

# Programmable Cryptography: The **Full Stack**

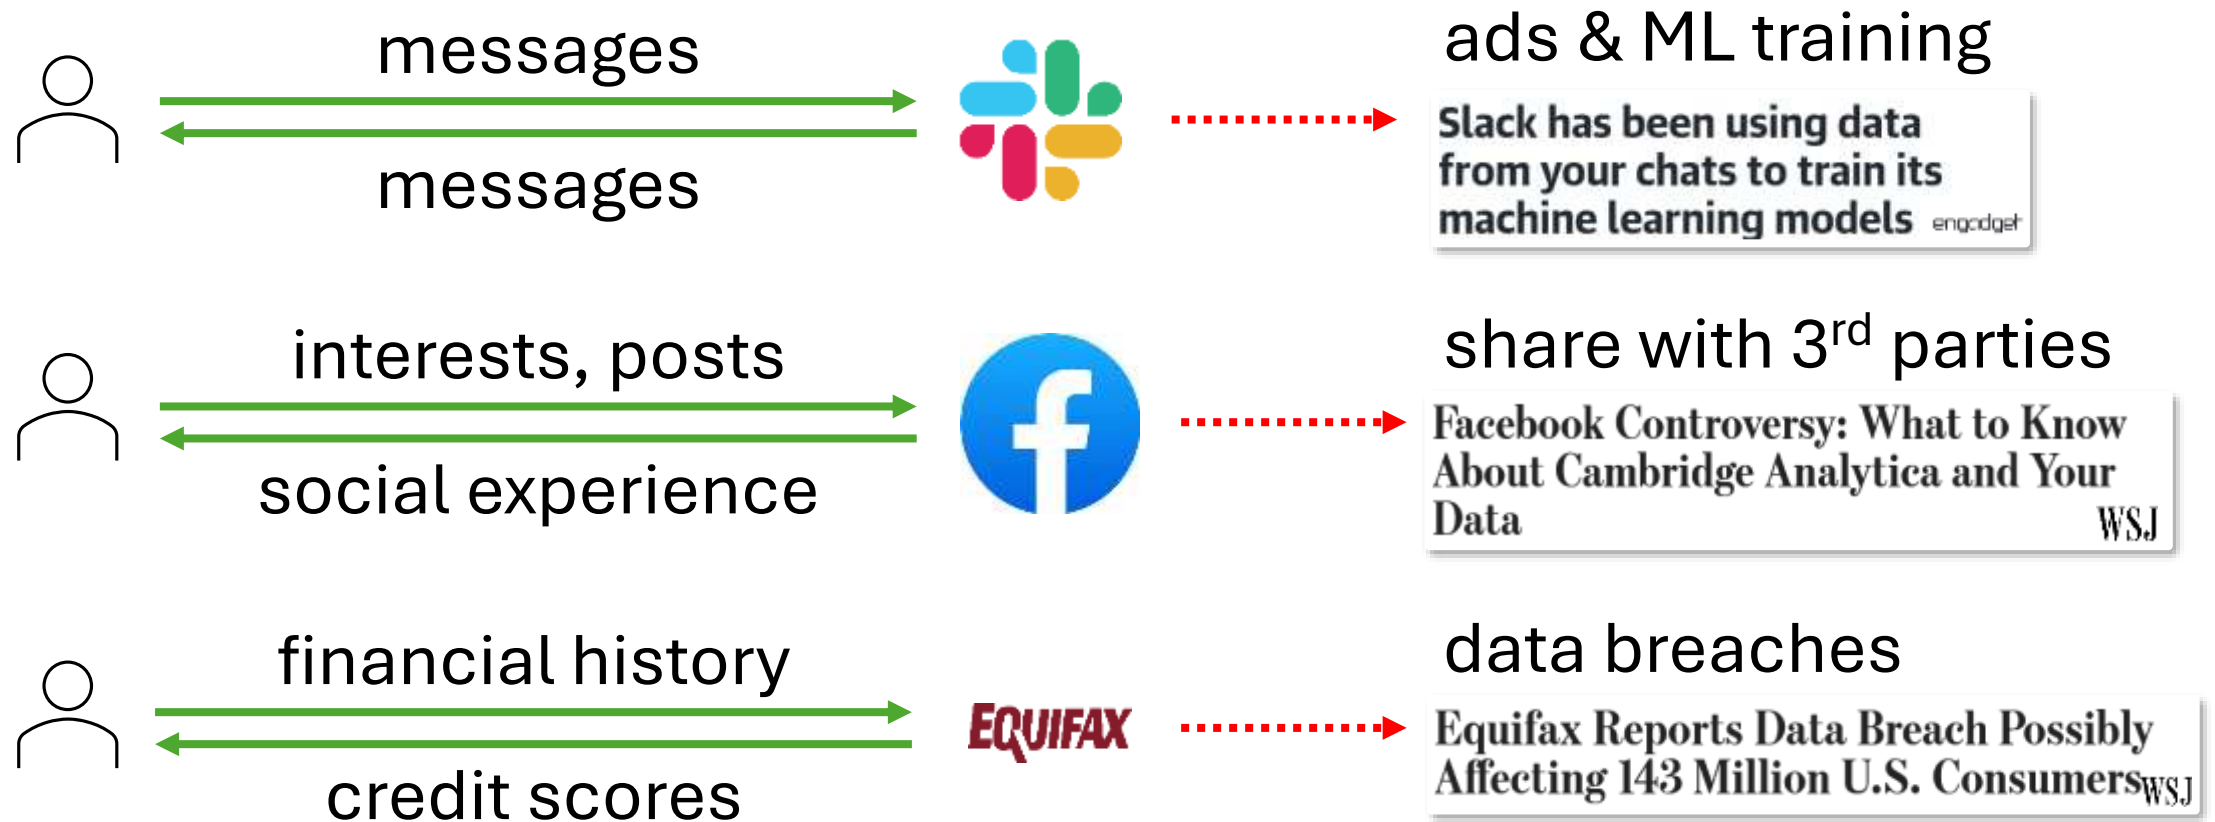
(via **cryptology**, **compilers**, and **verification**)

Alex Ozdemir


**May 30<sup>th</sup>, 1pm PT**

**Based on joint work with:** Clark Barrett, Dan Boneh, Fraser Brown,  
Gereon Kremer, Cesare Tinelli, Riad S. Wahby

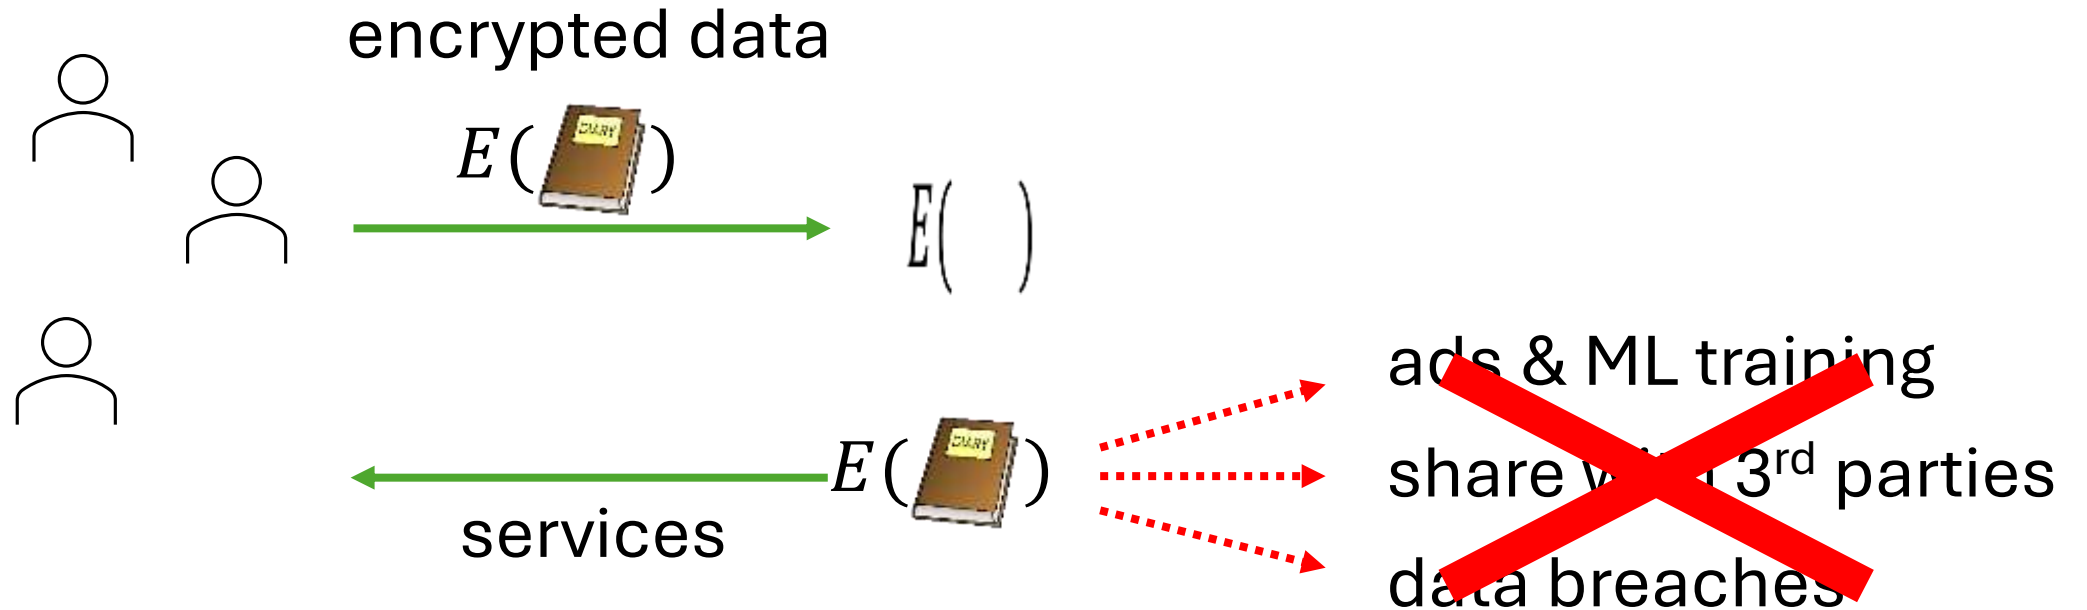
# To use the internet is to be vulnerable



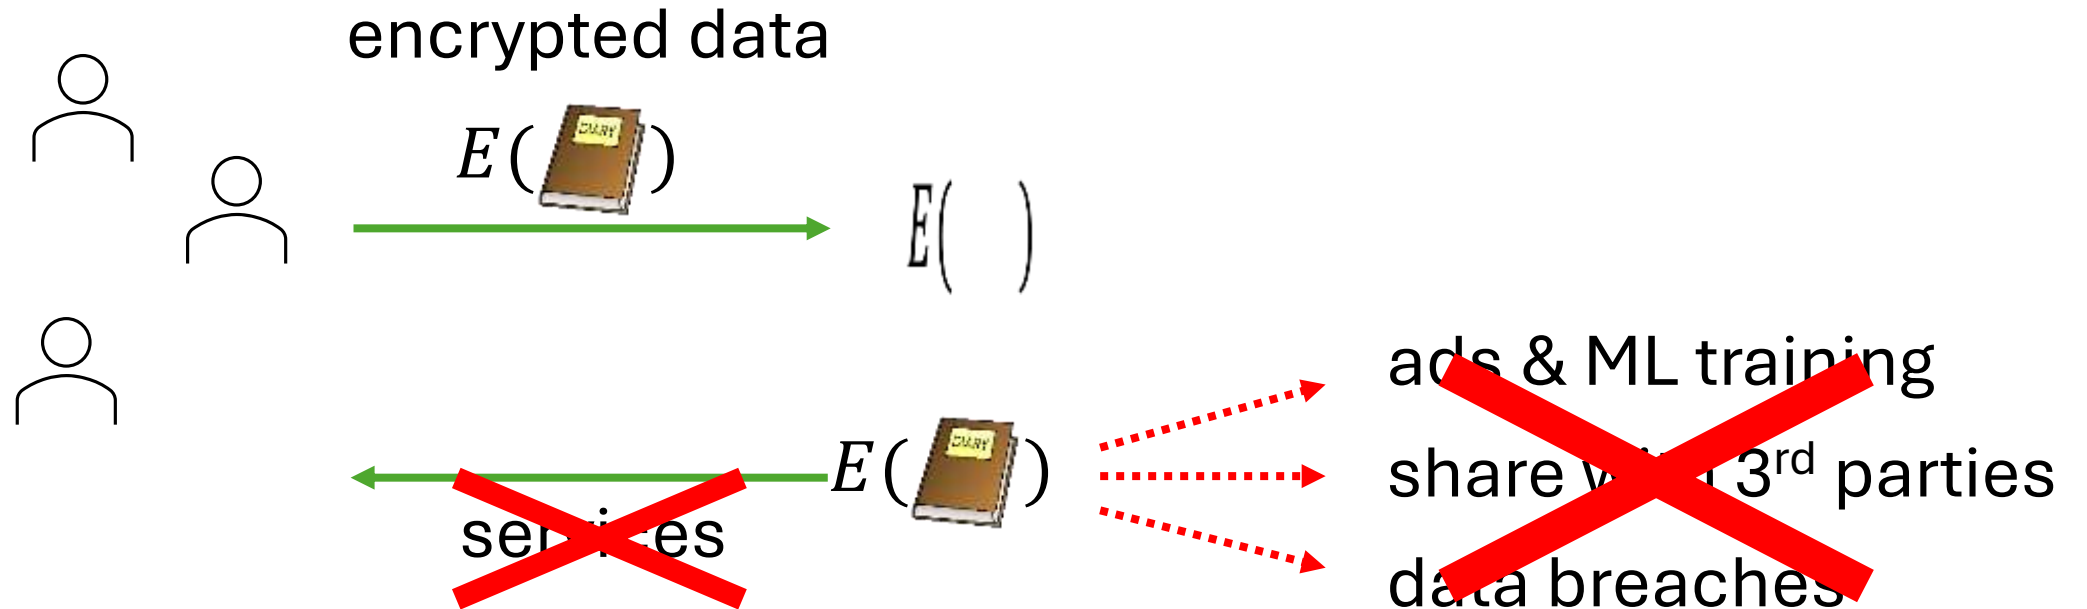
# To use the internet is to be vulnerable



# Encryption secures data



# Encryption secures data, but prevents use



# Programmable cryptography enables use

A **programmable cryptosystem** securely executes a developer-defined program on secret data.

Examples:

**FHE**  
Fully Homomorphic Encryption  
**compute** on encrypted data  
 $E(x) \rightarrow E(f(x))$

**ZKP**  
Zero Knowledge Proof  
**prove** properties of secret data  
 $E(x) \rightarrow \phi(x)$

**MPC**  
Multi Party Computation  
**compute** on distributed data  
 $x \parallel y \rightarrow f(x, y)$

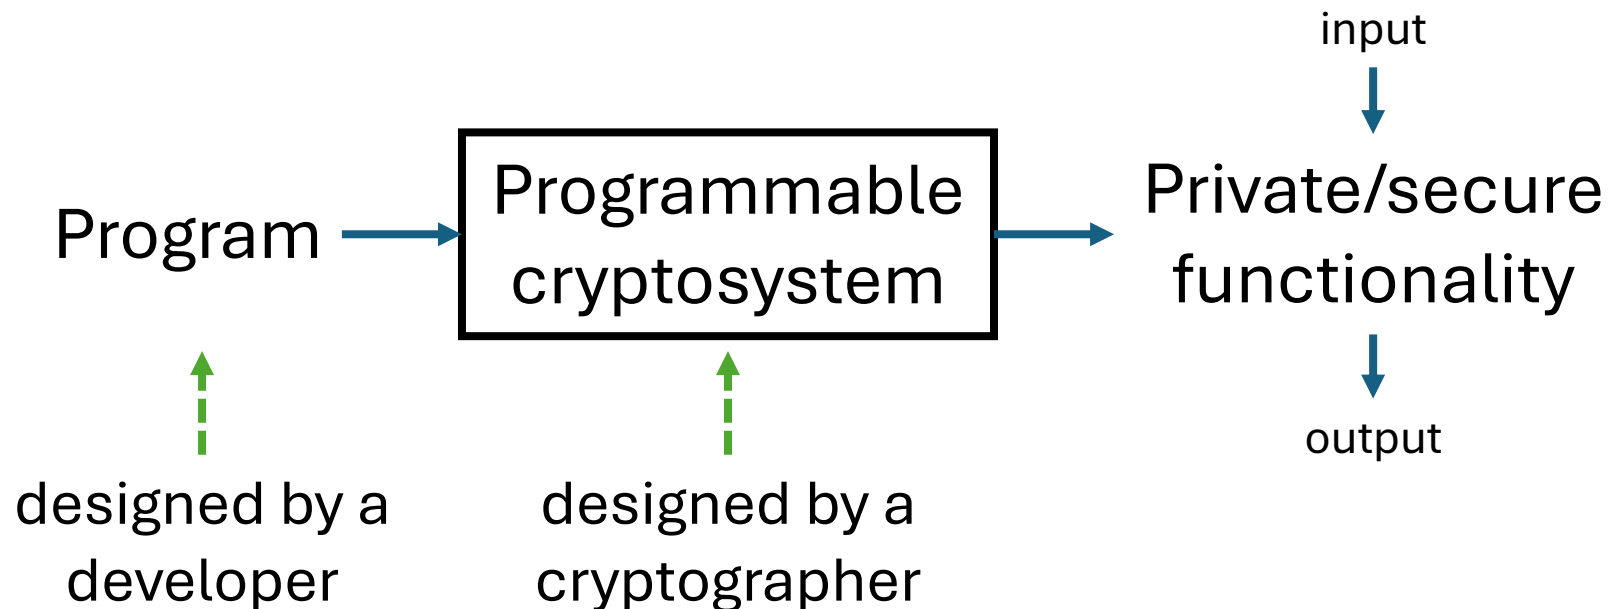
Encryption secures ***data***.

Programmable cryptography  
secures ***data in use***.

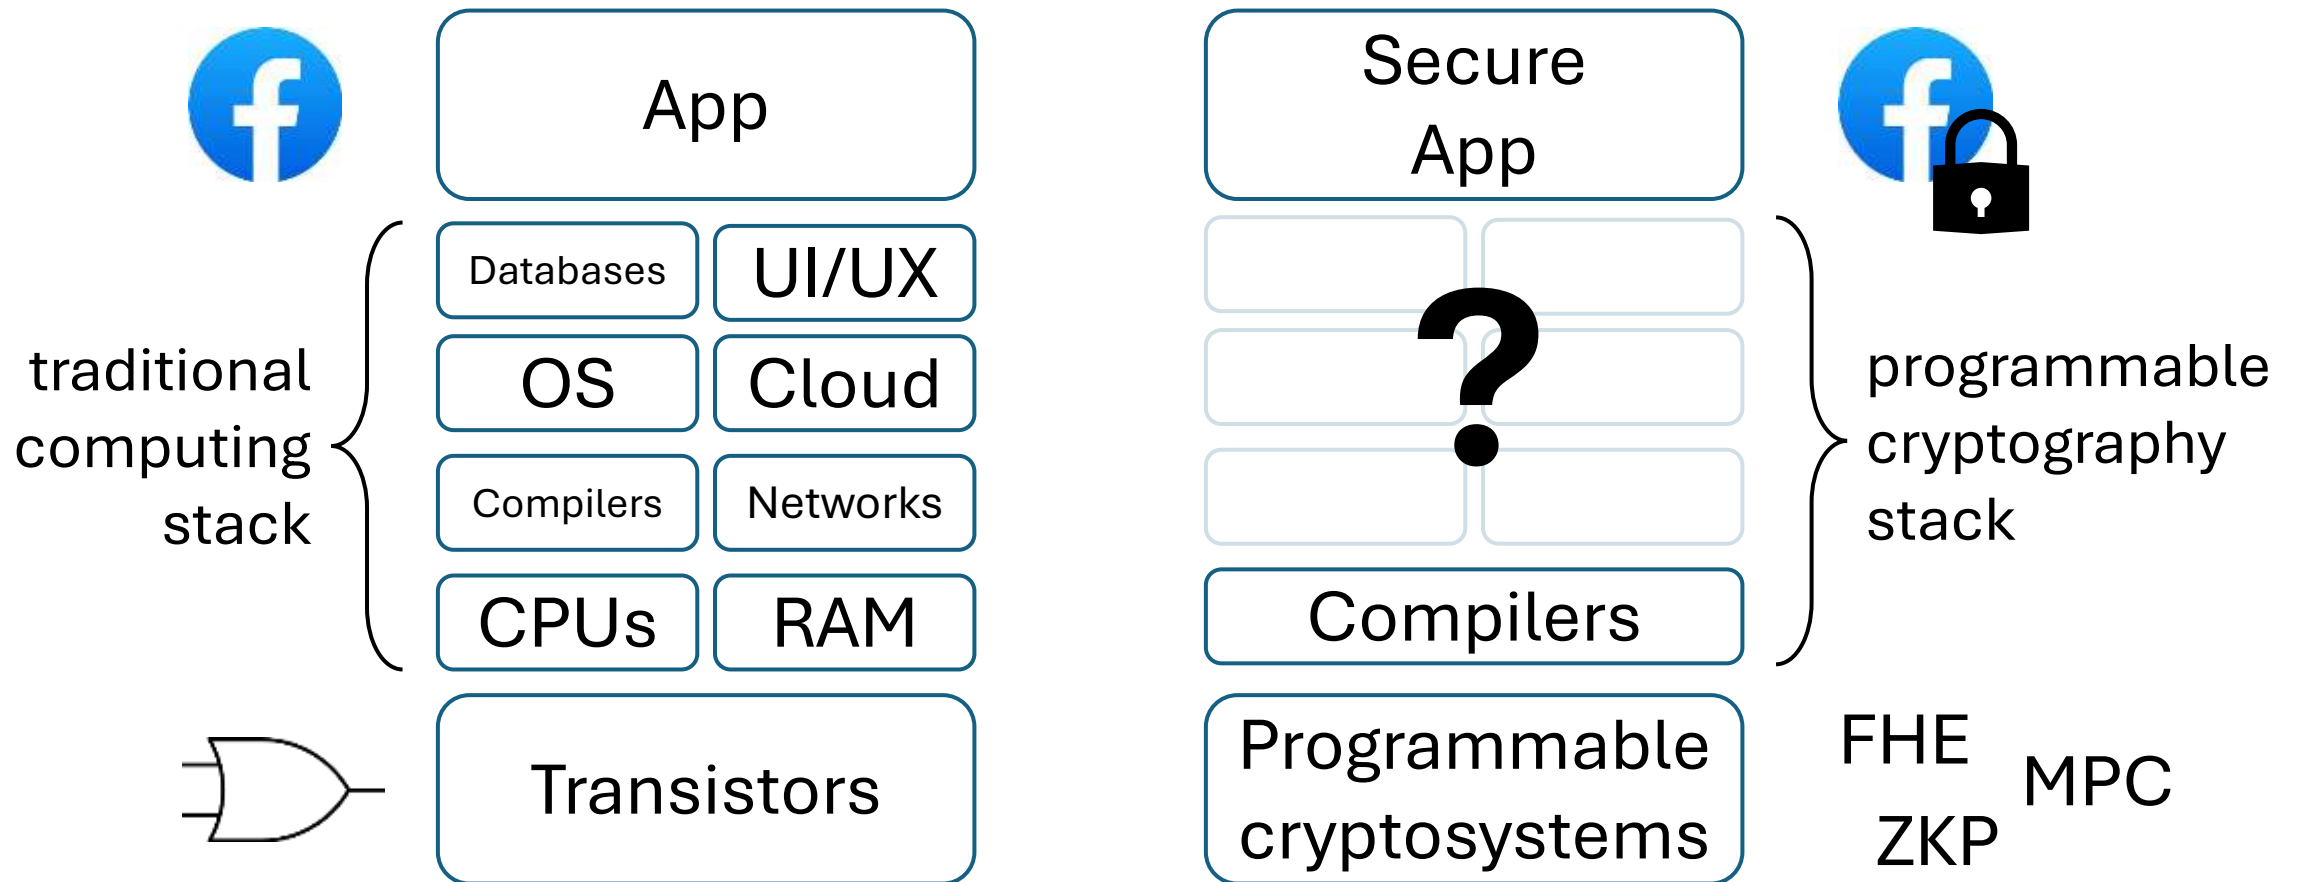


# Programmable cryptography is general

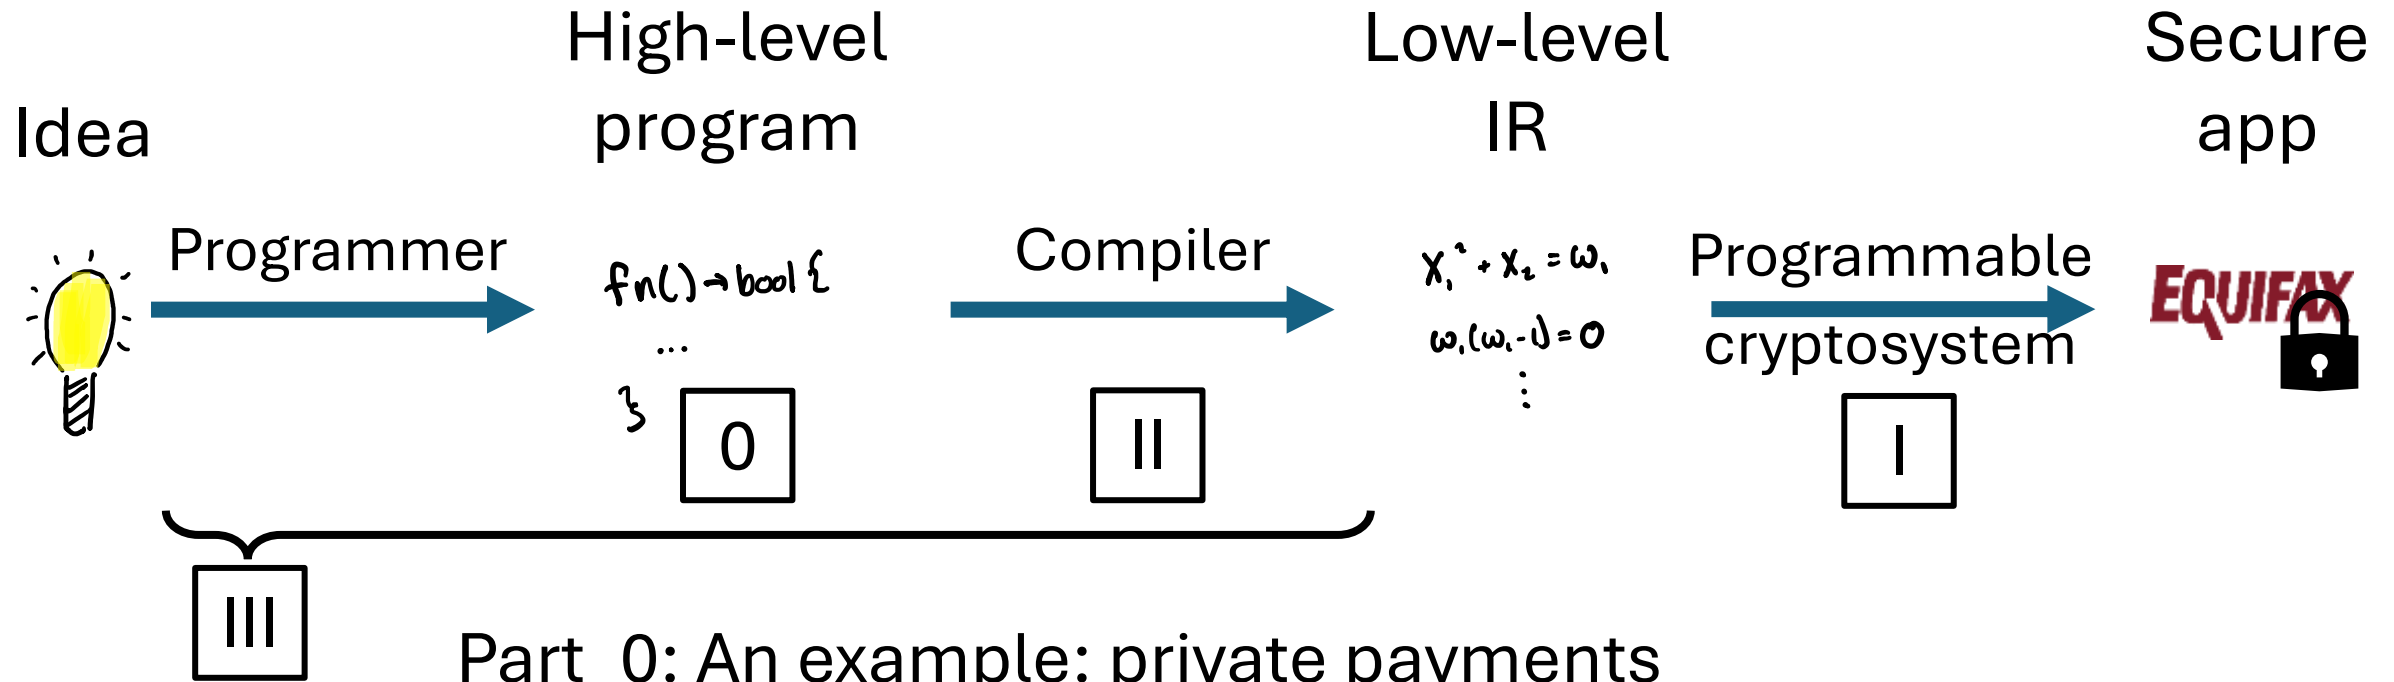
A **programmable cryptosystem** securely executes a **developer-defined program** on secret data.



# Programmable cryptography needs a **stack**



# My PhD work:



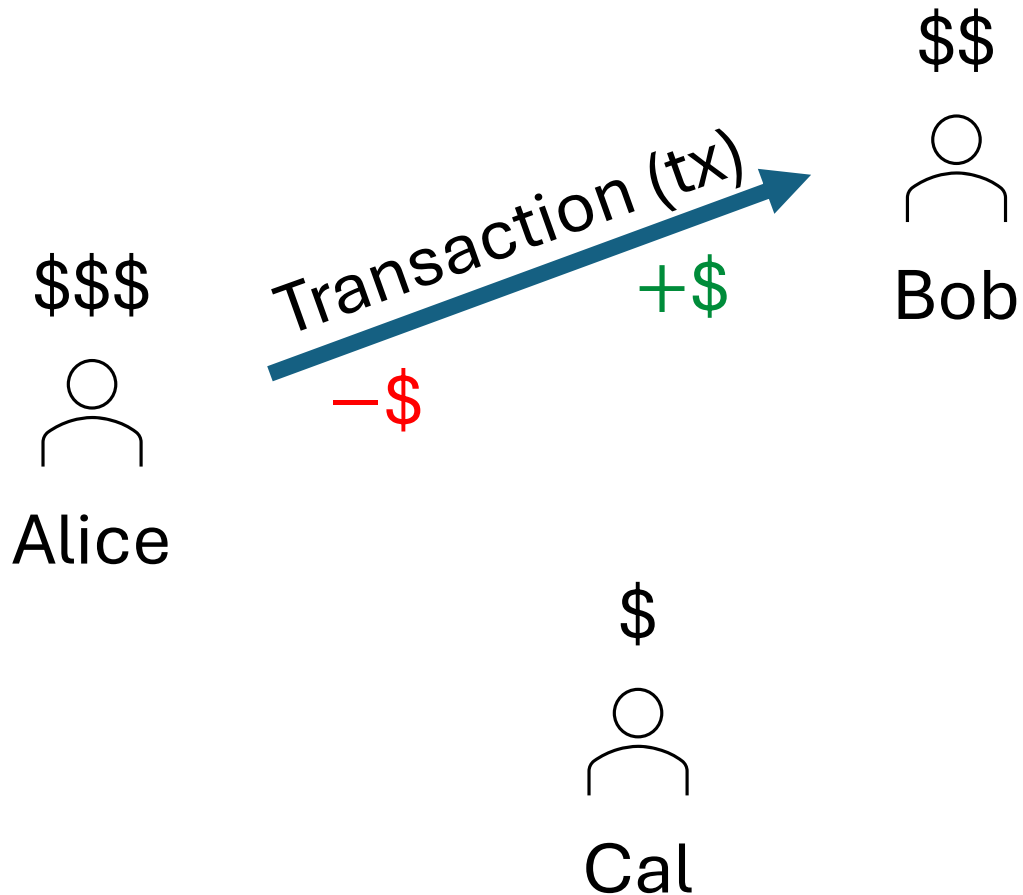
Part 0: An example: private payments

Part I: Collaborative zero knowledge (cryptography)

Part II: Common compiler infrastructure (compilers)

Part III: Verification via finite field solvers (verification)

# Running Example: Private Payments



Security properties:

## Privacy:

- only owners know balances
- only participants know tx amounts

## Integrity:

- money is conserved
- all balances are positive
- participants authorize every tx

# First attempt: payments from a public log

Alice has  $d_A$  \$.

Bob has  $d_B$  \$.

Alice sends  $\delta \leq [0, d_A]$  \$.

Public log

$A: d_A$

$B: d_B$

$A: d_A - \delta$

$B: d_B + \delta$

$\sigma_A, \sigma_B$

all data is public  
 $\Rightarrow$  no *privacy*

users check log  
 $\Rightarrow$  easy *integrity*

# Second attempt: just encrypt the log?

Alice has  $d_A$  \$.

$$e_A = E(k_A, d_A)$$

Alice sends  $\delta \leq [0, d_A]$  \$.

$$e'_A = E(k_A, d_A - \delta)$$

Public log

$e_A$

$e_B$

$e'_A$

$e'_B$

$\sigma_A, \sigma_B$

Bob has  $d_B$  \$.

$$e_B = E(k_B, d_B)$$

$$e'_B = E(k_B, d_B + \delta)$$

What if


$$e'_B = E(k_B, d_B + 2\delta)?$$

Balances is encrypted  
 $\Rightarrow$  easy *privacy*

Can't validate txs  
 $\Rightarrow$  no *integrity*

# Fixing it, with programmable cryptography

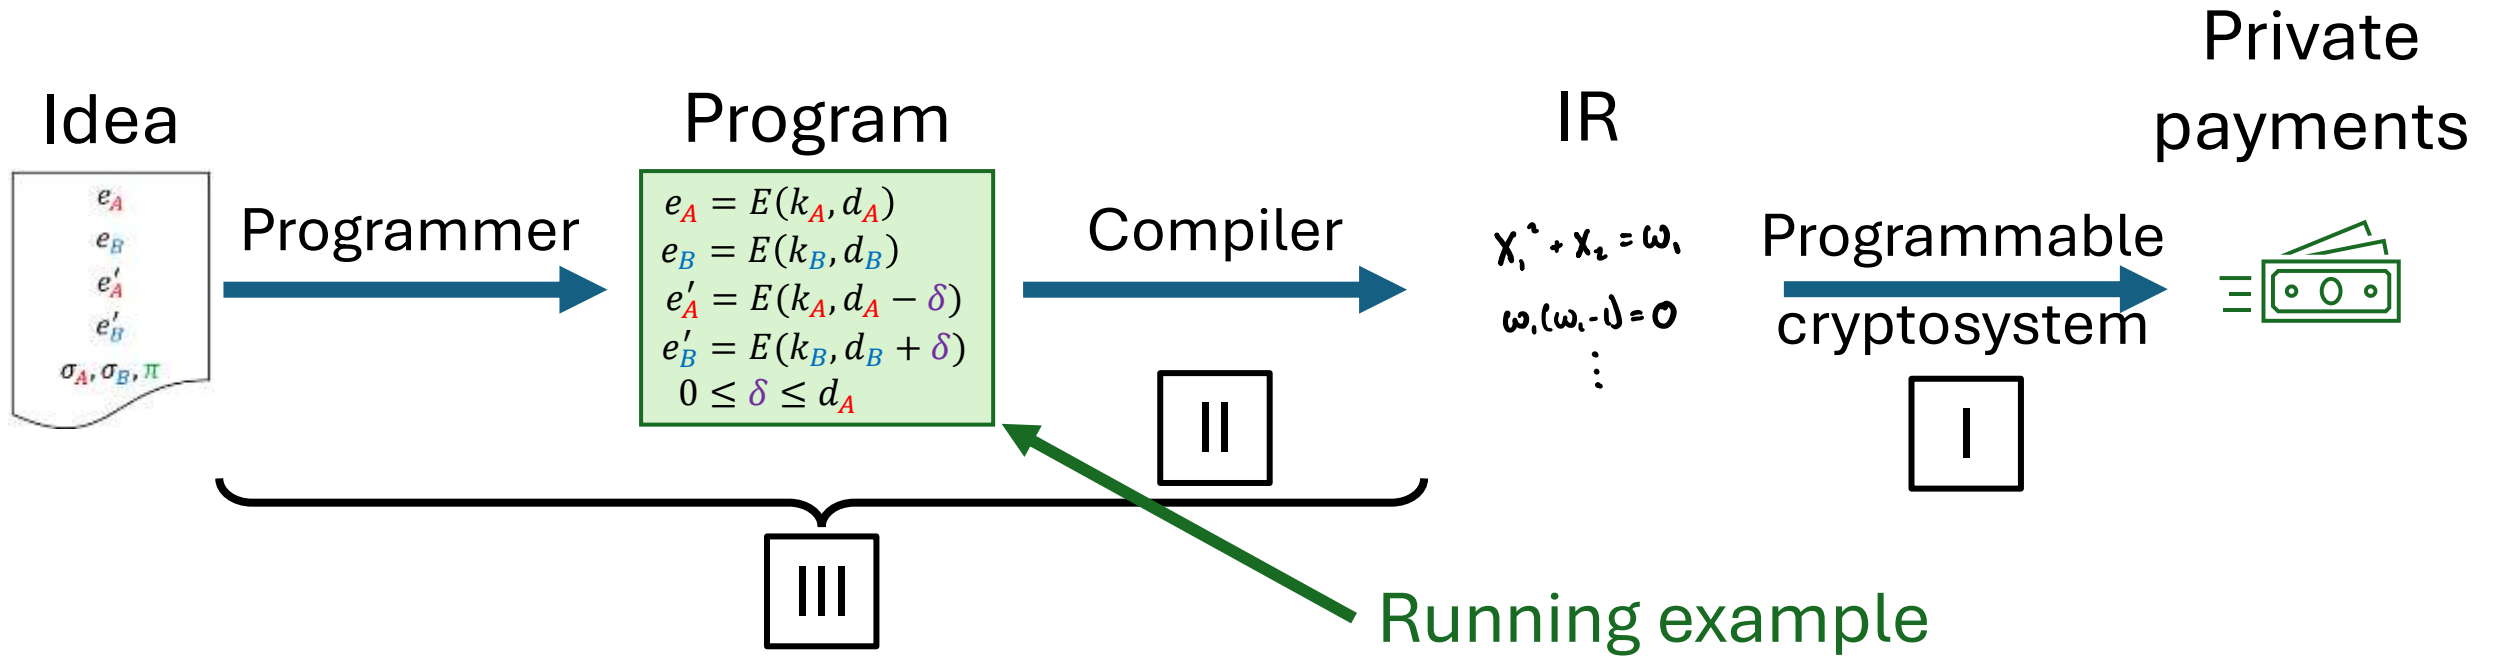
Public log



Idea: add a zero knowledge proof (ZKP)  $\pi$  that the updates are correct.

$$\begin{aligned}e_A &= E(k_A, d_A) \\e_B &= E(k_B, d_B) \\e'_A &= E(k_A, d_A - \delta) \\e'_B &= E(k_B, d_B + \delta) \\0 &\leq \delta \leq d_A\end{aligned}$$

# The full stack of programmable cryptography

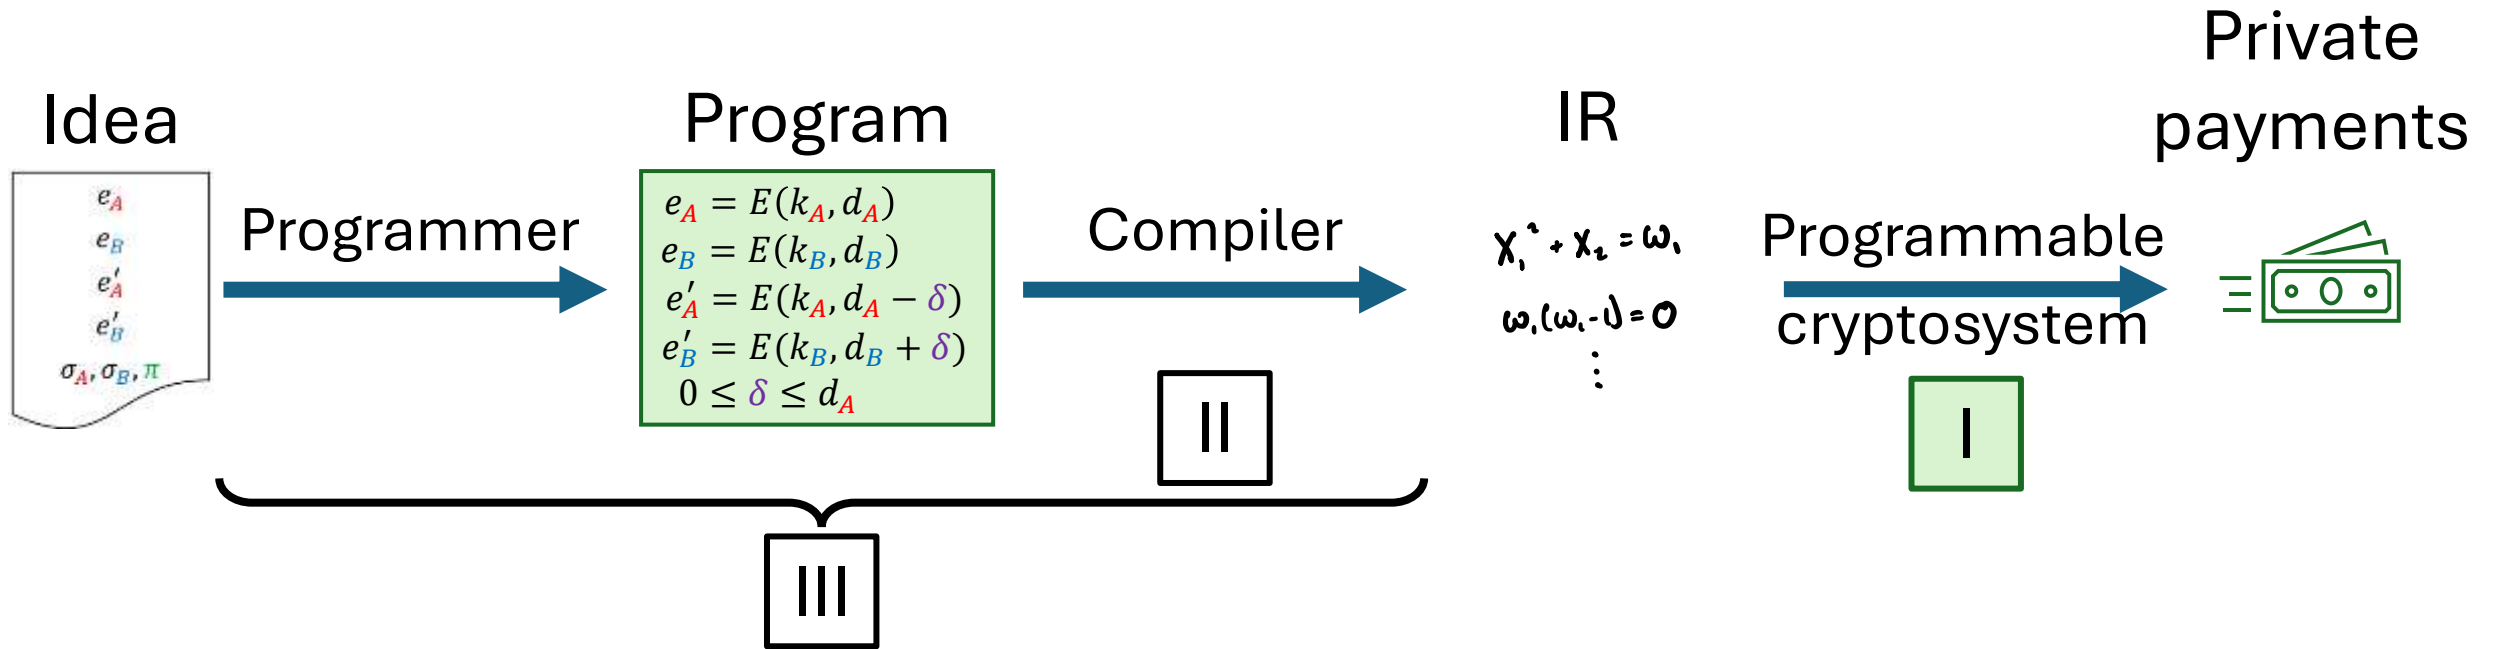


Part I: Collaborative zero knowledge

Part II: Common compiler infrastructure

Part III: Verification via finite field solving

# The full stack of programmable cryptography



Part I: Collaborative zero knowledge

[Ozdemir, Boneh; USENIX Security '22]

Part II: Common compiler infrastructure

Part III: Verification via finite field solving

# What is a traditional ZKP?

Consider a property  $P(x, w)$

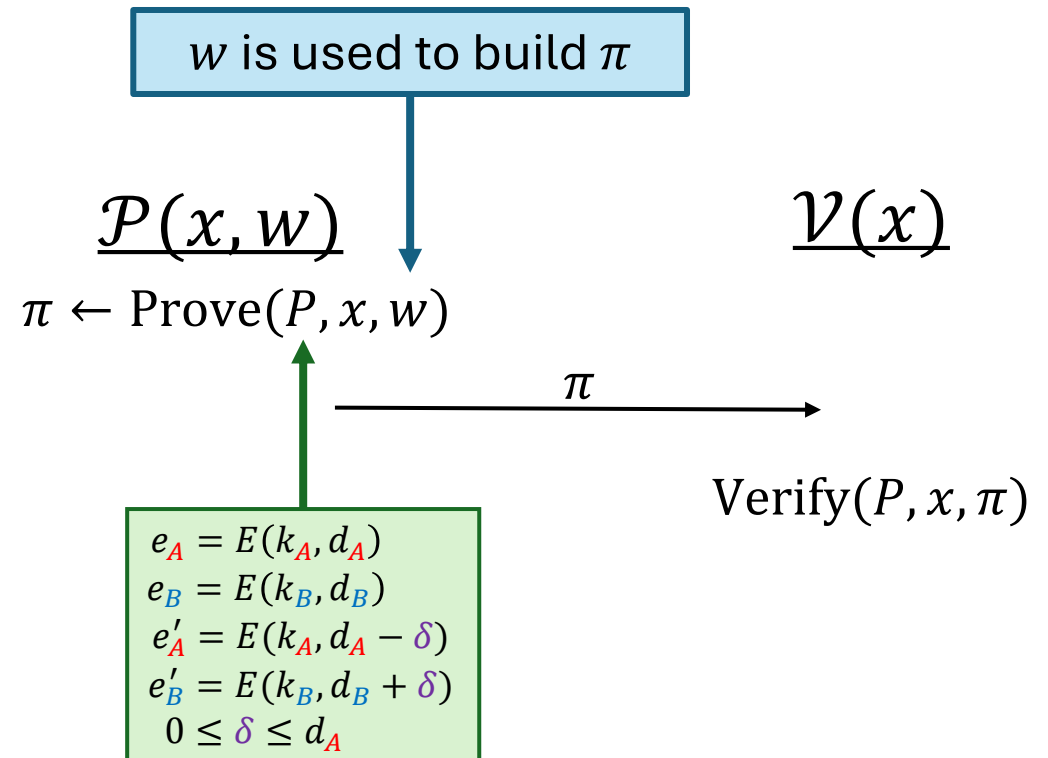
public data  $x$      $(e_A, e_B, e'_A, e'_B)$   
 secret data  $w$      $(d_A, d_B, \delta, k_A, k_B)$

A ZKP is two algorithms:

$\text{Prove}(P, x, w) \rightarrow \pi$   
 $\text{Verify}(P, x, \pi) \rightarrow \{\text{accept, reject}\}$

Security properties (informal):

complete:            a valid  $w$  yields a valid  $\pi$   
 sound:                constructing a valid  $\pi$  is infeasible without a valid  $w$   
 zero-knowledge:     $\pi$  reveals no information about  $w$



# Problem: who creates the ZKP?

$w = (d_A, d_B, \delta, k_A, k_B)$  is used to build  $\pi$



Alice knows some of  $w$ .

Bob knows some of  $w$ .

Problem: no **single party** can create a ZKP,  
because the secrets are **distributed**.


Sending  $w$  to one party is **not private**.

Contributions:

1. A new definition: the **collaborative ZKP**
2. A very efficient construction

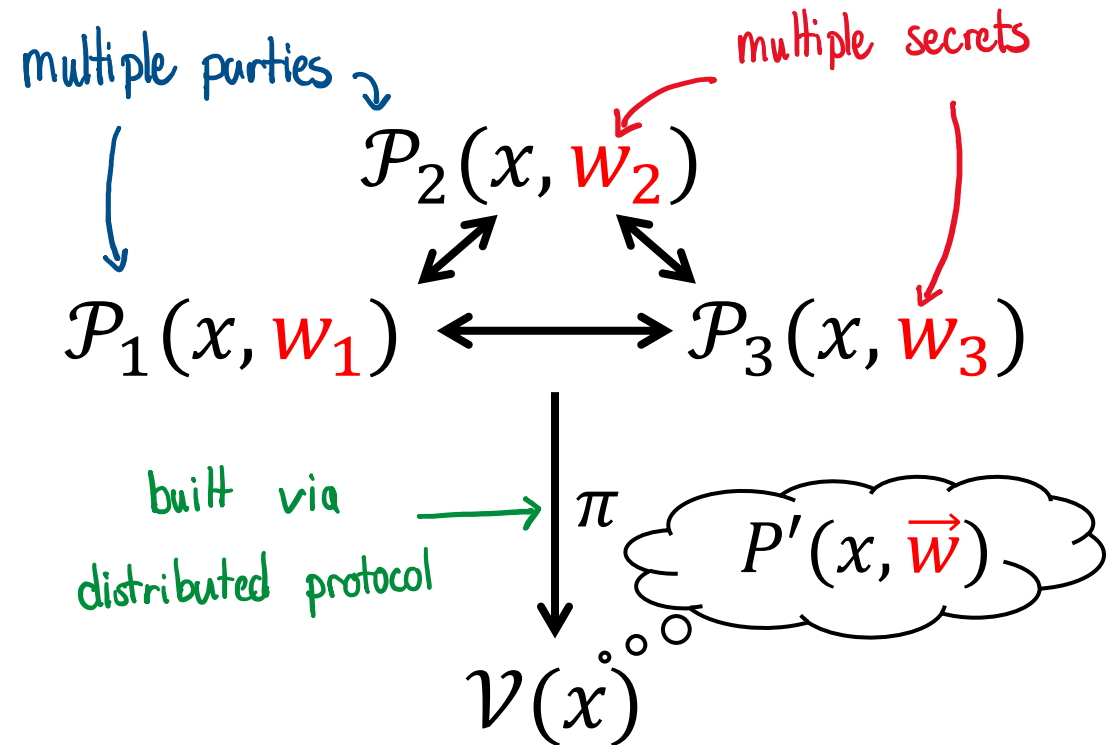
# Our definition: Collaborative ZKPs

Prior: ZKPs



sound, complete, hides  $w$  from  $\mathcal{V}$

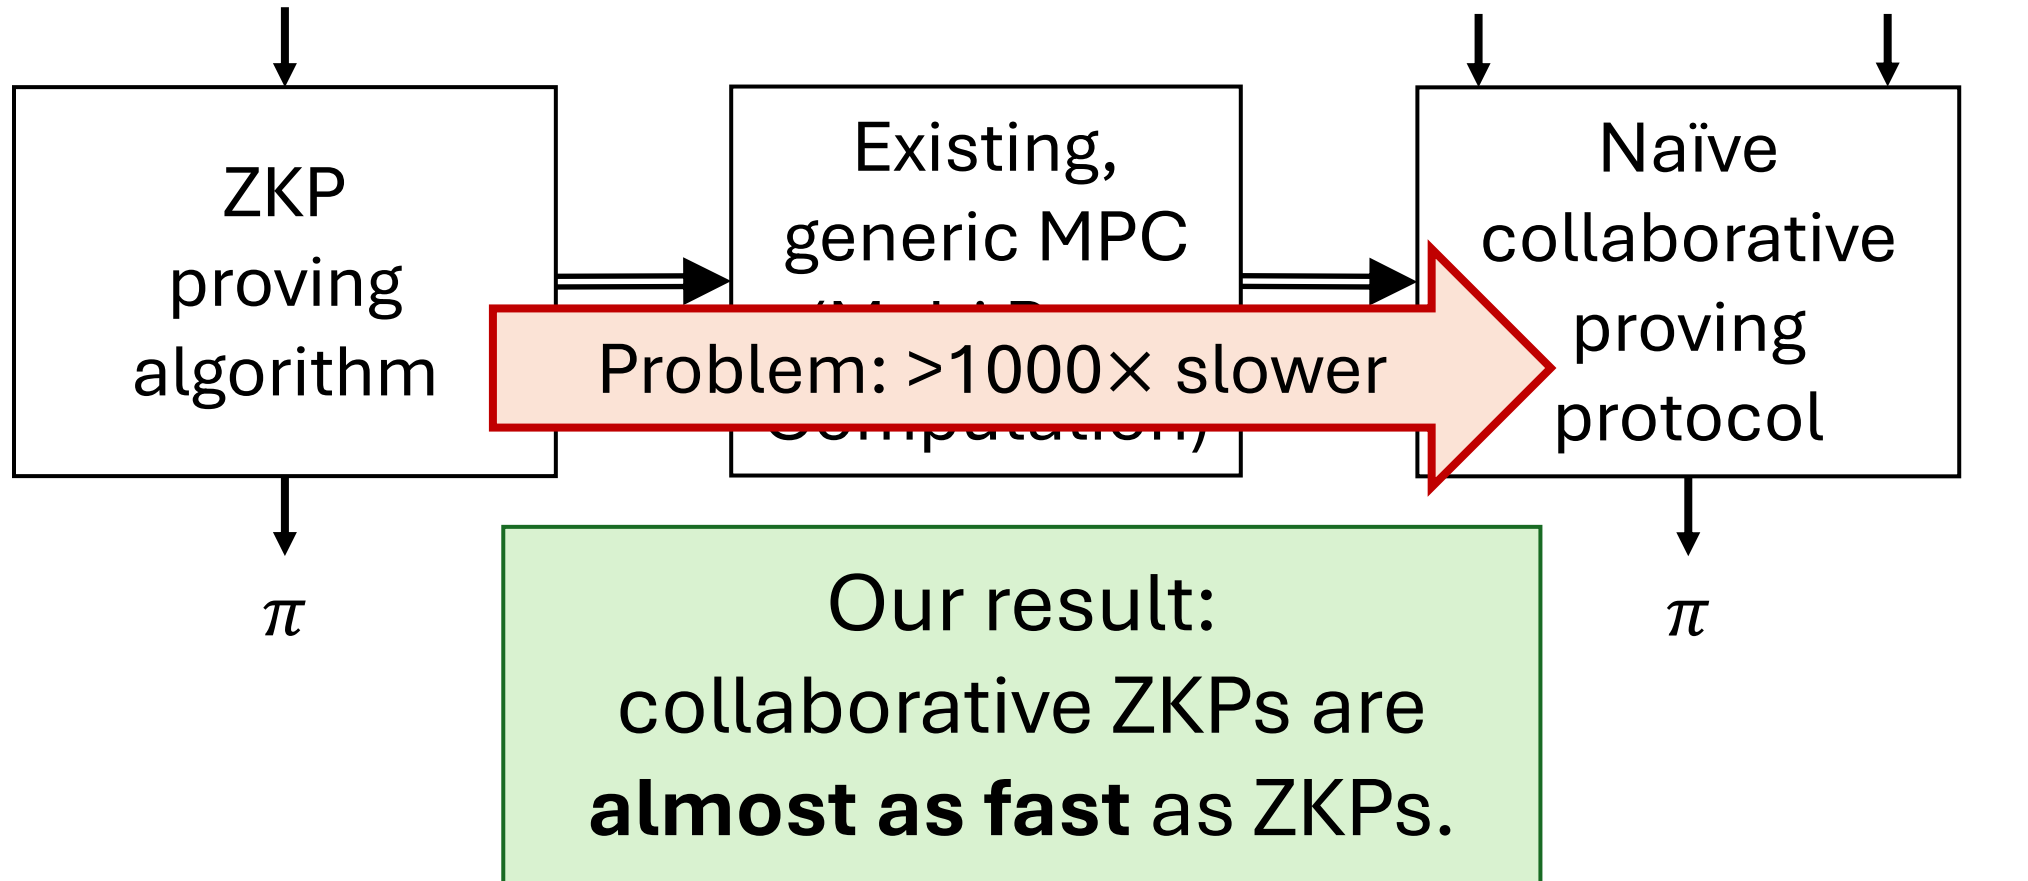
Collaborative ZKPs



also hides  $w_i$  from  $\mathcal{P}_j$

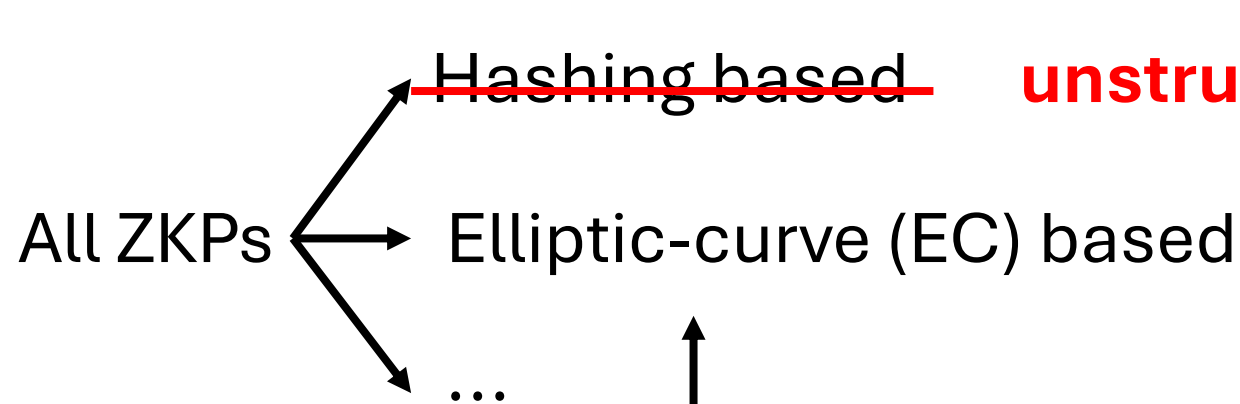
# A naïve protocol

$P', x, w = (w_1, w_2)$



# Choosing the right starting point

Step 1. Start from the right ZKP



An EC: a set of  $p$  points  $G, H, \dots$  with

- addition ( $G \boxplus H$ )
- multiplication by scalars ( $z \boxdot G$ )
  - $z \in \mathbb{Z}_p, \mathbb{Z}_p = \{0, \dots, p - 1\}$

the slowest part  
of proving

algebraic bottleneck

$$H \leftarrow \boxplus_{i=1}^N (z_i \boxdot G_i)$$

scalar  
**secret**  
(encodes  $\vec{w}$ )

EC point  
**public**  
(encodes  $P$ )

# Intuition: how to handle the bottleneck

Step 2. Solve the bottleneck with  $\mathbb{Z}_p$  **and EC** secret sharing

$$H \leftarrow \boxplus_{i=1}^N (z_i \boxdot G_i)$$

Local computation

Result:

Prover 1:

$\forall i$ , gets random  $z_{i,1}$

$$H_1 \leftarrow \boxplus_{i=1}^N (z_{i,1} \boxdot G_i)$$

Prover 2:

$\forall i$ , gets  $z_{i,2} \leftarrow z_i - z_{i,1}$

$$H_2 \leftarrow \boxplus_{i=1}^N (z_{i,2} \boxdot G_i)$$

$$H = H_1 \boxplus H_2$$

(note:  $z_i = z_{i,1} + z_{i,2}$ )

$\underbrace{\hspace{10em}}$   
A  $\mathbb{Z}_p$ -linear sharing of  $z_i$

**Evaluated the bottleneck  
without communication**

$\underbrace{\hspace{10em}}$   
An EC-linear sharing of  $H$

# Intuition: how to handle the bottleneck

Step 2. Solve the bottleneck with  $\mathbb{Z}_p$  and EC secret sharing

$$H \leftarrow \boxplus_{i=1}^N (z_i \boxtimes G_i)$$

## Zooming out:

- This is one (key) step in a big protocol
- Malicious security is necessary
- Sketch: lift SPDZ & GSZ'20 to ECs

Result:

Prover 1:  $\forall i$ , gets

Prover 2:  $\forall i$ , gets  $z_{i,2} \leftarrow z_i - z_{i,1}$   $H_2 \leftarrow \boxplus_{i=1}^N (z_{i,2} \boxtimes G_i)$

$$H = H_1 \boxplus H_2$$

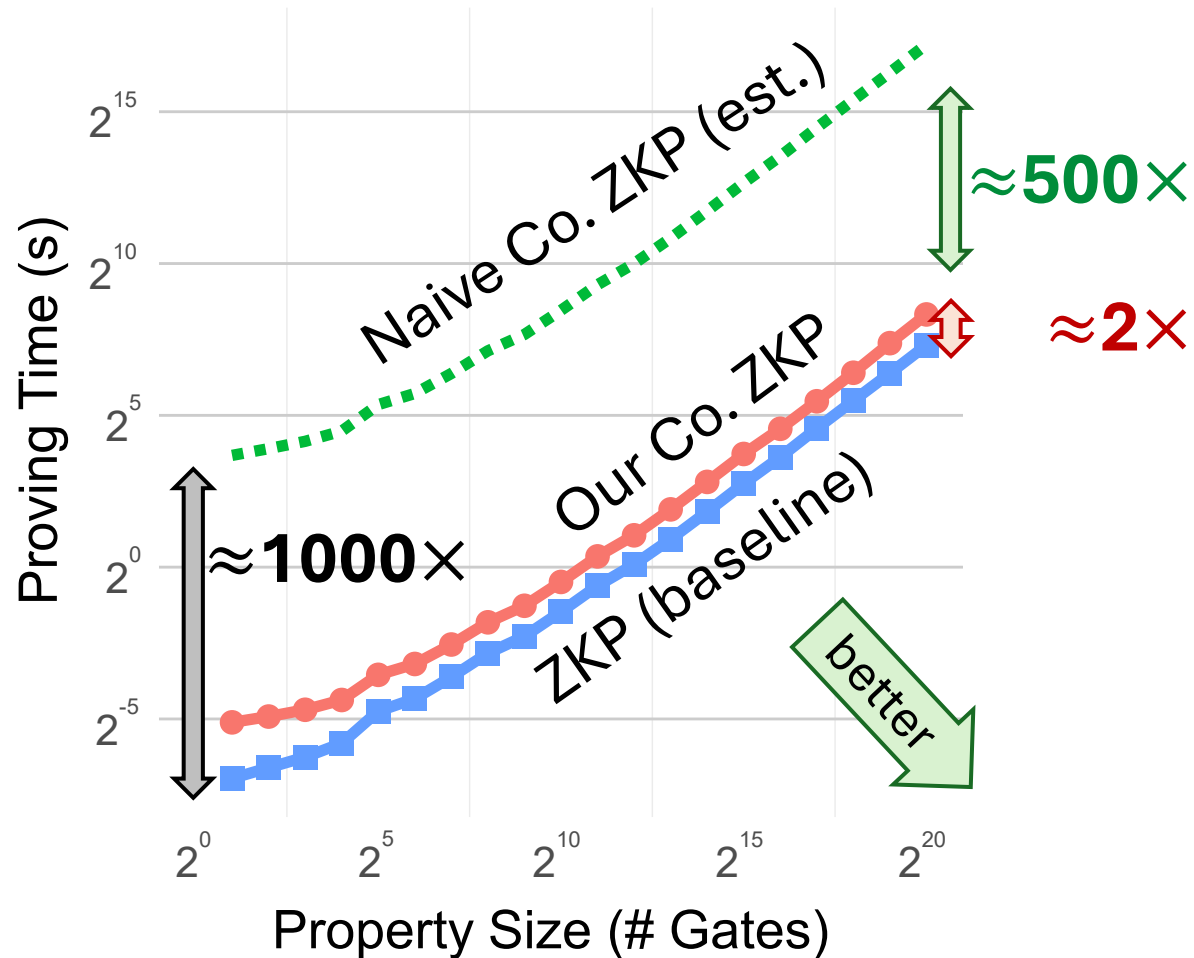
(note:  $z_i = z_{i,1} + z_{i,2}$ )

$\underbrace{\hspace{10em}}$   
A  $\mathbb{Z}_p$ -linear sharing of  $z_i$

Evaluated the bottleneck  
**without communication**

$\underbrace{\hspace{10em}}$   
An EC-linear sharing of  $H$

# Efficient collaborative ZKPs & their impact



Launched an active research area  
[HKLMRRV'22], [CLM'23], [GGJPS'23], [GJMW'24],  
[DKR'23], [ZFS'24], [GGSSZ'25], [RGGJS'25],...

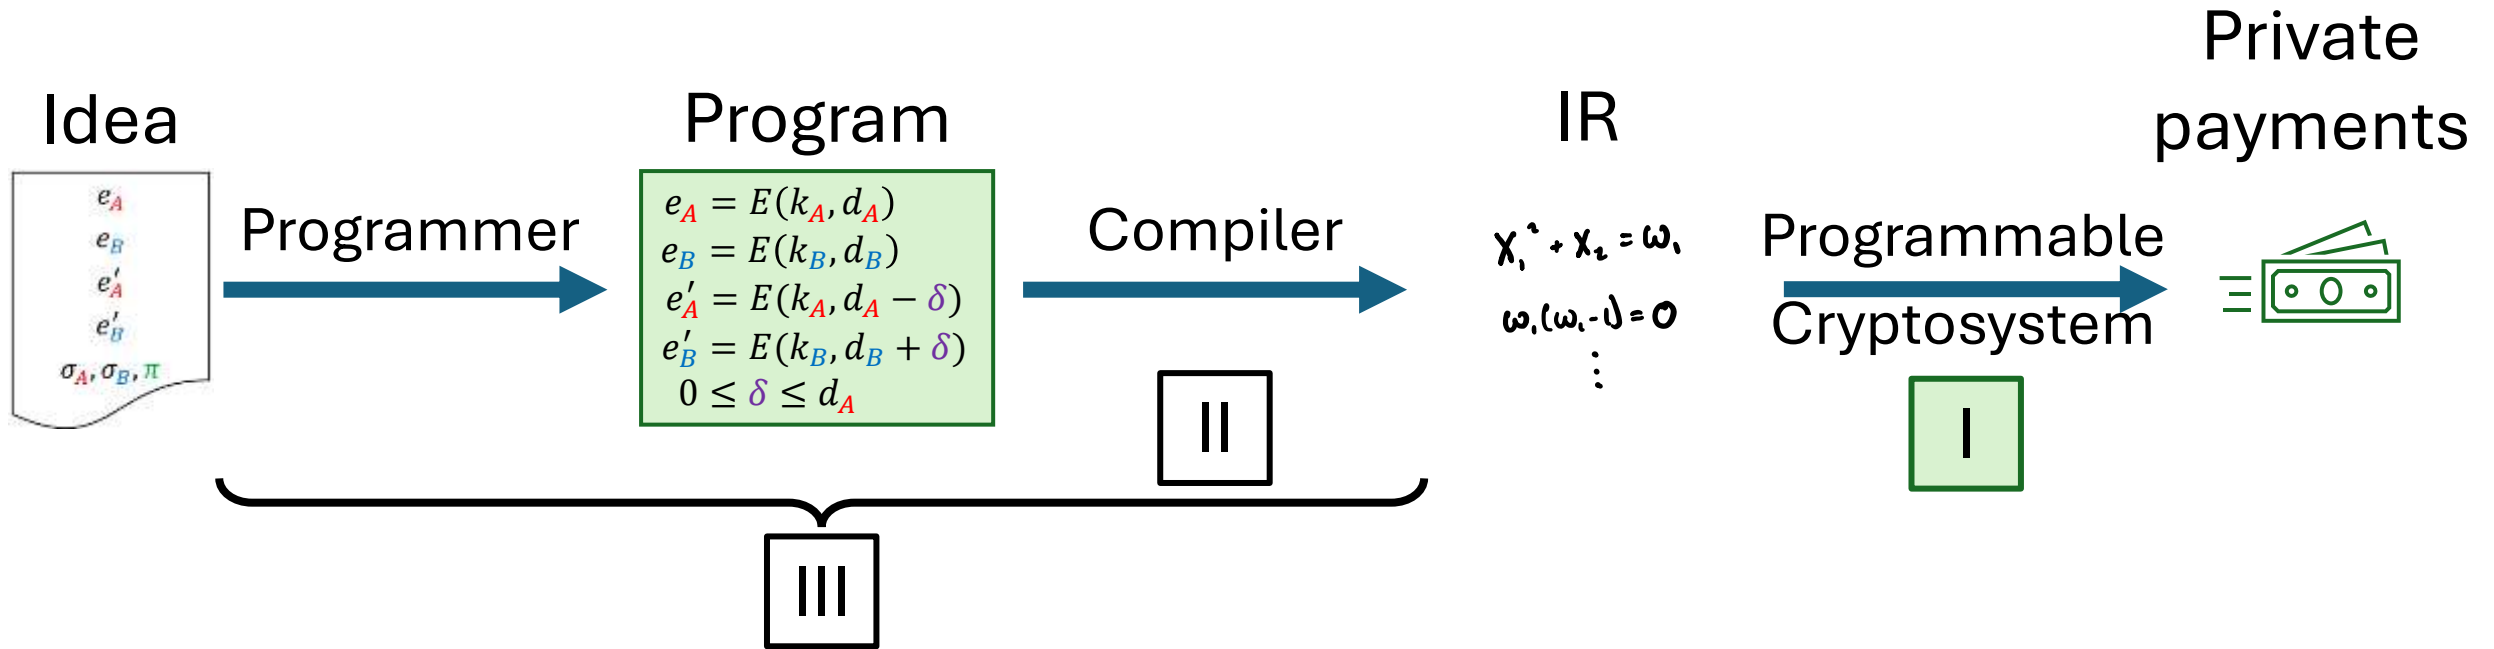
Multiple industrial  
implementations and compilers



Ongoing applications to private  
markets, ZKP outsourcing, ...



# The full stack of programmable cryptography



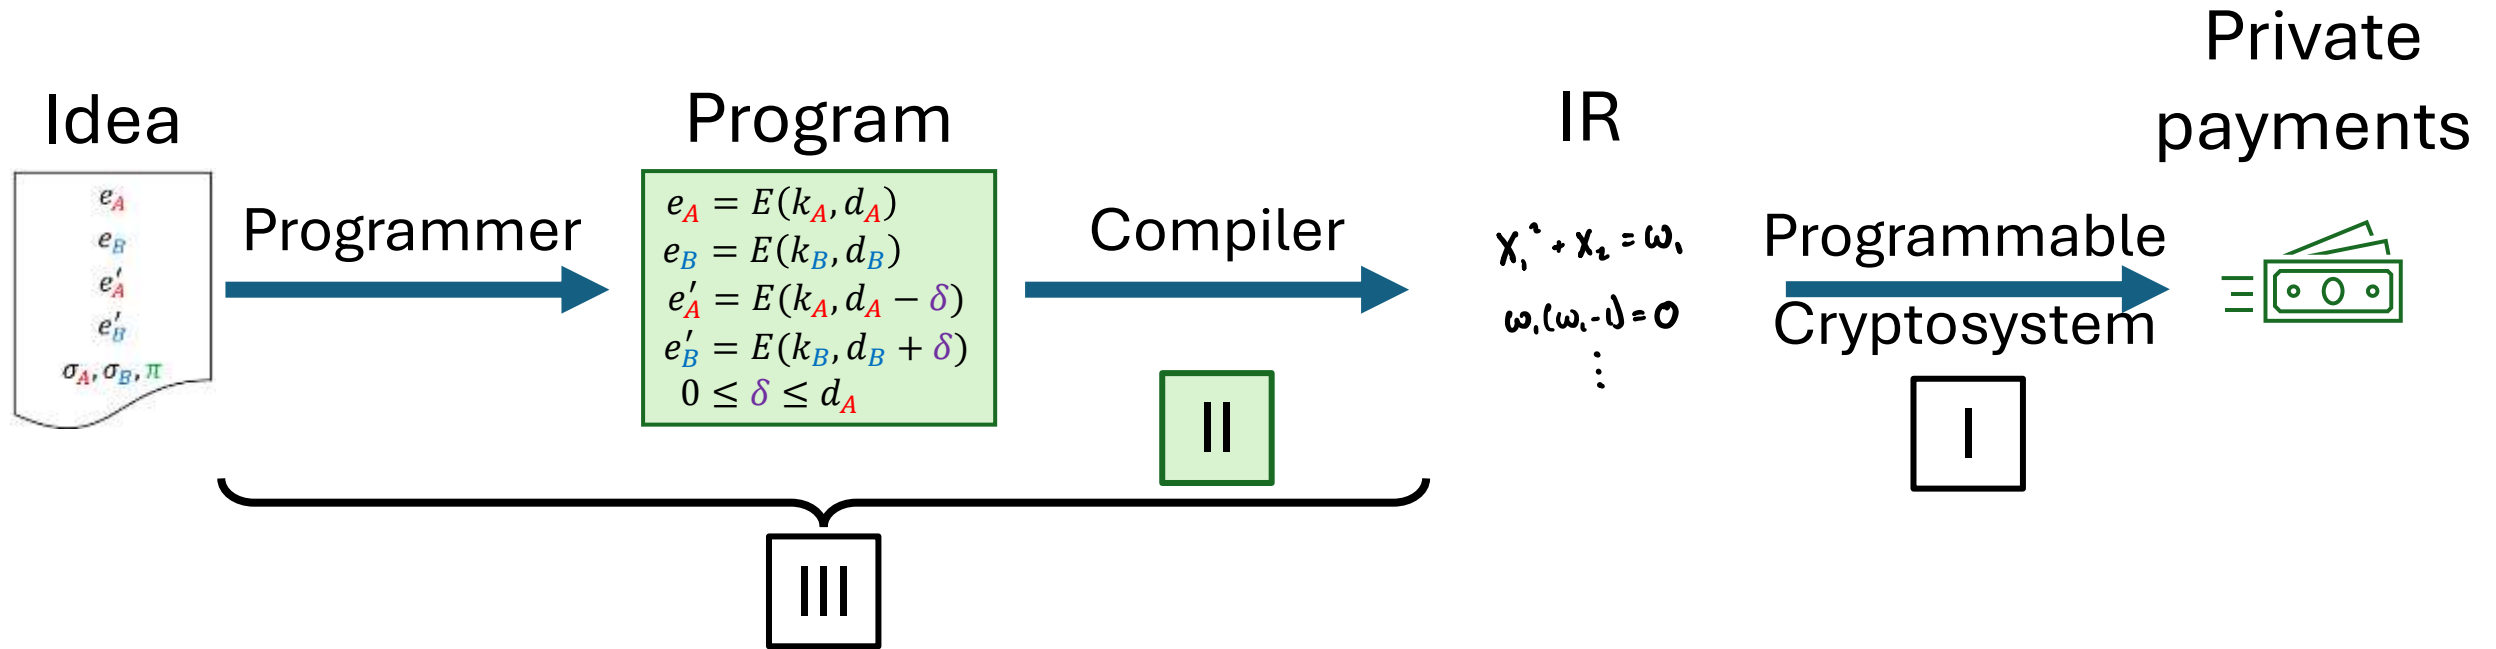
Part I: Collaborative zero knowledge

[Ozdemir, Boneh; USENIX Security '22]

Part II: Common compiler infrastructure

Part III: Verification via finite field solving

# The full stack of programmable cryptography



Part I: Collaborative zero knowledge

Part II: Common compiler infrastructure

[Ozdemir, Brown, Wahby; IEEE S&P '22]

Part III: Verification via finite field solving

# Problem: ZKPs for high-level programs

Complex applications require high-level languages

Even **simple** applications do too:

library function

if-statements

loops

arrays/RAM

custom types

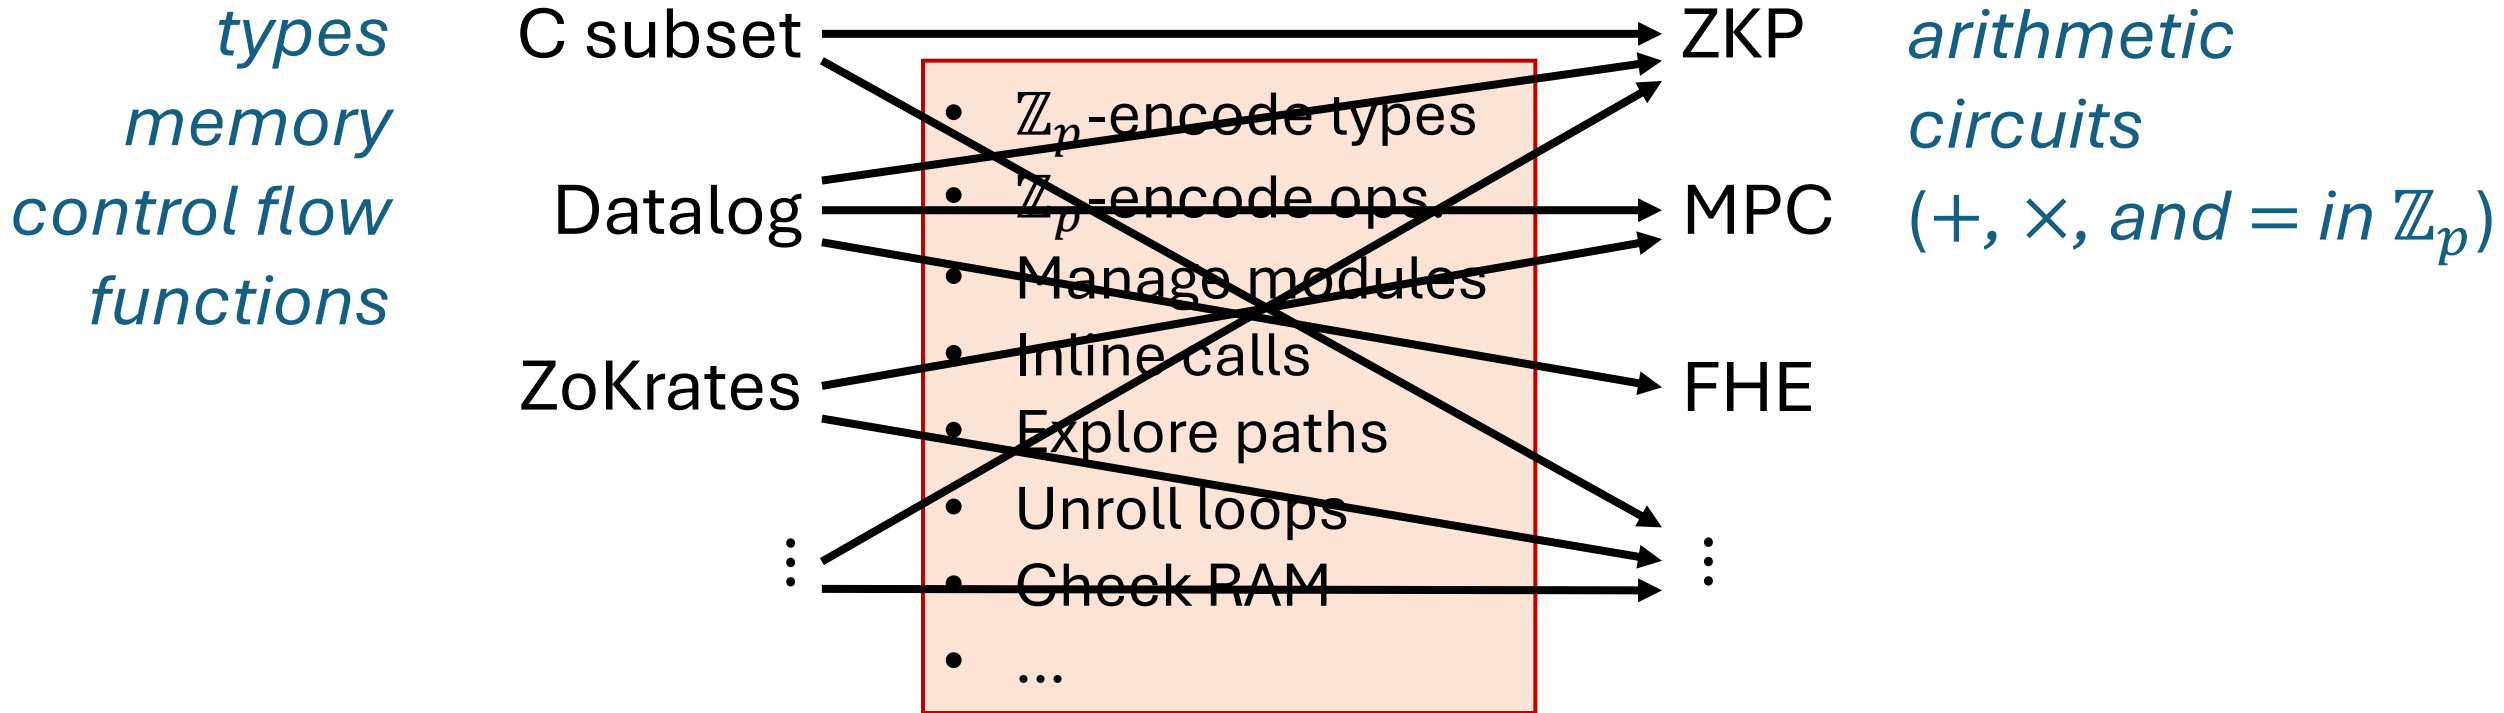
$$\begin{aligned} e_A &= E(k_A, d_A) \\ e_B &= E(k_B, d_B) \\ e'_A &= E(k_A, d_A - \delta) \\ e'_B &= E(k_B, d_B + \delta) \\ 1 &\leq \delta \leq d_A \end{aligned}$$

Booleans

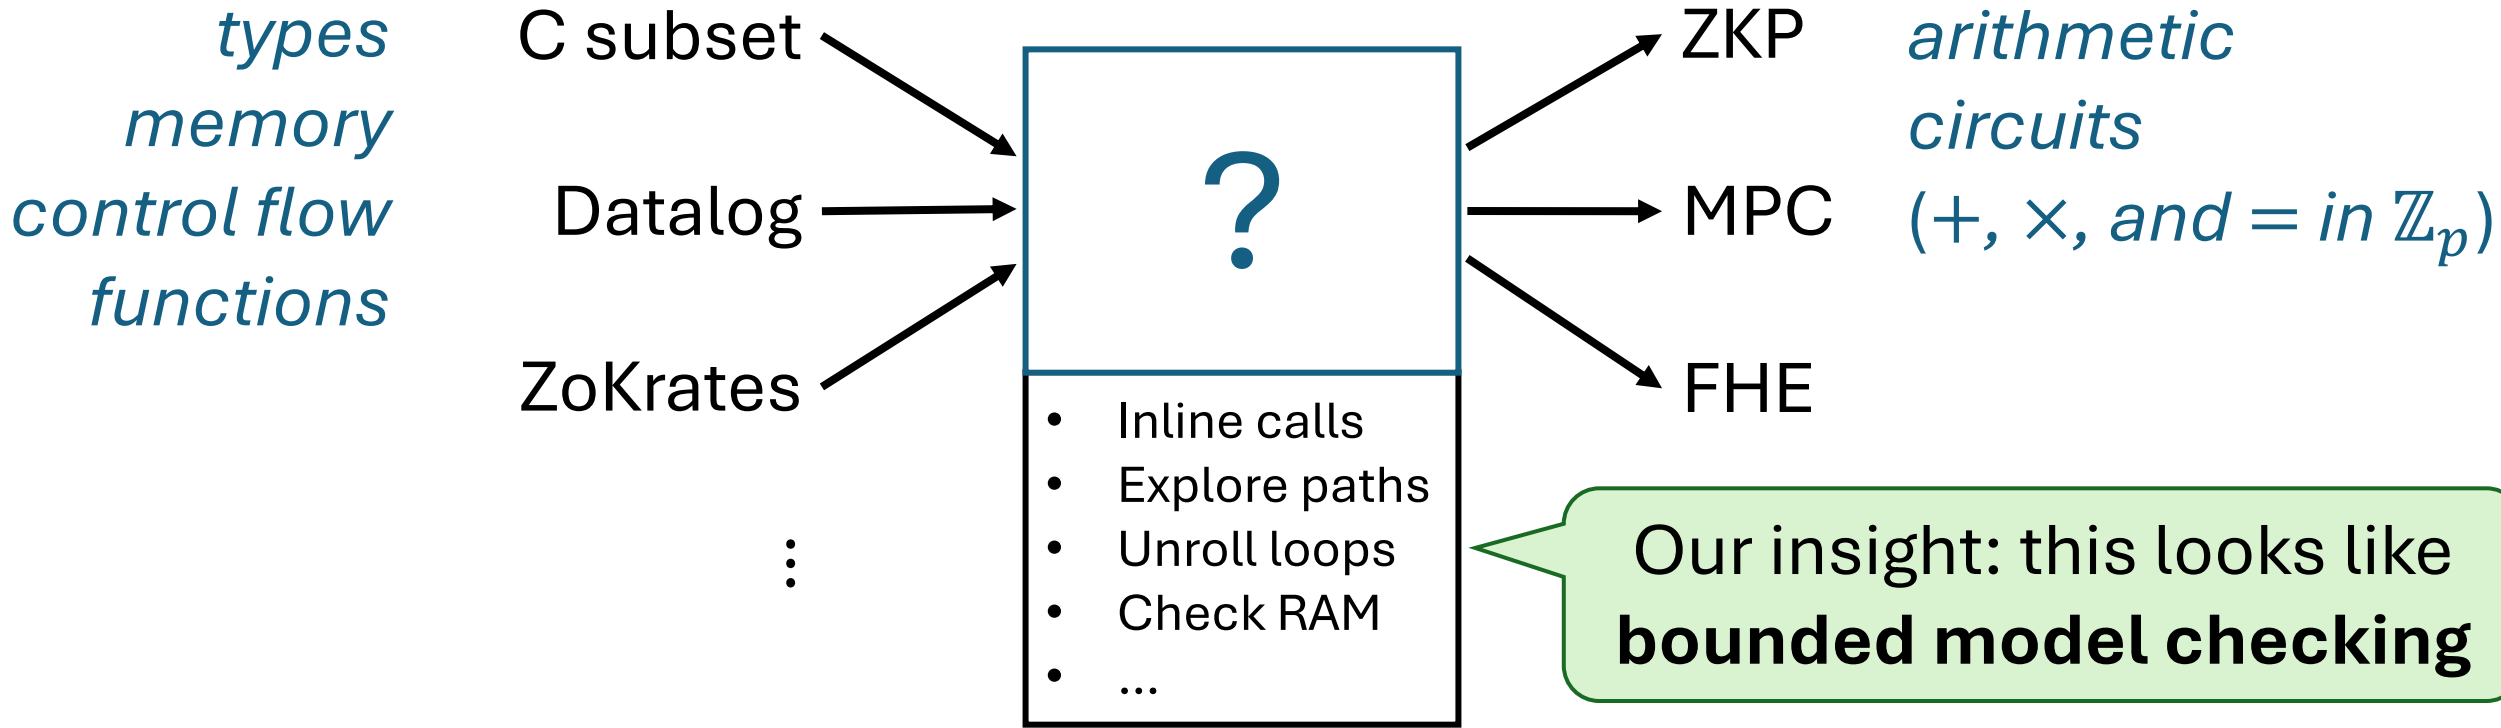
bounded  
integers

bitstrings

# Compilation is hard



# Can we make common infrastructure?



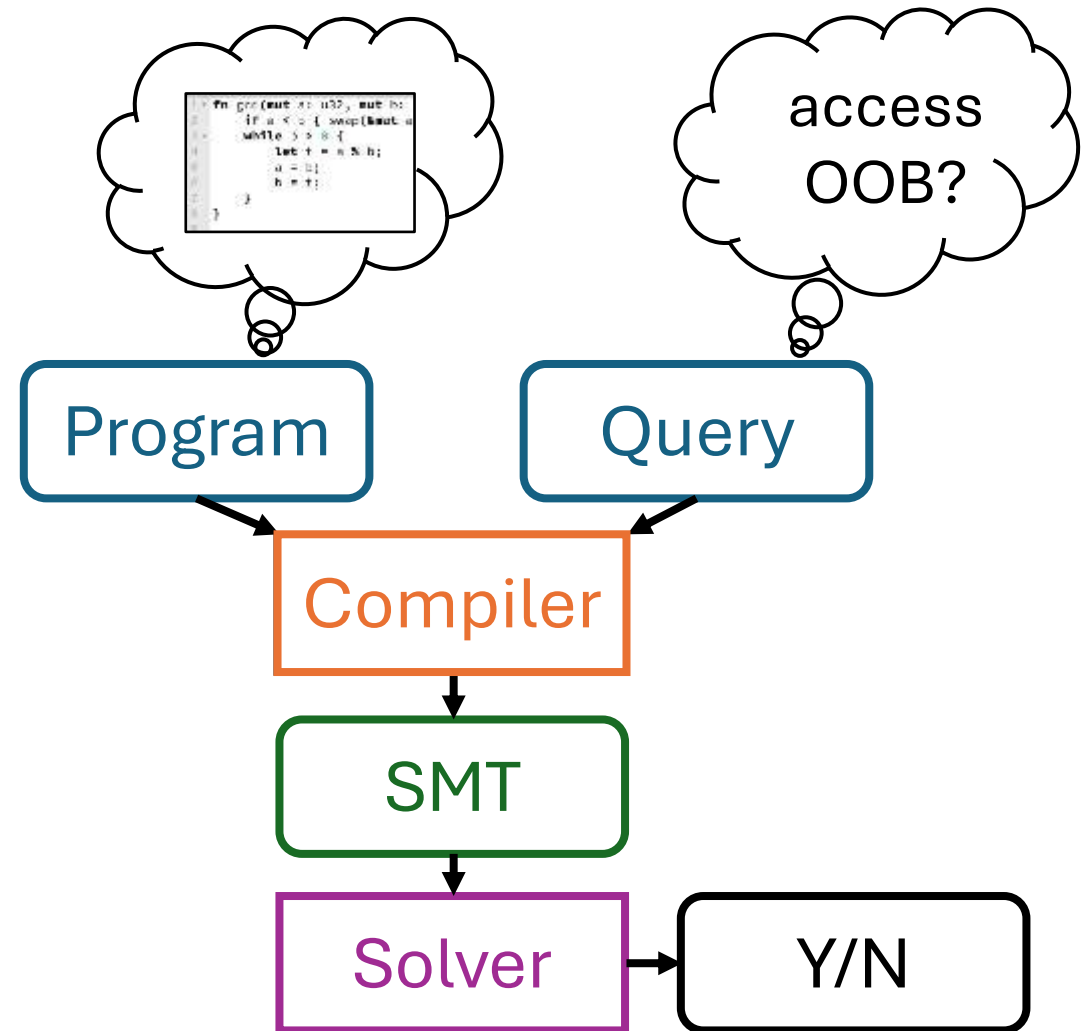
Key challenge: what is the common abstraction?

# BMC: bounded model checking is **compilation**

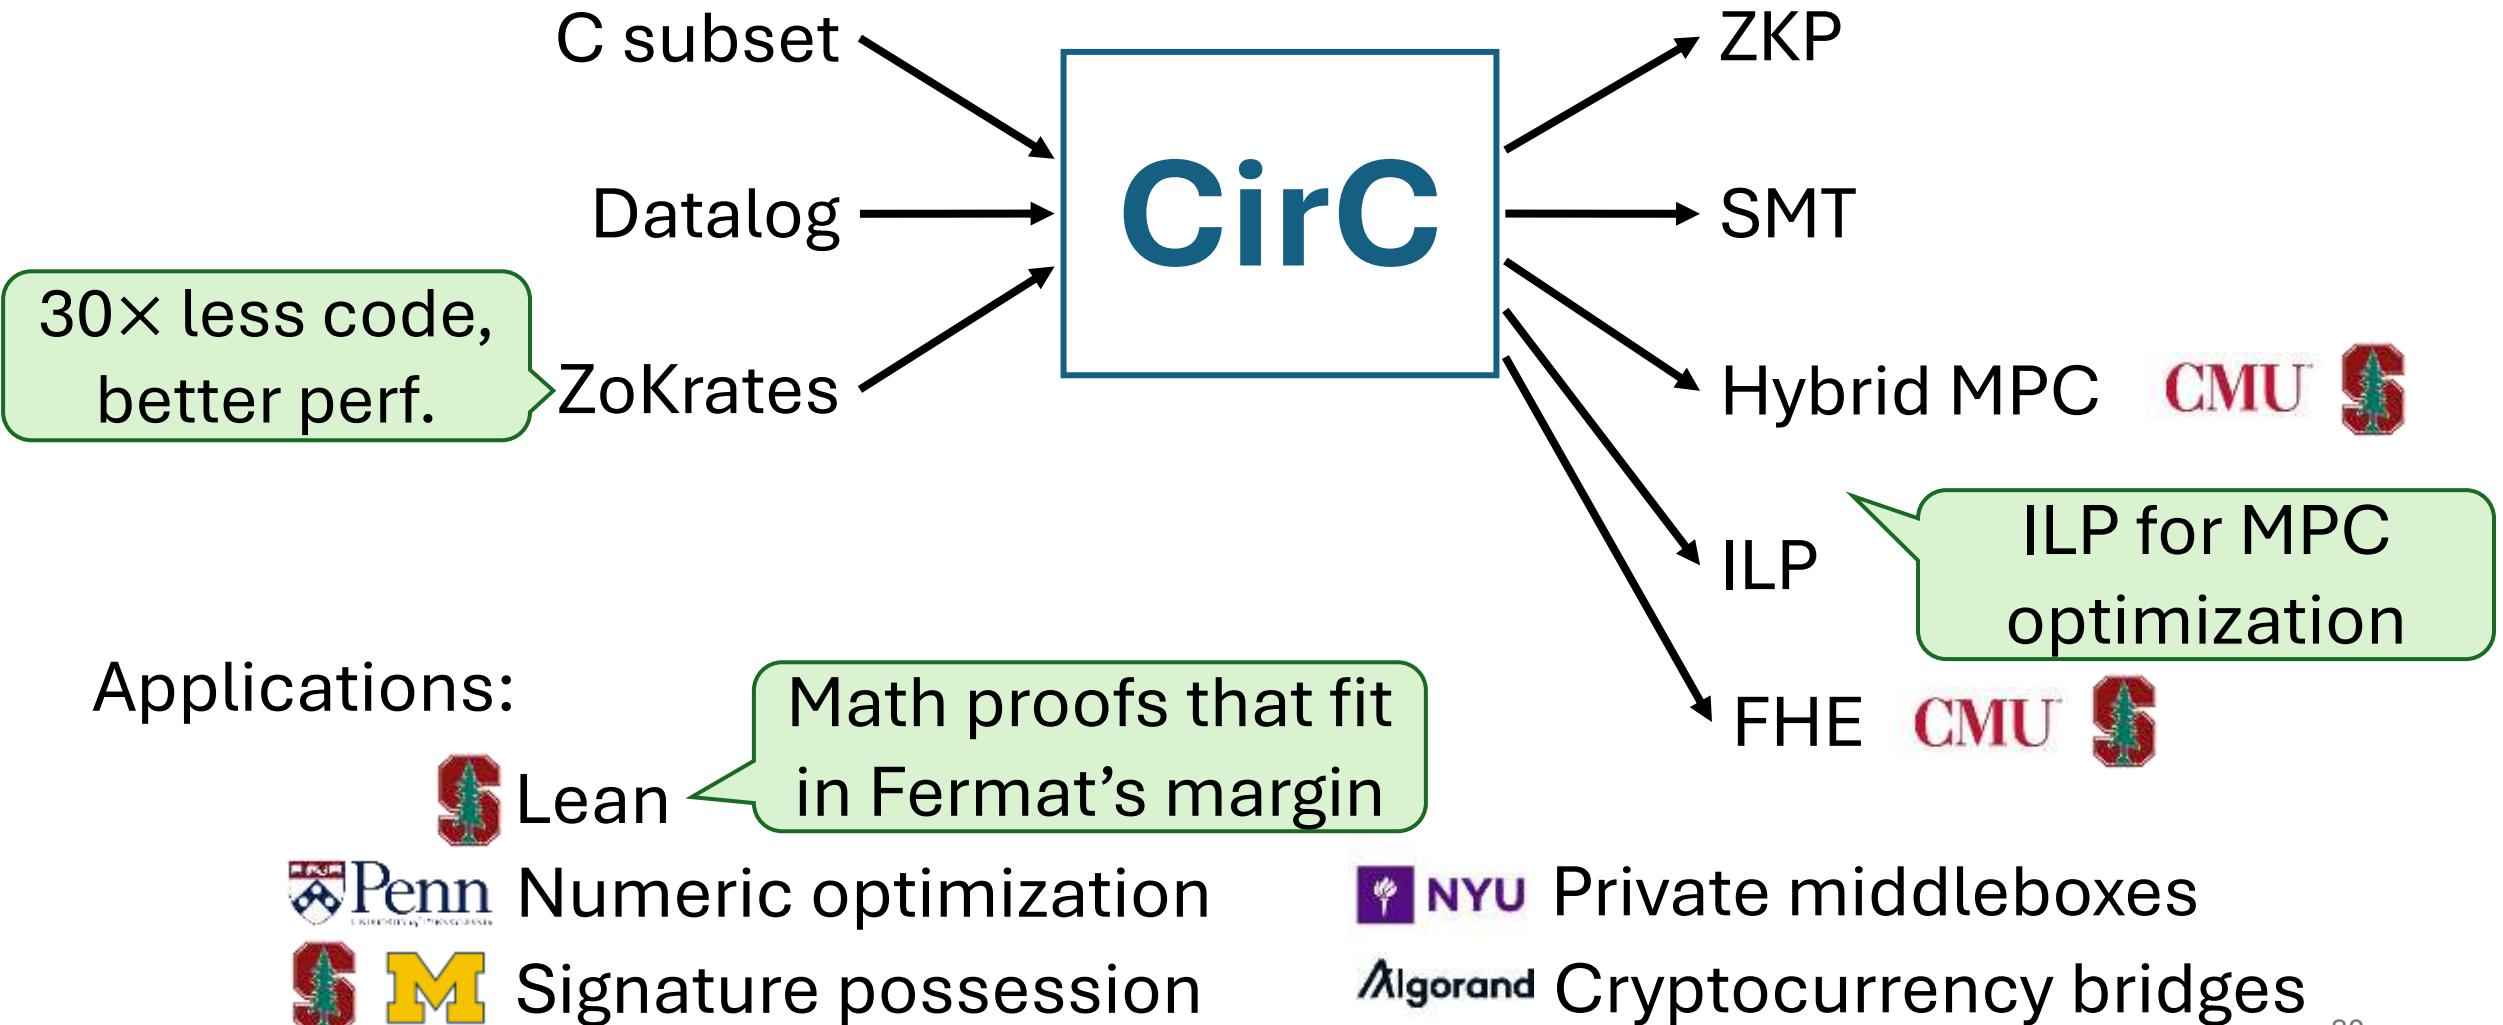
1. Model the **program** and **query** as an **SMT logical formula**

SMT: Satisfiability Modulo Theories  
(a many-typed formula)


2. Check the formula with a **solver**



# Building CirC (**C**ircuit **C**ompiler infrastructure)



# The full stack of programmable cryptography



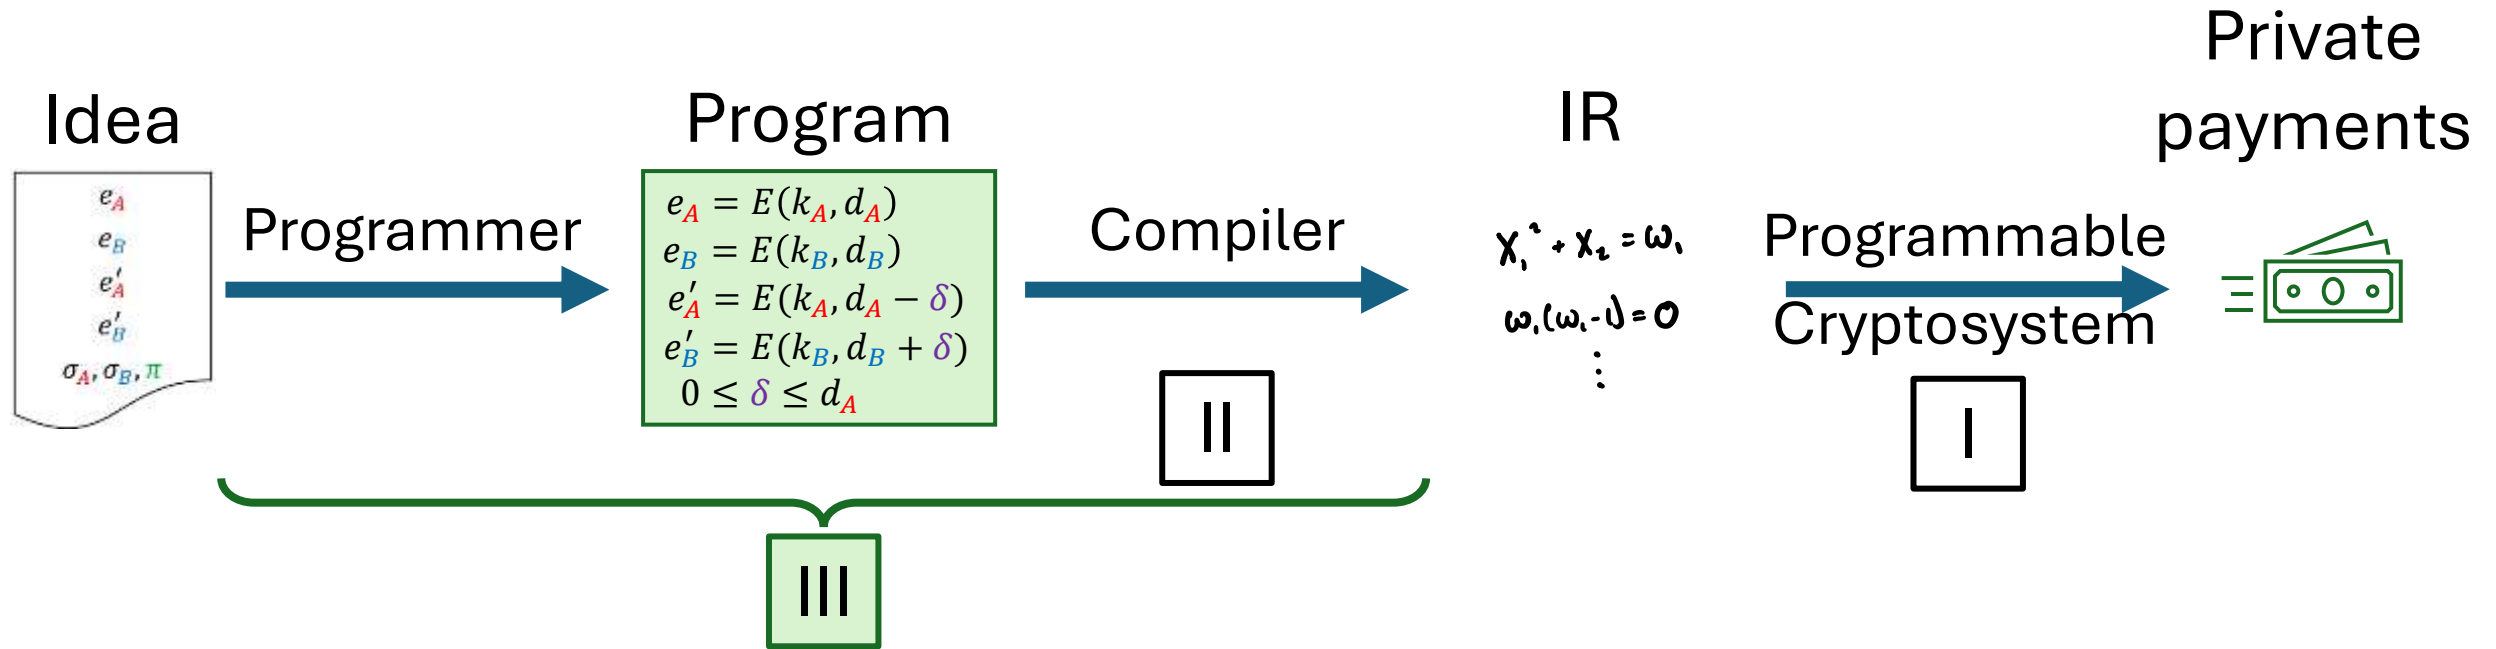
Part I: Collaborative zero knowledge

Part II: Common compiler infrastructure

[Ozdemir, Brown, Wahby; IEEE S&P '22]

Part III: SMT for finite fields

# The full stack of programmable cryptography

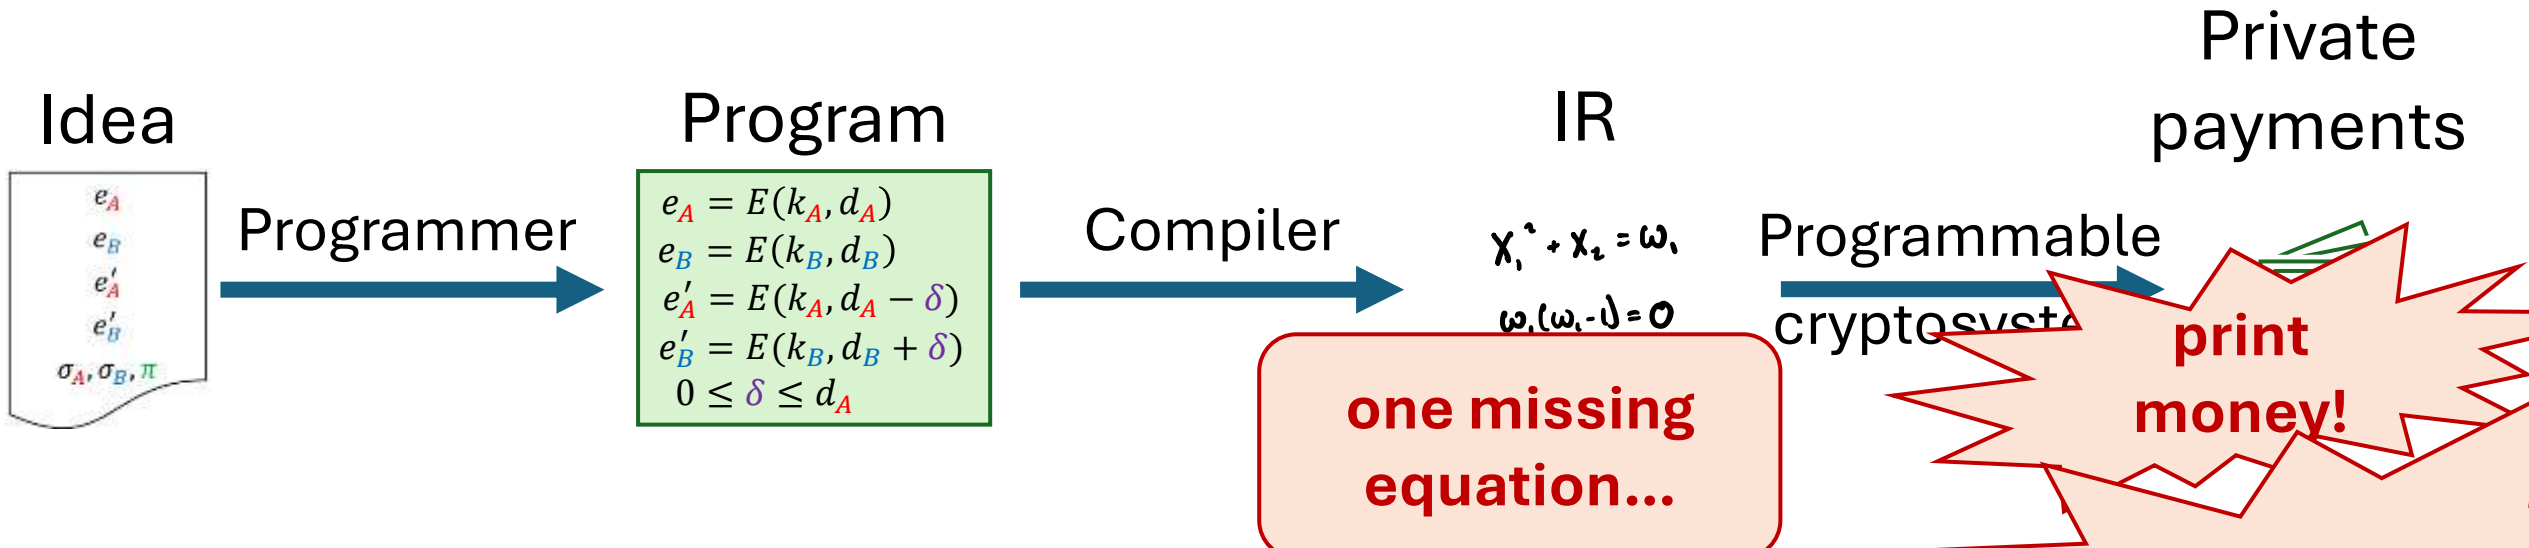


Part I: Collaborative zero knowledge

Part II: Common compiler infrastructure

Part III: SMT for finite fields

# Bugs have serious consequences



Correctness is a big deal:

real vulnerabilities



ZKP audit companies



# Definition: ZKP compiler soundness

High-level predicate

$$P(w \in \mathcal{W})$$



$\mathbb{Z}_p$ -equations

$$\phi(\vec{v} \in \mathbb{Z}_p^n)$$

Encoding

$$e: \mathcal{W} \rightarrow \mathbb{Z}_p^n$$

For simplicity,  
hardcode  $x$  into  $P$

For simplicity,  
assume  $e$  is a  
*bijection*

Definition of **soundness**:  
every  $\phi$ -solution encodes a  $P$ -solution.

$$\forall \vec{v} \in \mathbb{Z}_p^n, \quad \phi(\vec{v}) \Rightarrow P(e^{-1}(\vec{v}))$$

# Standard approach: reduction to satisfiability

$\phi$  is **sound** if:

$$\forall \vec{v} \in \mathbb{Z}_p^n, \quad \phi(\vec{v}) \Rightarrow P(e^{-1}(\vec{v}))$$

$\phi$  is **unsound** if

$$\neg \left( \phi(\vec{v}) \Rightarrow P(e^{-1}(\vec{v})) \right)$$

is **satisfiable**.

# We need the right kind of solver

$\phi$  is *unsound* if  $U$  is satisfiable:

$$\neg \left( \underbrace{\phi(\vec{v})}_{\mathbb{Z}_p} \Rightarrow \underbrace{P}_{\text{Booleans}} \left( \underbrace{e^{-1}(\vec{v})}_{\dots} \right) \right)$$

$\mathbb{Z}_p$

Booleans

...

$\Rightarrow$  SMT solver

Problem:

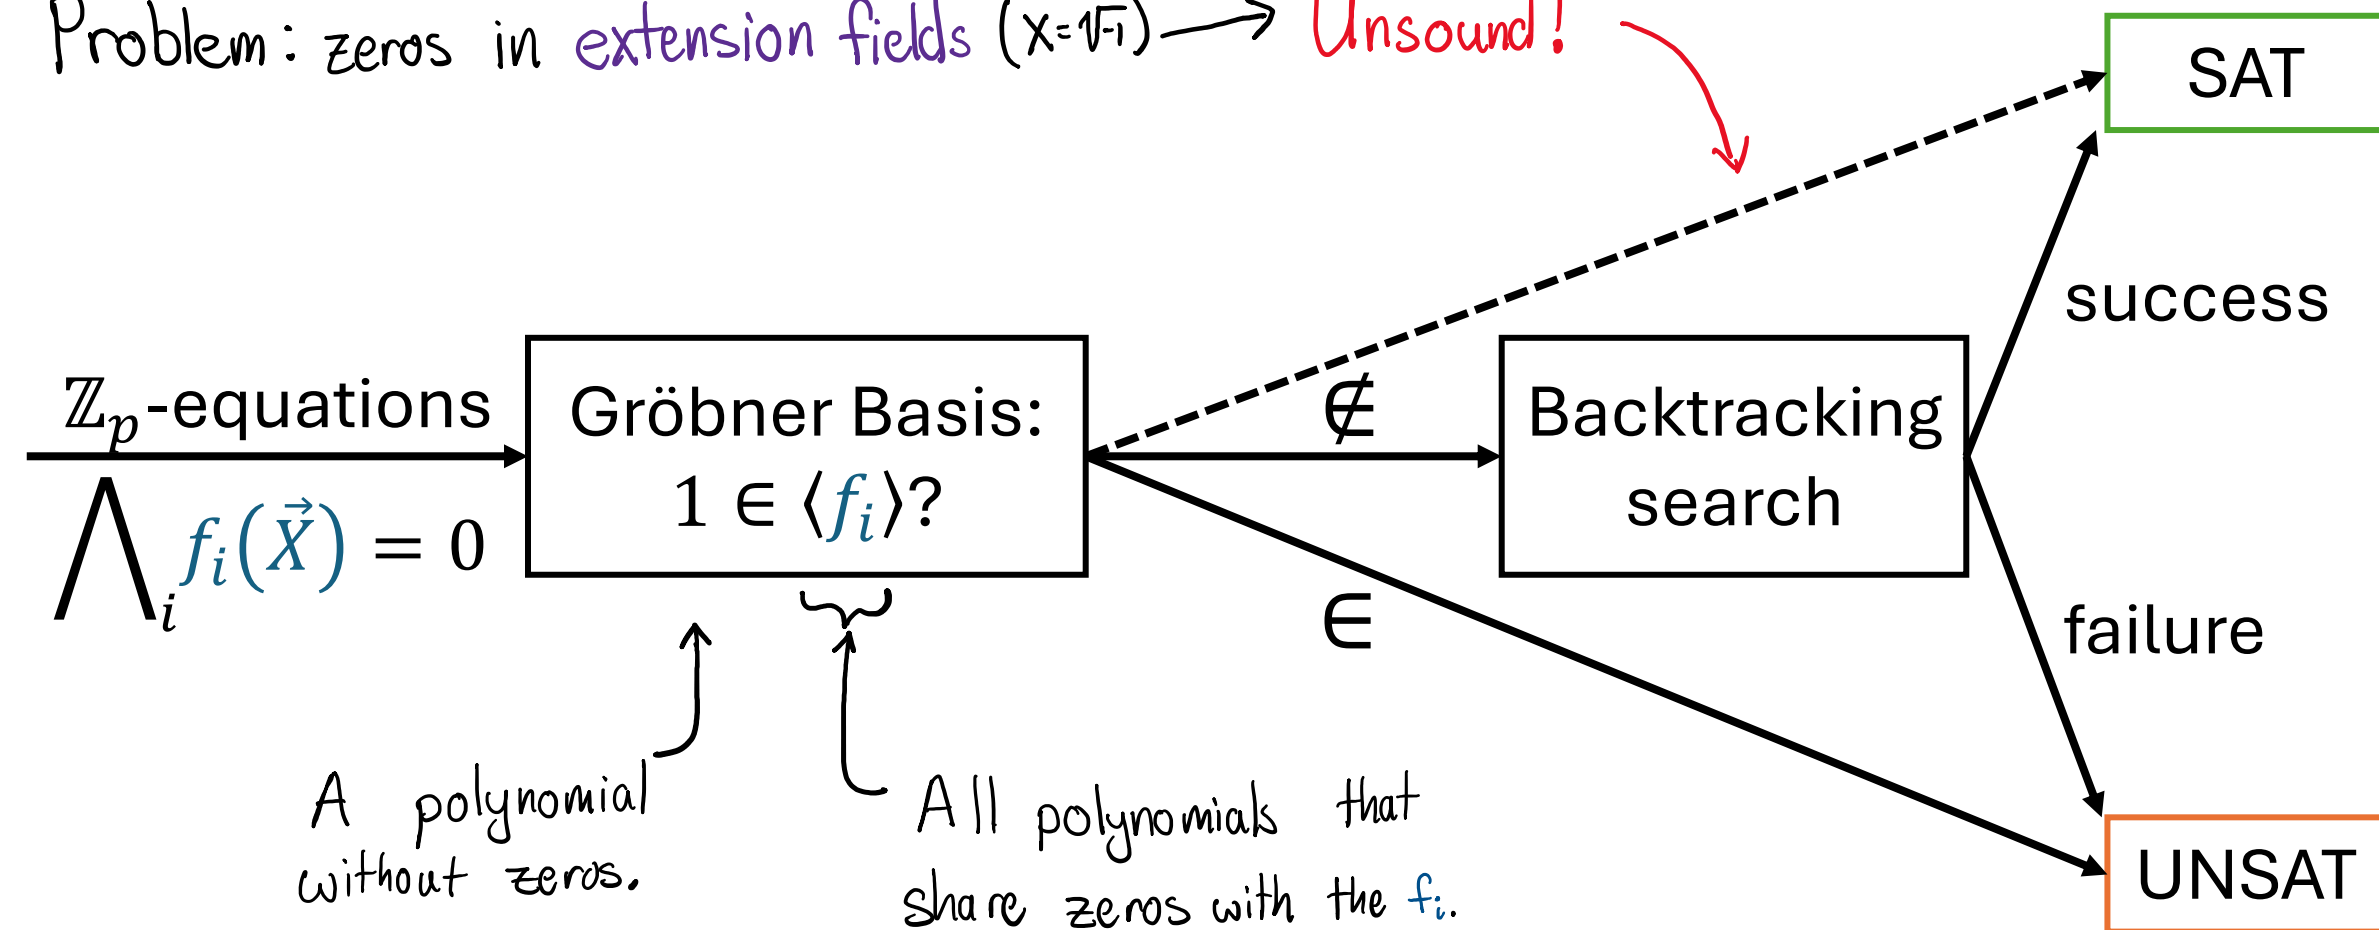
No SMT solvers natively support  $\mathbb{Z}_p$

Our contributions:


- An SMT theory solver for  $\mathbb{Z}_p$
- An implementation in **cvc5**

# $\mathbb{Z}_p$ -solving, Part I

Problem: zeros in extension fields ( $x=\sqrt{-1}$ )  $\rightarrow$  **Unsound!**



# $\mathbb{Z}_p$ -solving, Part II: backtracking search



Intuition:  $\rightarrow$   
 $\# \text{ solutions} \sim p^{\dim}$   
 (false over  $\overline{\mathbb{Z}_p}$  ;)

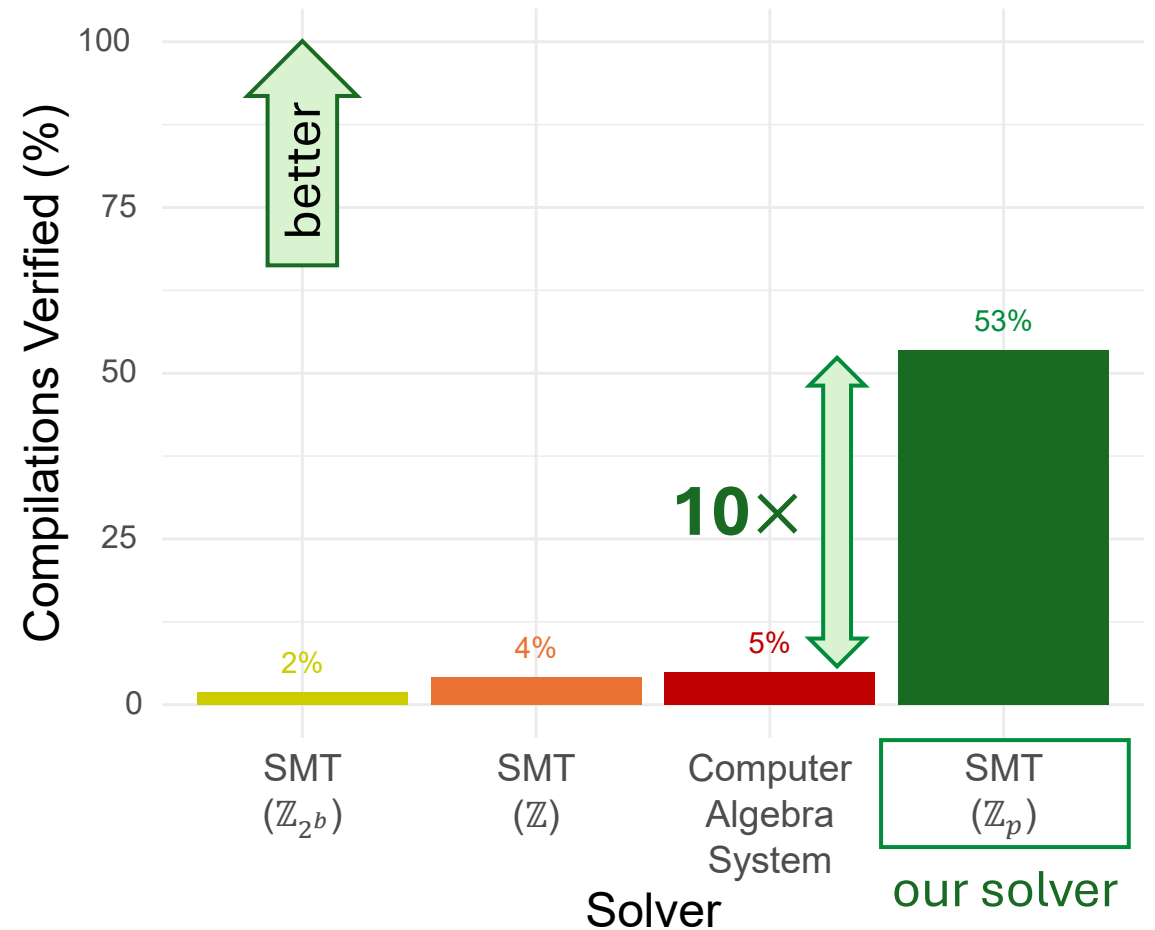
... but  $B$  is highly structured?

At each node:

- Compute a GB  $B$
- Branch in one of three ways:
  - If  $\exists$  univariate  $p(X_i) \in B$ :
    - factor  $p$  and branch on  $X_i$  roots
  - If  $\dim > 0$ : random search:
    - $X_1 = 1,$
    - $X_2 = 1,$
    - $X_1 = 2,$
    - $X_2 = 2, \dots$
  - Else, if  $\dim = 0$ :
    - Use  $B$  to compute a univariate  $p(X_i) \in \langle B \rangle$

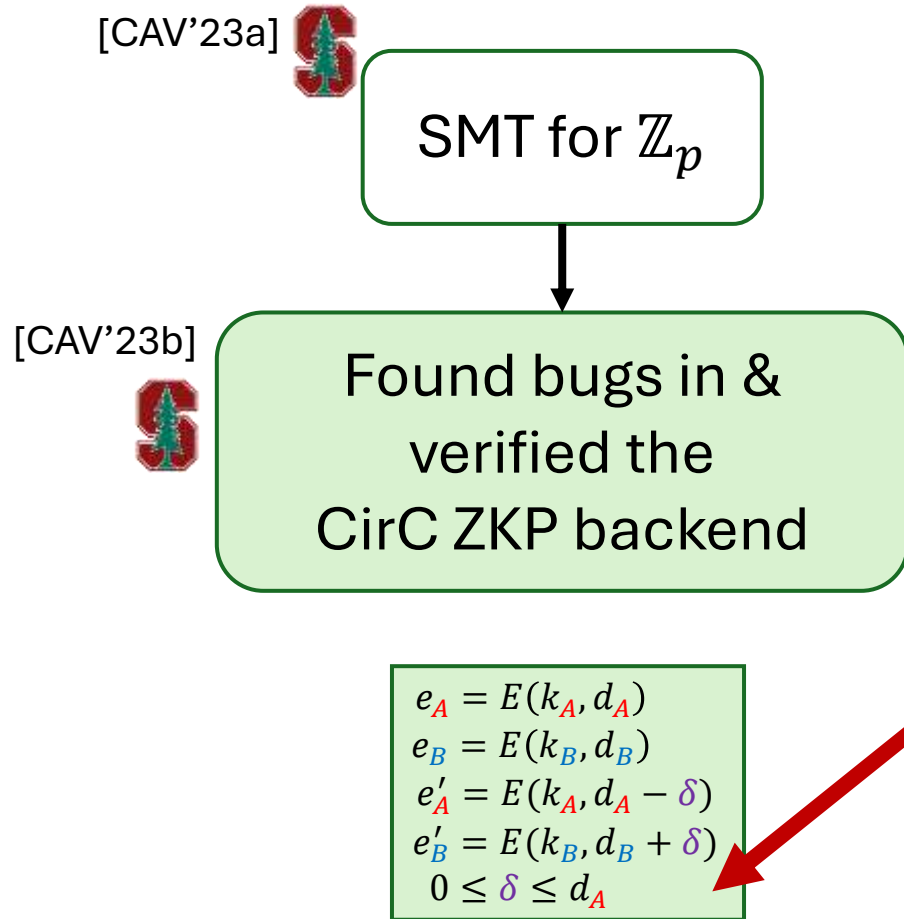
# Experiment: translation validation

Verifying compilations of random\* Boolean programs.



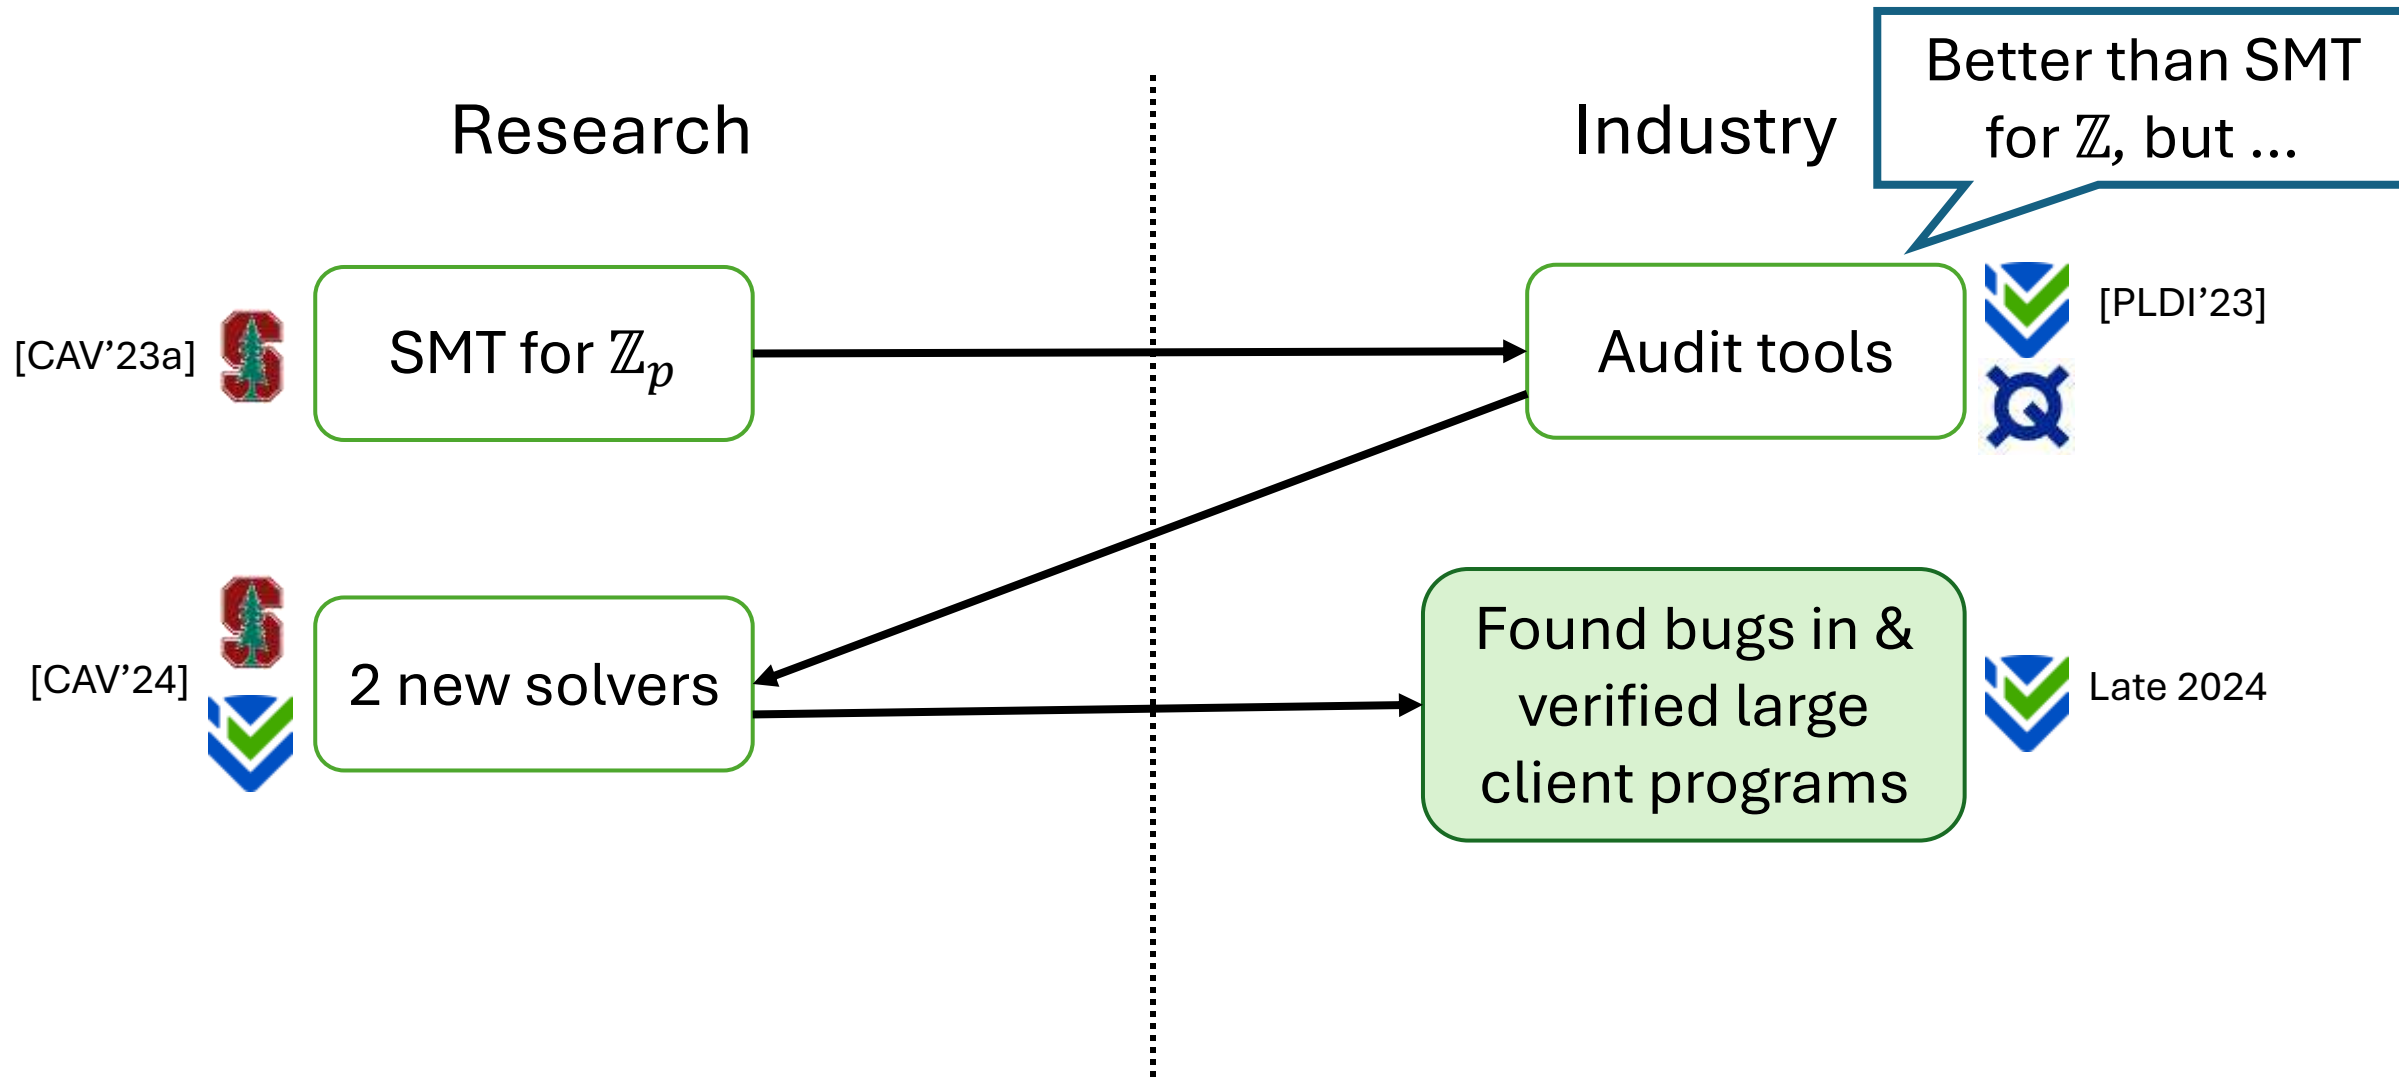
\* With 2-12 input variables and 16-128 intermediate operators.

# Verifying a compiler backend

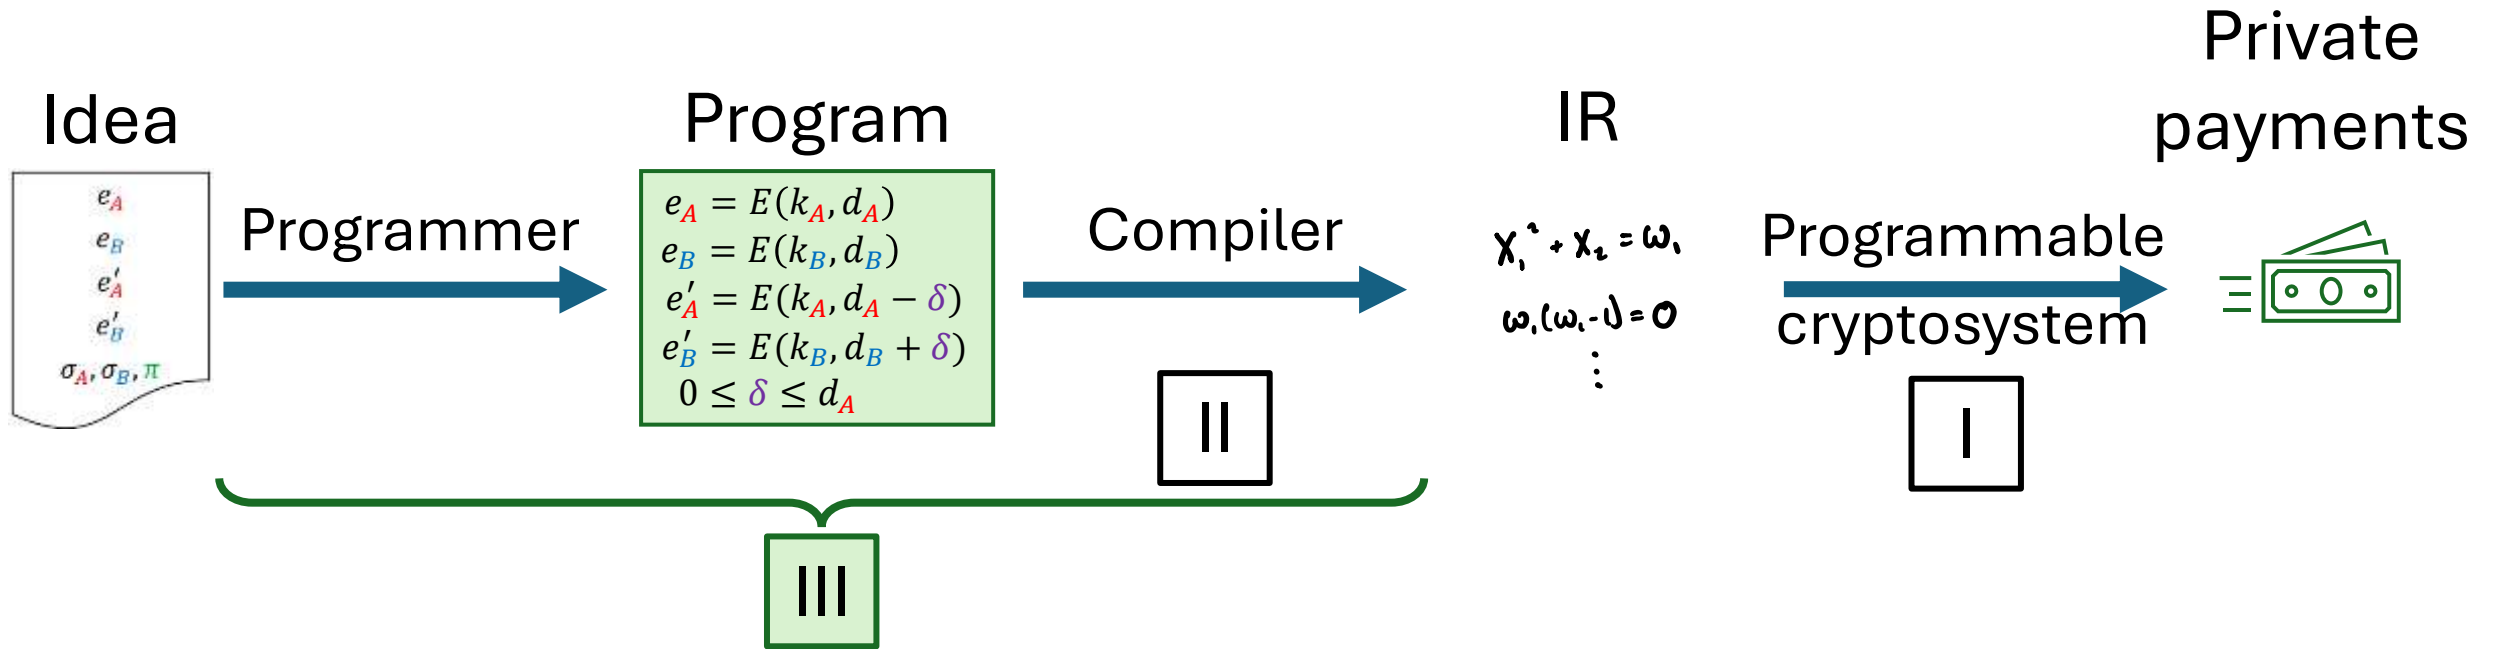


- 1<sup>st</sup> ZKP compiler correctness definition
  - **compilers** + **cryptology**
- 1<sup>st</sup> verified ZKP compiler pass
  - **compilers** + **verification**
  - **Verification failed!?**
    - A miscompilation in our example
- After patches, **verification passed**

# Industrial use



# The full stack of programmable cryptography



Part I: Collaborative ZKPs

Part II: Common compiler infrastructure

Part III: SMT for finite fields

[Ozdemir, Kremer, Tinelli, Barrett; CAV '23]

# Part IV: What's next?

# Now is the right time

Encryption is ubiquitous

- **Good news:** people want privacy
- **Next frontier:** data in use



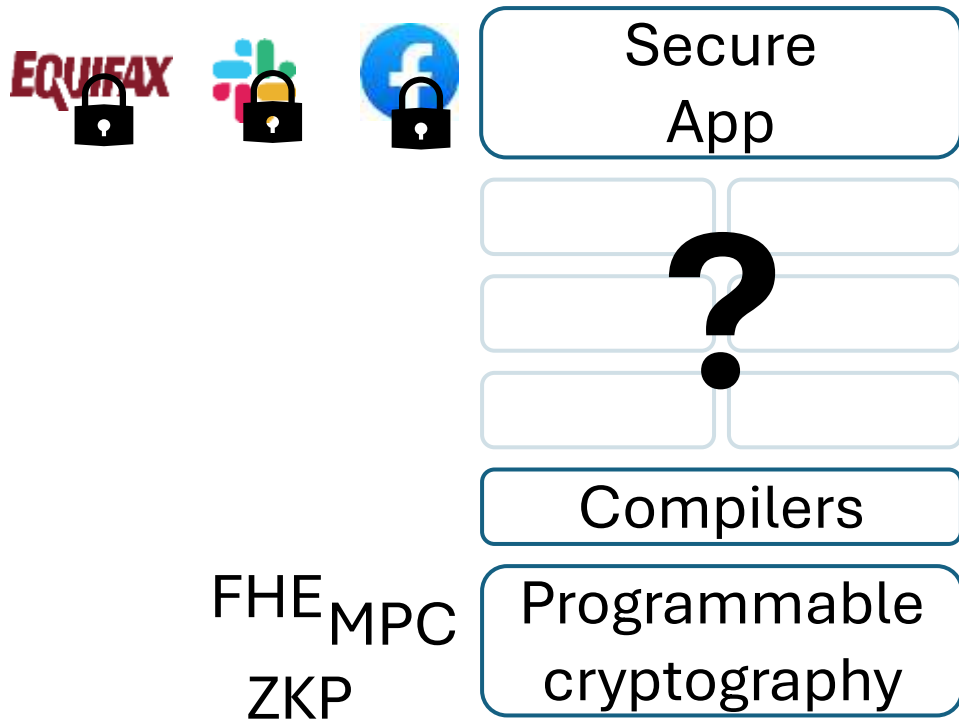
Programmable cryptography is becoming practical

- Deployments: payments (Monero/Zcash), statistics (iOS/Android), ...
- **Good news:** efficient *enough*
- **Bad news:** so far, it requires **application & cryptographic** expertise

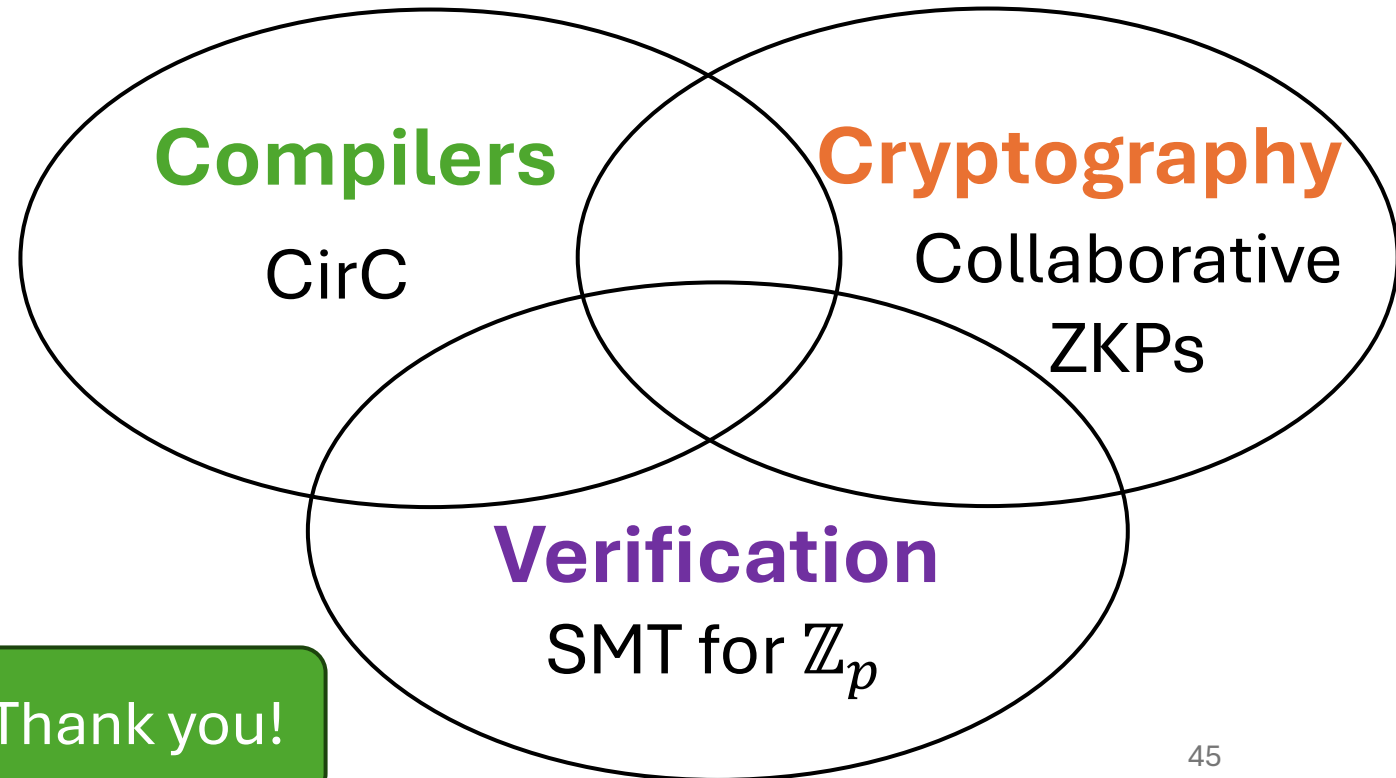
A stack that **separates concerns** will unlock more applications



# My vision: The **full stack** of programmable cryptography



My PhD work:



Thank you!

# Acknowledgements

# Advisors



# Committee



# Mentors




# Applied Cryptography Group



# Centaur: Center for Automated Reasoning



# Officemates



Aina  
Niemetz

Mathias  
Preiner




# Co-instructors and TAs: cs355 and cs343s



# Collaborators

- Abdalrhman Mohamed
- Abhishek Nair
- Ahmed Irfan
- Aina Niemetz
- Aleksandar Zeljic
- Alexander Ek
- Alireza Shirzad
- Alp Bassa
- Andres Nötzli
- Andrew Reynolds
- Anna P. Y. Woo
- Arash Mirzaei
- Arjun Arora
- Arjun Viswanathan
- Barry Whitehat
- Benedikt Bünz
- Brennan Shacklett
- Bruno Dutertre
- Byron Cook
- Cesare Tinelli
- Chad Sharp
- Clark Barrett
- Corina Pasareanu
- Damjan Vukcevic
- Dan Boneh
- Daniel Bork
- Deepti Raghavan
- Divya Gopinath
- Edward Chen
- Evan Laufer
- Fait Poms
- Fraser Brown
- Gereon Kremer
- Guy Katz
- Haniel Barbosa
- Hanna Lachnitt
- Haoze Wu
- Işıl Dillig
- Jean Sung
- Jess Woods
- Jincheng Wang
- Jinhao Zhu
- Jonathan Cogan
- Kayvon Fatahalian
- Keith Winstein
- Kostas Ferles
- Kyle Julian
- Lef Ioannidis
- Makai Mann
- Martin Brain
- Mathias Preiner
- Michael Sheely
- Mudathir Mohamed
- Olivier Pereira
- Pat Hanrahan
- Paul Grubbs
- Philip B. Stark
- Pratyush Mishra
- Ran Libeskind-Hadas
- Reyna Hulett
- Riad S. Wahby
- Ricson Cheng
- Sadjad Fouladi
- Saranyu Chattopadhyay
- Scott Viteri
- Sebastian Angel
- Shankara Pailoora
- Thomas Hader
- Thomas Pornin
- Vanessa Teague
- Wenting Zheng
- Wilson Nguyen
- Ying Sheng
- Yoni Zohar

# Stanford Community



Ann  
Harara

Ruth  
Harris



# Mira Loma HS, Harvey Mudd, & Loomis Chaffee



# Family



# Veronica





# My work:

## Compilers

First compiler infrastructure for multiple cryptosystems

[S&P'23]

[S&P'22]

[FMCAD'21]

[FMCAD'22]

[SAT'19]

[S&P'25a]

[S&P'25b]

[CAV'23b]

[CCS'24]

[CAV'23a]

[CAV'24]

[USENIX Sec.'20]

[USENIX Sec.'22]

[TACAS'22]

[IJCAR'22]



## Cryptography

Definitions and protocols for new kinds of security

## Verification

First SMT solver for finite fields