = Solver::new(&ctx);
9
10    // Create -1.5 as a 64-bit float
11    let x = Float::from_f64(&ctx, -100.5);
12    let y = Float::from_f64(&ctx, 100.55);
13
14    // Define rounding modes
15    let rtz = Float::round_towards_zero(&ctx);
16    let rtn = Float::round_towards_negative(&ctx);
17    let rtp = Float::round_towards_positive(&ctx);
18
19    // Apply rounding
20    let x_rtz = rtz.add(&x, &y);
21    let x_rtn = rtn.add(&x, &y);
22    let x_rtp = rtp.add(&x, &y);
23
24    match solver.check() {
25        z3::SatResult::Sat => {
26            let model = solver.get_model().unwrap();
27            println!("RTZ: {}", model.eval(&x_rtz, true).unwrap());
28            println!("RTN: {}", model.eval(&x_rtn, true).unwrap());
29            println!("RTP: {}", model.eval(&x_rtp, true).unwrap());
30        }
31        _ => println!("No solution"),
32    }
33}
```

Ironically, I always get the same result when using the rounding modes... why? The question arises as to why use `Float` over `Real` (or `BV`). First, special Values like `Nan`, $+\infty$, $-\infty$ are native to `Float`, but impossible in `Real` or `BV`. Also, rounding semantics like RTZ for truncate and RTP for ceiling are available in `Float`.

- (k) **Seq:** used to model vectors and ordered collections. In general, it indicates sequences of elements (often String theory).

```

1 use z3::{ast::String, Config, Context, Solver};
2
3 fn main() {
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
```

```

6     let solver = Solver::new(&ctx);
7
8     let s1 = String::from_str(&ctx, "abc").unwrap();
9     let s2 = String::from_str(&ctx, "bc").unwrap();
10
11    let substr = s1.contains(&s2);
12    solver.assert(&substr);
13
14    let starts_with = String::from_str(&ctx, "a").unwrap().prefix(&s1);
15    solver.assert(&starts_with);
16
17    assert_eq!(solver.check(), z3::SatResult::Sat);
18    let model = solver.get_model().unwrap();
19    println!("Valid: s1 = {}, s2 = {}", model.eval(&s1, true).unwrap(),
20            model.eval(&s2, true).unwrap());
21}

```

- (l) **RE:** Regular expression sort provides a representation of regular languages over sequences (often strings). Use it whenever you need to constrain a sequence variable by membership in a regular language—e.g., does this URL path match /api/v0-9+/items?

```

1 // is there a string that matches the regex "ab" one-or-more times and has length 1?
2 //   ↪ No!
3 fn main() {
4     use z3::{
5         ast::{Ast, Int, Regexp, String},
6         Config, Context, SatResult, Solver,
7     };
8
9     let cfg = Config::new();
10    let ctx = Context::new(&cfg);
11    let solver = Solver::new(&ctx);
12
13    let s = String::new_const(&ctx, "s");
14
15    // Build the regex: one-or-more "ab"
16    let re_literal = Regexp::literal(&ctx, "ab");
17    let one_or_more_ab = re_literal.plus();
18
19    // wraps Z3_mk_seq_in_re
20    let matches = s.regex_matches(&one_or_more_ab);
21
22    solver.assert(&matches);
23    solver.assert(&s.length()._eq(&Int::from_i64(&ctx, 1)));
24    match solver.check() {
25        SatResult::Sat => {
26            let model = solver.get_model().unwrap();
27            let s_value = model.eval(&s, true).unwrap();
28            println!("Found a string that matches the regex: {}", s_value);
29        }
30        SatResult::Unsat => {
31            println!("No string matches the regex.");
32        }
33        SatResult::Unknown => {
34            println!("The satisfiability of the regex is unknown.");
35        }
36    }
37}

```

- (m) **Unknown:** Should not be used unless there is a type that does not fall under any of the

previous categories. It represents an error condition rather than a legitimate sort type for users. Z3 uses numeric identifiers 0 – 12 & 1000 to represent different sort kinds in its internal AST representation.

```

1  typedef enum {
2      Z3_UNINTERPRETED_SORT = 0,
3      Z3_BOOL_SORT,           // 1
4      Z3_INT_SORT,           // 2
5      Z3_REAL_SORT,          // 3
6      Z3_BV_SORT,            // 4
7      Z3_ARRAY_SORT,          // 5
8      Z3_DATATYPE_SORT,       // 6
9      Z3_RELATION_SORT,       // 7
10     Z3FINITE_DOMAIN_SORT,   // 8
11     Z3FLOATING_POINT_SORT,  // 9
12     Z3_ROUNDING_MODE_SORT,  // 10
13     Z3SEQ_SORT,             // 11
14     Z3RE_SORT,              // 12
15     Z3UNKNOWN_SORT = 1000   // deliberately out-of-band
16 } Z3_sort_kind;

```

3. **DeclKind:** Represents the different types of interpreted function declaration. This enum corresponds directly to the Z3_decl_kind constants in Z3’s C API and defines the fundamental operations that Z3 can reason about in various mathematical theories. Specifically, DeclKind identifies the *semantic meaning* of function declarations in Z3; i.e., when you create expressions, each operation (like addition, conjunction, array access, etc.) has an associated declaration kind that tells Z3 how to reason about that operation. Z3’s DeclKind encompasses over 200 different operations. The following is a summary of all variants in the DeclKind type:

- **Boolean constants and connectives:** TRUE, FALSE, EQ, DISTINCT, ITE, AND, OR, IFF, XOR, NOT, IMPLIES, OEQ

```

1  fn main() {
2      use z3::fast::Bool, Config, Context;
3
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6
7      let p = Bool::new_const(&ctx, "p");
8      let q = Bool::new_const(&ctx, "q");
9
10     // DeclKind::AND - conjunction
11     let _and_expr = Bool::and(&ctx, [&p, &q]);
12
13     // DeclKind::OR - disjunction
14     let _or_expr = Bool::or(&ctx, [&p, &q]);
15
16     // DeclKind::NOT - negation
17     let _not_p = p.not();
18
19     // DeclKind::IMPLIES - implication
20     let _implies_expr = p.implies(&q);
21
22     // DeclKind::ITE - if-then-else
23     let _ite_expr = p.ite(&q, &Bool::from_bool(&ctx, false));
24 }

```

- **Arithmetic operators:** ANUM, AGNUM, UMINUS, ADD, SUB, MUL, DIV, IDIV, REM, MOD, LE, LT, GE, GT, TO_REAL, TO_INT, IS_INT, POWER

```

1 fn main() {
2     use z3::{ast::Ast, ast::Int, Config, Context, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let x = Int::new_const(&ctx, "x");
9     let y = Int::new_const(&ctx, "y");
10
11    // DeclKind::ADD - addition
12    let _sum = Int::add(&ctx, [&x, &y]);
13
14    // DeclKind::MUL - multiplication
15    let _product = Int::mul(&ctx, [&x, &y]);
16
17    // DeclKind::LE - less than or equal
18    let le_constraint = x.le(&Int::from_i64(&ctx, 10));
19
20    // DeclKind::GT - greater than
21    let gt_constraint = y.gt(&Int::from_i64(&ctx, 0));
22
23    solver.assert(&le_constraint);
24    solver.assert(&gt_constraint);
25    solver.assert(&_product._eq(&_sum));
26
27    match solver.check() {
28        z3::SatResult::Sat => {
29            let model = solver.get_model().unwrap();
30            println!("Model found:");
31            println!("x: {}", model.eval(&x, false).unwrap());
32            println!("y: {}", model.eval(&y, false).unwrap());
33        }
34        z3::SatResult::Unsat => {
35            println!("No solution found.");
36        }
37        z3::SatResult::Unknown => {
38            println!("Solver could not determine satisfiability.");
39        }
40    }
41 }

```

- **Arrays and sets:** STORE, SELECT, CONST_ARRAY, ARRAY_MAP, ARRAY_DEFAULT, AS_ARRAY, ARRAY_EXT, SET_UNION, SET_INTERSECT, SET_DIFFERENCE, SET_COMPLEMENT, SET_SUBSET

```

1 fn main() {
2     use z3::{
3         ast::{Array, Int},
4         Config, Context
5     };
6
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9
10    let array = Array::new_const(&ctx, "arr", &z3::Sort::int(&ctx),
11        &z3::Sort::int(&ctx));
12    let index = Int::new_const(&ctx, "i");
13    let value = Int::new_const(&ctx, "v");

```

```

14     // DeclKind::SELECT - array read
15     let read_value = array.select(&index);
16
17     // DeclKind::STORE - array write
18     let updated_array = array.store(&index, &value);
19
20     // DeclKind::CONST_ARRAY - constant array
21     let const_array = Array::const_array(&ctx, &z3::Sort::int(&ctx),
22                                         &Int::from_i64(&ctx, 42));
22 }

```

- **Bit-vector operations:** BNUM, BIT0, BIT1, BNEG, BADD, BSUB, BMUL, BSDIV, BDIV, BSREM, BUREM, BSMOD, CONCAT, SIGN_EXT, ZERO_EXT, EXTRACT, REPEAT, BSHL, BLSHR, BASHR, ROTATE_LEFT, ROTATE_RIGHT, BREDOR, BREDAND, BAND, BOR, BXOR, BNAND, BNOR, BXNOR, BCOMP

```

1 fn main() {
2     use z3::{ast::BV, Config, Context};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6
7     let a = BV::new_const(&ctx, "a", 8);
8     let b = BV::new_const(&ctx, "b", 8);
9
10    // DeclKind::BADD - bit-vector addition
11    let _sum = a.bvadd(&b);
12
13    // DeclKind::BAND - bitwise AND
14    let _and_result = a.bvand(&b);
15
16    // DeclKind::BSHL - left shift a for 2 bits
17    let _shifted = a.bvshl(&BV::from_u64(&ctx, 2, 8));
18
19    // DeclKind::ULT - unsigned less than
20    let _comparison = a.bvult(&b);
21
22    // DeclKind::EXTRACT - bit extraction
23    let _extract = a.extract(7, 4); // Extract bits 7-4 inclusive
24 }

```

- **Sequences and regexes:** SEQ_UNIT, SEQ_EMPTY, SEQ_CONCAT, SEQ_LENGTH, SEQ_CONTAINS, SEQ_PREFIX, SEQ_SUFFIX, SEQ_EXTRACT, SEQ_REPLACE, SEQ_AT, SEQ_INDEX, SEQ_TO_RE, SEQ_IN_RE, STR_TO_INT, INT_TO_STR, RE_PLUS, RE_STAR, RE_OPTION, RE_CONCAT, RE_UNION, RE_RANGE, RE_LOOP, RE_INTERSECT, RE_EMPTY_SET, RE_FULL_SET, RE_COMPLEMENT

```

1 fn main() {
2     use z3::{
3         ast::{Ast, String},
4         Config, Context,
5     };
6
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9
10    let s1 = String::from_str(&ctx, "hello").unwrap();
11    let space = String::from_str(&ctx, " ").unwrap();
12    let s2 = String::from_str(&ctx, "world").unwrap();
13

```

```

14 // DeclKind::SEQ_CONCAT - concatenation
15 let concat = z3::ast::String::concat(&ctx, &[&s1, &space, &s2]);
16
17 // DeclKind::SEQ_LENGTH - length
18 let length = concat.length();
19
20 // DeclKind::SEQ_CONTAINS - substring check
21 let _contains = s1.contains(&String::from_str(&ctx, "ell").unwrap());
22
23 let solver = z3::Solver::new(&ctx);
24 solver.assert(&length._eq(&z3::ast::Int::from_i64(&ctx, 11)));
25
26 match solver.check() {
27     z3::SatResult::Sat => println!("The concatenated string is 'hello world' with
28         length 11."),
29     z3::SatResult::Unsat => println!("The constraints are unsatisfiable."),
30     z3::SatResult::Unknown => println!("The satisfiability of the constraints is
31         unknown."),
32 }
31 }
```

- **Datatypes and relations:** DT_CONSTRUCTOR, DT_RECOGNISER, DT_ACCESSOR, DT_UPDATE_FIELD, RA_STORE, RA_EMPTY, RA_IS_EMPTY, RA_JOIN, RA_UNION, RA_WIDEN, RA_PROJECT, RA_FILTER, RA_NEGATION_FILTER, RA_RENAME, RA_COMPLEMENT, RA_SELECT, RA_CLONE

```

1 fn main() {
2     use z3::{ast::Int, Config, Context, DatatypeBuilder};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6
7     // Define a Point datatype
8     // Point = { mk_point: (x: Int, y: Int) }
9     let point = DatatypeBuilder::new(&ctx, "Point")
10        .variant(
11            "mk_point",
12            vec![
13                ("x", z3::DatatypeAccessor::Sort(z3::Sort::int(&ctx))),
14                ("y", z3::DatatypeAccessor::Sort(z3::Sort::int(&ctx))),
15            ],
16        )
17        .finish();
18
19     // DeclKind::DT_CONSTRUCTOR - create instance
20     // point.variants[0] is the constructor "mk_point"
21     let p = point.variants[0]
22        .constructor
23        .apply(&[&Int::from_i64(&ctx, 10), &Int::from_i64(&ctx, 20)]);
24
25     // DeclKind::DT_ACCESSOR - field access
26     // point.variants[0].accessors[0] is the accessor for "x"
27     // We can use it to get the value of x from the point instance p
28     let _x_val = point.variants[0].accessors[0].apply(&[&p]);
29 }
```

- **Pseudo-Boolean constraints:** Z3 represents linear constraints over Boolean literals as *pseudo-Boolean* (PB) atoms. Two special cases are **cardinality** constraints—“at most” and “at least” k true literals—while the general form allows integer weights. In the Rust API

the cardinality helpers PB_AT_MOST and PB_AT_LEAST do not appear as separate calls; they are expressed with the weighted predicates `pb_le` and `pb_ge` by giving every literal weight 1. Concretely:

- PB_AT_MOST — encoded as `Bool::pb_le(&ctx, &[(&b1,1), ...], k).`
- PB_AT_LEAST — encoded as `Bool::pb_ge(&ctx, &[(&b1,1), ...], k).`
- PB.LE — `Bool::pb_le(&ctx, &[(&b1,w1), ...], k).`
- PB.GE — `Bool::pb_ge(&ctx, &[(&b1,w1), ...], k).`
- PB.EQ — `Bool::pb_eq(&ctx, &[(&b1,w1), ...], k).`

```

1  use z3::{ast::Bool, Config, Context, Solver};
2
3  fn main() {
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let solver = Solver::new(&ctx);
7
8      // Three Boolean decision variables
9      let x = Bool::new_const(&ctx, "x");
10     let y = Bool::new_const(&ctx, "y");
11     let z = Bool::new_const(&ctx, "z");
12
13     /* At most two are true - x + y + z <= 2
14     let at_most_two = Bool::pb_le(&ctx, &[(&x, 1), (&y, 1), (&z, 1)], 2);
15     solver.assert(&at_most_two);
16
17     /* Weighted: 3·x + 2·y + 1·z >= 4
18     let weighted_ge = Bool::pb_ge(&ctx, &[(&x, 3), (&y, 2), (&z, 1)], 4);
19     solver.assert(&weighted_ge);
20
21     match solver.check() {
22         z3::SatResult::Sat => {
23             let m = solver.get_model().unwrap();
24             println!("x = {}", m.eval(&x, true).unwrap().as_bool().unwrap());
25             println!("y = {}", m.eval(&y, true).unwrap().as_bool().unwrap());
26             println!("z = {}", m.eval(&z, true).unwrap().as_bool().unwrap());
27         }
28         other => println!("Solver returned {:?}", other),
29     }
30 }
```

- **Floating-point operations and modes:** FPA_NUM, FPA_PLUS_INF, FPA_MINUS_INF, FPA_NAN, FPA_PLUS_ZERO, FPA_MINUS_ZERO, FPA_ADD, FPA_SUB, FPA_NEG, FPA_MUL, FPA_DIV, FPA_Rem, FPA_ABS, FPA_MIN, FPA_MAX, FPA_FMA, FPA_SQRT, FPA_ROUND_TO_INTEGRAL, FPA_EQ, FPA_LT, FPA_GT, FPA_LE, FPA_GE, FPA_IS_NAN, FPA_IS_INF, FPA_IS_ZERO, FPA_IS_NORMAL, FPA_IS_SUBNORMAL, FPA_IS_NEGATIVE, FPA_IS_POSITIVE, FPA_RM_NEAREST_TIES_TO_EVEN, FPA_RM_NEAREST_TIES_TO_AWAY, FPA_RM_TOWARD_POSITIVE, FPA_RM_TOWARD_NEGATIVE, FPA_RM_TOWARD_ZERO, FPA_FP, FPA_TO_FP, FPA_TO_FP_UNSIGNED, FPA_TO_UBV, FPA_TO_SBV, FPA_TO_REAL, FPA_TO_IEEE_BV, FPA_BVWRAP, FPA_BV2RM

```

1  fn main() {
2      use z3::{ast::Float, Config, Context};
3
4      let cfg = Config::new();
```

```

5     let ctx = Context::new(&cfg);
6
7     let x = Float::from_f64(&ctx, 1.5);
8     let y = Float::from_f64(&ctx, 2.3);
9
10    // DeclKind::FPA_RNE - rounding mode for floating-point arithmetic
11    let rne = Float::round_towards_zero(&ctx);
12
13    // DeclKind::FPA_ADD - floating-point addition
14    let sum = rne.add(&x, &y);
15
16    // DeclKind::FPA_MUL - floating-point multiplication
17    let product = rne.mul(&x, &y);
18
19    // DeclKind::FPA_LT - floating-point less than
20    let comparison = x.lt(&y);
21 }

```

- **Proof objects and special rules:** Z3 can generate detailed proof objects that justify why formulas are valid or unsatisfiable; for example: PR_UNDEF, PR_TRUE, PR_ASSERTED, PR_GOAL, PR_MODUS_PONENS, PR_REFLEXIVITY, PR_SYMMETRY, PR_TRANSITIVITY, PR_TRANSITIVITY_STAR, PR_MONOTONICITY, PR_QUANT_INTRO, PR_BIND, PR_DISTRIBUTIVITY, PR_AND_ELIM, PR_NOT_OR_ELIM, PR_REWRITE, PR_REWRITE_STAR, PR_PULL_QUANT, PR_PUSH_QUANT, PR_ELIM_UNUSED_VARS, PR_DER, PR_QUANT_INST, PR_HYPOTHESIS, PR_LEMMA, PR_UNIT_RESOLUTION, PR_IFF_TRUE, PR_IFF_FALSE, PR_COMMUTATIVITY, PR_DEF_AXIOM, PR_DEF_INTRO, PR_APPLY_DEF, PR_IFF_OEQ, PR_NNF_POS, PR_NNF_NEG, PR_SKOLEMIZE, PR_MODUS_PONENS_OEQ, PR_TH_Lemma, PR_HYPER_RESOLVE
- **Miscellaneous:** These are *catch-all* values that do not correspond to standard interpreted functions; for example INTERNAL, UNINTERPRETED. For UNINTERPRETED, Z3 treats the symbol purely as a fresh function over its domain and range sorts. For INTERNAL, it is used only by Z3's own re-writers and tactic pipelines; user code should never receive or create this kind. An example of UNINTERPRETED can be found in the following;

```

1 fn main() {
2     use z3::{
3         ast::{Ast, Dynamic, Int},
4         Config, Context, Solver,
5     };
6
7     let ctx = Context::new(&Config::new());
8     let solver = Solver::new(&ctx);
9
10    let f = z3::FuncDecl::new(&ctx, "f", &[&z3::Sort::int(&ctx)],
11        &z3::Sort::int(&ctx));
12    let x = Int::new_const(&ctx, "x");
13
14    solver.assert(&f.apply(&[&x])._eq(&Dynamic::from_ast(&Int::add(
15        &ctx,
16        &[&x, &Int::from_i64(&ctx, 1)],
17    )))); // f(x) = x+1
18
19    match solver.check() {
20        z3::SatResult::Sat => {
21            let model = solver.get_model().unwrap();
22            println!("Model found:");
23        }
24    }
25 }

```

```

22         println!("{}", model);
23     }
24     z3::SatResult::Unsat => {
25         println!("No solution found.");
26     }
27     z3::SatResult::Unknown => {
28         println!("Solver could not determine satisfiability.");
29     }
30 }
31 }
```

4. **GoalPrec:** Defines the precision of a given goal. Some goals can be transformed using over/under approximations. The following variants are present:

- **Precise:** No approximation has been applied—both satisfiable and unsatisfiable answers for the goal are exactly preserved.
- **Over:** In over-approximation, we make the formula looser. Therefore, unsat answers are preserved but sat answers are not; meaning that if we prove that the over-approximation is unsatisfiable, we can conclude that the original expression was unsatisfiable as well. This is useful if we want to keep the proofs valid. For example, adding a disjunct, eliminating a conjunct, or extending the range of numbers for an expression are examples of over-approximating the original formula.
- **Under:** In under-approximation, we make the formula tighter. Therefore, sat answers are preserved but unsat answers are not; i.e., proving the satisfiability of the under-approximation results into showing that the original expression is satisfiable. For example, adding a conjunct, eliminating a disjunct, or limiting the range of numbers for an expression are examples of under-approximating the original formula.
- **UnderOver:** Both under- and over-approximations have been applied (e.g. in different parts of the goal). Neither sat nor unsat answers are fully preserved, so the answer is *garbage*.

The following is a useful image I picked up from the following [link](#):

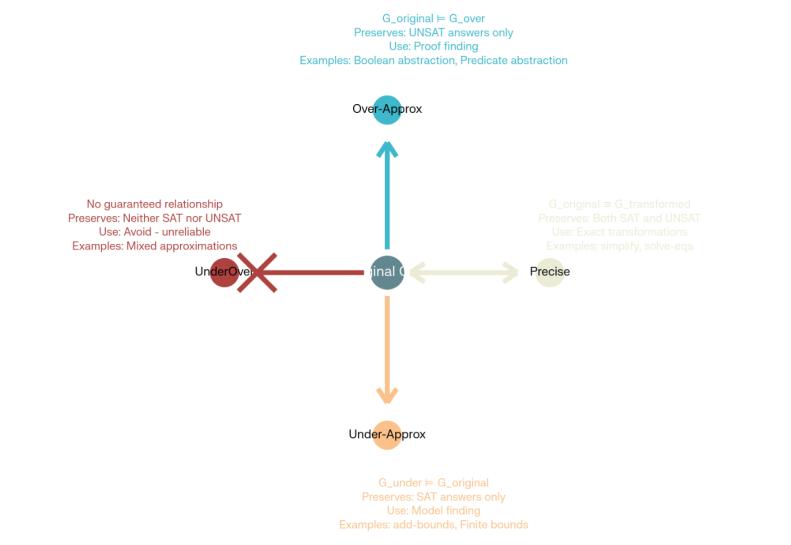


Figure 1: Z3 GoalPrec Approximation Types.

3 *params* re-exports

The `get_global_param`, `reset_all_global_params`, `set_global_param` from the `params` module are re-exported. Z3's *global parameters* are shared key-value store of configuration options that affect all created Z3 contexts; for example, `timeout`, `smt.string_solver`, etc.

- **Set** a parameter with `set_global_param(k, v)`, e.g. `set_global_param("timeout", "1000")` means *any new context will time out after 1000 ms*. This change applies to all subsequently created Z3 contexts. Existing contexts and solvers are unaffected.
- **Get** its current value with `get_global_param(k)`, which returns `Some(String)` if you've set it, or `None` to use Z3's built-in default.
- **Reset** everything to Z3's built-in defaults with `reset_all_global_params()`, after which every `get_global_param` returns `None` again. This does *not* affect any existing tactics, solvers, or contexts. Built-in defaults are hard-coded—e.g. the default for `timeout` is 4294967295 (no limit), for `model` is `false`, etc.

In short, `set_global_param` and `reset_all_global_params` manipulate a process-wide table that every new Config, Context, solver, tactic, etc. copies when it is created; once a Context exists, later changes to the global table do not reach it. Let's explain this with the following example:

```
1 use z3::{ast::Ast, ast::Int, reset_all_global_params, set_global_param, Config, Context,
2   ↵ Solver};
3
4 fn main() {
5     set_global_param("timeout", "1"); // 1ms timeout
6
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9     let solver = Solver::new(&ctx);
10
11    let x = Int::new_const(&ctx, "x");
12    let y = Int::new_const(&ctx, "y");
13
14    solver.assert(&(Int::mul(&ctx, &[&x, &y]))._eq(&Int::from_i64(&ctx, 100)));
15
16    match solver.check() {
17        z3::SatResult::Sat => println!("SAT"),
18        z3::SatResult::Unsat => println!("UNSAT"),
19        z3::SatResult::Unknown => println!("TIMEOUT/UNKNOWN"),
20    }
21
22    // Reset and try again with infinite timeout
23    reset_all_global_params();
24
25    let cfg2 = Config::new(); // inherits the now-empty table
26    let ctx2 = Context::new(&cfg2);
27    let solver2 = Solver::new(&ctx2);
28    let x2 = Int::new_const(&ctx2, "x");
29    let y2 = Int::new_const(&ctx2, "y");
30
31    solver2.assert(&(Int::mul(&ctx2, &[&x2, &y2]))._eq(&Int::from_i64(&ctx2, 100)));
32
33    match solver2.check() {
34        z3::SatResult::Sat => println!("SAT after reset"),
35        z3::SatResult::Unsat => println!("UNSAT after reset"),
36        z3::SatResult::Unknown => println!("TIMEOUT/UNKNOWN after reset"),
37    }
}
```

The output of the above code is the following:

```
TIMEOUT/UNKNOWN
SAT after reset
```

There are multiple factors to pinpoint here:

1. Global-parameter snapshotting.

- (a) `set_global_param("timeout", "1")` stores the pair `timeout = 1` in Z3's *process-wide* parameter table.
- (b) `let cfg1 = Config::new();` copies *all* global entries into `cfg1`. From now on that `Config` carries its own private copy `timeout = 1`.
- (c) `Context::new(&cfg1)` propagates the 1 ms limit to the first solver.
- (d) `reset_all_global_params()`; wipes the global table, but does *not* mutate `cfg1`.
- (e) Re-using `cfg1` for a second context therefore copies the same 1 ms limit again.
- (f) Both solvers inherit the same tiny time-out and report `unknown`.

Therefore, re-using the same `cfg1` will lead to both outputs showing `TIMEOUT/UNKNOWN`. Likewise, if we create a new `cfg2`, but erase the `reset_all_global_params()` line, the new configuration will inherit the process-wide parameter table, thus, inheriting the *1ms* timeout. *Fix:* after resetting the globals, create a fresh `Config` (call it `cfg2`) and build the second context from that. The new configuration will see an empty parameter table and therefore inherit the default “no time-out”.

- ### 2. Contexts are not interchangeable.
- Every AST node records the `Context` that created it. Mixing nodes from two different contexts in the same expression triggers an assertion inside the Rust bindings. In practice this means you must rebuild `x`, `y`, and any other ASTs from the first context, after you switch to `ctx2`.

4 *statistics* re-exports

The `StatisticsEntry`, `StatisticsValue` from the *statistics* module are re-exported. Z3 collects various runtime metrics in a `Statistics` object. Each entry is represented by:

`StatisticsEntry` A single key–value pair within `Statistics`:

`key: String` The name of the statistic.

`value: StatisticsValue` The associated value which can be either a non-negative integer counter (`UInt(u32)`) or a floating-point measurement (`Double(f64)`).

```
1 use z3::{Config, Context, Solver, StatisticsEntry, StatisticsValue};
2
3 fn main() {
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let x = z3::ast::Int::new_const(&ctx, "x");
9     solver.assert(&x.gt(&z3::ast::from_i64(&ctx, 0))); // x > 0
10    let _ = solver.check();
11
12
13    let stats = solver.get_statistics();
14
15    // Example keys to fetch
16    let example_keys = ["max memory", "time"];
```

```

17     for &key in &example_keys {
18         match stats.value(key) {
19             Some(StatisticsValue::UInt(n)) => println!("{} = {} (UInt)", key, n),
20             Some(StatisticsValue::Double(d)) => println!("{} = {} (Double)", key, d),
21             None => println!("{} = <not reported>", key),
22         }
23     }
24     println!("\nAll statistics:");
25     for StatisticsEntry { key, value } in stats.entries() {
26         println!("{} = {}", key, value);
27     }
28 }

```

5 *version* re-exports

The `full_version`, `version`, `Version` from the `version` module are re-exported. Z3 exposes the library version both as a structured tuple via the `version()` function and as a formatted string via `full_version()`. The Rust bindings provide:

`struct Version` Holds the four components of the Z3 library version:

```

major: u32 Major version number.
minor: u32 Minor version number.
build_number: u32 Internal build counter.
revision_number: u32 Revision identifier (patch/commit identifier).

```

`fn version() -> Version` Queries Z3 via `Z3_get_version`, filling each `Version` field. Returns a `Version` struct populated with the library's current version.

`fn full_version() -> &'static str` Calls `Z3_get_full_version` to obtain a C-string containing the full version string (including build date, git hash, etc.). Returns it as a Rust `&str`.

```

1 use z3::{full_version, version};
2
3 fn main() {
4     // Get structured version components
5     let ver = version();
6     println!("Structured version: {}", ver);
7
8     // Get full version string
9     let full_ver = full_version();
10    println!("Full version: {}", full_ver);
11
12    let ver1 = version();
13    let ver2 = version();
14    assert_eq!(ver1, ver2, "Version changed between calls!"); // Never fails
15 }

```

6 lib module

The following are the public enum and struct definitions in the *lib* module:

- *Config*

Purpose: Holds configuration parameters for a Z3 Context. Internally wraps a Z3_config and a list of key/value pairs.

Fields (private)

- kvs: Vec<(CString, CString)> of raw key/value pairs.
- z3_cfg: raw Z3_config handle.

When to use

- Before creating any Z3 Context to tune context or solver behavior.

How to use

```
1  use z3::{Config, Context};  
2  
3  fn main() {  
4      let mut cfg = Config::new();  
5  
6      cfg.set_param_value("timeout", "1000");           // 1 second timeout  
7      cfg.set_proof_generation(true);                 // enable proofs  
8      cfg.set_model_generation(true);                // enable model production  
9  
10     let ctx = Context::new(&cfg);  
11     let solver = z3::Solver::new(&ctx);  
12 }  
13 }
```

Tips

- All parameters must be set *before* creating the Context. In other words, calling setters after Context::new() has *no effect* on the already existing contexts.
- Use `set_param_value` for arbitrary flags not covered by the high-level API. Note that you need to define the configuration as *mutable* to change the parameters. The following is a set of the legal parameters:⁴
 - * **auto_config (bool) (default: true)** – Uses heuristics to automatically select the appropriate solver and configure it optimally. When enabled, Z3 analyzes the problem structure and chooses the best solving strategy.
 - * **ctrl_c (bool) (default: true)** – Although, `ctrl_c` is listed as a legal parameter, setting the parameter does not work. Specifically, `cfg.set_bool_param_value("ctrl_c", false);` does not work (maybe deprecated?).
 - * **debug_ref_count (bool) (default: false)** – Enables debugging support for AST reference counting, which may be useful for debugging memory management issues.
 - * **dot_proof_file (string) (default: proof.dot)** – Specifies the filename for outputting graphical proofs in DOT format. It is used when proof generation is enabled to visualize proof trees.

⁴A list of available parameters in Z3 can be obtained by running the `z3 -pd` command in the terminal.

- * **dump_models (bool) (default: false)** – Automatically dumps models whenever check-sat returns satisfiable.
- * **encoding (string) (default: unicode)** – Sets the internal string encoding used by Z3.
- * **model (bool) (default: true)** – Enables model generation for solvers when encountering satisfiable expressions. This is equivalent to `set_model_generation(true)`.
- * **model_validate (bool) (default: false)** – When enabled Z3 will *double-check* every model it produces against the original assertions before returning `Sat`. Internally, after finding a candidate assignment, Z3 re-evaluates all asserted formulas with that model. When model validation is enabled and a model fails to validate, Z3 will produce a warning message but still return the model.
- * **proof (bool) (default: false)** – Enables proof generation for unsatisfiable formulas. This is equivalent to `set_proof_generation(true)`.
- * **rlimit (unsigned int) (default: 0)** – This parameter imposes a hard cap on Z3’s *basic operation* count (e.g. AST construction, simplification, solving step). Once the solver has performed more than `rlimit` operations, it aborts the check and returns `Unknown`. Unlike `timeout`, this bound is completely deterministic and preferred for reproducible behavior. Lastly, 0 means *unlimited* operation cap.
- * **smtlib2_compliant (bool) (default: false)** – When this flag is `true`, Z3 enforces strict *SMT-LIB 2.0* syntax. Any unsupported language extensions are rejected, and only constructs defined in the official *SMT-LIB 2.0* standard may appear. This is useful when you need maximal portability of your benchmarks or want to ensure compatibility with other SMT-LIB 2-compliant solvers.

```

1  use z3::{ast::Int, Config, Context, Solver};
2
3  fn main() {
4      let mut cfg = Config::new();
5      cfg.set_param_value("smtlib2_compliant", "true"); // Enforce strict
6      // standard
6      let ctx = Context::new(&cfg);
7      let solver = Solver::new(&ctx);
8
9      let x = Int::new_const(&ctx, "x");
10     let y = Int::new_const(&ctx, "y");
11     solver.assert(&x.gt(&Int::from_i64(&ctx, 0))); // x > 0
12     solver.assert(&y.lt(&z3::ast::Int::add(&ctx, &[&x, &Int::from_i64(&ctx,
13     // 3)]))); // y < x + 3
14
15     let smt2_output = solver.to_smt2();
16     println!("{}", smt2_output);
17 }
```

- * **stats (bool) (default: false)** – Enables collection and reporting of solving statistics.
- * **timeout (unsigned int) (default: 4294967295)** – Sets maximum solving time in *milliseconds*. This is less deterministic than `rlimit` and can cause non-reproducible behavior. Note that if the solver does not produce a result before the timeout, the `Unknown` will be outputted. This is equivalent to `cfg.set_timeout_msec(5)`.
- * **trace (bool) (default: false)** – Enables trace message logging for debugging.
- * **trace_file_name (string) (default: z3.log)** – Specifies the output file for trace messages.

- * **type_check (bool) (default: true)** – Enables type checking of formulas and terms.
- * **unsat_core (bool) (default: false)** – Enables generation of unsatisfiable cores; i.e., it identifies minimal subsets of constraints that cause unsatisfiability. Without tracked assertion labels, `unsat_core()` returns empty even when enabled . Therefore, proper labeling is essential. For example the following code showcases this:

```

1  fn main() {
2      use z3::ast::Bool, Config, Context, SatResult, Solver;
3
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let solver = Solver::new(&ctx);
7      let p = Bool::from_bool(&ctx, true);
8      let q = Bool::from_bool(&ctx, false);
9
10     let lp = Bool::fresh_const(&ctx, "p_label");
11     let lq = Bool::fresh_const(&ctx, "q_label");
12
13     solver.assert_and_track(&p, &lp);
14     solver.assert_and_track(&q, &lq);
15
16     assert_eq!(solver.check(), SatResult::Unsat);
17
18     // Extract which labels caused the conflict
19     let core = solver.get_unsat_core();
20     for l in core {
21         println!("Conflict came from: {}", l);
22     }
23 }
```

- * **well_sorted_check (bool) (default: false)** – Performs additional type checking and more thorough type validation (alias for `type_check`).

```

1  use z3::ast, ast::Ast, Config, Context, Solver;
2
3  fn main() {
4      let mut cfg = Config::new();
5
6      cfg.set_timeout_msec(5000); // 5 second timeout
7      cfg.set_proof_generation(true); // Enable proof generation
8      cfg.set_model_generation(true); // Enable model generation
9
10     cfg.set_param_value("rlimit", "1000000"); // Resource limit
11     cfg.set_param_value("stats", "true"); // Enable statistics
12     cfg.set_param_value("unsat_core", "true"); // Enable unsat cores
13
14     let ctx = Context::new(&cfg);
15     let solver = Solver::new(&ctx);
16
17     let x = ast::Int::new_const(&ctx, "x");
18     let y = ast::Int::new_const(&ctx, "y");
19     solver.assert(&x.gt(&ast::Int::from_i64(&ctx, 0)));
20     solver.assert(&y.lt(&ast::Int::from_i64(&ctx, 10)));
21     solver.assert(&x._eq(&z3::ast::Int::mul(
22         &ctx,
23         &[&y, &ast::Int::from_i64(&ctx, 2)],
24     )));
25
26     // Check satisfiability
```

```

27     match solver.check() {
28         z3::SatResult::Sat => {
29             println!("Satisfiable!");
30             if let Some(model) = solver.get_model() {
31                 println!("Model: {}", model);
32             }
33         }
34         z3::SatResult::Unsat => println!("Unsatisfiable!"),
35         z3::SatResult::Unknown => println!("Unknown result"),
36     }
37 }
```

- **Context**

Purpose: Manages all Z3 objects (ASTs, solvers, models, etc.) and global configuration options.

Each **Context** defines an isolated logical environment. We may use multiple contexts. Contexts are independent; objects cannot be shared across them but can be translated.

Fields (private)

- z3_ctx: raw Z3_context pointer.

When to use

- After configuring **Config** and before creating ASTs or solvers.
- When you need multiple independent logical environments (e.g. parallel queries).

How to use

```

1  use z3::{ast, ast::Ast, Config, Context, Solver};
2
3  fn main() {
4      let mut cfg = Config::new();
5      cfg.set_model_generation(true);
6
7      // create context
8      let ctx = Context::new(&cfg);
9      let solver = Solver::new(&ctx);
10
11     let x = ast::Int::new_const(&ctx, "x");
12     solver.assert(&x.gt(&ast::Int::from_i64(&ctx, 0))); // x > 0
13
14     // new context
15     let ctx2 = Context::new(&cfg);
16     let solver2 = Solver::new(&ctx2);
17
18     let x_in_ctx2 = x.translate(&ctx2);
19     solver2.assert(&x_in_ctx2.lt(&ast::Int::from_i64(&ctx2, 10))); // x < 10
20 }
```

Tips

- Context is neither **Send** nor **Sync**: do not share it across threads.
- All ASTs, solvers, and models belong to their creating **Context**. Mixing terms will cause errors; instead, you can translate as seen in the above code.

- **ContextHandle**

Purpose: A thread-safe handle for interrupting Z3 computations (solving, model building, etc.) from another thread.

Fields (private)

- ctx: reference to the parent Context that this handle will interrupt.

When to use

- When you need to enforce timeouts or user-initiated cancellation from a different thread.
- In long-running queries or tactics where you want external control to abort.
- To recover from hangs in custom tactics or complex quantifier instantiations.

How to use

```
1  use std::{thread, time::Duration};
2  use z3::{ast::Int, Config, Context, SatResult, Solver};
3
4  fn main() -> Result<(), Box<dyn std::error::Error>> {
5      let mut cfg = Config::new();
6      cfg.set_param_value("timeout", "100000");
7      let ctx = Context::new(&cfg);
8      let solver = Solver::new(&ctx);
9      let x = Int::new_const(&ctx, "x");
10
11     for i in 0..50_000 {
12         solver.assert(&x.lt(&Int::from_i64(&ctx, i)));
13     }
14
15     thread::scope(|s| {
16         let handle = ctx.handle();
17         s.spawn(move || {
18             thread::sleep(Duration::from_millis(100));
19             handle.interrupt();
20             eprintln!("Interrupt requested");
21         });
22
23         match solver.check() {
24             SatResult::Unknown => println!("Solver was interrupted (Unknown)"),
25             SatResult::Sat     => println!("Solver unexpectedly returned Sat"),
26             SatResult::Unsat   => println!("Solver unexpectedly returned Unsat"),
27         }
28     });
29
30     Ok(())
31 }
```

Tips

- Calling `interrupt()` sets an internal flag; it does not force immediate exit, but the next check in Z3 will abort work.
- Always hold your `ContextHandle` alive as long as you need the ability to cancel.
- If you call `Context::interrupt()` directly, it affects all handles for that context. In other words, interruption state is stored at the context level, not per-handle.

• Symbol

Purpose: Represents the Z3 *symbol* type used to name term and type constructors (identifiers for constants, functions, and sorts).

Variants

- Int(u32)
- String(String)

When to use

- When declaring constants or functions: `Symbol::String("x".into())`.
- When inspecting symbols returned by Z3 APIs (e.g. `FuncDecl::name()`).

How to use:

```

1  use z3::{
2      ast::Bool,
3      Config, Context, FuncDecl, Sort, Symbol,
4  };
5
6  fn main() {
7      let cfg = Config::new();
8      let ctx = Context::new(&cfg);
9
10     let sym_str: Symbol = "Node".into();      // String variant
11     let sym_num: Symbol = 1u32.into();        // Int variant
12
13     // Declare an uninterpreted sort named "Node"
14     let node_sort: Sort = Sort::uninterpreted(&ctx, sym_str.clone());
15
16     // Declare a constant c: Node
17     let c_decl: FuncDecl = FuncDecl::new(&ctx, "c", &[], &node_sort);
18     let c = c_decl.apply(&[]); // c : Node
19
20     // Declare a function f1: Node -> Int
21     let f1_decl: FuncDecl =
22         FuncDecl::new(&ctx, sym_num.clone(), &[&node_sort], &Sort::int(&ctx));
23     let f1_app = f1_decl.apply(&[&c]).as_int().unwrap();
24
25     println!("f1(c) = {}", f1_app);
26     println!("Function name: {}", f1_decl.name()); // f1_decl.name() returns a
27     ↪ String
27 }
```

• Sort

Purpose: Wrapper around a Z3 sort (Z3_sort) in a given Context, representing a type in the logic.

Fields (private)

- ctx: reference to the owning Context.
- z3_sort: raw Z3 sort handle.

When to use

- When constructing AST nodes and types.

How to use

```

1  use z3::{Config, Context, Sort, SortKind};
2
3  fn main() {
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6
```

```

7      // Create some sorts
8      let bool_s = Sort::bool(&ctx);
9      let bv8_s = Sort::bitvector(&ctx, 8);
10
11     // Inspect
12     assert_eq!(bool_s.kind(), SortKind::Bool);
13     assert_eq!(bv8_s.kind(), SortKind::BV);
14 }
```

Tips

- Only for **in-crate users**, never construct `Sort` directly; always use `Sort::new()` to handle refcounting properly. For outside users, please use code as shown in the above example.

• *SortDiffers*

Purpose: Error type indicating two `Sort` values are not the same kind.

Fields (private)

- `left`, `right`: the mismatched `Sort` values.

When to use

- In custom type-checking to report meaningful errors on mismatch.

How to use

```

1  use z3::{Sort, SortDiffers};
2
3  fn check_same_sort<'a>(a: &'a Sort, b: &'a Sort) -> Result<(), SortDiffers<'a> {
4      if a.kind() == b.kind() {
5          Ok(())
6      } else {
7          Err(SortDiffers::new(a.clone(), b.clone()))
8      }
9  }
10
11 fn main() {
12     let cfg = z3::Config::new();
13     let ctx = z3::Context::new(&cfg);
14     let bo = z3::Sort::bool(&ctx);
15     let bv = z3::Sort::bitvector(&ctx, 8);
16     if let Err(e) = check_same_sort(&bo, &bv) {
17         eprintln!("Sorts differ {:?}", e);
18     }
19 }
```

• *IsNotApp*

Purpose: Error type indicating an `Ast` is not an `App` node; i.e., not a function application.

Fields (private)

- `kind`: the actual `AstKind` encountered.

When to use

- To signal that pattern-matching on `App` failed.

How to use

```

1  use z3::ast::{Ast, Bool, Int};
2  use z3::IsNotApp;
3
```

```

4  fn get_app_args<'a, T: Ast<'a>>(ast: &T) -> Result<&T, IsNotApp> {
5      if ast.is_app() {
6          return Ok(ast);
7      }
8      Err(IsNotApp::new(ast.kind()))
9  }
10
11 fn main() {
12     let cfg = z3::Config::new();
13     let ctx = z3::Context::new(&cfg);
14     let term = z3::ast::Int::from_i64(&ctx, 42); // constants are apps
15     match get_app_args(&term) {
16         Err(e) => eprintln!("Expected App, found {:?}", e.kind()),
17         _ => (),
18     };
19
20     // forall x. x > 0 this is a Quantifier node, not an App
21     let x = z3::ast::Int::new_const(&ctx, "x");
22     let body = x.gt(&Int::from_i64(&ctx, 0));
23     let forall: Bool = z3::ast::forall_const(&ctx, [&x], [], &body);
24     match get_app_args(&forall) {
25         Err(e) => eprintln!("Expected App, found {:?}", e.kind()),
26         _ => (),
27     };
28 }

```

- **Solver**

Purpose: It is a *combined solver*, internally maintaining both a *non-incremental* solver and an *incremental* solver, switching behavior based on usage. It is attached to a `Context` and wraps a raw `Z3_solver`. Also supports `push/pop`, `check`, and `proof / model extraction`.

Fields: (private)

- `ctx`: reference to the owning `Context`.
- `z3_slv`: raw solver handle (`Z3_solver`).

When to use

- General Sat solving and proof generation.
- Quick feasibility checks during symbolic execution.
- Running under a specific tactic or logic via `Solver::new_for_logic`.

How to use

```

1  use z3::{ast, Config, Context, Solver};
2
3  fn main() {
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let _solver = Solver::new(&ctx);
7      let solver2 = Solver::new_for_logic(&ctx, "QF_LIA").unwrap(); // if the logic is
8      // unknown, it will give None, otherwise it will return Some(Solver)
9
10     // Build a simple constraint: x > 0 ∧ x < 5
11     let x = ast::Int::new_const(&ctx, "x");
12     let c1 = x.gt(&ast::Int::from_i64(&ctx, 0));
13     let c2 = x.lt(&ast::Int::from_i64(&ctx, 5));

```

```

14     solver2.assert(&c1);
15     solver2.assert(&c2);
16
17     match solver2.check() {
18         z3::SatResult::Sat => {
19             let m = solver2.get_model().unwrap();
20             println!("Model: x = {}", m.eval(&x, true).unwrap());
21         }
22         z3::SatResult::Unsat => println!("Unsat"),
23         z3::SatResult::Unknown => println!("Unknown"),
24     }
25 }
```

- **Model**

Purpose: Finite mapping from constants/functions to values that satisfies the assertions proven `Sat` by a `Solver`. Models can also be obtained from an `Optimize` context via `Model::of_optimize`.

Fields: (private)

- `ctx`: owning `Context`.
- `z3_mdl`: raw model handle (`Z3_model`).

When to use

- After `solver.check() == Sat` to inspect witness values. It can also be used when the solver spits *unknown*. In the case of *unknown*, a model will not be generated or it will be generated; however, it is not guaranteed to satisfy the constraints.
- In counter-example generation or test-case extraction.

How to use

```

1  fn main() {
2     use z3::fast::Ast, ast::Int, Config, Context, Solver;
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let x = Int::new_const(&ctx, "x");
9     let y = Int::new_const(&ctx, "y");
10
11
12     let _sum = Int::add(&ctx, [&x, &y]);
13     let _product = Int::mul(&ctx, [&x, &y]);
14     let le_constraint = x.le(&Int::from_i64(&ctx, 10));
15     let gt_constraint = y.gt(&Int::from_i64(&ctx, 0));
16
17     solver.assert(&le_constraint);
18     solver.assert(&gt_constraint);
19     solver.assert(&_product.eq(&_sum));
20
21     match solver.check() {
22         z3::SatResult::Sat => {
23             let model = solver.get_model().unwrap();
24             println!("Model found:");
25             println!("x: {}", model.eval(&x, false).unwrap());
26             println!("y: {}", model.eval(&y, false).unwrap());
27         }
28     }
29 }
```

```

28         z3::SatResult::Unsat => {
29             println!("No solution found.");
30         }
31         z3::SatResult::Unknown => {
32             println!("Solver could not determine satisfiability.");
33         }
34     }
35 }
```

- ***Optimize***

Purpose: Optimization engine built on top of Z3's Z3_optimize object. It supports hard & soft constraints along with multi-objective (Pareto) optimization.⁵

Fields: (private)

- ctx: owning Context.
- z3_opt: raw optimize handle (Z3_optimize).

When to use

- Maximizing or minimizing arithmetic terms under constraints.
- For Pareto optimization of several objectives.
- Ability to parse SMT-LIB2 optimization problems via `from_string`.

How to use

```

1  use z3::{ast, Config, Context, Optimize, SatResult};
2
3  fn main() {
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let opt = Optimize::new(&ctx);
7
8      // Create integer variables
9      let x = ast::Int::new_const(&ctx, "x");
10     let y = ast::Int::new_const(&ctx, "y");
11
12     // Hard constraints (must be satisfied)
13     opt.assert(&x.ge(&ast::Int::from_i64(&ctx, 0))); // x >= 0
14     opt.assert(&y.ge(&ast::Int::from_i64(&ctx, 0))); // y >= 0
15     opt.assert(&(x + y).le(&ast::Int::from_i64(&ctx, 20))); // x + y <= 20
16
17     // Soft constraints (weighted penalties for violation)
18     opt.assert_soft(
19         &x.le(&ast::Int::from_i64(&ctx, 8)), // x <= 8 (preferred)
20         10, // Weight = 10 (higher weights make
21             // constraints "harder" to violate)
22         Some("group1".into()), // Constraint group name
23     );
24
25     opt.assert_soft(
26         &y.le(&ast::Int::from_i64(&ctx, 5)), // y <= 5 (preferred)
27         5, // Weight = 5
28         Some("group2".into()),
29     );
30 }
```

⁵Pareto optimization is a mathematical framework for solving problems with multiple conflicting objectives where improving one objective inevitably worsens another.

```

30     // Optimization objectives
31     let profit = &x * ast::Int::from_i64(&ctx, 3) + &y * ast::Int::from_i64(&ctx,
32         ↵ 2); // 3x + 2y
33     opt.maximize(&profit); // Maximize profit
34
35     match opt.check(&[]) {
36         SatResult::Sat => {
37             let model = opt.get_model().unwrap();
38
39             // Evaluate variables
40             let x_val = model.eval(&x, true).unwrap().as_i64().unwrap();
41             let y_val = model.eval(&y, true).unwrap().as_i64().unwrap();
42             let profit_val = model.eval(&profit, true).unwrap().as_i64().unwrap();
43
44             // Check which soft constraints were satisfied
45             let soft1_satisfied = x_val <= 8;
46             let soft2_satisfied = y_val <= 5;
47
48             println!("Optimal solution found:");
49             println!("- x = {}, y = {}", x_val, y_val);
50             println!("- Profit = {}", profit_val);
51             println!("Soft constraints:");
52             println!(
53                 "- x <= 8: {} (weight: 10)",
54                 if soft1_satisfied { "yes" } else { "no" }
55             );
56             println!(
57                 "- y <= 5: {} (weight: 5)",
58                 if soft2_satisfied { "yes" } else { "no" }
59             );
60             SatResult::Unsat => println!("No solution exists"),
61             SatResult::Unknown => println!("Could not determine solution"),
62         }
63     }

```

Tips

- Objectives are returned in the same order you added them.
- Combine hard constraints (`assert`) and soft constraints (`assert_soft`) to model *best effort* goals.

• *FuncDecl*

Purpose: Describes the signature of a constant or n-ary function in Z3 (name, argument sorts, and result sort). A constant is just a zero-arity FuncDecl.

Fields: (private)

- `ctx`: owning Context.
- `z3_func_decl`: raw Z3_func_decl handle.

When to use

- When declaring new symbols: constants, uninterpreted functions, or datatype constructors.

How to use

```

1  use z3::{
2      ast::{self, Ast},

```

```

3     Config, Context, FuncDecl, Sort, Symbol,
4 };
5
6 fn main() {
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9     let int = Sort::int(&ctx);
10    let bool = Sort::bool(&ctx);
11
12    // Declare a function f : Int * Int -> Bool
13    let f_sym = Symbol::String("f".into());
14    let f = FuncDecl::new(&ctx, f_sym, &[&int, &int], &bool);
15    let x = f.arity(); // f.arity() == 2
16    assert_eq!(x, 2);
17
18    // Apply it to concrete arguments
19    let a = ast::Int::new_const(&ctx, "a");
20    let b = ast::Int::new_const(&ctx, "b");
21    let app = f.apply(&[&a, &b]);
22
23    // Check the type of the application
24    let app_sort: Sort<'_> = app.get_sort();
25    assert_eq!(app_sort, Sort::bool(&ctx));
26 }

```

- ***FuncInterp***

Purpose: Holds the concrete interpretation of a function symbol inside a **Model** (i.e. which value it returns for every argument tuple).

Fields: (private)

- `ctx`: owning `Context`.
- `z3_func_interp`: raw `Z3_func_decl` handle.

When to use

- After checking that the solver is *Sat*, use `model.get_func_interp(&func)` to inspect how an uninterpreted function was satisfied.

How to use

```

1 use z3::{
2     ast::{self, Ast, Dynamic},
3     Config, Context, FuncDecl, Sort, Symbol,
4 };
5
6 fn main() {
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9     let solver = z3::Solver::new(&ctx);
10
11    let int_sort = Sort::int(&ctx);
12    let f_sym = Symbol::String("f".into());
13    let f = FuncDecl::new(&ctx, f_sym, &[&int_sort], &int_sort);
14
15    let zero = ast::Int::from_i64(&ctx, 0);
16    let one = ast::Int::from_i64(&ctx, 1);
17
18    let zero_dyn: Dynamic = zero.clone().into();
19    let one_dyn: Dynamic = one.clone().into();

```

```

20
21 // Example constraints: f(0) = 1 and f(1) = 0
22 solver.assert(&f.apply(&[&zero])._eq(&one_dyn));
23 solver.assert(&f.apply(&[&one])._eq(&zero_dyn));
24
25 assert_eq!(solver.check(), z3::SatResult::Sat);
26
27 let model = solver.get_model().unwrap();
28
29 if let Some(x) = model.get_func_interp(&f) {
30     println!("Interpretation of f: {:?}", x);
31 }
32 }
```

- ***FuncEntry***

Purpose: Represents one row in a `FuncInterp` table (a concrete argument tuple and its mapped value).

Fields: (private)

- `ctx`: owning `Context`.
- `z3_func_entry`: raw `Z3_func_entry` handle.

When to use

- Iterate over explicit mappings and the function interpretation using `x.get_entries()`, which returns a vector of `FuncEntry` representing the function's explicit input-output entries, where `x` is a `FuncInterp`. Note that when you inspect a function interpretation in Z3 using `get_entries()`, you may see only a single entry even if the function is defined for more inputs. This is expected behavior: Z3 models functions as a finite map (explicit entries) plus a default else value. *Only those input tuples that are explicitly constrained in your problem will appear as entries.*

How to use

```

1 use z3::{
2     ast::{self, Ast, Dynamic},
3     Config, Context, FuncDecl, Sort, Symbol,
4 };
5
6 fn main() {
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9     let solver = z3::Solver::new(&ctx);
10
11     let int_sort = Sort::int(&ctx);
12     let f_sym = Symbol::String("f".into());
13     let f = FuncDecl::new(&ctx, f_sym, &[&int_sort], &int_sort);
14
15     let zero = ast::Int::from_i64(&ctx, 0);
16     let one = ast::Int::from_i64(&ctx, 1);
17
18     let zero_dyn: Dynamic = zero.clone().into();
19     let one_dyn: Dynamic = one.clone().into();
20
21 // Example constraints: f(0) = 1 and f(1) = 0
22 solver.assert(&f.apply(&[&zero])._eq(&one_dyn));
23 solver.assert(&f.apply(&[&one])._eq(&zero_dyn));
```

```

25     assert_eq!(solver.check(), z3::SatResult::Sat);
26
27     let model = solver.get_model().unwrap();
28
29     if let Some(x) = model.get_func_interp(&f) {
30         let entries = x.get_entries();
31         for entry in entries {
32             println!("entry: {:?}", entry);
33         }
34         println!("Interpretation of f: {:?}", x);
35     };
36 }

```

- ***RecFuncDecl***

Purpose: Function declaration that can carry a (possibly recursive) definition inside the solver via `add_def`. Therefore using this type we can define functions that have a definition unlike `FuncDecl` which was used to define uninterpreted functions.

Fields: (private)

- `ctx`: owning Context.
- `z3_func_decl`: raw `Z3_func_decl` handle.

When to use

- Encoding recursive definitions such as factorial, list length, or user-defined predicates.
- Also usable for non-recursive definitions if you plan to give them a symbolic definition instead of adding assertions.

How to use

```

1  use z3::RecFuncDecl;
2  use z3::{
3      ast::{self, Ast, Dynamic},
4      Config, Context, Solver, Sort, Symbol,
5  };
6
7  fn main() {
8      let config = Config::new();
9      let ctx = Context::new(&config);
10     let solver = Solver::new(&ctx);
11     let n_sort = Sort::int(&ctx);
12     let fact_sym = Symbol::String("fact".into());
13     let fact = RecFuncDecl::new(&ctx, fact_sym, &[&n_sort], &n_sort);
14
15     let n = ast::Int::new_const(&ctx, "n");
16     let one = ast::Int::from_i64(&ctx, 1);
17     let cond = n.le(&one);
18     let n_minus_1 = ast::Int::sub(&ctx, &[&n, &one]);
19     let rec = fact.apply(&[&n_minus_1]).as_int().unwrap();
20     let n_times_fact = ast::Int::mul(&ctx, &[&n, &rec]);
21     let body = cond.ite(&one, &n_times_fact);
22
23     fact.add_def(&[&n], &body); // attach definition
24     let five = ast::Int::from_i64(&ctx, 5);
25     let one_twenty = ast::Int::from_i64(&ctx, 120);
26     let one_twenty_dynamic: Dynamic = one_twenty.clone().into();
27     solver.assert(&fact.apply(&[&five])._eq(&one_twenty_dynamic));
28     match solver.check() {

```

```

29         z3::SatResult::Sat => {
30             println!("The factorial of 5 is 120.");
31         }
32         z3::SatResult::Unsat => {
33             println!("The assertion failed.");
34         }
35         z3::SatResult::Unknown => {
36             println!("The result is unknown.");
37         }
38     }
39 }
```

Tips

- After `add_def`, the solver treats calls to the function like an unfoldable definition (no axiom explosion).

• *DatatypeBuilder*

Purpose: Builder pattern for declaring an algebraic data type (ADT) in Z3 with variants, constructors, testers, and field accessors. It is used to model structured data such as enums, records, recursive trees, or mutually recursive types.

Fields: (private)

- `ctx`: owning `Context`.
- `name`: `Symbol` for the datatype.
- `constructors`: specification of each variant (`variant_name, [(field, DatatypeAccessor)]`).⁶

When to use

- Modelling Rust-style enums, structs, recursive trees, or mutually recursive types.

How to use

```

1  use z3::{Config, Context, DatatypeAccessor, DatatypeBuilder, DatatypeSort, Sort};
2
3  fn main() {
4      let config = Config::new();
5      let ctx = Context::new(&config);
6
7      // equivalent to:
8      // enum OptionInt {
9      //     IntNone,
10     //     IntSome { int_value: i32 },
11     // }
12    let option_int: DatatypeSort<'_> = DatatypeBuilder::new(&ctx, "OptionInt")
13        .variant("IntNone", vec![])
14        .variant(
15            "IntSome",
16            vec![("int_value", DatatypeAccessor::Sort(Sort::int(&ctx))]),
17        )
18        .finish(); // returns DatatypeSort
19
20    let option_int_sort: Sort<'_> = option_int.sort;
21    println!("OptionInt sort: {:?}", option_int_sort);
22
23    let v: &Vec<z3::DatatypeVariant<'_>> = &option_int.variants;
```

⁶More precisely, `Vec < (String, Vec < (String, DatatypeAccessor) >) >`.

```

24     println!("variants: {:?}", v);
25
26     let c1: &z3::DatatypeVariant<'_> = &v[0];
27     let c1_constructor = c1.constructor.apply(&[]);
28     let c1_tester = &v[0].tester;
29     let c1_accessor = &v[0].accessors;
30     println!("constructor for IntNone: {:?}", c1_constructor);
31     println!("tester for IntNone: {:?}", c1_tester);
32     println!("accessors for IntNone: {:?}", c1_accessor);
33
34     let c2 = &v[1];
35     let c2_constructor = c2.constructor.apply(&[&z3::ast::Int::from_i64(&ctx, 42)]);
36     let c2_tester = &v[1].tester;
37     let c2_accessor = &v[1].accessors;
38     println!("constructor for IntSome(42): {:?}", c2_constructor);
39     println!("tester for IntSome: {:?}", c2_tester);
40     println!("accessors for IntSome: {:?}", c2_accessor);
41
42     let option_bool = DatatypeBuilder::new(&ctx, "OptionBool")
43         .variant("BoolNone", vec![])
44         .variant(
45             "BoolSome",
46             vec![("bool_value", DatatypeAccessor::Sort(Sort::bool(&ctx))]),
47         )
48         .finish(); // returns DatatypeSort
49
50     let option_bool_sort = option_bool.sort();
51     println!("OptionBool sort: {:?}", option_bool_sort);
52
53     let w = &option_bool.variants;
54     println!("variants: {:?}", w);
55
56     let is_bool = &w[0].tester; // recogniser for OptionBool
57     println!("is_bool recogniser: {:?}", is_bool);
58 }

```

Tips

- Call `finish()` exactly once; it registers the sort and yields a `DatatypeSort` with ready-made constructors/testers.
- Variants may be mutually recursive: use `DatatypeAccessor::Datatype(other_symbol)` for fields.

• *DatatypeAccessor*

Purpose: Describes the type of a field inside a variant — either a plain `Sort` or a reference to another (possibly recursive) datatype.

Variants (public)

- `Sort(Sort)` – field has this concrete sort.
- `Datatype(Symbol)` – field is another ADT identified by symbol.

When to use

- While constructing a datatype with `DatatypeBuilder`.
- To express recursive/self-referential fields.

• *DatatypeVariant*

Purpose: Holds the Z3 artifacts for one variant of the ADT.

Fields: (public)

- **constructor**: FuncDecl to build the variant.
- **tester**: FuncDecl predicate *is* to check if the a value is of the respective type.
- **accessors**: field-projection FuncDecls.

When to use

- Constructing values: DatatypeSortType.variants[1].constructor.apply([&value]).
- Pattern-matching: DatatypeSortType.variants[0].tester.apply([&value]).
- Accessing fields inside the variant:
DatatypeSortType.variants[1].accessors[0].apply([&value]).

```
1  use z3::{
2      ast::{self, Ast},
3      Config, Context, DatatypeAccessor, DatatypeBuilder, DatatypeSort, Sort,
4  };
5  fn main() {
6      let cfg = Config::new();
7      let ctx = Context::new(&cfg);
8      let solver = z3::Solver::new(&ctx);
9
10     // enum OptionInt { IntNone, IntSome { int_value: Int } }
11     let option_int: DatatypeSort<'_> = DatatypeBuilder::new(&ctx, "OptionInt")
12         .variant("IntNone", vec![])
13         .variant(
14             "IntSome",
15             vec![("int_value", DatatypeAccessor::Sort(Sort::int(&ctx)))]),
16         )
17     .finish(); // register the sort DatatypeSort.sort and .variants[...]
18
19     // Create two constants of this sort
20     let x = ast::Datatype::new_const(&ctx, "x", &option_int.sort);
21     let y = ast::Datatype::new_const(&ctx, "y", &option_int.sort);
22
23     // Assert x is None and y == Some(3)
24     let none_tester = &option_int.variants[0].tester;           //
25     //   tester for IntNone
26     solver.assert(&none_tester.apply(&[&x]).as_bool().unwrap());
27
28     let some_ctor = &option_int.variants[1].constructor;          //
29     //   constructor for IntSome
30     let three = ast::Int::from_i64(&ctx, 3);
31     let some_three = some_ctor.apply(&[&three]);
32     solver.assert(&y._eq(&some_three.as_datatype().unwrap()));
33
34     assert_eq!(solver.check(), z3::SatResult::Sat);               //
35
36     // Extract the value inside the Some variant
37     let accessor = &option_int.variants[1].accessors[0];          //
38     //   the int_value accessor
39     let extracted = accessor.apply(&[&y]).as_int().unwrap();
40     let model = solver.get_model().unwrap();
41     let v = model.eval(&extracted, true).unwrap().as_i64().unwrap();
42     assert_eq!(v, 3);
43 }
```

- *DatatypeSort*

Purpose: Result of DatatypeBuilder::finish(); represents the actual ADT sort plus all its variants.

Fields: (public)

- `sort`: Sort object for typing terms.
- `variants`: `Vec<DatatypeVariant>`.

```

1  use z3::*;
2      ast::{self, Ast},
3      datatype_builder::create_datatypes,
4      Config, Context, DatatypeAccessor, DatatypeBuilder, DatatypeSort, Sort,
5  };
6
7  fn main() {
8      let cfg = Config::new();
9      let ctx = Context::new(&cfg);
10
11     run_z3_example(&ctx);
12 }
13
14 fn run_z3_example(ctx: &Context) {
15     let solver = z3::Solver::new(ctx);
16
17     // Build: enum Tree { Leaf(i32), Node { left: Tree, right: Tree } }
18     let tree_dt: DatatypeSort<'_> = DatatypeBuilder::new(ctx, "Tree")
19         .variant(
20             "Leaf",
21             vec![("value", DatatypeAccessor::Sort(Sort::int(ctx))]),
22         )
23         .variant(
24             "Node",
25             vec![
26                 ("left", DatatypeAccessor::Datatype("Tree".into())),
27                 ("right", DatatypeAccessor::Datatype("Tree".into())),
28             ],
29         )
30         .finish();
31
32     // Constants
33     let t1 = ast::Datatype::new_const(ctx, "t1", &tree_dt.sort);
34     let t2 = ast::Datatype::new_const(ctx, "t2", &tree_dt.sort);
35
36     // Construct t1 = Node(Leaf(1), Leaf(2))
37     let leaf_ctor = &tree_dt.variants[0].constructor;
38     let one = ast::Int::from_i64(ctx, 1);
39     let two = ast::Int::from_i64(ctx, 2);
40     let leaf1 = leaf_ctor.apply(&[&one]).as_datatype().unwrap();
41     let leaf2 = leaf_ctor.apply(&[&two]).as_datatype().unwrap();
42
43     let node_ctor = &tree_dt.variants[1].constructor;
44     let node12 = node_ctor.apply(&[&leaf1, &leaf2]).as_datatype().unwrap();
45     solver.assert(&t1.eq(&node12));
46
47     // Assert t2 is a Leaf
48     let is_leaf = &tree_dt.variants[0].tester;
49     solver.assert(&is_leaf.apply(&[&t2]).as_bool().unwrap());
50

```

```

51     assert_eq!(solver.check(), z3::SatResult::Sat);
52
53     // Access fields of Node
54     if let Some(m) = solver.get_model() {
55         // t1 should be a Node
56         let is_node = &tree_dt.variants[1].tester;
57         assert!(m
58             .eval(&is_node.apply(&[&t1]).as_bool().unwrap(), true)
59             .unwrap()
60             .as_bool()
61             .unwrap());
62     }
63
64     // and mutually recursive datatypes
65     let even_b = DatatypeBuilder::new(&ctx, "Even")
66         .variant("Zero", vec![]) // Even::Zero
67         .variant(
68             // Even::ESucc { pred : Odd }
69             "ESucc",
70             vec![("pred", DatatypeAccessor::Datatype("Odd".into()))],
71         );
72
73     // Odd datatype
74     let odd_b = DatatypeBuilder::new(&ctx, "Odd").variant(
75         // Odd::OSucc { pred : Even }
76         "OSucc",
77         vec![("pred", DatatypeAccessor::Datatype("Even".into()))],
78     );
79
80     let [_even_sort, _odd_sort]: [DatatypeSort; 2] =
81         create_datatypes(vec![even_b, odd_b]).try_into().unwrap();
82 }
```

Tips

- Use `constructor` to build values, `tester` to discriminate, and `accessors` to project fields.

• Params

Purpose: Key-value map of configurations used to configure many Z3 components (simplifiers, tactics, solvers, optimize, etc.)

Fields: (private)

- `ctx`: owning Context.
- `z3_params`: raw `Z3_params` handle.

When to use

- Fine-tune solver behaviour.
- Pass custom parameters to a tactic chain.

How to use

```

1  use z3::{Config, Context, Params, Tactic};
2
3  fn main() {
4      let config = Config::new();
5      let ctx = Context::new(&config);
6
7      let mut p = Params::new(&ctx);
```

```

8     p.set_u32("timeout", 2_000); // 2-second timeout
9     p.set_bool("mbqi", false); // disable MBQI
10
11    let tactics = Tactic::new(&ctx, "smt");
12    let solver = tactics.solver();
13    solver.set_params(&p);
14 }

```

Tips

- Call all `set_*`() before passing the params to a tactic/solver.

- ***SatResult***

Purpose: Represents the three possible outcomes of a satisfiability check – `Solver::check()` / `Optimize::check()`. `Unknown` covers timeout, interrupt, or incomplete theory.

Variants: (public) `Unsat` — `Unknown` — `Sat`

How to use

```

1  use z3::{Config, Context, Params, SatResult, Tactic};
2
3  fn main() {
4      let config = Config::new();
5      let ctx = Context::new(&config);
6
7      let mut p = Params::new(&ctx);
8      p.set_u32("timeout", 2_000); // 2-second timeout
9      p.set_bool("mbqi", false); // disable MBQI
10
11     let tactics = Tactic::new(&ctx, "smt");
12     let solver = tactics.solver();
13     solver.set_params(&p);
14
15     match solver.check() {
16         SatResult::Sat => println!("Model = {:?}", solver.get_model().unwrap()),
17         SatResult::Unsat => println!("Proof  = {:?}", solver.get_proof().unwrap()),
18         SatResult::Unknown => println!("Why?   = {:?}", solver.get_reason_unknown().unwrap()),
19     }
20 }

```

Regarding `Optimize::check()` the following code snippet presents an example:

```

1  use z3::ast::Ast;
2  use z3::{ast::Int, Config, Context, Optimize, SatResult};
3
4  fn main() {
5      let cfg = Config::new();
6      let ctx = Context::new(&cfg);
7      let opt = Optimize::new(&ctx);
8
9      let x = Int::new_const(&ctx, "x");
10     let y = Int::new_const(&ctx, "y");
11
12     // 0 <= x <= 10, 0 <= y <= 10, x + y <= 11
13     opt.assert(&x.ge(&Int::from_i64(&ctx, 0)));
14     opt.assert(&x.le(&Int::from_i64(&ctx, 10)));
15     opt.assert(&y.ge(&Int::from_i64(&ctx, 0)));
16     opt.assert(&y.le(&Int::from_i64(&ctx, 10)));
17     opt.assert(&Int::add(&ctx, &[&x, &y]).le(&Int::from_i64(&ctx, 11)));

```

```

18
19     // maximize x, maximize y
20     opt.maximize(&x);
21     opt.maximize(&y);
22
23     match opt.check(&[]) {
24         SatResult::Sat => {
25             let model = opt.get_model().unwrap();
26             let x_val = model.eval(&x, true).unwrap();
27             let y_val = model.eval(&y, true).unwrap();
28             println!("Optimal solution: x = {}, y = {}", x_val, y_val);
29         }
30         SatResult::Unsat => {
31             println!("No solution exists");
32         }
33         SatResult::Unknown => {
34             println!("Solver returned unknown: {:?}", opt.get_reason_unknown());
35         }
36     }
37 }
```

- ***Pattern***

Purpose: Explicit trigger for quantifier instantiation; influences the heuristic used by Z3's E-matching engine.

Fields: (private)

- `ctx`: owning Context.
- `z3_pattern`: raw Z3_pattern handle.

When to use

- Manually control which subterms trigger quantifier instantiation.
- Improve performance by avoiding automatic (and sometimes suboptimal) pattern inference.

How to use

```

1  use z3::{ast, ast::Ast, Config, Context, Pattern, Solver};
2
3  fn main() {
4      let config = Config::new();
5      let ctx = Context::new(&config);
6      let solver = Solver::new(&ctx);
7
8      let x = ast::Int::new_const(&ctx, "x");
9      let y = ast::Int::new_const(&ctx, "y");
10     let pat = Pattern::new(&ctx, &[&ast::Int::add(&ctx, [&x, &y])]);
11
12     let forall = ast::forall_const(
13         &ctx,
14         &[&x, &y],
15         &[&pat],
16         &ast::Int::add(&ctx, [&x, &y])._eq(&ast::Int::add(&ctx, [&y, &x])), );
17
18     solver.assert(&forall);
19 }
```

- ***ApplyResult***

Purpose: Captures the list of sub-goals produced by applying a Tactic to a Goal.

Fields: (private)

- ctx: owning Context.
- z3_apply_result: raw Z3_apply_result handle.

When to use

- Inspect intermediate goals after rewriting / simplification.
- Feed the resulting sub-goals to different tactics or solvers.

How to use

```
1  use z3::{
2      ast::Bool, Int,
3      ApplyResult, Config, Context, Goal, Solver, Tactic,
4  };
5
6  fn main() {
7      let cfg = Config::new();
8      let ctx = Context::new(&cfg);
9
10     // (produce_models = true, produce_unsat_cores = true, produce_proofs = false)
11     let goal = Goal::new(&ctx, true, true, false);
12
13     // x + y > 10 AND x < 0
14     let x = Int::new_const(&ctx, "x");
15     let y = Int::new_const(&ctx, "y");
16     let c1 = z3::ast::Int::add(&ctx, [&x, &y]).gt(&Int::from_i64(&ctx, 10));
17     let c2 = x.lt(&Int::from_i64(&ctx, 0));
18
19     goal.assert(&c1);
20     goal.assert(&c2);
21
22     println!("==== Original Goal ====");
23     for f in goal.get_formulas::<Bool>() {
24         println!("  {:?}", f);
25     }
26
27     // simplify (rewrites equations) then solve-eqs (solves simple equalities)
28     let simplify = Tactic::new(&ctx, "simplify");
29     let solve_eqs = Tactic::new(&ctx, "solve-eqs");
30
31     // first simplify, then solve-eqs
32     let pipeline = simplify.and_then(&solve_eqs);
33
34     // apply the tactic pipeline to the goal
35     let ar: ApplyResult<'_> = pipeline
36         .apply(&goal, None)
37         .expect("tactic application failed");
38
39     // solve each sub-goal separately
40     println!("\n==== Solving each sub-goal ====");
41     for (i, subg) in ar.list_subgoals().enumerate() {
42         println!("  Sub-goal {} formulas:", i);
43         let solver = Solver::new(&ctx);
44         for f in subg.get_formulas() {
45             println!("  {:?}", f);
46             solver.assert(&f);
```

```

47     }
48     let res = solver.check();
49     println!(" Sub-goal {} is {:?}", i, res);
50   }
51 }
```

- **Tactic**

Purpose: Encapsulates a reasoning strategy. Tactics can be composed to form pipelines or converted into a solver.

Fields: (private)

- ctx: owning Context.
- z3_tactic: raw Z3_tactic handle.

When to use

- Custom preprocessing before feeding a core solver.
- Domain-specific strategies (bit-blasting, omega, nnf, ...).

How to use

```

1  use std::time::Duration;
2  use z3::{Config, Context, Tactic};
3
4  fn main() {
5    let cfg = Config::new();
6    let ctx = Context::new(&cfg);
7
8    let t1 = Tactic::new(&ctx, "simplify");
9    let t2 = Tactic::new(&ctx, "smt");
10   let combo = t1.and_then(&t2).try_for(Duration::from_secs(5)); // 5-second limit
11
12  let _solver: z3::Solver<'_> = combo.solver(); // tactic-backed solver
13 }
```

Tips

- When you create a solver with tactics, the specified tactics are the *only* strategies used in checking the satisfiability of the formula. These tactics are used in order. Whereas if you created a solver with the *new* function, all default tactics would have been available.
- Use *Tactic::list_all()* to see a list of all tactics and use *Tactic::new()* to create a single one.

- **Goal**

Purpose: Represents an (immutable) set of formulas that can be solved or transformed by tactics and solvers.⁷

Fields: (private)

- ctx: owning Context.
- z3_goal: raw Z3_goal handle.

When to use

- Break a large set of assertions into independent sub-goals.
- Analyse goals with Probe before choosing a tactic.

⁷See the code example in the `ApplyResult` section for how `Goal` is used.

Tips

- Use `Goal::new(ctx, models, unsat_cores, proofs)` to create a new goal.
 - * `models` – **true** tells Z3 to store a *model* whenever the goal (or the solver that consumes it) is `Sat`. Set to **false** if you never need models and want to save memory/time.
 - * `unsat_cores` – **true** enables collection of an *unsat-core* (minimal subset of assertions causing `Unsat`) when the goal is unsatisfiable. Disable if you do not care about explanations.
 - * `proofs` – **true** enables production of full formal *proof objects*. Proof generation can be expensive; keep **false** unless you need proof terms.

- **Probe**

Purpose: A fast metric that inspects a goal (number of variables, is-linear, etc.). A Probe is essentially a little *sensor* you can stick on a Goal to ask Z3 for a quantitative property of your formulas. Think of it like asking, “How big is this goal?” or “How many bit-vector terms does it contain?” or “Does it mix arithmetic and arrays?”

Fields: (private)

- `ctx`: owning Context.
- `z3_probe`: raw `Z3_probe` handle.

When to use

- Choose between tactics; i.e., selecting a tactic based on goal-characteristics.
- Collect statistics for adaptive pipelines.

How to use

```
1  use z3::{ApplyResult, Config, Context, Goal, Probe, Tactic};  
2  
3  fn main() {  
4      let ctx = Context::new(&Config::new());  
5      let goal = Goal::new(&ctx, true, false, false);  
6  
7      // Probe the size of the goal  
8      let size = Probe::new(&ctx, "size").apply(&goal);  
9  
10     let size_tactic = Tactic::new(&ctx, "bit-blast");  
11     let base = Tactic::new(&ctx, "simplify");  
12     let pipeline = if size > 10.0 {  
13         base.and_then(&size_tactic)  
14     } else {  
15         base  
16     };  
17  
18     let ar: ApplyResult = pipeline.apply(&goal, None).unwrap();  
19  
20     for g in ar.list_subgoals() {  
21         println!("Subgoal: {}", g);  
22     }  
23 }
```

Tips

- Write `Probe::list_all(&ctx)` to get a list of all the probes.

- *Statistics*

Purpose: Holds profiling data (counters (UInt) or timings (Double)) produced by a solver or tactic pipeline.

Fields: (private)

- `ctx`: owning Context.
- `z3_stats`: raw Z3_stats handle.

When to use

- After `solver.check()` to see time spent in each module.
- Compare tactic pipelines quantitatively.

How to use

```

1  use z3::{ast::Int, Config, Context, SatResult, Solver};
2
3  fn main() {
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let solver = Solver::new(&ctx);
7
8      let x = Int::new_const(&ctx, "x");
9      let y = Int::new_const(&ctx, "y");
10
11     // x + y > 10, x < 5
12     solver.assert(&Int::add(&ctx, [&x, &y]).gt(&Int::from_i64(&ctx, 10)));
13     solver.assert(&x.lt(&Int::from_i64(&ctx, 5)));
14
15     match solver.check() {
16         SatResult::Sat => println!("SAT"),
17         SatResult::Unsat => println!("UNSAT"),
18         SatResult::Unknown => println!("UNKNOWN"),
19     }
20
21     // Print solver statistics
22     let stats = solver.get_statistics();
23     println!("\n==== Z3 Solver Statistics ====");
24     for x in stats.entries() {
25         println!("{}: {}", x);
26     }
27 }
```

Tips

- After running the above code, the following statistics will be provided in the terminal:

```

StatisticsEntry { key: "elim-unconstrained", value: UInt(5) }
StatisticsEntry { key: "solve-eqs-steps", value: UInt(2) }
StatisticsEntry { key: "solve-eqs-elim-vars", value: UInt(2) }
StatisticsEntry { key: "num allocs", value: UInt(29699) }
StatisticsEntry { key: "rlimit count", value: UInt(177) }
StatisticsEntry { key: "max memory", value: Double(17.76) }
StatisticsEntry { key: "memory", value: Double(17.76) }
StatisticsEntry { key: "time", value: Double(0.007) }
```

- The statistic `elim-unconstrained` (5) indicates that Z3 removed five variables deemed irrelevant to the final result; `solve-eqs-steps` (2) and `solve-eqs-elim-vars` (2) show it took two elimination steps to remove two variables via equation solving; `num allocs` (29706) tracks the total number of internal memory allocations; `rlimit count` (171) records how many times Z3 hit its resource-limit checks (used for timeouts and fairness); `max memory` and `memory` (17.74 MB) report peak and current memory usage; and `time` (0.004 s) is the total solve time in seconds. These entries help you diagnose performance bottlenecks and resource consumption in your queries.
- Depending on the logic and the query, the statistics can present different keys; e.g., *propagations*, *mk clause*, etc.

Before moving on to the individual modules of the `z3` crate, we will provide extra explanation of the core abstractions: `Goal`, `Tactic`, `Solver`, `Probe`, and `Statistics`.

- **Goal vs. Tactic vs. Solver**

- **Goal** A *bag of formulas*. Think “workspace”: you `assert` formulas into it, but it does not solve anything.
- **Tactic** A *transformer* that consumes a goal and returns an `ApplyResult` (usually a list of simpler sub-goals). Examples: “`simplify`”, “`qe`”.
- **Solver** A complete decision procedure that ultimately returns `Sat`/`Unsat`/`Unknown`. Internally a solver *contains* a tactic pipeline (either the Z3 default or one you supply via `tactic.solver()`).

- **Relation**

`Goal` –> `Tactic` => new `Goal` ... –> finally fed to a `Solver`

- **Probe vs. Statistics**

- **Probe** Query *before* solving. A probe inspects a goal and accordingly based on the data is used to drive tactic selection: `Tactic::when(p, heavy, light)`.
- **Statistics** Collected *after* solving. A statistics object is a map of counters/timers. Useful for profiling and performance debugging.⁸

7 *version* module

This module has been extensively discussed in Section 5. To reiterate, the following public items are accessible:

- **Version:** This struct holds the four-part Z3 version. The first part is *major*, which is the first (left-most) number, and represents significant changes that break backward compatibility or represent major architectural changes. The second part is equivalent to *minor* in *semantic versioning*, and represents significant enhancements that maintain backward compatibility. The third part is *Build number*, which represents a recompilation of the same source for different platforms or compilers. The last number is *Revision number*, often linked to fully interchangeable fixes to shipped versions. This is defined as follows:

```

1  pub struct Version {
2      major: u32,
3      minor: u32,
4      build_number: u32,
5      revision_number: u32, }
```

⁸**Mnemonic:** *Probe* = “peek first”; *Statistics* = “summary after”.

- pub fn `version()` -> `Version`: Internally calls `Z3_get_version` via FFI (Z3 C API) and returns the Z3 `Version`. Specifically, it first creates a zero initialized version with `Version::default()`. Then, passes the four field addresses to the raw C API `Z3_get_version` and returns the now-populated version.
- pub fn `full_version()` -> `&'static str`: Internally calls `Z3_get_full_version()` and returns the full version string as `&'static str`.

```

1 fn main() {
2     use z3::version;
3     use z3::full_version;
4
5     let ver = version();
6     println!("Z3 version: {:?}", ver);
7
8     let ver2 = full_version();
9     println!("Z3 full version: {:?}", ver2);
10 }
```

The previous code gives the following output on my computer:

```
Z3 version: Version { major: 4, minor: 14, build_number: 1, revision_number: 0 }
Z3 full version: "4.14.1.0"
```

8 tactic module

The following public items are accessible:

- `impl ApplyResult:`
 - pub fn `list_subgoals(self)` -> `impl Iterator<Item = Goal>`: Iterate over the sub-goals produced by the application of a tactic on a *goal*. The following code snippet produces multiple sub-goals after applying the tactic on the original goal:

```

1 use z3::{
2     ast::{Bool, Int},
3     Config, Context, Goal, Tactic,
4 };
5
6 fn main() {
7     let cfg = Config::new();
8     let ctx = Context::new(&cfg);
9
10    let x = Int::new_const(&ctx, "x");
11    let y = Int::new_const(&ctx, "y");
12    let z = Int::new_const(&ctx, "z");
13    let one = Int::from_i64(&ctx, 1);
14    let zero = Int::from_i64(&ctx, 0);
15
16    // (x > 1) ∨ (y > 1) ∨ (z > 1)
17    let big_or = Bool::or(&ctx, &[&x.gt(&one), &y.gt(&one), &z.gt(&one)]);
18
19    // (x < 0) ∨ (y < 0)
20    let other_or = Bool::or(&ctx, &[&x.lt(&zero), &y.lt(&zero)]);
21
22    // Goal
23    let g = Goal::new(
24        &ctx, /*models=*/ false, /*unsat_core=*/ false, /*proofs=*/ false,
25    );
26    g.assert(&big_or);
27    g.assert(&other_or);
```

```

28
29     // simplify then split the disjunctions
30     let tactic = Tactic::new(&ctx, "simplify").and_then(&Tactic::new(&ctx,
31         "split-clause"));
32
33     // apply the tactic to the goal to get the apply result sub-goals
34     let result = tactic.apply(&g, None).unwrap();
35
36     // Print every sub-goal produced
37     for (i, sg) in result.list_subgoals().enumerate() {
38         println!("Sub-goal {}:{}\n", i, sg);
39     }

```

- **impl Tactic:** The `tactic` API lets you programmatically compose Z3 *tactics*.⁹

Tactics vs. Logic: In short, **Logics** are *declarative*; i.e., they specify what theories to use whereas **Tactics** are *procedural*; i.e., they specify how to solve problems.

Logic: A *declaration of the theory fragment* you intend to use. For example, QF_BV (quantifier-free bit-vectors), QF_LIA (quantifier-free linear integer arithmetic), AUFLIRA (Arrays + Uninterpreted Functions + Linear Real Arithmetic), etc. Choosing a logic tells Z3 *which sorts, functions, and decision procedures* are allowed and will enable specialized back-ends.

Tactic: A *programmable strategy* that rewrites or solves goals step by step. Examples include `simplify`, `bit-blast`, `qe`, `solve-eqs`, or whole pipelines like `qfbv`. They let you *control exactly how Z3 transforms and dispatches* your formulas to its core solver. Without a logic declaration, Z3 falls back to a general-purpose `smt` tactic.

- pub fn `list_all(ctx: &Context) -> impl Iterator<Item = Result<&str, Utf8Error>>`¹⁰: Enumerate every tactic name known to Z3.

```

1  fn main() {
2      use z3::{Config, Context, Tactic};
3
4      let ctx = Context::new(&Config::new());
5      for name in Tactic::list_all(&ctx).flatten() {
6          println!("{}{}", name);
7      }
8  }

```

A more detailed explanation about each of the tactics can be found on the [Z3](#) guide.

- pub fn `new(ctx: &Context, name: &str) -> Tactic`: Create a tactic by its name (e.g. "smt", "solve-eqs"). Please look at the example at the end of the previous page. Note that the string must be a valid tactic name.
- pub fn `create_skip(ctx: &Context) -> Tactic`: A tactic that returns the goal unchanged. This no-op may sound useless, but it is essential whenever building compound strategies. Combinators such as `or_else`, `cond`, or `if` need a safe fall-back that leaves the goal intact when a probe fails or a branch condition is false.
- pub fn `create_fail(ctx: &Context) -> Tactic`: A tactic that always fails. It gives a way to deliberately abort a pipeline when a probe or condition is not met. Specifically, because tactics are combined, sometimes we need to deliberately *short-circuit* the pipeline.

⁹Algorithmic recipes for preprocessing or solving goals.

¹⁰Note that these signatures are meant as a detailed overview and may not exactly match the function definitions; they may omit lifetime annotations and other details.

```

1 fn main() {
2     use z3::{ApplyResult, Config, Context, Goal, Probe, Tactic};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let g = Goal::new(&ctx, true, false, false);
7
8     let p = Probe::new(&ctx, "size");
9     let p_le_10 = p.le(&Probe::constant(&ctx, 10.0));
10
11    let simplify_heavy = Tactic::new(&ctx, "ctx-solver-simplify");
12
13    // if size <= 10 then run heavy simplifier, else skip
14    let pipeline = Tactic::cond(&ctx, &p_le_10, &simplify_heavy,
15                                &Tactic::create_skip(&ctx));
16
17    let ar: ApplyResult = pipeline.apply(&g, None).unwrap();
18    for sub in ar.list_subgoals() {
19        println!("Sub-goal: {:?}", sub);
20    }
}

```

- pub fn repeat(ctx: &Context, t: &Tactic, max: u32) -> Tactic: Keep applying tactic t until fix-point (until the goal is not modified anymore) or *max* iterations has been reached.
- pub fn try_for(&self, timeout: Duration) -> Tactic: Wrap self with a wall-clock timeout. Abort tactic if it doesn't terminate within the period specified by *timeout*.
- pub fn and_then(&self, then_tactic: &Tactic) -> Tactic: Sequential composition - first self, then *then_tactic* is applied on each sub-goal produced by the application of the original tactic on the original goal.
- pub fn or_else(&self, else_t: &Tactic) -> Tactic: Fallback - use *self* on the given goal, only on failure apply *else_t* on the given goal.
- pub fn probe_or_else(&self, p: &Probe, t: &Tactic) -> Tactic: Probe-guarded choice between *self* and *t*. If probe *p* holds on the goal, use *self*, otherwise *t*.
- pub fn when(&self, p: &Probe) -> Tactic: Run *self* only if *p* is true, else behave like *skip*.
- pub fn cond(ctx: &Context, p: &Probe, t1: &Tactic, t2: &Tactic) -> Tactic: Probe-based if-then-else between two tactics. *if p then t1 else t2*.
- pub fn fail_if(ctx: &Context, p: &Probe) -> Tactic: Produce a failing tactic when *p* is false.
- pub fn apply(&self, goal: &Goal, params: Option<&Params>) -> Result<ApplyResult, String>: Run the tactic on a goal (with optional parameters). If the tactic succeeds, returns *Ok(ApplyResult)*.
- pub fn solver(&self) -> Solver: Convert the tactic into a solver that allows adding assertions and checking their satisfiability. Internally a solver contains a tactic pipeline as its processing strategy.

```

1 use std::time::Duration;
2 use z3::{
3     ast::{Ast, Bool, Int},
4     Config, Context, Goal, Params, Probe, Tactic,
5 };
6
7 fn main() {
8     let cfg = Config::new();

```

```

9   let ctx = Context::new(&cfg);
10
11  let x = Int::new_const(&ctx, "x");
12  let y = Int::new_const(&ctx, "y");
13
14  for i in Tactic::list_all(&ctx) {
15      println!("Tactic: {:?}", i);
16  }
17  // Build a goal and assert constraints
18  let goal = Goal::new(&ctx, false, false, false);
19  goal.assert(&Int::add(&ctx, &x, &y).eq(&Int::from_i64(&ctx, 10))); // x + y = 10
20  goal.assert(&x.ge(&Int::from_i64(&ctx, 0))); // x >= 0
21  goal.assert(&y.ge(&Int::from_i64(&ctx, 0))); // y >= 0
22
23  // tactics
24  let simp = Tactic::new(&ctx, "simplify");
25  let solve_eqs = Tactic::new(&ctx, "solve-eqs");
26
27  // try simplification for 100 ms
28  let simp_timeout = simp.try_for(Duration::from_millis(100));
29
30  // repeat simplification up to 5 times or until fix-point
31  let simp_repeat = Tactic::repeat(&ctx, &simp_timeout, 5);
32
33  // sequentially simplify then solve equations
34  let simplify_then_solve = simp_repeat.and_then(&solve_eqs);
35
36  // fallback: if that pipeline fails, just use the default SMT tactic
37  let default_smt = Tactic::new(&ctx, "smt");
38  let pipeline = simplify_then_solve.or_else(&default_smt);
39
40  // check if the goal size is larger than a threshold
41  let size = Probe::new(&ctx, "size");
42  let threshold = Probe::constant(&ctx, 2.0);
43  let size_large = size.gt(&threshold);
44  // if size_large holds, use pipeline; otherwise, skip (no-op)
45  let smart_tactic = pipeline.when(&size_large);
46
47  // if the goal is not a linear real arithmetic problem, use a different tactic
48  let non_linear = Probe::new(&ctx, "is-lra");
49  let nlsat = Tactic::new(&ctx, "nlsat");
50  let cond_tactic = Tactic::cond(&ctx, &non_linear, &nlsat, &smart_tactic);
51
52  // fail if already solved (size == 0), to demonstrate fail_if
53  let size_zero = Probe::new(&ctx, "size").eq(&Probe::constant(&ctx, 0.0));
54  let guarded_tactic = Tactic::fail_if(&ctx, &size_zero).and_then(&cond_tactic);
55
56  // apply the composed tactic to our goal
57  let result = guarded_tactic
58      .apply(&goal, Some(&Params::new(&ctx)))
59      .expect("Tactic application failed");
60
61  let all_formulas: Vec<Bool> = result
62      .list_subgoals()
63      .flat_map(|goal| {
64          let formulas: Vec<Bool> = goal.get_formulas().iter().cloned().collect();
65          formulas.into_iter()
66      })
67      .collect();

```

```

68
69 // Convert the tactic into a solver and check satisfiability
70 let solver = guarded_tactic.solver();
71 solver.assert(&Bool::and(&ctx, &all_formulas));
72 match solver.check() {
73     z3::SatResult::Sat => {
74         let model = solver.get_model().unwrap();
75         println!(
76             "sat; model: x = {}, y = {}",
77             model.eval(&x, true).unwrap(),
78             model.eval(&y, true).unwrap(),
79         );
80     }
81     z3::SatResult::Unsat => println!("unsat"),
82     z3::SatResult::Unknown => println!("unknown"),
83 }
84 }
```

9 symbol module

Z3 uses a unified `Symbol` type to represent identifiers for functions, sorts, constants, and datatypes. Internally, a symbol may be an integer (e.g. `Symbol::Int(0)`) or a string (e.g. `Symbol::String("x")`). The following methods are accessible:

- `pub fn as_z3_symbol(&self, ctx: &Context) -> Z3_symbol`: This method is called internally whenever you invoke a Z3 API function that needs a symbol argument. It turns your Rust-side `Symbol` into a Z3-internal symbol that the Z3 C API expects.

```

1 use z3::{ast::Int, Config, Context, FuncDecl, Solver, Sort, Symbol};
2
3 fn main() {
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6
7     let sym_id: Symbol = 42u32.into(); // -> Symbol::Int(42)
8     let sym_f = Symbol::from("f"); // -> Symbol::String("f")
9     let _sym_x = Symbol::from(String::from("x")); // -> Symbol::String("x")
10
11    let raw_f = sym_f.as_z3_symbol(&ctx); // Z3_mk_string_symbol
12    let raw_id = sym_id.as_z3_symbol(&ctx); // Z3_mk_int_symbol
13    println!("Raw symbols: {raw_f:?}, {raw_id:?}");
14
15    let int = Sort::int(&ctx);
16    let bool = Sort::bool(&ctx);
17
18    // FuncDecl::new(ctx, name, domain_sorts, &range_sort)
19    let f_decl = FuncDecl::new(&ctx, sym_f.clone(), &[&int, &int], &bool);
20
21    let x = Int::new_const(&ctx, "x"); // uses a string literal
22    let y = Int::new_const(&ctx, "y");
23
24    let fxy = f_decl.apply(&[&x, &y]).as_bool().unwrap();
25
26    let forty_two = Int::from_i64(&ctx, 42);
27    let f42y = f_decl.apply(&[&forty_two, &y]).as_bool().unwrap();
28
29    let solver = Solver::new(&ctx);
30    solver.assert(&fxy); // assert f(x, y)
31    solver.assert(&f42y.not()); // assert not f(42, y)
32 }
```

```

33     println!("sat? {:#?}", solver.check());
34     if solver.check() == z3::SatResult::Sat {
35         let model = solver.get_model().unwrap();
36         println!("Model: {model:#?}");
37     } else {
38         println!("No solution found.");
39     }
40 }
```

- The following From implementations are provided:
 - From<u32> → Symbol::Int
 - From<String> → Symbol::String
 - From<&str> → Symbol::String

10 *statistics* module

The Statistics, StatisticsValue, and StatisticsEntry types have been extensively discussed before, specifically in section 4. The Statistics holds profiling data produced by a solver or tactic pipeline. The following are additional items not covered in the previous sections:

- **impl Statistics**
 - pub fn value(&self, key: &str) → Option<StatisticsValue> Look up the statistic named key, returning its value if present.
 - pub fn entries(&self) → impl Iterator<Item = StatisticsEntry> Iterate over every (key, value) pair in the statistics set.

```

1  fn main() {
2      use z3::{Config, Context, Solver, StatisticsValue};
3
4      let ctx = Context::new(&Config::new());
5      let solver = Solver::new(&ctx);
6      let stats = solver.get_statistics();
7      if let Some(StatisticsValue::Double(t)) = stats.value("memory") {
8          println!("Solver uses {:.2} memory", t);
9      }
10     for entry in stats.entries() {
11         println!("{}: {:#?}", entry.key, entry.value);
12     }
13 }
```

11 *sort* module

Sort is the Z3 *type* object. The different kind of Z3 types has been extensively covered in section 2. The following represent the methods provided by the high-level Rust Z3 API for working with sorts:

- pub fn get_z3_sort(&self) → Z3_sort: Exposes the raw C Z3_sort handle.
- pub fn uninterpreted(ctx: &Context, name: Symbol) → Sort: Build a user-defined sort such as *type parameters* & *generics*. For full details, please visit section 2.
- pub fn bool(ctx: &Context) → Sort: Boolean sort.
- pub fn int(ctx: &Context) → Sort: Mathematical integers.
- pub fn real(ctx: &Context) → Sort: Mathematical real numbers.
- pub fn float(ctx: &Context, ebits: u32, sbits: u32) → Sort: IEEE-754 generic float (ebits exponent, sbits significand).

- pub fn float32(ctx: &Context) -> Sort: 32-bit IEEE float (ebits = 8 & sbits = 24).
- pub fn double(ctx: &Context) -> Sort: 64-bit IEEE double (ebits = 11 & sbits = 53).
- pub fn string(ctx: &Context) -> Sort: Z3's sequence/Unicode string sort.
- pub fn bitvector(ctx: &Context, sz: u32) -> Sort: Fixed-width bit-vector of n bits.
- pub fn array(ctx: &Context, domain: &Sort, range: &Sort) -> Sort : domain \rightarrow range array sort.
- pub fn set(ctx: &Context, elt: &Sort) -> Sort: Finite set (*array-like*) with Boolean range.
- pub fn seq(ctx: &Context, elt: &Sort) -> Sort: Sequence (*ordered collections*) of type elt.
- pub fn enumeration(ctx: &Context, name: Symbol, enum_names: &[Symbol],) -> (Sort, Vec<FuncDecl>, Vec<FuncDecl>): Creates an enum along with its associated constructors and testers. This function is fundamental for modeling discrete, finite sets of values where each value is distinct and mutually exclusive. It receives a symbol representing the name of the enumeration sort itself and an array of symbols representing the names of individual variants as inputs. The outputs are the enum type, constructor function declarations, and testers.

```

1  use z3::{
2      ast::{Ast, Dynamic},
3      Config, Context, SatResult, Solver, Sort, Symbol,
4  };
5
6  fn main() {
7      let cfg = Config::new();
8      let ctx = Context::new(&cfg);
9
10     let (color_sort, constructors, testers) = Sort::enumeration(
11         &ctx,
12         Symbol::from("color"),
13         &[
14             Symbol::from("Red"),
15             Symbol::from("Green"),
16             Symbol::from("Blue"),
17         ],
18     );
19
20     let red_constructor = &constructors[0]; // Creates Red values
21     let green_constructor = &constructors[1]; // Creates Green values
22     let blue_constructor = &constructors[2]; // Creates Blue values
23
24     let is_red = &testers[0];
25     let is_green = &testers[1];
26     let is_blue = &testers[2];
27
28     let red_value = red_constructor.apply(&[]);
29     let green_value = green_constructor.apply(&[]);
30     let blue_value = blue_constructor.apply(&[]);
31
32     let x = Dynamic::new_const(&ctx, "x", &color_sort);
33
34     let solver = Solver::new(&ctx);
35     solver.assert(&x._eq(&red_value));
36
37     match solver.check() {

```

```

38     SatResult::Sat => {
39         let model = solver.get_model().unwrap();
40         println!("x = {:?}", model.eval(&x, true));
41         let is_red = is_red.apply(&[&x]);
42         println!("x is Red: {:?}", model.eval(&is_red, true));
43     }
44     - => {
45         println!("No solution found");
46     }
47 }
48 }
```

- pub fn kind(&self) -> SortKind: Return the low-level Z3 C API type.
- pub fn float_exponent_size(&self) -> Option<u32>: Returns the exponent width wrapped inside Some if this is a float sort and if not, None.
- pub fn float_significand_size(&self) -> Option<u32>: Returns the significand width wrapped inside Some if this is float sort and if not, None.
- pub fn is_array(&self) -> bool: True for Array and Set sorts and false otherwise.
- pub fn array_domain(&self) -> Option<Sort>: Returns the domain sort of an array/set.
- pub fn array_range(&self) -> Option<Sort>: Returns the range sort of an array/set (Bool is always returned for sets).

SortDiffer Utility: The SortDiffer struct is used to represent and report a sort mismatch error when two Sort instances cannot be meaningfully compared or combined. This typically arises in typed AST construction when an operation (e.g., equality, function application, or conditional branches) involves terms of incompatible sorts. It holds two fields: left and right, both of type Sort, and implements Display to format a clear error message. For example, comparing an Int and a Bool will produce a SortDiffer error: “Cannot compare nodes, Sort does not match. Nodes contain types Int and Bool.” This type is useful for debugging and enforcing sort correctness in higher-level Z3 abstractions.

- impl SortDiffer:
 - pub fn new(left: Sort, right: Sort) -> Self: Package two mismatched sorts into an error.
 - left(&self) -> &Sort, right(&self) -> &Sort: Access the offending sorts for diagnostics.

```

1 fn main() {
2     use z3::{Config, Context, Sort, SortKind};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let s = Sort::array(&ctx, &Sort::int(&ctx), &Sort::bool(&ctx));
7
8     assert_eq!(s.kind(), SortKind::Array);
9     assert!(!s.is_array());
10    assert_eq!(s.array_domain(), Some(Sort::int(&ctx)));
11    assert_eq!(s.array_range(), Some(Sort::bool(&ctx)));
12
13    let fp = Sort::double(&ctx);
14    assert_eq!(fp.float_exponent_size(), Some(11));
15    assert_eq!(fp.float_significand_size(), Some(53));
16
17    match assert_equal(&s, &fp) {
18        Ok(_) => println!("Sorts are equal"),
19        Err(e) => println!("Sorts differ: {:?}", e),
20    }
21 }
```

```

20     };
21 }
22
23 use z3::{Sort, SortDiffers};
24 fn assert_equal<'a>(a: &'a Sort, b: &'a Sort) -> Result<(), SortDiffers<'a>> {
25     if a != b {
26         return Err(SortDiffers::new(a.clone(), b.clone()));
27     }
28     Ok(())
29 }
```

Tips on numeric sorts in Z3

- Integer-like vs. Bit-Vector vs. Real vs. Float
 - **Int** — the sort of *mathematical integers* ($\dots, -2, -1, 0, 1, 2, \dots$). Operations follow pure number-theory rules: addition never overflows and division is exact.
 - **BitVector n** — an n -bit *machine word*. Arithmetic wraps modulo 2^n exactly like CPU registers; you can also extract individual bits or slices. Use this sort for low-level purposes for example when modeling hardware.
 - **Real** — the sort of *rationals*. Represents mathematical real numbers with exact algebraic precision, allowing you to write constraints that reason about fractions, square-roots, and polynomial roots without the rounding errors of floating-point arithmetic – i.e., Z3 internally treats *reals* as exact rational values (like $1/3$ stays as $1/3$, not 0.333).
 - **Float(e,s)** — an *IEEE-754 floating point* value with a fixed number of exponent bits (e) and significand bits (s). Arithmetic follows IEEE rounding rules and can distinguish $+0$ from -0 , NaN , Inf , etc.
- **float** vs **float32** vs **double**
 - **float(e,s)** — generic constructor; you choose any exponent-significand width.
 - **float32** — shorthand for the 32-bit IEEE single precision format ($e = 8$ bits, $s = 24$ bits).
 - **double** — shorthand for 64-bit IEEE double precision format ($e = 11$, $s = 53$).
- Exponent vs. Significand (IEEE-754) A floating-point number is stored as

$$(\text{sign}) \times 2^{\text{exponent}} \times (1.\text{significand bits})$$
 - **Exponent bits** (e) decide *how large or small* the number can be (range of magnitudes).
 - **Significand bits** (s) decide *how many significant digits* you can keep (precision of number).

In short, more exponent bits \Rightarrow wider range; more significand bits \Rightarrow finer precision.

12 solver module

This module provides the associated functions to the **Solver** struct data type defined in the *lib.rs* module:

- **pub fn new(ctx: &Context) -> Solver:** Create Z3’s *combined* solver. When you call **Solver::new(&ctx)** Z3 actually builds *two* engines under the hood; an *incremental* and *non-incremental* solver. The default solver is the *non-incremental* one; i.e., **solver1**. The moment you start working incrementally—calling **Solver::push()**/**Solver::pop()** or asserting new constraints after a satisfiability check with **Solver::assert()** or **Solver::assert_and_track()** without an intervening **Solver::reset()**—Z3 transparently switches to **solver2**, a fully incremental SMT engine. You can tweak this behaviour via the parameters **solver2_timeout**, **solver2_unknown**, and **ignore_solver1**. More often than not, using the *new* function is how you want to build a *solver*, unless you want a solver with custom tactic pipeline.

What is *non-incremental*? You just call `assert()` and `check()` once without modifying the solver again. `solver1` chooses a tactic based on the logic of your assertions. If it matches a known logic, Z3 picks a specialized tactic (e.g., `qflia`, `qfbv`, etc.) Otherwise, it runs the general tactic chain: (`and-then simplify smt`). This means: first simplify, then run the general SMT solver.

What is *incremental*? If you use `push()`, `pop()`, or call `assert()/assert_and_track` again after a check without a `reset`, the solver upgrades to `solver2`, which behaves like a stateful SMT solver with support for backtracking. It always uses the general-purpose `smt` tactic and maintains the assertion stack. The solver cannot go back to the non-incremental version.

One-shot solving (uses `solver1`)

```

1 fn main() {
2     use z3::{ast::Int, Config, Context, SatResult, Solver};
3     let cfg = Config::new();
4     let ctx = Context::new(&cfg);
5     let solver = Solver::new(&ctx);
6
7     let x = Int::new_const(&ctx, "x");
8     solver.assert(&x.gt(&Int::from_i64(&ctx, 10)));
9     solver.assert(&x.lt(&Int::from_i64(&ctx, 20))); // not an assertion after the check
10    assert_eq!(solver.check(), SatResult::Sat);
11 }
```

Incremental solving (switches to `solver2`)

```

1 fn main() {
2     use z3::{ast::Ast, ast::Int, Config, Context, Solver};
3     let cfg = Config::new();
4     let ctx = Context::new(&cfg);
5     let solver = Solver::new(&ctx);
6
7     let x = Int::new_const(&ctx, "x");
8     let y = Int::new_const(&ctx, "y");
9     let five = Int::from_i64(&ctx, 5);
10    let ten = Int::from_i64(&ctx, 10);
11    solver.push();
12    solver.assert(&x.eq(&y));
13    solver.assert(&x.gt(&ten));
14    solver.assert(&y.lt(&five));
15    match solver.check() {
16        z3::SatResult::Sat => println!("x == y"),
17        z3::SatResult::Unsat => println!("x != y"),
18        z3::SatResult::Unknown => println!("unknown"),
19    }
20    solver.pop(1); // pop everything after the last push before (until line 11)
21    match solver.check() {
22        z3::SatResult::Sat => {
23            println!("x == y");
24            let model = solver.get_model().unwrap();
25            println!("x = {}", model.eval(&x, true).unwrap());
26            println!("y = {}", model.eval(&y, true).unwrap());
27        }
28        z3::SatResult::Unsat => println!("x != y"),
29        z3::SatResult::Unknown => println!("unknown"),
30    }
31 }
```

Use Params to override default behaviour

```
1 fn main() {
2     use z3::{Config, Context, Params, SatResult, Solver};
3     let ctx = Context::new(&Config::new());
4     let solver = Solver::new(&ctx);
5
6     let mut params = Params::new(&ctx);
7
8     params.set_bool("ignore_solver1", true); // always use solver2 (force incremental
9     // engine)
10    params.set_u32("solver2_timeout", 5000); // 5 second timeout
11
12    // solver2_unknown:
13    //     0: If the check returns "Unknown" with solver2, simply return "Unknown"
14    //     without attempting fallback
15    //     1: Default behavior - If the check returns "Unknown" with solver2, retry
16    //     with solver1 only if the problem is quantifier-free
17    //     2: If the check returns "Unknown" with solver2, always retry with solver1
18    //     regardless of whether the problem contains quantifiers
19    params.set_u32("solver2_unknown", 2);
20    solver.set_params(&params);
21    // continue ... with your logic
22
23    match solver.check() {
24        SatResult::Sat => {
25            println!("Satisfiable");
26        }
27        SatResult::Unsat => {
28            println!("Unsatisfiable");
29        }
30        SatResult::Unknown => {
31            println!("Unknown");
32        }
33    }
34 }
```

You can skip all this and just create a solver from a tactic:

```
1 fn main() {
2     use z3::{Config, Context, Tactic};
3     let ctx = Context::new(&Config::new());
4     let t: Tactic<'_> = Tactic::new(&ctx, "qflia");
5     let _s: z3::Solver<'_> = t.solver();
6     // continue ... with your logic
7 }
```

- **pub fn from_string<T: Into<Vec<u8>>>(&self, source_string: T):** This method parses a raw SMT-LIB2 string and directly adds all resulting assertions, soft constraints, and optimization goals to the current solver instance. Use this when you already have SMT formulas in SMT-LIB2 format (as strings) and want to quickly inject them into a solver. This is especially helpful for testing the behavior of Z3 on known benchmarks or formulas generated by other tools. Note that all variables and declarations inside the string are local to Z3. They are not linked to any `z3::ast` objects in Rust.

```
1 fn main() {
2     use z3::{Config, Context, Solver};
3     let cfg = Config::new();
4     let ctx = Context::new(&cfg);
5     let solver = Solver::new(&ctx);
```

```

6
7     let smt = r#"
8         (declare-const x Int)
9         (declare-const y Int)
10        (assert (> x y))
11        (assert (> y 0))
12        (maximize x)
13    "#;
14
15    solver.from_string(smt);
16    solver.from_string("(declare-sort A)");
17    solver.from_string("(declare-const z A)"); // 'A' is still available
18    solver.from_string("(assert (= z z))"); // 'z' is still available
19
20    println!(":{:?}\n", solver.check());
21    println!(":{:?}\n", solver.get_assertions());
22 }

```

- pub fn new_for_logic<S: Into<Symbol>>(ctx: &Context, logic: S) -> Option<Solver>: This function creates a new solver that is explicitly specialized for a particular SMT-LIB logic, such as "QF_LIA" (quantifier-free linear integer arithmetic), "QF_BV" (quantifier-free bitvectors), or "AUFLIA" (arrays and uninterpreted functions with linear integer arithmetic). If the logic string is invalid or unsupported by Z3, the function returns None. *Why use this?* Using Solver::new_for_logic bypasses Z3's internal logic auto-detection (used in Solver::new). This gives you full control over which logical fragment is being solved, and can result in faster performance. Use new_for_logic() when you're solving formulas that fall into a well-known logical fragment.

```

1 fn main() {
2     use z3::{ast::Ast, ast::Int, Config, Context, SatResult, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new_for_logic(&ctx, "QF_LIA").expect("logic not supported");
7
8     let x = Int::new_const(&ctx, "x");
9     let y = Int::new_const(&ctx, "y");
10    solver.assert(&x._eq(&y));
11    match solver.check() {
12        SatResult::Sat => println!("Satisfiable"),
13        SatResult::Unsat => println!("Unsatisfiable"),
14        SatResult::Unknown => println!("Unknown"),
15    }
16 }

```

Tips:

- The logic string must match exactly what Z3 expects — case sensitive, e.g., "QF_LIA", not "qf_lia".
- This avoids the combined solver behavior (i.e., it does not switch between incremental/non-incremental under the hood). The tactic-based approach and logic-specialized solver uses the *same tactic throughout its lifetime*, without switching between different solving strategies based on usage patterns.
- pub fn translate(&self, dest: &Context) -> Solver: This copies the full solver state (with its assertions) into another Context.

```

1 fn main() {

```

```

2   use z3::{ast::Ast, ast::Int, Config, Context, Solver};
3   let ctx = Context::new(&Config::new());
4   let solver = Solver::new_for_logic(&ctx, "QF_LIA").expect("logic not supported");
5
6   let x = Int::new_const(&ctx, "x");
7   let y = Int::new_const(&ctx, "y");
8   solver.assert(&x.eq(&y));
9
10  let other_ctx = Context::new(&Config::new());
11  let solver2 = solver.translate(&other_ctx);
12
13  println!(":{}?", solver2.get_assertions());
14 }

```

- `pub fn get_context(&self) -> &Context:` This method simply returns a reference to the Context that the solver was created with. Since all Z3 values (sorts, variables, expressions, etc.) are tied to a specific context, this is useful when you need to construct new expressions in the same context that a solver operates in. Sometimes you only have access to the solver (e.g., passed as a parameter), and you want to create a new variable or sort. Instead of passing the context separately, you can just call `get_context()` to retrieve it.

```

1  use z3::ast::Int;
2  use z3::Solver;
3  fn add_constraint_to_solver(solver: &Solver) {
4      let ctx = solver.get_context();
5      let x = Int::new_const(ctx, "x");
6      solver.assert(&x.gt(&Int::from_i64(ctx, 0)));
7  }
8
9  fn main() {
10     let ctx = z3::Context::new(&z3::Config::new());
11     let solver = Solver::new(&ctx);
12
13     add_constraint_to_solver(&solver);
14
15     match solver.check() {
16         z3::SatResult::Sat => println!("Satisfiable"),
17         z3::SatResult::Unsat => println!("Unsatisfiable"),
18         z3::SatResult::Unknown => println!("Unknown"),
19     }
20 }

```

- `pub fn assert(&self, ast: &ast::Bool):` This method adds a hard constraint (i.e., an assertion that must hold) into the solver. All constraints you assert are logically **ANDed together**. Once you have asserted all relevant formulas, you can use `check()` or `check_assumptions()` to determine whether the entire set is satisfiable or not. Each call to `assert()` adds a new formula to the solver's internal goal. All assertions are accumulated unless you `reset()` or use `push()/pop()` for backtracking.

```

1  fn main() {
2      use z3::{ast::Bool, Config, Context, SatResult, Solver};
3
4      let cfg = Config::new();
5      let ctx = Context::new(&cfg);
6      let mut solver = Solver::new(&ctx);
7
8      solver.assert(&Bool::from_bool(&ctx, true)); // TRUE is OK
9      solver += &Bool::from_bool(&ctx, false); // now conflicting
10     solver += Bool::fresh_const(&ctx, "p"); // another unknown prop

```

```

11
12     assert_eq!(solver.check(), SatResult::Unsat); // system is inconsistent
13     solver.get_assertions()
14         .iter()
15         .for_each(|a| println!("assertion: {}", a));
16 }

```

Tips:

- If you want to label assertions for unsat-core extraction, use `assert_and_track()` instead.
- You can retrieve all current assertions with `get_assertions()`.
- There is no way to remove a specific assertion — use `push()/pop()` or `reset()` instead.
- pub fn `assert_and_track(&self, ast: &ast::Bool, p: &ast::Bool)`: This method *asserts* and *labels* the constraint with a Boolean constant *p* for unsat-core extraction. Adds a Boolean constraint to the solver, just like `assert()`, but attaches a fresh `label` so that Z3 can return it later if the formula causes unsatisfiability. This is used to extract **unsat cores**. When Z3 says the problem is **Unsat**, the *unsat core* is the minimal subset of the asserted constraints that is already inconsistent. It helps explain *why* the problem has no solution. *Why use labels?* You can only get the unsat core if you explicitly label your assertions using `assert_and_track`. The returned core will be a list of those labels.

```

1 fn main() {
2     use z3::{ast::Bool, Config, Context, SatResult, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let p = Bool::from_bool(&ctx, true);
9     let q = Bool::from_bool(&ctx, false);
10    let lp = Bool::fresh_const(&ctx, "p_label");
11    let lq = Bool::fresh_const(&ctx, "q_label");
12
13    solver.assert_and_track(&p, &lp);
14    solver.assert_and_track(&q, &lq);
15
16    assert_eq!(solver.check(), SatResult::Unsat);
17
18    // Extract which labels caused the conflict
19    let core = solver.get_unsat_core();
20    for l in core {
21        println!("Conflict came from: {}", l);
22    }
23 }

```

- pub fn `reset(&self)`: `solver.reset()` removes all assertions.
- pub fn `push(&self)`: Push a *backtracking* point. Create backtracking frames for incremental solving. Please refer to the example related to incremental solving.¹²
- pub fn `pop(&self, n: u32)`: Pop *n* *backtracking* levels. Remove backtracking frames for incremental solving.

```

1 solver.push(); solver.assert(&phi1);
2 solver.push(); solver.assert(&phi2);
3 solver.pop(2); // both phi1 and phi2 removed

```

- pub fn check(&self) -> SatResult: This method asks Z3: “Are all the current assertions consistent?” It returns one of three possible results:
 - SatResult::Sat — the formula is satisfiable (there exists a model).
 - SatResult::Unsat — no model exists that satisfies all assertions (there is a proof).
 - SatResult::Unknown — Z3 cannot determine satisfiability (e.g., from incomplete theory support, resource limit, timeout, unsupported logic, etc.).

Z3 takes all formulas added via `assert()` or `assert_and_track()` and tries to solve them as a conjunction. If the formula is satisfiable, you can retrieve the model using `get_model()` (if model generation is enabled in the config). If the formula is unsatisfiable and proof generation was enabled, you can retrieve a proof via `get_proof()`.

Proof Generation:

```

1  fn main() {
2      use z3::{ast::Int, Config, Context, SatResult, Solver};
3
4      let mut cfg = Config::new();
5      cfg.set_proof_generation(true); // Enable proof generation
6      let ctx = Context::new(&cfg);
7      let solver = Solver::new(&ctx);
8
9      let x = Int::new_const(&ctx, "x");
10     let y = Int::new_const(&ctx, "y");
11
12     solver.assert(&x.gt(&y)); // x > y
13     solver.assert(&y.gt(&x)); // y > x (conflict)
14
15     let result = solver.check();
16     assert_eq!(result, SatResult::Unsat);
17     println!("Solver result: {:?}", result);
18     println!("proof: {:?}", solver.get_proof().unwrap());
19 }
```

Model Generation:

```

1  fn main() {
2      use z3::{ast::Int, Config, Context, SatResult, Solver};
3
4      let mut cfg = Config::new();
5      cfg.set_model_generation(true); // Enable model generation
6      let ctx = Context::new(&cfg);
7      let solver = Solver::new(&ctx);
8
9      let x = Int::new_const(&ctx, "x");
10     let y = Int::new_const(&ctx, "y");
11
12     solver.assert(&x.gt(&y)); // x > y
13
14     let result = solver.check();
15     assert_eq!(result, SatResult::Sat);
16     println!("Solver result: {:?}", result);
17     println!("model: {:?}", solver.get_model().unwrap());
18 }
```

What if result is Unknown? This means Z3 couldn’t solve it. After `SatResult::Unknown`, use `get_reason_unknown()` for a human-readable explanation to inspect why. Also, a model may still

be returned on `Unknown` using `get_model()`, but it is not guaranteed to satisfy assertions.

Tips

- Use `get_model()` only if the result is `Sat` or `Unknown`. In case it is `unknown`, `get_model()` will either not produce a model or it will produce a model, under which the satisfiability of the assertions is not guaranteed.
 - Use `get_proof()` only if the result is `Unsat` and proof generation was enabled.
 - Use `check_assumptions()` if you want to solve under temporary literals.
- `pub fn check_assumptions(&self, assumptions: &[ast::Bool]) -> SatResult:` Solve under one-shot assumptions. This method works like `check()`, but temporarily assumes that the given Boolean literals are true **only during this call**. These are not added permanently to the solver — they exist just for the duration of this check. This is useful when you want to:
 - test the satisfiability of your current assertions *under certain assumptions*,
 - selectively enable constraints via assumptions,
 - extract **unsat cores** through trial & error when the problem becomes inconsistent.

```
1 fn main() {
2     use z3::{ast::Bool, Config, Context, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let a = Bool::fresh_const(&ctx, "a");
9     let b = Bool::fresh_const(&ctx, "b");
10
11    solver.assert(&a.not()); // a = false
12    let result = solver.check_assumptions(&[a, b]); // assume a = true (conflict)
13    println!("Result: {:?}", result);
14 }
```

Here, the solver would normally be satisfiable (`a = false`), but we temporarily assumed `a = true`, which contradicts the assertion.

Comparison to `assert_and_track()`:

- `assert_and_track()` permanently adds a constraint and labels it for core tracking.
 - `check_assumptions()` keeps the assumption temporary.
 - You can mix both in the same solver — Z3 will combine them in the core.
- `pub fn get_assertions(&self) -> Vec<Bool>:` Return the current list of asserted formulas. This does *not* include assumptions from `check_assumptions()`.

```
1 fn main() {
2     use z3::{ast::Bool, Config, Context, Params, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let mut p = Params::new(&ctx);
9     p.set_u32("timeout", 10_000);
10    solver.set_params(&p);
```

```

11
12     let a = Bool::fresh_const(&ctx, "a");
13     let b = Bool::fresh_const(&ctx, "b");
14     let x = Bool::fresh_const(&ctx, "x");
15
16     solver.assert(&a.not()); // a = false
17     solver.assert_and_track(&b, &x); // track b
18     let _result = solver.check_assumptions(&[a]); // assume a = true (conflict)
19
20     for i in solver.get_assertions() {
21         println!("Assertion: {}", i);
22     }
23 }
```

The snippet prints the following lines:

```

Assertion: (not a!0)
Assertion: (=> x!2 b!1)
```

- `pub fn get_unsat_core(&self) -> Vec<Bool>`: Returns a subset of assumptions after the last call to `check_assumptions()` or a sequence of calls to `assert_and_track` followed by `check()`. These are the assumptions Z3 used in the unsatisfiability proof.¹¹ This method returns a subset of the *tracked* literals and/or one-shot assumptions that together are already inconsistent. Each element of the vector is a `Bool` constant that you previously supplied as a label in `assert_and_track()`, or as an assumption in `check_assumptions()`.

```

1 fn main() {
2     use z3::{ast::Bool, Config, Context, Solver};
3
4     let cfg = Config::new();
5     let ctx = Context::new(&cfg);
6     let solver = Solver::new(&ctx);
7
8     let a = Bool::fresh_const(&ctx, "a");
9     let b = Bool::fresh_const(&ctx, "b");
10    let x = Bool::fresh_const(&ctx, "x");
11    let y = Bool::fresh_const(&ctx, "y");
12    let z = Bool::fresh_const(&ctx, "z");
13
14    solver.assert_and_track(&a.not(), &x); // a = false
15    solver.assert_and_track(&b.not(), &y); // b = false
16    solver.assert_and_track(&Bool::or(&ctx, &[&a, &b]), &z); // a or b must be true
17
18    let c = solver.check();
19    println!("Check result: {:?}", c);
20
21    let unsat_core = solver.get_unsat_core();
22    println!("Unsat core: {:?}", unsat_core);
23 }
```

A typical run prints:

```

Check result : Unsat
Unsat core   : [x!2, y!3, z!4]
```

How it works. Z3 proves the contradiction internally and then asks which of the *labels* (`x`, `y`, `z`) it needed for the proof. Those labels are returned as the unsatisfiable core. If you omit a label (e.g.

¹¹By default, unsat cores are not minimized. To enable minimization, set `sat.core.minimize` or `smt.core.minimize` to `true` (for SAT or SMT respectively). This increases solving cost.

you used plain `assert()`), that clause can never appear in the core—even if it is the actual cause of the inconsistency.

Tips

- Always call `check()` before querying the core; otherwise the vector is empty.
- Use either `assert_and_track()` or `check_assumptions()` to create *every* constraint you want to diagnose; un-labelled `assert()`s are invisible to the core.
- By default, Z3 does not guarantee that the unsat core it returns is unique (or even minimal).
- pub fn `get_consequences(&self, assumptions: &[ast::Bool], variables: &[ast::Bool],) -> Vec<ast::Bool>:` Under given assumptions, return all literals about `vars` that are logically implied. With the current `assertions` plus a temporary list of Boolean `assumptions`, compute every literal over the given propositional `vars` that is *logically implied*. Each literal is returned as a `Bool` constant, possibly negated internally (e.g. $\neg b$).

```

1  fn main() {
2      use z3::{ast::Bool, Config, Context, Solver};
3
4      let cfg     = Config::new();
5      let ctx     = Context::new(&cfg);
6      let solver = Solver::new(&ctx);
7
8      let a = Bool::fresh_const(&ctx, "a");
9      let b = Bool::fresh_const(&ctx, "b");
10     let c = Bool::fresh_const(&ctx, "c");
11
12     solver.assert(&a.implies(&b));    // a -> b
13     solver.assert(&b.implies(&c));    // b -> c
14
15     let consequences = solver.get_consequences(&[a.clone()], &[b.clone(), c.clone()]);
16
17     println!("Consequences under a = true:");
18     for lit in consequences {
19         println!("  {}", lit);        // prints: b, c
20     }
21 }
```

Internally Z3 searches for all literals over the provided variables that follow from `assertions` \wedge `assumptions`. The result is *complete*: any Boolean combination over `vars` that is entailed will be reported in the output.

Tips

- Variables and assumptions must be `Bool constants`, not arbitrary formulas.
- The solver state is *not* modified; you can call `get_consequences()` repeatedly with different sets.
- pub fn `get_model(&self) -> Option<Model>:` Obtain a satisfying assignment when the last `Solver::check()` or `Solver::check_assumptions()` returned `Sat` or `Unknown`. Nevertheless, in the case of `Unknown`, Z3 does not ensure that calls to `get_model()` succeed and any models produced in this case are not guaranteed to satisfy the assertions. Note that on the configuration level, `model generation` is enabled by default and calling `cfg.set_model_generation(false)` is intended to override this.¹² However, if the model generation is set to false using `cfg.set_model_generation(false)`, when the solver is built (`Solver::new(&ctx)`), the model

¹²Thus, there is no need to set `cfg.set_model_generation(true)` explicitly for model generation.

generation value resets to `true` for the corresponding solver. The `configuration` setting therefore does not have a lasting effect on the models, while it remains in force for the proofs on the `solver` level. This is because models are considered cheap and very useful and proofs carry a significant memory/time cost. `get_model()` invokes an error handler if `check()` / `check_assumptions()` have not been called, or if the last result was `Unsat`.

```

1 fn main() {
2     use z3::{ast::Bool, Params, Config, Context, Solver};
3
4     let mut cfg = Config::new();
5     // cfg.set_model_generation(false);
6     let ctx = Context::new(&cfg);
7     let solver = Solver::new(&ctx);
8
9     // let mut p = Params::new(&ctx);
10    // p.set_bool("model", false); // solver-local override
11    // solver.set_params(&p);
12
13    let a = Bool::fresh_const(&ctx, "a");
14    solver.assert(&a.not()); // a = false
15
16    match solver.check() {
17        z3::SatResult::Sat => {
18            println!("SAT");
19            let model = solver.get_model().unwrap();
20            println!("Model: {}", model);
21        }
22        z3::SatResult::Unsat => {
23            println!("UNSAT");
24            let proof = solver.get_proof().unwrap();
25            println!("Proof: {:?}", proof);
26        }
27        z3::SatResult::Unknown => {
28            println!("UNKNOWN");
29        }
30    }
31 }
```

The above code returns the following:

```
SAT
Model: a!0 -> false
```

If we un-comment line 5 or lines 9-11 - or both - the output will be as following:

```
SAT
Model:
```

This is what we mean by *"if model generation is disabled at either the configuration level or the solver level, then the table of constant and function interpretations will not be populated."* Lastly the following is an example where the solver parameter overrides the configuration parameter:

```

1 fn main() {
2     use z3::{ast::Bool, Params, Config, Context, Solver};
3
4     let mut cfg = Config::new();
5     cfg.set_model_generation(false);
6     let ctx = Context::new(&cfg);
7     let solver = Solver::new(&ctx);
8
9     let mut p = Params::new(&ctx);
```

```

10 p.set_bool("model", true); // solver-local override
11 solver.set_params(&p);
12
13 let a = Bool::fresh_const(&ctx, "a");
14 solver.assert(&a.not()); // a = false
15
16 match solver.check() {
17     z3::SatResult::Sat => {
18         println!("SAT");
19         let model = solver.get_model().unwrap();
20         println!("Model: {}", model);
21     }
22     z3::SatResult::Unsat => {
23         println!("UNSAT");
24         let proof = solver.get_proof().unwrap();
25         println!("Proof: {:?}", proof);
26     }
27     z3::SatResult::Unknown => {
28         println!("UNKNOWN");
29     }
30 }
31 }

```

The above code returns the following:

```

SAT
Model: a!0 -> false

```

- pub fn get_proof(&self) -> Option<impl Ast>: Get a proof object for the last Solver::check() or Solver::check_assumptions() when the result was Unsat and proofs were enabled using cfg.set_proof_generation(true). Note that proof generation is disabled by default at the configuration level, and the solver always inherits this setting. Unlike model generation, the proof option cannot be enabled at the solver level—it must be set to true in the configuration before creating the context. get_proof() invokes the error handler if proof generation was not enabled using cfg.set_proof_generation(true), or if check() / check_assumptions() have not been called, or if the last result was not Unsat.

```

1 fn main() {
2     use z3::{ast::Bool, Config, Context, Solver};
3
4     let mut cfg = Config::new();
5     cfg.set_proof_generation(true); // commenting this line will disable proof generation
6     // → and will give an error if you try to get a proof
7
8     let ctx = Context::new(&cfg);
9     let solver = Solver::new(&ctx);
10
11    let a = Bool::fresh_const(&ctx, "a");
12
13    solver.assert(&a.not()); // a = false
14    solver.assert(&a); // a = true
15
16    match solver.check() {
17        z3::SatResult::Sat => {
18            println!("SAT");
19            let model = solver.get_model().unwrap();
20            println!("Model: {}", model);
21        }
22        z3::SatResult::Unsat => {

```

```

22     println!("UNSAT");
23     let proof = solver.get_proof().unwrap();
24     println!("Proof: {:_?}" , proof);
25   }
26   z3::SatResult::Unknown => {
27     println!("UNKNOWN");
28     let unknown = solver.get_reason_unknown().unwrap();
29     println!("Reason unknown: {}", unknown);
30   }
31 }
32 }
```

- pub fn `get_reason_unknown(&self) -> Option<String>`: Brief textual reason when the solver answered Unknown for the most recent `check()/check_assumptions()`.
- pub fn `set_params(&self, params: &Params)`: Update solver-level parameters. Note that when the parameter is set on the configuration level it affects all the child solvers. Also note that some parameters like setting the proof generation are *configuration-level* and must be done before creating the context.
- pub fn `get_statistics(&self) -> Statistics`: Returns statistics from the most recent `check`.
- pub fn `to_smt2(&self) -> String`: Dump the current solver state as SMT-LIB2 text for inspection or replay.
- Trait impls
 - `AddAssign<&Bool>` and `AddAssign<Bool>`: Enables the use of the shorthand notation `solver += &constraint` for `solver.assert(&constraint)`.

13 *rec_func_decl* module

This module provides the associated functions for the `RecFuncDecl` data type defined in the `lib.rs` module.

- pub fn `new<S: Into<Symbol>>(&ctx: &Context, name: S, domain: [&Sort], range: &Sort,) -> Self`: This function creates a fresh `RecFuncDecl` in `ctx` with the given `name`, argument `domain`, and `range` sort.
- ```

1 let mut f = RecFuncDecl::new(
2 &ctx,
3 "fact", // function name
4 [&Sort::int(&ctx)], // domain: a single Int
5 &Sort::int(&ctx) // range: Int
6);
7 // Now `f` exists in Z3; no body yet.
8 // Next step: call `add_def()` to give it a potentially recursive equation.
```
- pub fn `add_def(&self, args: [&dyn Ast::Ast], body: &dyn Ast)`: Adds the body to a recursive function. Note that `args` should have the types corresponding to the `domain` of the `RecFuncDecl`. For more information please refer to section 6, specifically to the part devoted to recursive functions.

## 14 *probe* module

A probe is a tiny function that inspects a given Goal (i.e. a set of formulas) and returns a double (a 64-bit float). A probe never mutates a goal; feel free to call it multiple times. They are indispensable for writing *tactics that adapt to the structure of a goal at runtime*. If you think of a Goal as *some formulas I want to solve*, a probe asks a question like, *How many OR nodes are in this goal?* or *Are all top-level assertions pseudo-Boolean constraints?*. Z3 by default has dozens of built-in probes. When a probe returns a non-zero double, we treat that as “true”; when it returns 0.0, we treat it as “false.” But a probe

can also return arbitrary positive reals (e.g. the actual count of something).

- pub fn list\_all(ctx:&Context) -> impl Iterator<Item = Result<&str, Utf8Error>>: Iterate over every probe name compiled into the current Z3 build.

```

1 fn main() {
2 use z3::Config, Context, Probe;
3
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let probes: Vec<_> = Probe::list_all(&ctx).filter_map(|r| r.ok()).collect();
7 for p in probes {
8 println!("Probe: {}", p);
9 }
10 }
```

- pub fn describe(ctx: &Context, name: &str) -> Result<&str, Utf8Error>: Fetch the human-readable documentation string attached to a probe.

```

1 fn main() {
2 use z3::Config, Context, Probe;
3
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 for p in Probe::list_all(&ctx).filter_map(|r| r.ok()) {
7 match Probe::describe(&ctx, p) {
8 Ok(desc) => println!("{}: {}", p, desc),
9 Err(e) => println!("{}: <bad UTF-8 or invalid name>", p),
10 }
11 }
12 }
```

The code prints the following:

```

1 ackr-bound-probe: A probe to give an upper bound of Ackermann congruence lemmas that a
 ↳ formula might generate.
2 is-unbounded: true if the goal contains integer/real constants that do not have
 ↳ lower/upper bounds.
3 is-pb: true if the goal is a pseudo-boolean problem.
4 arith-max-deg: max polynomial total degree of an arithmetic atom.
5 arith-avg-deg: avg polynomial total degree of an arithmetic atom.
6 arith-max-bw: max coefficient bit width.
7 arith-avg-bw: avg coefficient bit width.
8 is-qflia: true if the goal is in QF_LIA.
9 is-qfauflia: true if the goal is in QF_AUFLIA.
10 is-qflra: true if the goal is in QF_LRA.
11 is-qflira: true if the goal is in QF_LIRA.
12 is-ilp: true if the goal is ILP.
13 is-qfnia: true if the goal is in QF_NIA (quantifier-free nonlinear integer arithmetic).
14 is-qfnra: true if the goal is in QF_NRA (quantifier-free nonlinear real arithmetic).
15 is-nia: true if the goal is in NIA (nonlinear integer arithmetic, formula may have
 ↳ quantifiers).
16 is-nra: true if the goal is in NRA (nonlinear real arithmetic, formula may have
 ↳ quantifiers).
17 is-nira: true if the goal is in NIRA (nonlinear integer and real arithmetic, formula may
 ↳ have quantifiers).
18 is-lia: true if the goal is in LIA (linear integer arithmetic, formula may have
 ↳ quantifiers).
19 is-lra: true if the goal is in LRA (linear real arithmetic, formula may have
 ↳ quantifiers).
```

```

20 is-lira: true if the goal is in LIRA (linear integer and real arithmetic, formula may
 ↳ have quantifiers).
21 is-qfufnra: true if the goal is QF_UFNRA (quantifier-free nonlinear real arithmetic with
 ↳ other theories).
22 is-qfbv-eq: true if the goal is in a fragment of QF_BV which uses only =, extract,
 ↳ concat.
23 is-qffp: true if the goal is in QF_FP (floats).
24 is-qffpbv: true if the goal is in QF_FPBV (floats+bit-vectors).
25 is-qffplra: true if the goal is in QF_FPLRA.
26 memory: amount of used memory in megabytes.
27 depth: depth of the input goal.
28 size: number of assertions in the given goal.
29 num-exprs: number of expressions/terms in the given goal.
30 num-consts: number of non Boolean constants in the given goal.
31 num-bool-consts: number of Boolean constants in the given goal.
32 num-arith-consts: number of arithmetic constants in the given goal.
33 num-bv-consts: number of bit-vector constants in the given goal.
34 produce-proofs: true if proof generation is enabled for the given goal.
35 produce-model: true if model generation is enabled for the given goal.
36 produce-unsat-cores: true if unsat-core generation is enabled for the given goal.
37 has-quantifiers: true if the goal contains quantifiers.
38 has-patterns: true if the goal contains quantifiers with patterns.
39 is-propositional: true if the goal is in propositional logic.
40 is-qfbv: true if the goal is in QF_BV.
41 is-qfaufbv: true if the goal is in QF_AUFBV.
42 is-quasi-pb: true if the goal is quasi-pb.

```

- pub fn new(ctx: &Context, name: &str) -> Probe: Lookup an existing probe by its literal name. Call Probe::new only with names obtained from Probe::list\_all(&ctx) to avoid invalid-name panics. This does not create a new probe; it merely returns a handle to the built-in probe.
- pub fn apply(&self, goal: &Goal) -> f64: Execute the probe on a given Goal (a set of Z3 formulas). Use apply when you want to inspect a Goal and obtain its probe-based property. It returns a f64 result:
  - If the probe is Boolean (e.g. "is-pb"), nonzero means true, zero means false.
  - If the probe returns a numeric measure (e.g. "num-clauses"), that integer (as a float) is returned.

## Tips

- Combine with Tactic::cond to branch a tactic pipeline.
- pub fn constant(ctx: &Context, val: f64) -> Probe: Create a probe that always evaluates to the given floating-point val. Useful as “literals” in probe arithmetic.

## Tips

- Use constant to build composite probes in combination with relational or logical combinators (e.g. le, and, etc.), for example:

```

1 let small_probe = Probe::new(&ctx, "num-clauses")
2 .le(&Probe::constant(&ctx, 10.0));

```

- Create constant true and false probes. For example:

```

1 let always_true = Probe::constant(&ctx, 1.0);
2 let always_false = Probe::constant(&ctx, 0.0);

```

- **Relational probes:** `lt`, `gt`, `le`, `ge`, `eq`, `ne`: Produce composite probes that compare the numeric results of two probes. All relational operators return a *Boolean* probe; i.e., any non-zero value is interpreted as *true* and a zero value is interpreted as *false*.

### Signatures

```

1 pub fn lt(&self, p: Probe) -> Probe
2 pub fn gt(&self, p: &Probe) -> Probe
3 pub fn le(&self, p: &Probe) -> Probe
4 pub fn ge(&self, p: &Probe) -> Probe
5 pub fn eq(&self, p: &Probe) -> Probe
6 pub fn ne(&self, p: &Probe) -> Probe

```

- **Logical combinators:** `and`, `or`, `not`: Boolean connectives lifted to probes with the conventional semantics.

### Signatures

```

1 pub fn and(&self, p: &Probe) -> Probe
2 pub fn or (&self, p: &Probe) -> Probe
3 pub fn not(&self) -> Probe

```

### Example

```

1 use z3::ast::Int;
2
3 fn main() {
4 use z3::{ast::Ast, Config, Context, Goal, Probe, Tactic};
5
6 let cfg = Config::new();
7 let ctx = Context::new(&cfg);
8
9 // x + y = 10 and x > 0
10 let x = Int::new_const(&ctx, "x");
11 let y = Int::new_const(&ctx, "y");
12 let goal = Goal::new(&ctx, true, false, false);
13 goal.assert(&Int::add(&ctx, &[&x, &y])._eq(&Int::from_i64(&ctx, 10)));
14 goal.assert(&x.gt(&Int::from_i64(&ctx, 0)));
15
16 // check if the goal has <= 2 assertions
17 let size = Probe::new(&ctx, "size");
18 let small = size.le(&Probe::constant(&ctx, 2.0));
19
20 let heavy = Tactic::new(&ctx, "ctx-solver-simplify");
21 let skip = Tactic::create_skip(&ctx);
22 let pipeline = Tactic::cond(&ctx, &small, &heavy, &skip);
23
24 let result = pipeline.apply(&goal, None).unwrap();
25 for (i, sub) in result.list_subgoals().enumerate() {
26 println!("Sub-goal {}: {}", i, sub);
27 }
28}

```

## 15 pattern module

In Z3, a *pattern* (also called a *trigger*) is a syntactic cue that tells the instantiation engine *when* to fire a quantified formula. If a pattern is not provided for a quantifier, then Z3 will automatically compute a set of patterns for it. However, for optimal performance, the user should provide the patterns. Whenever

Z3 encounters a ground term that **matches** the pattern, it instantiates the quantifier with the matching substitution. Good patterns accelerate solving; bad patterns can do nothing or explode the search space.

- **Single-pattern** = a non-empty list with one term. A single match of that term triggers instantiation.
- **Multi-pattern** = a list with  $n > 1$  terms. *All* terms must match simultaneously for the trigger to fire.
- `pub fn new(ctx: &Context, terms: &[&dyn Ast]) -> Pattern`: The returned `Pattern` can then be supplied to `ast::forall_const` or `ast::exists_const`. Note that `terms` must be non-empty.<sup>13</sup>

## Tips

- Use constructor-shaped terms rather than arbitrary arithmetic to avoid unification loops.
- **Coverage:** Every bound variable that actually occurs in the body *must* also occur in at least one pattern—otherwise Z3 rejects the quantifier as ill-formed.

### Example 1 — single-pattern trigger

```
1 use z3::{
2 ast::{forall_const, Ast, Int},
3 Config, Context, Pattern,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9
10 let x = Int::new_const(&ctx, "x");
11 let zero = Int::from_i64(&ctx, 0);
12 let lhs = Int::add(&ctx, &[&x, &zero]); // x + 0
13 let pat = Pattern::new(&ctx, &[&lhs]); // single-pattern
14
15 let forall = forall_const(
16 &ctx,
17 &[&x],
18 &[&pat],
19 &lhs._eq(&x), // body: x + 0 = x
20);
21
22 println!("{}", forall);
23 }
```

**When to use.** Whenever an axiom has a “canonical” left-hand side (here  $x + 0$ ), firing on that term is usually enough to prove equalities quickly.

### Example 2 — multi-pattern trigger

```
1 use z3::{
2 ast::{forall_const, Ast, Int},
3 Config, Context, FuncDecl, Pattern, Sort,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
```

<sup>13</sup>If empty, the `assertion failed: !terms.is_empty()` error will be invoked.

```

9
10 let int = Sort::int(&ctx);
11 let f = FuncDecl::new(&ctx, "f", [&int], &int);
12
13 // f(x) = f(y) -> x = y
14 let x = Int::new_const(&ctx, "x");
15 let y = Int::new_const(&ctx, "y");
16 let fx = f.apply(&[&x]).as_int().unwrap();
17 let fy = f.apply(&[&y]).as_int().unwrap();
18
19 // multi-pattern ⟨f(x), f(y)⟩
20 let pat = Pattern::new(&ctx, [&fx, &fy]);
21
22 let lemma = forall_const(&ctx, [&x, &y], [&pat], &fx._eq(&fy).implies(&x._eq(&y)));
23 println!("{}", lemma);
24 }

```

**Why multi-pattern?** We only want the lemma fired when two *distinct* applications of  $f$  appear together; the conjunction  $\langle f(x), f(y) \rangle$  prevents useless self-matches like  $f(t)$  with itself. In general, a pattern that matches too often (e.g. the variable alone) can blow up search exponentially - *over-general patterns*. Multi-patterns tame this by requiring simultaneous matches. But be aware that if your pattern rarely matches, the quantifier never fires and proofs fail - *over-specific patterns*.

## 16 params module

This module provides the related functions for the `Params` type defined in the `lib.rs` module. A `Params` object is essentially a map from parameter names (`Symbol`) to values (`Symbol`, `bool`, `f64`, or `u32`). The parameters in a `Params` object can then be passed to tactics or solvers to tune Z3's behavior.

- pub fn `new(ctx: &Context) -> Params`: Construct an empty parameter set.
- pub fn `set_symbol<K: Into<Symbol>, V: Into<Symbol>>(&mut self, k: K, v: V)`: Associate the parameter named `k` with the symbol value `v`.
- pub fn `set_bool<K: Into<Symbol>>(&mut self, k: K, v: bool)`: Associate the parameter named `k` with the Boolean value `v`.
- pub fn `set_f64<K: Into<Symbol>>(&mut self, k: K, v: f64)`: Associate the parameter named `k` with the double value `v`.
- pub fn `set_u32<K: Into<Symbol>>(&mut self, k: K, v: u32)`: Associate the parameter named `k` with the unsigned integer value `v`.

```

1 use z3::{Config, Context, Solver, Params, ast::Int, SatResult};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6
7 let mut params = Params::new(&ctx);
8 params.set_u32("timeout", 1000);
9 params.set_bool("auto_config", false);
10 params.set_f64("relevancy", 0.7);
11 params.set_u32("sat.random_seed", 42);
12
13 let solver = Solver::new(&ctx);
14 solver.set_params(¶ms);
15
16 // 10 < x < 20
17 let x = Int::new_const(&ctx, "x");

```

```

18 solver.assert(&x.gt(&Int::from_i64(&ctx, 10)));
19 solver.assert(&x.lt(&Int::from_i64(&ctx, 20)));
20
21 match solver.check() {
22 SatResult::Sat => {
23 println!("SAT");
24 println!("Model: {}", solver.get_model().unwrap());
25 }
26 SatResult::Unsat => println!("UNSAT"),
27 SatResult::Unknown => println!("UNKNOWN"),
28 }
29 }
```

- pub fn `get_global_param(k: &str) -> Option<String>`: Retrieve the current value of the global or module parameter named k. Returns `Some(value)` if the parameter exists, else `None`.
- pub fn `set_global_param(k: &str, v: &str)`: Set a global or module-level parameter k to the string value v across all Z3 contexts.
- pub fn `reset_all_global_params()`: Restore all global and module parameters to their default values. Does not affect already created contexts or objects.

```

1 use z3::{
2 Config, Context, Params, Solver, {get_global_param, reset_all_global_params,
3 set_global_param},
4 };
5
5 fn main() {
6 if let Some(old) = get_global_param("timeout") {
7 println!("Global timeout before: {}", old);
8 } else {
9 println!("Global timeout not set yet.");
10 }
11
12 set_global_param("timeout", "500");
13 println!(
14 "Global timeout now: {}",
15 get_global_param("timeout").unwrap()
16);
17
18 reset_all_global_params();
19 println!(
20 "After reset, global timeout: {:?}",
21 get_global_param("timeout")
22);
23
24 let cfg = Config::new();
25 let ctx = Context::new(&cfg);
26 let solver = Solver::new(&ctx);
27
28 let mut local_params = Params::new(&ctx);
29 local_params.set_u32("timeout", 250);
30 local_params.set_bool("auto_config", false);
31 println!(
32 "global timeout: {:?}",
33 get_global_param("timeout")
34);
35 solver.set_params(&local_params);
36
37 let a = z3::ast::Bool::fresh_const(&ctx, "a");
```

```

38 solver.assert(&a);
39 assert_eq!(solver.check(), z3::SatResult::Sat);
40 }

```

## 17 context module

The `Context` is the foundational object in Z3's Rust API. It encapsulates the Z3 logical context (`Z3_context`) and manages the lifetime of all AST nodes, solvers, and other Z3 objects created within it. Note that objects from different contexts cannot be mixed.

- `pub fn new(cfg: &Config) -> Context`: Create a new Z3 context with the provided configuration.
- `pub fn get_z3_context(&self) -> Z3_context`: Expose the raw C `Z3_context` handle for low-level operations or interfacing with Z3's C API directly.
- `pub fn interrupt(&self)`: Interrupt a solver performing a satisfiability test, a tactic processing a goal, or simplify functions. This method delegates to `ContextHandle::interrupt()`.
- `pub fn handle(&self) -> ContextHandle`: Obtain a handle that can be used to interrupt computation from another thread. The returned `ContextHandle` is both `Send` and `Sync`, allowing safe cross-thread interruption.
- `pub fn update_param_value(&mut self, k: &str, v: &str)`: Update a context-local parameter `k` to the string value `v`. This affects only the current context.
- `pub fn update_bool_param_value(&mut self, k: &str, v: bool)`: Helper function to update a parameter `k` with the Boolean value `v`.

```

1 use std::thread;
2 use std::time::Duration;
3 use z3::{Config, Context, SatResult, Solver};
4
5 fn main() {
6 let cfg = Config::new();
7 let mut ctx = Context::new(&cfg);
8 ctx.update_param_value("timeout", "5000");
9 ctx.update_bool_param_value("auto_config", false);
10
11 let raw_ctx = ctx.get_z3_context();
12 println!("Raw context handle: {:p}", raw_ctx);
13
14 let solver = Solver::new(&ctx);
15 let a = z3::ast::Bool::fresh_const(&ctx, "a");
16 solver.assert(&a);
17
18 let handle = ctx.handle();
19 thread::scope(|s| {
20 s.spawn(|| {
21 std::thread::sleep(Duration::from_millis(0));
22 handle.interrupt();
23 println!("Computation interrupted!");
24 });
25
26 // do the actual check
27 match solver.check() {
28 SatResult::Sat => println!("SAT"),
29 SatResult::Unsat => println!("UNSAT"),
30 SatResult::Unknown => println!("UNKNOWN or interrupted"),
31 }
32 });
33
34 // do the actual check
35 match solver.check() {
36 SatResult::Sat => println!("SAT"),
37 SatResult::Unsat => println!("UNSAT"),
38 SatResult::Unknown => println!("UNKNOWN or interrupted"),
39 }
40 }

```

```

32 });
33 }

```

- **pub fn interrupt(&self):** This method for `ContextHandle`, interrupts any ongoing computation in the associated context. This method is thread-safe.

## 18 *optimize* module

The `Optimize` API turns Z3 from a mere satisfiability checker into a *Max-SMT / Optimization engine*. In other words, instead of asking merely *is the formula satisfiable?*, an optimizer lets you add *objectives* to minimize, maximize, or soft-satisfy. Then the optimizer returns a model that meets the hard constraints **and** is optimal w.r.t. those objectives. Accordingly, *soft constraints* receive rational weights so you can trade them off, and multiple objectives are solved lexicographically by default; i.e., earlier objectives are given a higher priority.<sup>14</sup>

- **impl Optimize:**
  - **pub fn new(ctx: &Context) -> Optimize:** Construct an empty optimizer attached to `ctx`. Equivalent to `sat/smt` solvers but with optimization support.
  - **pub fn from\_string(&self, smtlib2: impl Into<Vec<u8>>):** Parse an SMT-LIB2 string and add the parsed constraints and objectives to the optimizer. This method parses a raw SMT-LIB2 string and directly adds all resulting assertions, soft constraints, and optimization goals to the current optimizer instance. Use this when you already have SMT formulas in SMT-LIB2 format (as strings) and want to quickly inject them into an optimizer. This is especially helpful for testing the behavior of Z3 on known benchmarks or formulas generated by other tools. Note that all variables and declarations inside the string are local to Z3. They are not linked to any `z3::ast` objects in Rust.
  - **pub fn get\_context(&self) -> &Context:** This method simply returns a reference to the `Context` that the optimizer was created with. Since all Z3 values (sorts, variables, expressions, etc.) are tied to a specific context, this is useful when you need to construct new expressions in the same context that a optimizer operates in. Sometimes you only have access to the optimizer (e.g., passed as a parameter), and you want to create a new variable or sort. Instead of passing the context separately, you can just call `get_context()` to retrieve it.
  - **pub fn assert(&self, f: &impl Ast):** Add a *hard* constraint. Violating it makes the problem `Unsat`.
  - **pub fn assert\_and\_track(&self, f: &Bool, p: &Bool):** Like `assert` but annotates the clause with a fresh propositional variable `p` so you can later extract an unsat core.
  - **pub fn assert\_soft(&self, f: &impl Ast, weight: impl Weight, group: Option<Symbol>):** Add a *soft* constraint with weight. Weight is a positive, rational penalty for violating the constraint. On `check`, Z3 minimises the *total penalty* of violated soft clauses.
  - **pub fn maximize(&self, term: &impl Ast) / minimize:** Optimise a numeric term (maximize or minimize). Only `Int`, `Real`, or `BitVec` are allowed to be optimized; Booleans must be turned into 0/1 integers.
  - **pub fn push(&self) / pop(&self):** Create and roll back back-tracking scopes, exactly like in `Solver`. Note that the number of calls to `pop` cannot exceed the number of calls to `push`.
  - **pub fn check(&self, assumptions: &[Bool]) -> SatResult:** Checks for satisfiability and produces optimal values.

---

<sup>14</sup>Many of the methods and their respective descriptions for the `Optimize` data structure are similar to `Solver` as they essentially provide the same interface for asserting formulas, backtracking, and checking satisfiability — with the addition of optimization-specific features.

- pub fn `get_model(&self) -> Option<Model>`: Retrieve the model from the last successful optimize check so you can query optimal values. The error handler is invoked if a model is not available because the `Optimize::check()` command was not invoked for the given optimization solver, or if the result of `Optimize::check()` was `SatResult::Unsat`.
  - pub fn `get_objectives(&self) -> Vec<Dynamic>`: Retrieve the objectives for the last `Optimize::check()`. This contains maximize/minimize objectives and grouped soft constraints.
  - pub fn `get_unsat_core(&self) -> Vec<Bool>`: Returns a subset of assumptions after the last call to `check()` or a sequence of calls to `assert_and_track` followed by `check()`. These are the assumptions Z3 used in the unsatisfiability proof.<sup>15</sup> This method returns a *subset* of the *tracked* literals and/or one-shot assumptions that together are already inconsistent. Each element of the vector is a `Bool` constant that you previously supplied as a label in `assert_and_track()`.
  - pub fn `get_reason_unknown(&self) -> Option<String>`: Explain why Z3 gave up when it returns `Unknown`.
  - pub fn `set_params(&self, params: &Params)`: Tune the optimiser just like any solver.
  - pub fn `get_statistics(&self) -> Statistics`: Retrieve the statistics for the last `Optimize::check()` just like any solver.
- **Weight:** The trait is sealed; you cannot implement it for other types. It is only implemented for `u8`, `u16`, `u32`, `u64`, `u128`, `usize`, `i8`, `i16`, `i32`, `i64`, `i128`, `isize`, `BigInt`, `BigUint`, `BigRational`. In other words, these types can act as the weight in soft constraints.

## Example 1 – Lexicographic minimisation

```

1 use z3::{ast::Int, Config, Context, Optimize, SatResult};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let opt = Optimize::new(&ctx);
7
8 let x = Int::new_const(&ctx, "x");
9 let y = Int::new_const(&ctx, "y");
10
11 // Hard constraints
12 opt.assert(&x.ge(&Int::from_i64(&ctx, 0))); // x >= 0
13 opt.assert(&y.ge(&Int::from_i64(&ctx, 0))); // y >= 0
14 opt.assert(&Int::add(&ctx, &[&x, &y]).le(&Int::from_i64(&ctx, 10))); // x + y <= 10
15
16 // Objectives (lexicographic)
17 opt.minimize(&x); // 1st priority: minimise x
18 opt.maximize(&y); // 2nd priority: then maximise y
19
20 match opt.check(&[]) {
21 SatResult::Sat => {
22 let model = opt.get_model().unwrap();
23 println!(
24 "x = {}, y = {}", // prints: x = 0, y = 10
25 model.eval(&x, true).unwrap(),
26 model.eval(&y, true).unwrap()
27);
28 }
29 }
30}

```

<sup>15</sup>By default, unsat cores are not minimized. To enable minimization, set `sat.core.minimize` or `smt.core.minimize` to `true` (for SAT or SMT respectively). This increases solving cost.

```

28 }
29 SatResult::Unsat => println!("unsat!"),
30 SatResult::Unknown => println!("unknown! {:?}", opt.get_reason_unknown()),
31 }
32 }
```

## Example 2 – Max-SAT via soft constraints

```

1 use z3::{ast::Bool, Config, Context, Optimize, SatResult};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let opt = Optimize::new(&ctx);
7
8 let a = Bool::new_const(&ctx, "a");
9 let b = Bool::new_const(&ctx, "b");
10 let c = Bool::new_const(&ctx, "c");
11
12 // Hard constraint: at most one can be true
13 opt.assert(&a.implies(&Bool::and(&ctx, &[&b.not(), &c.not()])));
14 opt.assert(&b.implies(&Bool::and(&ctx, &[&a.not(), &c.not()])));
15 opt.assert(&c.implies(&Bool::and(&ctx, &[&a.not(), &b.not()])));
16
17 // Soft constraints with different weights
18 opt.assert_soft(&a, 5, None); // worth 5 points
19 opt.assert_soft(&b, (1, 2), None); // worth 0.5 points
20 opt.assert_soft(&c, 1u32, None); // worth 1 point
21
22 // Solve
23 assert_eq!(opt.check(&[]), SatResult::Sat);
24 println!("violated weight sum = {:?}", opt.get_objectives());
25 }
```

*Why use soft constraints?* If the hard constraints cannot all be met simultaneously, Z3 drops a subset whose *total weight is minimal*.

## Example 3 – Extracting an unsat core

```

1 use z3::{ast::Bool, Config, Context, Optimize, SatResult};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let opt = Optimize::new(&ctx);
7
8 let p = Bool::new_const(&ctx, "p");
9 let q = Bool::new_const(&ctx, "q");
10
11 // Track each hard clause with a literal
12 opt.assert_and_track(&p, &p); // (p) tagged by p
13 opt.assert_and_track(&q, &q); // (q) tagged by q
14 opt.assert_and_track(
15 &Bool::and(&ctx, &[&p, &q]).not(),
16 &Bool::new_const(&ctx, "r"),
17);
18
19 if let SatResult::Unsat = opt.check(&[]) {
20 for culprit in opt.get_unsat_core() {
21 println!("in core: {}", culprit);
22 }
23 }
```

```

23 }
24 }
```

The optimiser returns a subset of tracked literals that is itself inconsistent, allowing you to pinpoint which assumptions caused `Unsat`.<sup>16</sup>

## 19 ops module

This module enables operator overloading for Z3’s Rust AST types—specifically for `BV`, `Int`, `Real`, `Float`, and `Bool`—by implementing the corresponding traits from `std::ops`. The macros defined in this module automate trait implementations for arithmetic and logic operations, making Z3 expressions ergonomic and Rust-like. Concretely, it means you can write expressions such as `let expr = (&x + 5u64) * &y - 1i64;` instead of the more verbose `Int::add(...)` calls. Even references (`&Int`) and mixed forms (`Int + &Int`) are handled automatically. A user can simply *import* the desired sorts and rely on the familiar syntax—from that point on, every operator compiles down to the correct Z3 primitive with no extra work.<sup>17</sup>

```

1 use z3::{
2 ast::{Bool, Int, BV},
3 Config, Context,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9
10 let x = Int::new_const(&ctx, "x");
11 let y = Int::new_const(&ctx, "y");
12 let sum = &x + 5u64 + &y; // Int + u64 + &Int
13 let diff = 10i64 - &x; // i64 - &Int
14
15 // Bit-vector
16 let a = BV::new_const(&ctx, "a", 32);
17 let expr = (&a << 3u64) ^ 0xffu64; // (a << 3) XOR 0xFF
18
19 // Boolean
20 let p = Bool::new_const(&ctx, "p");
21 let q = Bool::new_const(&ctx, "q");
22 let formula = !p & (q ^ true); // NOT p AND (q XOR true)
23
24 println!("sum : {sum}");
25 println!("diff : {diff}");
26 println!("expr : {expr}");
27 println!("logic: {formula}");
28 }
29
```

---

<sup>16</sup>These three patterns—*objectives*, *soft constraints*, and *unsat-core extraction*—cover 95 % of practical optimisation tasks you will encounter when combining Rust with Z3.

<sup>17</sup>Most of the patterns here mirror what you would write by hand against `std::ops`, but the macros shrink hundreds of lines of duplicated impls into a maintainable few dozen.

## Bit-Vector (BV) Operators

| Operator                         | Z3 Method          |
|----------------------------------|--------------------|
| <code>+, +=</code>               | <code>bvadd</code> |
| <code>-, -=</code>               | <code>bvsub</code> |
| <code>*, **=</code>              | <code>bvmul</code> |
| <code>&amp;, &amp;=</code>       | <code>bvand</code> |
| <code> ,  =</code>               | <code>bvor</code>  |
| <code>^, ^=</code>               | <code>bvxor</code> |
| <code>&lt;&lt;, &lt;&lt;=</code> | <code>bvshl</code> |
| <code>!</code>                   | <code>bvnot</code> |
| <code>- (unary)</code>           | <code>bvneg</code> |

## Integer (Int) Operators

| Operator               | Z3 Method                |
|------------------------|--------------------------|
| <code>+, +=</code>     | <code>add</code>         |
| <code>-, -=</code>     | <code>sub</code>         |
| <code>*, *=</code>     | <code>mul</code>         |
| <code>/, /=</code>     | <code>div</code>         |
| <code>%, %=</code>     | <code>rem</code>         |
| <code>- (unary)</code> | <code>unary_minus</code> |

## Real (Real) Operators

| Operator               | Z3 Method                |
|------------------------|--------------------------|
| <code>+, +=</code>     | <code>add</code>         |
| <code>-, -=</code>     | <code>sub</code>         |
| <code>*, *=</code>     | <code>mul</code>         |
| <code>/, /=</code>     | <code>div</code>         |
| <code>- (unary)</code> | <code>unary_minus</code> |

## Float (Float) Operators

| Operator               | Z3 Method              |
|------------------------|------------------------|
| <code>- (unary)</code> | <code>unary_neg</code> |

## Boolean (Bool) Operators

| Operator                   | Z3 Method        |
|----------------------------|------------------|
| <code>&amp;, &amp;=</code> | <code>and</code> |
| <code> ,  =</code>         | <code>or</code>  |
| <code>^, ^=</code>         | <code>xor</code> |
| <code>!</code>             | <code>not</code> |

## 20 *model* module

A `Model` represents one satisfying assignment (or optimum witness) produced by a `Z3 solver`. In other words, a `model` in Z3 represents a satisfying assignment for a set of constraints. When Z3 successfully finds that a formula is satisfiable (`SatResult::Sat`), it can provide a model that assigns concrete values to the variables and functions in your formula, demonstrating why the formula is true.<sup>18</sup>

- `impl Model:`
  - `pub fn of_solver(slv: &Solver) -> Option<Model>:` Extract a model from a solver after a successful `check()` call that returned `SatResult::Sat`. Returns `None` if no model is available (e.g., if the last check was unsatisfiable). This is the most common way to obtain a model.

---

<sup>18</sup> Again if the formula is `SatResult::Unknown`, there is no guarantee that a model will be generated or the generated model may no satisfy the constraints.

```

1 use z3::{
2 ast::{self, Ast, Int},
3 Config, Context, Model, SatResult, Solver,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 let x = ast::Int::new_const(&ctx, "x");
12 let y = ast::Int::new_const(&ctx, "y");
13
14 // x + y = 10 and x > 5
15 solver.assert(&(&Int::add(&ctx, &[&x, &y]))._eq(&ast::Int::from_i64(&ctx, 10)));
16 solver.assert(&x.gt(&ast::Int::from_i64(&ctx, 5)));
17
18 match solver.check() {
19 SatResult::Sat => {
20 let m1 = Model::of_solver(&solver).unwrap();
21 let m2 = solver.get_model().unwrap();
22 println!("Verification with first method:");
23 println!("x = {}", m1.eval(&x, true).unwrap());
24 println!("y = {}", m1.eval(&y, true).unwrap());
25
26 // Verify the model with the second method
27 println!("Verification with second method:");
28 println!("x = {}", m2.eval(&x, true).unwrap());
29 println!("y = {}", m2.eval(&y, true).unwrap());
30 }
31 SatResult::Unsat => println!("No solution exists"),
32 SatResult::Unknown => println!("Z3 couldn't determine satisfiability"),
33 }
34 }

```

- pub fn of\_optimize(opt: &Optimize) -> Option<Model>: Same idea, but for Optimize but used where you’re not just looking for any satisfying assignment, but an optimal one according to some objective function. Lets you read the optimum assignment after pushing objectives.

```

1 use z3::{ast, ast::Int, Config, Context, Optimize, SatResult};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let opt = Optimize::new(&ctx);
7
8 let x = ast::Int::new_const(&ctx, "x");
9 let y = ast::Int::new_const(&ctx, "y");
10
11 // x + y >= 10, x >= 0, y >= 0
12 opt.assert(&Int::add(&ctx, &[&x, &y]).ge(&ast::Int::from_i64(&ctx, 10)));
13 opt.assert(&x.ge(&ast::Int::from_i64(&ctx, 0)));
14 opt.assert(&y.ge(&ast::Int::from_i64(&ctx, 0)));
15
16 // minimize x + y
17 opt.minimize(&Int::add(&ctx, &[&x, &y]));
18
19 match opt.check(&[]) {
20 SatResult::Sat => {

```

```

21 if let Some(model) = opt.get_model() {
22 println!("Optimal solution:");
23 println!("x = {}", model.eval(&x, true).unwrap());
24 println!("y = {}", model.eval(&y, true).unwrap());
25 }
26 }
27 - => println!("No optimal solution found"),
28 }
29 }
```

- pub fn translate(&self, dest: &Context) -> Model: Create a copy of this model in a different Z3 context. This is useful when you need to work with models across different Z3 contexts, as Z3 objects are tied to the context in which they were created.

```

1 use z3::{ast, ast::Ast, Config, Context, SatResult, Solver};
2
3 fn main() {
4 let cfg1 = Config::new();
5 let ctx1 = Context::new(&cfg1);
6 let solver1 = Solver::new(&ctx1);
7
8 let x1 = ast::Int::new_const(&ctx1, "x");
9 solver1.assert(&x1._eq(&ast::Int::from_i64(&ctx1, 42)));
10
11 match solver1.check() {
12 SatResult::Sat => {
13 let model1 = solver1.get_model().unwrap();
14
15 let cfg2 = Config::new();
16 let ctx2 = Context::new(&cfg2);
17 let model2 = model1.translate(&ctx2);
18
19 println!("Original model: {}", model1);
20 println!("Translated model: {}", model2);
21 }
22 - => println!("Unsatifiable"),
23 }
24 }
```

- pub fn get\_const\_interp<T: Ast>(&self, c: &T) -> Option<T>: Get the interpretation (assigned value) of a constant in the model. This works for any constant (zero-arity function), including variables you declared. Returns None if the constant has no interpretation in the model.

```

1 use z3::{ast, ast::Ast, Config, Context, SatResult, Solver};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let solver = Solver::new(&ctx);
7
8 let x = ast::Int::new_const(&ctx, "x");
9 let y = ast::Bool::new_const(&ctx, "y");
10
11 // x = 100 and y = true
12 solver.assert(&x._eq(&ast::Int::from_i64(&ctx, 100)));
13 solver.assert(&y._eq(&ast::Bool::from_bool(&ctx, true)));
14
15 if solver.check() == SatResult::Sat {
16 let model = solver.get_model().unwrap();
```

```

17 if let Some(x_val) = model.get_const_interp(&x) {
18 println!("x = {}", x_val);
19 }
20
21 if let Some(y_val) = model.get_const_interp(&y) {
22 println!("y = {}", y_val);
23 }
24
25 // Check for a constant that doesn't have an interpretation
26 let z = ast::Int::new_const(&ctx, "z");
27 match model.get_const_interp(&z) {
28 Some(z_val) => println!("z = {}", z_val),
29 None => println!("z has no interpretation in this model"),
30 }
31 }
32 }
33 }
```

- pub fn get\_func\_interp(&self, f: &FuncDecl) -> Option<FuncInterp>: Get the interpretation of a function declaration in the model. This handles both zero-arity functions (constants) and higher-arity functions. For arrays, this method provides access to the array’s interpretation as a function.
- pub fn eval<T: Ast>(&self, t: &T, model\_completion: bool) -> Option<T>: Evaluate any expression in the context of this model. The `model_completion` parameter determines whether Z3 should assign default values to undefined parts of the model. This is the most flexible way to extract values from a model. When `model_completion` is `true`, Z3 will provide interpretations for symbols that don’t have explicit values in the model, essentially “completing” the model. When `false`, evaluation may fail for undefined symbols.

```

1 use z3::{
2 ast::{self, Ast, Int},
3 Config, Context, SatResult, Solver,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 let x = ast::Int::new_const(&ctx, "x");
12 let y = ast::Int::new_const(&ctx, "y");
13
14 solver.assert(&x._eq(&ast::Int::from_i64(&ctx, 7))); // x = 7
15
16 if solver.check() == SatResult::Sat {
17 let model = solver.get_model().unwrap();
18
19 println!("x = {}", model.eval(&x, false).unwrap());
20
21 // (x + y) * 2
22 let expr = Int::mul(
23 &ctx,
24 &[&Int::add(&ctx, &[&x, &y]), &ast::Int::from_i64(&ctx, 2)],
25);
26
27 match model.eval(&expr, false) {
28 Some(val) => println!("2*(x+y) without completion = {}", val),
29 None => println!("Cannot evaluate 2*(x+y) without completion"),
30 }
31 }
32 }
```

```

30 }
31 if let Some(val) = model.eval(&expr, true) {
32 println!("2*(x+y) with completion = {}", val);
33 }
34
35 if let Some(y_val) = model.eval(&y, true) {
36 println!("y (with completion) = {}", y_val);
37 }
38 }
39 }
```

- `pub fn iter(&self) -> impl Iterator<Item = FuncDecl>:` Iterate over every constant and function that *does* have an interpretation.

## Tips:

- **Context Lifetime:** Models are tied to their Z3 context. Make sure the context lives long enough for your model usage.
- **Checking for Interpretations:** Always check if `get_const_interp` or `eval` returns `Some` before using the result, as not all symbols may have interpretations.
- **Array Handling:** Arrays are represented as functions in Z3 models. Use `get_func_interp` to examine array contents.

## 21 *goal* module

Goals are collections of constraints that tactics can process, simplify, or solve. Unlike direct solver usage, goals allow you to apply *preprocessing* transformations before solving.

**Goals vs. Solvers:** Goals are *containers for formulas* that can be transformed by tactics, while **Solvers** directly check satisfiability. Goals provide more fine-grained control over the solving process by allowing tactical preprocessing.

The following public items are accessible:

- `pub fn new(ctx: &Context, models: bool, unsat_cores: bool, proofs: bool) -> Goal:` Create a new goal with specified capabilities. The boolean parameters control what additional information Z3 will track:

- `models`: Generate models for satisfiable goals
- `unsat_cores`: Track unsatisfiable cores for debugging
- `proofs`: Generate proofs (requires context with proof support otherwise a segmentation error will occur)

```

1 use z3::{
2 ast::Int,
3 Config, Context, Goal,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let goal = Goal::new(&ctx, true, false, false);
10
11 let x = Int::new_const(&ctx, "x");
12 let y = Int::new_const(&ctx, "y");
13 goal.assert(&x.gt(&Int::from_i64(&ctx, 0))); // x > 0
14 goal.assert(&y.lt(&Int::from_i64(&ctx, 10))); // y < 10
```

```

15 goal.assert(&x.le(&y)); // x <= y
16
17 println!("Goal: {}", goal);
18 println!("Goal size: {}", goal.get_size());
19 }

```

- pub fn assert(&self, ast: &impl Ast) -> (): Add a formula to the goal. This is the primary way to build up a collection of constraints that will be processed together.
- pub fn is\_inconsistent(&self) -> bool: Check if the goal contains the formula `false`, indicating it's unsatisfiable. This will only return `true` if the `goal.assert(&Bool::from_bool(&ctx, false))` is explicitly added as a term to the formula. In short, *it is only a syntactic test that asks does this goal already contain the literal false?* It does not do reasoning.

```

1 use z3::{
2 ast::{Bool, Int},
3 Config, Context, Goal,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let goal = Goal::new(&ctx, true, false, false);
10
11 let x = Int::new_const(&ctx, "x");
12
13 goal.assert(&x.gt(&Int::from_i64(&ctx, 5)));
14 goal.assert(&x.lt(&Int::from_i64(&ctx, 0)));
15 goal.assert(&Bool::from_bool(&ctx, false)); // comment this line to make the goal
16 → consistent even thoug there is a contradiction!!!
17 println!("Goal is inconsistent: {}", goal.is_inconsistent());
18
19 let consistent_goal = Goal::new(&ctx, true, false, false);
20 let y = Int::new_const(&ctx, "y");
21 consistent_goal.assert(&y.gt(&Int::from_i64(&ctx, 0)));
22 consistent_goal.assert(&y.lt(&Int::from_i64(&ctx, 10)));
23
24 println!("Consistent goal is inconsistent: {}", consistent_goal.is_inconsistent());
25 }

```

- pub fn get\_depth(&self) -> u32: Return the depth of transformations applied to the goal. This tracks how many tactical transformations have been applied.
- pub fn get\_size(&self) -> u32: Return the number of formulas in the goal. This counts top-level assertions, not subformulas.
- pub fn get\_num\_expr(&self) -> u32: Return the total number of expressions (formulas, subformulas, and terms) in the goal. This provides a measure of the goal's complexity.

```

1 use z3::{
2 ast::{Bool, Int},
3 Config, Context, Goal,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let goal = Goal::new(&ctx, true, false, false);
10
11 let x = Int::new_const(&ctx, "x");
12 let y = Int::new_const(&ctx, "y");

```

```

13
14 println!(

15 "Empty goal - Size: {}, Expressions: {}",

16 goal.get_size(),

17 goal.get_num_expr()

18);

19

20 goal.assert(&x.gt(&Int::from_i64(&ctx, 0)));

21 println!(

22 "After x > 0 - Size: {}, Expressions: {}",

23 goal.get_size(),

24 goal.get_num_expr()

25);

26

27 let complex = Bool::and(

28 &ctx,

29 &[

30 Int::add(&ctx, &[&x, &y]).gt(&Int::from_i64(&ctx, 5)), // x + y > 5

31 Int::sub(&ctx, &[&x, &y]).lt(&Int::from_i64(&ctx, 2)), // x - y < 2

32],

33);

34 goal.assert(&complex);

35 println!(

36 "After complex constraint - Size: {}, Expressions: {}",

37 goal.get_size(),

38 goal.get_num_expr()

39);

40 }

```

- pub fn `is_decided_sat(&self)` -> bool: Return true if the goal is empty and the precision is precise or under—cases where an empty set of constraints really is a guaranteed Sat instance. This indicates the goal has been solved and is satisfiable. *It does not try to solve the goal; it merely asks whether the goal is already in a trivially decidable syntactic form; i.e., the goal is empty (no formulas left) and its precision tag is precise or under.*
- pub fn `is_decided_unsat(&self)` -> bool: Return true if the goal contains false and the precision is precise or over—situations where false is a trustworthy Unsat certificate. This indicates the goal has been solved and is unsatisfiable. *It does not try to solve the goal; it merely asks whether the goal is already in a trivially decidable syntactic form; i.e., the goal contains the literal false and its precision is precise or over.*

```

1 use z3::{ast::Int, Config, Context, Goal, Tactic};

2

3 fn main() {

4 let cfg = Config::new();

5 let ctx = Context::new(&cfg);

6

7 let sat_goal = Goal::new(&ctx, true, false, false);

8 let x = Int::new_const(&ctx, "x");

9 sat_goal.assert(&x.gt(&Int::from_i64(&ctx, 0)));

10

11 let unsat_goal = Goal::new(&ctx, true, false, false);

12 unsat_goal.assert(&x.gt(&Int::from_i64(&ctx, 5)));

13 unsat_goal.assert(&x.lt(&Int::from_i64(&ctx, 0)));

14

15 println!("Sat goal decided sat: {}", sat_goal.is_decided_sat());

16 println!("Sat goal decided unsat: {}", unsat_goal.is_decided_unsat());

17

18 let tactic = Tactic::new(&ctx, "simplify")

```

```

19 .and_then(&Tactic::new(&ctx, "propagate-values"))
20 .and_then(&Tactic::new(&ctx, "smt"));
21 if let Ok(result) = tactic.apply(&sat_goal, None) {
22 for subgoal in result.list_subgoals() {
23 println!("Subgoal decided sat: {}", subgoal.is_decided_sat());
24 println!("Subgoal decided unsat: {}", subgoal.is_decided_unsat());
25 }
26 };
27 }

```

- `pub fn reset(&self) -> ()`: Erase all formulas from the goal, returning it to an empty state. This is useful for reusing goal objects.
- `pub fn translate(self, ctx: &Context) -> Goal`: Move a goal from one Z3 context to another. This is necessary when working with multiple contexts.

```

1 use z3::{
2 ast::Int,
3 Config, Context, Goal,
4 };
5
6 fn main() {
7 let cfg1 = Config::new();
8 let ctx1 = Context::new(&cfg1);
9 let goal1 = Goal::new(&ctx1, true, false, false);
10
11 let x = Int::new_const(&ctx1, "x");
12 goal1.assert(&x.gt(&Int::from_i64(&ctx1, 0))); // x > 0
13
14 println!("Original goal: {}", goal1);
15
16 let cfg2 = Config::new();
17 let ctx2 = Context::new(&cfg2);
18 let goal2 = goal1.translate(&ctx2);
19
20 println!("Translated goal: {}", goal2);
21
22 let y = Int::new_const(&ctx2, "y");
23 goal2.assert(&y.lt(&Int::from_i64(&ctx2, 10))); // y < 10
24 println!("Goal after adding constraint in new context: {}", goal2);
25 }

```

- `pub fn get_precision(&self) -> GoalPrec`: Return the precision level of the goal. Goals can be transformed using over-approximations or under-approximations, and this method indicates the current precision level.
- `pub fn iter_formulas<T>(&self) -> impl Iterator<Item = T>`: Create an iterator over all formulas in the goal.

- `pub fn get_formulas<T>(&self) -> Vec<T>`: Return a vector containing all formulas in the goal.

```

1 use z3::{
2 ast::{Ast, Bool, Int},
3 Config, Context, Goal,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let goal = Goal::new(&ctx, true, false, false);
10

```

```

11 let x = Int::new_const(&ctx, "x");
12 let y = Int::new_const(&ctx, "y");
13
14 goal.assert(&x.gt(&Int::from_i64(&ctx, 0))); // x > 0
15 goal.assert(&y.lt(&Int::from_i64(&ctx, 10))); // y < 10
16 goal.assert(&Int::add(&ctx, &[&x, &y])._eq(&Int::from_i64(&ctx, 5))); // x + y = 5
17
18 println!("Formulas via iterator:");
19 for (i, formula) in goal.iter_formulas::<Bool>().enumerate() {
20 println!(" {}: {}", i, formula);
21 }
22
23 println!("Formulas via get_formulas:");
24 let formulas: Vec<Bool> = goal.get_formulas();
25 for (i, formula) in formulas.iter().enumerate() {
26 println!(" {}: {}", i, formula);
27 }
28
29 let combined = Bool::and(&ctx, &formulas);
30 println!("Combined formula: {}", combined);
31 }

```

## 22 func\_interp module

The FuncInterp API allows you to work with *function interpretations* in Z3 models. When Z3 solves constraints involving function declarations, it produces a *function interpretation* that specifies how the function maps inputs to outputs.<sup>19</sup>

**What is a Function Interpretation?** A function interpretation is essentially a *lookup table* that tells you what value a function returns for specific inputs, plus a default *else* value for all other inputs. For example, if you have a function  $f(x)$  and Z3 determines that  $f(1) = 10$ ,  $f(2) = 20$ , and  $f(x) = 0$  for all other values of  $x$ , then the function interpretation contains two explicit entries and an else value of 0.

The following public items are accessible:

- `pub fn get_arity(&self) -> usize`: Returns the number of arguments the function takes. For example, a function  $f(x, y)$  has arity 2, while  $g(x)$  has arity 1.

```

1 use z3::*;
2 ast::{Ast, Int},
3 Config, Context, FuncDecl, Solver, Sort,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 let f = FuncDecl::new(&ctx, "f", &[&Sort::int(&ctx)], &Sort::int(&ctx));
12 let input1 = Int::from_i64(&ctx, 1);
13 let input2 = Int::from_i64(&ctx, 2);
14 let value1 = Int::from_i64(&ctx, 100);
15 let value2 = Int::from_i64(&ctx, 200);
16
17 // f(input1) = value1 and f(input2) = value2
18 solver.assert(&f.apply(&[&input1])._eq(&value1.into()));

```

---

<sup>19</sup>Arrays in Z3 are internally represented as functions using the (`_as_array f`) construct, but when extracting interpretations from models, arrays typically require `get_const_interp` rather than `get_func_interp`, as Z3 often represents array models as constant values rather than explicit function interpretations.

```

19 solver.assert(&f.apply(&[&input2])._eq(&value2.into()));
20
21 if solver.check() == z3::SatResult::Sat {
22 let model = solver.get_model().unwrap();
23 if let Some(func_interp) = model.get_func_interp(&f) {
24 println!("Function arity: {}", func_interp.get_arity()); // Function arity: 1
25 }
26 }
27 }
```

- `pub fn get_num_entries(&self) -> u32`: Returns the number of explicit entries in the function interpretation. This tells you how many input-output pairs Z3 has explicitly defined, excluding the default `else` value.

```

1 use z3::{
2 ast::{Ast, Int},
3 Config, Context, FuncDecl, Solver, Sort,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 let f = FuncDecl::new(&ctx, "f", [&Sort::int(&ctx)], &Sort::int(&ctx));
12 let input1 = Int::from_i64(&ctx, 1);
13 let input2 = Int::from_i64(&ctx, 2);
14 let value1 = Int::from_i64(&ctx, 100);
15 let value2 = Int::from_i64(&ctx, 200);
16
17 // f(input1) = value1 and f(input2) = value2
18 solver.assert(&f.apply(&[&input1])._eq(&value1.into()));
19 solver.assert(&f.apply(&[&input2])._eq(&value2.into()));
20
21 if solver.check() == z3::SatResult::Sat {
22 let model = solver.get_model().unwrap();
23 if let Some(func_interp) = model.get_func_interp(&f) {
24 println!("Function interpretation: {:?}", func_interp);
25 println!("Function arity: {}", func_interp.get_arity()); // Function arity: 1
26 let num_entries = func_interp.get_num_entries();
27 println!("Number of entries: {}", num_entries); // Number of entries: 1
28 }
29 }
30 }
```

- `pub fn add_entry(&self, args: &[Dynamic], value: &Dynamic)`: Adds a new entry to the function interpretation, specifying that the function should return `value` when called with `args`. This is useful when you're building function interpretations programmatically.

```

1 use z3::{
2 ast::{Ast, Int},
3 Config, Context, FuncDecl, Solver, Sort,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 let f = FuncDecl::new(&ctx, "f", [&Sort::int(&ctx)], &Sort::int(&ctx));
```

```

12 let input1 = Int::from_i64(&ctx, 1);
13 let input2 = Int::from_i64(&ctx, 2);
14 let value1 = Int::from_i64(&ctx, 100);
15 let value2 = Int::from_i64(&ctx, 200);
16
17 // f(input1) = value1 and f(input2) = value2
18 solver.assert(&f.apply(&[&input1])._eq(&value1.into()));
19 solver.assert(&f.apply(&[&input2])._eq(&value2.into()));
20
21 if solver.check() == z3::SatResult::Sat {
22 let model = solver.get_model().unwrap();
23 if let Some(func_interp) = model.get_func_interp(&f) {
24 println!("Function interpretation: {:?}", func_interp);
25 println!("Function arity: {}", func_interp.get_arity()); // Function arity: 1
26 let num_entries = func_interp.get_num_entries();
27 println!("Number of entries: {}", num_entries); // Number of entries: 1
28 // add entries
29 for i in 0..3 {
30 func_interp.add_entry(
31 &[Int::from_i64(&ctx, i as i64).into()],
32 &Int::from_i64(&ctx, i as i64 * 10).into(), // overwrites existing
33 // entries (2 -> 20)
34);
35 }
36 println!("Updated function interpretation: {:?}", func_interp);
37 let num_entries = func_interp.get_num_entries();
38 println!("Number of entries: {}", num_entries);
39 }
40 }

```

- `pub fn get_entries(&self) -> Vec<FuncEntry>`: Returns all explicit entries in the function interpretation (not the `else`). Each entry contains the input arguments and corresponding output value.
- `pub fn get_else(&self) -> Dynamic`: Returns the default `else` value that the function returns for all inputs not explicitly specified in the entries. This is crucial for understanding complete function behavior.
- `pub fn set_else(&self, ast: &Dynamic)`: Sets the default value for the function interpretation. This value will be returned for any input arguments not explicitly covered by the entries.

Note that regarding *efficiency*, function interpretations are more memory-efficient than storing every possible input-output pair.

## 23 *func\_entry* module

The `FuncEntry` API represents individual entries in a function interpretation.<sup>20</sup> The following public items are accessible:

**Function Entries vs. Function Interpretations:** A `FuncEntry` is a single mapping (e.g., `f(2, 3) = 7`), while a `FuncInterp` is the complete function definition containing all entries plus a default `else` value.

- `pub fn get_value(&self) -> Dynamic`: Returns the output value that the function produces for this entry's arguments. This is the result of applying the function to the specific argument combination represented by this entry.

---

<sup>20</sup>A function entry is a single mapping from arguments to a value within a function interpretation, essentially representing one row in a function's lookup table.

```

1 use z3::{
2 ast::{Ast, Int},
3 Config, Context, FuncDecl, SatResult, Solver,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9 let solver = Solver::new(&ctx);
10
11 // f: Int x Int -> Int
12 let int_sort = Int::new_const(&ctx, "dummy").get_sort();
13 let f = FuncDecl::new(&ctx, "f", &[&int_sort, &int_sort], &int_sort);
14
15 let x = Int::new_const(&ctx, "x");
16 let y = Int::new_const(&ctx, "y");
17 let z = Int::new_const(&ctx, "z");
18 let two = Int::from_i64(&ctx, 2);
19 let three = Int::from_i64(&ctx, 3);
20 let seven = Int::from_i64(&ctx, 7);
21
22 // f(2, 3) = 7
23 solver.assert(&f.apply(&[&two, &three])._eq(&seven.into()));
24
25 // f(x, y) = z where x + y = 5 and z = 10
26 let five = Int::from_i64(&ctx, 5);
27 let ten = Int::from_i64(&ctx, 10);
28 solver.assert(&Int::add(&ctx, &[&x, &y])._eq(&five));
29 solver.assert(&f.apply(&[&x, &y])._eq(&z.clone().into()));
30 solver.assert(&z._eq(&ten));
31
32 if solver.check() == SatResult::Sat {
33 let model = solver.get_model().unwrap();
34 println!("{} is satisfiable", model);
35 if let Some(func_interp) = model.get_func_interp(&f) {
36 for entry in func_interp.get_entries() {
37 let value = entry.get_value();
38 let num_args = entry.get_num_args();
39 let args = entry.get_args();
40 println!("Function entry value: {}", value);
41 println!("Number of arguments: {}", num_args);
42 for (i, arg) in args.iter().enumerate() {
43 println!("Argument {}: {}", i, arg);
44 }
45 }
46 }
47 }
48 }

```

- pub fn get\_num\_args(&self) -> u32: Returns the number of arguments (arity) for this function entry. This tells you how many input parameters the function takes for this specific mapping.
- pub fn get\_args(&self) -> Vec<Dynamic>: Returns the input arguments for this function entry as a vector of Dynamic values. These are the specific values that, when passed to the function, produce the output returned by get\_value().

## 24 *func\_decl* module

The FuncDecl API provides a way to create and work with function declarations in Z3. Function declarations define the signature of functions (their name, domain types, and range type) without specifying their behavior. They are essential for creating uninterpreted functions.

**Function Declaration:** A *signature specification* that defines a function's name and type information (domain and range sorts). It declares what the function looks like but not what it does. For example, declaring  $f: \text{Int} \times \text{Int} \rightarrow \text{Bool}$  creates a function that takes two integers and returns a boolean, but doesn't specify the actual mapping behavior.

**Function Interpretation:** A *concrete implementation* that defines how a function behaves by providing explicit input-output mappings. This is what you get in a model after Z3 solves constraints involving the function declaration.

The following public items are accessible:

- pub fn new<S: Into<Symbol>>(ctx: &Context, name: S, domain: &[&Sort], range: &Sort) -> FuncDecl: Create a new function declaration with the given name, domain types, and range type. The domain is an array of sorts representing the argument types, and the range is the return type. If the domain is empty, this creates a *constant* declaration.

```
1 use z3::{Config, Context, FuncDecl, Sort};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6
7 // constant declaration (0-arity function)
8 let x = FuncDecl::new(&ctx, "x", &[], &Sort::int(&ctx));
9 println!("Constant x: {}", x);
10
11 // unary function f: Int → Bool
12 let f = FuncDecl::new(&ctx, "f", &[&Sort::int(&ctx)], &Sort::bool(&ctx));
13 println!("Function f: {}", f);
14
15 // binary function add: Int × Int → Int
16 let add = FuncDecl::new(
17 &ctx,
18 "add",
19 &[&Sort::int(&ctx), &Sort::int(&ctx)],
20 &Sort::int(&ctx),
21);
22 println!("Function add: {}", add);
23
24 // ternary function: Real × Real × Real → Real
25 let distance3d = FuncDecl::new(
26 &ctx,
27 "distance3d",
28 &[&Sort::real(&ctx), &Sort::real(&ctx), &Sort::real(&ctx)],
29 &Sort::real(&ctx),
30);
31 println!("Function distance3d: {}", distance3d);
32 println!("Function distance3d arity: {}", distance3d.arity());
33 }
```

- pub fn arity(&self) -> usize: Return the number of arguments (arity) of the function declaration. For constants, this returns 0. For unary functions, this returns 1, and so on.
- pub fn apply(&self, args: &[&dyn Ast]) -> Dynamic: Create a function application by ap-

plying the function declaration to the given arguments. If no arguments are provided (empty slice), this creates a constant. The number and types of arguments must match the function's domain specification.

```

1 use z3::{
2 ast::{Ast, Bool, Dynamic, Int},
3 Config, Context, FuncDecl, SatResult, Solver, Sort,
4 };
5
6 fn main() {
7 let cfg = Config::new();
8 let ctx = Context::new(&cfg);
9
10 let x_decl = FuncDecl::new(&ctx, "x", &[], &Sort::int(&ctx));
11 let y_decl = FuncDecl::new(&ctx, "y", &[], &Sort::int(&ctx));
12
13 // Create function declarations
14 let f_decl = FuncDecl::new(&ctx, "f", &[&Sort::int(&ctx)], &Sort::bool(&ctx));
15 println!("kind of f: {:?}", f_decl.kind());
16 let add_decl = FuncDecl::new(
17 &ctx,
18 "add",
19 &[&Sort::int(&ctx), &Sort::int(&ctx)],
20 &Sort::int(&ctx),
21);
22
23 let x: Dynamic = x_decl.apply(&[]); // Constant x
24 let y: Dynamic = y_decl.apply(&[]); // Constant y
25 let five = Int::from_i64(&ctx, 5);
26 let f_5: Dynamic = f_decl.apply(&[&five]); // f(5)
27 let ten = Int::from_i64(&ctx, 10);
28 let add_5_10: Dynamic = add_decl.apply(&[&five, &ten]); // add(5, 10)
29
30 let solver = Solver::new(&ctx);
31 // Assert f(5) is true
32 solver.assert(&Bool::from_bool(&ctx, true).eq(&f_5.as_bool().unwrap()));
33 // Assert add(5, 10) = 15
34 solver.assert(&Int::from_i64(&ctx, 15).eq(&add_5_10.as_int().unwrap()));
35 // Assert add(x, y) = 20
36 let twenty = Int::from_i64(&ctx, 20);
37 solver.assert(&twenty.eq(&add_decl.apply(&[&x, &y]).as_int().unwrap()));
38
39 match solver.check() {
40 SatResult::Sat => {
41 let model = solver.get_model().unwrap();
42 println!("Model found!");
43 println!("x = {}", model.eval(&x, true).unwrap());
44 println!("y = {}", model.eval(&y, true).unwrap());
45 println!("f(5) = {}", model.eval(&f_5, true).unwrap());
46 println!("add(5, 10) = {}", model.eval(&add_5_10, true).unwrap());
47 }
48 SatResult::Unsat => println!("No solution"),
49 SatResult::Unknown => println!("Unknown"),
50 }
51 }
```

- pub fn kind(&self) -> DeclKind: Return the declaration kind of this function declaration. This tells you what type of declaration it is (e.g., uninterpreted function, built-in operator, etc.).

```
1 use z3::{Config, Context, FuncDecl, Sort};
```

```

2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6
7 let f = FuncDecl::new(&ctx, "f", [&Sort::int(&ctx)], &Sort::bool(&ctx));
8 let g = FuncDecl::new(&ctx, "g", [], &Sort::int(&ctx));
9 let h = FuncDecl::new(
10 &ctx,
11 "binary_op",
12 [&Sort::int(&ctx), &Sort::int(&ctx)],
13 &Sort::int(&ctx),
14);
15
16 println!("Function f kind: {:?}", f.kind());
17 println!("Constant g kind: {:?}", g.kind());
18 println!("Function h kind: {:?}", h.kind());
19
20 }

```

All of the above are UNINTERPRETED.

- `pub fn name(&self) -> String`: Return the name of this function declaration as a string. For string symbols, this returns the symbol directly. For integer symbols, this returns the symbol with a "k!" prefix.

## 25 *datatype\_builder* module

The `DatatypeBuilder` API lets you create custom *algebraic data types* in Z3. Datatypes are composite types that can represent structures like lists, trees, structs, enums, and other recursive data structures. Datatypes are ADTs that combine multiple values into structured data. Unlike basic sorts (integers, booleans), datatypes have *constructors* (functions that create values) and *accessors* (functions that extract components). Each datatype can have multiple *variants*, each with different fields.

- `pub fn new<S: Into<Symbol>>(ctx: &Context, name: S) -> Self`: Create a new datatype builder with the given name. The name will be used as the sort name for the final datatype.
- `pub fn variant(mut self, name: &str, fields: Vec<(&str, DatatypeAccessor)>) -> Self`: Add a variant (constructor) to the datatype. Each variant represents a different way to construct values of this datatype. The `name` parameter is the constructor name, and `fields` specify the components that this variant contains. Each field has a name and a `DatatypeAccessor` that specifies its type (either another datatype or a regular sort).
- `pub fn finish(self) -> DatatypeSort`: Finalize the datatype construction and return a `DatatypeSort` that can be used to create values and constraints. This creates the actual Z3 datatype sort with all its *constructors, testers, and accessors*.

```

1 use z3::{ast::Datatype, Config, Context, DatatypeBuilder, SatResult, Solver};
2
3 fn main() {
4 let cfg = Config::new();
5 let ctx = Context::new(&cfg);
6 let solver = Solver::new(&ctx);
7
8 let color_sort = DatatypeBuilder::new(&ctx, "Color")
9 .variant("Red", vec![])
10 .variant("Green", vec![])
11 .variant("Blue", vec![])
12 .finish();
13

```

```

14 let red = color_sort.variants[0].constructor.apply(&[]);
15 let green = color_sort.variants[1].constructor.apply(&[]);
16 let blue = color_sort.variants[2].constructor.apply(&[]);
17
18 // Create a variable of Color type
19 let x = Datatype::new_const(&ctx, "x", &color_sort.sort);
20
21 // Assert that x is red
22 let is_red = color_sort.variants[0].tester.apply(&[&x]);
23 solver.assert(&is_red.as_bool().unwrap());
24
25 assert_eq!(solver.check(), SatResult::Sat);
26 let model = solver.get_model().unwrap();
27 println!("x = {}", model.eval(&x, true).unwrap());
28 }

```

For more details, please refer to the section on `DatatypeBuilder` in `lib.rs`.

- pub fn `create_datatypes(datatype_builders: Vec<DatatypeBuilder>) -> Vec<DatatypeSort>`: Create multiple datatypes at once. This is required when you have datatypes that reference each other (mutual recursion). Z3 needs to create all mutually recursive datatypes in a single operation to properly handle the circular dependencies.

```

1 use z3::{
2 ast::{Datatype, Int},
3 datatype_builder::create_datatypes,
4 Config, Context, DatatypeAccessor, DatatypeBuilder, SatResult, Solver, Sort,
5 };
6
7 fn main() {
8 let cfg = Config::new();
9 let ctx = Context::new(&cfg);
10 let solver = Solver::new(&ctx);
11
12 let expr_builder = DatatypeBuilder::new(&ctx, "Expr")
13 .variant(
14 "Num",
15 vec![("value", DatatypeAccessor::Sort(Sort::int(&ctx)))] ,
16)
17 .variant(
18 "Add",
19 vec![
20 ("left", DatatypeAccessor::Datatype("Expr".into())),
21 ("right", DatatypeAccessor::Datatype("Expr".into())),
22],
23)
24 .variant(
25 "Let",
26 vec![
27 ("binding", DatatypeAccessor::Datatype("Binding".into())),
28 ("body", DatatypeAccessor::Datatype("Expr".into())),
29],
30);
31
32 let binding_builder = DatatypeBuilder::new(&ctx, "Binding").variant(
33 "Bind",
34 vec![
35 ("name", DatatypeAccessor::Sort(Sort::int(&ctx))), // simplified with int
36 ("value", DatatypeAccessor::Datatype("Expr".into())),
37],

```

```

38);
39
40 let datatypes = create_datatypes(vec![expr_builder, binding_builder]);
41 let expr_sort = &datatypes[0];
42 let binding_sort = &datatypes[1];
43
44 let five = Int::from_i64(&ctx, 5);
45 let x_name = Int::from_i64(&ctx, 1);
46
47 let num5 = expr_sort.variants[0].constructor.apply(&[&five]);
48
49 // let x = 5
50 let binding = binding_sort.variants[0]
51 .constructor
52 .apply(&[&x_name, &num5]);
53 let body_expr = Datatype::new_const(&ctx, "body", &expr_sort.sort);
54 let let_expr = expr_sort.variants[2]
55 .constructor
56 .apply(&[&binding, &body_expr]);
57
58 // Assert that body_expr is an Add expression
59 let is_add = expr_sort.variants[1].tester.apply(&[&body_expr]);
60 solver.assert(&is_add.as_bool().unwrap());
61
62 // Check satisfiability
63 assert_eq!(solver.check(), SatResult::Sat);
64 let model = solver.get_model().unwrap();
65 println!("Let expression = {}", model.eval(&let_expr, true).unwrap());
66 println!(
67 "Body expression = {}",
68 model.eval(&body_expr, true).unwrap()
69);
70 }

```

### Tips:

**Constructors:** Functions that create values of the datatype. Each variant has its own constructor function.

**Testers:** Functions that check if a value was created with a specific constructor (e.g., `is-Red`, `is-Node`).

**Accessors:** Functions that extract field values from datatype instances. Only valid when the tester for the corresponding constructor returns true.

**Mutual Recursion:** When datatypes reference each other, they must be created together using `create_datatypes` to handle circular dependencies properly.

## 26 config module

The `Config` API allows you to customize Z3's behavior before creating a context. Configuration objects are used to set parameters that control proof generation, model generation, timeouts, and other Z3 engine settings. `Config` objects are created first to specify Z3's operational parameters, then passed to `Context::new()` to create the actual Z3 context. Once a context is created, its configuration cannot be changed. The following public items are accessible:

- `pub fn new() -> Config`: Create a new configuration object with default settings. This is the starting point for customizing Z3's behavior.
- `pub fn set_param_value(&mut self, k: &str, v: &str)`: Set any Z3 configuration parameter by name. This is the most general method for configuring Z3, allowing access to all available

parameters.

- `pub fn set_bool_param_value(&mut self, k: &str, v: bool)`: Helper method for setting boolean parameters. This automatically converts the boolean value to the appropriate string representation ("true" or "false").
- `pub fn set_proof_generation(&mut self, b: bool)`: Enable or disable proof generation. When enabled, Z3 will generate proofs for unsatisfiable results, which can be retrieved using `Solver::get_proof()`.
- `pub fn set_model_generation(&mut self, b: bool)`: Enable or disable model generation. When enabled, Z3 will generate models for satisfiable results, which can be retrieved using `Solver::get_model()`.
- `pub fn set_debug_ref_count(&mut self, b: bool)`: Enable or disable reference counting debugging. This is useful for debugging memory leaks in Z3 applications by tracking reference counts of Z3 objects.
- `pub fn set_timeout_msec(&mut self, ms: u64)`: Set a timeout in milliseconds for Z3 operations. This prevents Z3 from running indefinitely on difficult problems.

#### Tips:

- **Set configuration before context creation:** All configuration parameters must be set before creating the Z3 context. Configuration cannot be changed after context creation.
- **Enable features as needed:** Model and proof generation have performance overhead. Only enable them when you need to inspect results.
- **Use timeouts for robustness:** Set reasonable timeouts to prevent your application from hanging on difficult problems.

The `Config` type implements `Default`, so you can use `Config::default()` as an alternative to `Config::new()`:

```
1 use z3::{Config, Context};
2
3 fn main() {
4 // These are equivalent:
5 let cfg1 = Config::new();
6 let cfg2 = Config::default();
7
8 // Both create contexts with default settings
9 let ctx1 = Context::new(&cfg1);
10 let ctx2 = Context::new(&cfg2);
11}
```

## 27 ast module

This module provides the `Ast` trait and nodes representing various data values in Z3. The key AST node types include:

- **Bool:** `new_const` & `fresh_const` for creating boolean constants.<sup>21</sup> `from_bool` & `as_bool` for conversion to/from Rust Booleans. `ite` for *if-then-else expressions*. Logical operations: `and`, `or`, `xor`, `iff`, `implies`, `not`. Pseudo-boolean operations: `pb_le`, `pb_ge`, `pb_eq` for linear integer programming constraints. Use for boolean logic and satisfiability problems.

---

<sup>21</sup>`new_const` creates a constant with a user-provided name and returns the same symbolic variable if reused, making it ideal for defining and referencing specific variables. In contrast, `fresh_const` creates a constant with a unique, auto-generated name prefixed by the given string from the user, ensuring distinct variables each time.

- **Int:** `new_const` & `fresh_const` for creating integer constants. `new_const` creates a constant with a user-provided name and returns the same symbolic variable if reused, making it ideal for defining and referencing specific variables. In contrast, `fresh_const` creates a constant with a unique, auto-generated name prefixed by the given string, ensuring distinct variables each time. `from_big_int` & `from_str` for parsing from `BigInt` and `String`. `from_i64`, `from_u64`, `as_i64`, `as_u64` for conversion to/from Rust primitives. `from_real`, `to_real` for real number conversion. `from_bv`, `to_ast` for bitvector conversion. Arithmetic operations: `add`, `sub`, `mul`, `div`, `rem`, `modulo`, `power`, `unary_minus`. Comparison operations: `lt`, `le`, `gt`, `ge`.
- **Real:** `new_const` & `fresh_const` for creating real number constants. `from_big_rational` & `from_real_str` for parsing from rational numbers. `from_real` for creating from numerator/denominator pairs. `as_real` for extraction as (numerator, denominator) tuple. `approx`, `approx_f64` for decimal approximation. `from_int`, `to_int` for integer conversion. `is_int` predicate to check if a real is actually an integer. Arithmetic operations: `add`, `sub`, `mul`, `div`, `power`, `unary_minus`. Comparison operations: `lt`, `le`, `gt`, `ge`.
- **Float:** `new_const`, `new_const_float32`, `new_const_double` for creating floating-point constants. `fresh_const`, `fresh_const_float32`, `fresh_const_double` for fresh constants. `from_f32` & `from_f64` for conversion from Rust primitives. `as_f64` for extraction to Rust `f64`. Rounding modes: `round_towards_zero`, `round_towards_negative`, `round_towards_positive`. Arithmetic with explicit rounding: `add_towards_zero`, `sub_towards_zero`, `mul_towards_zero`, `div_towards_zero`. `to_ieee_bv` for IEEE bitvector representation. Unary operations: `unary_abs`, `unary_neg`. Predicates: `is_infinite`, `is_normal`, `is_subnormal`, `is_zero`, `is_nan`. Comparison operations: `lt`, `le`, `gt`, `ge`. General arithmetic: `add`, `sub`, `mul`, `div`.
- **String:** `new_const` & `fresh_const` for creating string constants. `from_str` & `as_string` for conversion to/from Rust strings. `at` for single character access. `substr` for substring extraction. `regex_matches` for regular expression matching. `concat` for string concatenation. `length` for string length. String predicates: `contains`, `prefix`, `suffix`.
- **BV:** `new_const` & `fresh_const` for creating bitvector constants. `from_str`, `from_i64`, `from_u64` for conversion from various formats. `as_i64`, `as_u64` for extraction to Rust primitives. `from_int`, `to_int` for integer conversion. `get_size` for bitwidth information. Bitwise operations: `bvnot`, `bvneg`, `bvredand`, `bvredor`, `bvand`, `bvor`, `bvxor`, `bvnand`, `bvnor`, `bvxnor`. Arithmetic operations: `bvadd`, `bvsub`, `bvmul`, `bvudiv`, `bvsdiv`, `bvurem`, `bvsrem`, `bvsmod`. Comparison operations: `bvult`, `bvslt`, `bvule`, `bvsle`, `bvuge`, `bvsge`, `bvugt`, `bvsge`. Shift operations: `bvshl`, `bvlshr`, `bvashr`, `bvrotl`, `bvrotr`. Manipulation operations: `concat`, `extract`, `sign_ext`, `zero_ext`. Overflow checking: `bvneg_no_overflow`, `bvadd_no_overflow`, etc.
- **Array:** `new_const` & `fresh_const` for creating array constants. `const_array` for creating arrays with constant default values. `select` & `select_n` for element access (single and n-ary). `store` for functional array update. `is_const_array` predicate to check if array has constant values.
- **Set:** `new_const` & `fresh_const` for creating set constants. `empty` for creating empty sets. `add`, `del` for element insertion and removal. `member` for membership testing. Set operations: `intersect`, `set_union`, `complement`, `difference`. Set relations: `set_subset`.
- **Seq:** `new_const` & `fresh_const` for creating sequence constants. `unit` for creating singleton sequences. `at` for unit sequence access at position. `nth` for element access at position. `length` for sequence length. `concat` for sequence concatenation.
- **Datatype:** `new_const` & `fresh_const` for creating datatype constants. Used for algebraic data types, enumerations, and user-defined structures.
- **Dynamic:** `from_ast` for creating from any AST node. `new_const` & `fresh_const` for creating constants with explicit sort. `sort_kind` for runtime type inspection. Type conversion methods:

`as_bool`, `as_int`, `as_real`, `as_float`, `as_string`, `as_bv`, `as_array`, `as_set`, `as_seq`, `as_datatype`. Use when the exact AST node type is unknown at compile time or for generic programming.

- **Regexp:** node representing a regular expression. `literal` for creating a regular expression that recognizes the string given as parameter. `range` for character range patterns (e.g., `[a-z]`). `loop`, `power` for repetition patterns. Special patterns: `full` (matches everything), `allchar` (matches any single character), `empty` (matches nothing). Quantifiers: `plus` (one or more), `star` (zero or more), `option` (zero or one). `complement` for pattern negation. `diff` for pattern difference. Composition operations: `concat`, `union`, `intersect`.

Each of the datatypes also share the methods provided by the `Ast` trait<sup>22</sup>:

- `get_ctx(&self) -> &Context` – Returns the Z3 context associated with this AST node.
- `get_z3_ast(&self) -> Z3_ast` – Returns the raw Z3 AST pointer.
- `wrap(ctx: &Context, ast: Z3_ast) -> Self` – Unsafe constructor for wrapping raw Z3 AST pointers.
- `_eq(&self, other: &Self) -> Bool` – Creates an equality constraint between two AST nodes of the *same* type.
- `_safe_eq(&self, other: &Self) -> Result<Bool, SortDiffers>` – Safe equality that returns an error if sorts don't match.
- `distinct(ctx: &Context, values: &[impl Borrow<Self>]) -> Bool` – Creates a constraint that all values with the same type are pairwise distinct.
- `get_sort(&self) -> Sort` – Returns the sort (type) of this AST node.
- `simplify(&self) -> Self` – Applies algebraic simplification rules and returns a simplified equivalent AST.
- `substitute<T: Ast>(&self, substitutions: &[(&T, &T)]) -> Self` – Performs substitution of sub-expressions with replacement values.
- `num_children(&self) -> usize` – Returns the number of child nodes (0 for leaf nodes). Use for AST traversal and analysis.
- `nth_child(&self, idx: usize) -> Option<Dynamic>` – Returns the *n*th child node as `Dynamic`. Use for AST traversal.
- `children(&self) -> Vec<Dynamic>` – Returns all child nodes as a vector of `Dynamic`. Use for complete AST traversal.
- `kind(&self) -> AstKind` – Returns the kind of AST node (App, Numeral, Var, etc.). Use for pattern matching on AST structure.
- `is_app(&self) -> bool` – Returns `true` if this is a function application (including constants).
- `is_const(&self) -> bool` – Returns `true` if this is a constant (0-arity function application);
- `decl(&self) -> FuncDecl` – Returns the function declaration (panics if not an app). Use to get function symbol information.
- `safe_decl(&self) -> Result<FuncDecl, IsNotApp>` – Safe version of `decl` that returns an error if not an app.
- `translate(&self, dest: &Context) -> Self` – Translates an AST node to a different Z3 context. Use when moving expressions between contexts.

---

<sup>22</sup>Abstract syntax tree (AST) nodes represent terms, constants, or expressions.

Additional utility functions:

- `forall_const(ctx, bounds, patterns, body) -> Bool` – Creates a universal quantifier over bound variables. Use for universal quantification in logic formulas.
- `exists_const(ctx, bounds, patterns, body) -> Bool` – Creates an existential quantifier over bound variables. Use for existential quantification in logic formulas.
- `quantifier_const(ctx, is_forall, weight, quantifier_id, skolem_id, bounds, patterns, no_patterns, body) -> Bool` – Creates a quantifier with additional attributes and control.
- `lambda_const(ctx, bounds, body) -> Array` – Creates a lambda expression represented as an array. Use for functional programming constructs.
- `atmost(ctx, args, k) -> Bool` – Creates an at-most- $k$  constraint over Boolean variables. It generates a Boolean constraint that at most  $k$  of the Boolean variables in args must be true
- `atleast(ctx, args, k) -> Bool` – Creates an at-least- $k$  constraint over Boolean variables. It generates a Boolean constraint that at least  $k$  of the Boolean variables in args must be true

#### Tips:

- All `Ast` types can be converted to the `Dynamic` data type and vice versa.
- `IsNotApp` is a trait that indicates that a node is not a function declaration. It has two methods, `new`, `kind`.