

LEOPARD: Accelerating Cloud-based Access Control Policy Verification Using Logical Encoding Optimization

Xing Fang¹, Feiyan Ding¹, Mingyuan Song¹, Yuntao Zhao¹, Lizhao You¹,
Qiao Xiang¹, Linghe Kong², Jiwu Shu^{1, 3}, Xue Liu⁴,

¹ Xiamen University, ² Shanghai Jiao Tong University, ³ Minjiang University, ⁴ McGill University

2025/07/03



Access Control Policy Configuration Errors Are Common

Misconfigured Azure Blob Storage Exposed the Data of 65K Companies and 548K Users

According to SOCRadar, “the amount and scale of the leaked data make it the most significant B2B data leak in the recent history of cybersecurity.”

McGraw Hill's S3 bucket exposed students' grades and other data

Educator gets an F for security

Salesforce cloud outage caused by security change

By [Richard Chirgwin](#)

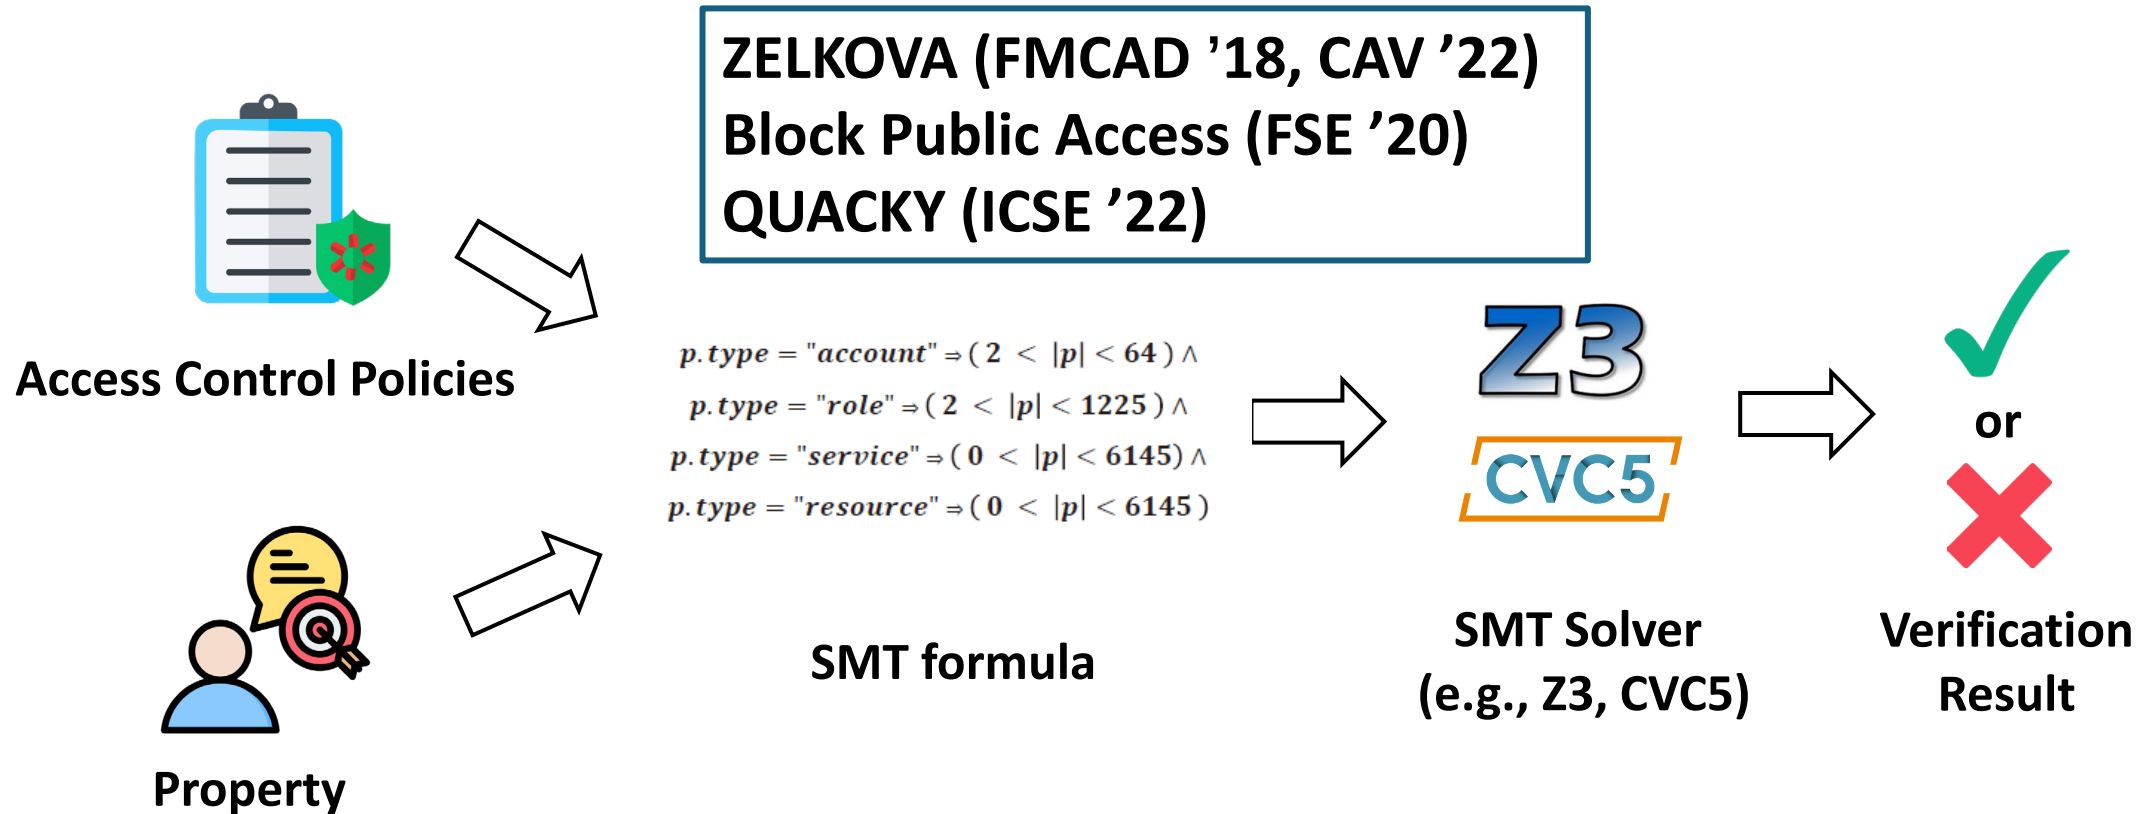
Sep 27 2023 12:34PM



October 20, 2022

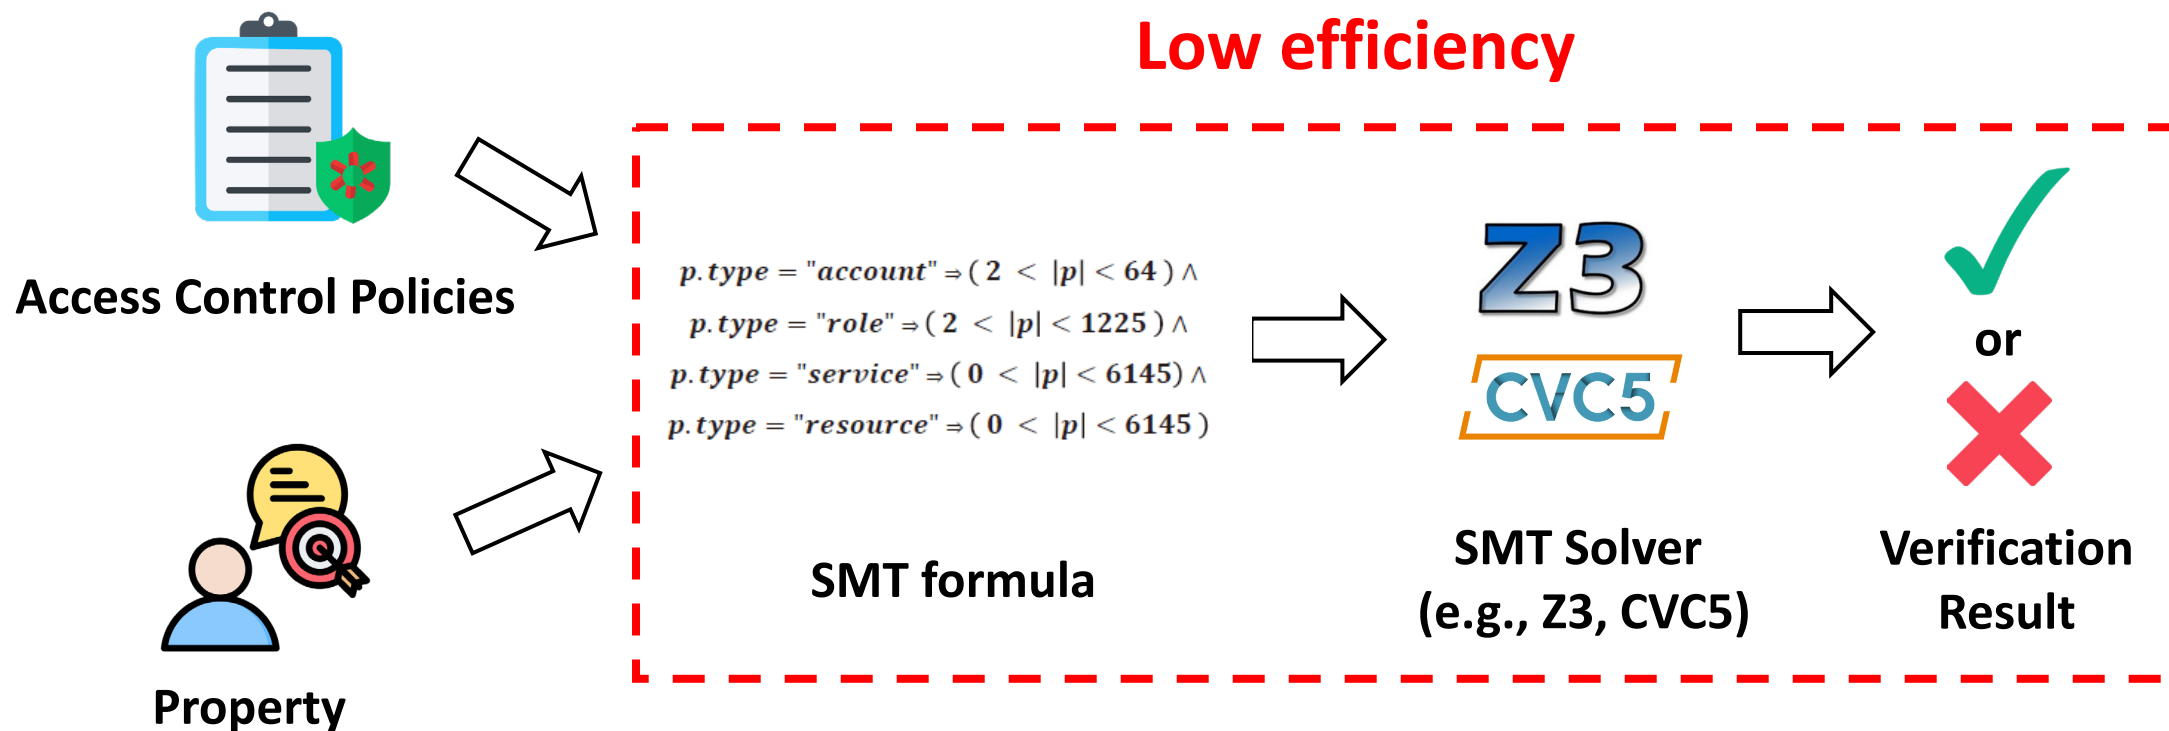
Alibaba Cloud outage sees apps taken offline

Policy Verification Helps Prevent Misconfigurations



Existing works use *SMT-based* verification to check whether access control policies satisfy desired properties.

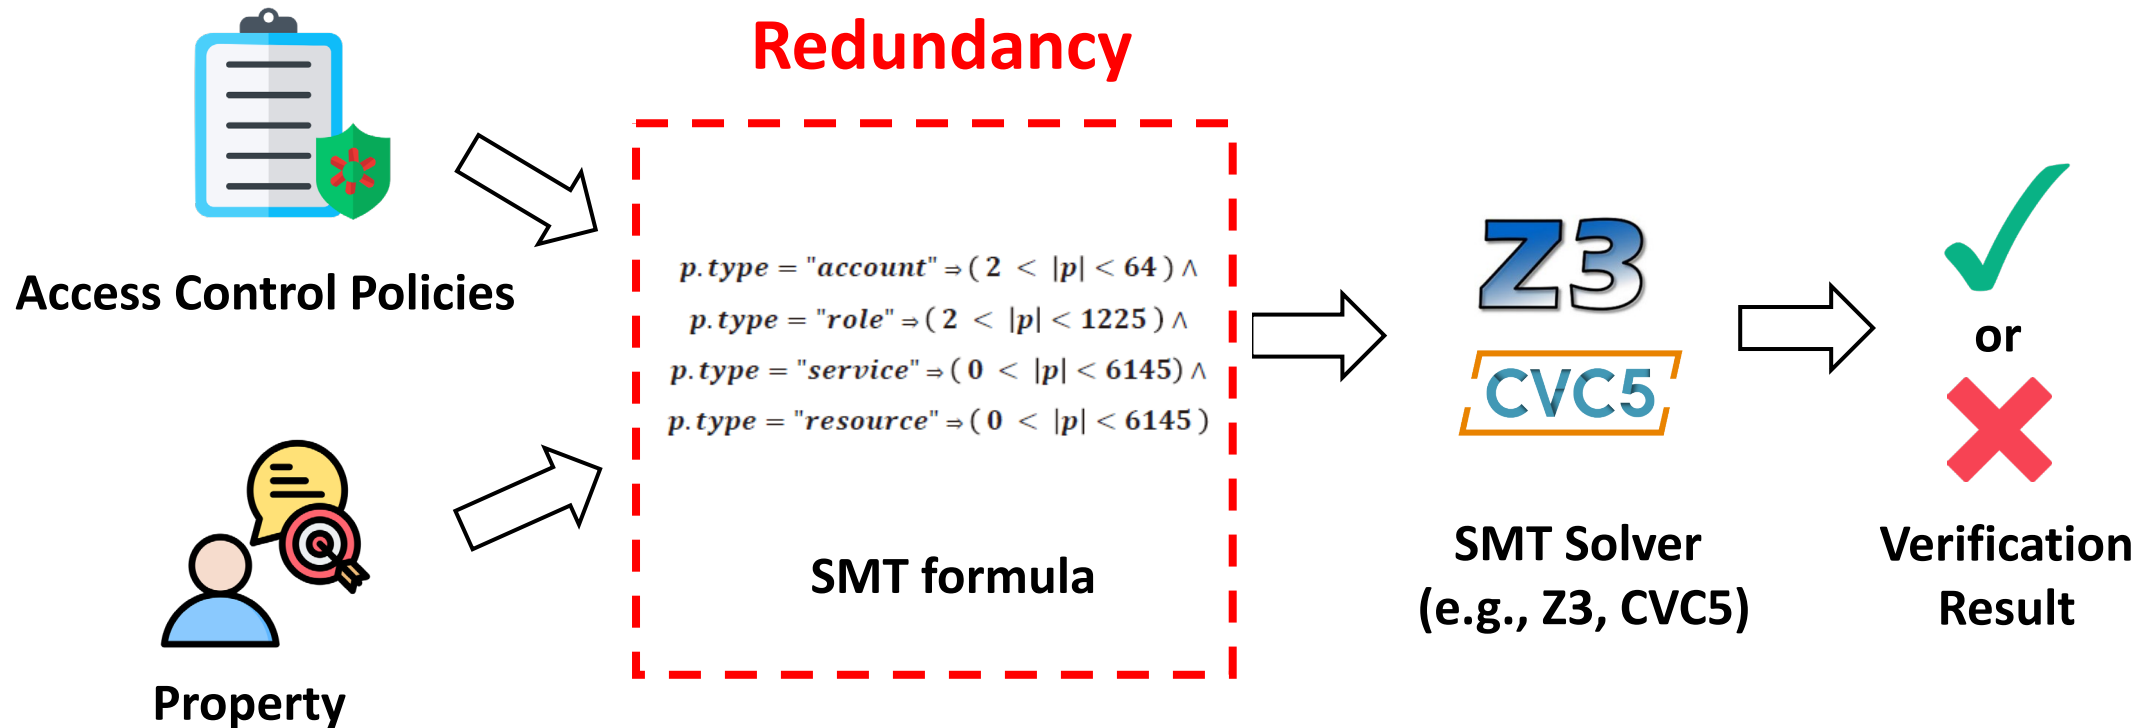
Challenge: SMT Verification at Scale



Existing tools lack the efficiency needed to support billions of SMT queries per day^[1].

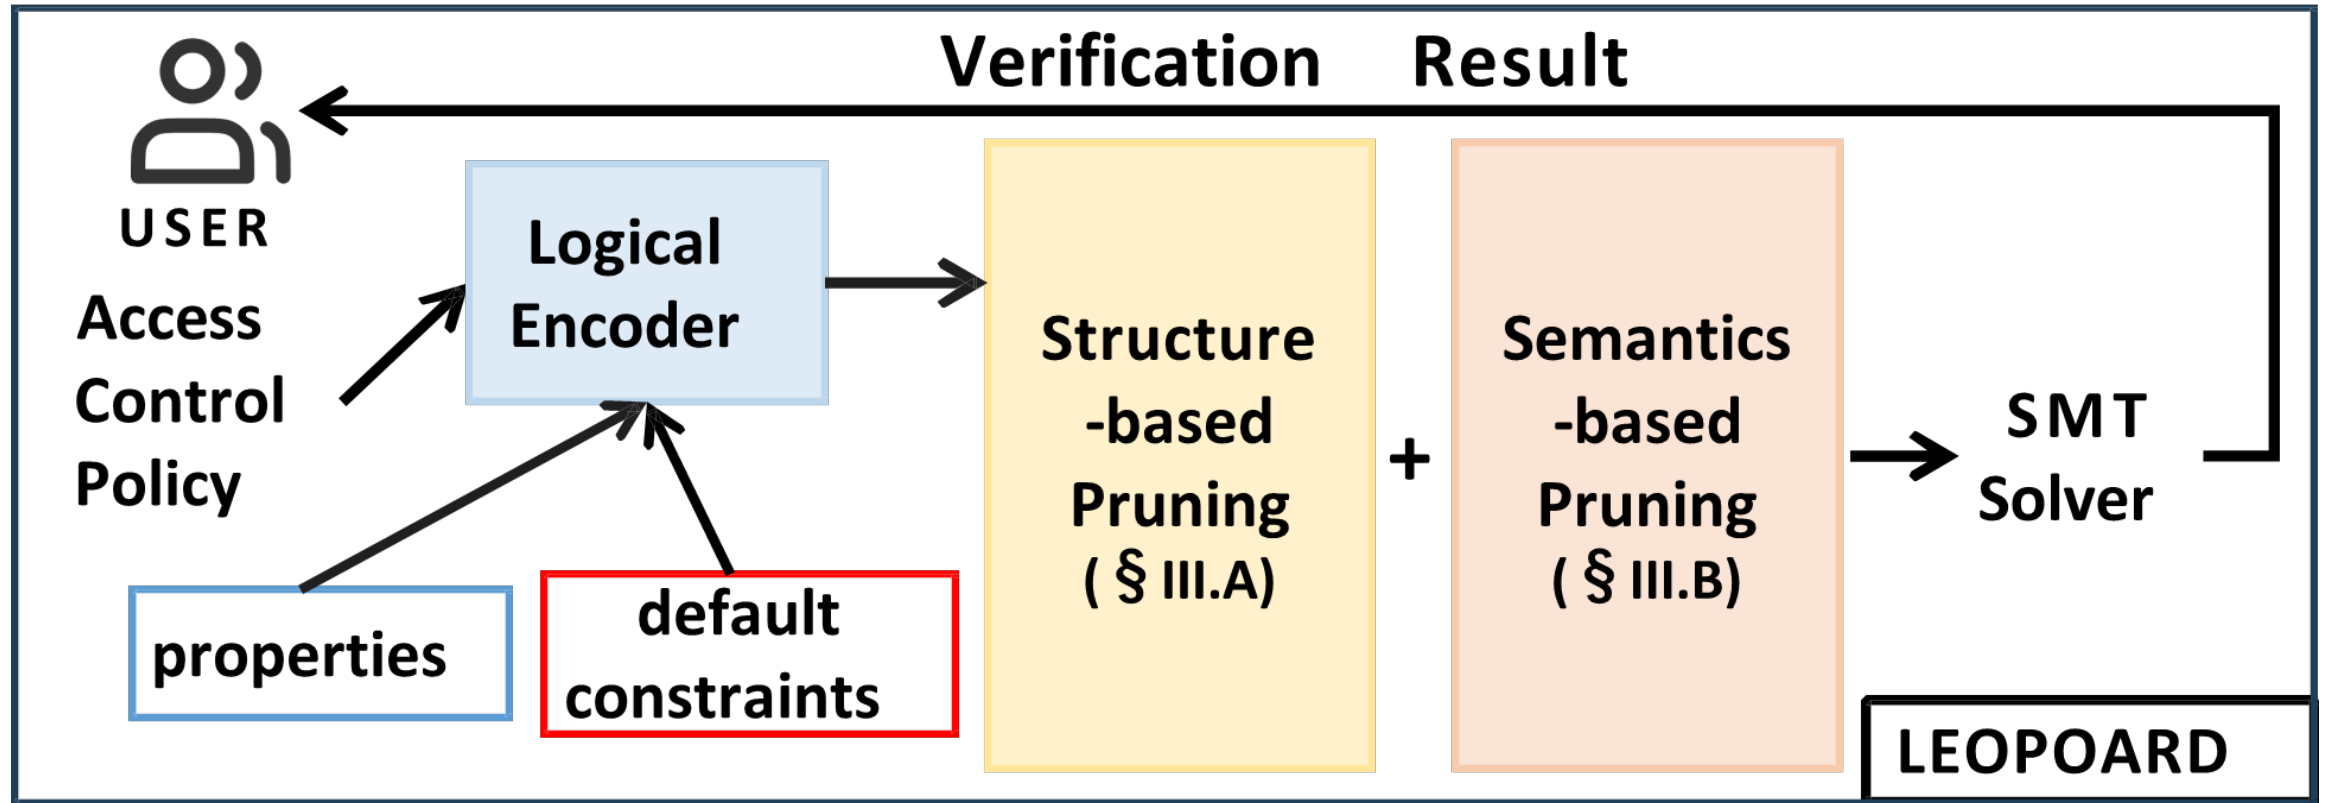
[1] Rungta, N. "A Billion SMT Queries a Day (Invited Paper)". CAV'22

Motivation: SMT Formula Encoding Redundancy



Existing works *encode redundant constraints* without considering their relevance to the verification property.

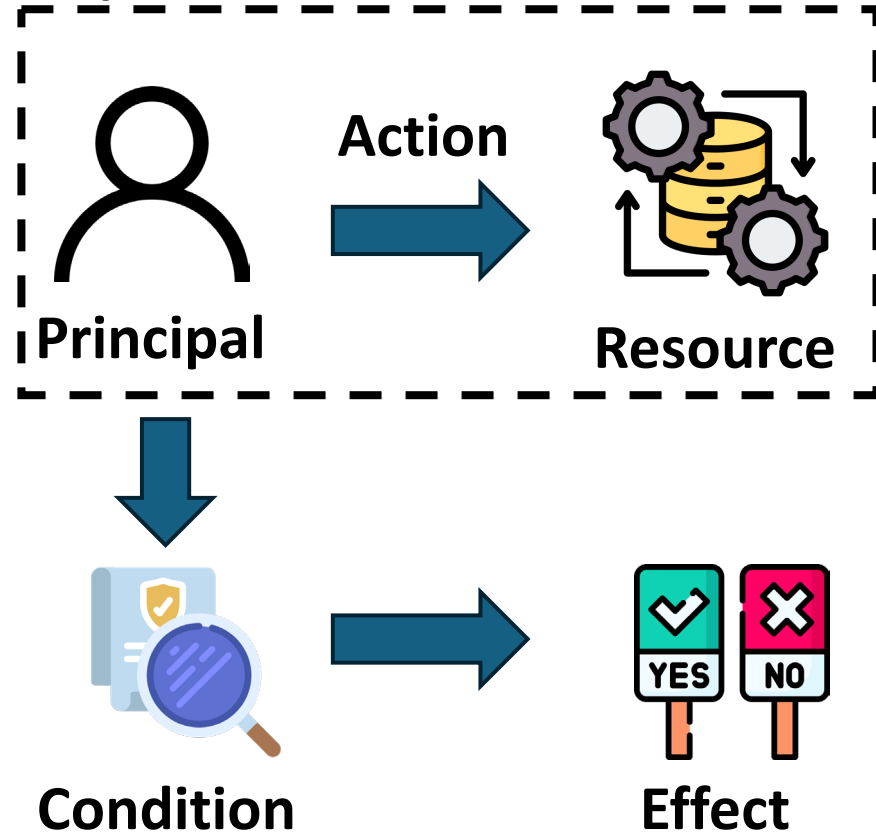
Our Approach: LEOPARD



Optimizing logical encoding to accelerate policy verification.

Cloud-based Access Control Policy

Request



- Effect: **Allow or Deny** the request
- Principal: **Who** makes the request
- Action: **What operation** is performed
- Resource: **Which resource** is targeted
- Condition: Under **what conditions** it applies

Access control policies define who can perform which actions on which resources.

Cloud-based Access Control Policy

```
{{Allow
  Principal : "*",
  Action    : "GetObject",
  Resource  : "examplebucket/*",
  Condition: ""},
(Deny
  Principal : "*",
  Action    : "GetObject",
  Resource  : "examplebucket/*",
  Condition: (StringNotLike, UserId, EXAMPLEID:*))}}
```

Policy X:

Allow all users to get objects from the bucket

Deny users whose UserId does not match
EXAMPLEID:* to get objects from the bucket

```
{{Allow
  Principal : "*",
  Action    : "*",
  Resource  : "examplebucket/*",
  Condition: ""}}
```

Policy Y:

Allow all users to perform all actions on the
bucket

Policy Example

Cloud-based Access Control Policy

Actions	Resources
GetObject	examplebucket/*
GetBucketPolicy	
PutBucketAcl	
PutBucketCors	

Action Table Example

String	Length
UserId	(0, 64)
UserName	(0, 64)
AgencyId	(0, 64)
AgencyName	(0, 64)

String Table Example

Default constraints define which requests are considered valid by the cloud platform.

Logical Encoding for Access Control Policy

$$[P] = \left(\bigvee_{S \in Allow} [S] \right) \wedge \neg \left(\bigvee_{S \in Deny} [S] \right)$$

Policy Encoding:

A request is allowed if it satisfies any Allow statement and does not match any Deny statement.

$$[S] = \left(\bigvee_{v \in S(P)} p = v \right) \wedge \left(\bigvee_{v \in S(A)} a = v \right) \wedge \left(\bigvee_{v \in S(R)} r = v \right) \wedge \left(\bigwedge_{O \in S(C)} O \right).$$

Statement Encoding:

A request matches a statement if the principal, action, and resource each match, and all conditions are satisfied.

Logical Encoding for Access Control Policy

```
{(Allow
  Principal : "*",
  Action    : "GetObject",
  Resource   : "examplebucket/*",
  Condition  : ""),
 (Deny
  Principal : "*",
  Action    : "GetObject",
  Resource   : "examplebucket/*",
  Condition  : (StringNotLike, UserId, EXAMPLEID:*))}
```

Policy X

$$\begin{aligned} & a = \text{"GetObject"} \wedge \\ & \text{"examplebucket/" prefixOf } r \wedge \\ & \neg(a = \text{"GetObject"} \wedge \\ & \text{"examplebucket/" prefixOf } r \wedge \\ & \wedge \neg(\text{"EXAMPLEID:" prefixOf } UserId)) \end{aligned}$$

SMT encoding of Policy X

```
{(Allow
  Principal : "*",
  Action    : "*",
  Resource   : "examplebucket/*",
  Condition  : "")}
```

Policy Y

$$\text{"examplebucket/" prefixOf } r$$

SMT encoding of Policy Y

Logical Encoding for Default Constraints

$$[D_A] = \bigvee_{(T_A, T_R) \in \text{ActionTable}} \left(\left(\bigvee_{v \in T_A} a = v \right) \wedge \left(\bigvee_{v \in T_R} r = v \right) \right)$$

Action Table Encoding:

An action–resource pair must match an entry in the Action Table.

$$[D_S] = \bigwedge_{(s, \min, \max) \in \text{StringTable}} \min \leq |s| \leq \max$$

String Table Encoding:

A string field must satisfy its type-specific length range.

Logical Encoding for Default Constraints

Actions	Resources
GetObject	examplebucket/*
GetBucketPolicy	
PutBucketAcl	
PutBucketCors	

Action Table Example


$$\begin{aligned} & ((a = \text{"GetObject"} \vee \\ & a = \text{"GetBucketPolicy"} \vee \\ & a = \text{"PutBucketAcl"} \vee \\ & a = \text{"PutBucketCors"}) \wedge \\ & \text{"examplebucket/" prefixOf } r) \end{aligned}$$

SMT encoding of Action Table

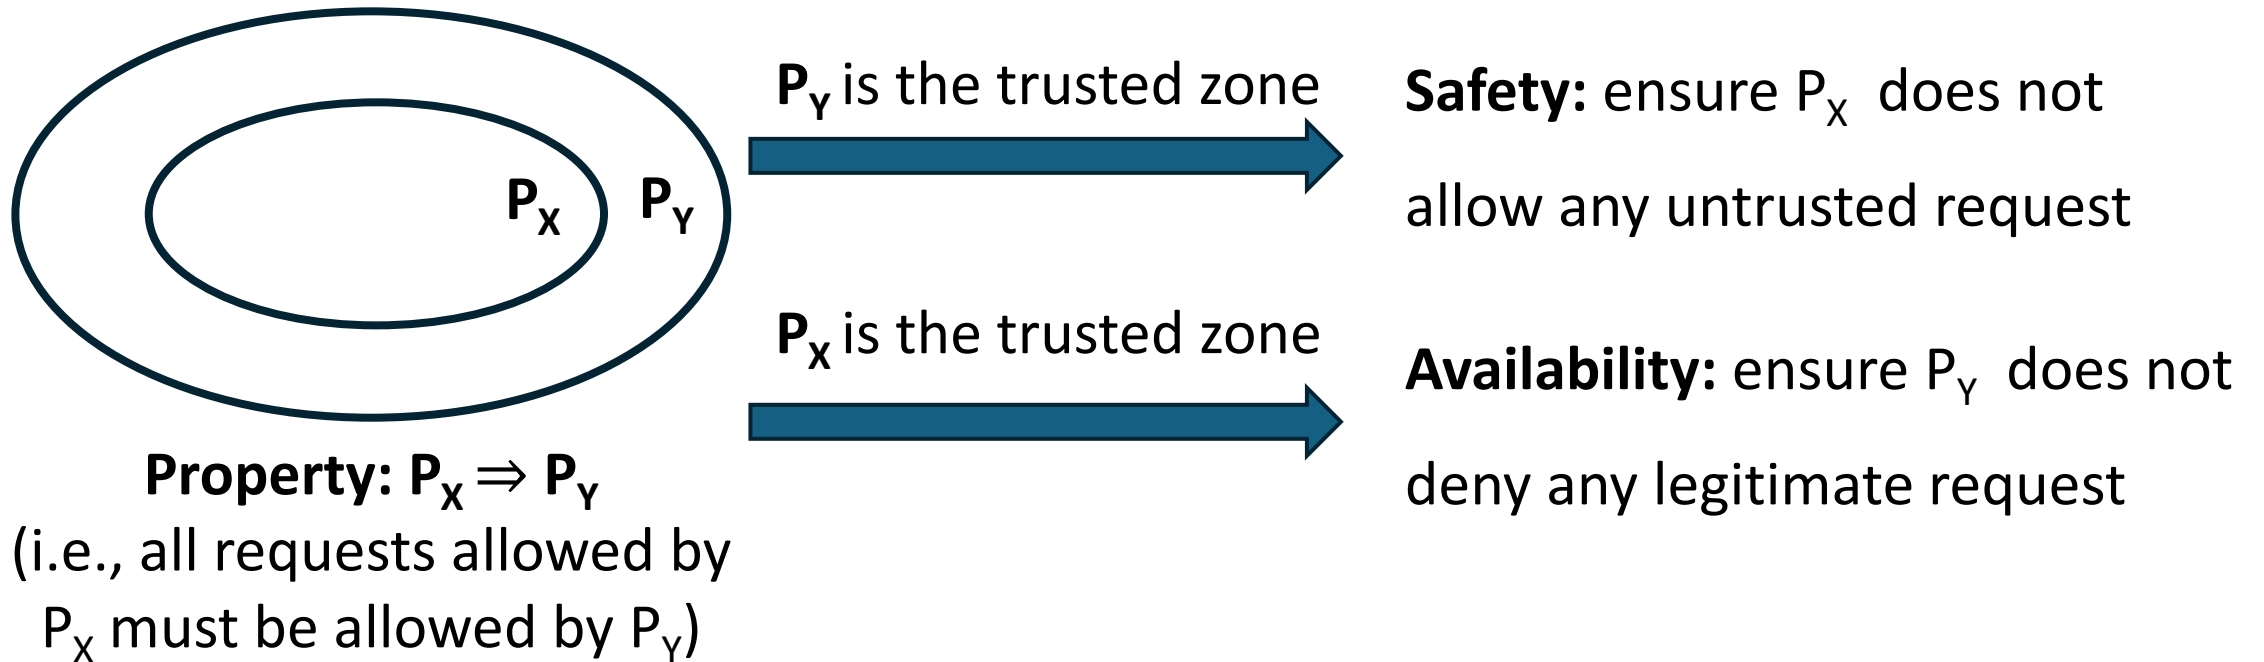
String	Length
UserId	(0, 64)
UserName	(0, 64)
AgencyId	(0, 64)
AgencyName	(0, 64)

String Table Example


$$\begin{aligned} & (0 < |UserId| < 64) \wedge \\ & (0 < |UserName| < 64) \wedge \\ & (0 < |AgencyId| < 64) \wedge \\ & (0 < |AgencyName| < 64) \end{aligned}$$

SMT encoding of String Table

Logical Encoding for Verification Problem



SMT formula: $([P_X] \wedge \neg[P_Y] \wedge [D])$

Result: If unsatisfiable, the property holds: $P_X \Rightarrow P_Y$.

Observation: Redundancy in Default Constraint Encoding

Property: $[P_X] \wedge \neg[P_Y]$

Default Constraint: $[D]$



Encode **all** constraints

SMT formula: $([P_X] \wedge \neg[P_Y] \wedge [D])$

Existing works encode redundant constraints into the SMT formula.

Structural Redundancy

$$\begin{aligned} & a = \text{"GetObject"} \wedge \\ & \text{"examplebucket/" } \textit{prefixOf} \ r \wedge \\ & \neg(a = \text{"GetObject"} \wedge \\ & \text{"examplebucket/" } \textit{prefixOf} \ r \wedge \\ & \wedge \neg(\text{"EXAMPLEID:" } \textit{prefixOf} \ \textit{UserId})) \end{aligned}$$

SMT encoding of Policy X

$$\text{"examplebucket/" } \textit{prefixOf} \ r$$

SMT encoding of Policy Y

Policy encoding

$$\begin{aligned} & (a = \text{"GetObject"} \vee \\ & a = \text{"GetBucketPolicy"} \vee \\ & a = \text{"PutBucketAcl"} \vee \\ & a = \text{"PutBucketCors"}) \wedge \\ & \text{"examplebucket/" } \textit{prefixOf} \ r) \wedge \\ & (0 < |\textit{UserId}| < 64) \wedge \\ & (0 < |\textit{UserName}| < 64) \wedge \\ & (0 < |\textit{AgencyId}| < 64) \wedge \\ & (0 < |\textit{AgencyName}| < 64) \end{aligned}$$

Default constraints encoding

Constraints are structurally redundant if they have no dependency on variables used in the verification property.

Semantic Redundancy

$a = \text{"GetObject"} \wedge$
 $\text{"examplebucket/" prefixOf } r \wedge$
 $\neg(a = \text{"GetObject"} \wedge$
 $\text{"examplebucket/" prefixOf } r \wedge$
 $\wedge \neg(\text{"EXAMPLEID:" prefixOf } UserId))$

SMT encoding of Policy X

$\text{"examplebucket/" prefixOf } r$

SMT encoding of Policy Y

Policy encoding

$(a = \text{"GetObject"} \vee$
 $a = \text{"GetBucketPolicy"} \vee$
 $a = \text{"PutBucketAcl"} \vee$
 $a = \text{"PutBucketCors"}) \wedge$
 $\text{"examplebucket/" prefixOf } r) \wedge$
 $(0 < |UserId| < 64) \wedge$
 $(0 < |UserName| < 64) \wedge$
 $(0 < |AgencyId| < 64) \wedge$
 $(0 < |AgencyName| < 64)$

Default constraints encoding

Constraints are semantically redundant if they are unsatisfiable under the variable values fixed by the verification property.


Structure-based Pruning

Algorithm 1 StrucutreBasedPruning(D , $Prop$)


Input: Formulas D , $Prop$

Output: Pruned formulas D_p

```
1: Initialize a queue  $Q$ 
2:  $Q.enqueue(Prop)$ 
3:  $D_p = \{\emptyset\}$ 
4: while NotEmpty( $Q$ ) do
5:    $f = Q.dequeue()$ 
6:   for formula  $d$  in  $D$  do
7:     if  $d$  and  $f$  share a variable and  $d \notin D_p$  then
8:        $Q.enqueue(d)$ 
9:        $D_p = D_p \cup d$ 
10: return  $D_p$ 
```



Property-related variables



Irrelevant variables

Perform dependency analysis to prune default constraints that have no dependency on variables used in the verification property.

Semantic-based Pruning

Algorithm 2 SemanticsBasedPruning(D , $Prop$)

Input: Formulas D_A , $Prop$

Output: pruned formulas D_p

```
1:  $D_p = \{\emptyset\}$ 
2: for formula  $d$  in  $D_A$  do
3:    $intersect = d.actions \cap Prop.allowedActions$ 
4:   if  $intersect \neq \emptyset$  then
5:      $d.actions = intersect$ 
6:      $D_p = D_p \cup d$ 
7: return  $D_p$ 
```



$a = \textit{"GetObject"}$

Allowed actions

$a = \textit{"GetBucketPolicy"}$

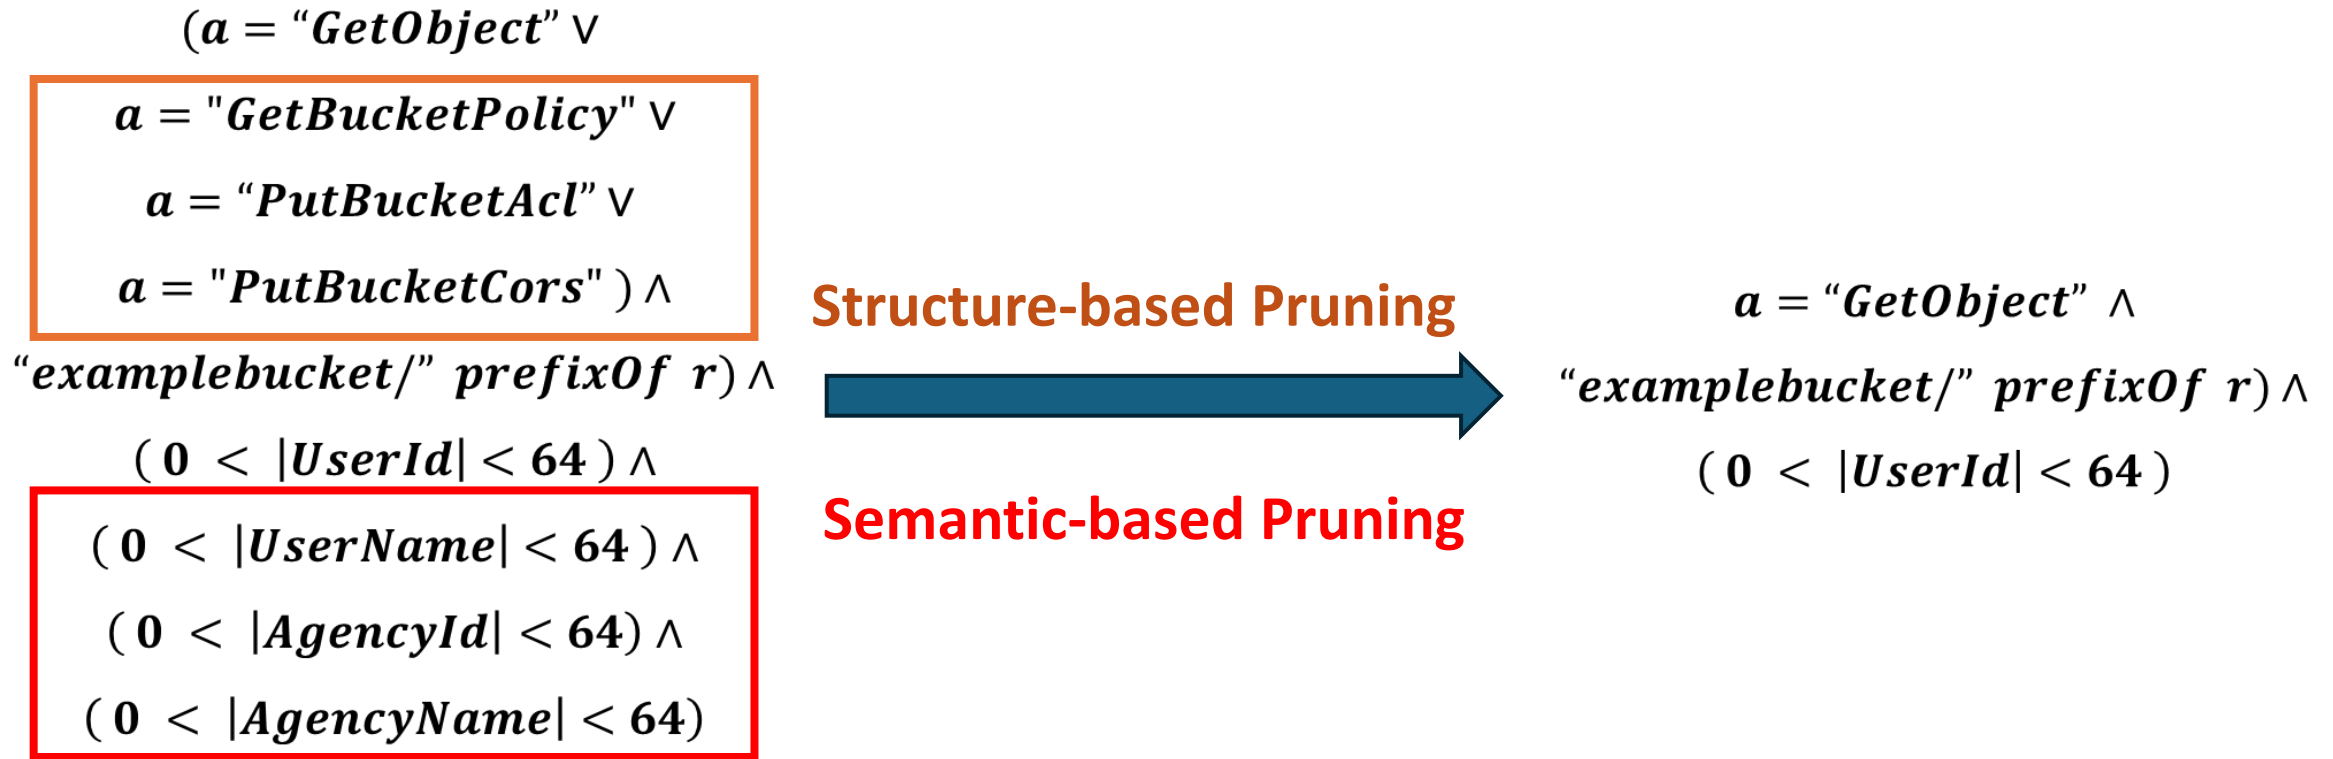
$a = \textit{"PutBucketAcl"}$

$a = \textit{"PutBucketCors"}$

Semantically invalid actions

Remove actions that are semantically invalid under the property's constraints.

Combined Pruning Yields Significant Simplification



The SMT formula is significantly simplified by pruning redundant constraints, while preserving correctness.

Evaluation

- Experiment Setting
 - **Datasets:** 2 synthesized datasets from AWS and a large cloud provider.
 - **Properties:** binary and quantitative analysis.
 - **Metric:** solving time before and after pruning.
 - **Baseline:** advanced SMT solvers (Z3, Z3str3RE, OSTRICH, ABC) and the state-of-the-art tool QUACKY for quantitative analysis.

Evaluation

TABLE I: The Number of the Original and Pruned Instances on Datasets.

Dataset	<u>#Instance</u>		
	ORIGINAL	LEOPARD	RATIO(%)
Dataset1	2792	1267	45%
Dataset2	3018	1836	61%

**A significant portion of formulas can be pruned:
45% in Dataset1 and 61% in Dataset2.**

Evaluation




Fig. 5: CDF of solving times for binary analysis of original and pruned SMT instances in Dataset1 across three solvers

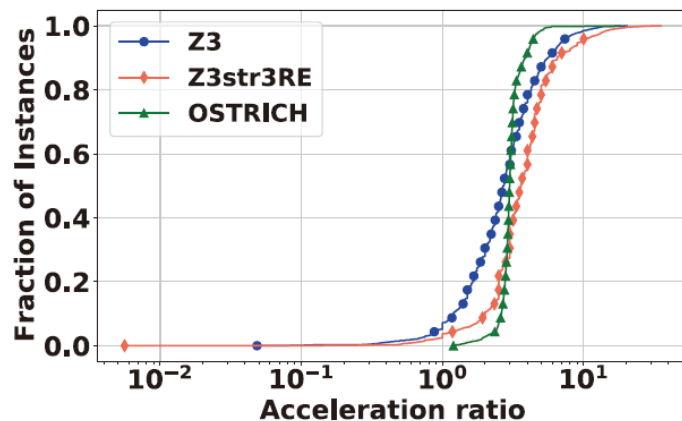


Fig. 6: The ratio of solving time for binary analysis of the original and pruned instances on Dataset1 for all solvers

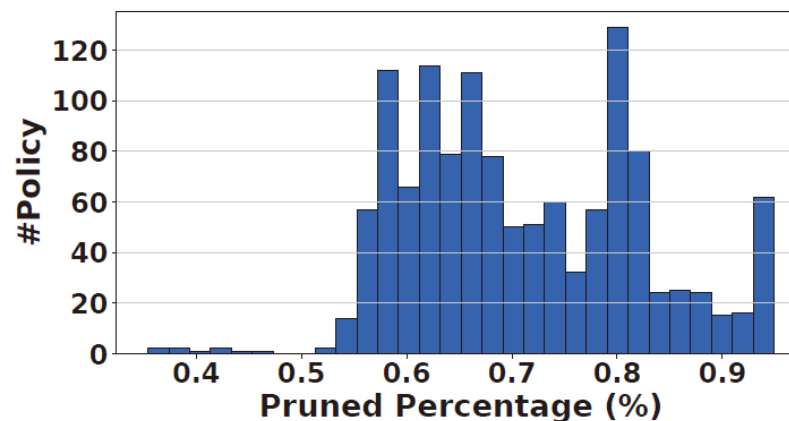


Fig. 7: The distribution of the percentage of atomic formulas pruned by LEOPARD across all pruned instances in Dataset1

Dataset1 binary analysis

Evaluation




Fig. 8: CDF of solving times for binary analysis of original and pruned SMT instances in Dataset2 across three solvers

Consistent
Speedup





Fig. 9: The ratio of solving time for binary analysis of the original and pruned instances on Dataset2 for all solvers



Significant
Pruning

Fig. 10: The distribution of the percentage of atomic formulas pruned by LEOPARD across all pruned instances in Dataset2

Dataset2 binary analysis

Evaluation

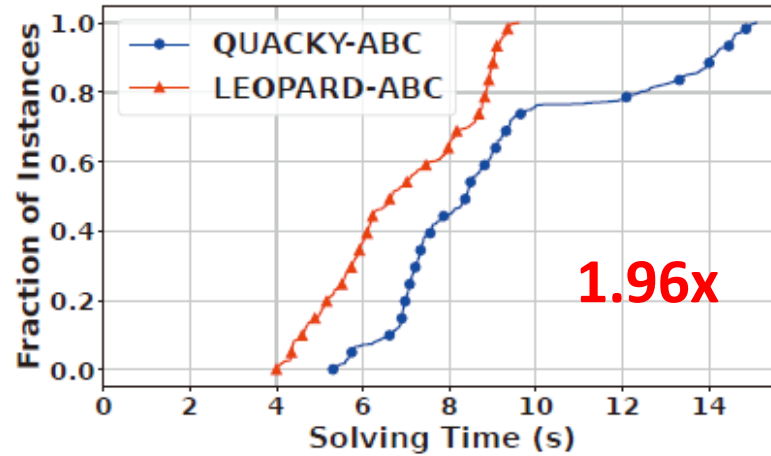


Fig. 11: CDF of solving times for quantitative analysis of original and pruned SMT instances in Dataset1 using ABC




Fig. 12: CDF of solving times for quantitative analysis of original and pruned SMT instances in Dataset1 using Z3

Dataset1 quantitative analysis

Conclusion

- LEOPARD: a method to accelerate access control policy verification via logical encoding optimization
 - Structure-based Pruning
 - Semantic-based Pruning
- Extensive evaluation