What is a “Good” Ruleset?

Comparing Rulesets Using Equality Saturation

Cynthia Richey, gannet@cs.washington.edu

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

PLSE

- ➔ An **e-graph** is a data structure that represents sets of equivalent terms.
- ➔ **Equality saturation**: a technique that uses an e-graph to apply rewrite rules *non-destructively*.
- ➔ For equality saturation to be effective, a compact, useful set of rewrite rules (a “ruleset”) is necessary.
- ➔ ***How do we decide if a ruleset is good?***

A rule, *r* *(x → y)*, is **derivable** by a ruleset *R* if *x* and *y* are determined to be equivalent following the application of rules in *R*.

- ➔ We can define different versions of derivability by changing the initial state of the e-graph.
- ➔ Resource limits make a big difference!

Domain	R1 → R2 (lhs, lhs-rhs)	R2 → R1
bool	100%, 100%	87.5%, 96.9%
bv4	100%, 100%	38.3%, 41.1%
bv32	100%, 100%	58.3%, 60.0%
rational	97.3%, 100%	52.0%, 58.5%

LHS

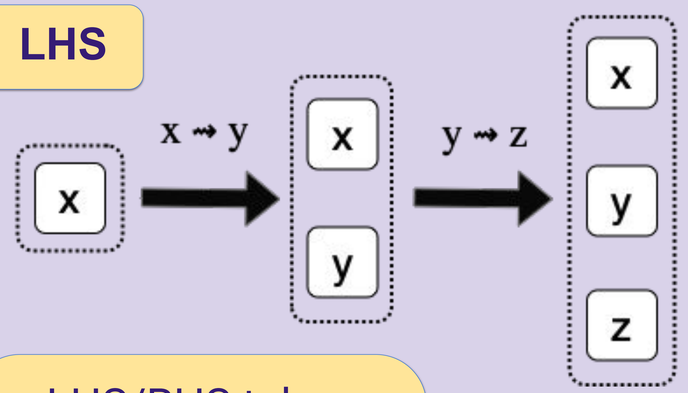

LHS/RHS

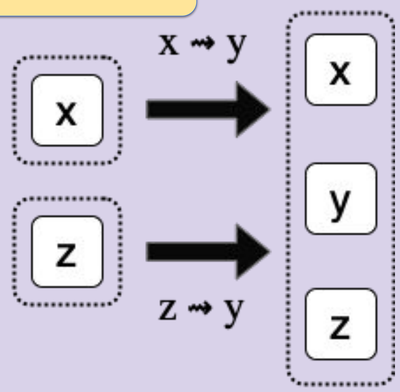

LHS/RHS takes fewer iterations of equality saturation! … but performance is only half the story.

Initializing the e-graph with the left *and* right-hand side of *r* results in y → x deriving x → y; using just the LHS, we can *never* derive it!

LHS

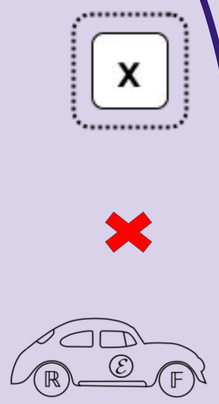

LHS/RHS

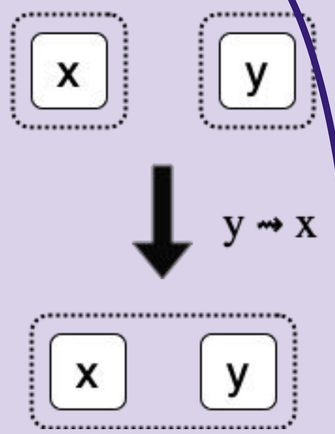

Good for optimization!

Good for equivalence-checking!

Comparing rulesets using LHS & LHS/RHS!