

CirC


Compiler Infrastructure for Cryptosystems and Verification

Alex Ozdemir, Fraser Brown, Riad S. Wahby

Cryptosystems are computers

53+/103 CRYPTO'21 papers!

Multi-Party Computation



$f(x, y)$

Fully Homomorphic Encryption

$$\text{Enc}(x) \rightarrow \text{Enc}(f(x))$$

functionality and performance depend on f 's representation

Zero-Knowledge

prove that

$$f(x) = 1$$

Cryptosystems are low-level

Problems:

1. How do we *express* f ?
2. How do we *optimize* f ?

4 bit
comparison

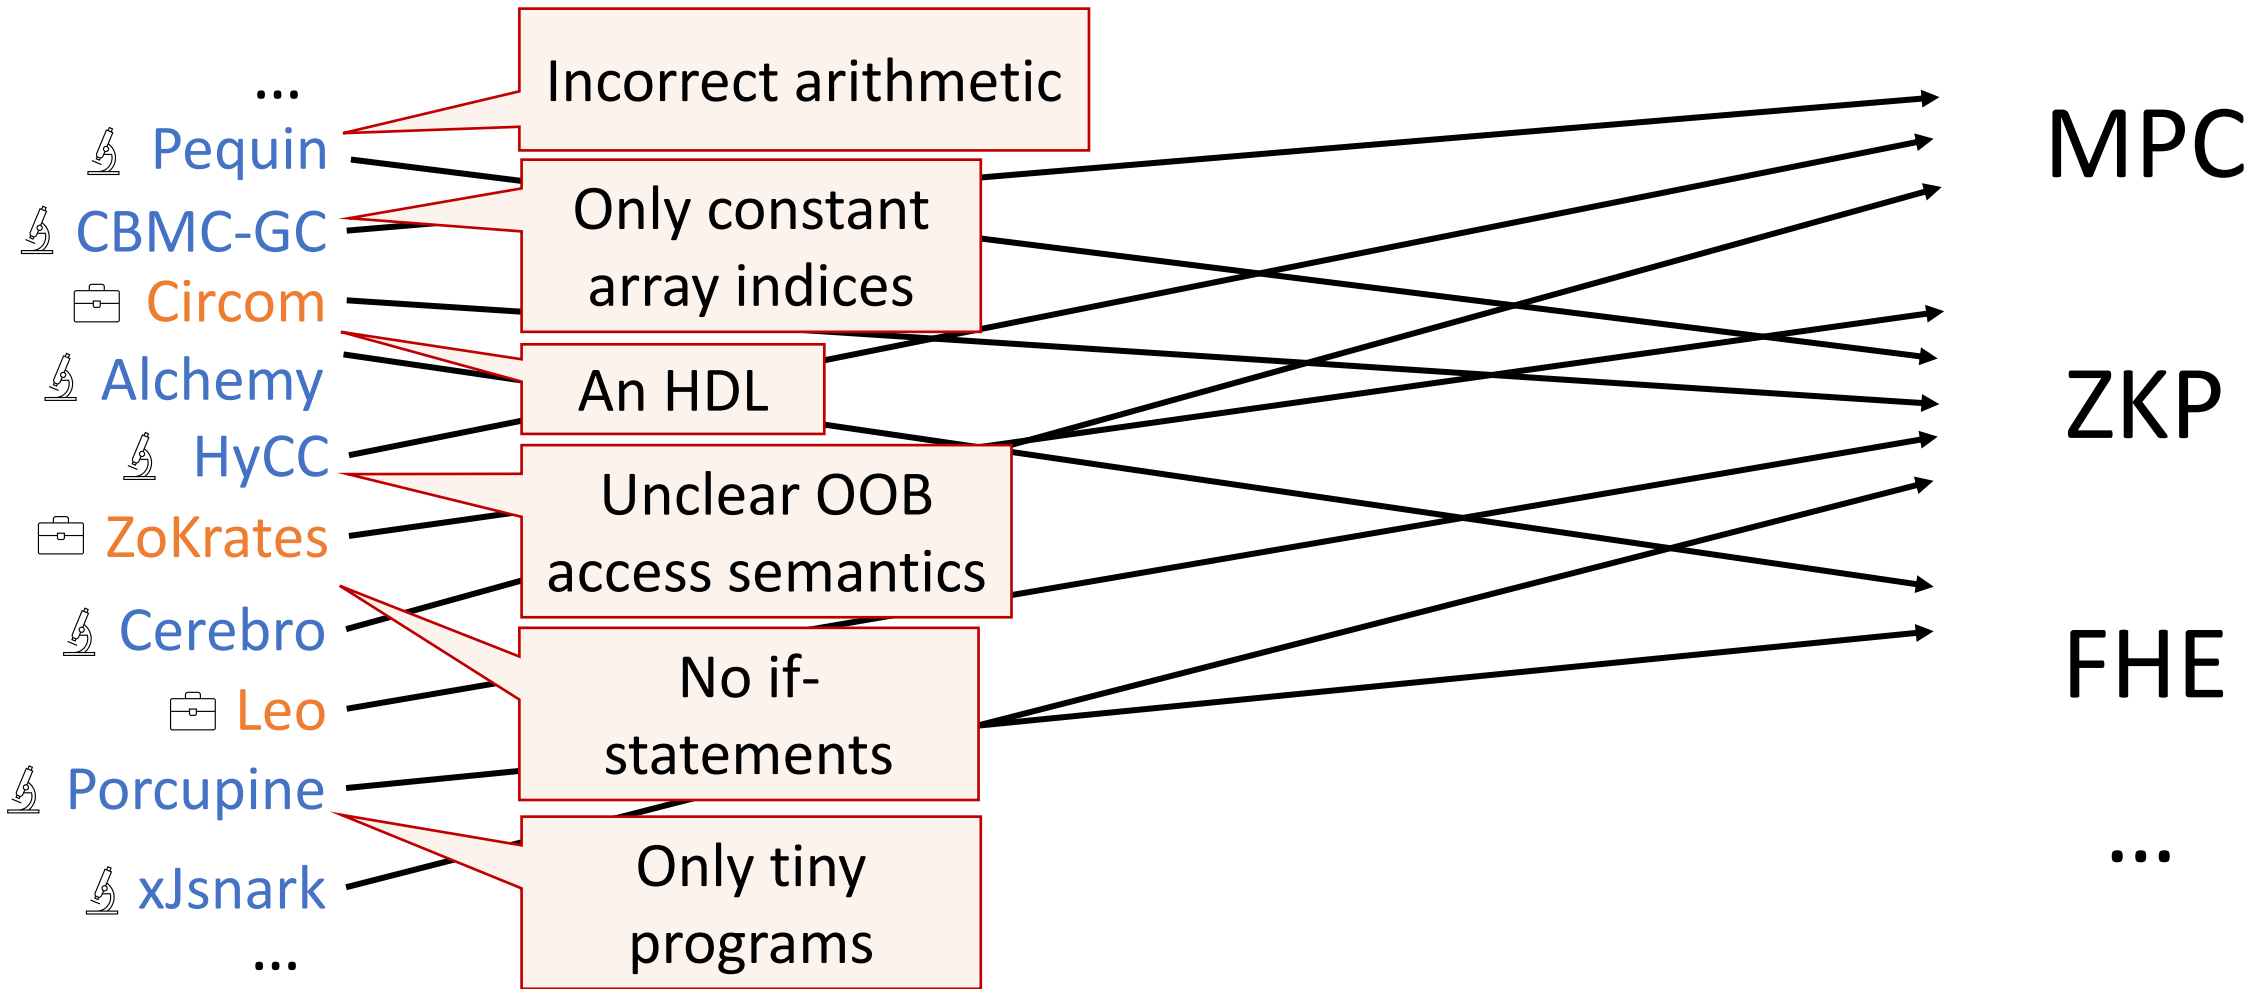
f :

```
a0(a0 - 1) = 0 ; a1(a1 - 1) = 0 ;  
a2(a2 - 1) = 0 ; a3(a3 - 1) = 0 ;  
b0(b0 - 1) = 0 ; b1(b1 - 1) = 0 ;  
b2(b2 - 1) = 0 ; b3(b3 - 1) = 0 ;  
a0(b0 - 1) = w0 ; 2a1(b1) = e1 -  
a1 - b1; a0(b0 - 1) = w0 ; ....
```

f :

```
a0(a0 - 1) = 0  
; a1(a1 - 1) =  
0 ; a2(a2 - 1)  
= 0 ; a3(a3 -  
1) = 0 ; ....
```

Cryptosystems need compilers



Can we share infrastructure?

Benefits:

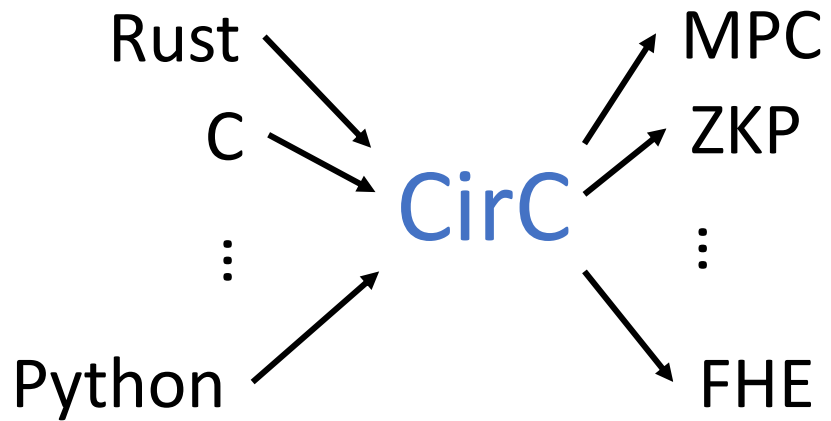
1. Save time

→ more research

→ better product

2. Produce better compilers

This Talk: CirC




1. Design & computational model
2. Output performance
3. Extensibility
4. Cross-over opportunities

CirC: Compiler Infrastructure for Cryptosystems and Verification

Cryptosystems process arithmetic circuits

- Directed acyclic graph
 - inputs
 - wires
 - gates
- Domain: finite field
 - wires \rightarrow field values
 - gates \rightarrow field operations
- Optional **equality assertions**
- Fewer gates \rightarrow better perf



$$f(\vec{x}, \vec{w}) \approx$$



How do we compile programs to circuits?

```
x = y  
x = x * y  
if y > 0:  
    x = x * x
```

```
for (z=1; z<3; z++):  
    y = y + z
```



1. Eliminate mutation

$$\begin{array}{l} x = y \\ x = x * y \end{array} \xrightarrow{\hspace{10em}} \begin{array}{l} x_1 = y_1 \\ x_2 = x_1 * y_1 \end{array}$$

2. Eliminate branches

```
x = y
x = x * y
if y > 0:
    x = x * x
```



```
x1 = y1
x2 = x1 * y1
x3 = y1 > 0
      ? x2 * x2
      : x2
```



C “ternary”
operator: a MUX.

3. Unroll Loops

```
x = y
x = x * y
if y > 0:
    x = x * x
```

```
for (z=1; z<3; z++):
    y = y + z
```



```
x1 = y1
x2 = x1 * y1
x3 = y1 > 0
      ? x2 * x2
      : x2
z1 = 1
y2 = y1 > 0 && z1 < 3
      ? y1 + z1
      : y1
z2 = z1 + 1
y3 = y1 > 0 && z1 < 3 && z2 < 3
      ? y2 + z2
```

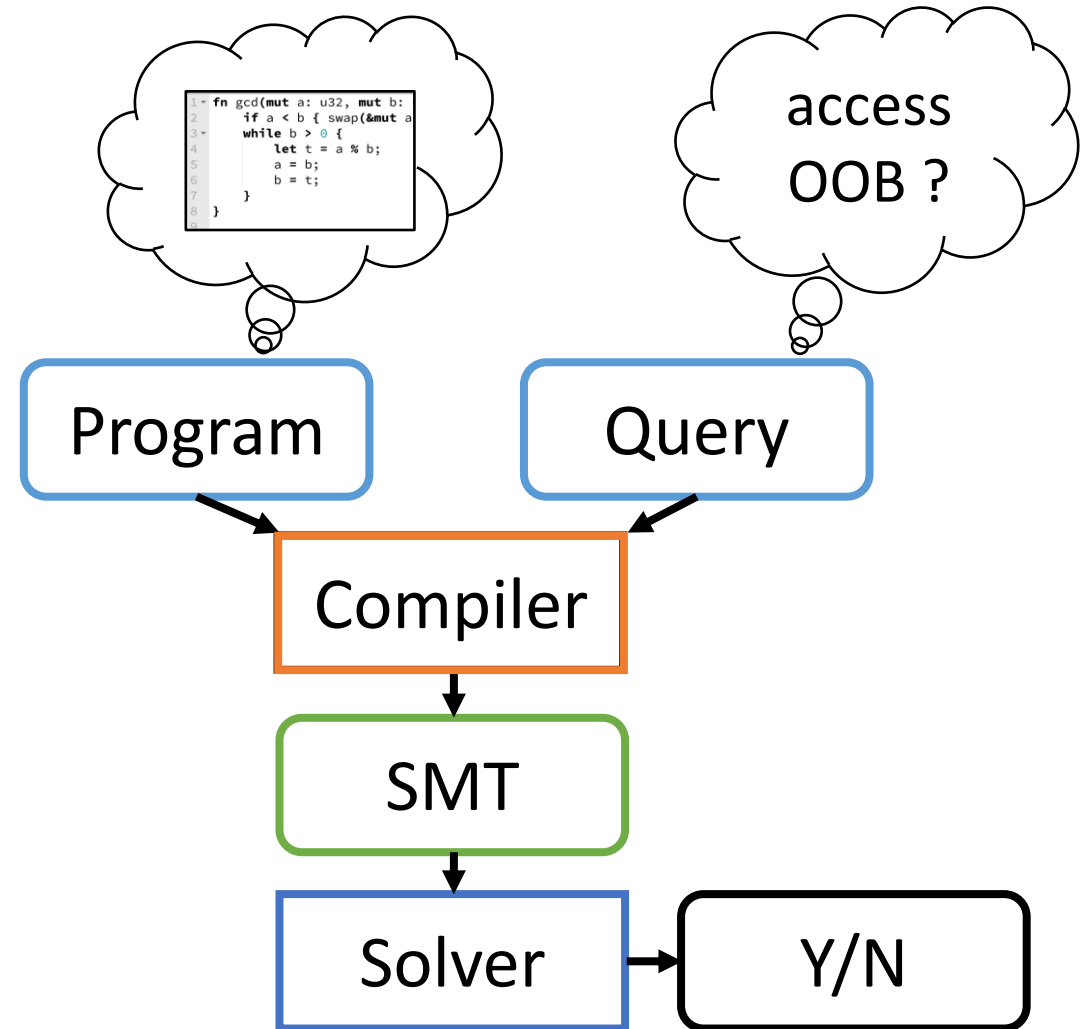
How do we compile programs to circuits?

1. Eliminate mutation
2. Eliminate branches
3. Unroll loops
4. Inline function calls
5. Functionalize arrays
6. ...

Similar to
Bounded Model
Checking!

Bounded Model Checking is compilation


1. Model the **program*** and **query** as a **logical formula**
2. Check the model with a **solver**



**the BMC connection is latent in projects like CMBC-GC

*consider bounded traces 14

Computational Model: Existentially Quantified Circuits (EQCs)



- Mutation free
- Control-flow free
- Fixed-size (non-uniform)
- Variables of *restricted knowledge*:
 - ZKP: known to prover
 - FHE: known only to user
 - MPC: distributed knowledge
 - SMT: existentially quantified

} *circuits*

} *existential
quantifiers
(with
metadata)*

Computational Model: Existentially Quantified Circuits (EQCs)

Problem:
Existentially Quantified Circuits (EQCs)
are not easy to program!

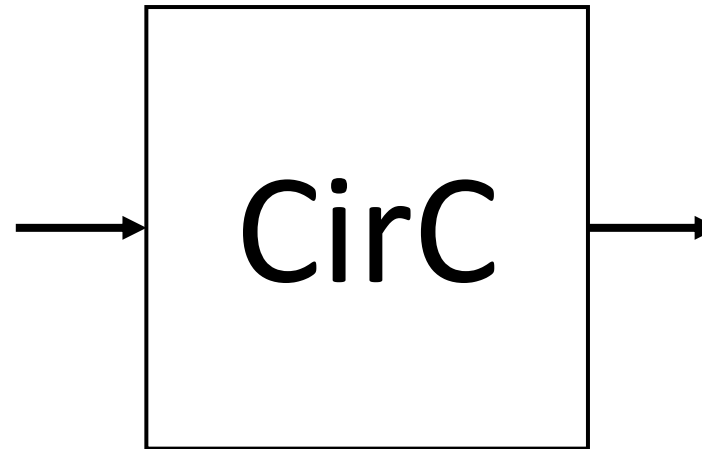
CirC compiles high-level languages to circuits

High-Level Language

- Mutation
- RAM
- Control-Flow
- Uniform


```
x = y
x = x * y
if y > 0:
    x = x * x

for (z=1; z<3; z++):
    y = y + z
```



EQC

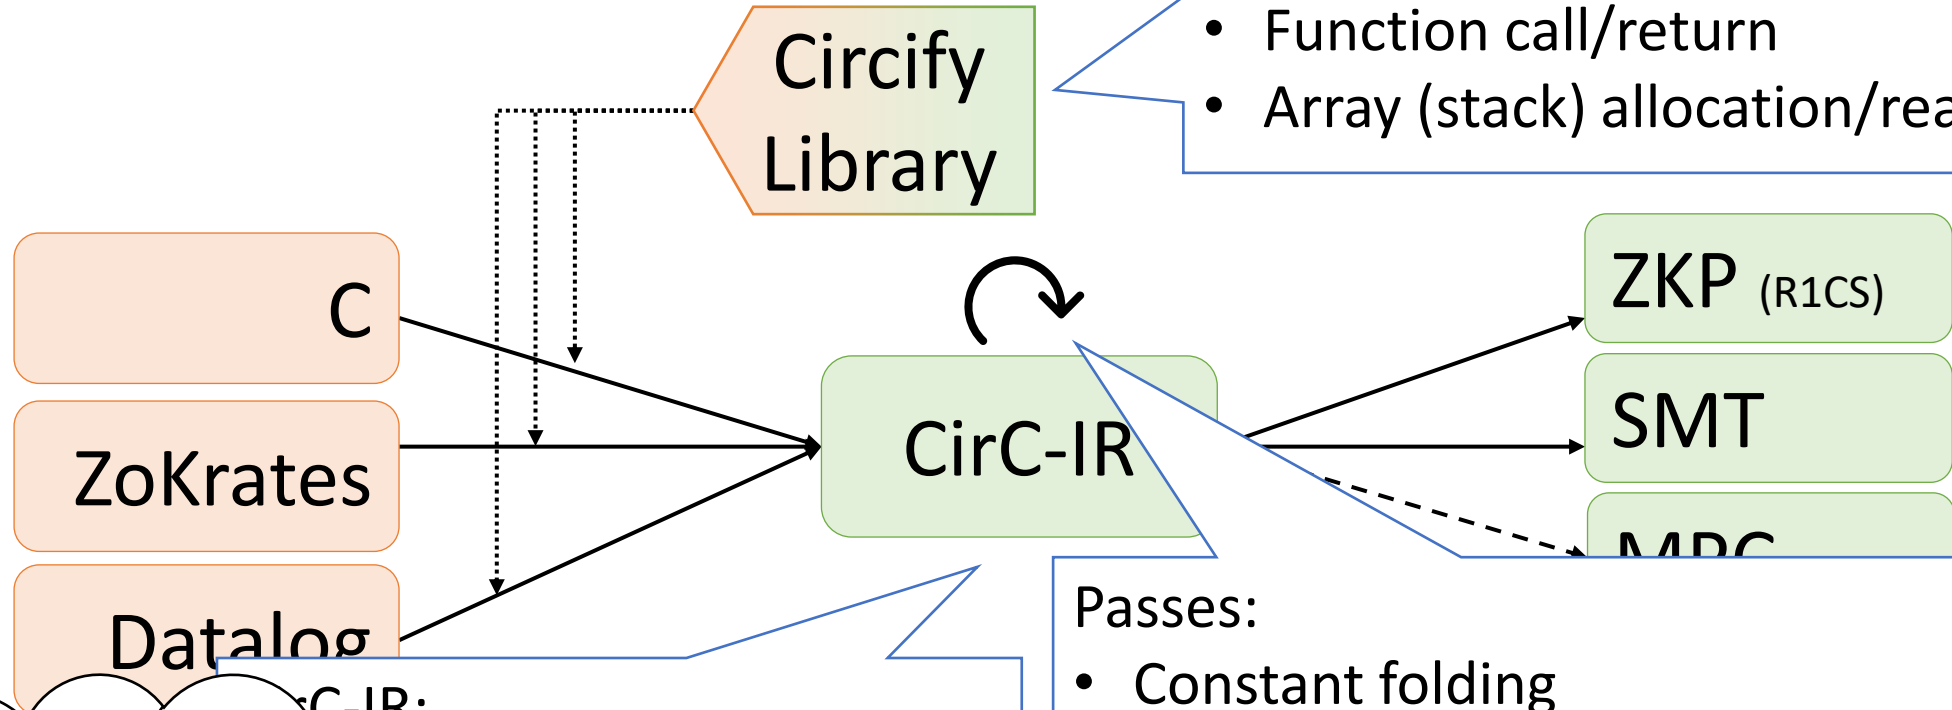
- No mutation
- No RAM
- No control-flow
- Non-uniform



CirC's Architecture

Manages:

- Variable scope/decl./read/write
- Control flow: branch, break,...
- Function call/return
- Array (stack) allocation/read/write

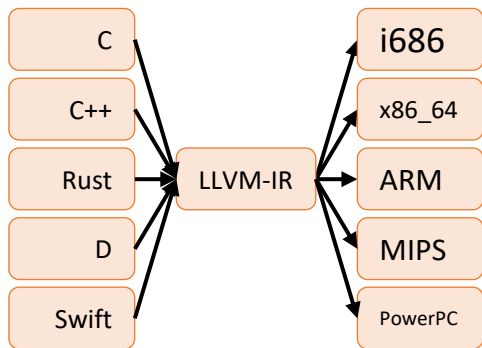


CirC-IR:

cleans, **harder:**
 the fields, floating-point
pgm → cirC
 nded arrays
 es

Passes:

- Constant folding
- Peephole (e.g. bitwise ternary)
- Tuple elimination
- Array elimination (oblivious, linear scan, transcript-checking)




Case Study: ZoKrates to ZKP

With CirC, it's *easy* to build a *performant* compiler.

ZoKrates (v0.6.2)

- Targets ZKP
- Types: Booleans, machine-integers, fields, arrays, structures,...
- Good tooling
- Supported by the Ethereum Foundation



Stdlib: ecc/edwardsAdd.zok:

```
from "ecc/babyjubjubParams" import BabyJubJubParams as P
```

```
def main(field[2] pt1, field[2] pt2, P pp) -> field[2]:
```


```
    field a = pp.JUBJUB_A  
    field d = pp.JUBJUB_D
```

```
    field u1 = pt1[0]  
    field v1 = pt1[1]  
    field u2 = pt2[0]  
    field v2 = pt2[1]
```

```
    field uOut = (u1*v2 + v1*u2) / (1 + d*u1*u2*v1*v2)  
    field vOut = (v1*v2 - a*u1*u2) / (1 - d*u1*u2*v1*v2)
```


```
    return [uOut, vOut]
```

ZoKrates' reference compiler is big




Compiler	Reference	
Lines of Code	~28,000	
Development Time	3 years	
Contributors	36	
Output Size	good	

Compiling ZoKrates is easier with CirC




Compiler	Reference	CirC
Lines of Code	~28,000	~700
Development Time	3 years	1 week
Contributors	36	1
Output Size	good	better

CirC-ZoKrates outperforms ZoKrates



Constraint Ratio
CirC/ZoKrates
(lower is better)



File	CirC		ZoKrates	
	Constraint	Time (s)	Constraint	Time (s)
ec-add	10	0.08	10	0.01
ec-scalar-mult	21	0.07	20	0.01
eddsa	363	1.06	366	0.11
field-to-bools	2	0.06	2	0.02
mimc7	11	0.12	9	0.02
mux3signed	40	0.09	60	0.03
pedersen	8744	0.16	8884	0.06
sha2round	6100	1.25	6488	0.08
u32s-to-bools	40	0.46	42	0.02
ec-add	300	1.31	282	0.19
ec-scalar-mult	300	1.26	302	0.08
eddsa	460	1.26	462	1.18
field-to-bools	3612	0.46	3643	0.19
mimc7	3622	0.46	3600	0.16
mux3signed	8422	0.11	7975	0.11
pedersen	12062	0.12	14614	0.42
sha2round	17791	0.46	17407	0.36
u32s-to-bools	1714	0.12	1688	0.16
ec-add	1554	12.11	14804	8.76
ec-scalar-mult	2840	0.46	2740	1.14
eddsa	2840	0.46	2870	2.47
field-to-bools	4100	0.16	4100	0.16
mimc7	3040	0.16	3040	0.16
mux3signed	4120	0.16	4120	0.16
pedersen	264	0.65	264	0.05
sha2round	244	0.62	260	0.02
u32s-to-bools	312	0.16	320	0.02
ec-add	4	0.16	4	0.02
ec-scalar-mult	4	0.16	4	0.02
eddsa	500	0.07	4	0.02
field-to-bools	512	0.08	512	0.08
mimc7	256	1.10	256	0.06
mux3signed	256	1.10	256	0.06
pedersen	256	2.05	263	0.11
sha2round	256	0.16	256	0.06
u32s-to-bools	256	0.16	256	0.06

Full comparison
in appendix

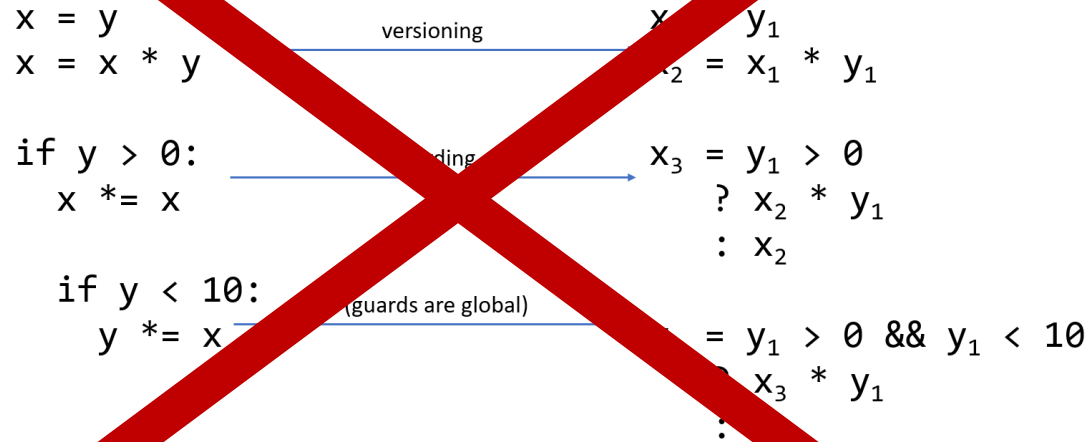
CirC's ZKP compilers perform well

Input Language	Baseline Compiler	CirC Improvement
ZoKrates	ZoKrates	~1x - 1.8x
C	Pequin	~1x - 8x
Circom	Circom	<i>identical (HDL)</i>

A Tour of the ZoKrates Front-End

Building a circuit compiler in “three easy pieces”

Writing a New Front-End



Key Steps:

1. Represent language values using CirC-IR
2. Hook into Circify
3. Translate AST (using Circify)

Mapping Language Values to CirC-IR

- All ZoKrates types have CirC-IR analogs:
 - `bool` → booleans
 - `u8, u16, u32` → bit-vectors
 - `field` → prime fields
 - `[]` → field-indexed arrays
 - `structures` → tuples


```
/src/front/zokrates/term.rs
27  pub enum Ty {
28      Uint(usize),
29      Bool,
30      Field,
31      Struct(String, FieldList<Ty>),
32      Array(usize, Box<Ty>),
33  }

123 pub struct T {
124     pub ty: Ty,
125     pub term: ir::Term,
126 }
```



Hooking Into Circify

- Circify must *merge states*



- Circify needs *if-then-else* expressions for language values

/src/front/zokrates/term.rs

```
465 fn ite(c: Term, a: T, b: T) -> Result<T, String> {  
466     if &a.ty != &b.ty {  
467         Err(format!("Cannot perform ITE on {} and {}", a, b))  
468     } else {  
469         Ok(T::new(a.ty.clone(), term![Op::Ite; c, a.term, b.term]))  
470     }  
471 }
```

x = Op::Ite(b > 0, y, x)



“Interpreting” the AST Using Circify

```
/src/front/zokrates/mod.rs*  
impl Interp {  
    fn stmt(&mut self, s: &ast::Statement<'ast>) {           // "interpret" a statement
```

*w/o error handling

Map values to CirC-IR



Hook into Circify



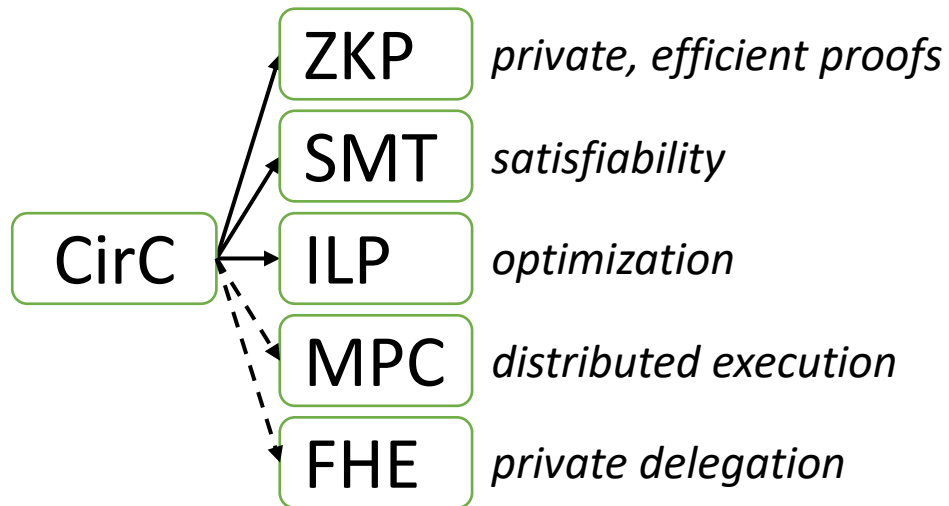
“Interpret” the AST

Stretch break!

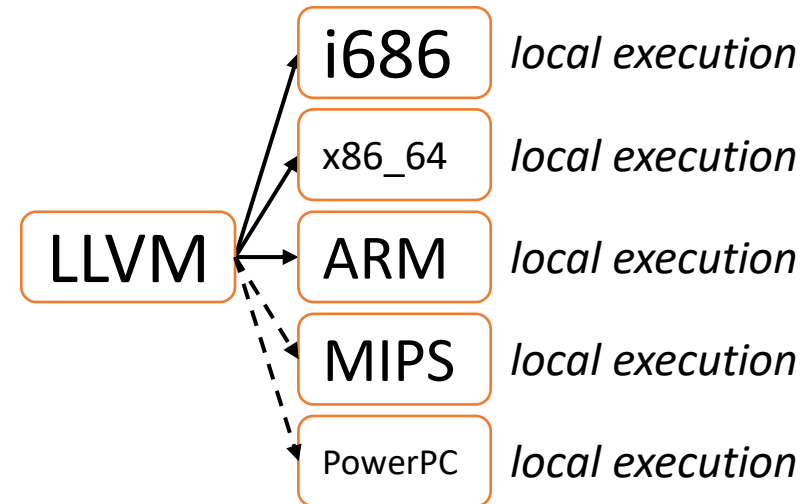
Cross-Over Opportunities

Could applications use *multiple* targets?

CirC: targets have *different* purposes



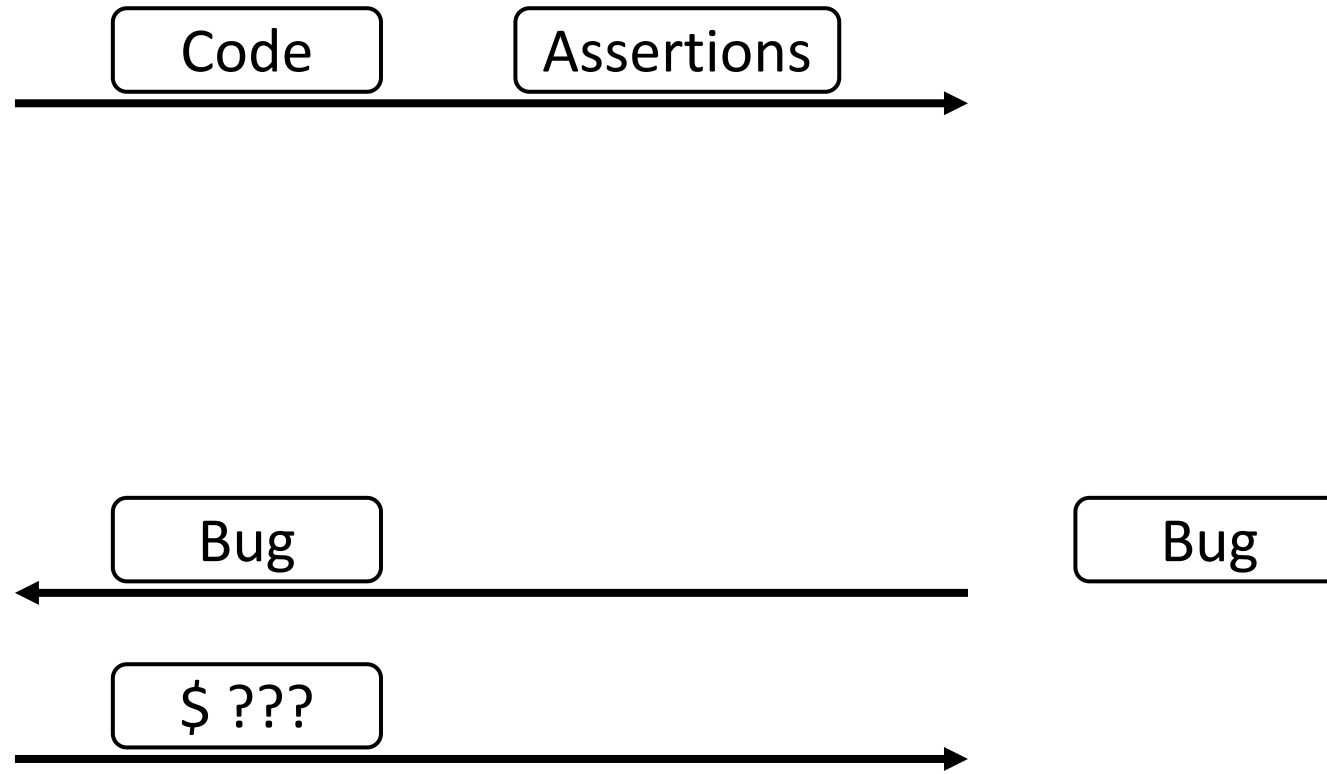
LLVM: targets have the same purpose



Supporting ZK Bug Bounties

Codebase Owner

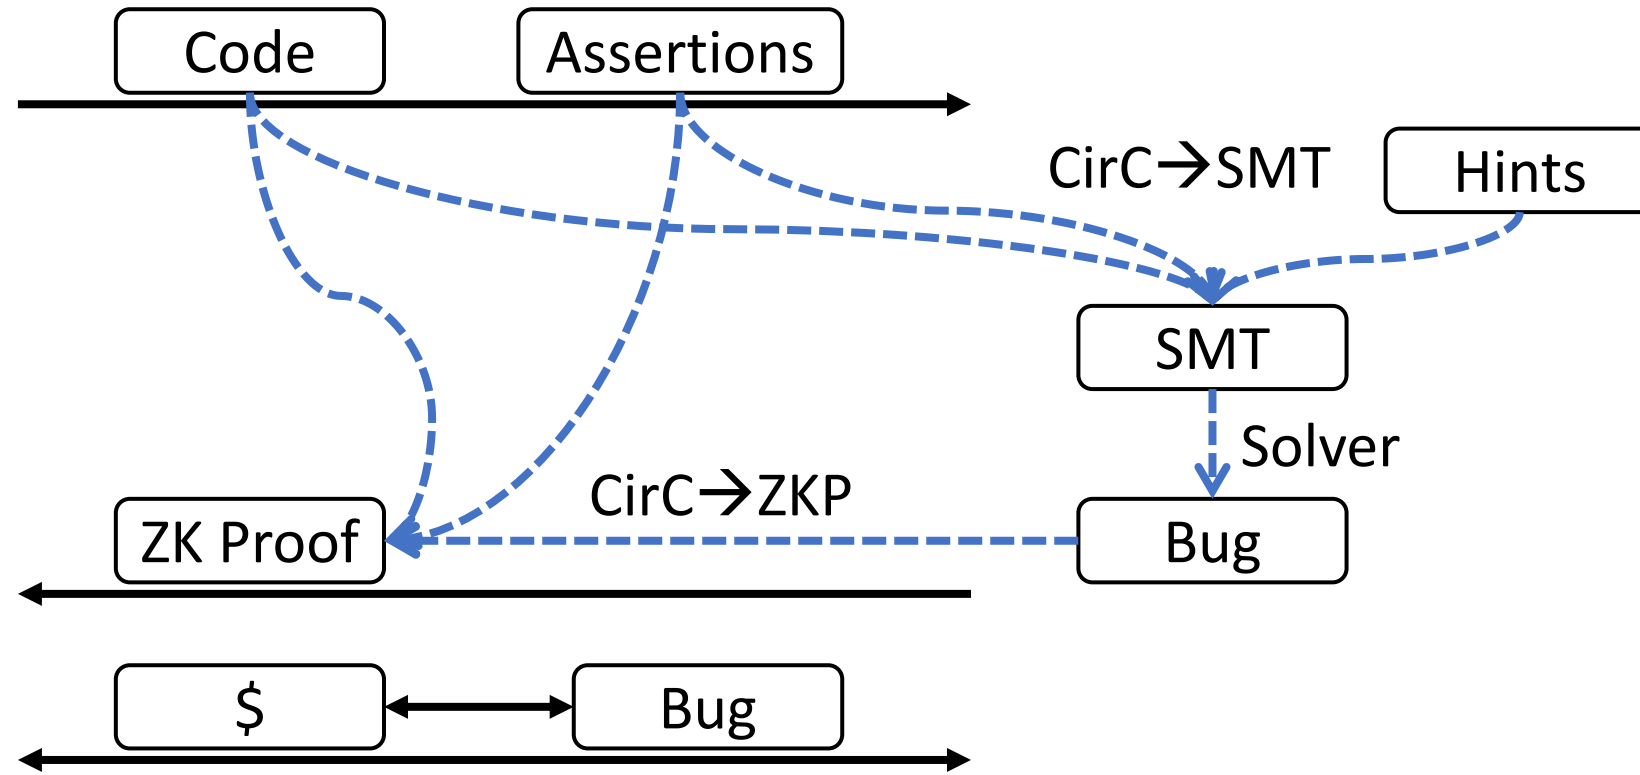
Analyst



Supporting ZK Bug Bounties

Codebase Owner

Analyst



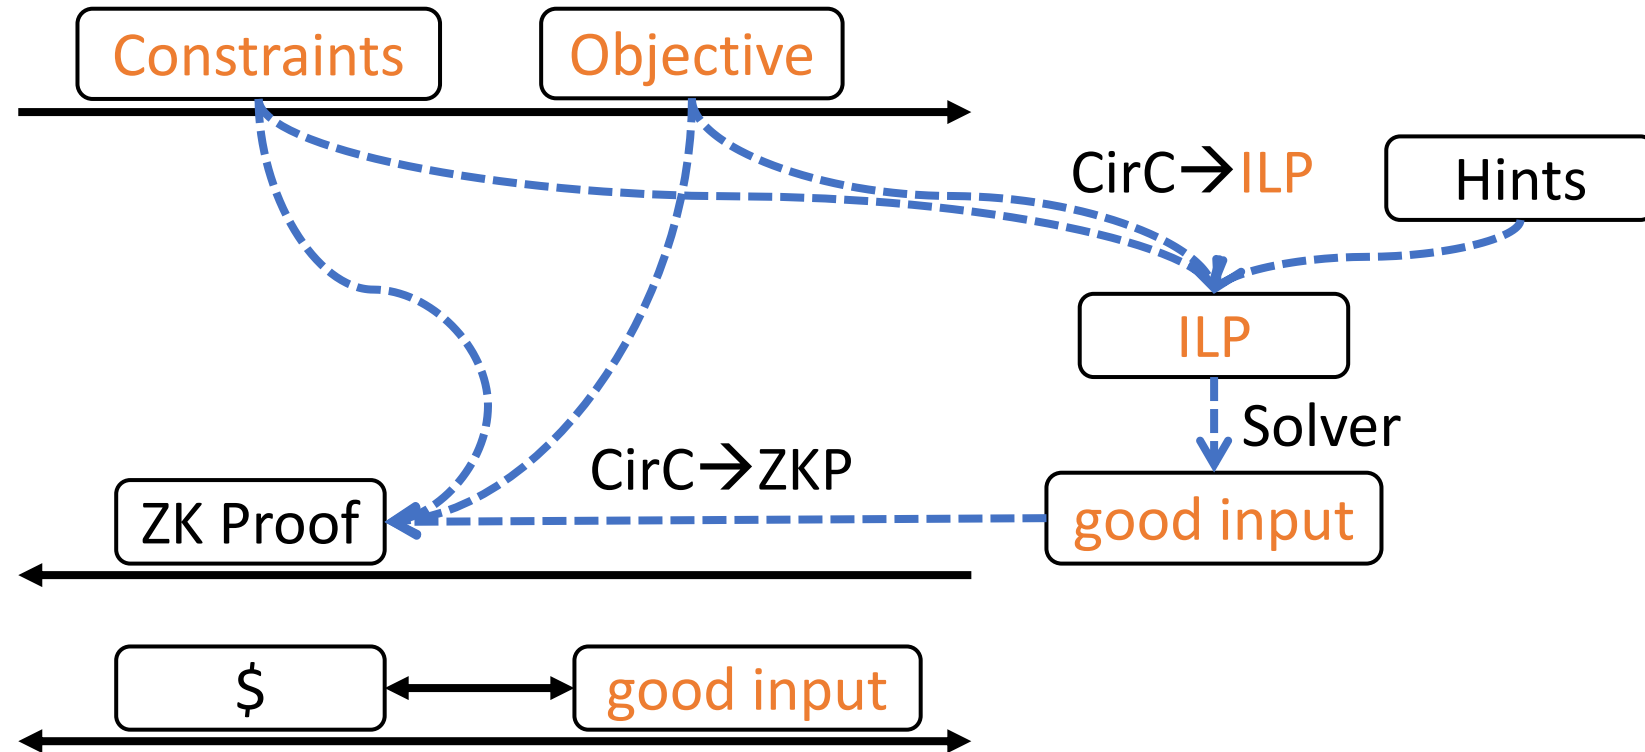
Proof of concept: CVE-2014-3570 (OpenSSL)

≈ 60 LOC

Supporting ZK Optimization Competitions

Operator

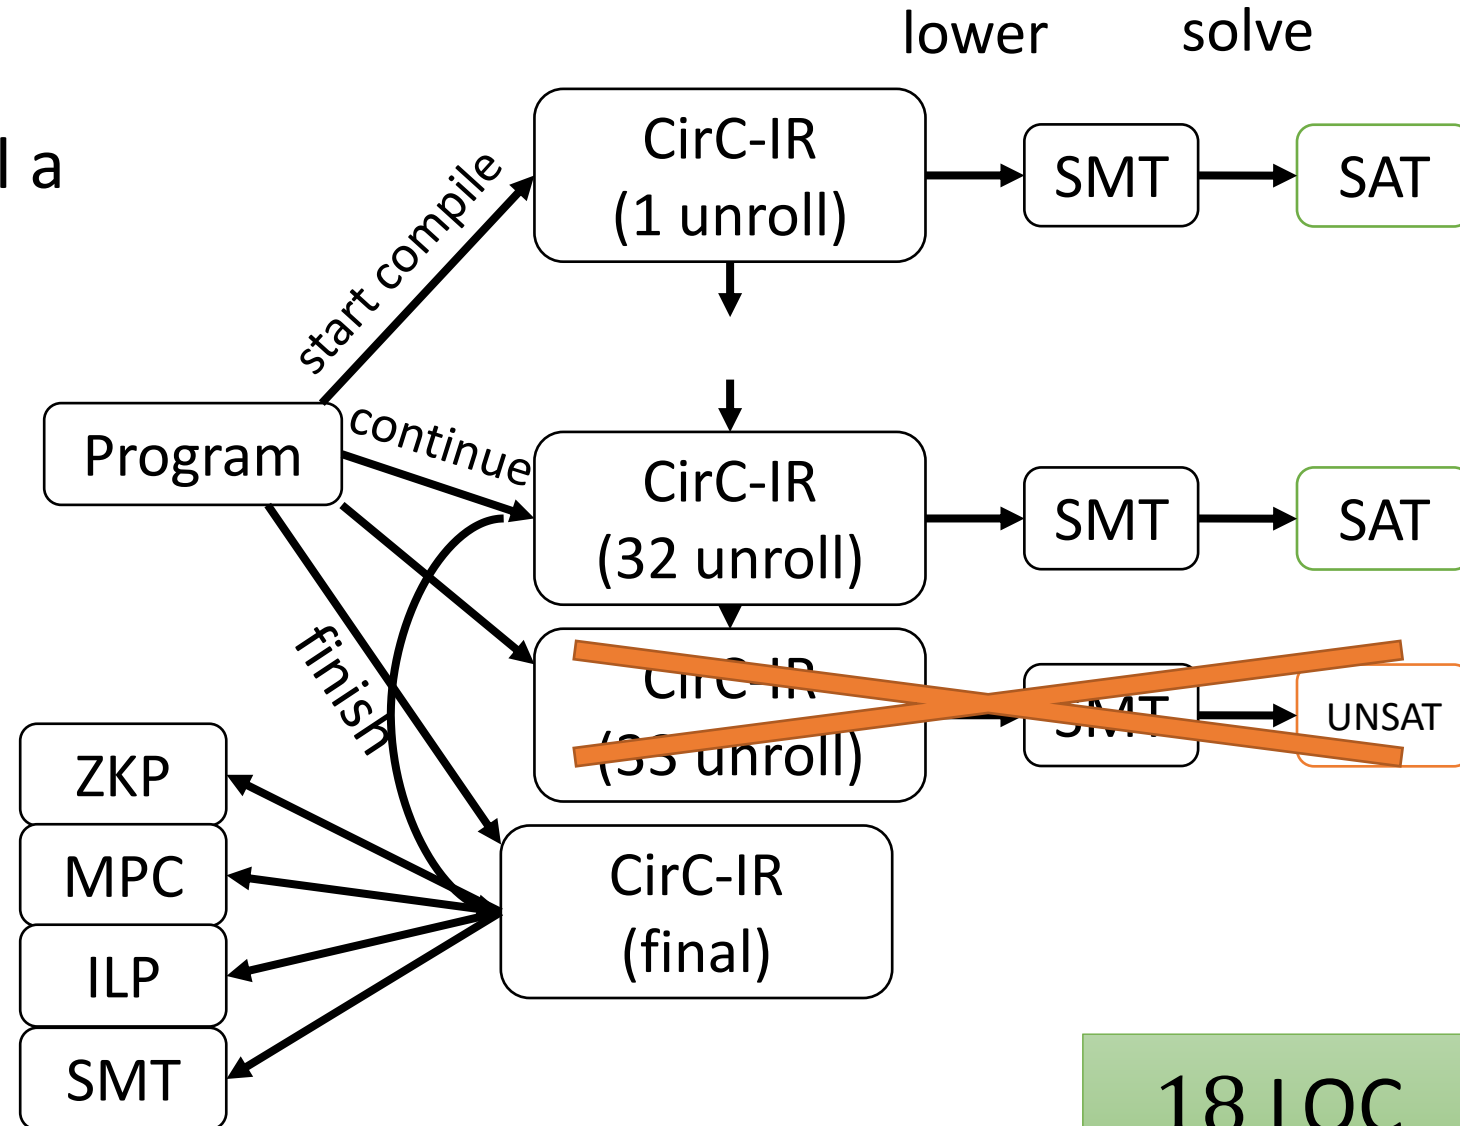
Competitor



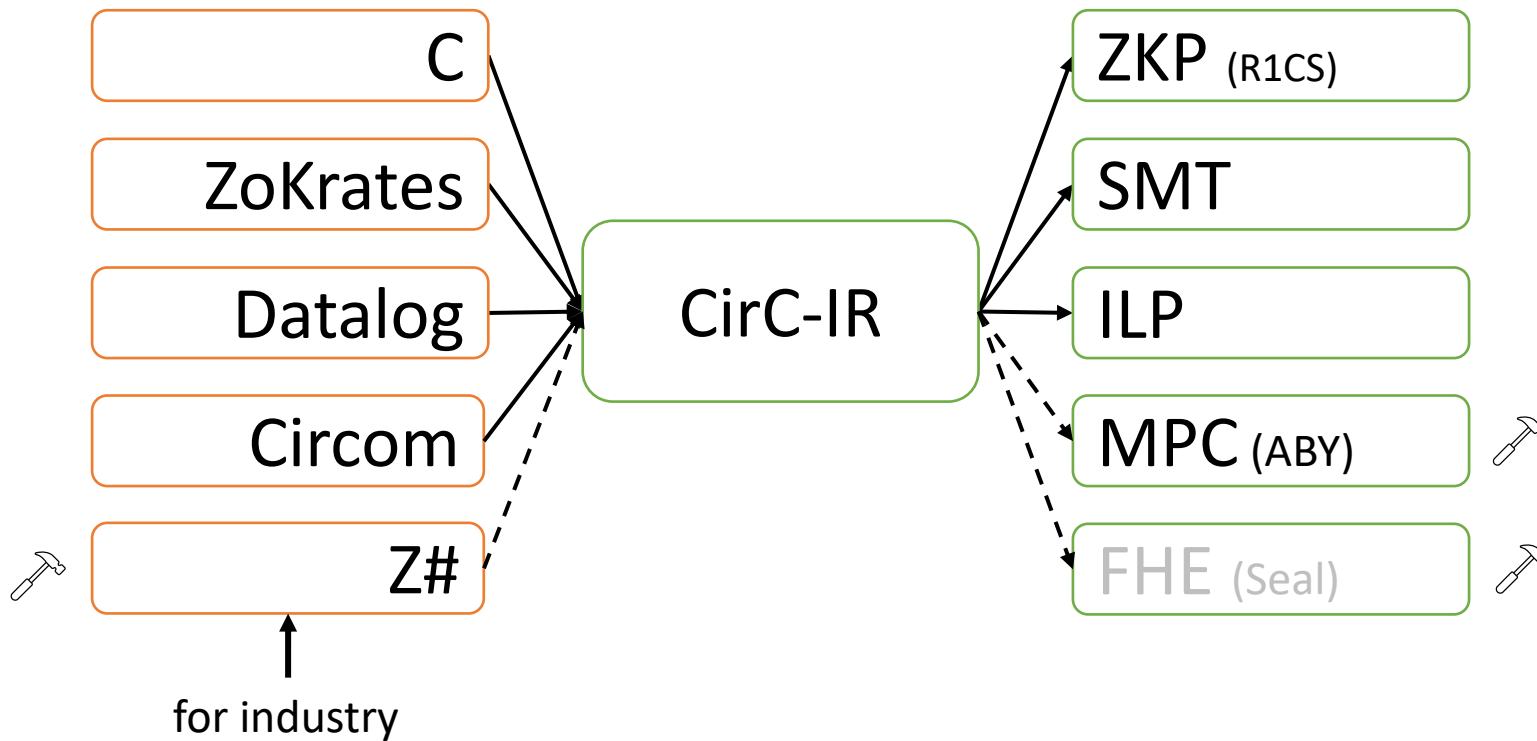
SMT-Assisted Optimization

How many times do we unroll a loop?

- Too many → inefficient
- Too few → incomplete



Ongoing & Future Work



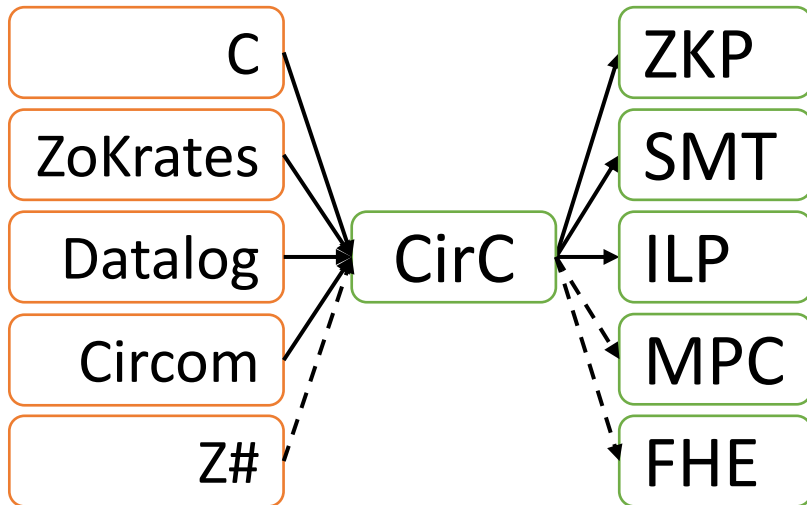
- Optimize for “stacked garbling” MPCs?
- Batching for MPC protocols?
- Compiling to “layered circuits”?

Builds on (Haskell) CirC: “Efficient Representation of Numerical Optimization Problems for SNARKs” <https://ia.cr/2021/1436>

Sebastian Angel and Andrew J. Blumberg and Eleftherios Ioannidis and Jess Woods, USENIX Security’22

CirC: Circuit Compiler Infrastructure



Different circuits can **share** compiler infrastructure




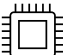
Based on **EQCs**

(existentially quantified circuits)

Benefits:

-  easy extension
-  shared optimizations

Future directions:

-  new front-ends
-  new backends (and opportunities)

? come talk to us!