


# Advancing Scalability in Decentralized Storage: a Novel Approach to Proof-of-Replication via Polynomial Evaluation


Giuseppe Ateniese<sup>1</sup>, Foteini Baldimtsi<sup>1</sup>,  
Matteo Campanelli<sup>2</sup>, **Danilo Francati**<sup>1</sup>, Ioanna Karantaidou<sup>1</sup>

<sup>1</sup>George Mason University, <sup>2</sup>Matter Labs

# Proof of Space (PoS)




Verifier



Prover

Commit: I dedicate  
1GB of storage space


# Proof of Space (PoS)



(eco-friendly) Alternative to Proof-of-Work

**Applications:** spam prevention, DDoS attack resistance, Sybil-resistant blockchain consensus

# Proof of Space (PoS)




no specific/useful data are being stored by the prover

(eco-friendly) Alternative to Proof-of-Work



**Applications:** spam prevention, DDoS attack resistance, Sybil-resistant blockchain consensus

# Proof of Replication (PoRep)

Setup  
Phase



Verifier

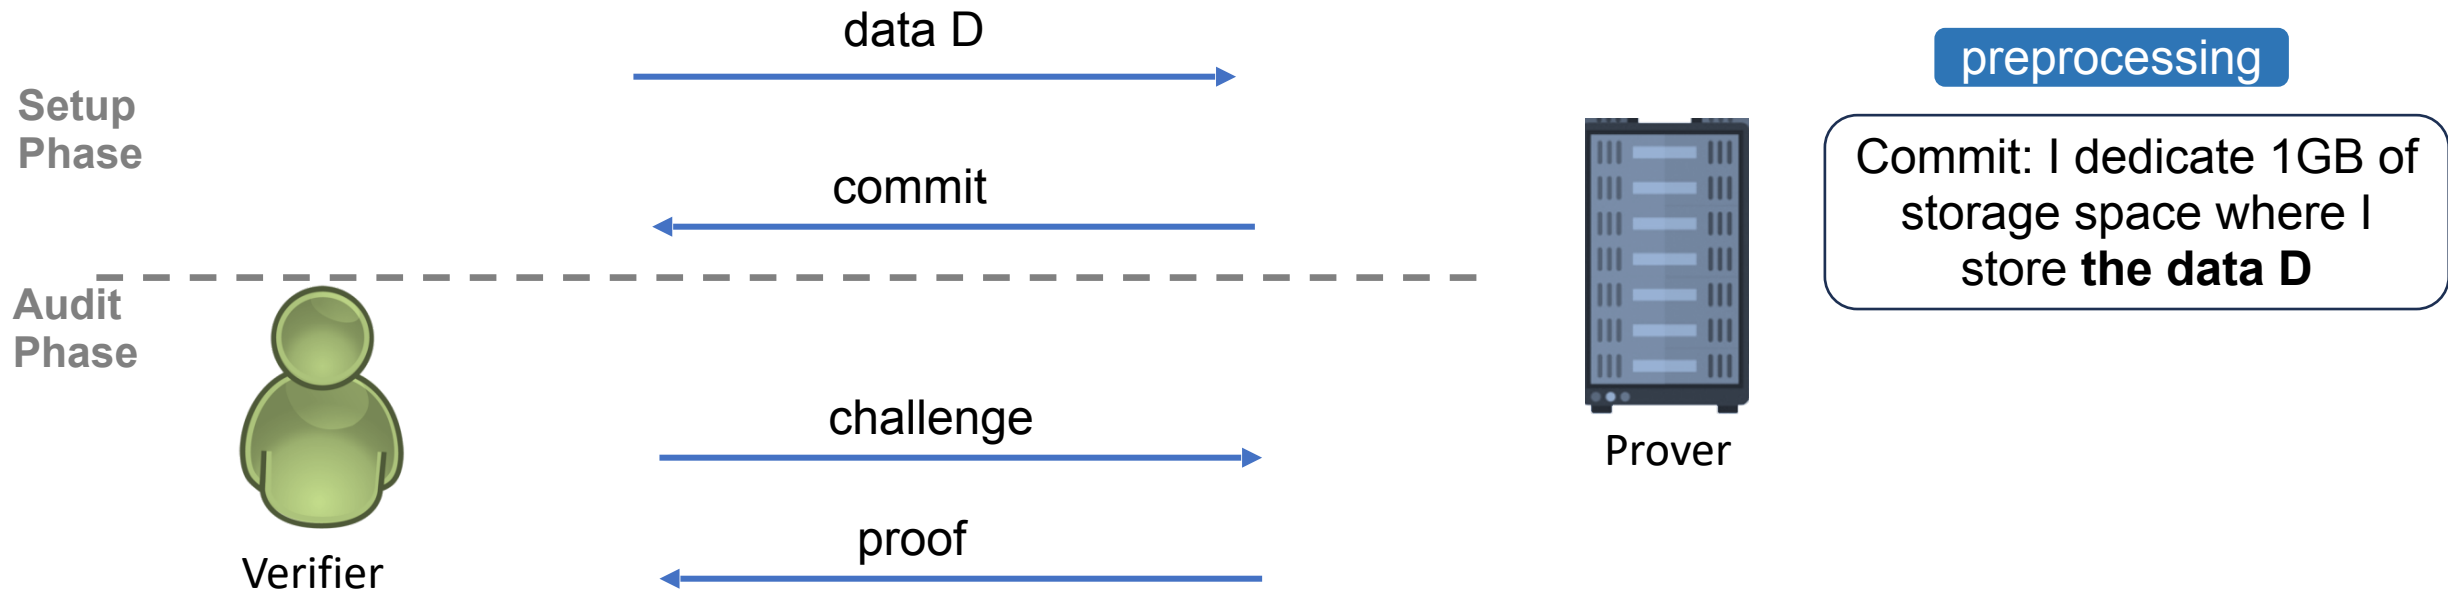


Prover


preprocessing

Commit: I dedicate 1GB of storage space where I store **the data D**

# Proof of Replication (PoRep)



# Proof of Replication (PoRep)



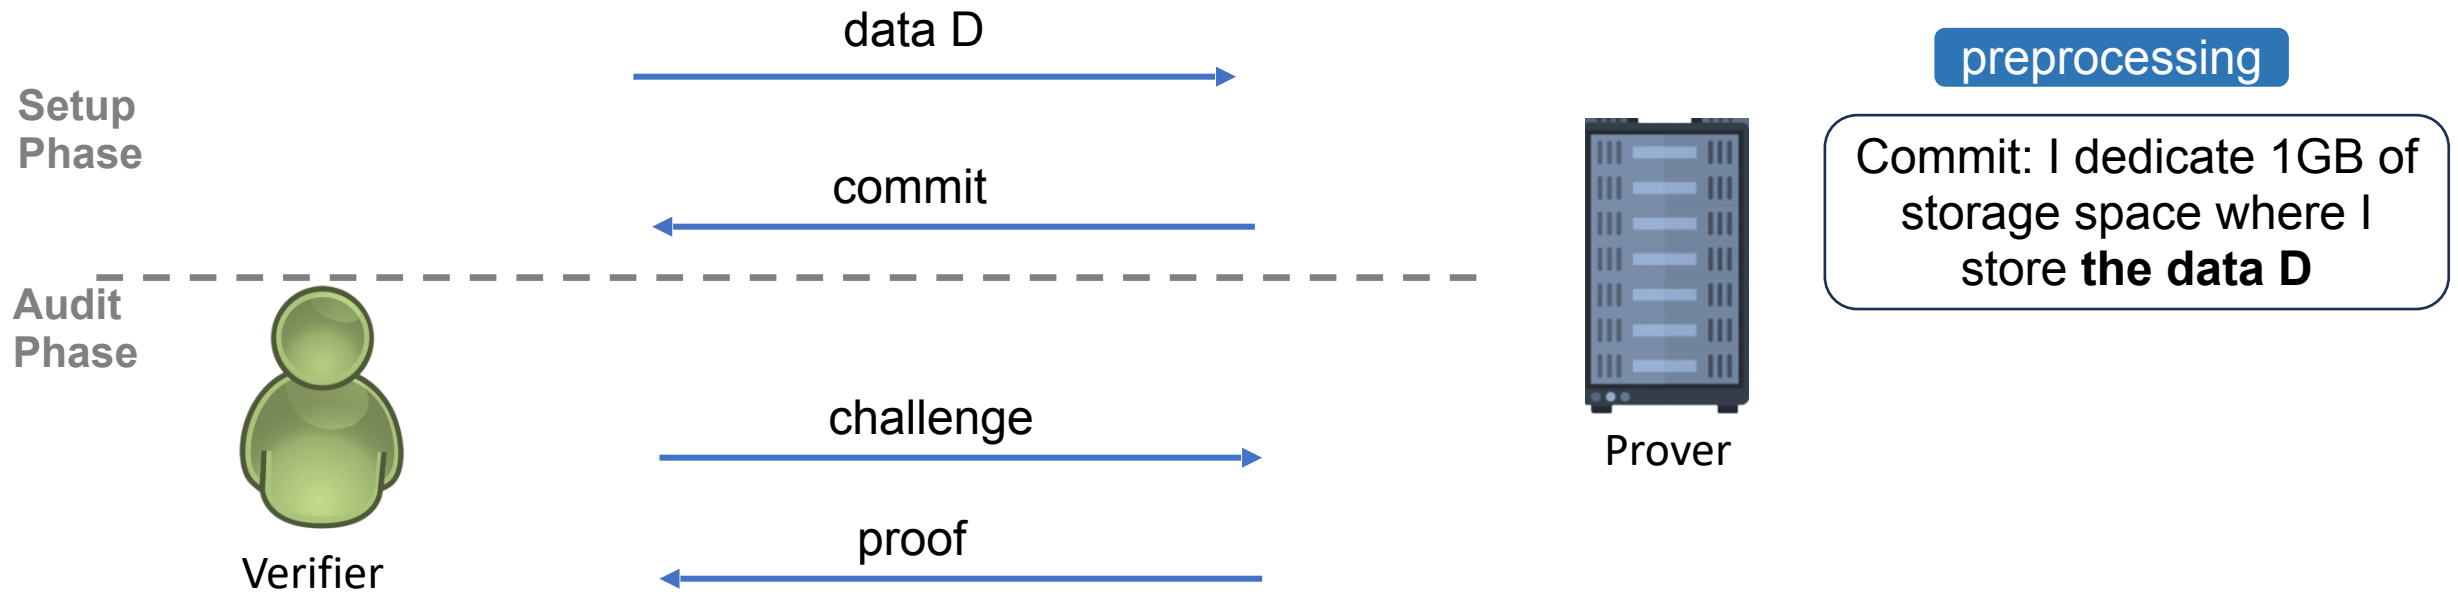
No waste of the dedicated space!

**PoRep** guarantees that the prover is dedicating unique storage resources per replica of the data



decentralized and verifiable file storage

# Proof of Replication (PoRep)



No waste of the dedicated space!

**PoRep** guarantees that the prover is dedicating unique storage resources per replica of the data



decentralized and verifiable file storage



“Filecoin (FIL) is an open-source, public cryptocurrency and digital payment system intended to be a blockchain-based cooperative digital storage and data retrieval method.”



# PoRep Definition

# PoRep Definition

preprocessing



audit phase

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$  public keys  $ek, pk, vk$
- $\text{Encode}(m, ek, id) \rightarrow$  encoding  $c$  and digest  $h$

Publish  $h$

---

preprocessing

audit phase

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$  public keys  $ek, pk, vk$
- $\text{Encode}(m, ek, id) \rightarrow$  encoding  $c$  and digest  $h$

preprocessing

Publish  $h$

---

- $\text{Prove}(pk, \text{challenge}, c) \rightarrow \pi$

audit phase

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$  public keys  $ek, pk, vk$
- $\text{Encode}(m, ek, id) \rightarrow$  encoding  $c$  and digest  $h$

preprocessing

Publish  $h$

---

- $\text{Prove}(pk, \text{challenge}, c) \rightarrow \pi$
- $\text{Verify}(vk, h, \pi) \rightarrow 0/1$

audit phase

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$   $\text{ek, pk, vk}$  public keys
- $\text{Encode}(m, \text{ek}, \text{id}) \rightarrow$  encoding  $c$  and digest  $h$

preprocessing

Publish  $h$

---

- $\text{Prove}(\text{pk}, \text{challenge}, c) \rightarrow \pi$
- $\text{Verify}(\text{vk}, h, \pi) \rightarrow 0/1$
- $\text{Decode}(\text{ek}, \text{id}, c) \rightarrow m$

audit phase

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow \text{ek, pk, vk}$  public keys
- $\text{Encode}(m, \text{ek}, \text{id}) \rightarrow \text{encoding } c \text{ and digest } h$

preprocessing

Publish  $h$

---

- $\text{Prove}(\text{pk}, \text{challenge}, c) \rightarrow \pi$
- $\text{Verify}(\text{vk}, h, \pi) \rightarrow 0/1$
- $\text{Decode}(\text{ek}, \text{id}, c) \rightarrow m$

audit phase

**Properties:** encoding correctness, proof completeness, replication and extraction

# Proof of Replication (PoRep)





# Proof of Replication (PoRep)

Proof of Storage  
PoS



Proof of Data Possession  
PDP



Proof of Retrievability  
PoR

Memory usage

File size  $n$  then  
memory usage is  $n$

**Even if file is compressible**

The data/file is stored

Data/file can be retrieved

# PoRep Definition

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow \text{ek, pk, vk}$  public keys
- $\text{Encode}(m, \text{ek}, \text{id}) \rightarrow \text{encoding } c \text{ and digest } h$



In some settings, messages might be highly compressible, i.e.  $m=(F_k(1), F_k(2), \dots)$ . There, a prover could avoid storing  $m$  and generate as needed when challenged.

Publish  $h$

---

- $\text{Prove}(\text{pk}, \text{challenge}, c) \rightarrow \pi$
- $\text{Verify}(\text{vk}, h, \pi) \rightarrow 0/1$
- $\text{Decode}(\text{ek}, \text{id}, c) \rightarrow m$

Slow Encode()




Run time of computing proof <  
run time of Encode()

**Properties:** encoding correctness, proof completeness, replication and extraction


# Proof of Replication (PoRep) – [Fisch'18]

preprocessing

Setup  
Phase




Verifier



Prover


# Proof of Replication (PoRep) – [Fisch'18]

Setup  
Phase



Verifier

data D




Prover

preprocessing

# Proof of Replication (PoRep) – [Fisch'18]

Setup  
Phase



Verifier


data D



Prover

preprocessing


$c_i$  : data encodings



Depth Robust Graphs (DRGs)

# Proof of Replication (PoRep) – [Fisch'18]

Setup Phase



Verifier

data D


commit to the DRG (MT on last layer)



Prover

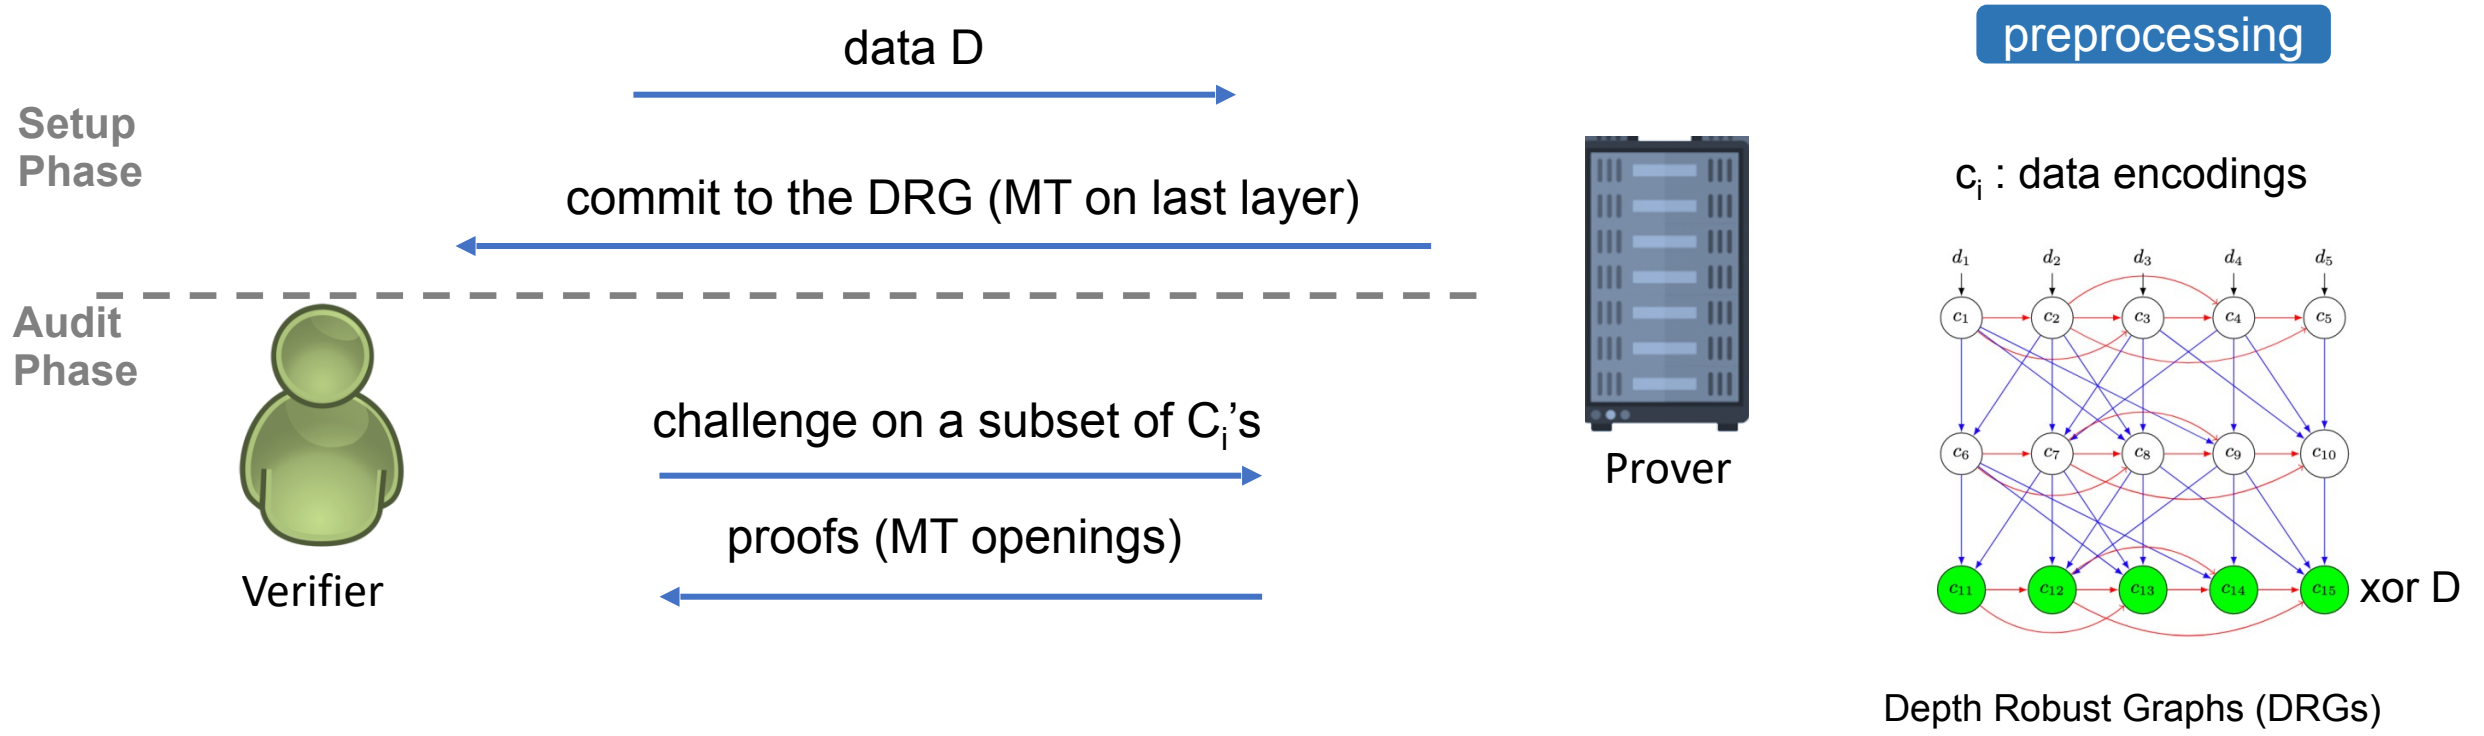
preprocessing

$c_i$  : data encodings

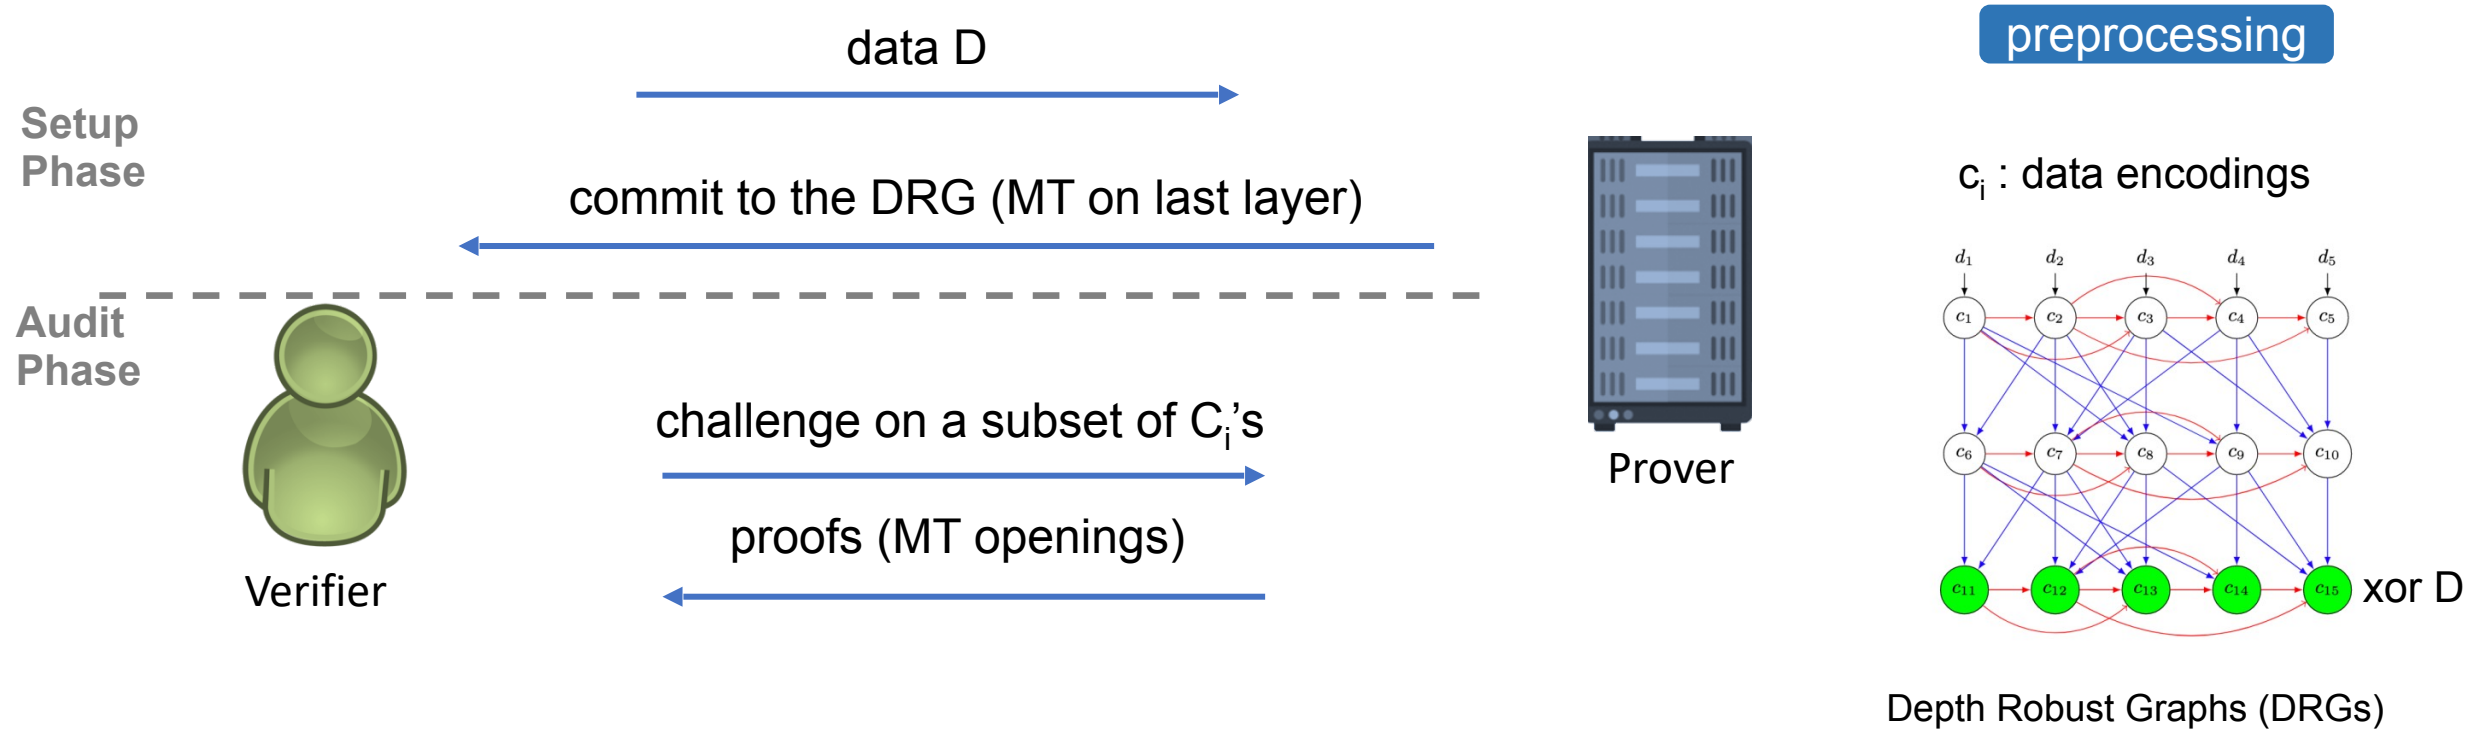


Depth Robust Graphs (DRGs)

# Proof of Replication (PoRep) – [Fisch'18]



# Proof of Replication (PoRep) – [Fisch'18]



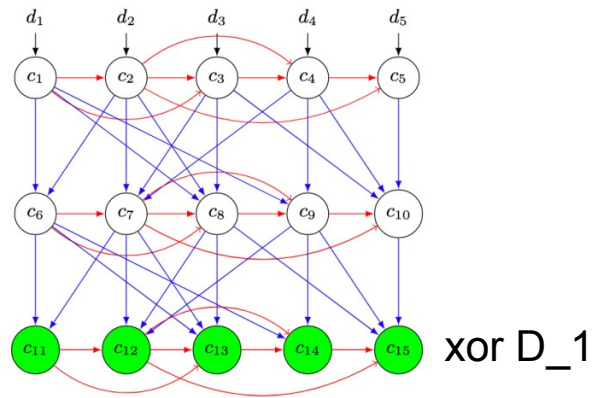
The auditing phase is probabilistic – a large set of challenges is needed to get a good level of security.

File of  $n$  blocks -  $\rightarrow$  Memory usage is  $n(1-\epsilon)$  where  $\epsilon$  is constant



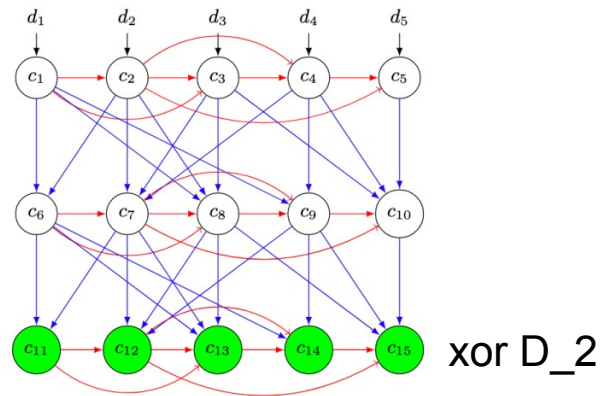
# Proof of Replication (PoRep) – [Fisch'18]

$c_i$  : data encodings




Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

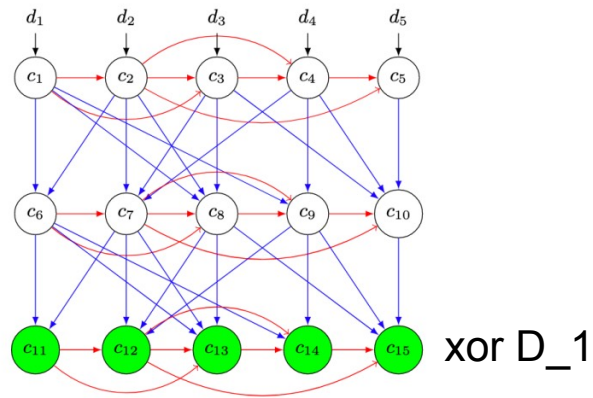
$c_i$  : data encodings



Depth Robust Graphs (DRGs)


# Proof of Replication (PoRep) – [Fisch'18]

$c_i$  : data encodings




Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

$c_i$  : data encodings




Depth Robust Graphs (DRGs)

u file of  $n$  blocks -  $\rightarrow$  Memory usage is  $u \cdot n(1-\epsilon)$  where  $\epsilon$  is constant.


# Proof of Replication (PoRep) – [Fisch'18]

$c_i$  : data encodings




Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

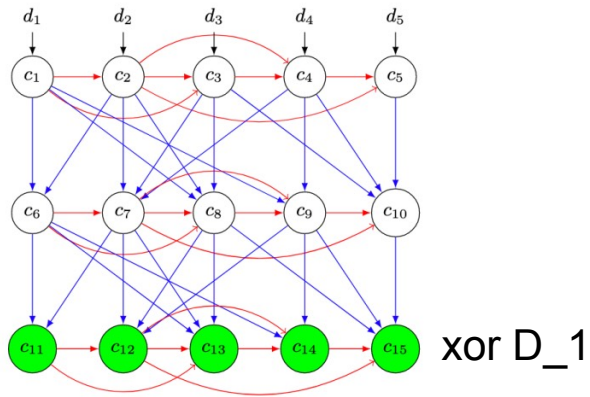
u file of n blocks -  $\rightarrow$  Memory usage is  $u \cdot n \cdot (1 - \epsilon)$  where  $\epsilon$  is constant.

Loss scales with u (the number of files stored), i.e., loss is  $u \cdot n \cdot \epsilon$

The problem is the probabilistic check on the last layer (the encoding)


# Proof of Replication (PoRep) – [Fisch'18]

$c_i$  : data encodings




Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

$c_i$  : data encodings



Depth Robust Graphs (DRGs)

u file of n blocks -  $\rightarrow$  Memory usage is  $u \cdot n \cdot (1 - \epsilon)$  where  $\epsilon$  is constant.

Loss scales with u (the number of files stored), i.e., loss is  $u \cdot n \cdot \epsilon$

The problem is the probabilistic check on the last layer (the encoding)



Remove the loss proportional to u  $\rightarrow$  Scales the challenges proportional to u



# Our work

•  $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$   $\text{ek, pk, vk}$  public keys

preprocessing

•  $\text{Encode}(m, \text{ek}, \text{id}) \rightarrow$  encoding  $c$  and digest  $h$

Publish  $h$

---

•  $\text{Prove}(\text{pk}, \text{challenge}, c) \rightarrow \pi$

audit phase

•  $\text{Verify}(\text{vk}, h, \pi) \rightarrow 0/1$

•  $\text{Decode}(\text{ek}, c) \rightarrow m$

**Properties:** encoding correctness, proof completeness, replication and extraction

# Our work

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow$  **public keys**  $ek, pk, vk$

In this work we assume **the encoding is honestly executed** (i.e., digest  $h$  is honest).  
Worst case: Use a SNARK to verify  $h$

preprocessing

- $\text{Encode}(m, ek, id) \rightarrow$  encoding  $c$  and digest  $h$

**Publish  $h$**

---

- $\text{Prove}(pk, \text{challenge}, c) \rightarrow \pi$
- $\text{Verify}(vk, h, \pi) \rightarrow 0/1$
- $\text{Decode}(ek, c) \rightarrow m$

audit phase

**Properties:** encoding correctness, proof completeness, replication and extraction

# Our work

- $\text{Setup}(1^\lambda, 1^t, 1^n) \rightarrow \text{ek, pk, vk}$  public keys

- $\text{Encode}(m, \text{ek}, \text{id}) \rightarrow \text{encoding } c \text{ and digest } h$

Publish  $h$

- $\text{Prove}(\text{pk}, \text{challenge}, c) \rightarrow \pi$

- $\text{Verify}(\text{vk}, h, \pi) \rightarrow 0/1$

- $\text{Decode}(\text{ek}, c) \rightarrow m$

In this work we assume **the encoding is honestly executed** (i.e., digest  $h$  is honest).

Worst case: Use a SNARK to verify  $h$

preprocessing

Our focus is to **reduce complexity** and **increase security** of this phase

audit phase

**Properties:** encoding correctness, proof completeness, replication and extraction

**Our construction**



# Polynomial Evaluation


degree  $d$  polynomial  $f(X) \leftarrow \mathbb{Z}_p[X]$

$$f(X) = a_0 + a_1 \cdot x + a_2 \cdot x^2 + \dots + a_d \cdot x^d$$

# Polynomial Evaluation

degree  $d$  polynomial  $f(X) \leftarrow_{\$} \mathbb{Z}_p[X]$

$$f(X) = a_0 + a_1 \cdot x + a_2 \cdot x^2 + \dots + a_d \cdot x^d$$



random  $x$




$f(x)$



# Polynomial Evaluation

degree  $d$  polynomial  $f(X) \leftarrow_{\$} \mathbb{Z}_p[X]$

$$f(X) = a_0 + a_1 \cdot x + a_2 \cdot x^2 + \dots + a_d \cdot x^d$$



random  $x$



$f(x)$



space  
requirements?

# Incompressibility of Random Polynomials

randomly sampled  
polynomial of degree  $d$

$$f(X) \leftarrow \$ \mathbb{Z}_p[X]$$

The goal: evaluating  $f(X)$  should require memory close to  $|f(X)|$

# Incompressibility of Random Polynomials

randomly sampled  
polynomial of degree  $d$

$$f(X) \leftarrow \$ \mathbb{Z}_p[X]$$

The goal: evaluating  $f(X)$  should require memory close to  $|f(X)|$



- pre-processes  $f(X)$  to compute a memory  $\alpha$  smaller than  $|f(X)|$
- shouldn't be possible to evaluate  $f(x)$  on a random point  $x$ , unless  $|\alpha| \approx |f(X)|$

# Incompressibility of Random Polynomials

randomly sampled  
polynomial of degree  $d$

$$f(X) \leftarrow \mathbb{Z}_p[X]$$

The goal: evaluating  $f(X)$  should require memory close to  $|f(X)|$



- pre-processes  $f(X)$  to compute a memory  $\alpha$  smaller than  $|f(X)|$
- shouldn't be possible to evaluate  $f(x)$  on a random point  $x$ , unless  $|\alpha| \approx |f(X)|$


$$[\text{ACFPT}'23]: |\alpha| \approx |f(X)| - d^*|x|$$

# Multiple polynomials?

$$f_1(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$

...

$$f_u(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$



random x



$f_1(x) \dots f_u(x)$

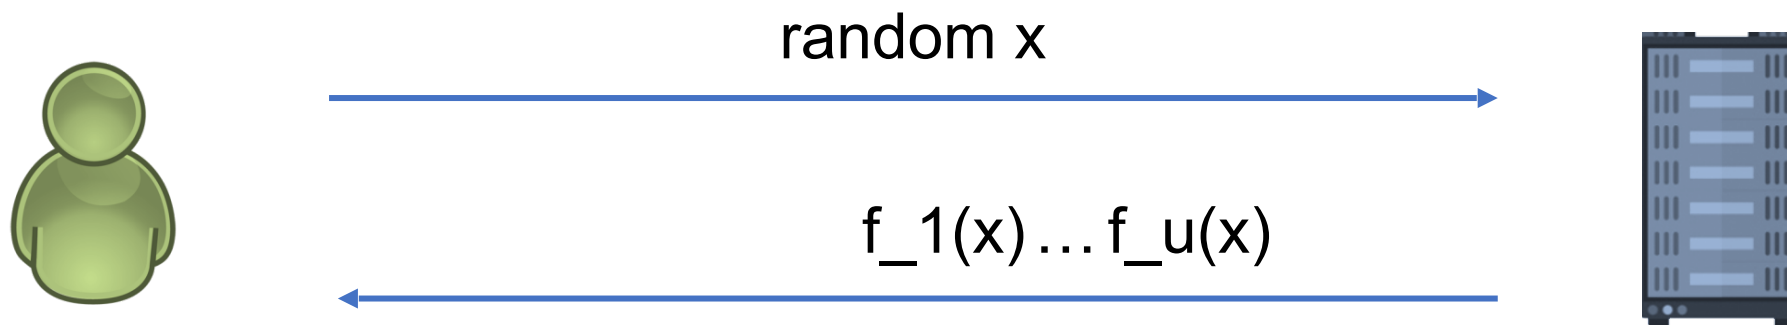


# Multiple polynomials?

$$f_1(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$

...

$$f_u(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$



single polynomial

$$[\text{ACFPT}'23]: |\alpha| \approx |f(X)| - d^*|x|$$




# Multiple polynomials?

$$f_1(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$

...

$$f_u(X) \leftarrow_{\$} \mathbb{Z}_p[X]$$



random  $x$



$f_1(x) \dots f_u(x)$



single polynomial


$$[\text{ACFPT}'23]: |\alpha| \approx |f(X)| - d^*|x|$$

multiple polynomials


$$|\alpha| \approx u^*|f(X)| - d^*|x|$$

# Partial construction

Setup  
Phase





Verifier




Prover

preprocessing





# Partial construction

Setup Phase





Verifier



Prover

preprocessing




# Partial construction

Setup Phase

data D

Commit on ???

Audit Phase



Verifier


x

$f_1(x), \dots, f_u(x)$




Prover

preprocessing

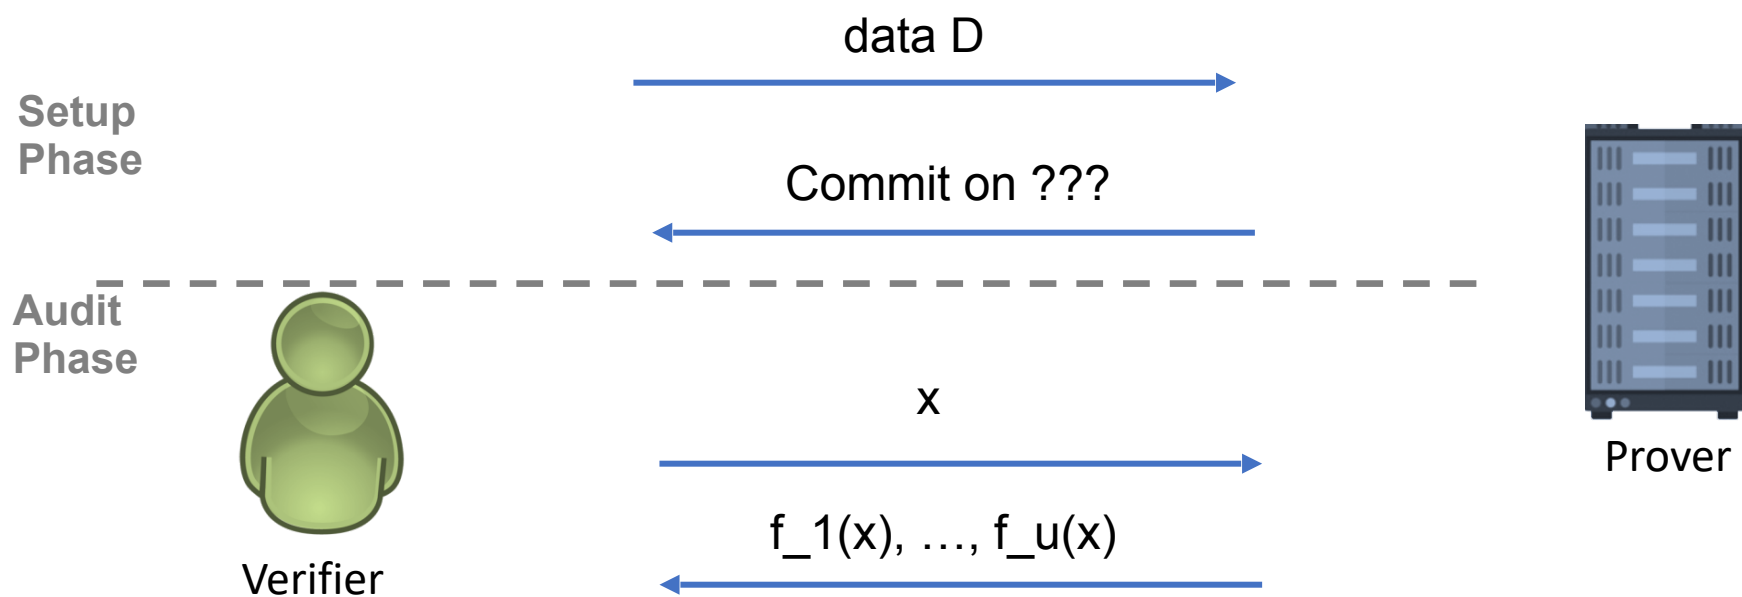


$\text{xor } D_1 = f_1(X)$

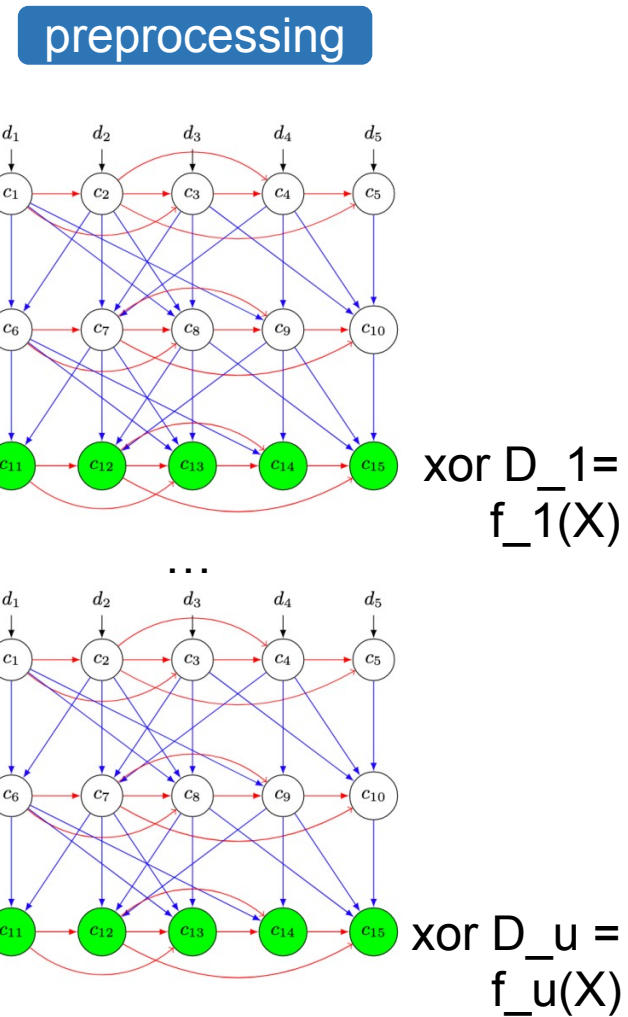


$\text{xor } D_u = f_u(X)$

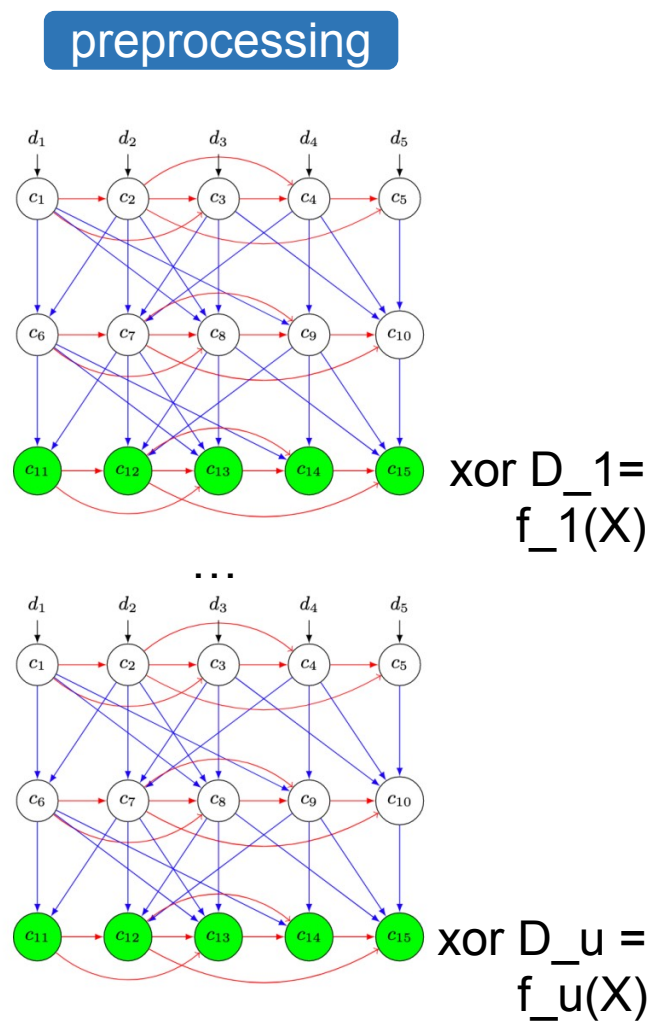
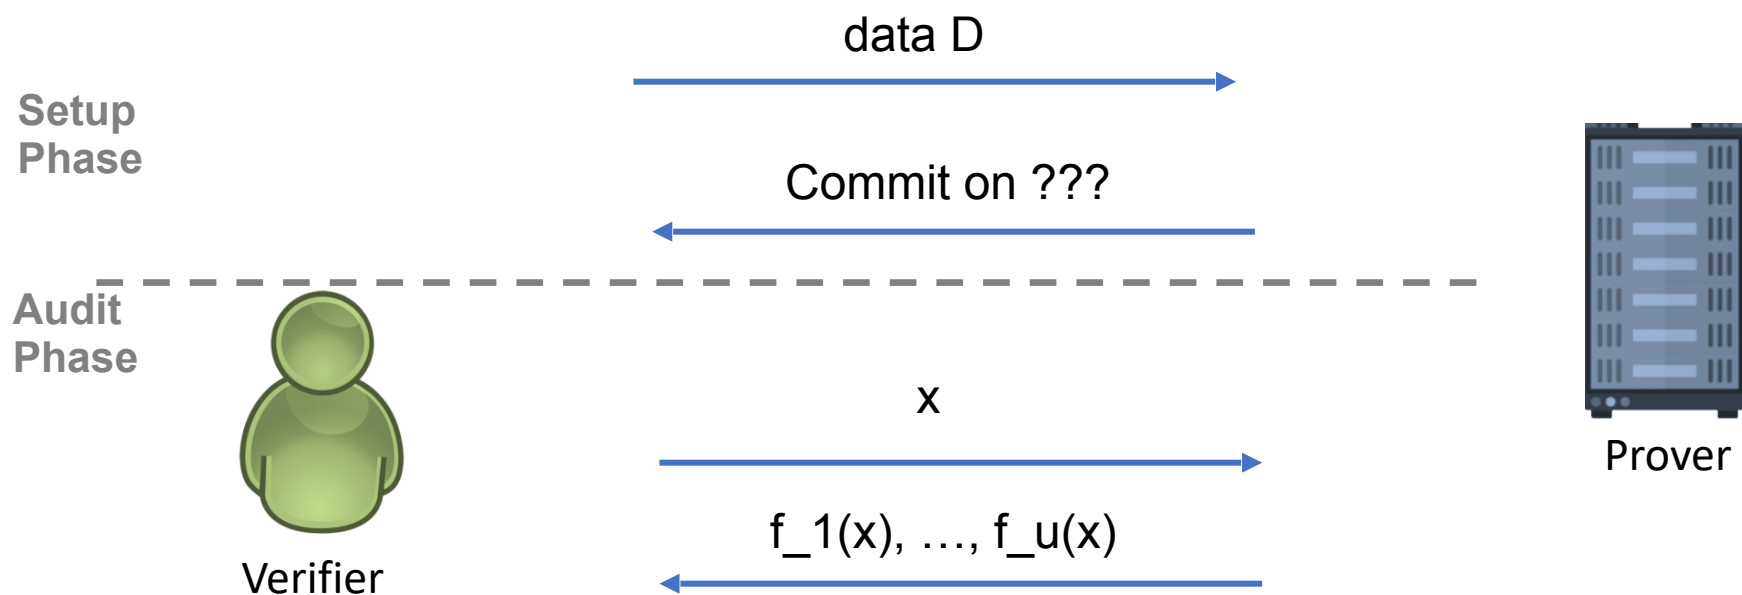
# Partial construction



Memory usage:  $u \cdot |f(x)| - d \cdot |x| \approx u \cdot |f(x)|$  for large values of  $u$



# Partial construction



Memory usage:  $u \cdot |f(x)| - d \cdot |x| \approx u \cdot |f(x)|$  for large values of  $u$





Prover needs to read  $f_1(X), \dots, f_u(X)$  entirely - Problem for efficiency



How verify that  $f_1(x), \dots, f_u(x)$  are correct evaluations (efficiently)?

# Poly-logarithmic prover's running time

Kedlaya and Uman 2011: "Fast polynomial factorization and modular composition"



Prover


# Poly-logarithmic prover's running time

Kedlaya and Uman 2011: "Fast polynomial factorization and modular composition"





uses a RAM data structure to expedite polynomial evaluations

(computes  $f(x)$  in poly-log time in  $d$ )



...



Prover




# Poly-logarithmic prover's running time

Kedlaya and Uman 2011: "Fast polynomial factorization and modular composition"




uses a RAM data structure to expedite polynomial evaluations

(computes  $f(x)$  in poly-log time in  $d$ )



xor  $D_1 = f_1(X) \rightarrow$  Compute data structure  $D_1$

...



xor  $D_u = f_u(X) \rightarrow$  Compute data structure  $D_u$



Prover


# Poly-logarithmic prover's running time

Kedlaya and Uman 2011: "Fast polynomial factorization and modular composition"




uses a RAM data structure to expedite polynomial evaluations

(computes  $f(x)$  in poly-log time in  $d$ )




xor  $D_1 = f_1(X) \rightarrow$  Compute data structure  $D_1$

...



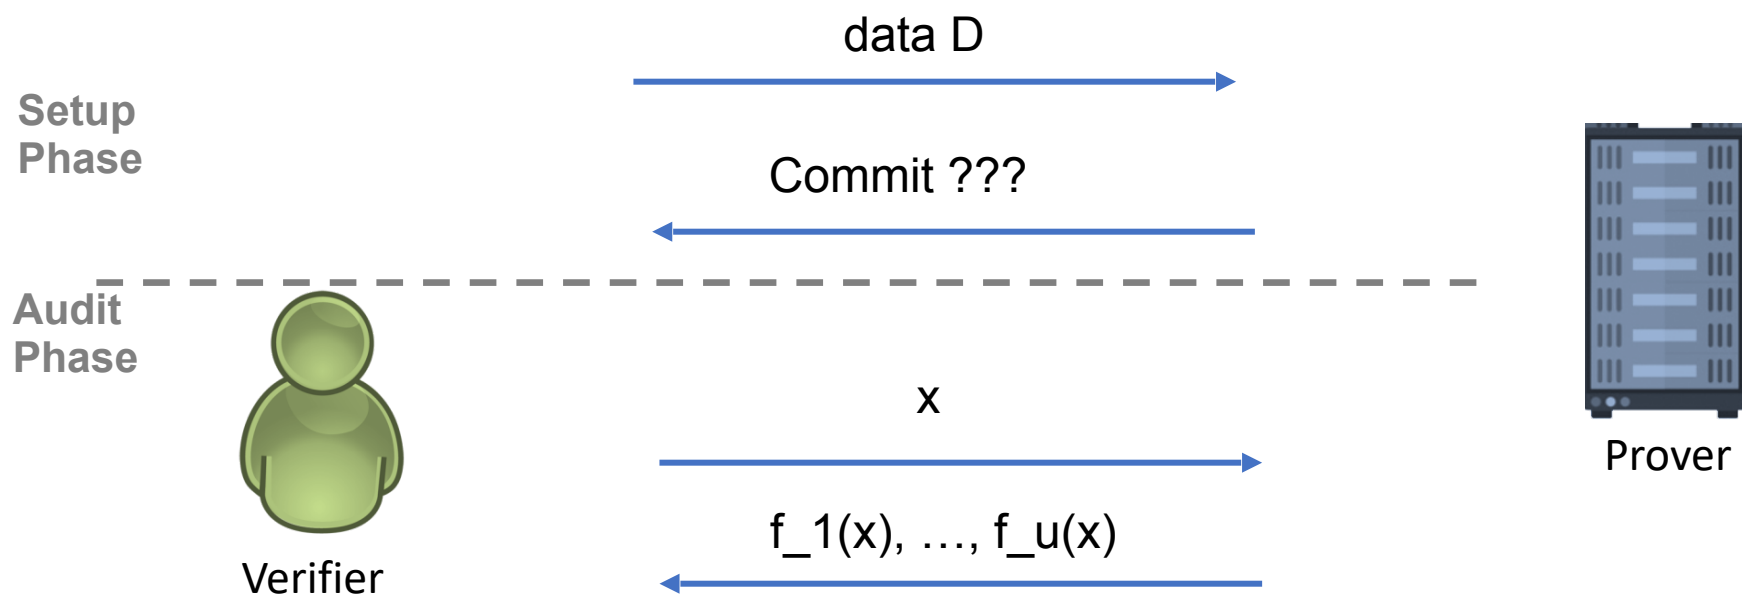
xor  $D_u = f_u(X) \rightarrow$  Compute data structure  $D_u$

**$f_i(X)$  can be evaluated in time  $\text{polylog}(d)$**



Prover

# Partial construction




Memory usage:  $u \cdot |f(x)| - d \cdot |x| \approx u \cdot |f(x)|$  for large values of  $u$




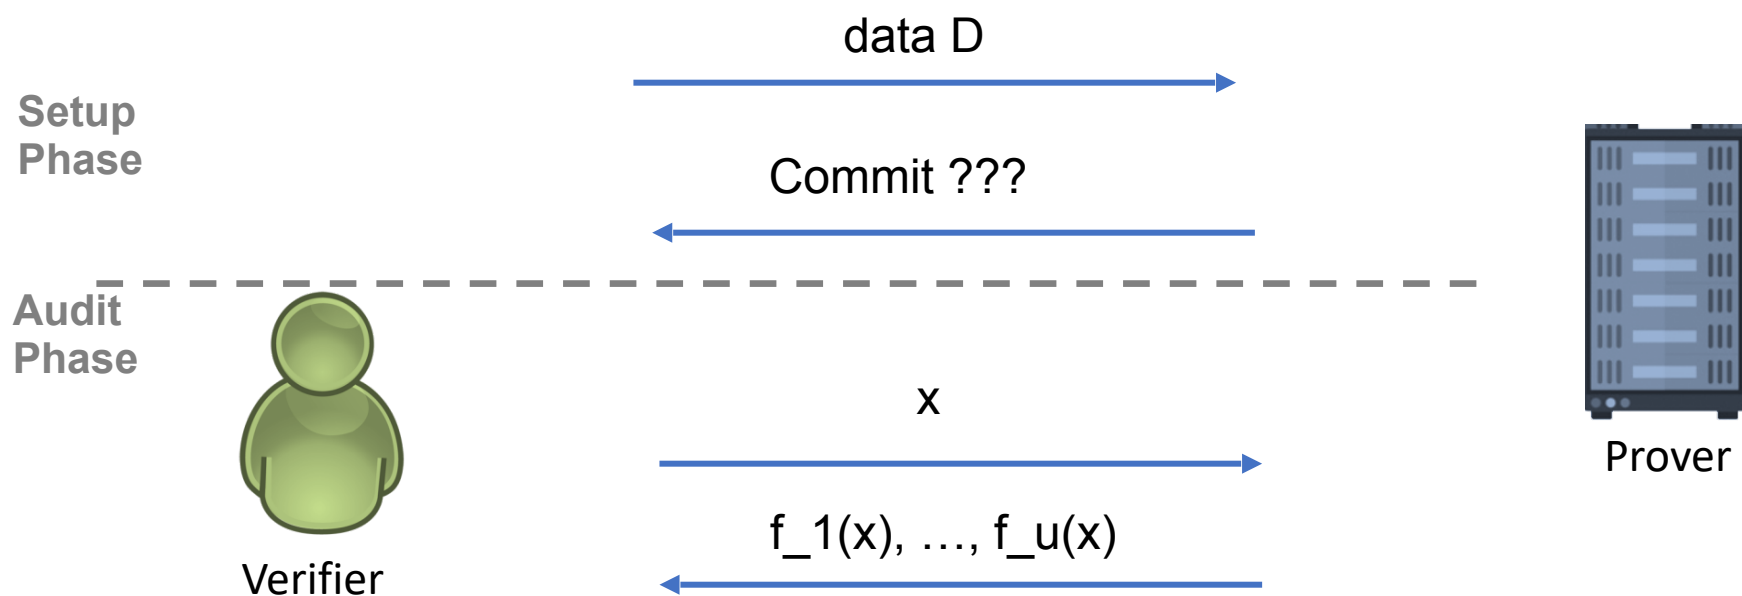
Prover needs to read  $f_1(X), \dots, f_u(X)$  entirely - Problem for efficiency



How verify that  $f_1(x), \dots, f_u(x)$  are correct evaluations?



# Partial construction



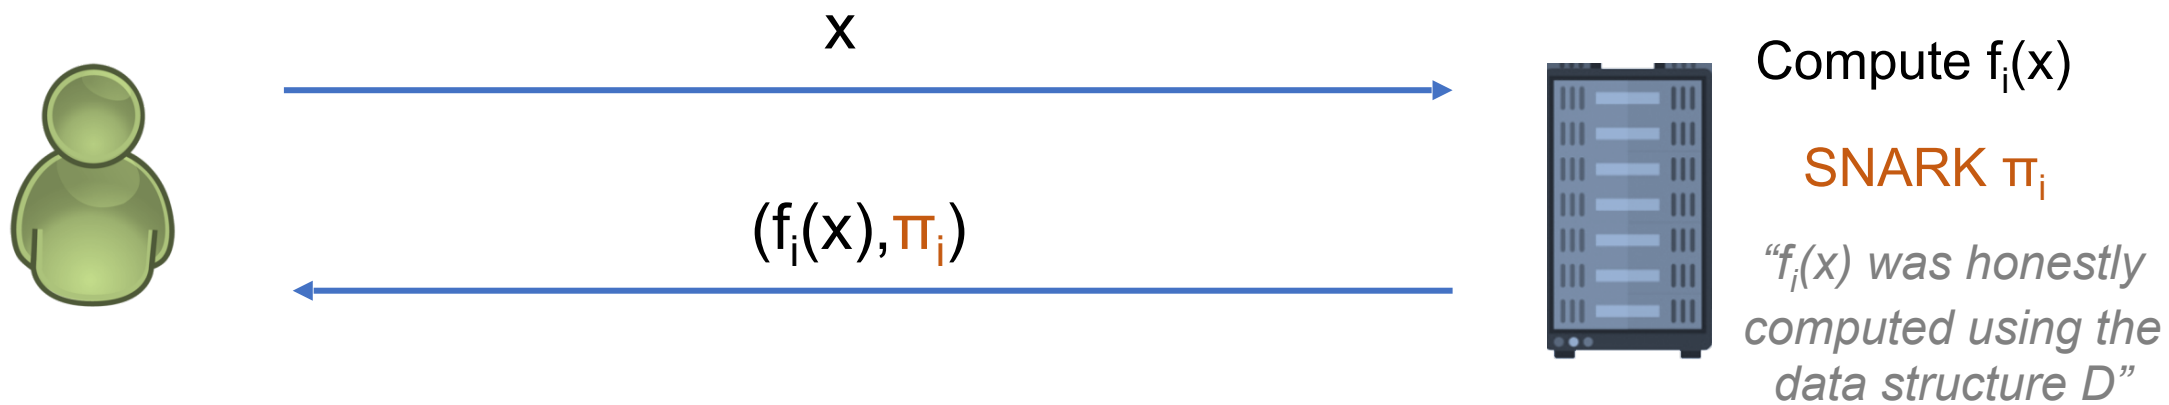
Memory usage:  $u \cdot |f(x)| - d \cdot |x| \approx u \cdot |f(x)|$  for large values of  $u$

~~⚠ Prover needs to read  $f_1(X), \dots, f_u(X)$  entirely. Problem for efficiency.~~

⚠ How verify that  $f_1(x), \dots, f_u(x)$  are correct evaluations?


# Poly-log verification

Let's use SNARKs!



# Poly-log verification

~~Let's use SNARKs!~~



Compute  $f_i(x)$

SNARK  $\pi_i$

*" $f_i(x)$  was honestly  
computed using the  
data structure  $D$ "*



prover time becomes  
linear to  $D$

# Poly-log verification

What if the verifier also computes  $f(x)$ ?

# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✘ no access to D!



# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✘ no access to  $D$ !



Prover only needs to access a poly-log number of blocks  $D' \subset D$

# Poly-log verification


What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$

preprocessing



# Poly-log verification

What if the verifier also computes  $f(x)$ ?


✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$

preprocessing

D




# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$



preprocessing

D

Merkle  
Tree

h


# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$



Audit phase

$x$



preprocessing

$D$

Merkle  
Tree

$h$




# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$



Audit phase

$x$



preprocessing

$D$



Merkle  
Tree

$h$

Compute  $f(x)$   
from  $D$ .

Let  $D'$  the blocks  
read.


# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$




Audit phase

$x$

$f(x),$

blocks  $D'$  + corresponding set of MT openings



preprocessing

$D$

Merkle  
Tree

$h$

Compute  $f(x)$   
from  $D$ .

Let  $D'$  the blocks  
read.


# Poly-log verification

What if the verifier also computes  $f(x)$ ?

✗ no access to D!



Prover only needs to access a poly-log number of blocks  $D' \subset D$



Audit phase

$x$

$f(x),$

blocks  $D'$  + corresponding set of MT openings



preprocessing

$D$

Merkle  
Tree

$h$


Compute  $f(x)$   
from  $D$ .  
Let  $D'$  the blocks  
read.

- check proofs
- compute  $y$  from  $D'$
- check  $y \stackrel{?}{=} f(x)$


**poly-log verification**



# Final Construction (Example: 1 File)



Verifier




Prover


# Final Construction (Example: 1 File)

Setup  
Phase

data D



Verifier




Prover


# Final Construction (Example: 1 File)

Setup Phase

data D




Verifier



Prover


preprocessing




# Final Construction (Example: 1 File)

Setup Phase

data D




Verifier



Prover

preprocessing




$$\text{xor } D = f(X)$$


# Final Construction (Example: 1 File)

Setup Phase

data D




Verifier



Prover

preprocessing




xor D =  
 $f(X)$

Compute from  $f(X)$   
data structure D


# Final Construction (Example: 1 File)

Setup Phase

data D




Verifier



Prover

preprocessing



xor D = f(X)

Compute from f(X)  
data structure D


Merkle-tree on D's  
blocks

# Final Construction (Example: 1 File)


Setup Phase

data D

Publish merkle-tree root h (digest)




Verifier



Prover

preprocessing




Compute from  $f(X)$   
data structure D

Merkle-tree on D's  
blocks


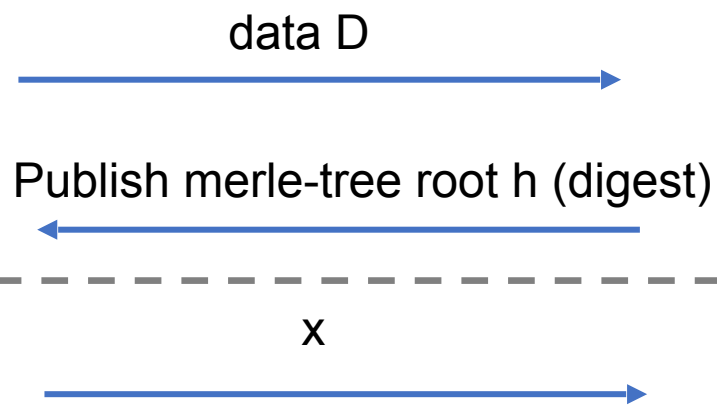
# Final Construction (Example: 1 File)

Setup Phase

Audit Phase




Verifier



Prover

preprocessing



Compute from  $f(X)$   
data structure D

Merkle-tree on D's  
blocks



# Final Construction (Example: 1 File)

Setup Phase

data D



Publish merle-tree root h (digest)




Audit Phase

x



$\pi = (f(x), \text{blocks } D', \text{ merle-tree openings})$




Verifier



Prover

preprocessing



xor D = f(X)

Compute from f(X)  
data structure D

Merkle-tree on D's  
blocks

# Final Construction (Example: 1 File)

Setup Phase

data D



Publish merle-tree root h (digest)




Audit Phase


x



$\pi = (f(x), \text{blocks } D', \text{ merle-tree openings})$




Verifier



Prover

- check proofs
- compute  $y$  from  $D'$
- check  $y \stackrel{?}{=} f(x)$

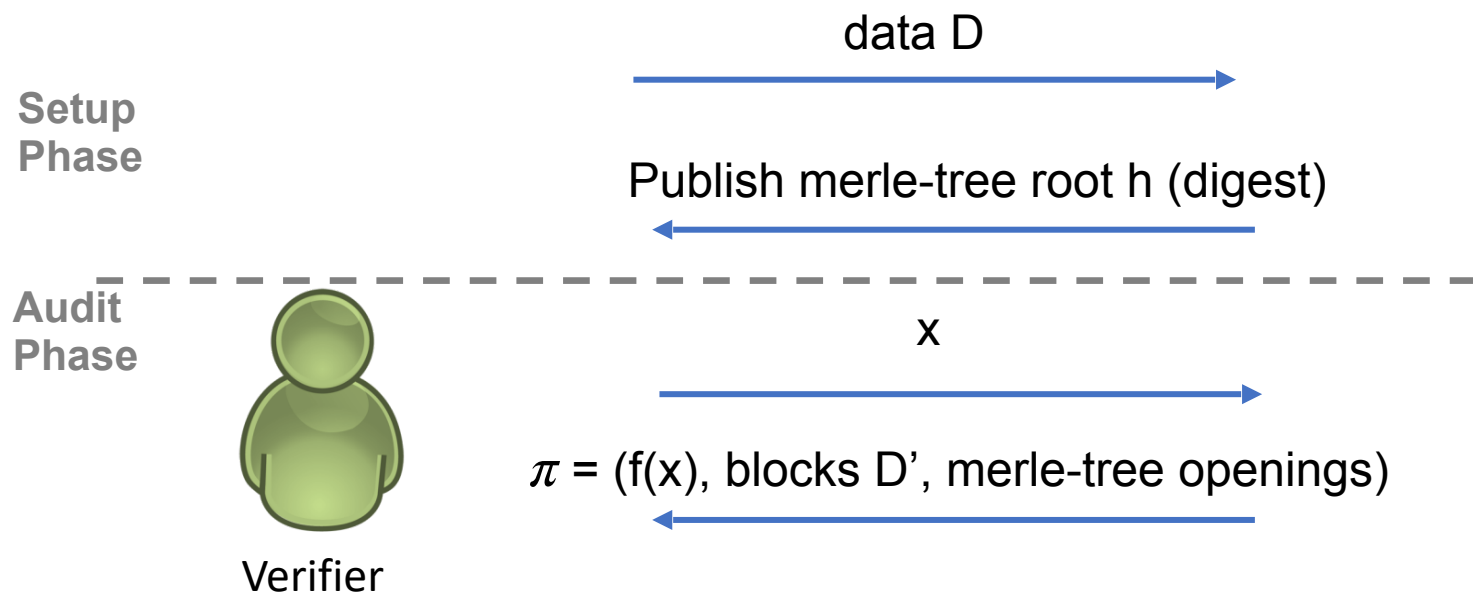
preprocessing



Compute from  $f(X)$   
data structure D

Merkle-tree on  $D$ 's  
blocks

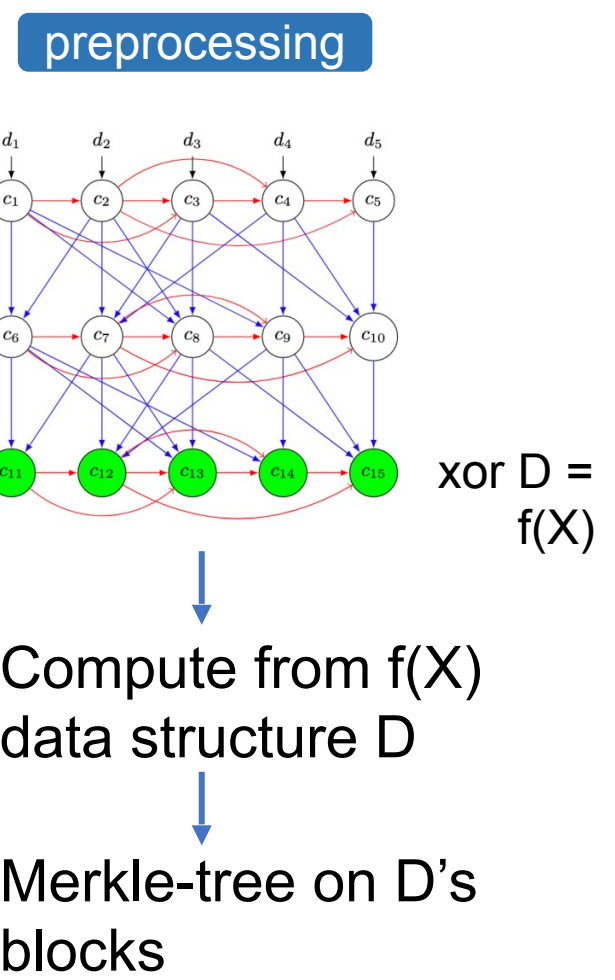
# Final Construction (Example: 1 File)



- check proofs
- compute  $y$  from  $D'$
- check  $y \stackrel{?}{=} f(x)$

✓ Memory usage  $\approx u \cdot |f(X)|$  for large number of files

✓ Efficient proving and verification time



# Absolute memory usage

---

[Fisch'18]

---

Ours

# Absolute memory usage


[Fisch'18]

Files sizes  $u \cdot n$

Ours

Files sizes  $u \cdot n$

# Absolute memory usage

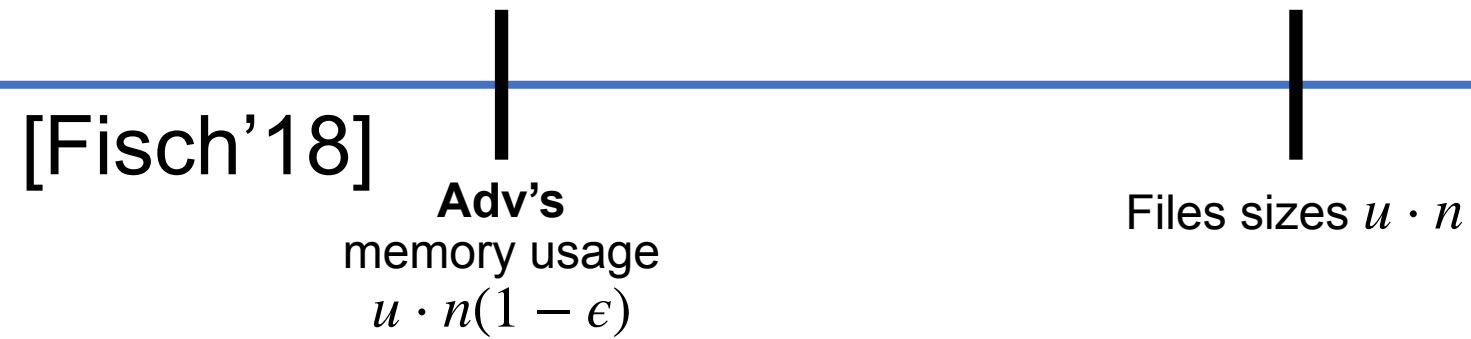


# Absolute memory usage

[Fisch'18]

Adv's memory usage  
 $u \cdot n(1 - \epsilon)$

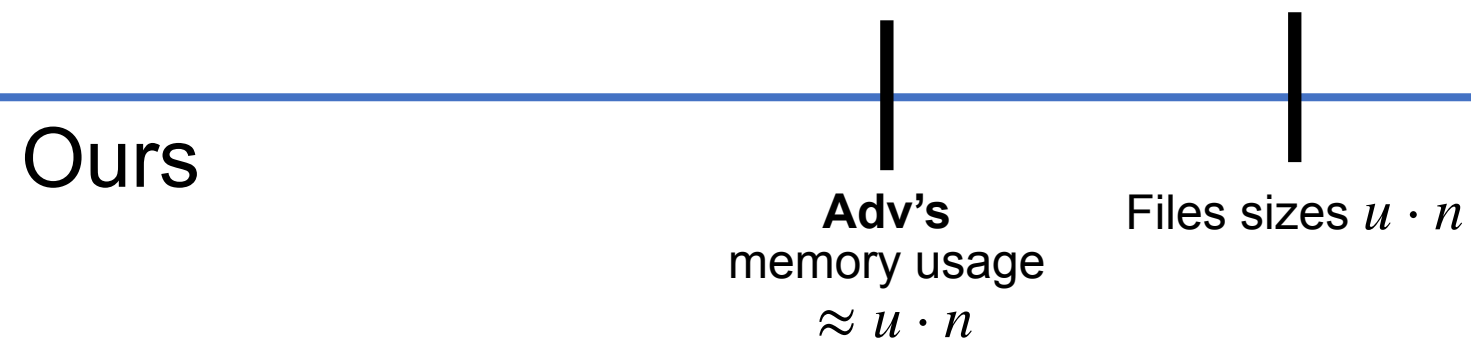
Files sizes  $u \cdot n$

A horizontal blue line represents a baseline. Two vertical black bars are positioned above this line. The left bar is labeled 'Adv's memory usage' with the formula  $u \cdot n(1 - \epsilon)$  below it. The right bar is labeled 'Files sizes' with the formula  $u \cdot n$  below it. The right bar is taller than the left bar.

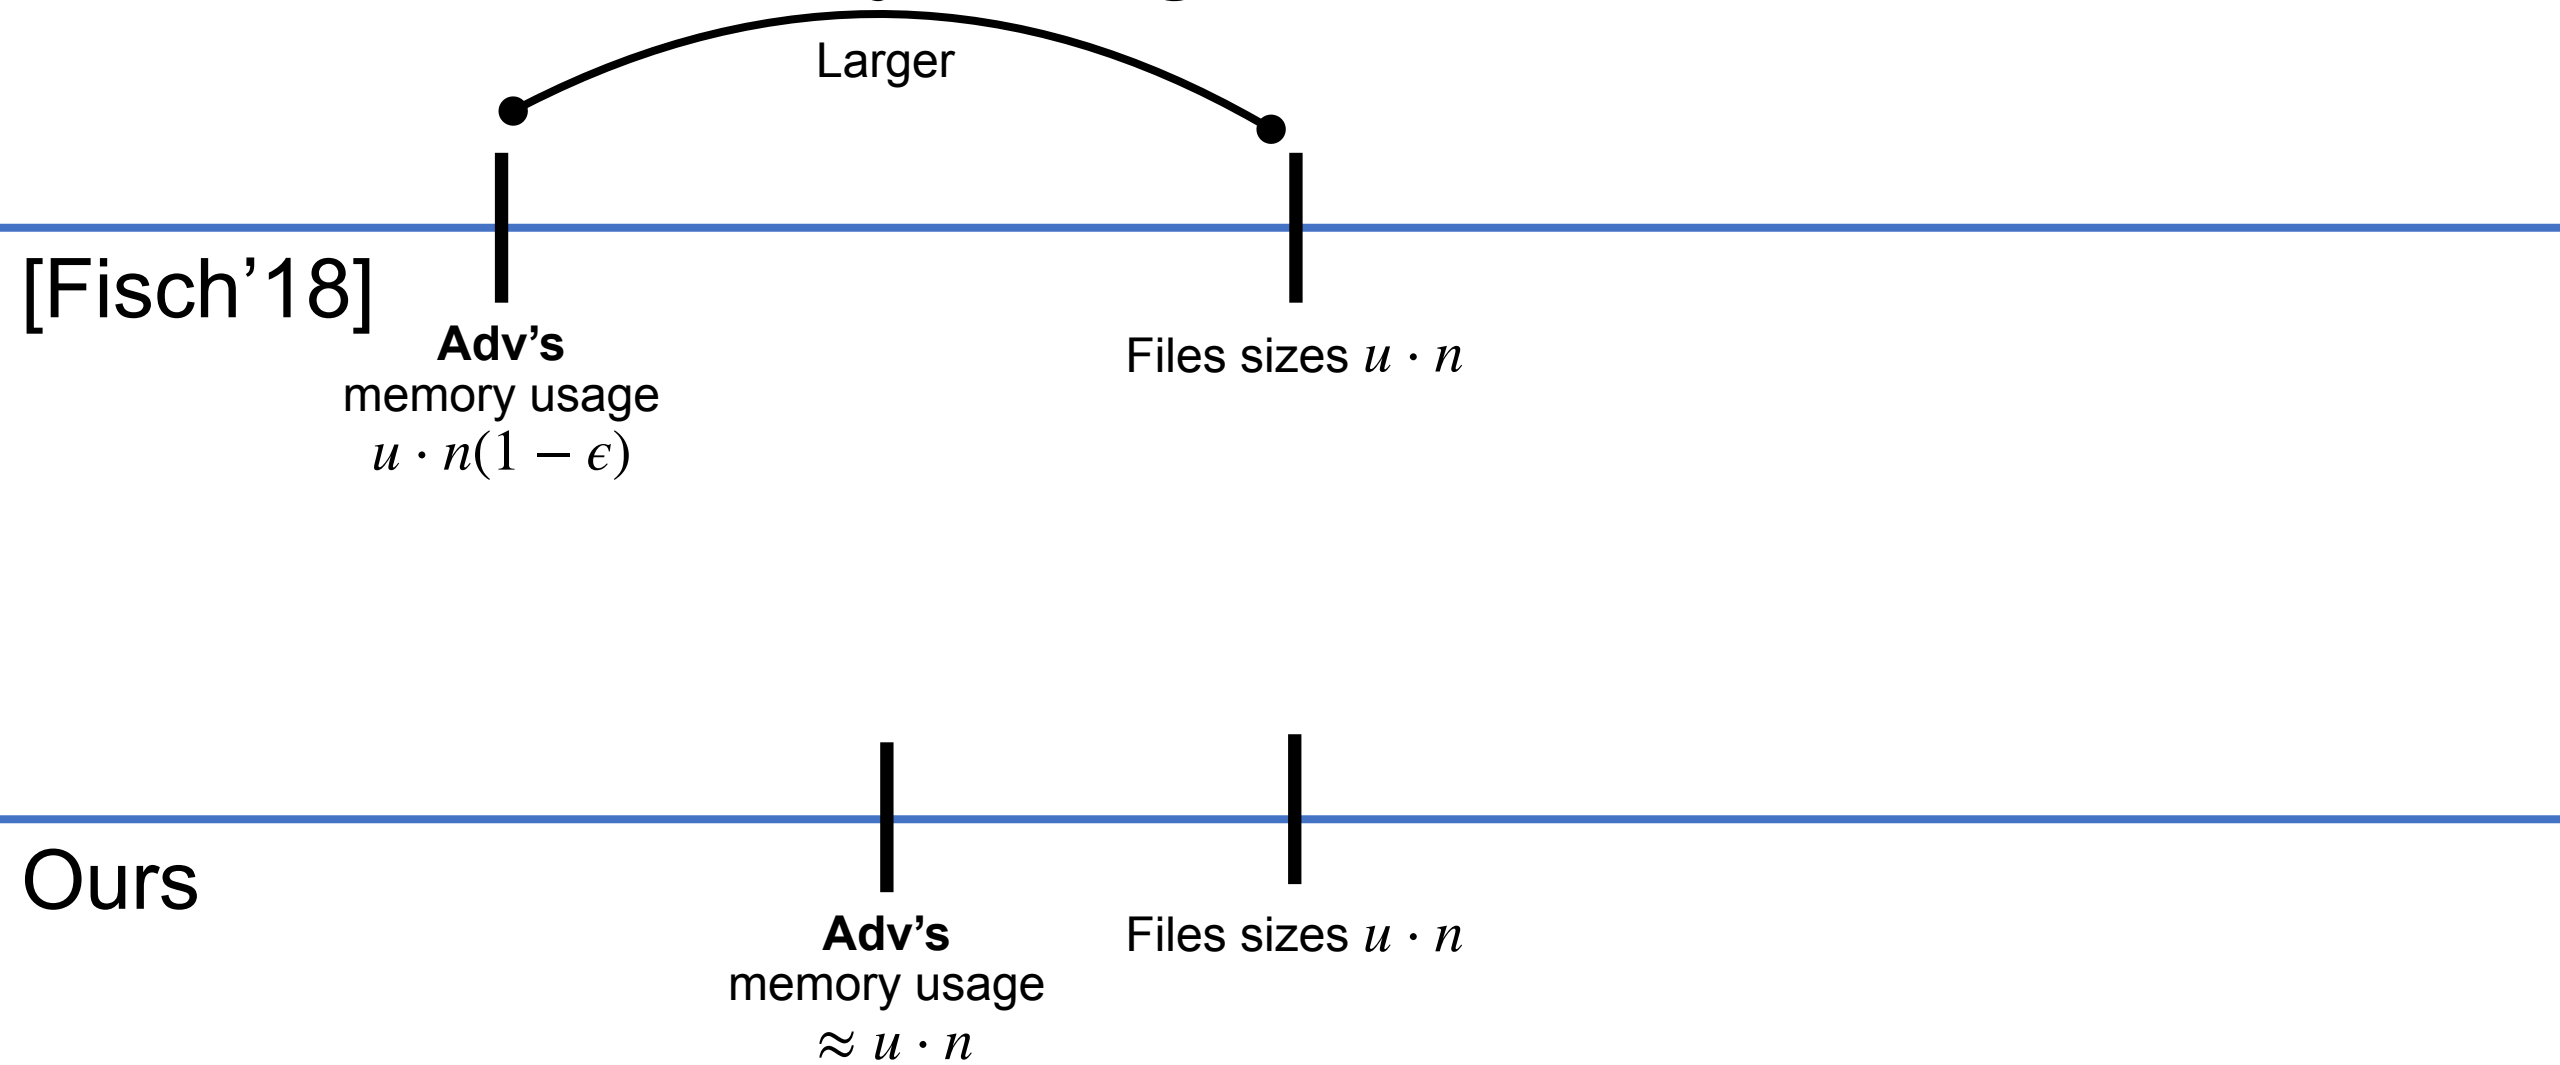
Ours

Adv's memory usage  
 $\approx u \cdot n$

Files sizes  $u \cdot n$



A horizontal blue line represents a baseline. Two vertical black bars are positioned above this line. The left bar is labeled 'Adv's memory usage' with the formula  $\approx u \cdot n$  below it. The right bar is labeled 'Files sizes' with the formula  $u \cdot n$  below it. The two bars are approximately the same height.

# Absolute memory usage





# Absolute memory usage



# Limitation = Space gap

[Fisch'18]

Adv's memory usage  
 $u \cdot n(1 - \epsilon)$

Files sizes  $u \cdot n$

The diagram for [Fisch'18] features a horizontal blue line. To the left of the line, the text "[Fisch'18]" is written. To the right of the line, there are two vertical black bars. The left bar is significantly taller than the right bar. Below the left bar, the text "Adv's memory usage" and the formula  $u \cdot n(1 - \epsilon)$  are displayed. Below the right bar, the text "Files sizes  $u \cdot n$ " is displayed. The large vertical distance between the two bars indicates a significant space gap.

Ours

Adv's memory usage  
 $\approx u \cdot n$

Files sizes  $u \cdot n$

The diagram for "Ours" features a horizontal blue line. To the left of the line, the text "Ours" is written. To the right of the line, there are two vertical black bars of equal height. Below the left bar, the text "Adv's memory usage" and the formula  $\approx u \cdot n$  are displayed. Below the right bar, the text "Files sizes  $u \cdot n$ " is displayed. The equal height of the two bars indicates a minimal space gap.

# Limitation = Space gap

[Fisch'18]

Adv's  
memory usage  
 $u \cdot n(1 - \epsilon)$

Files sizes  $u \cdot n$   
**Honest**  
memory usage  
 $u \cdot n$

The diagram for [Fisch'18] features two horizontal blue lines. The top line is labeled [Fisch'18]. Below it, a vertical black line is positioned to the left of the center, with the text 'Adv's memory usage' and the formula  $u \cdot n(1 - \epsilon)$  below it. To the right of the center, another vertical black line is positioned, with the text 'Files sizes  $u \cdot n$ ' and '**Honest** memory usage  $u \cdot n$ ' below it. The vertical distance between the top blue line and the bottom blue line is significantly larger than the vertical distance between the two vertical lines, indicating a large space gap.

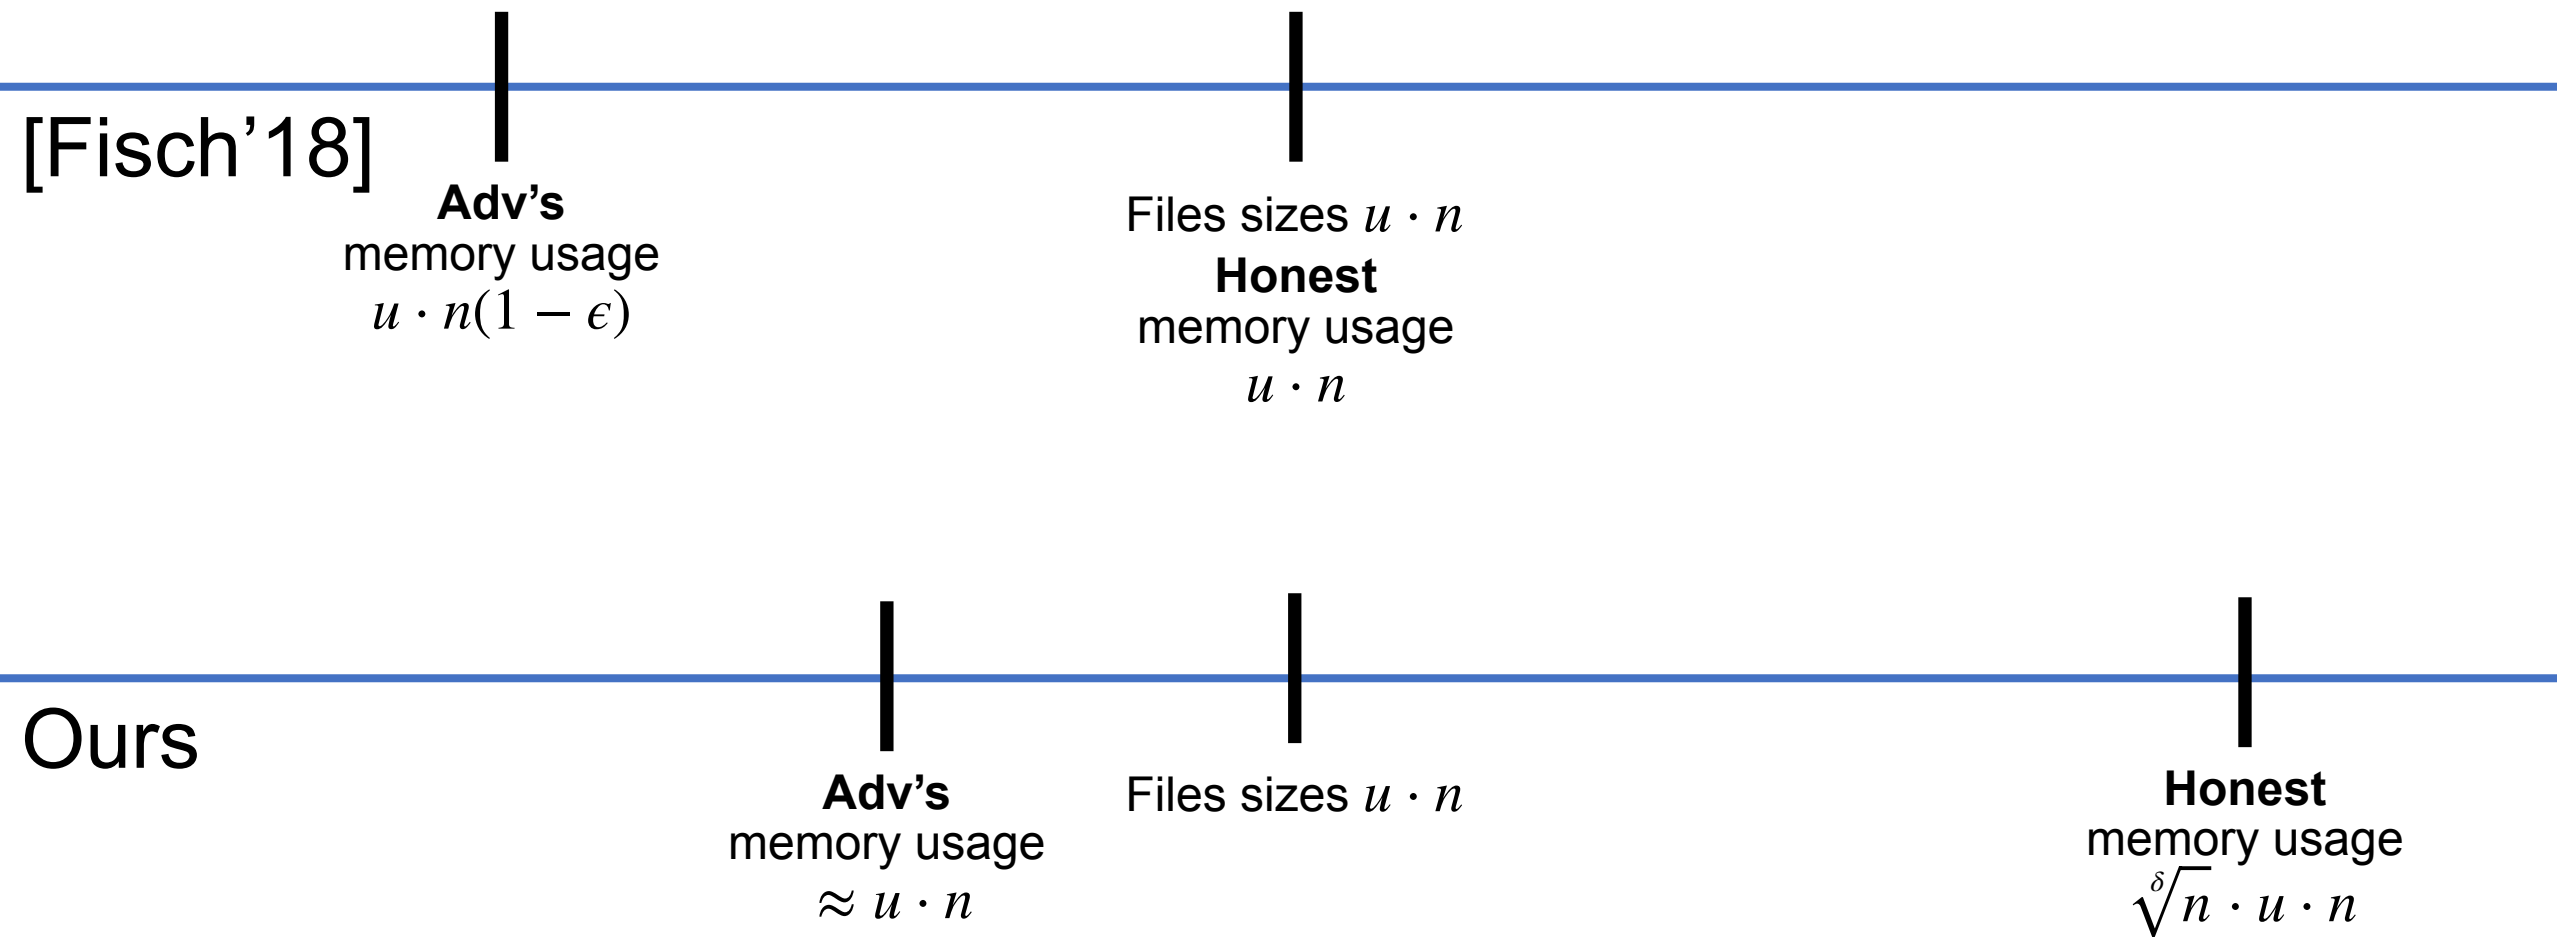
Ours

Adv's  
memory usage  
 $\approx u \cdot n$

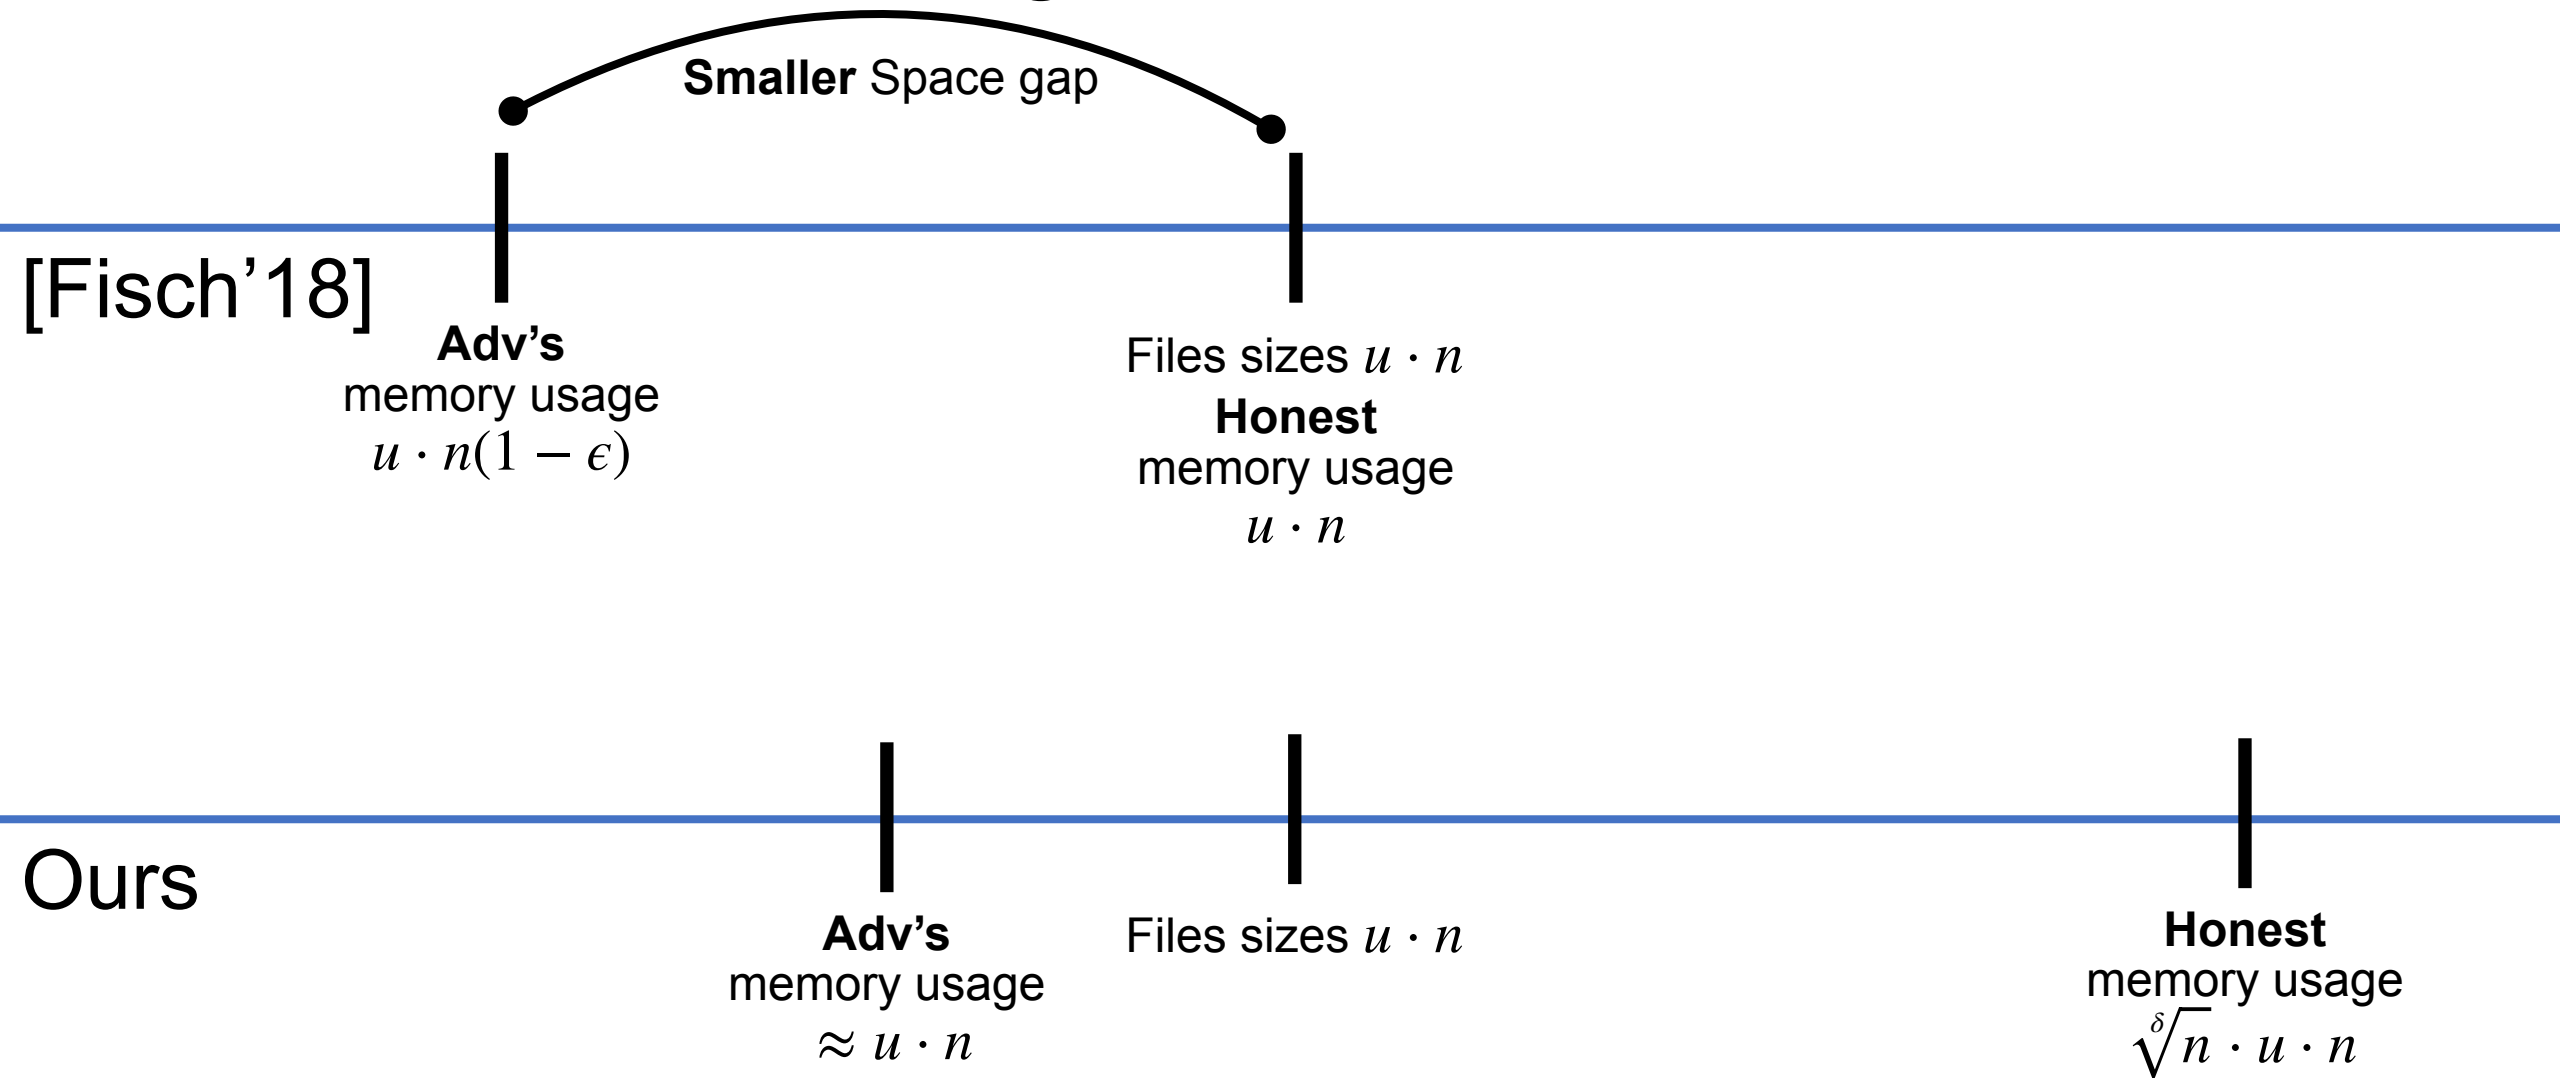
Files sizes  $u \cdot n$

The diagram for Ours features two horizontal blue lines. The bottom line is labeled Ours. Below it, a vertical black line is positioned to the left of the center, with the text 'Adv's memory usage' and the formula  $\approx u \cdot n$  below it. To the right of the center, another vertical black line is positioned, with the text 'Files sizes  $u \cdot n$ ' below it. The vertical distance between the two blue lines is smaller than in the [Fisch'18] diagram, and the two vertical lines are much closer to each other, indicating a smaller space gap.

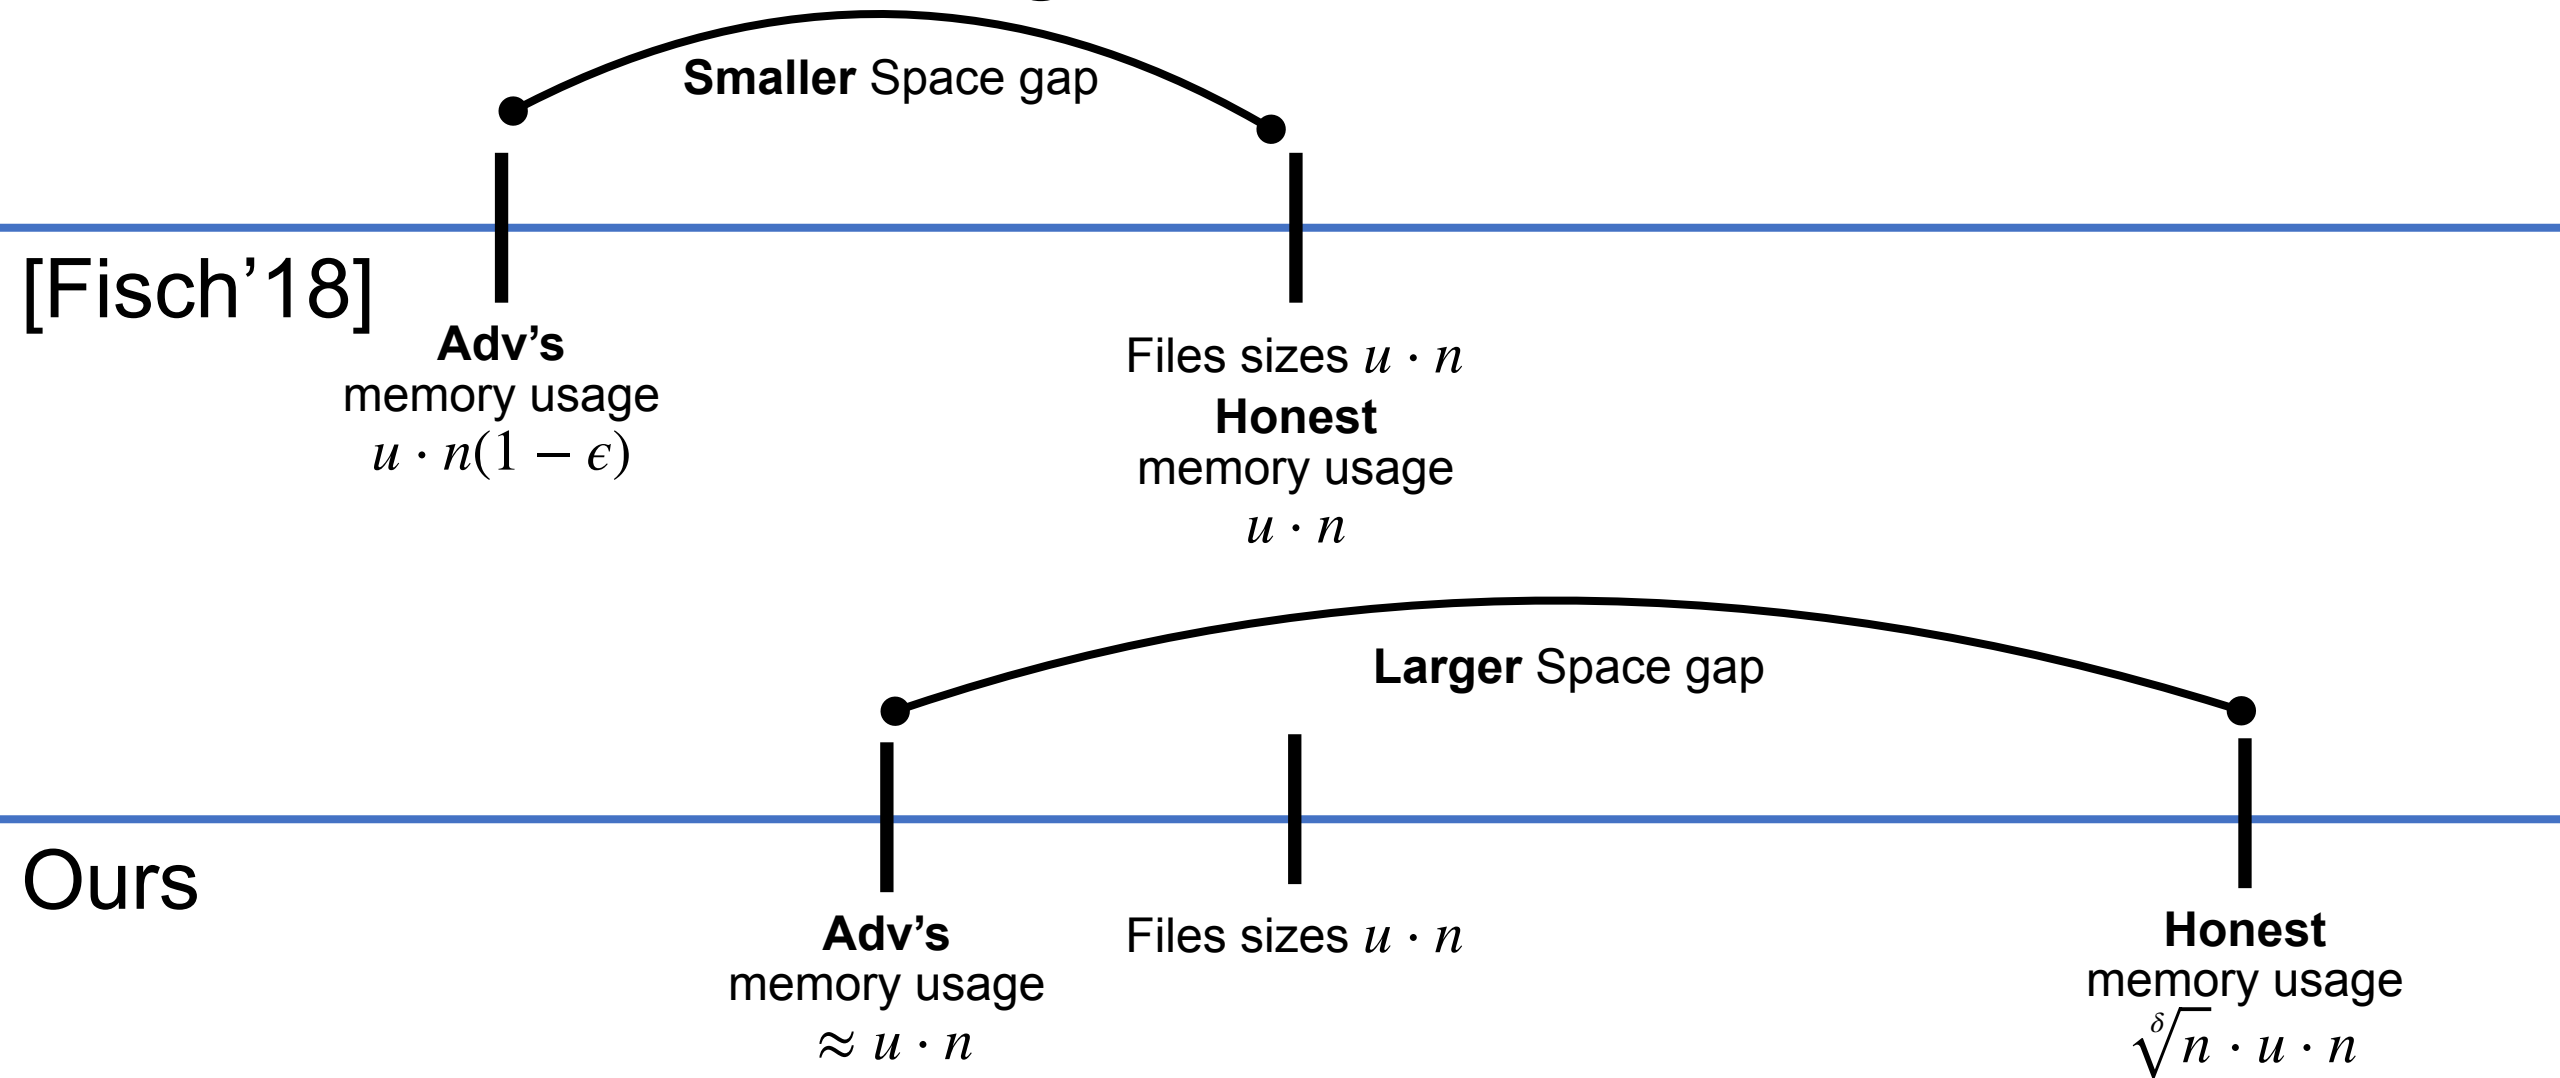
# Limitation = Space gap



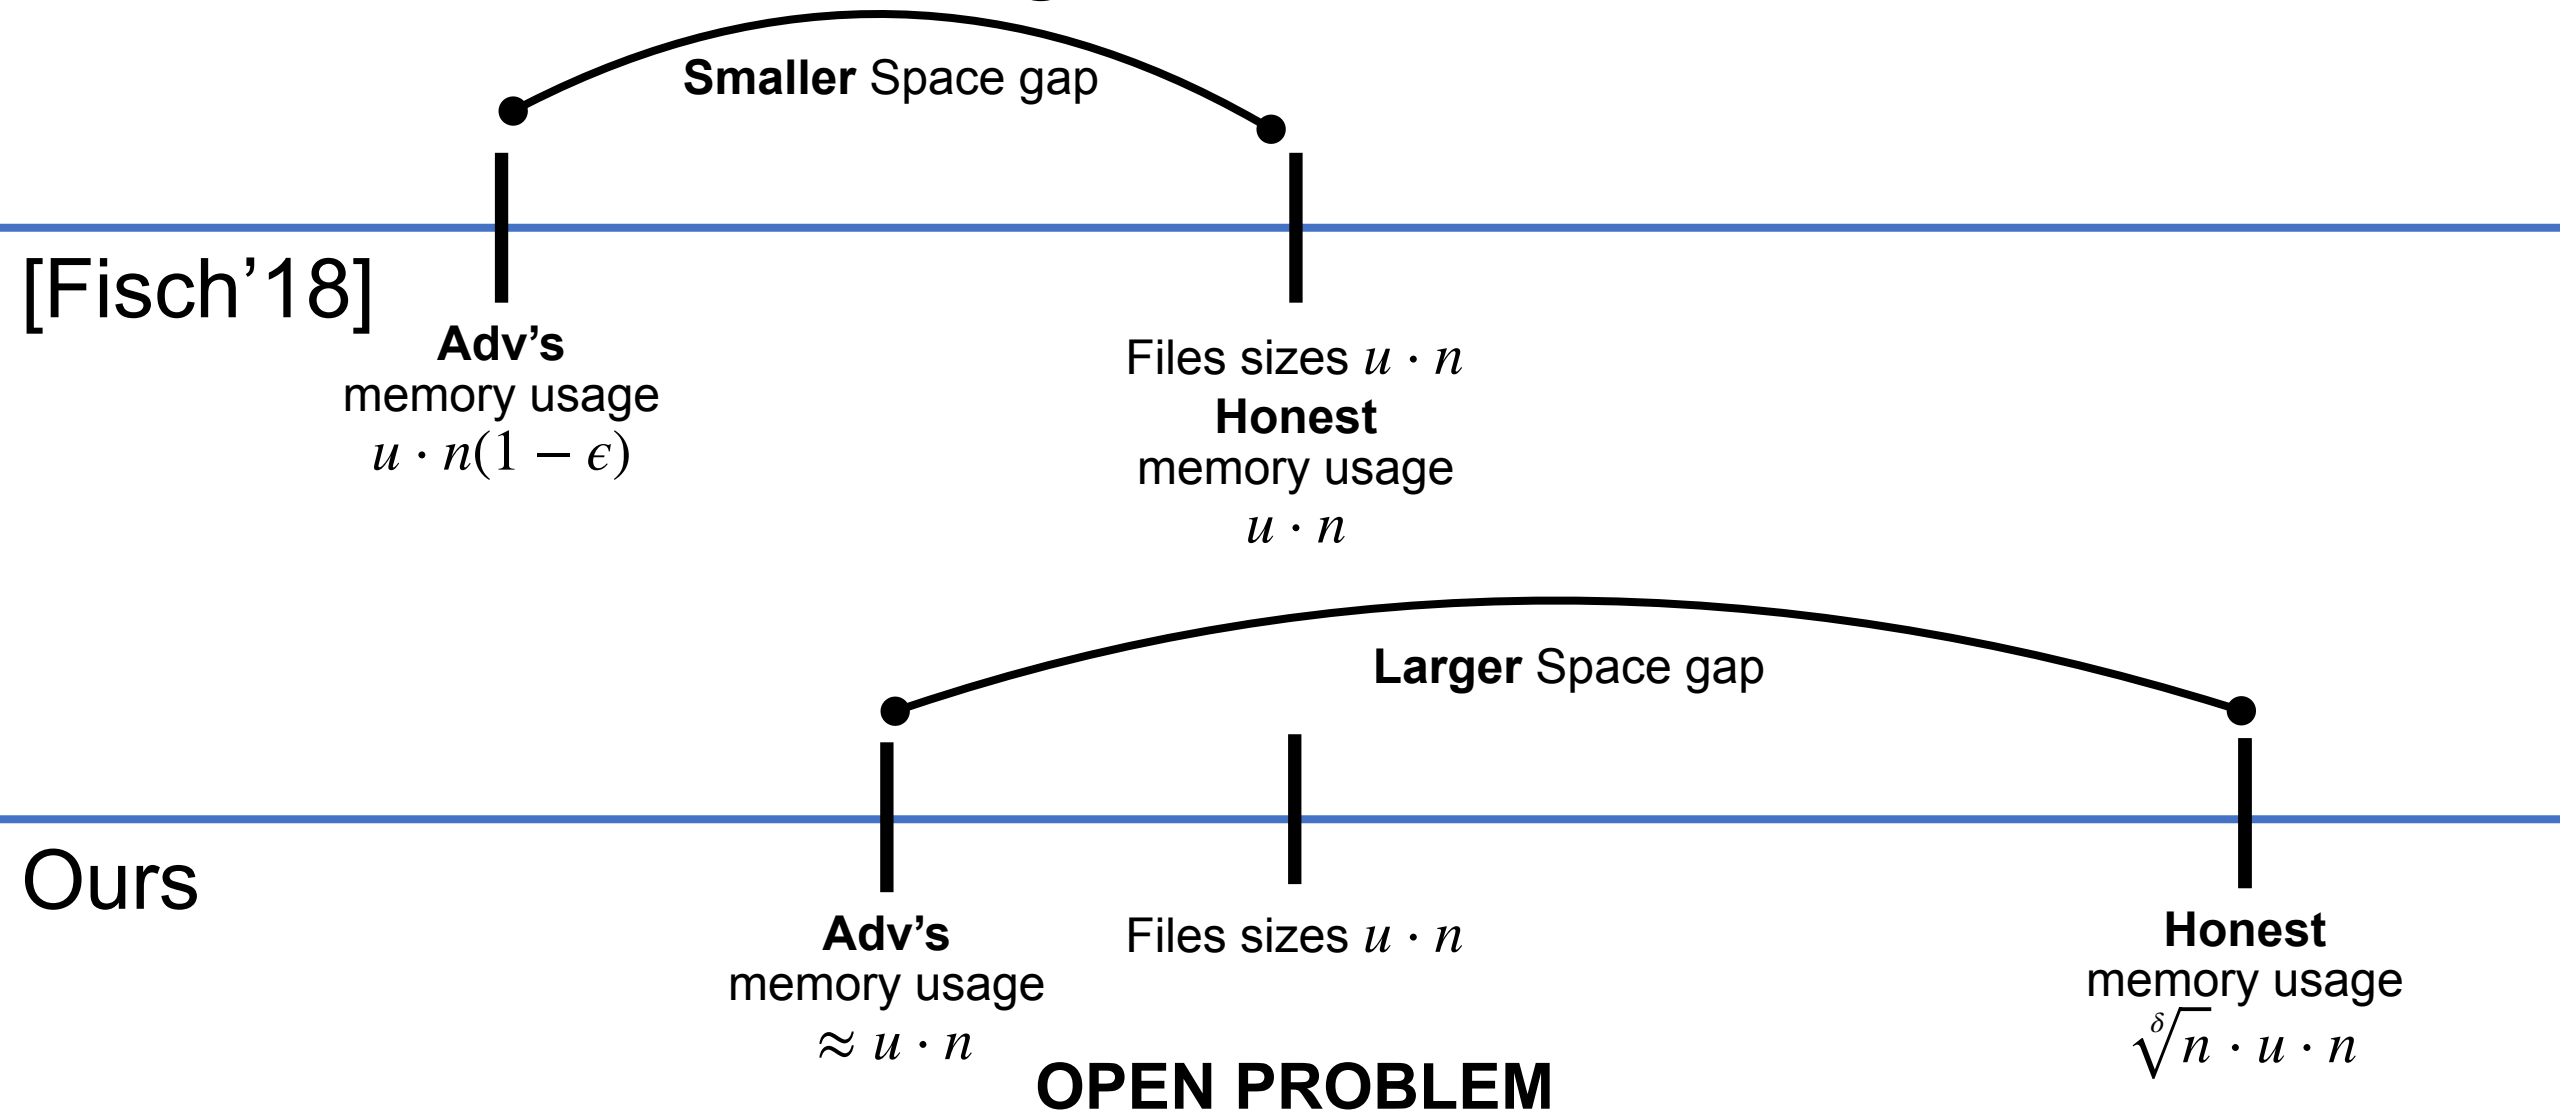
# Limitation = Space gap



# Limitation = Space gap



# Limitation = Space gap



Keep **efficiency** and **high (absolute) memory usage** while **minimising space gap**

**THANK YOU!**