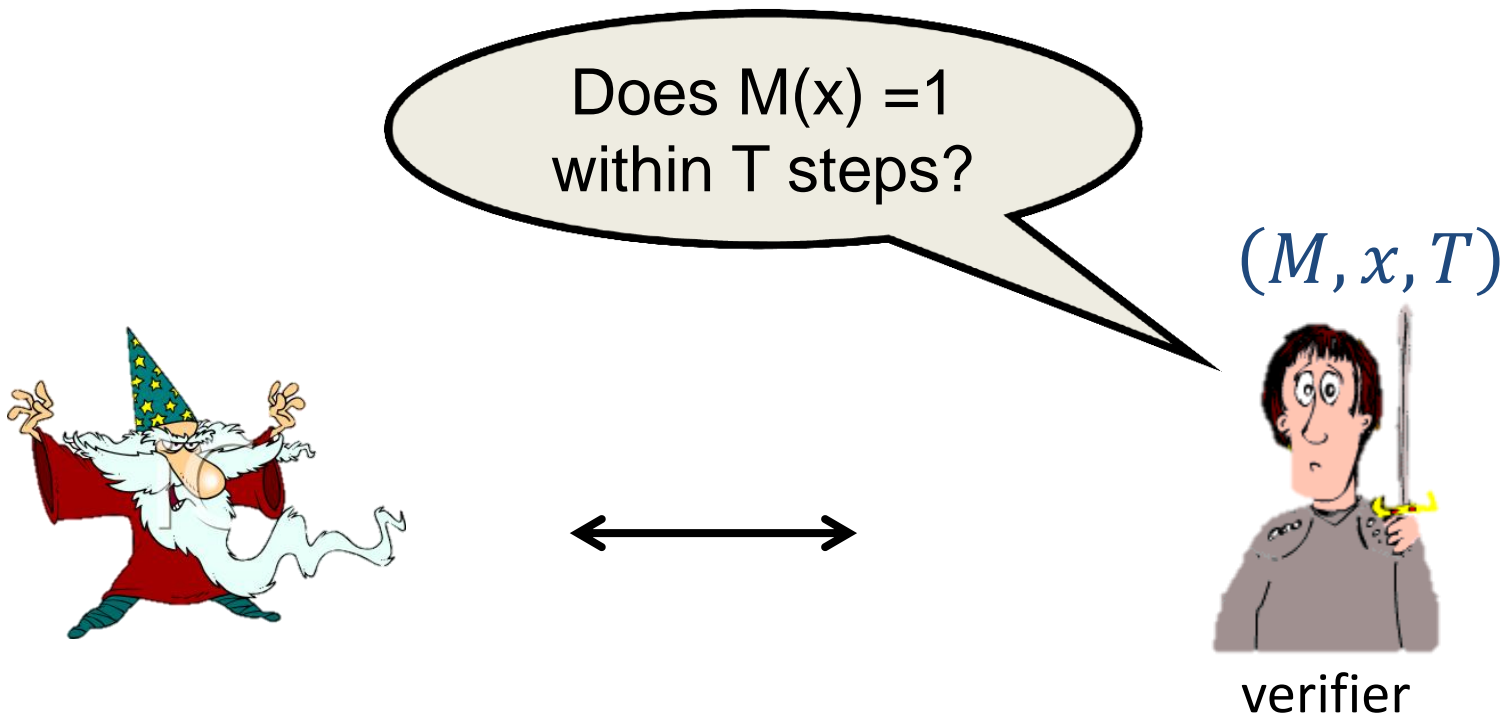


How to Bootstrap a SNARK *in Public*

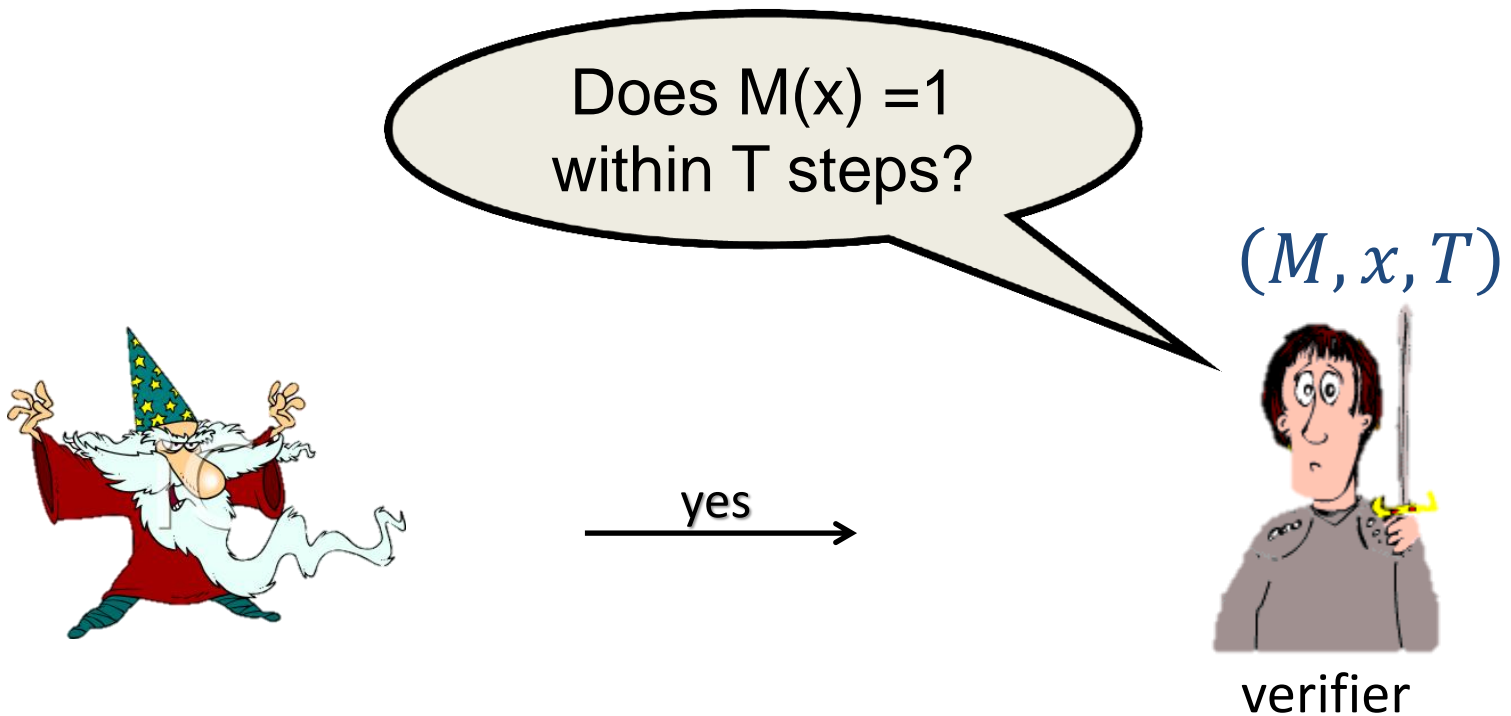
Nir Bitansky, Ran Canetti, Alessandro Chiesa, Eran Tromer

How quickly can we verify the
result of long computations?

How quickly can we verify the result of long computations? (Plain version)

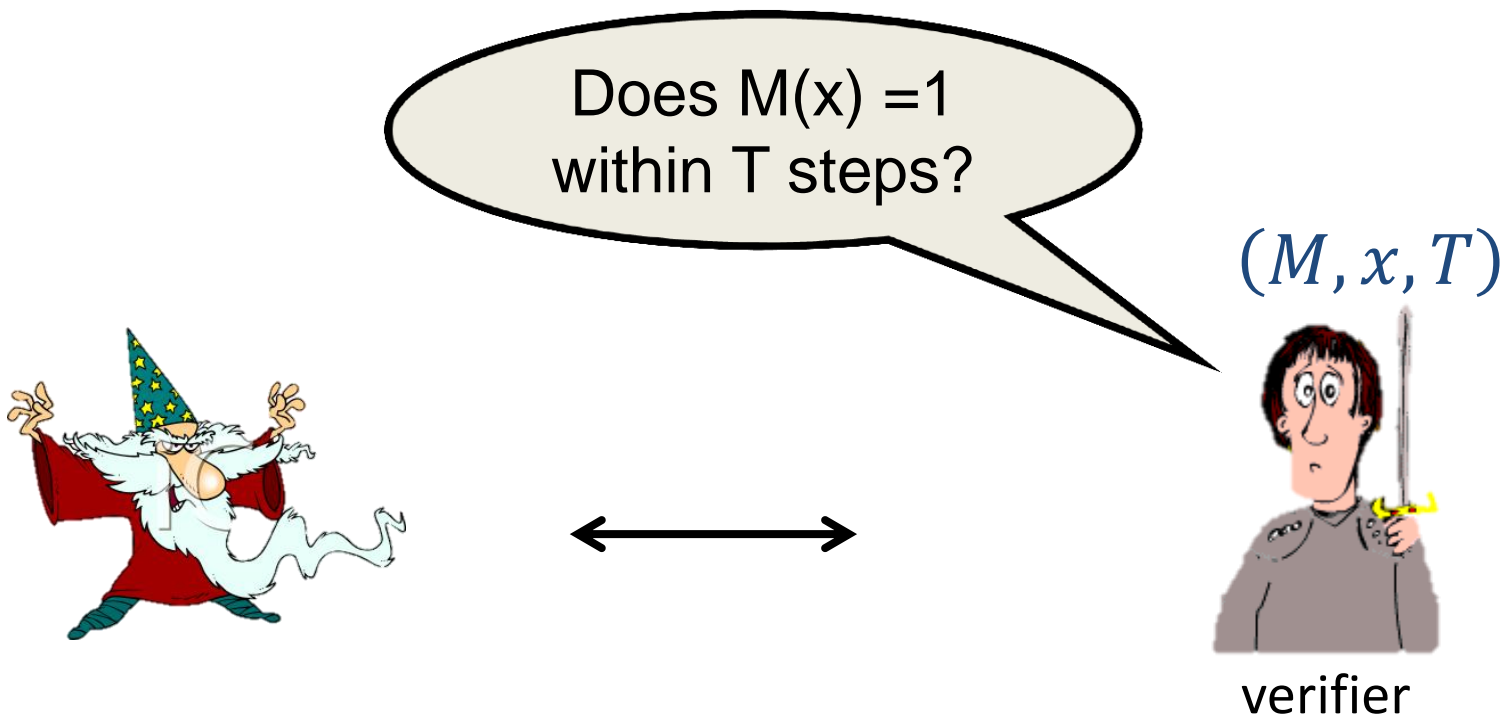


How quickly can we verify the result of long computations? (Plain version)



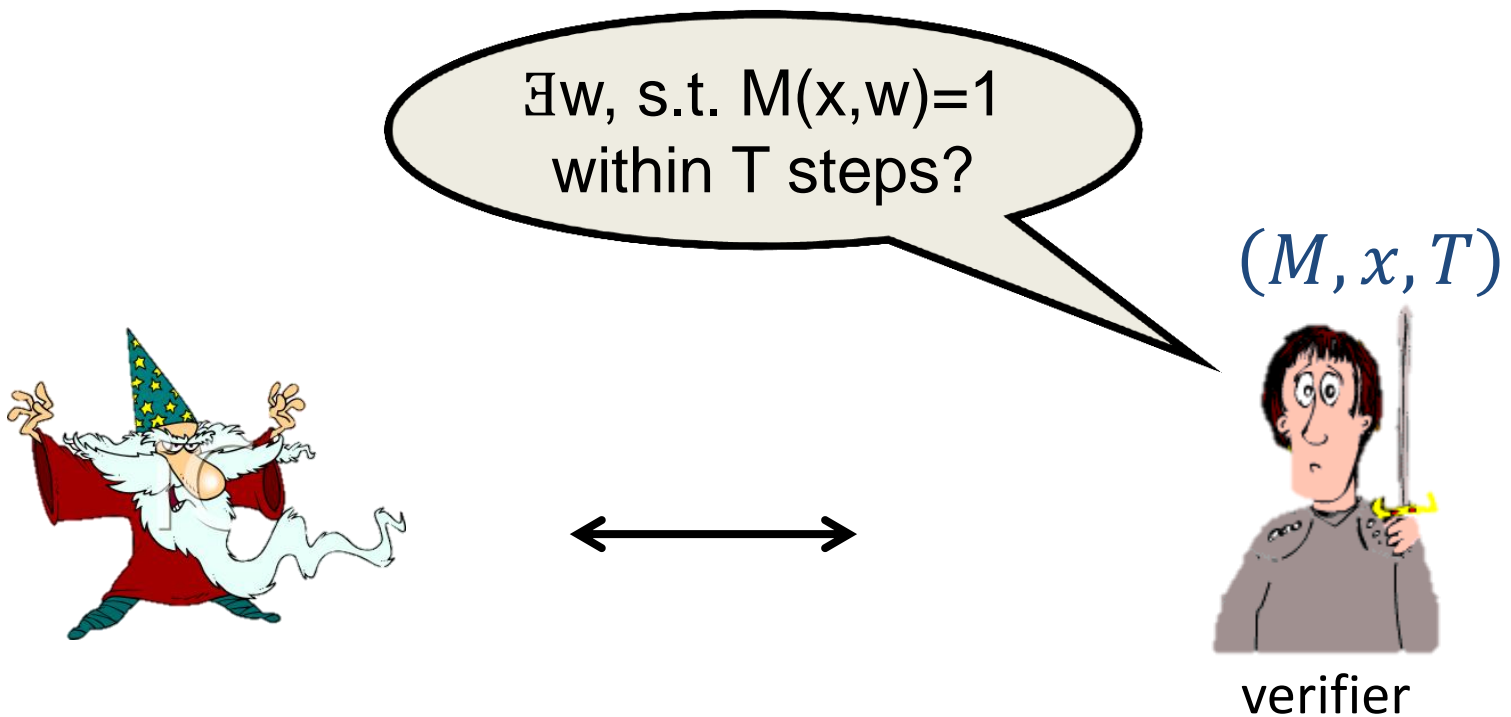
Verify by running $M(x)$ for T steps.

How quickly can we verify the result of long computations? (Plain version)

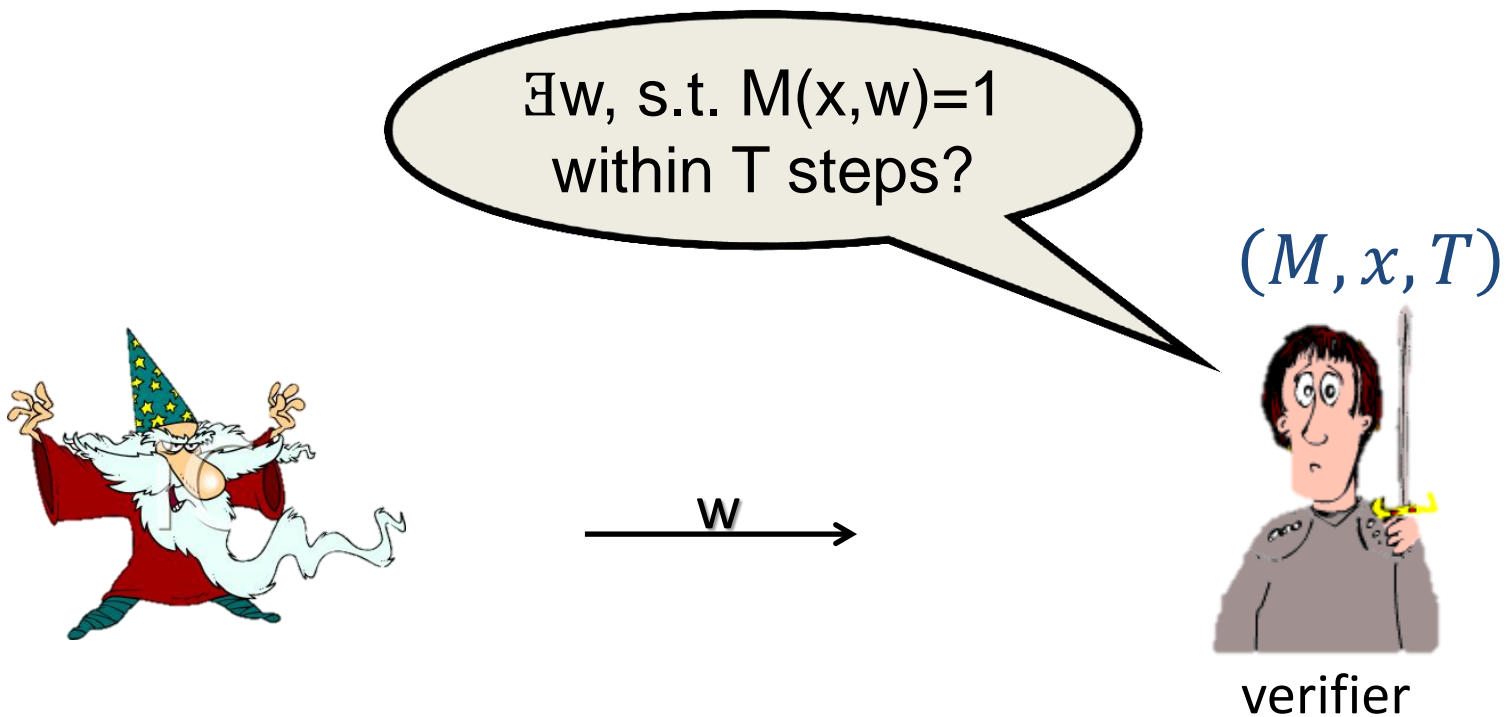


Can we do better?

How quickly can we verify the
result of long computations?
(with prover input – “NP version”)




How quickly can we verify the
result of long computations?
(with prover input – “NP version”)



Verify by running $M(x, w)$ for T steps.

How quickly can we verify the
result of long computations?
(with prover input – “NP version”)



Can we do better?

Succinct Proofs with incomplete input (" for NP ")

possibly long

$(M, x, T), w$



$\exists w$ s.t. $M(x, w) = 1$
in $\leq T$ steps?



(M, x, T)



$\text{poly}(|x|, k)$

universal poly, e.g.

$|x| \cdot k$


independent of T !

security

parameter

Succinct Proofs with incomplete input (“ for NP ”)


- Statistical soundness is unlikely [BHZ87, GH98, GVW02]. Thus we settle for computational soundness.
- However, we require extractability:
 - Natural in real-life applications (databases...)
 - Crucial for this work



How many rounds do succinct arguments require?


How many rounds do succinct arguments require?

[Kilian 92]: can do 4-message
(assuming CRH)




How many rounds do succinct arguments require?

[Kilian 92]: can do 4-message
(assuming CRH)



[Micali 94]: one message!
with a random oracle
(aka “CS proofs”)



Non-interactive in the plain model?




Non-interactive in the plain model?



Totally non-interactive protocols
(against non-uniform provers
for “hard enough languages”)
Are unlikely [BHZ87, GH98, GVW02].

With a verifier initial message (reference string)?



reference string σ
sent before statements

Succinct Non-Interactive Argument of Knowledge (SNARK):

A protocol (P,V) such that:

- V sends an initial message σ to P
- Repeat:
 - P sends $(M,x,T), \pi$ to V
 - $V(M,x,T, \pi, \sigma)=\text{acc/rej}$

Succinct Non-Interactive Argument of Knowledge (SNARK):

A protocol (P,V) such that:

- V sends an initial message σ to P
- Repeat:
 - P sends $(M,x,T), \pi$ to V
 - $V(M,x,T, \pi, \sigma)=\text{acc/rej}$

Completeness: If $\exists w$ s.t. $M(x,w)=1$ within T steps, then V accepts.

Extractability: \forall pt P' \exists pt E , such that when (P',V) accepts (M,x,t, π) , E outputs w s.t. $M(x,w)=1$ within T steps (except w.p. $\text{negl}(k)$).

Designated verifier SNARKs

Same as (publicly verifier) SNARKs except:

- V keeps secret state τ associated with σ .
- V uses τ in each verification.

Disadvantages:

- Vulnerable to leakage on verifier (even the verifier's decision)
- Proofs are no longer transferrable or publicly verifiable ("publishable").
- Harder to compose (later on)

Can we construct SNARKs?

No SNARK can be proven secure via “black-box reduction to an efficiently falsifiable assumption” [Gentry-Wichs11].

- even for designated verifier SNARKs
- even if we only require plain soundness (without knowledge extraction)

Can we construct SNARKs?

No SNARK can be proven secure via “black-box reduction to an efficiently falsifiable assumption” [Gentry-Wichs11].

- even for designated verifier SNARKs
- even if we only require plain soundness (without knowledge extraction)

What can we do?

- Option 1: Use non BB reductions
- Option 2: Use other assumptions

SNARKs from “non-falsifiable assumptions”

- Replace the RO in [Micali94] with a “sufficiently complicated” hash function and assume security.

Disadvantages: Implementation specific, doesn't teach us much...

- Based on “extractable collision resistant hash functions” [Bitansky Canetti Chiesa Tromer 11 , Goldwasser Lin Rubinfeld 11, Damgard Faust Hazay 11]

Disadvantage: Only designated verifier.

PV SNARKs with long reference string (“with pre-processing”)

In the initial stage, V “works hard”:

generates (σ, τ) where:

- τ is $\text{poly}(k)$
- σ is $\text{poly}(T, k)$

In proof stage, V is still succinct - only uses τ .

PV SNARKs with long reference string (“with pre-processing”)

In the initial stage, V “works hard”:

generates (σ, τ) where:

- τ is $\text{poly}(k)$
- σ is $\text{poly}(T, k)$

In proof stage, V is still succinct - only uses τ .

Note: τ is public!

Can realize based on a Knowledge-of-exponent assumption in bilinear groups

[Groth10, Lipmaa12, Gennaro-Gentry-Parno-Raykova12]

Another advantage of [G10,L12,GGPR12]
(Following [Ishai-Kushilevitz-Ostrovsky07])

Very different techniques – alternative to PCPs

Potentially better efficiency (for prover).

Prover efficiency is important !
(e.g. cloud computing)

Another advantage of [G10,L12,GGPR12]
(Following [Ishai-Kushilevitz-Ostrovsky07])

Very different techniques – alternative to PCPs

Potentially better efficiency (for prover).

But...

For computations with time T , space S

Prover needs $T \text{ poly}(k)$ **space!**

Would like to preserve time and space individually.

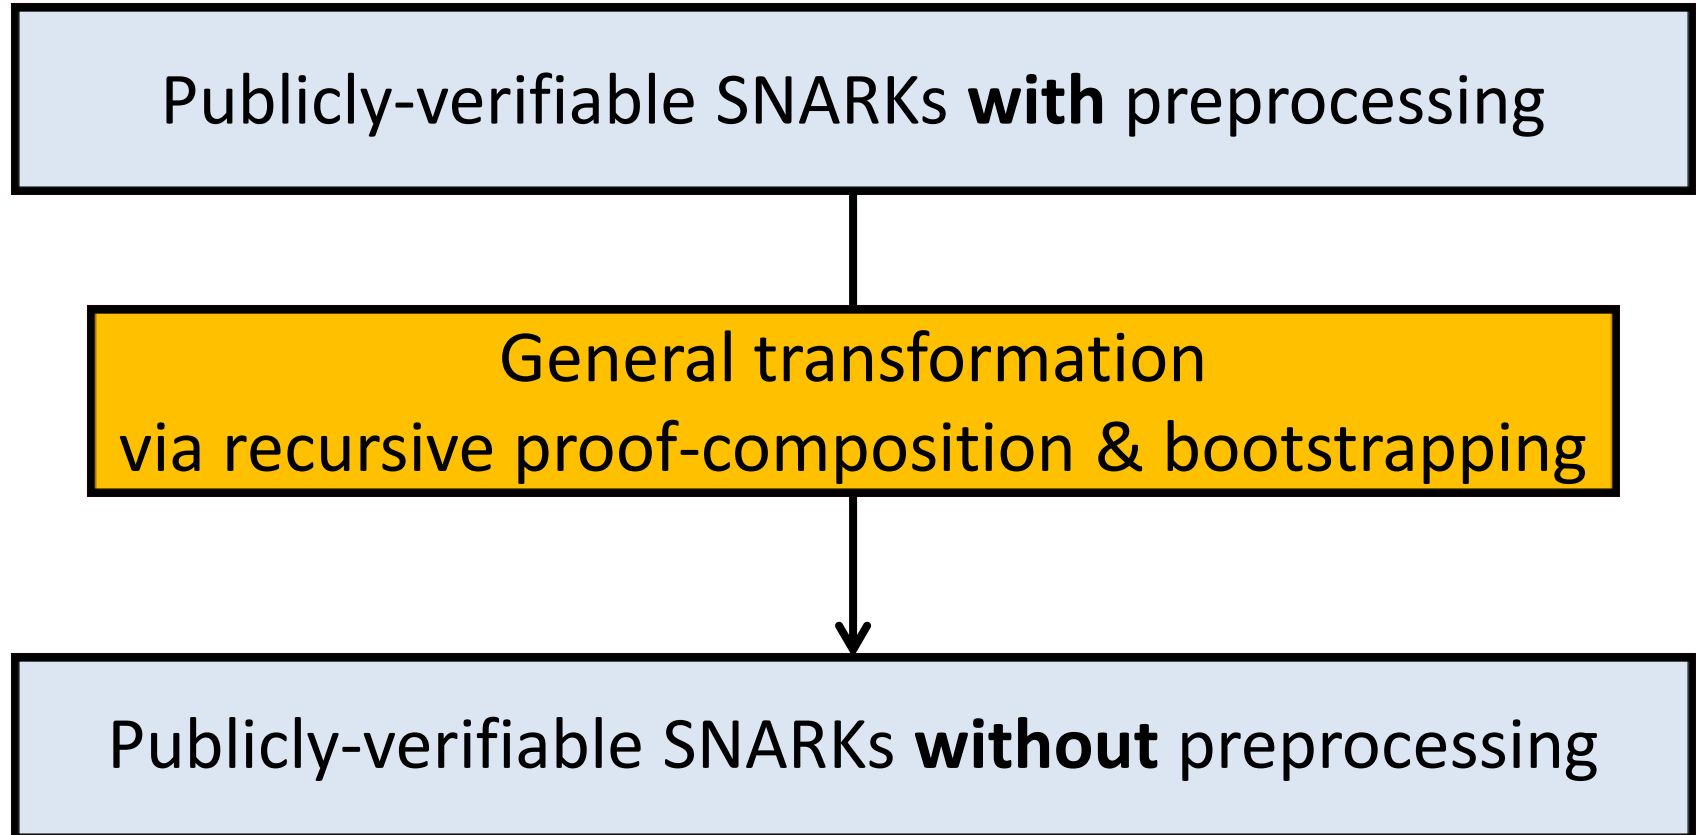
First Main Result

Publicly-verifiable SNARKs **with** preprocessing

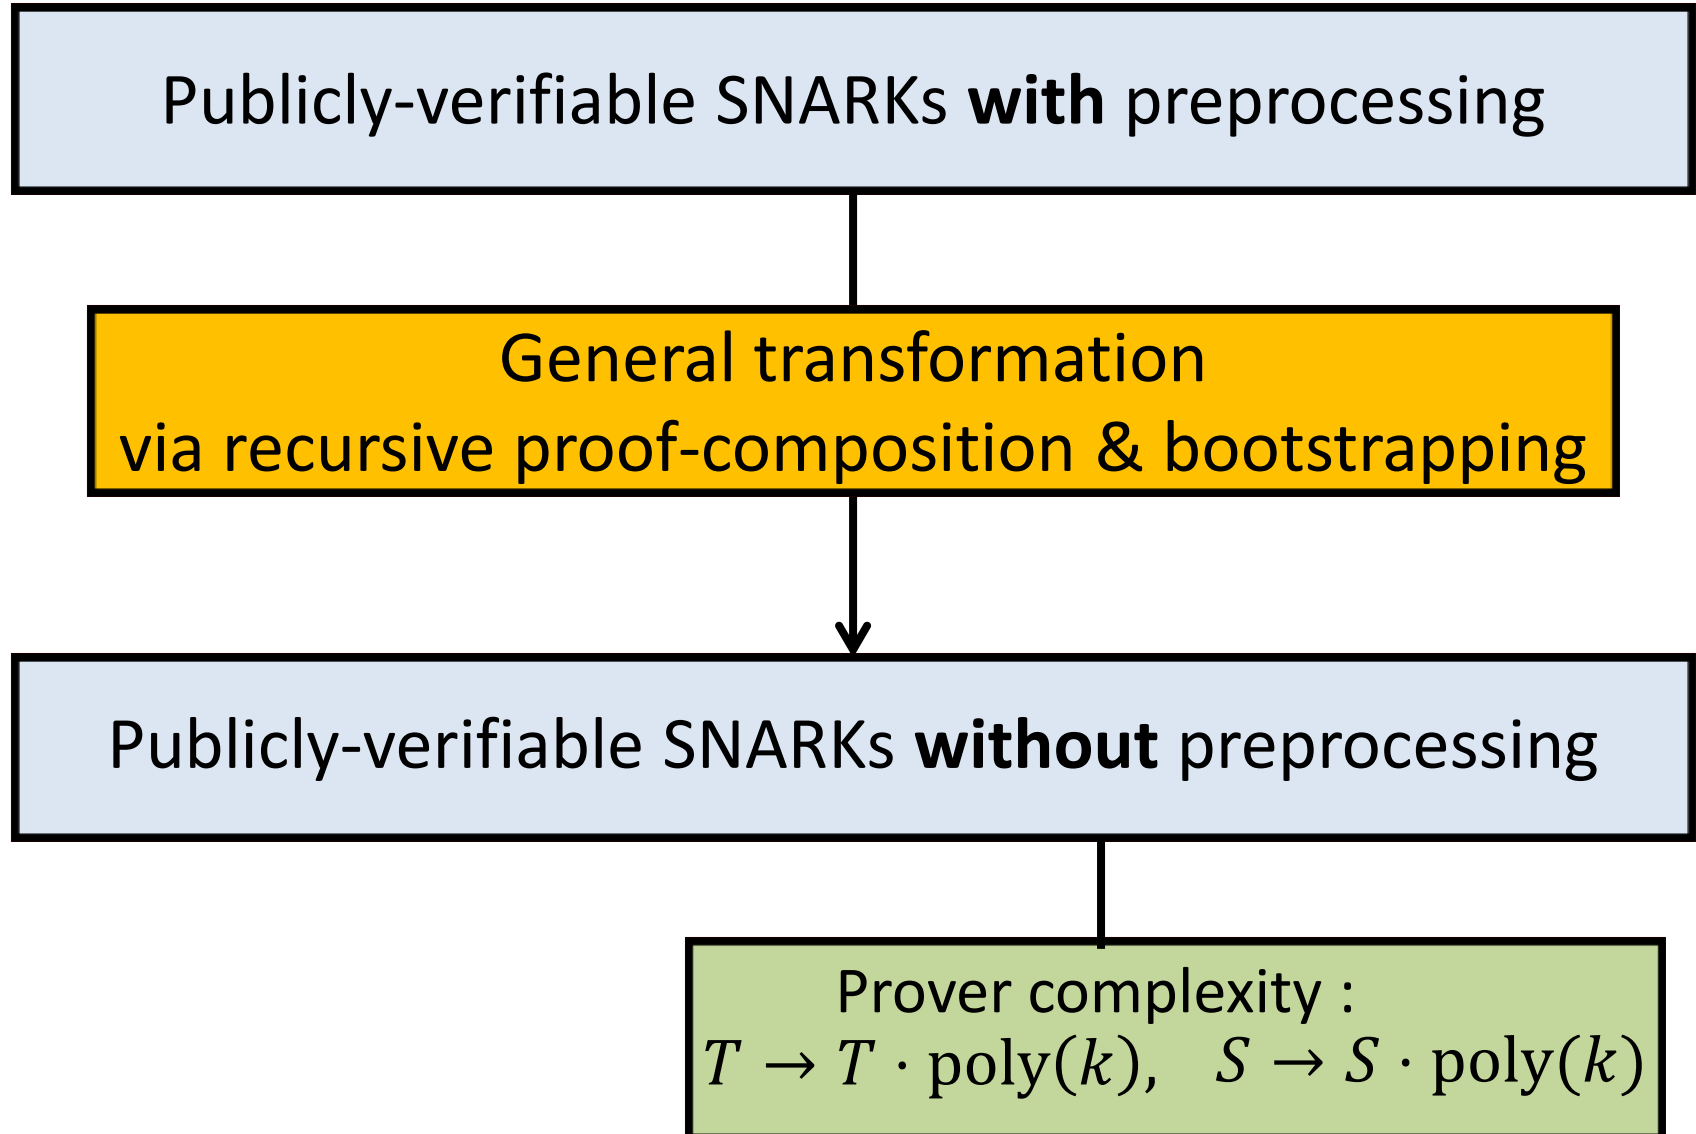


Publicly-verifiable SNARKs **without** preprocessing

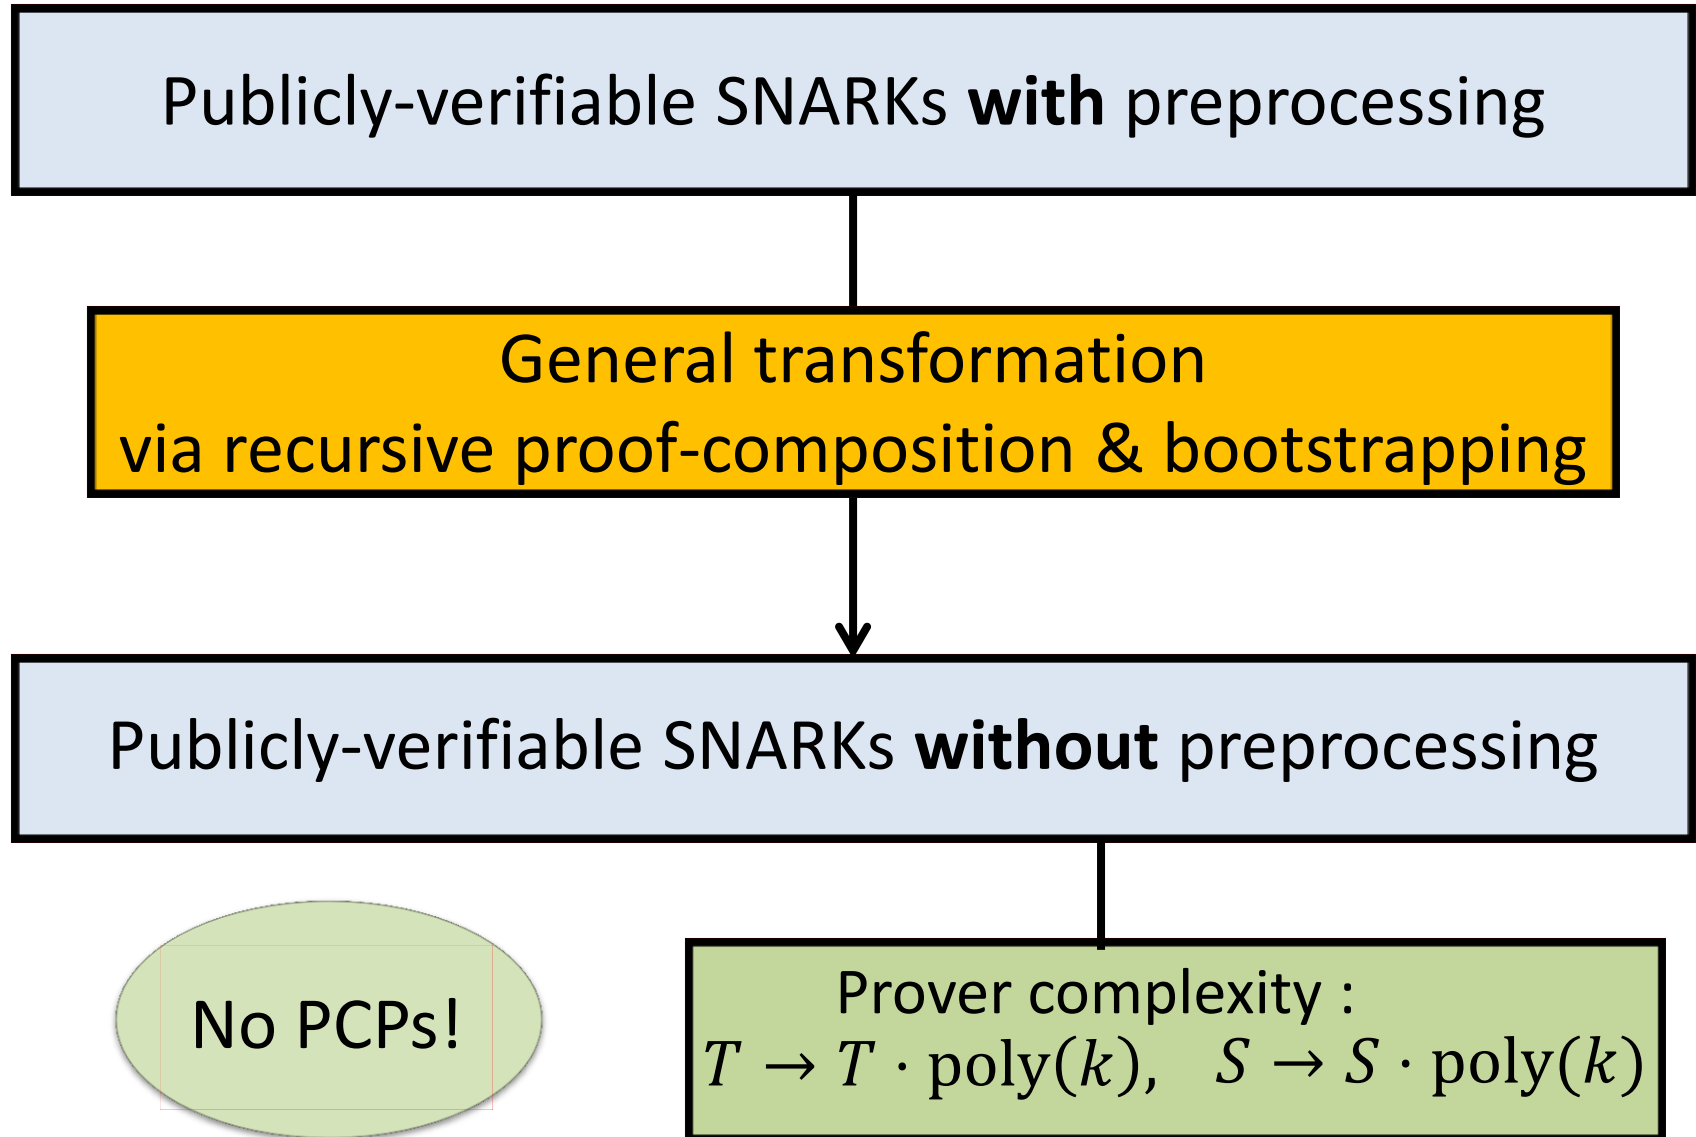
First Main Result



First Main Result



First Main Result



Corollaries

Assuming KEA in a bilinear group, there exist *fully succinct* publicly-verifiable SNARKs .

Any SNARK can be transformed into a SNARK where:

- Prover time is $T \cdot \text{poly}(k)$
- Prover space is $S \cdot \text{poly}(k)$

(T,S are time and space of original M)

The Core Idea: Bootstrapping a SNARK

Only need to be able to prove correctness of
(many) small computations.

The Core Idea: Bootstrapping a SNARK

Only need to be able to prove correctness of (many) small computations.

How small?

as small as SNARK verification (and a bit more)

- The preprocessing becomes cheap ($\text{poly}(k)$)
- Prover overhead becomes $\text{poly}(k)$
(both in time and in space)

Part I

How to Bootstrap a SNARK:
a Bare-Bones Description

Part II

Using the Proof Carrying Data
(PCD) abstraction

Part I:
Bare-Bones Description

Incremental Computation [Valiant08]

a possibly useful idea

Compile a computation $M(x, w)$ to a new one that after each step spits a **short** proof of its correctness **so far**

Incremental Computation [Valiant08]

a possibly useful idea

Compile a computation $M(x, w)$ to a new one that after each step spits a **short** proof of its correctness **so far**

but... (implicitly) assumes fully-succinct SNARKs

Incremental Computation [Valiant08]

a possibly useful idea

Still uses SNARKs in a non-trivial way:
proofs only involve “small” computations:
proportional to the space S used by M .

Can use preprocessing SNARKs, where
preprocessing is as cheap as S

Problems:

In general, S may be as large as T

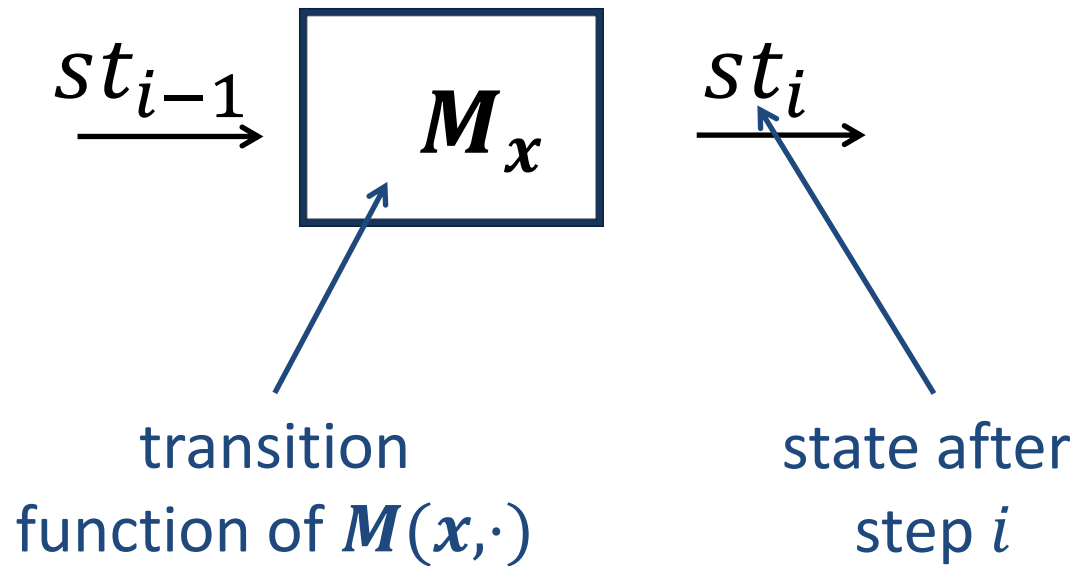
Need to carefully aggregate proofs by composition

Incremental Computation

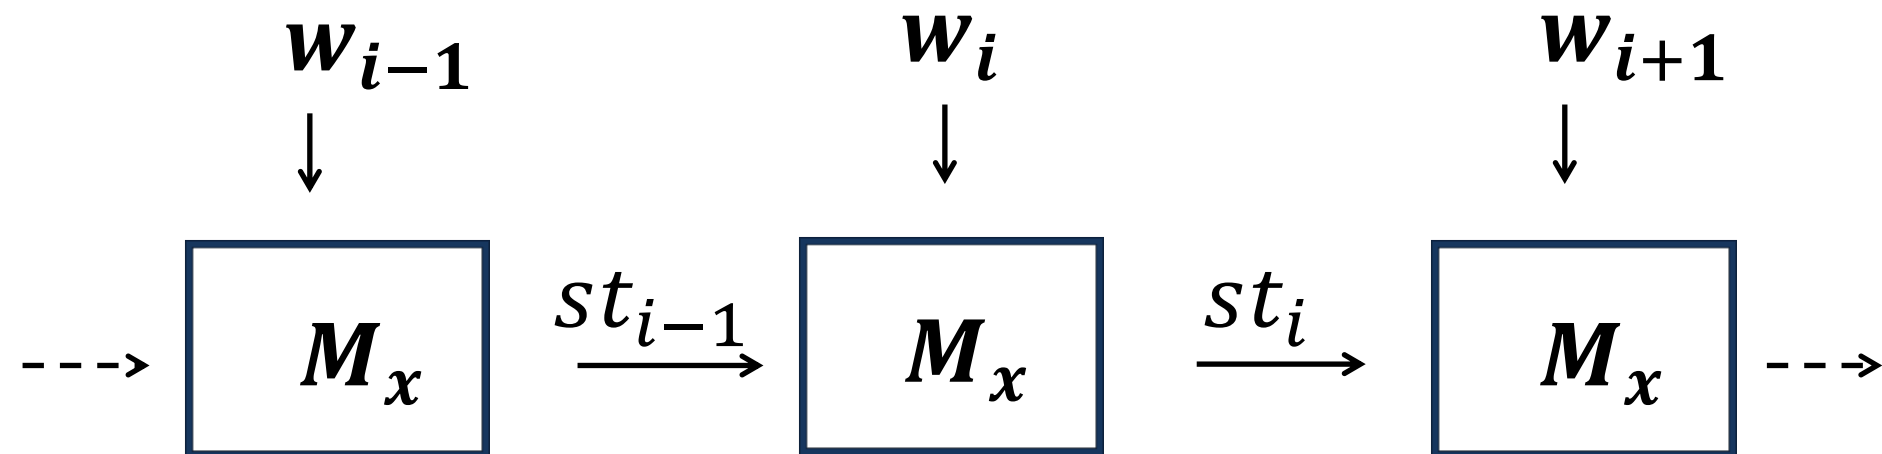
More Concretely


Split a T -step computation $M(x, w)$ to T single-step computations

Potential additional
input bit read at step i $\rightarrow w_i$



Split a T -step computation $M(x, w)$
to T single-step computations

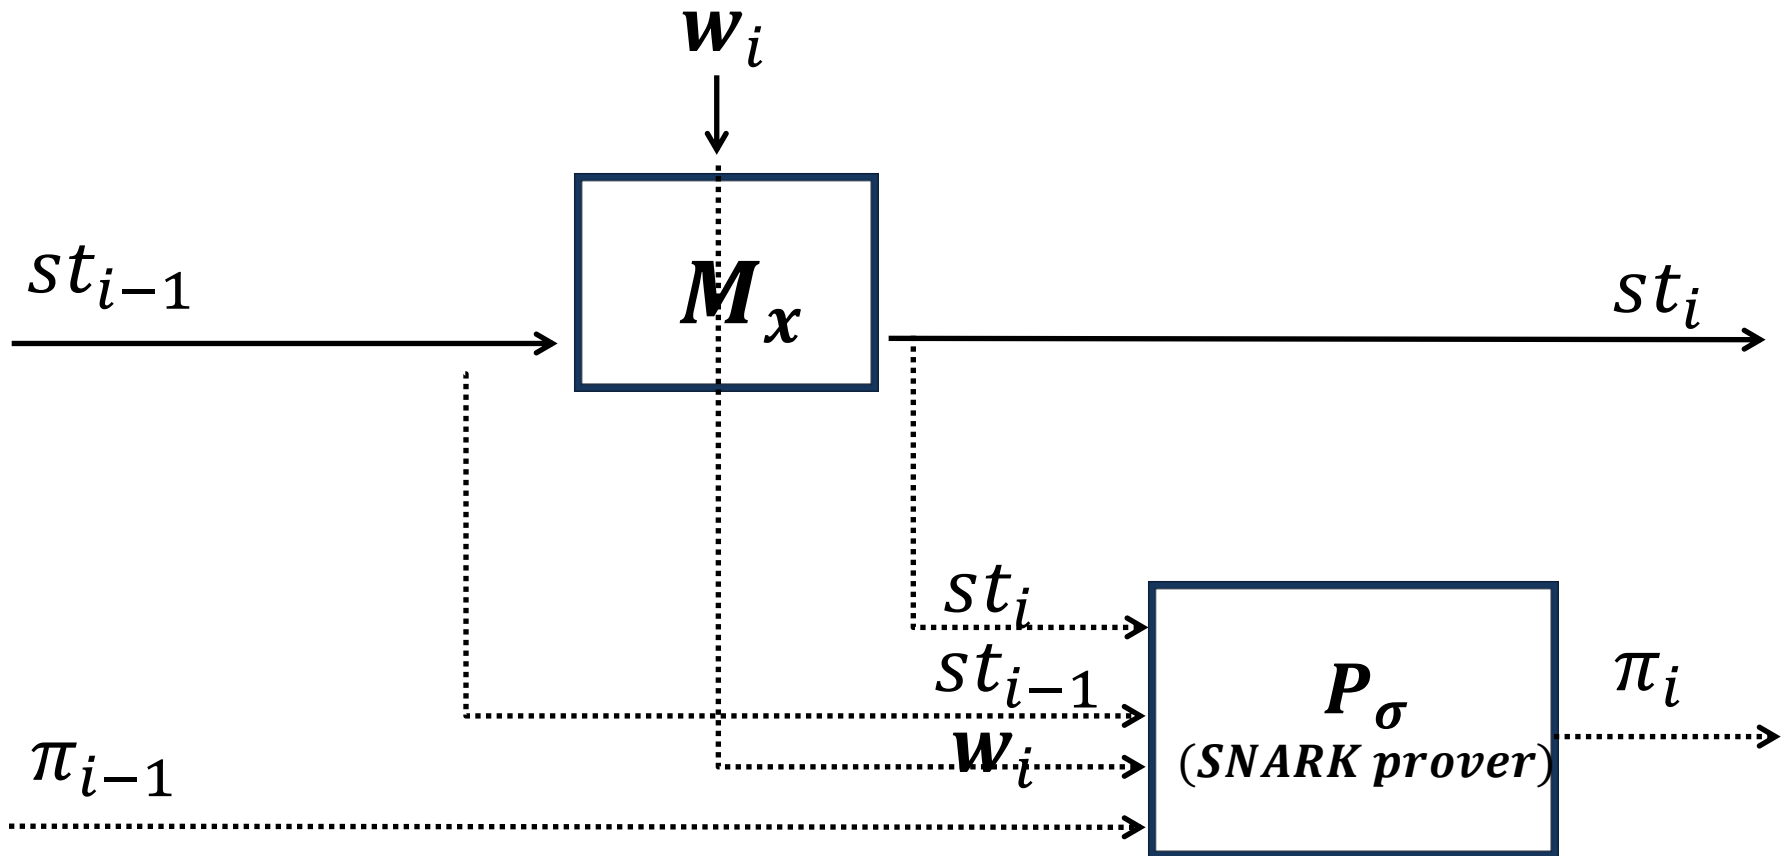




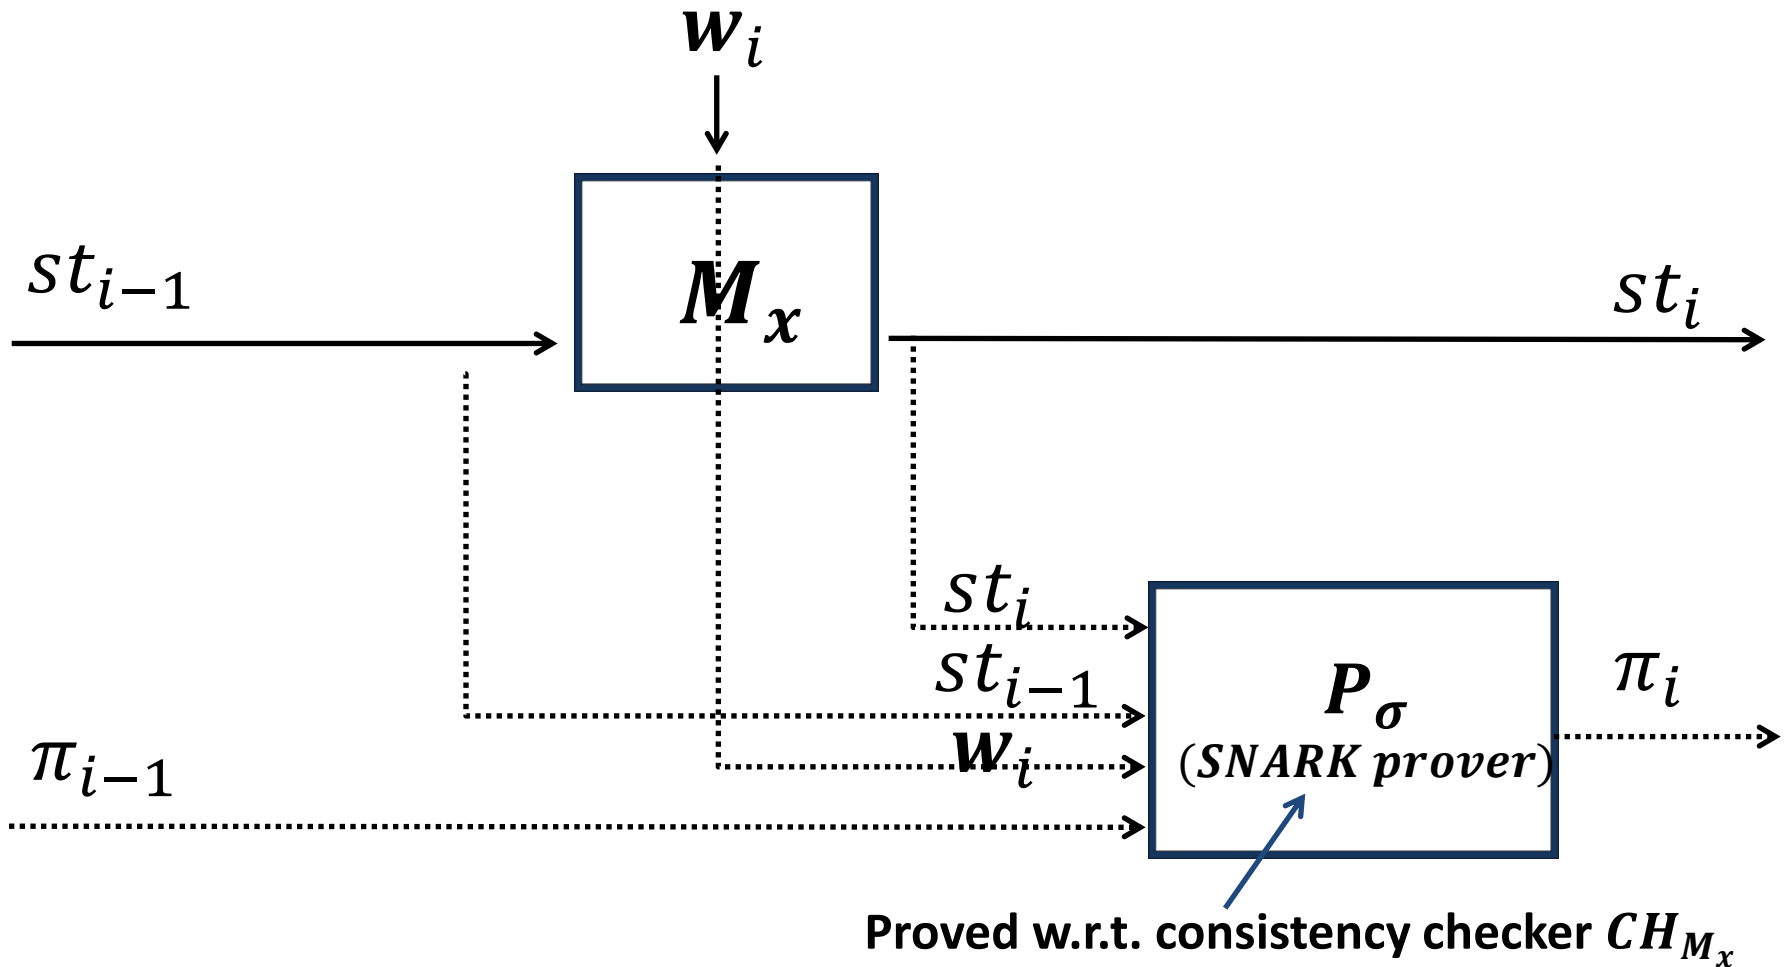
Compose short proof for current step with short proof for previous steps:


1. performed step i correctly
2. verified a proof π_{i-1} for correctness of steps $1 \dots i - 1$

Augment computation $M(x, w)$ with consistency proofs




Augment computation $M(x, w)$ with consistency proofs





Recursive Consistency Checker CH_{M_x}

Input: st_i witness: $(st_{i-1}, \pi_{i-1}, w_i)$



Recursive Consistency Checker CH_{M_x}

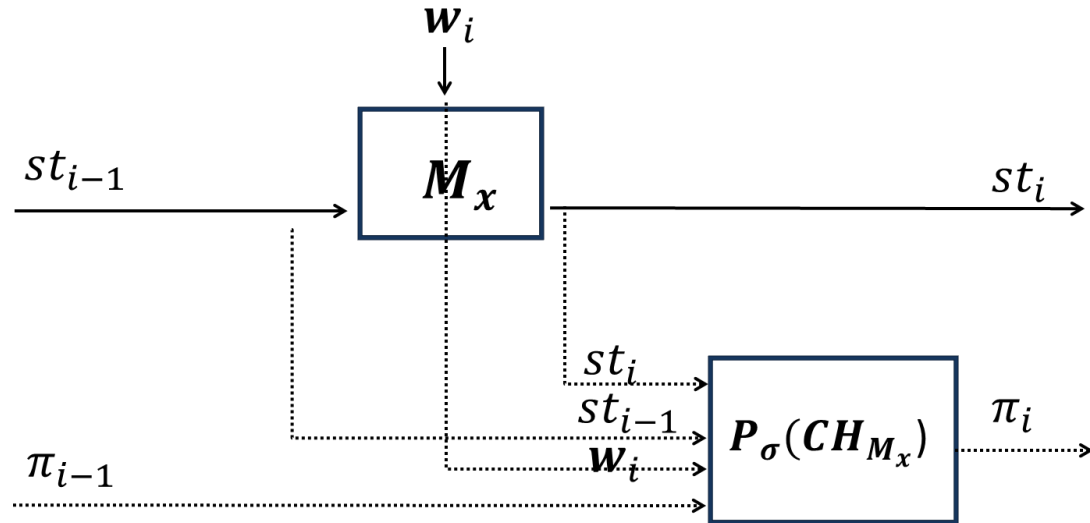
Input: st_i witness: $(st_{i-1}, \pi_{i-1}, w_i)$

If $i = 0$ and $st_0 = \text{initial state}$, accept.

else check:

$$M_x(st_{i-1}; w_i) = st_i$$

$$V_\tau(CH_{M_x}, st_{i-1}, \pi_{i-1}) = \text{acc}$$



Recursive Consistency Checker CH_{M_x}

Input: st_i witness: $(st_{i-1}, \pi_{i-1}, w_i)$

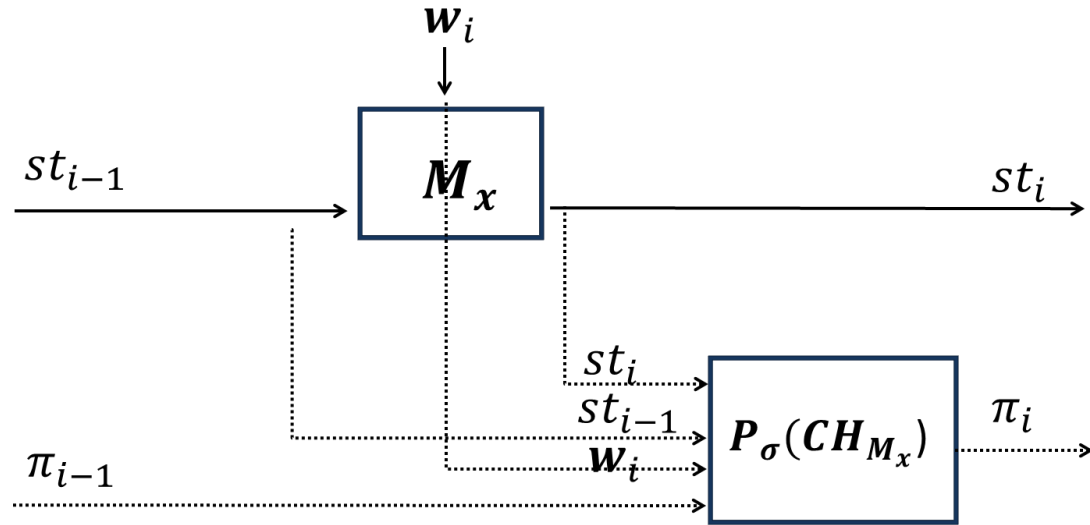
If $i = 0$ and $st_0 = \text{initial state}$, accept.

else check:

$$M_x(st_{i-1}; w_i) = st_i$$

$$V_\tau(CH_{M_x}, st_{i-1}, \pi_{i-1}) = \text{acc}$$


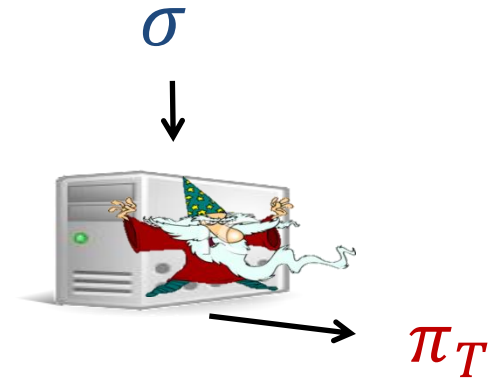
need recursion



Is the resulting proof sound?

$$V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1$$

$st_T = \text{"accept"}$



Recursive Consistency Checker \mathbf{CH}_{M_x}

Input: st_i witness: $(st_{i-1}, \pi_{i-1}, w_i)$

If st_i is initial state of M_x accept.

else check:


$$M_x(st_{i-1}; w_i) = st_i$$

V_τ accepts π_{i-1} for statement " \mathbf{CH}_{M_x} accepts st_{i-1} "

$$\begin{aligned} &\exists(\pi_{T-1}, st_{T-1}, w_T): \\ &M_x(st_{T-1}, w_T) = st_T \\ &V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1 \end{aligned}$$



$$\begin{aligned} &V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1 \\ &st_T = \text{"accept"} \end{aligned}$$



Recursive Consistency Checker \mathbf{CH}_{M_x}

Input: st_i

witness: $(st_{i-1}, \pi_{i-1}, w_i)$

If st_i is initial state of M_x accept.

else check:

$$M_x(st_{i-1}; w_i) = st_i$$

V_τ accepts π_{i-1} for statement " \mathbf{CH}_{M_x} accepts st_{i-1} "


$$\begin{aligned} & \exists(\pi_{T-1}, st_{T-1}, w_T): \\ & M_x(st_{T-1}, w_T) = st_T \\ & V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1 \end{aligned}$$



$$\begin{aligned} & V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1 \\ & st_T = \text{"accept"} \end{aligned}$$



$$\begin{aligned} & \exists(\pi_{T-2}, st_{T-2}, w_{T-1}): \\ & M_x(st_{T-2}, w_{T-1}) = st_{T-1} \\ & V_\tau(\mathbf{CH}_{M_x}; st_{T-2}; \pi_{T-2}) = 1 \end{aligned}$$



Recursive Consistency Checker \mathbf{CH}_{M_x}

Input: st_i

witness: $(st_{i-1}, \pi_{i-1}, w_i)$

If st_i is initial state of M_x accept.

else check:

$$M_x(st_{i-1}; w_i) = st_i$$

V_τ accepts π_{i-1} for statement " \mathbf{CH}_{M_x} accepts st_{i-1} "

$$\begin{aligned} &\exists(\pi_{T-1}, st_{T-1}, w_T): \\ &M_x(st_{T-1}, w_T) = st_T \\ &V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1 \end{aligned}$$


$$\begin{aligned} &V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1 \\ &st_T = \text{"accept"} \end{aligned}$$



$$\begin{aligned} &\exists(\pi_{T-2}, st_{T-2}, w_{T-1}): \\ &M_x(st_{T-2}, w_{T-1}) = st_{T-1} \\ &V_\tau(\mathbf{CH}_{M_x}; st_{T-2}; \pi_{T-2}) = 1 \end{aligned}$$



$$\begin{aligned} &\exists(\pi_0, st_0, w_1): \\ &M_x(st_0, w_1) = st_1 \\ &st_0 = \text{"start"} \end{aligned}$$



Recursive Consistency Checker \mathbf{CH}_{M_x}

Input: st_i

witness: $(st_{i-1}, \pi_{i-1}, w_i)$

If st_i is initial state of M_x accept.

else check:

$$M_x(st_{i-1}; w_i) = st_i$$

V_τ accepts π_{i-1} for statement " \mathbf{CH}_{M_x} accepts st_{i-1} "

$$\begin{aligned} & \exists(\pi_{T-1}, st_{T-1}, \mathbf{w}_T): \\ & M_x(st_{T-1}, \mathbf{w}_T) = st_T \\ & V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1 \end{aligned}$$

$$\begin{aligned} & V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1 \\ & st_T = \text{"accept"} \end{aligned}$$



$$\begin{aligned} & \exists(\pi_{T-2}, st_{T-2}, \mathbf{w}_{T-1}): \\ & M_x(st_{T-2}, \mathbf{w}_{T-1}) = st_{T-1} \\ & V_\tau(\mathbf{CH}_{M_x}; st_{T-2}; \pi_{T-2}) = 1 \end{aligned}$$




$$\begin{aligned} & \exists(\pi_0, st_0, \mathbf{w}_1): \\ & M_x(st_0, \mathbf{w}_1) = st_1 \\ & st_0 = \text{"start"} \end{aligned}$$

- ➔ **Computational** soundness isn't enough
- ➔ Need knowledge extraction
- ➔ Need to apply the extraction recursively.

The extraction guarantee of SNARKs

\forall prover P^*

ref string




st_i, π_i that
 V_τ accepts

\exists extractor E_{P^*}




$wit = st_{i-1}, \pi_{i-1}, w_i$
s.t. $CH_{M_x}(st_i, wit) = 1$



π_T

$$V_{\tau}(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1$$

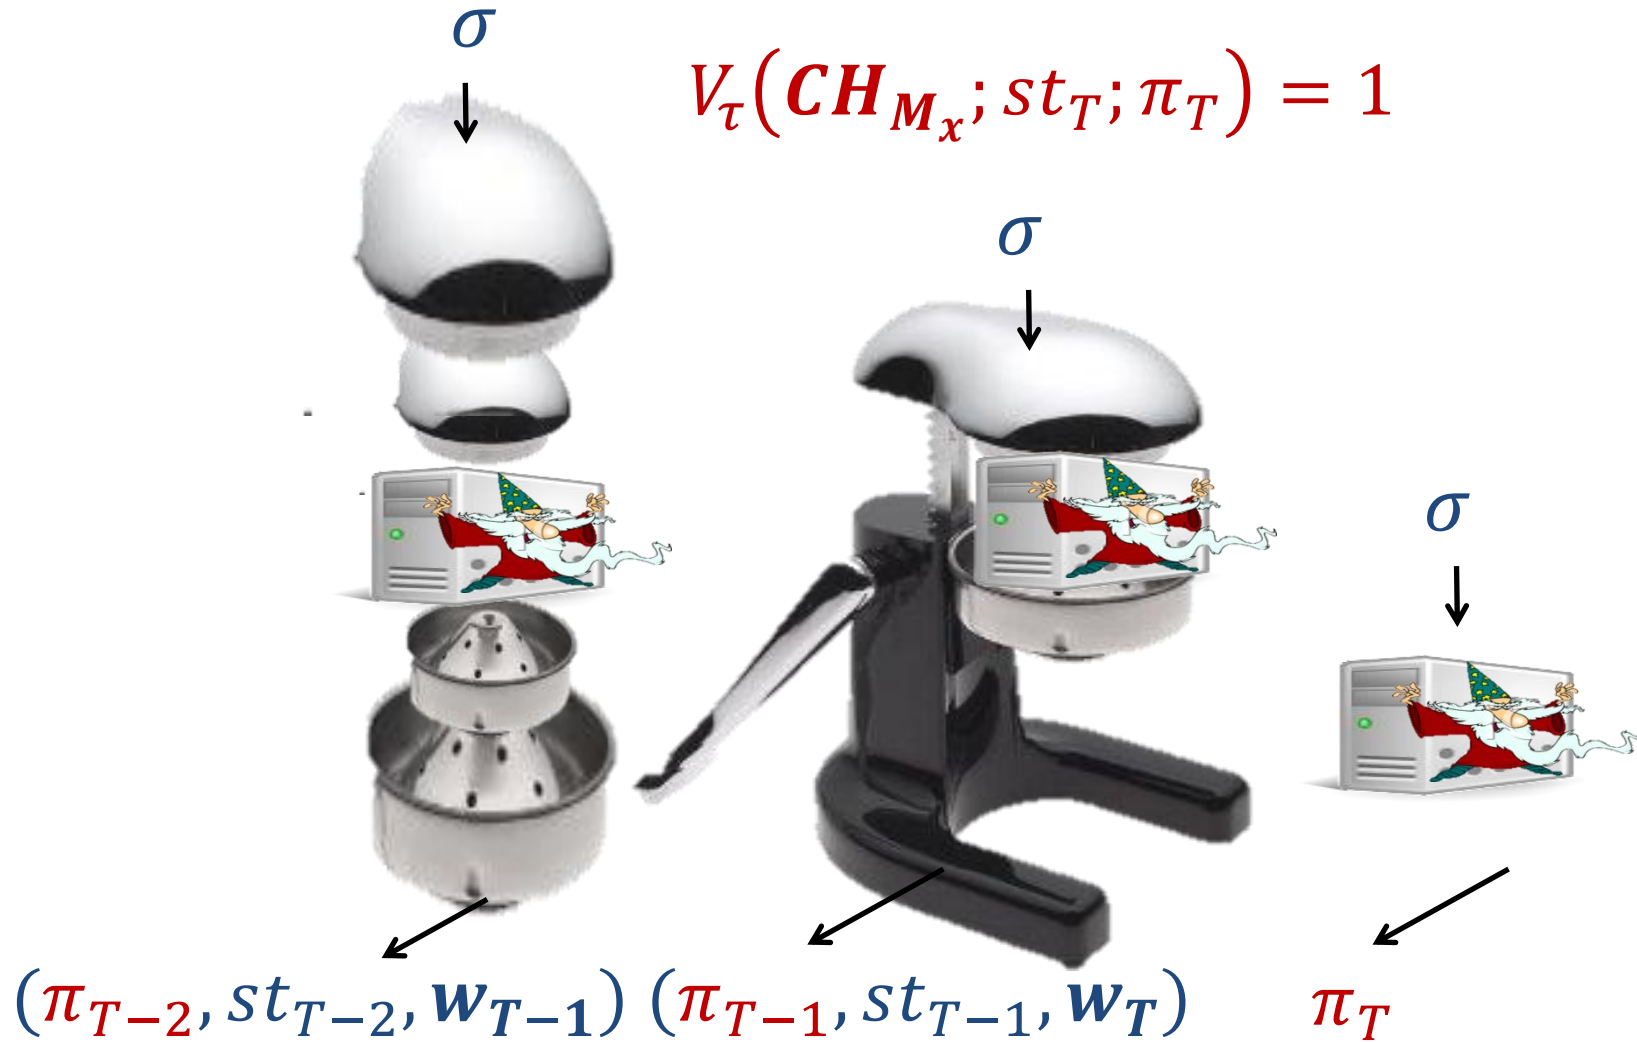
$$V_{\tau}(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1$$



$$V_{\tau}(\mathbf{CH}_{M_x}; st_{T-2}; \pi_{T-2}) = 1$$

$$V_{\tau}(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1$$

$$V_{\tau}(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1$$



$$V_\tau(\mathbf{CH}_{M_x}; st_{T-2}; \pi_{T-2}) = 1$$



$$V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1$$



$$V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1$$



$$(\pi_{T-2}, st_{T-2}, \mathbf{w}_{T-1}) \quad (\pi_{T-1}, st_{T-1}, \mathbf{w}_T)$$

$$\pi_T$$



$$V_\tau(\mathbf{CH}_{M_x}; st_{T-1}; \pi_{T-1}) = 1$$



$$V_\tau(\mathbf{CH}_{M_x}; st_T; \pi_T) = 1$$



$$(\pi_{T-2}, st_{T-2}, \mathbf{w}_{T-1}) \quad (\pi_{T-1}, st_{T-1}, \mathbf{w}_T)$$


$$\pi_T$$



V_T

How large is this extractor?
(or mcextractor)

$(\pi_T) = 1$



$(\pi_{T-2}, st_{T-2}, w_{T-1})$ $(\pi_{T-1}, st_{T-1}, w_T)$

π_T



V_T

How large is this extractor?
(or mcextractor)

$(\pi_T) = 1$



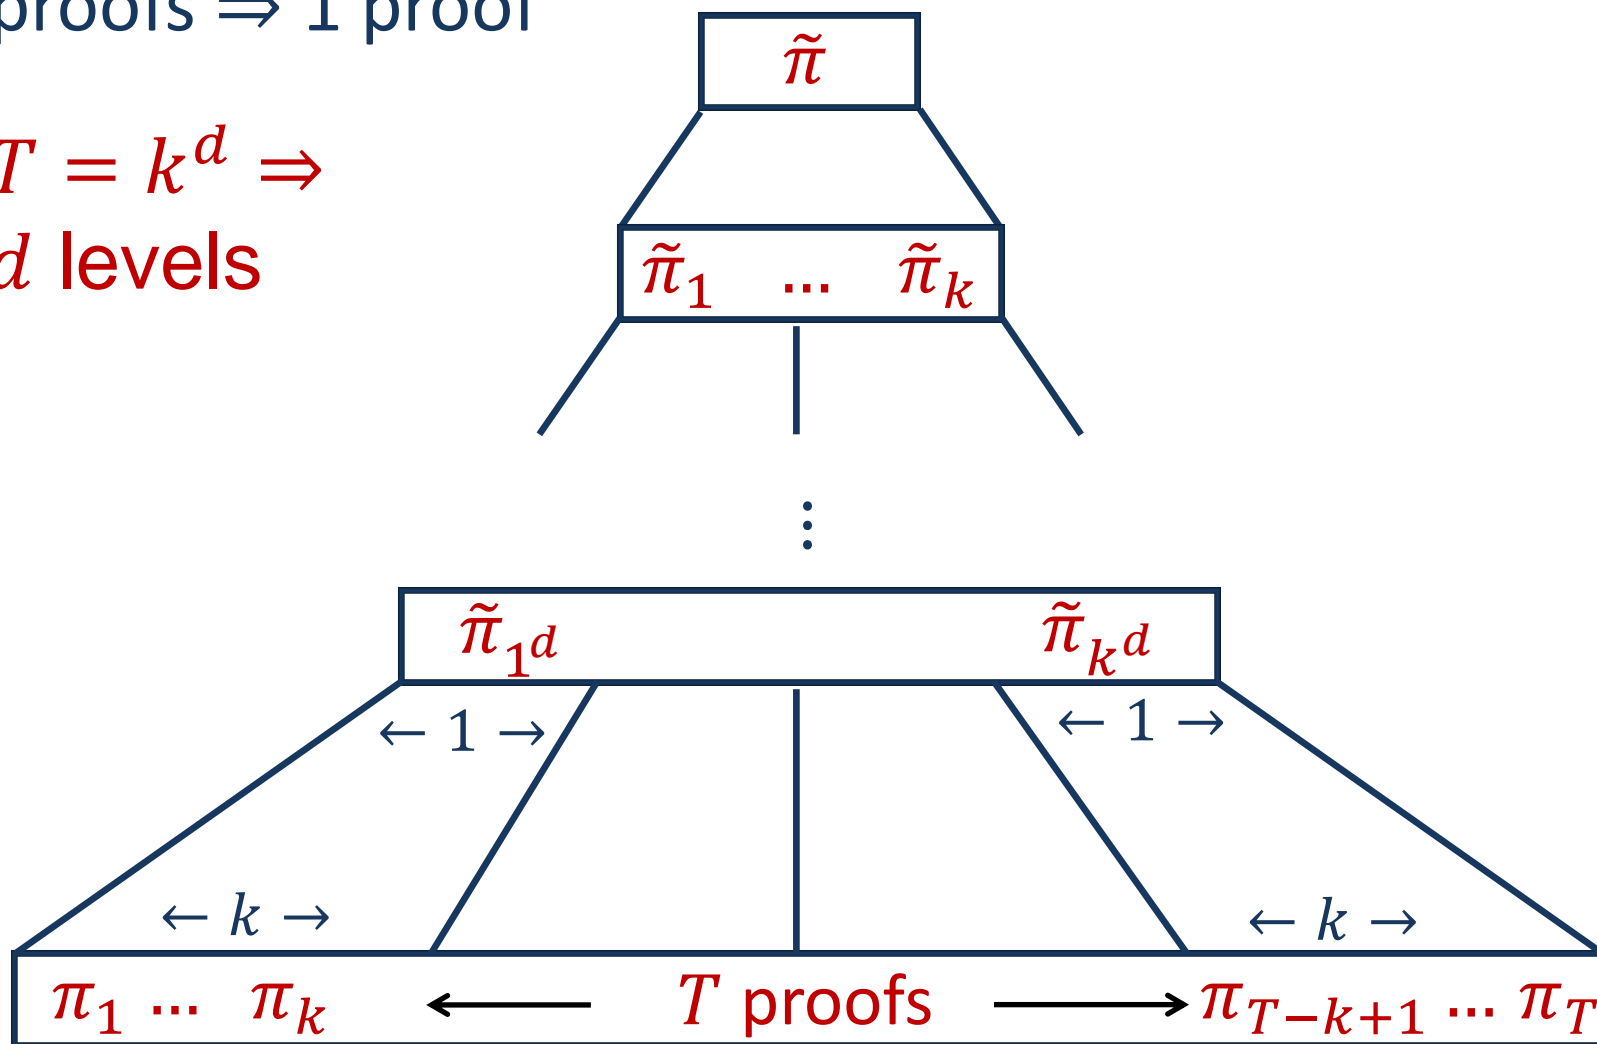
Typical recursive extraction problem:
each level incurs a poly blowup! $\Rightarrow |E| = k^{2^{O(d)}}$
can only deal with $O(1)$ levels

$(\pi_{T-2}, st_{T-2}, w_{T-1})$ $(\pi_{T-1}, st_{T-1}, w_T)$ π_T

A solution: aggregate proofs in a wide tree

k proofs \Rightarrow 1 proof

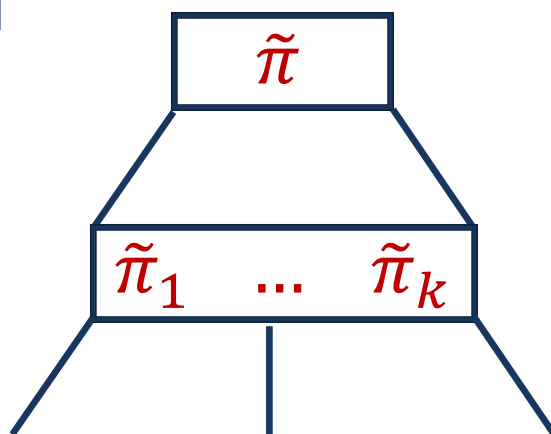
If $T = k^d \Rightarrow$
 d levels




A solution: aggregate proofs in a wide tree

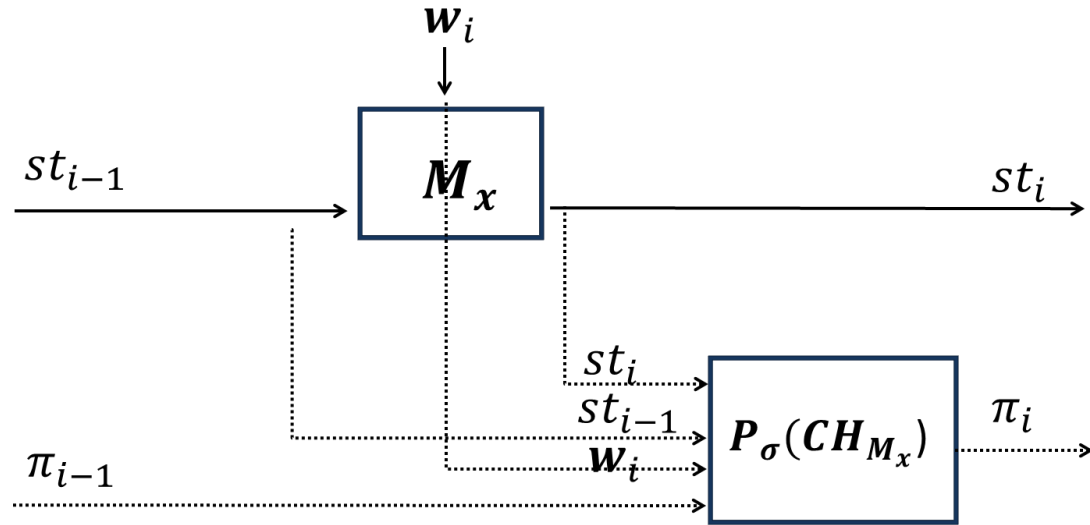
k proofs \Rightarrow 1 proof

If $T = k^d \Rightarrow$
 d levels




The tree is constructed dynamically with only $\text{poly}(k)$ overhead






Is the resulting proof sound?






So Far: Preprocessing cost is proportional to single-step computation.



So Far: Preprocessing cost is proportional to single-step computation.

But how large is a single-step computation?



So Far: Preprocessing cost is proportional to single-step computation.

But how large is a single-step computation?



→ bounded only by S , which can be as large as T ...

→ preprocessing stage can still be $\text{poly}(T)$...

Idea:

Move from machines with large memory to machines with:

- “small” trusted memory
- “big” untrusted memory



witness



M

trusted (big) mem



witness

memory transcript of entire execution of M

+ consistency info


\tilde{M}

trusted mem

← *only poly(k)*

A computational reduction using CRH

[Blum Evans Gemmell Kannan Naor 94,
Ben-Sasson Chiesa Genkin Tromer 12]




Or

h -collisions can be (eff.)
extracted from mem

A computational reduction using CRH

[Blum Evans Gemmell Kannan Naor 94,
Ben-Sasson Chiesa Genkin Tromer 12]




Or

h -collisions can be (eff.)
extracted from mem \rightarrow $\exists mem$ not enough
need knowledge

A computational reduction using CRH

[Blum Evans Gemmell Kannan Naor 94,
Ben-Sasson Chiesa Genkin Tromer 12]



Dynamic Merkle-hashing
of memory

\tilde{M}_h runs in time $T_M \cdot \text{poly}(k)$, space $\text{poly}(k)$
and *mem* computed from (\mathbf{x}, \mathbf{w}) in
time $T_M \cdot \text{poly}(k)$ & space $S_M \cdot \text{poly}(k)$

→ A single-step computation is now of size
 $\text{poly}_h(k)$

(subsequent steps can be computed dynamically
preserving time and space of original computation)

what's left?
...SNARK verification

Input: st_i witness: $(\pi_{i-1}, st_{i-1}, \mathbf{w}_i)$

If st_i is initial state of \tilde{M}_x **accept**.

else check:

$$\tilde{M}_x(st_{i-1}; \mathbf{w}_i) = st_i$$

V_τ **accepts** π_{i-1} for statement $CH_{\tilde{M}_x}, st_{i-1}$

only $\text{poly}_V(k)$,
independently of preprocessing limit

what's left?
...SNARK verification

Input: st_i witness: $(\pi_{i-1}, st_{i-1}, w_i)$

If st_i is initial state of \tilde{M}_x accept.

else check:

$$\tilde{M}_x(st_{i-1}; w_i) = st_i$$

V_τ accepts π_{i-1} for statement $CH_{\tilde{M}_x}, st_{i-1}$

only $\text{poly}_V(k)$,
independently of preprocessing limit
 \Rightarrow budget only for $\text{poly}_V(k) + \text{poly}_h(k)$

Bye Bye

Long Preprocessing...

Part I:

How to Bootstrap a SNARK
in Public



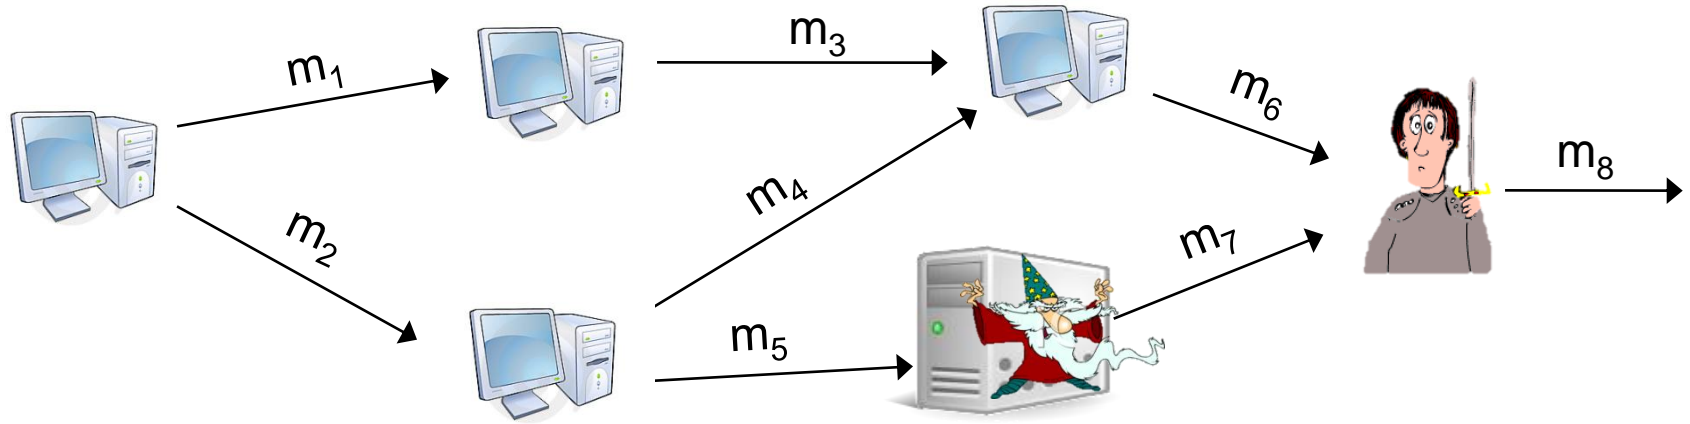
Part II:

Part I (again) and Beyond
with Proof Carrying Data

In SNARKs: one prover and one verifier

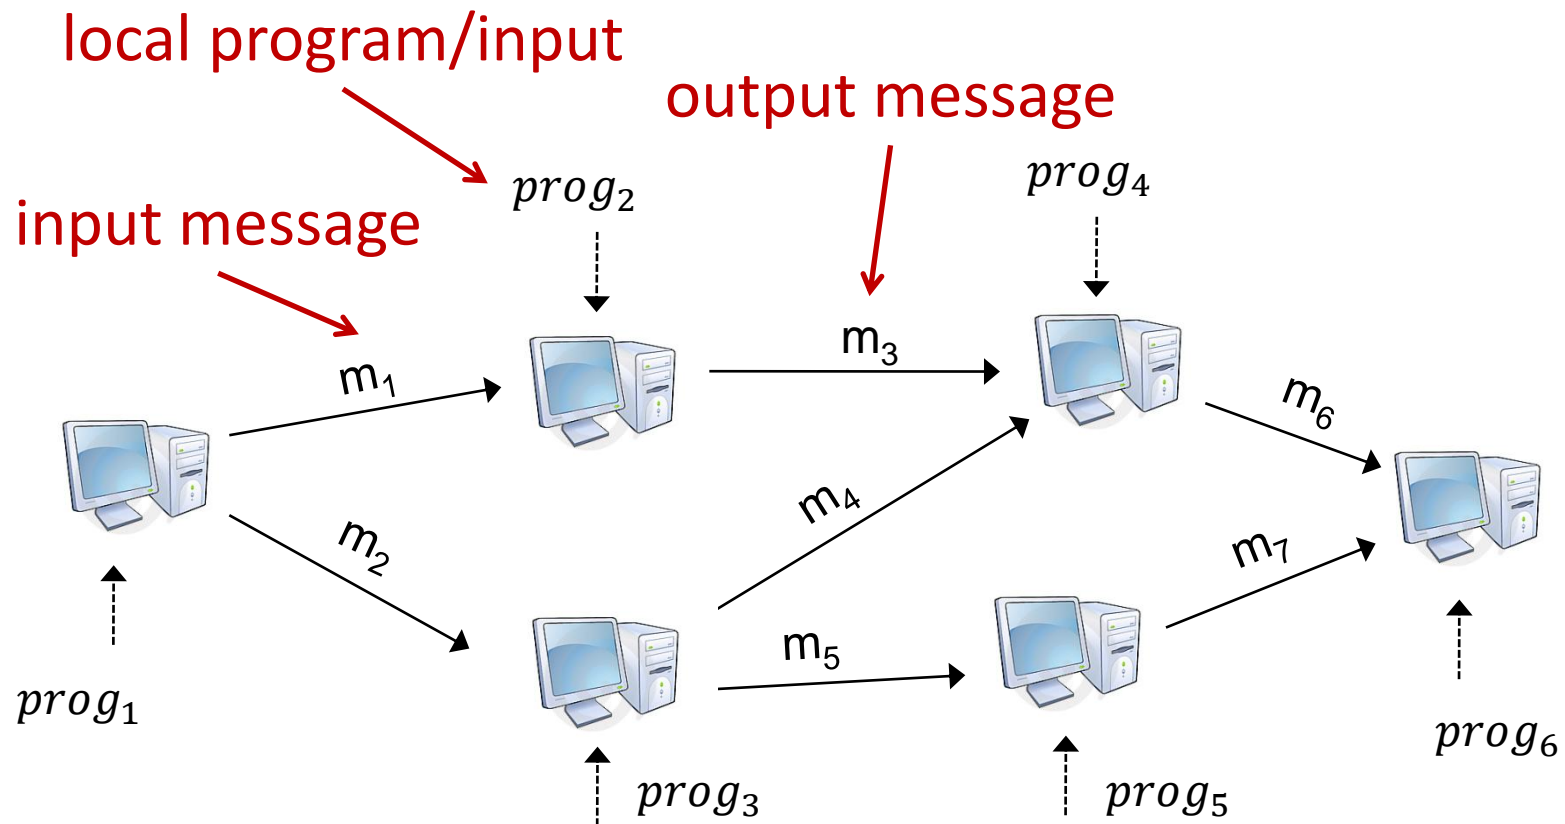


But sometimes in life...



Computations involve many parties
each party has its own:
role, capabilities, friends, enemies,...

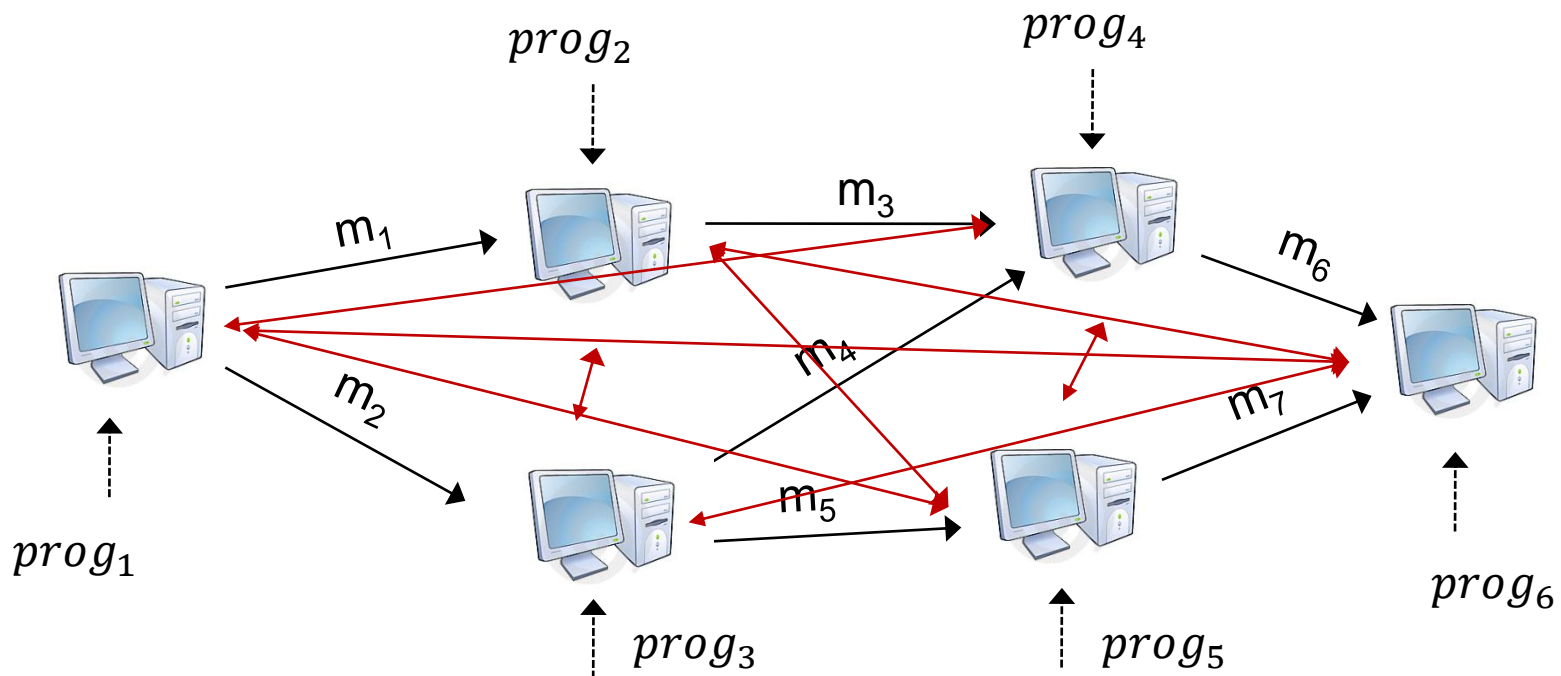
How can we enforce general correctness properties of distributed computations?



Use MPC?

enforce **any** property of **all** the inputs/outputs of all parties

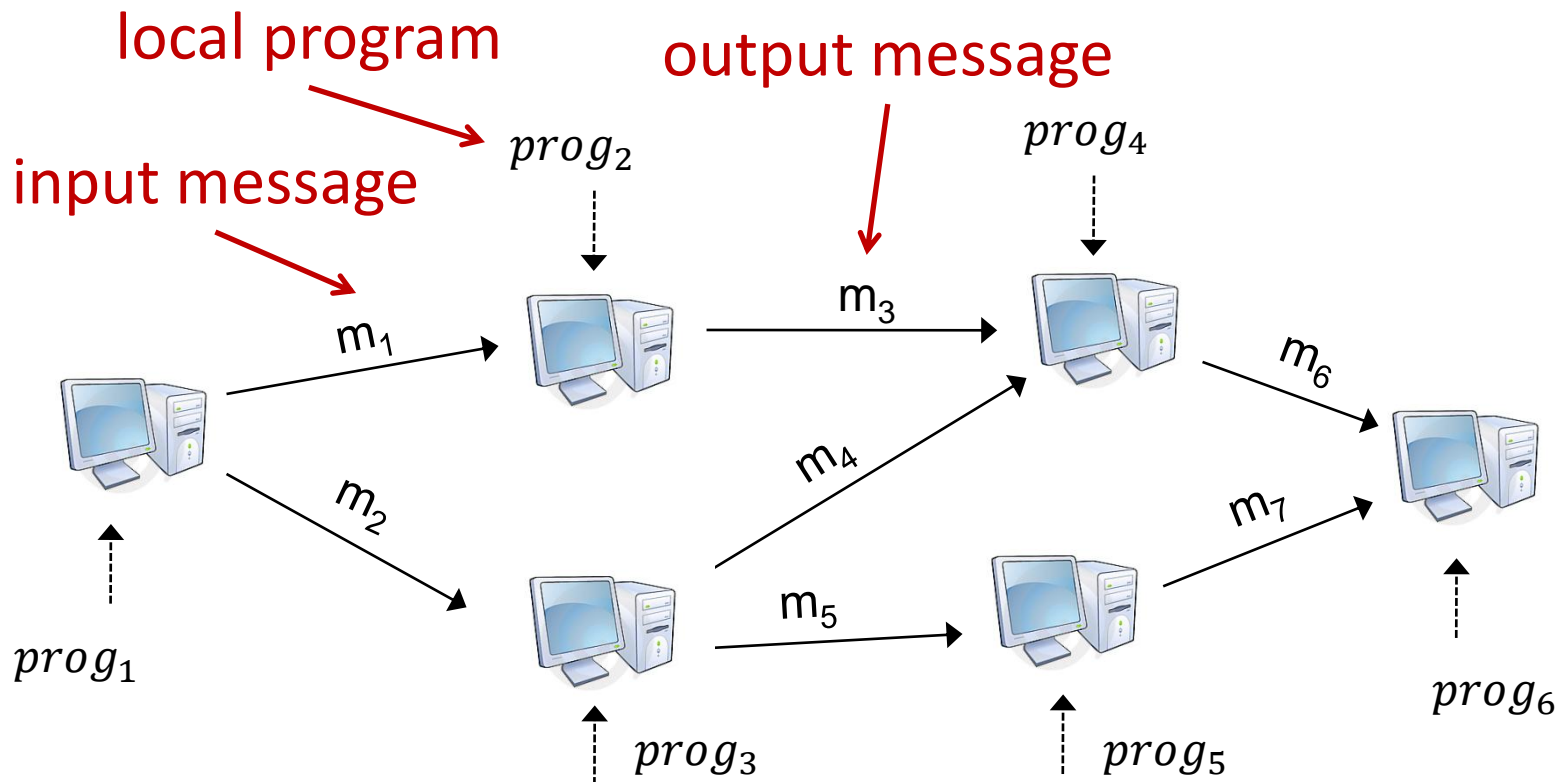
but: large overhead: all parties must communicate with each other
(necessary, e.g. Byzantine agreement)



A relaxed question:
how to enforce **local** properties?

Local property = property of the view of a single node

Example: ensure that the program executed at every node was signed by system admin
if property holds **everywhere** → global meaning



Proof Carrying Data (PCD)

[Chiesa Tromer 10]

Goal:

Guarantee “local properties” **while respecting the original computation:**

- preserve communication graph
- minimal computational overhead

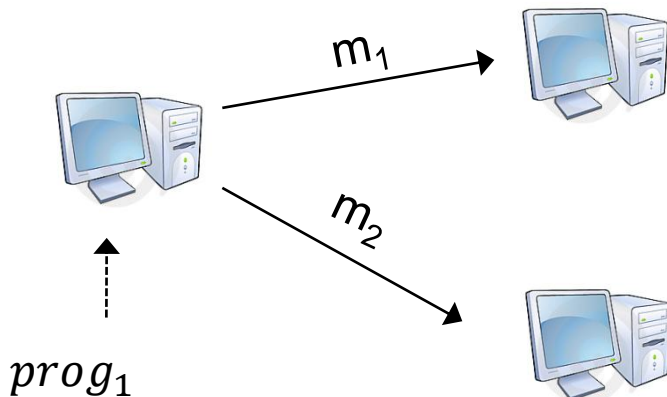
The original computation

- Can be viewed as a DAG evolving over time
- nodes have input and output messages
+ a local program (with embedded inputs).




The original computation

- Can be viewed as a DAG evolving over time
- nodes have input and output messages + a local program (with embedded inputs).



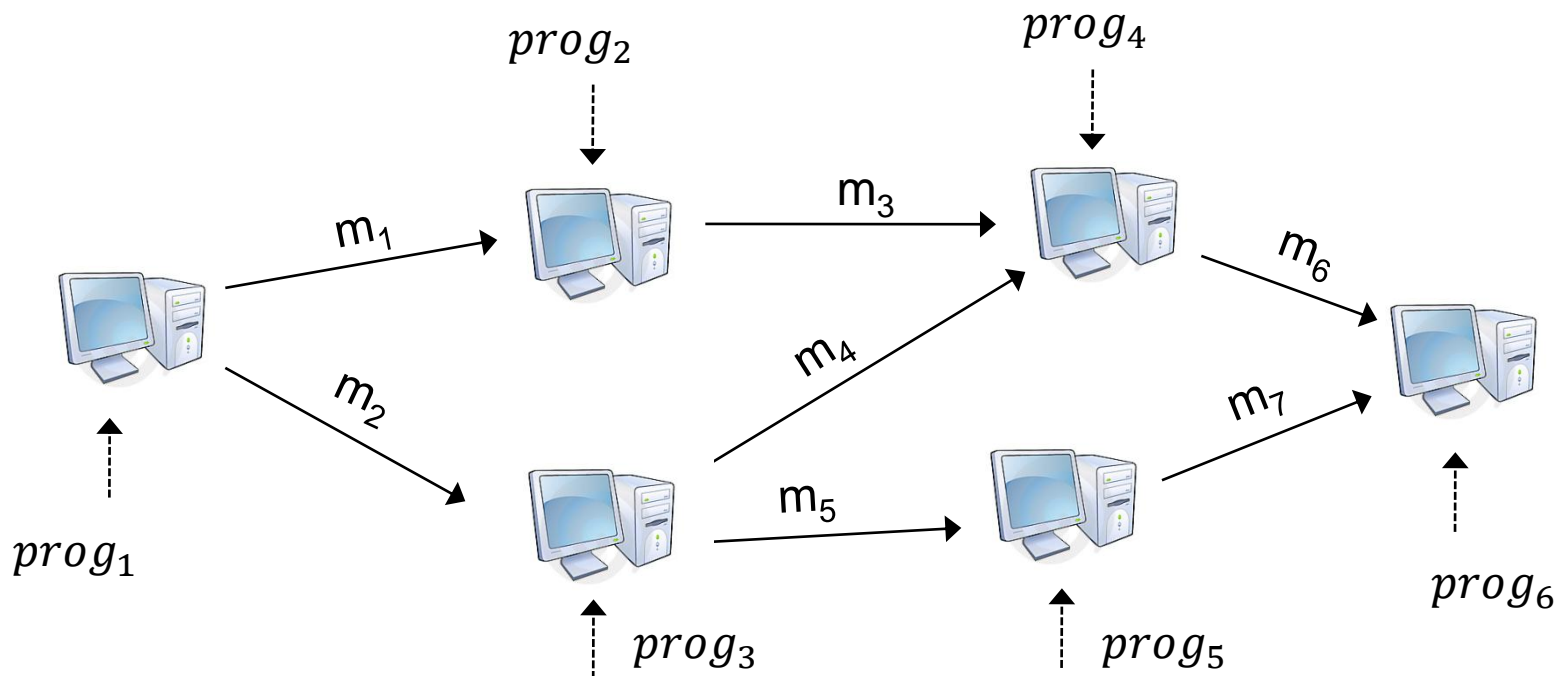
The original computation

- Can be viewed as a DAG evolving over time
- nodes have input and output messages + a local program (with embedded inputs).



The original computation

- Can be viewed as a DAG evolving over time
- nodes have input and output messages + a local program (with embedded inputs).



Local properties as \mathcal{C} -compliance

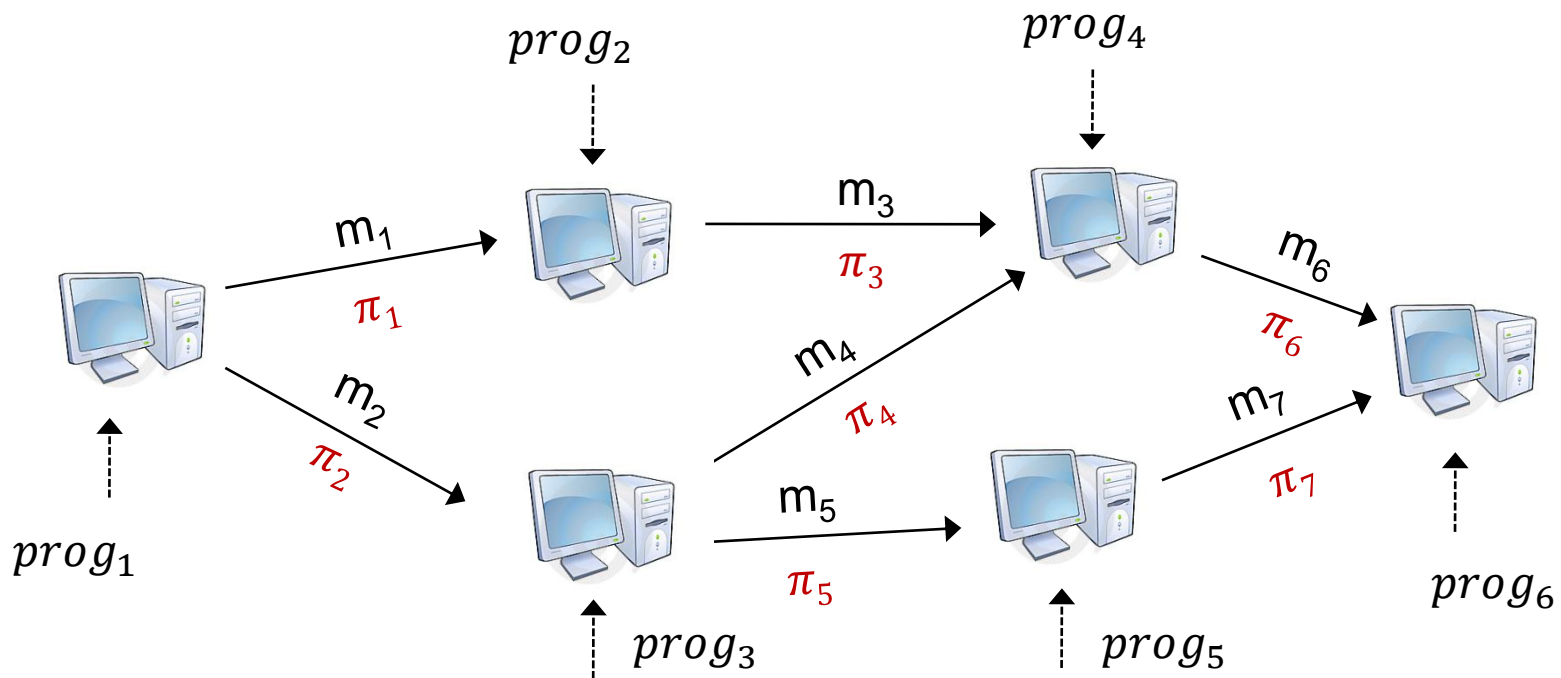
$\mathcal{C}(prog, m_{in}, m_{out})$ is a predicate specifying a local property, e.g.:

- \mathcal{C}_{adm} : “ $prog = (M, s)$ where s is an admin signature on M and $M(m_{in}) = m_{out}$ ”
- \mathcal{C}_{JVM} : “ $prog$ is a JAVA program and $JVM(prog, m_{in}) = m_{out}$ ”
- \mathcal{C}_{M_x} : “ $prog = w_i, m_{in} = st_{i-1}, m_{out} = st_i$ and $M_x(st_{i-1}, w_i) = st_i$ ”



A PCD system

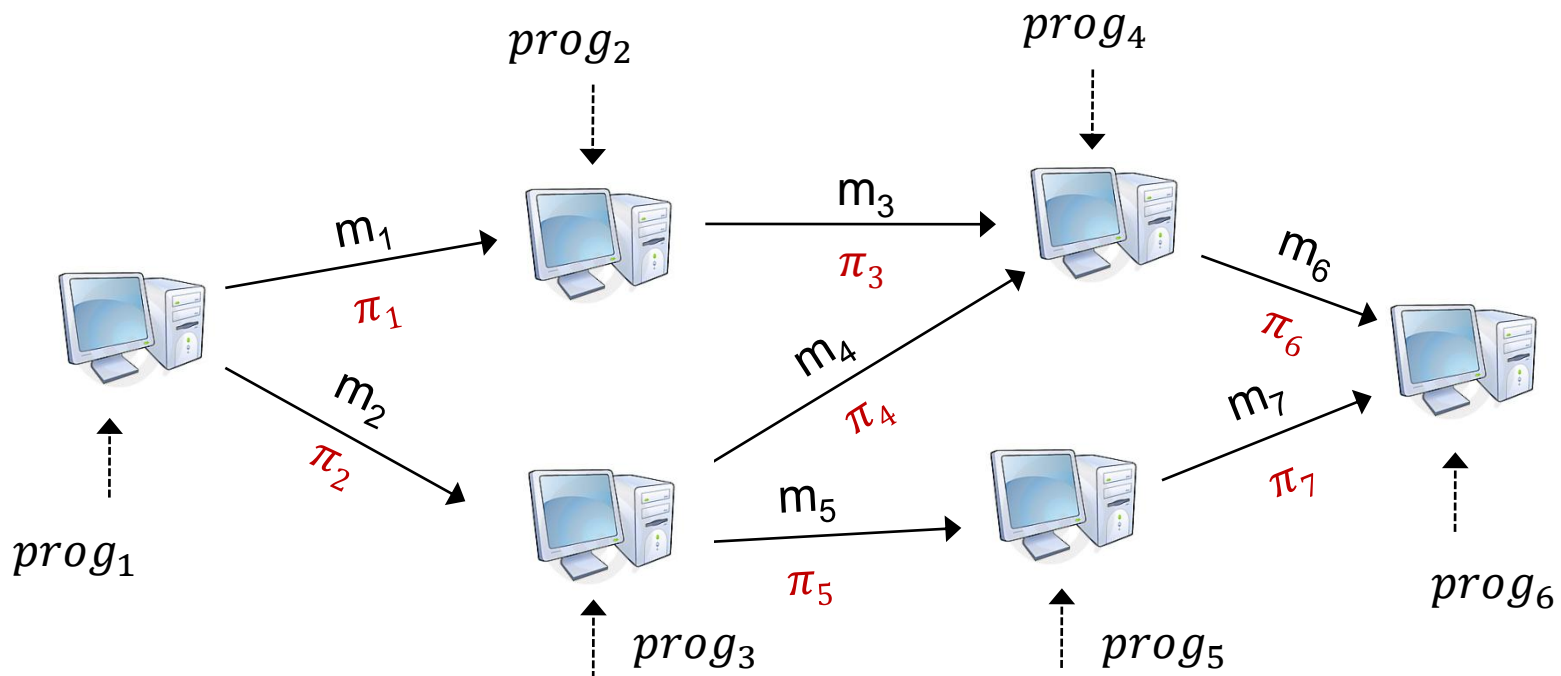
- compile on-the-fly original computation
- (short) proofs are appended to messages



A PCD system

- compile on-the-fly original computation
- (short) proofs are appended to messages

Note: not all properties can be verified this way.
Eg, verifying that $m_1 = m_2$ requires additional interaction.



How to construct PCDs?

[CT10]: Using an abstract
signature card

How to construct PCDs?

This work: SNARK composition

Results (revisited): General transformations

Publicly-verifiable SNARKs
with (resp. without) preprocessing

SNARK recursive composition

Publicly-verifiable PCDs **for constant depth graphs**
with (resp. without) preprocessing

Results (revisited): General transformations

Publicly-verifiable SNARKs
with (resp. without) preprocessing

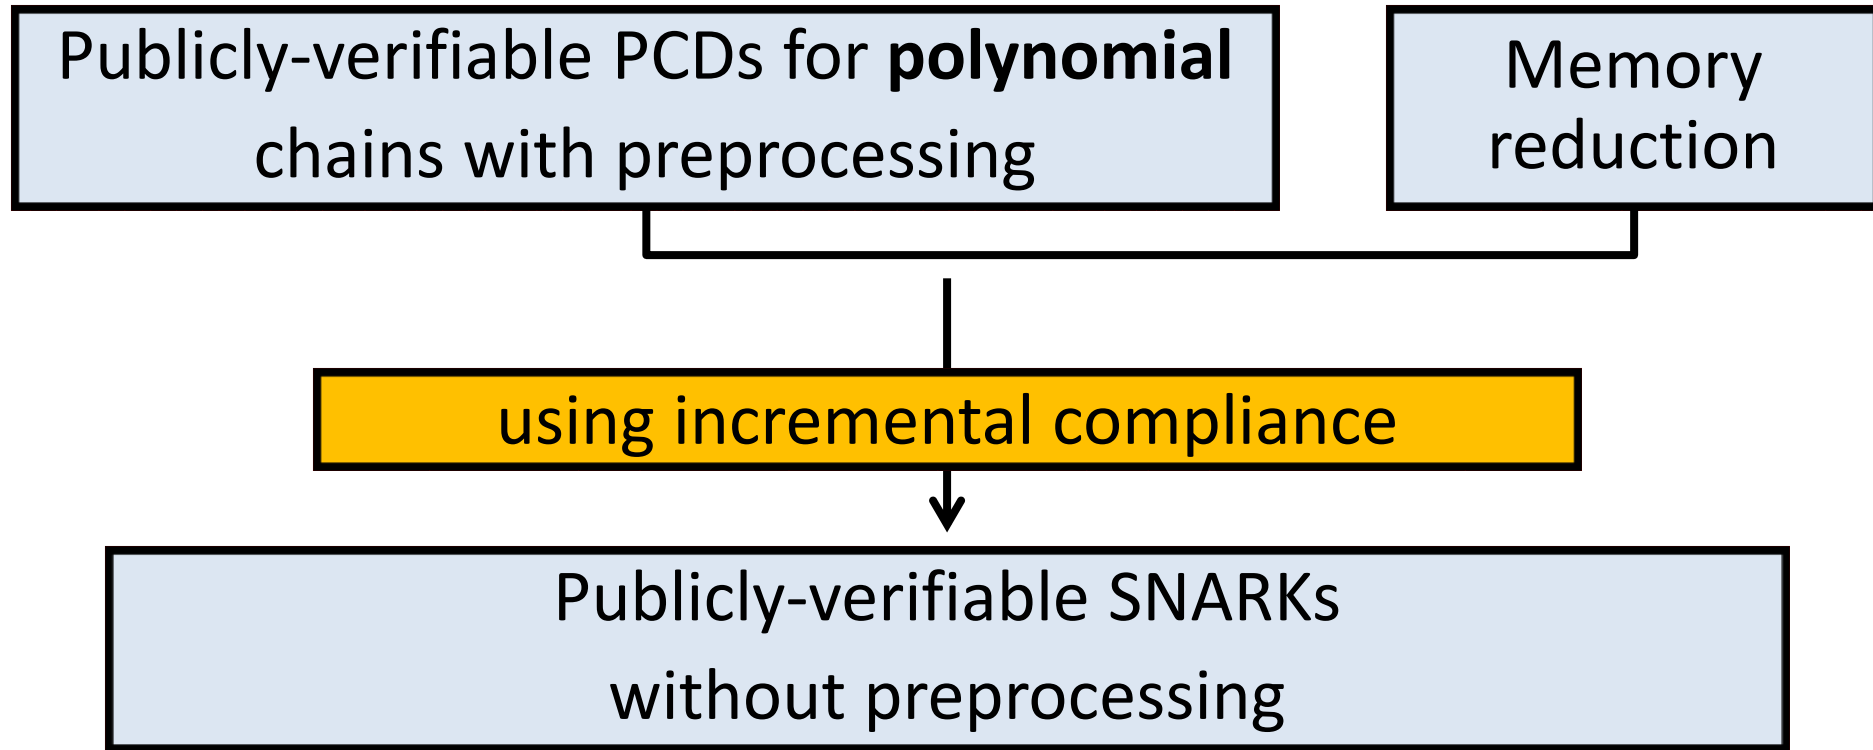
SNARK recursive composition

Publicly-verifiable PCDs **for constant depth graphs**
with (resp. without) preprocessing

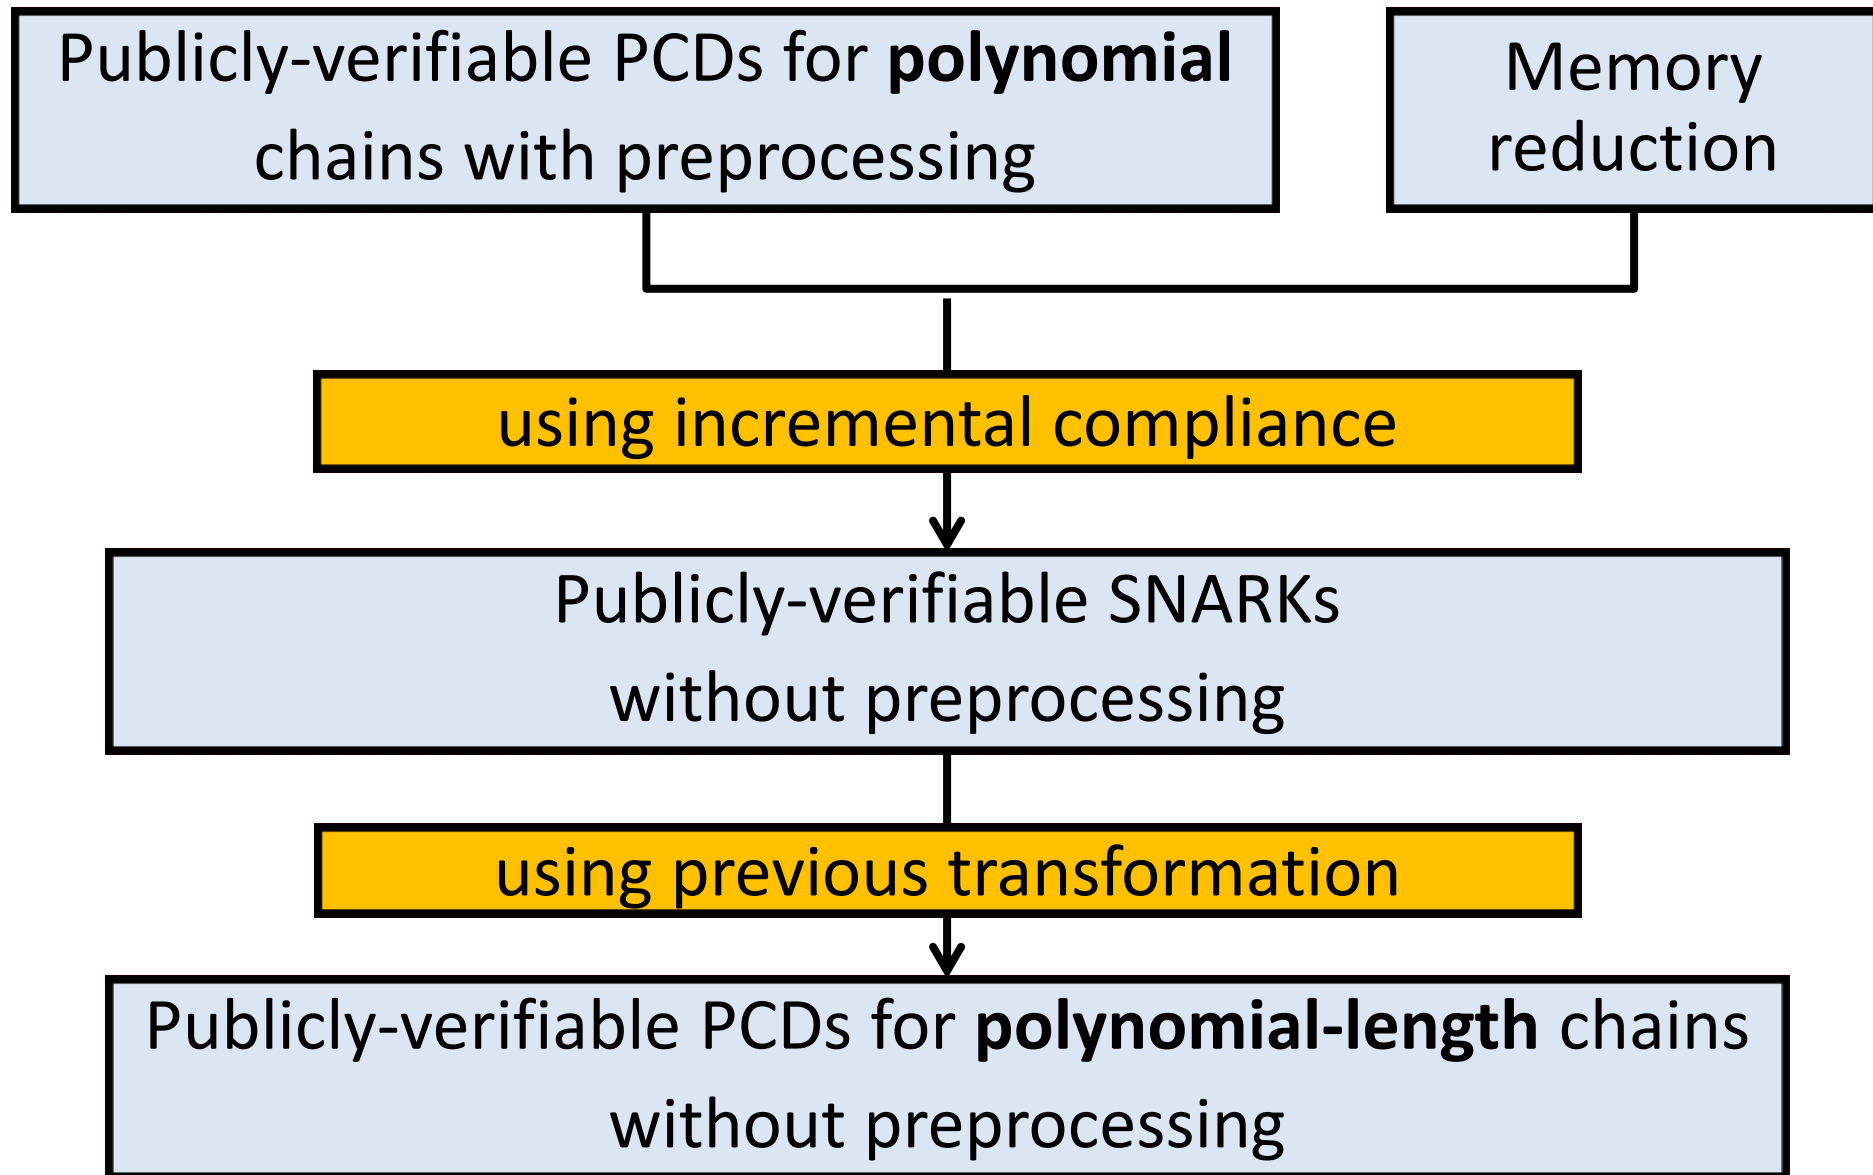
General transformation
of path-compliance to tree-compliance

Publicly-verifiable PCDs for **polynomial-length** chains
with (resp. without) preprocessing

Results (revisited): Eliminating expensive preprocessing



Results (revisited): Eliminating expensive preprocessing



A bonus:

privately-verifiable SNARKs also compose!

A bonus:
privately-verifiable SNARKs also compose!

To compose SNARKs we used public-verifiability
proved “I verified a SNARK”

A bonus:
privately-verifiable SNARKs also compose!

To compose SNARKs we used public-verifiability
proved “I verified a SNARK”

Looks surprising... but doable (using FHE).

→ All the PCD results have their
privately-verifiable analogs

Question:

which security goals we express
using the PCD language?

We've seen some examples
others include: targeted-malleability [BSW11],
computing on authenticated Data,...

Other properties?