




Supercharging Virtual Plant Configurations using Z3

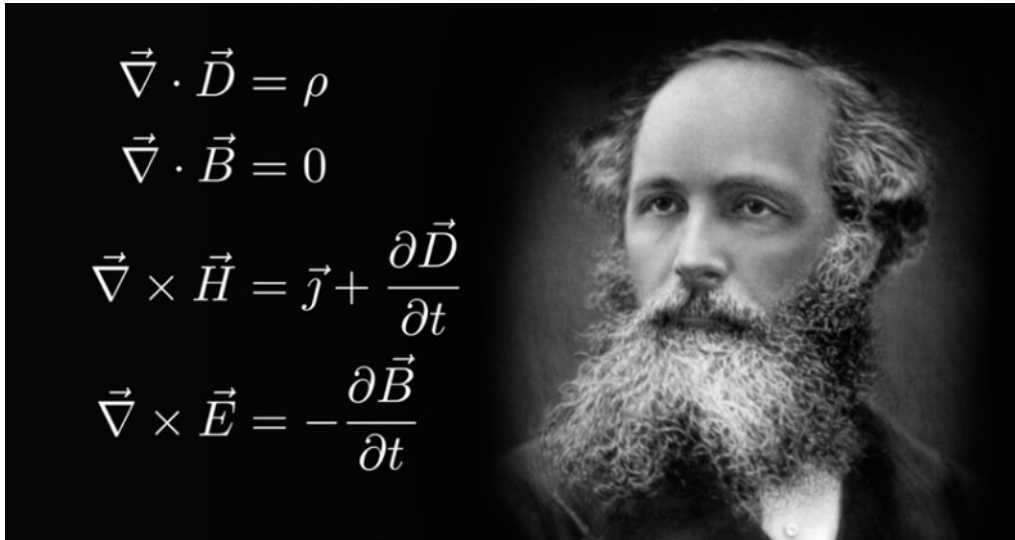
Nikolaj Bjørner, Max Levatich,
Nuno Lopes, Andrey Rybalchenko,
Chandra Vuppalapati

Microsoft



- 1 | Introduction and overview of **Z3**
- 2 | **Z3** for virtual plant design automation

Logic: The Calculus of Computation



Differential, Integral Calculus

Dynamics, Conduction,..
Matlab, Mathematica, Simulink

Invariants

Flowchart illustrating a program with invariants:

- Initial condition: $0 < n$
- Block 1: $r := 1$, $u := 1$
- Block 2: $v := u$ (Invariant: $r \leq n$, $u = r!$)
- Decision 1: $r - n \geq 0$ (True path: $v = n!$, STOP; False path: $s := 1$)
- Block 3: $u := u + v$ (Invariant: $s \leq r < n$, $u = sr!$, $v = r!$)
- Block 4: $s := s + 1$ (Invariant: $s \leq r < n$, $u = (s+1)r!$, $v = r!$)
- Decision 2: $(s-1) - r < 0$ (True path: $r := r + 1$; False path: $s-1 \leq r < n$)
- Block 5: $r := r + 1$ (Invariant: $r < n$, $u = (r+1)r!$)

Friday, 24th June [1949]
Checking a large routine by Dr A. Turing.

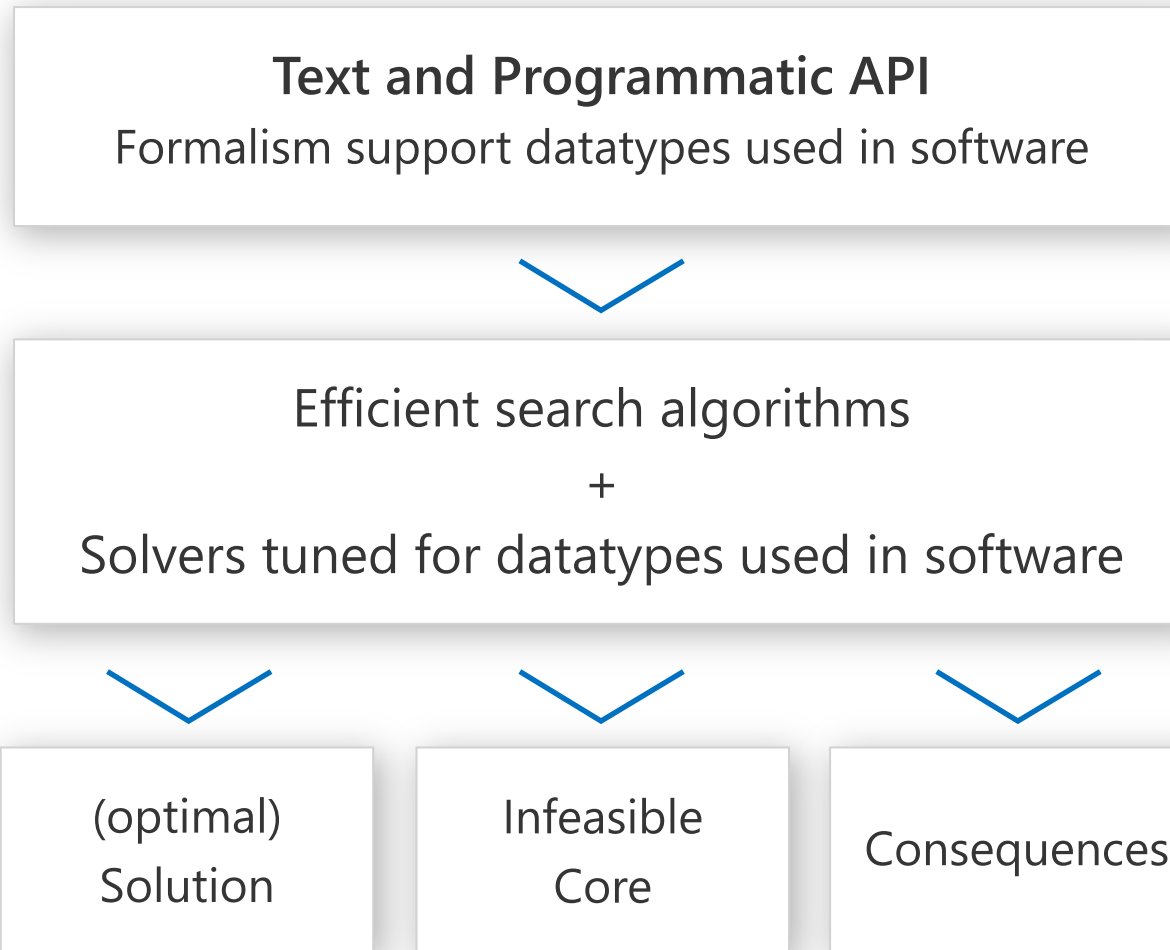
Logic

Computation




Claim: Practically all modern program analysis tools involve solving logical formulas


Z3 Efficient Solver for Symbolic Logic



Symbolic Solving: Foundations and Engineering

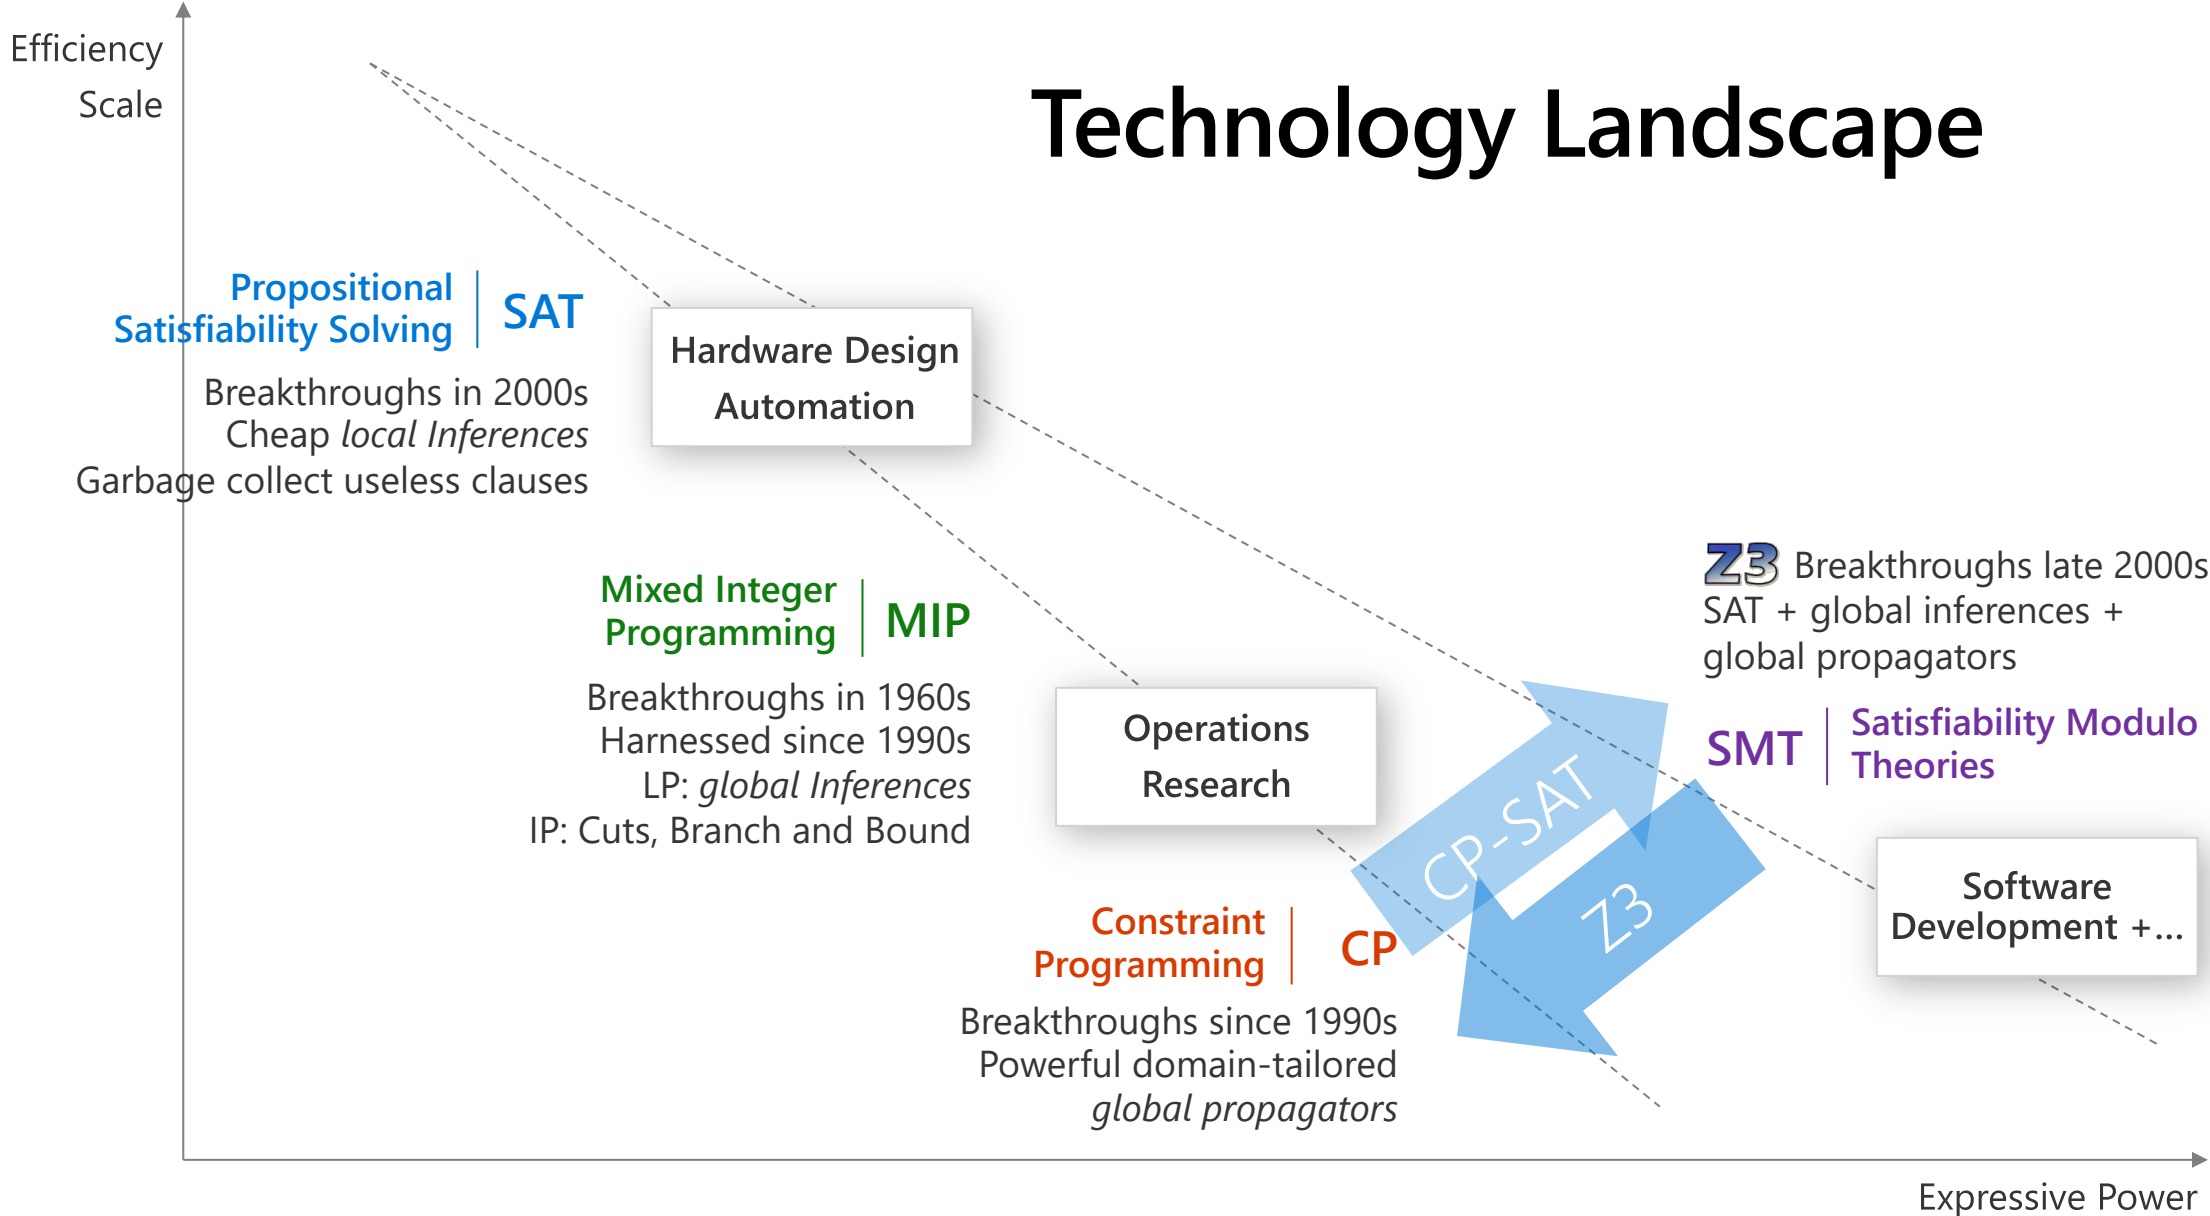


Core solvers in Z3 founded on duality between solution witness / infeasibility proof



Advances guided by application domains and evaluated by extensive benchmarking

Technology Landscape



SMT Solvers – Main Services

Support domains that are natural in Software and Hardware analysis

- Make it easy to translate program assertions into SMT

Holy grail of SMT: modularity + efficiency

- Combine disjoint Theory Solvers by reconciling equalities between shared variables

Quantifier-Free First Order Theories

- Int, Real, Bit-vectors, IEEE floating point numbers, arrays, algebraic data-types

Quantified First-Order, Higher-Order Logics

- As aid to proof assistants

Base theory: Uninterpreted Functions



$$a = f(f(a)), \quad a = f(f(f(a))), \quad a \neq f(a)$$

$$a = v_2, a = v_3, a \neq v_1,$$

$$v_1 \equiv f(a), v_2 \equiv f(v_1), v_3 \equiv f(v_2)$$

- Produce Proofs
- Incremental Updates
- Propagate Literals


Step 1: Equivalence classes from equalities



Union Find

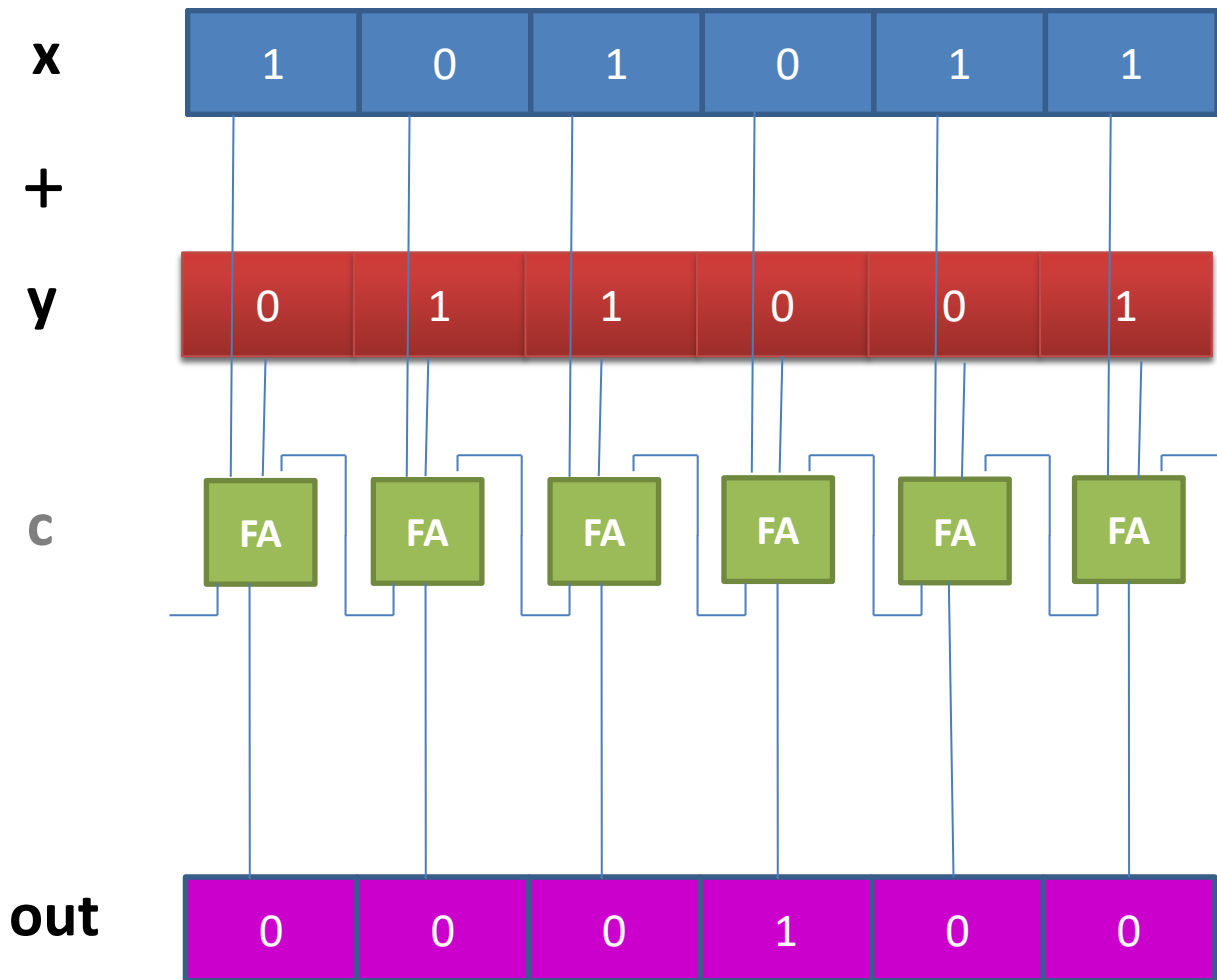
Step 2: Apply Congruence Rule:

$$a \simeq v_2 \text{ implies } f(a) \simeq f(v_2): \quad v_1 \simeq v_3$$



E-graph

Compiled into SAT: Bit-vectors



Bit-vector addition is expressible using bit-wise operations and bit-vector equalities.

$$\begin{aligned} \text{out} &= \text{xor}(\mathbf{x}, \mathbf{y}, \mathbf{c}) \\ \mathbf{c}' &= (\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{c}) \vee (\mathbf{y} \wedge \mathbf{c}) \\ \mathbf{c}[0] &= 0 \\ \mathbf{c}'[N-2:0] &= \mathbf{c}[N-1:1] \end{aligned}$$

Benefits:

- Efficient finite domain reasoning

Limitations:

- Not suitable for heavy use of linear arithmetic
- Bit-vector multiplication is super expensive

Note:



$$\begin{aligned} \text{out} &\leftrightarrow \text{xor}(\mathbf{x}, \mathbf{y}, \mathbf{c}) \\ \mathbf{c}' &\leftrightarrow (\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{c}) \vee (\mathbf{y} \wedge \mathbf{c}) \end{aligned}$$

Combining Theories in the age of CDCL(T)

Foundations

1979 Nelson, Oppen: Framework

1996 Tinelli & Harindi: N.O Fix

2000 Barrett et al: N.O + Rewriting

2002 Zarba & Manna: "Nice" Theories

2004 Ghilardi et al: N.O. Generalized

Efficiency using rewriting

1984 Shostak: Theory solvers

1996 Cyrluk et al: Shostak Fix #1

1998 B: Shostak with Constraints

2001 Rueß & Shankar: Shostak Fix #2

2004 Ranise et al: N.O + Superposition



1998 de Silva, Sakallah; 2001 Moskewicz et al: DPLL → CDCL made guessing cheap

2006 Bruttomesso et al: Delayed Theory Combination

2007 de Moura & B: Model-based Theory Combination

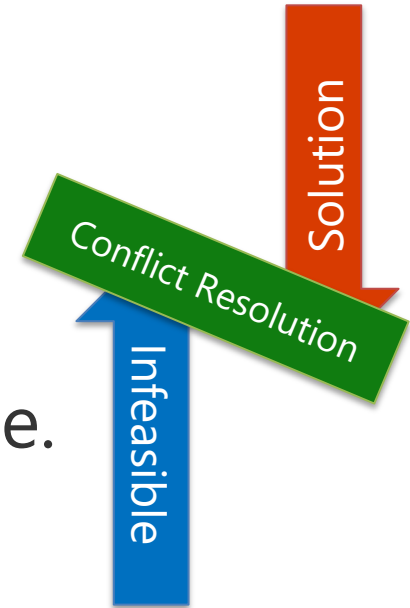
Model-based theory combination

Pre-existing methods

- Propagate all implied equalities – complicated costly.
- Delayed theory combination – $O(n^2)$ equalities, " 2^{n^2} " time.

Model-based theory combination

- Each theory constructs a candidate model.
- Propagate all equalities implied by candidate model, hedging that other theories will agree.
- If not, use backtracking to fix the model.



Propagating Equalities

Asserted inequalities

$$x + u \leq z \quad \frac{x + 1 = y}{z - 1 \leq y} \quad \frac{y - 1 = z}{x \leq z} \quad 1 \leq u \leq 1$$

Equality inferences require addition/subtraction operations

How the solver sees the constraints

$$x + u + s_1 = z \quad z - 1 + s_2 = y \quad y + s_3 = x \quad 1 \leq u \leq 1 \quad 0 \leq s_1 \quad 0 \leq s_2 \quad 0 \leq s_3$$

After pivoting

$$x = z - u - s_1 \quad y = z - 1 + s_2 \quad s_3 = -s_2 - u + 1 - s_1 \quad 1 \leq u \leq 1 \quad 0 \leq s_1 \quad 0 \leq s_2 \quad 0 \leq s_3$$

After propagating bounds on s_1, s_2, s_3

$$x = z - u - s_1 \quad y = z - 1 + s_2 \quad s_3 = -s_2 - u + 1 - s_1 \quad 1 \leq u \leq 1 \quad 0 \leq s_1 \leq 0 \quad 0 \leq s_2 \leq 0 \quad 0 \leq s_3 \leq 0$$

Subtract first two equalities to infer


$$x = y$$

Subtle complexity: Every row can have many fixed variables. Adding values of constant bounds requires significant runtime.

Propagating Equalities - Efficiently

$$\begin{array}{ccccccccc} x = z - u - s_1 & y = z - 1 + s_2 & 1 \leq u \leq 1 & 0 \leq s_1 \leq 0 & 0 \leq s_2 \leq 0 & 0 \leq s_3 \leq 0 & & & \\ \hline & & & x = y & & & & & \end{array}$$

Instead of adding up rows to *prove* implied equality



Use fact: all variables have values assigned by Simplex solver

Then two variables are equal if

- They are connected through offset equalities
- They have the same value

offset equality

$$x = y + a_1 z_1 + a_2 z_2 + \dots$$


$$\underbrace{b_1 \leq z_1 \leq b_1, b_2 \leq z_2 \leq b_2, \dots}$$

Example: If solver assigns $x = 3$ then $z = 4, y = 3$

x, y are connected (over $z - 1$).

x, y have the same value (3)


Z3 for Software + ...




Azure Network
Verification



Verifying C
Compiler



Quantum
Compilation



Biological
Computations




Verified Crypto
Libraries & Protocols



Dynamics
AX



SVACE
Static Analysis Engines



Artificial
Life




Security Risk
Detection



Smart Contract
Verification



ALIVE2
Translation Validation
for LLVM & Visual C++

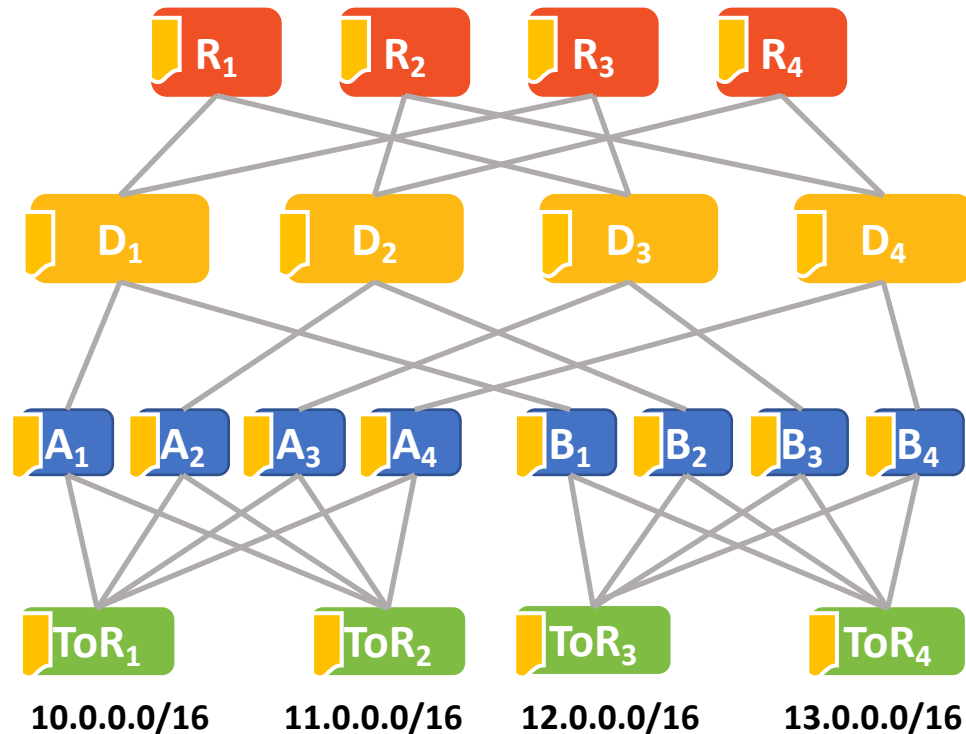


Axiomatic
Economics



NFL
Scheduling

Live Monitoring of Forwarding Behavior



Global reachability as **local contracts**




- ✓ Each router has a fixed rule for a set of addresses
- ✓ Enough to verify rule is enforced on each router

5 Billion Z3 queries per day

[Jayaraman et al, Sigcomm 2019]

Alive2: Integration with LLVM



```
clang w/ alive2 plugin:  
$ alivecc file.c
```

```
opt plugin:  
$ opt -tv -instcombine -tv file.ll
```

[Nuno Lopes, Juneyoung Lee, Chung-Kil Hur, Zhengyang Liu, John Regehr, PLDI 2021]

Tools and internals developed
in a feedback loop

Tools

BOOGIE

HAVOC

SLAYER



SLAM SAGE

TERMINATOR

VccYogi

FORMULA
Modeling Foundations.



SecGuru

Efficient E-matching for SMT solvers

Relevancy Propagation

Model-based Theory Combination

Effectively Propositional Logic

Engineering DPLL(T) + Saturation

Generalized, Efficient Array Decision Procedures

Model Based Quantifier Instantiation

Linear Quantifier Elimination

Quantified Bit-Vectors

CutSAT: Linear Integer Formulas

Model Constructing SAT

Existential Reals

Generalized PDR

μ Z: Datalog

vZ: Opt+MaxSMT

SLS, floats

Z3 Internals

Z3 for Virtual Plant Design Automation

An ongoing collaboration

Solving Virtual Plants (in a nutshell)

Solve for:

- Assign every task to a station and an operator

Subject to:

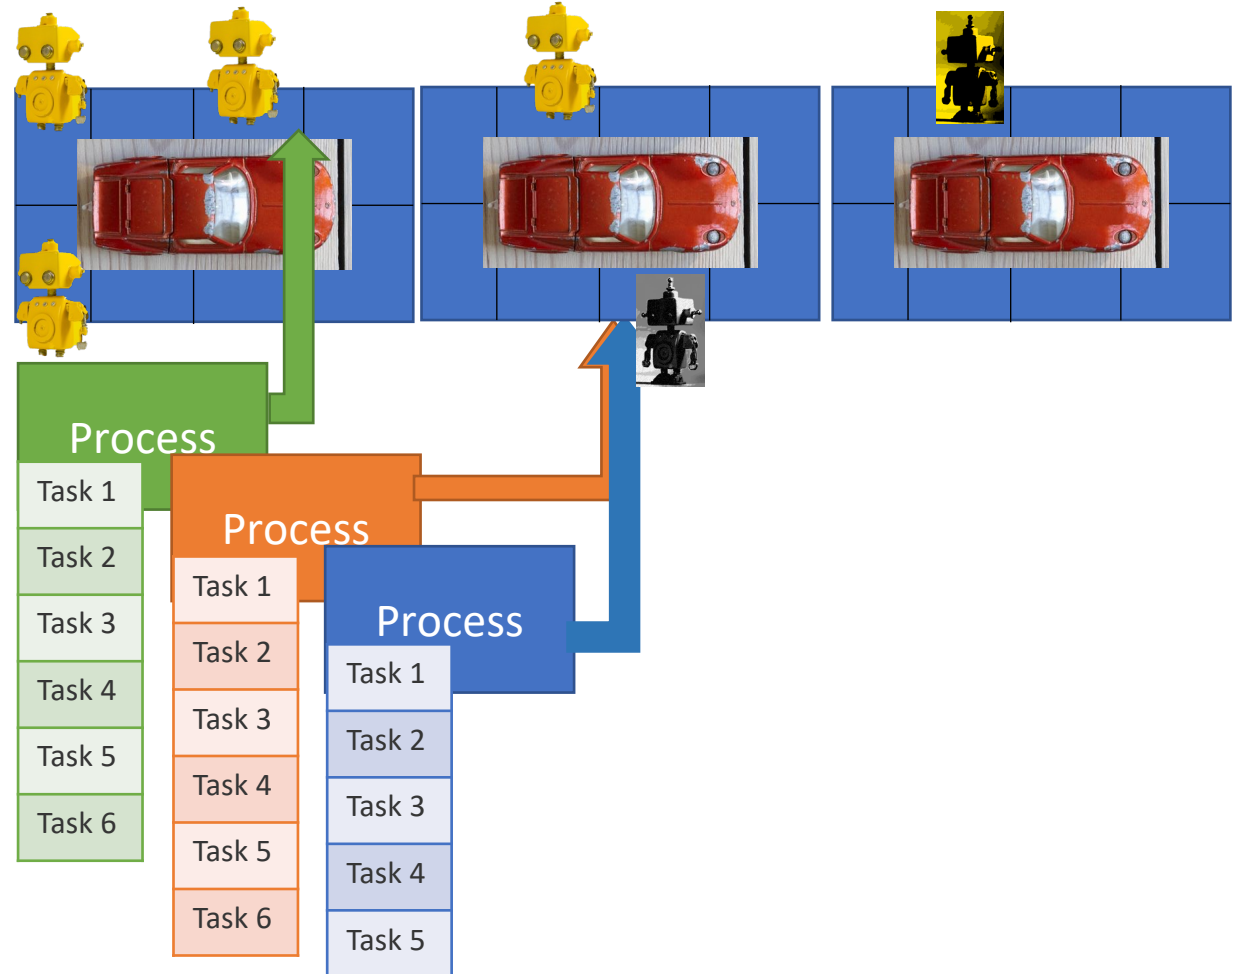
- Bounded completion time
- Partial order of stations and processes
- What operations stations can perform

Objectives:

- Minimize resource consumption
- Minimize operator congestion

Enable:

- Automate manual puzzle
- Optimize over design space
- Scale and be nimble: new factories, new models
- Track and manage inventory



Experiences Summary

Domain Engineering

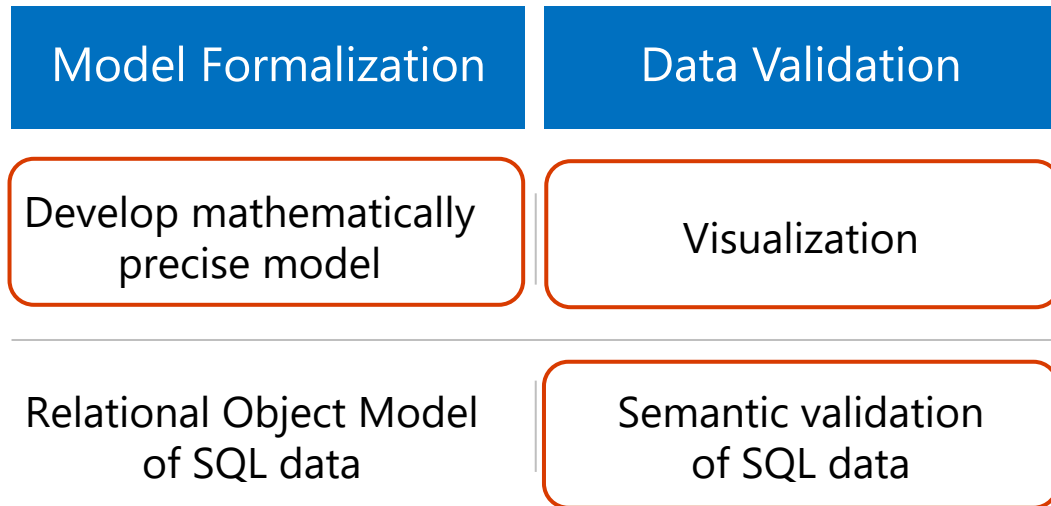
Model Formalization	Data Validation
Develop mathematically precise model	Visualization
Relational Object Model of SQL data	Semantic validation of SQL data

Solver Engineering

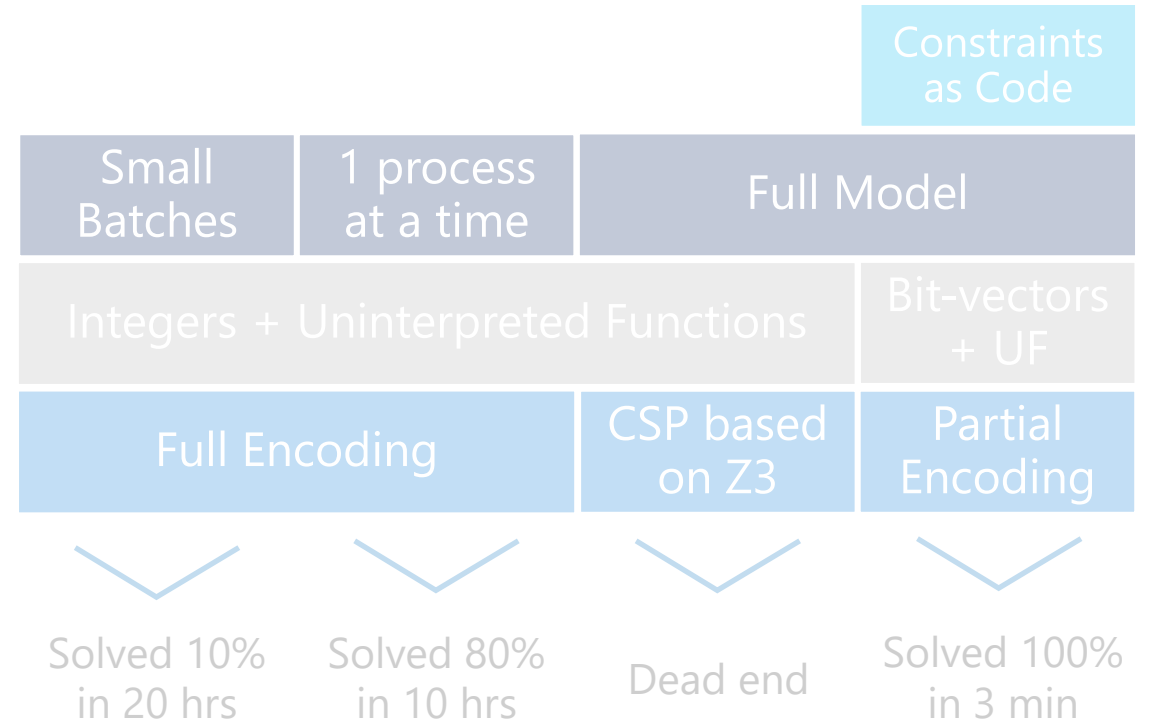
		Constraints as Code	
Small Batches	1 process at a time	Full Model	
Integers + Uninterpreted Functions		Bit-vectors + UF	
Full Encoding		CSP based on Z3	Partial Encoding
Solved 10% in 20 hrs	Solved 80% in 10 hrs	Dead end	Solved 100% in 3 min

Experiences Summary

Domain Engineering



Solver Engineering



Domain Engineering – Mathematical Modeling

Solve for:

$assign_p: Station \rightarrow$ each process p

Auxiliary Functions:

$maxHeight: Station \rightarrow Nat$

$operator: Station \times Zone \rightarrow Operator$


Assignment Constraints:

$operator(assign_p, z) \in \{op_1, op_2\}$

$maxHeight(assign_p) \geq height_p$

Mapping to Z3 at *same* the level of model

- Uninterpreted functions
- Nested formulas (no tuning for big Ms)
- Finite domains using bit-vectors




Research Angles

- Mathematical concepts matured in tools for requirements capture (TLA+)
- A sweet spot for Formal Methods skillsets
- Uncovered many subtle implicit assumptions

Domain Engineering – Semantic Validation

Grand Goal



Deep Solving

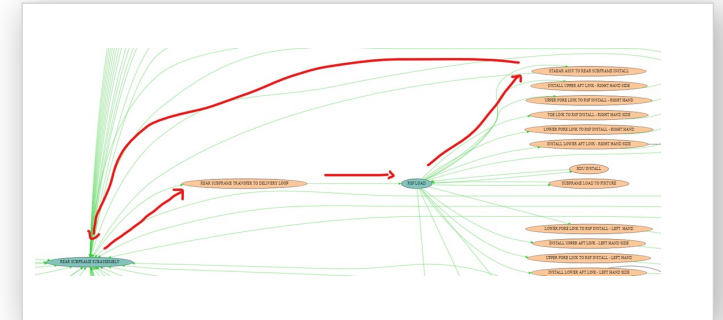
Humble Path



Deep Validation

Validation Approaches

1. Scripts that check invariants of database entries
2. Provenance information using infeasible cores from Z3




Z3 Features

- Native core minimization in SAT solver
- Core and correction set enumeration
- Software bug localization and repair


Practical impact

- Invariant checker and provenance tools in hands of collaborators
- Used to fix a significant set of data entry bugs

Domain Engineering – Visualization



Aid to understand model stored in database and spot bugs by simple inspection



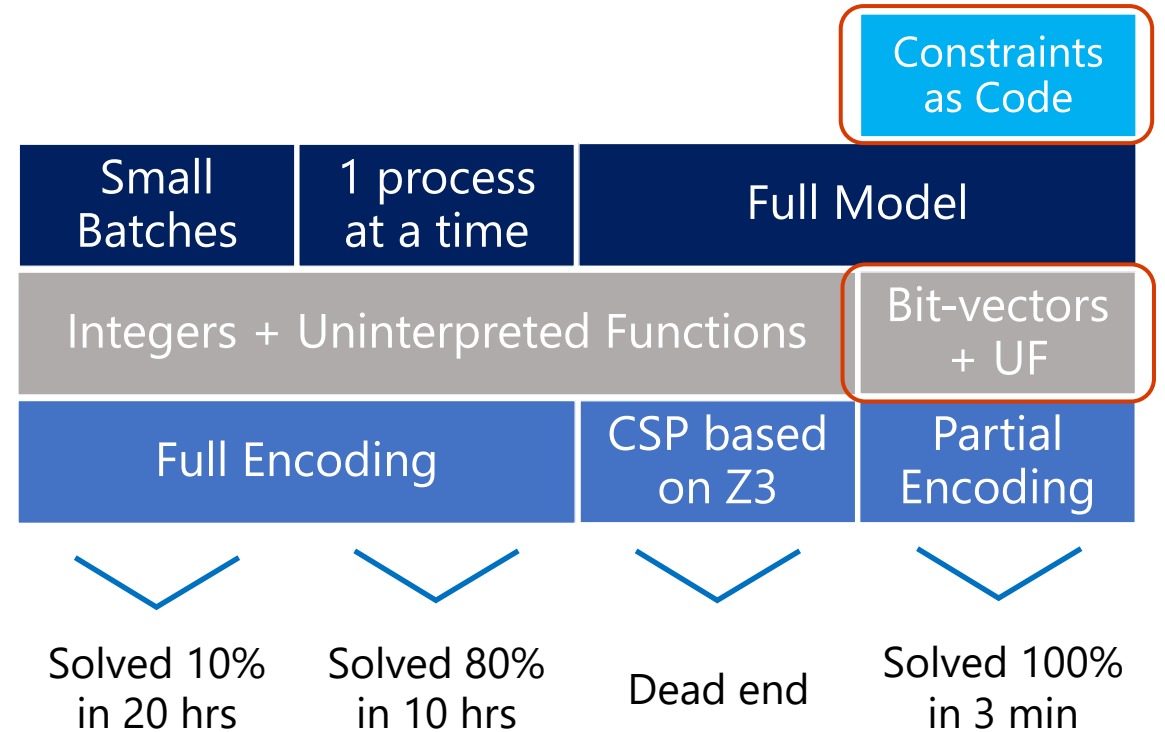
MSAGL – Automated Graph Layout engine

Experiences Summary

Domain Engineering

Model Formalization	Data Validation
Develop mathematically precise model	Visualization
Relational Object Model of SQL data	Semantic validation of SQL data

Solver Engineering





A Fly in the Ointment and a Wasp in the Rose



Domain Overload

Number of Processes = $O(1K)$

Number of Stations = $O(1K)$

Number of Tasks = $O(10K)$

Up to $O(10)$ different operators per station

Direct MIP-style encoding: $t_{i,s,op}$ - Task i is at station s using operator op

$$10K \times 1K \times 10 = 100M \text{ variables}$$

Our approach: Use uninterpreted functions for “symbolic indices”

Constraint Overload

Cycle Time

Cycle times for each station s and $op \in s.operators$:

$$time(p, z) := \sum \{ t.time \mid t \in p.tasks \wedge t.zone = z \}$$

$$preTime(p, z) := \sum \{ t.time \mid t \in p.preTasks \wedge t.zone = z \}$$

$$postTime(p, z) := \sum \{ t.time \mid t \in p.postTasks \wedge t.zone = z \}$$

$$Full := \{ time(p, z) \mid \neg isSplit(p) \wedge station(p) = s \wedge op = wzOp(station(p), z) \}$$

$$Pre := \{ preTime(p, z) \mid isSplit(p) \wedge station(p) = s \wedge op = wzOp(station(p), z) \}$$

$$Post := \{ postTime(p, z) \mid isSplit(p) \wedge station(p) + 1 = s \wedge op = wzOp(station(p), z) \}$$

$$\sum Full + \sum Pre + \sum Post \leq s.time$$

a polluting, nasty side constraint

Comprehension Full, Pre, Post is over $p \in Process, z \in \{t.zone \mid t \in p.tasks\}$.

Constrained multi-knapsack:

A set of items, each is added to one knapsack, subject to side-constraints

Our approach: program constraint as an ad-hoc theory

Solver Engineering – Mathematical Modeling

Solve for:

$assign_p: Station \rightarrow$ each process p

Auxiliary Functions:

$maxHeight: Station \rightarrow Nat$

$operator: Station \times Zone \rightarrow Operator$

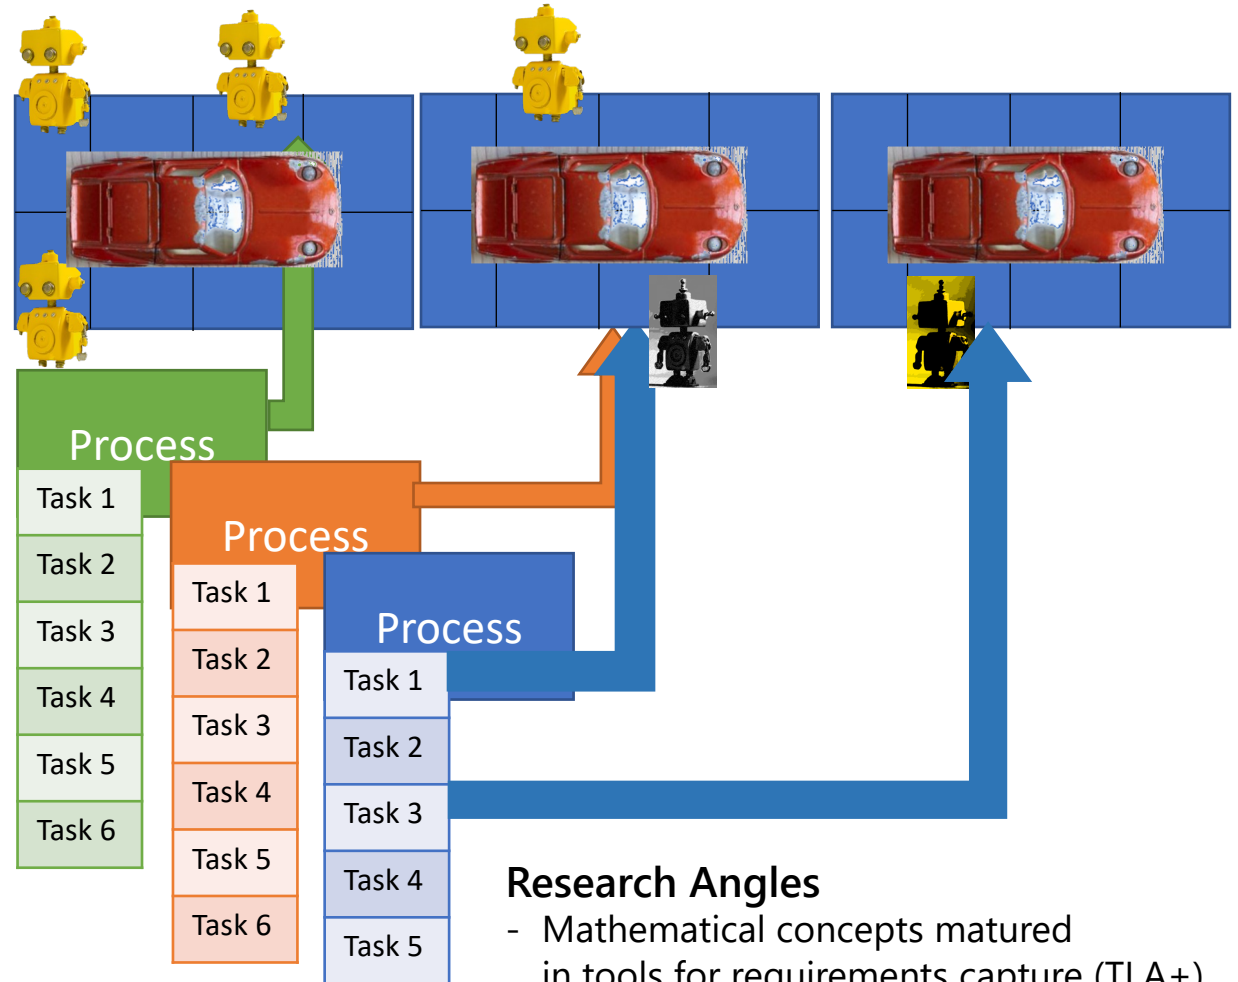
Assignment Constraints:

$operator(assign_p, z) \in \{op_1, op_2\}$

$maxHeight(assign_p) \geq height_p$

Mapping to Z3 at *same* the level of model

- Uninterpreted functions
- Nested formulas (no tuning for big Ms)
- Finite domains using bit-vectors



Research Angles

- Mathematical concepts matured in tools for requirements capture (TLA+)
- A sweet spot for Formal Methods skillsets
- Uncovered many subtle implicit assumptions

Solver Engineering – Mathematical Modeling

Solve for:

$assign_p: Station \rightarrow Station$ each process p

Auxiliary Functions:

$maxHeight: Station \rightarrow Nat$

$operator: Station \times Zone \rightarrow Operator$

Assignment Constraints:

$operator(assign_p, z) \in \{op_1, op_2\}$

$maxHeight(assign_p) \geq height_p$

Mapping to Z3 at *same* the level of model

- Uninterpreted functions
- Nested formulas (no tuning for big Ms)
- Finite domains using bit-vectors

Z3 solves for *functions* not just (integer/real) *variables*

Allows succinct encodings of constraints

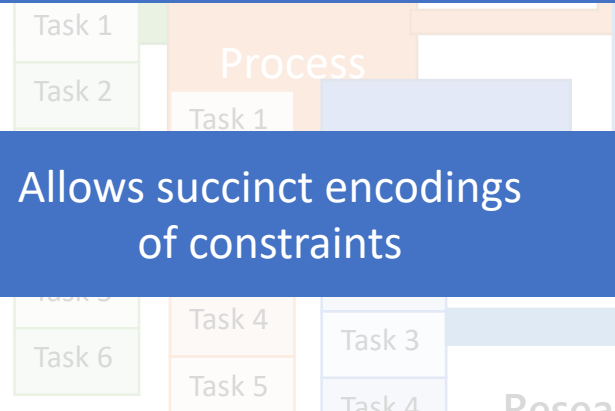

Bit-vectors map to SAT solver technology
Integers map to MIP solver technology

Maps to specialized solver
Incremental Congruence Closure

$$\begin{aligned} x &= f(g(f(x))) \\ \frac{x &= g(f(x))}{x &= g(x)} \end{aligned}$$


Research Angles

- Concepts matured
- Implementations capture (TLA+)
- Formal Methods skillsets
- Uncovered many subtle implicit assumptions



Solver Engineering – Constraints as Code

When supported formalisms lag, or encoding is impractical



```
def on_x_is_fixed_to_value_v(self, x, v):
    old_sum = self.sum
    self.trail.append(lambda : self.undo(old_sum, x))
    self.sum += len(w for w in self.xvalues.values() if v + w == 42 and (v > 30 or w > 30))
    self.xvalues[x] = v
    if self.sum > 1000:
        self.conflict([self.x2id[x] for x in self.xvalues])
```

Solving Strategy

v1: Pre-solving

- Assign a small batch of 4-10 processes to stations at a time
- Stations and processes, station heights and task time are integers
- Solved 100 processes very slowly.

v2: Pre-solving take two

- Assign just one process at a time and only encode process constraints when a process gets assigned.
- The resulting solver can assign 950 out of 1050 processes in a few minutes.

v3: A custom CDCL / CSP solver

- Perform branching and propagation of cycle time constraints on top of repeated calls to Z3.
- Maintain backtracking stack and add lemmas based on the chosen branches.
- This was complex to engineer and only exercised in preliminary form.

v4: with Custom Propagator and Bit-vectors

- With bit-vectors, without cycle-time: solvable in 30 seconds.
- With bit-vectors and cycle-time: solvable for 300 processes in a few minutes, but not all processes.
- With bit-vectors, programmable-propagator for cycle-time: patching + solving
- Initial: a few hours
- Current: a few minutes.

SMT for OR?

- Already happened: CP-SAT uses CDCL(T) for OR domains
- Approach here: Uninterpreted Functions, Bit-Vectors, Constraints as Code
- From experiences to tuning:
 - LNS for Modulo Theories?
 - A modernized core solver for Z3: In-processing for SMT?
 - Sound MIP is too costly for CP: Specialized LP for modular machine arithmetic

Summary

Z3 – an efficient SMT solver

The screenshot shows the GitHub repository for the Z3 Theorem Prover, highlighting various components like 'cmake', 'contrib', 'doc', 'examples', 'nomi', 'NSO', 'SCRIPT', 'src', '.doc', '.git', '.gitignore', '.gitmodules', '.travis', 'CMake', 'LICENSE', and 'README.cmake.md'. To the right, a flow diagram illustrates the solver's architecture: source code (C++, Python, .Net, Java, Ocaml) is processed through C and SMTLIB2, then through Preprocessing (Tactics: Cube & Conquer) and Solvers (SMT, Fixedpoint, NLSat, SAT, QSAT), leading to Optimization.

Theories

Congruence Closure

$$x = f(g(f(x)))$$

$$\frac{x = g(f(x))}{x = g(x)}$$

Constraints as Code

```
def on_x_is_fixed_to_value_v(self, x, v):
    old_sum = self.sum
    self.trail.append(lambda : self.undo(old_sum, x))
    self.sum += len(w for w in self.xvalues.values() if v + w == 42 and (v > 30 or w > 30))
    self.xvalues[x] = v
    if self.sum > 1000:
        self.conflict([self.x2id[x] for x in self.xvalues])
```

Bit-vectors

```
def encoding(self):
    line_bits = math.ceil(math.log(len(self.model.lines), 2))
    self.Line = BitVecSort(line_bits)
    self.Station = BitVecSort(station_bits)
    self.Operator = BitVecSort(operator_bits)
    self.Segment = BitVecSort(segment_bits)
    self.op_used = Function('op_used', self.Station, self.Operator, BoolSort()) # Is operator used
    self.wz_used = Function('wz_used', self.Station, self.Zone, BoolSort()) # Which workzones used
    self.min_height = Function('min_height', self.Station, self.Height) # Min height of station

...
yield Implies(p.is_split, self.min_height(p.to_station + 1) <= min_height), E.suf_height_lo(p, min_height)
yield Implies(p.is_split, self.max_height(p.to_station + 1) >= max_height), E.suf_height_hi(p, max_height)
```

Experiences

Domain Engineering

Model Formalization	Data Validation
Develop mathematically precise model	Visualization
Relational Object Model of SQL data	Semantic validation of SQL data

Solver Engineering

		Constraints as Code	
Small Batches	1 process at a time	Full Model	
Integers + Uninterpreted Functions		Bit-vectors + UF	
Full Encoding	CSP based on Z3	Partial Encoding	
Solved 10% in 20 hrs	Solved 80% in 10 hrs	Dead end	Solved 100% in 3 min

Humble Path



Data Validation

Fly in the Ointment



Nimble Constraints


Extra Slides

Tools used as part of collaboration

Contact at
Manufacturer



Azure SQL



Shared VM



Azure Bastion

Snapshot



Requirements



Azure Repos

Object Model




python

Constraints

Z3



Weekly + ad-hoc sync



Microsoft
Automated
Graph Layout

Visualization

Optimization Objectives

Currently at early stage

Understanding what best serves our scenario

Likely main objective

Reduce number of operators, reduce number of tools used overall.

First approach is by programmable Branch & Bounding to find a Pareto Front per run.

Research Angles

- Pareto Strategies
- Any-time optimization
- Local Neighborhood search
- MaxSAT based on:
 - Cores
 - Hitting sets
 - Correction sets
 - Branch and bound

Z3 Technologies

- Core based MaxSAT
- Primal Simplex
- Multi-objective optimization: Pareto, Lex, Box

Some years ago

Used Azure cloud scaling (cube & conquer) and large neighborhood search to optimize NFL schedules

Z3 – an efficient SMT solver

master 5 branches 20 tags

Go to file Add file Code

About

The Z3 Theorem Prover

Readme View license

Releases 20

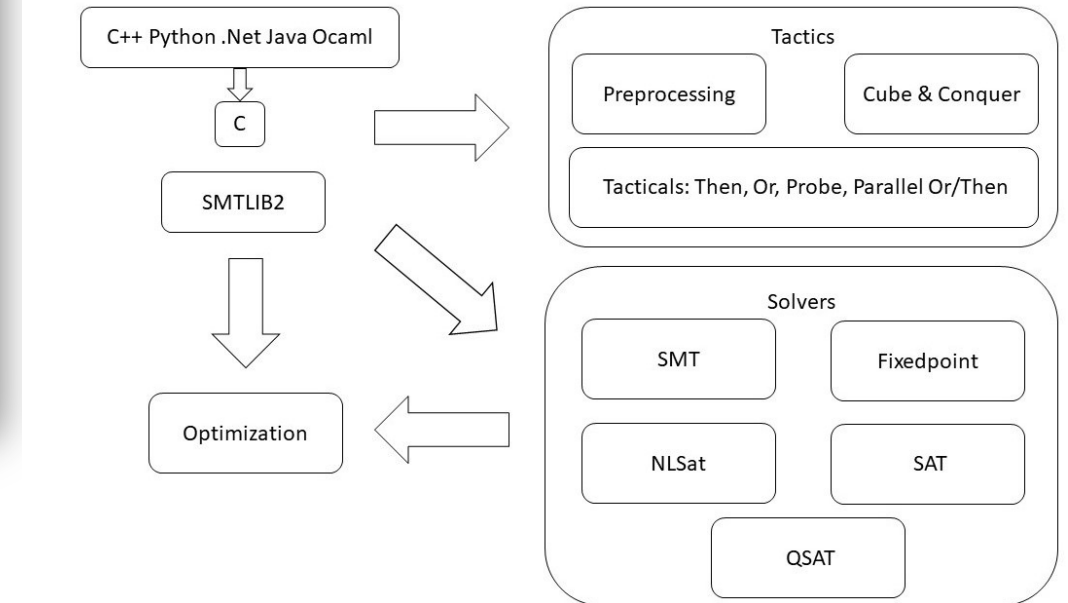
z3-4.8.9 Latest on Sep 10

+ 19 releases

NikolajBjorner add recfuns to Java #4820 3 hours ago 14,647 commits

cmake	regex pattern per #2986	10 months ago
contrib	modular Axiom Profiler (#4619)	4 months ago
doc	align readme-cmake and cmakefiles.txt according to current state #2732	12 months ago
examples	c++ example: call Z3_finalize_memory() so that the buildbot leak chec...	last month
noar	cmd_context fix #4790	14 days ago
reso	math add a comment in nla_order	15 days ago
scrip	model fix #4812	3 days ago
src	muz adding dt-solver (#4739)	last month
	nlsat pass algebraic manager to arith-plugin mk-numeral because rational ch...	4 months ago
	opt add recfuns to Java #4820	4 hours ago
	params fix #4808	4 days ago
	parsers DRAT debugging updates	3 days ago
	qe redo purification	27 days ago
	sat add recfuns to Java #4820	4 hours ago
	shell DRAT debugging updates	3 days ago
	smt z3str3: reject certain unhandled expressions (#4818)	4 hours ago
	solver debug arith/mbi	23 days ago
	tactic include order	3 days ago
	test adding dt-solver (#4739)	last month

README-CMake.md Use Z3_option prefix in cmake with Java bindings build command (#4612) 4 months ago



Congruence Closure

$$\frac{x = f(g(f(x)))}{x = g(x)}$$

Constraints as Code

```
def on_x_is_fixed_to_value_v(self, x, v):
    old_sum = self.sum
    self.trail.append(lambda : self.undo(old_sum, x))
    self.sum += len(w for w in self.xvalues.values() if v + w == 42 and (v > 30 or w > 30))
    self.xvalues[x] = v
    if self.sum > 1000:
        self.conflict([self.x2id[x] for x in self.xvalues])
```

Bit-vectors

```
def encoding(self):|
line_bits      = math.ceil(math.log(len(self.model.lines), 2))
self.Line      = BitVecSort(line_bits)
self.Station   = BitVecSort(station_bits)
self.Operator  = BitVecSort(operator_bits)
self.Segment   = BitVecSort(segment_bits)
self.op_used   = Function( 'op_used',    self.Station, self.Operator, BoolSort()    ) # Is operator used
self.wz_used   = Function( 'wz_used',    self.Station, self.Zone,      BoolSort()    ) # Which workzones used
self.min_height = Function( 'min_height', self.Station, self.Height   ) # Min height of station
...
yield Implies(p.is_split, self.min_height(p.to_station + 1) <= min_height), E.suf_height_lo(p, min_height)
yield Implies(p.is_split, self.max_height(p.to_station + 1) >= max_height), E.suf_height_hi(p, max_height)
```