The Complexity of Verifying Boolean Programs as Differentially Private

Ludmila Glinskih

Joint work with Mark Bun and Marco Gaboardi

To appear in the proceedings of CSF'22

Harvard University, November 22

Plan of the talk

1. Prior work
   - How hard is it to verify whether a program is DP for
     - Turing-complete languages
     - Boolean languages with bounded memory without loops
2. Our results and proof ideas
   - BPWhile: Boolean language with loops and finite memory
   - PSPACE-completeness of the verification of the DP for BPWhile
   - PSPACE-hardness: reduction from TQBF
   - PSPACE algorithm based on computing hitting probabilities in a Markov chain

Differential Privacy

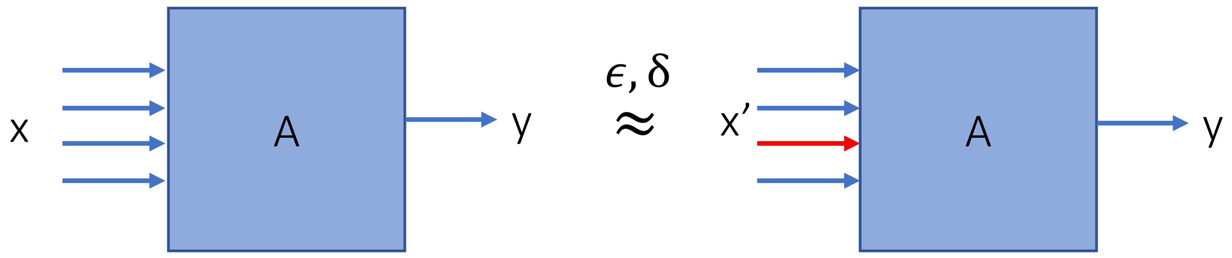

Non-termination is one of the possible outputs

C is (ϵ, δ)-differentially private if for every set of possible outputs *O*, and for every neighboring *x, x'*:

$$\mathrm{P}[C(x) \in O] \leq e^{\epsilon} \cdot \mathrm{P}[C(x') \in O] + \delta$$

Neighboring relation we consider:

- Inputs differ in one bit
- Can be extended to any bounded polyspace computable relation

Verification of Differential Privacy

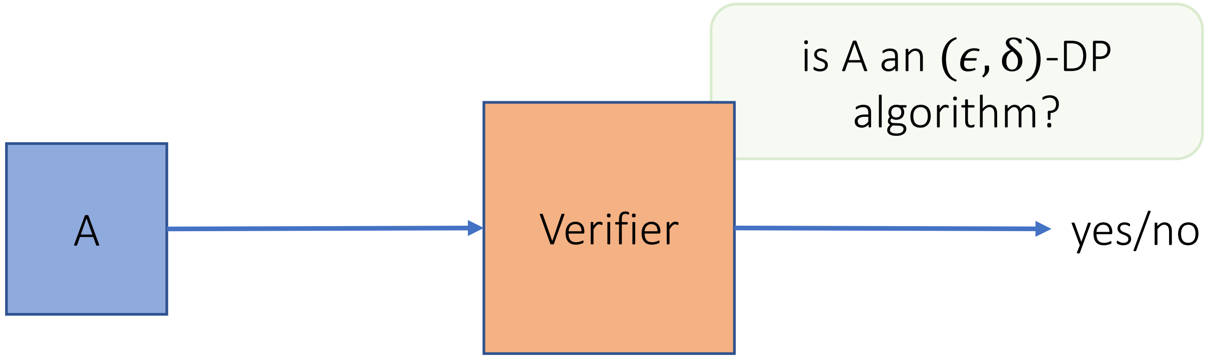

Complexity of verification depends on the expressivity of the language:

- For languages working with infinite data
  - Undecidable
    - $\Rightarrow$ Undecidable for Turing-complete languages

[Barthe, Chadha, Jagannath, Sistla, Viswanathan'20]

- For a simple Boolean language with **bounded memory,** if statements and **random assignments**, but **without loops**
  - $coNP^{\#P}$-completeness for $(\epsilon, 0)$-DP
    - Reduction from All-Min-SAT
  - $coNP^{\#P}$ and in $coNP^{\#P^{\#P}}$ for (ϵ, δ)-DP

[Gaboardi, Nissim, Purser'20]

Verification of DP: Black box vs White box

Complexity also depends on the type of access to the code:

- No information about the algorithm, query access
  - Impossible to verify $(\epsilon, 0)$-DP

[Gilbert, McMillan'19]

- Full access to the code/representation:
  - Linear algorithm in the size of automaton
    - For pure-DP
  - #P-hardness for approximating parameters in labelled Markov chains
    - For approximate-DP
    - Undecidable to compute exactly

[Chadha, Sistla, Viswanathan'21]

[Chistikov, Murawski, Purser'19]

In this work: **white-box** model

BPWhile: Boolean language with While loops

We design the language for the following goals:

- Captures classical computations on real computers
- Simple to analyze

x ::= $[a-z]^+$ — Variable identifiers

b ::= *true* | *false* | random | x | b ∧ b | b ∨ b | !b — Boolean expressions

c ::= skip | x := b | c; c | if b then c else c | **while b then c** — Commands

t ::= x | t, x — List of Boolean variables

p::= input(t); c; return(t) — Programs

Numbers of variables and input bits are fixed in the definition of the program

⇒ Length of the program is an upper bound on the size of the memory that the program uses

Example of a BPWhile program:

```
input(c⃗, n, ϵ);
k⃗ := ⌈log(2/ϵ)⌉;
d⃗ := (2^(k⃗+1) + 1)(2^k⃗ + 1)^(n−1);
u⃗ := uniform(0, d⃗];
z⃗ := 0;
r⃗ := n;
while z⃗ < n⃗ ∧ r⃗ = n then
  if z⃗ < c⃗ then
    if u⃗ ≤ 2^(k⃗(c⃗−z⃗))(2^k⃗ + 1)^(n−(c⃗−z⃗))
    then r⃗ := z⃗
    else skip
  else
    if u⃗ ≤ d − 2^(k⃗(z⃗−c⃗+1))(2^k⃗ + 1)^(n−1−(z⃗−c⃗))
    then r⃗ := z⃗
    else skip
  z⃗ = z⃗ + 1;
return(z⃗);
```

Implementation of the Bounded Geometric Mechanism in finite precision arithmetic

[Ghosh, Roughgarden, Sundararajan'09]

[Balcer, Vadhan'17]

Our results

Main result: if A is a BPWhile program, then the problem of verifying whether A is differentially private is PSPACE-complete.

It holds for the following notions of differential privacy:

- $(\epsilon, 0)$-DP
- (ϵ, δ)-DP
- (ϵ, δ)-DP parameters approximation
- Renyi-DP
- Zero-Concentrated-DP
- Truncated Concentrated-DP

PSPACE-completeness

PSPACE-completeness of a problem A implies:

- A is solvable by a TM that uses polynomial space
- A is solvable in exponential time
- A is at least as hard as any problem solvable in polyspace
- No polytime algorithm for A, unless P = PSPACE
  - That is widely believed not to be true

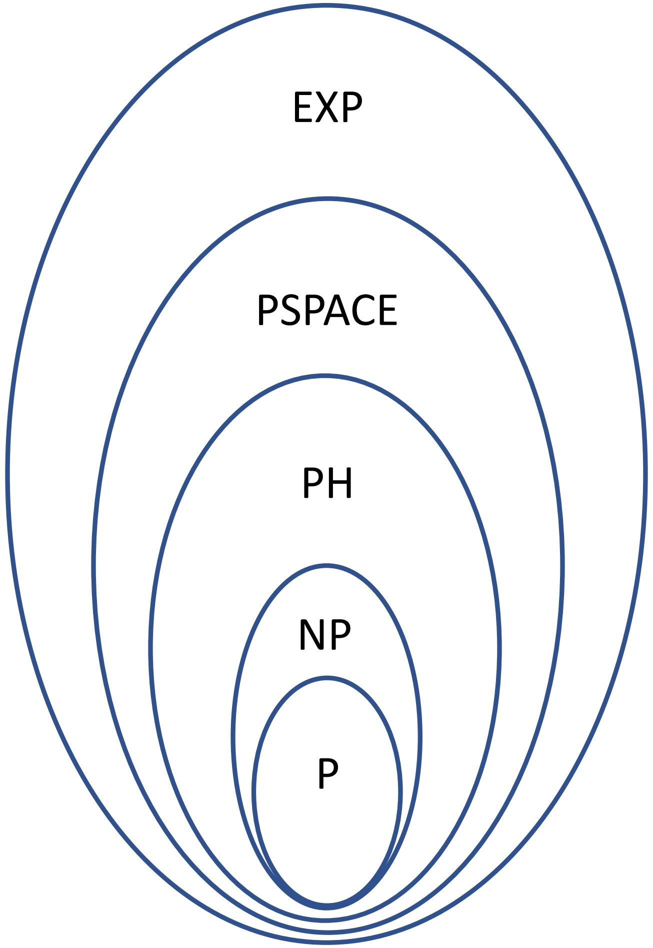

To show PSPACE completeness we need:

1. Show hardness: construct sequence of reductions from TQBF
2. Construct polynomial-space algorithm: analyze Markov chain based on the state graph of the program

PSPACE-hardness

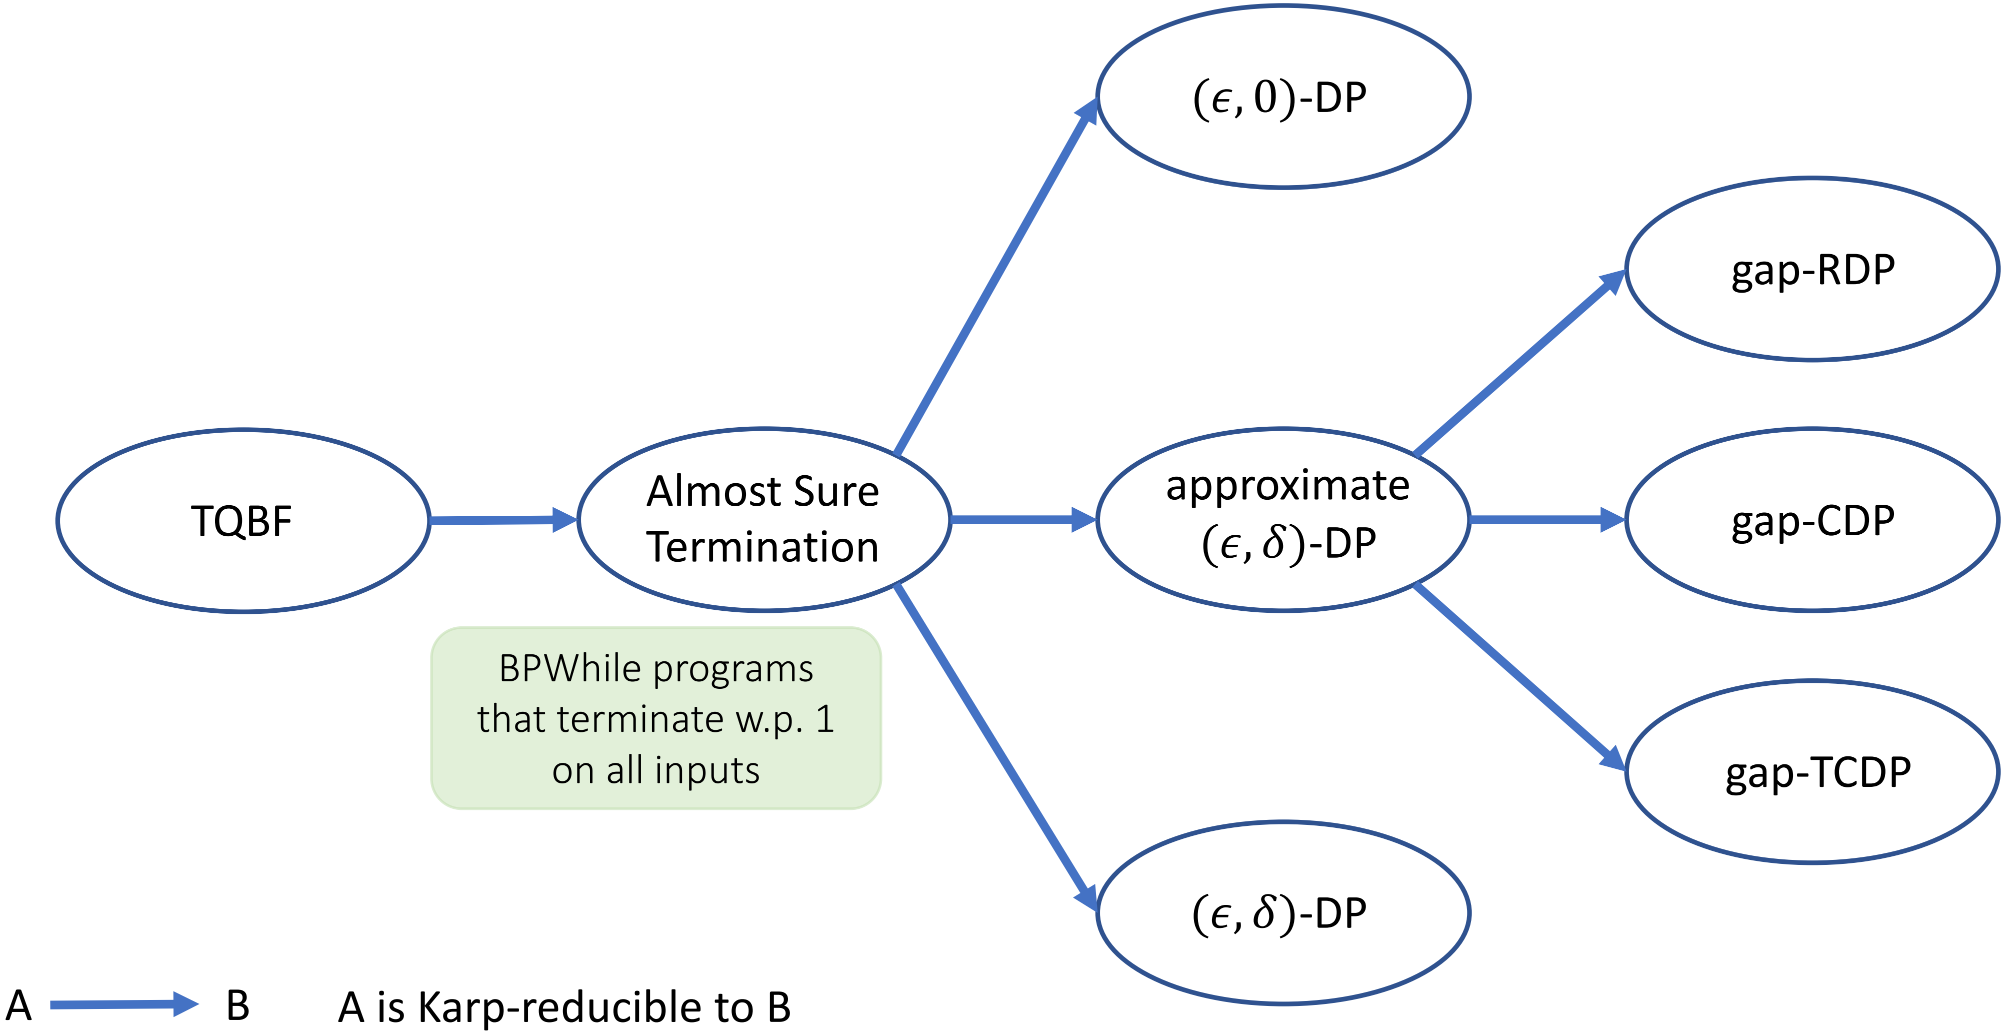

PSPACE-hardness: illustrative reduction

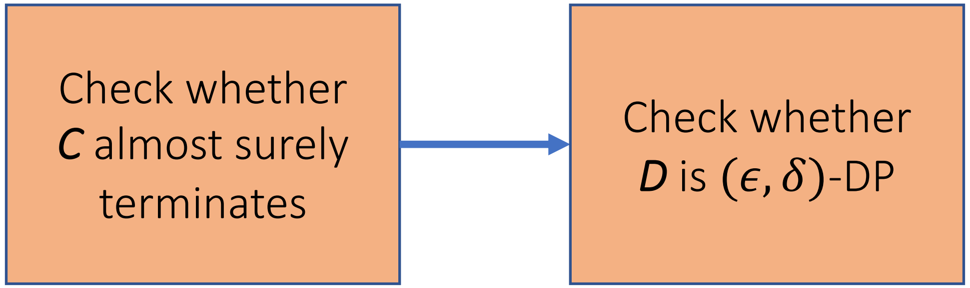

We need to show that
C almost surely terminates $\Leftrightarrow$ D is (ϵ, δ)-DP

C almost surely terminates $\Rightarrow$ D is (ϵ, δ)-DP

If C **terminates w.p.** 1, then for all x:

- $D(x, 1)$ doesn't terminate w.p. δ
- $D(x, 1)$ terminates and outputs 1 w.p. $1 - \delta$
- $D(x, 0)$ doesn't terminate w.p. 0
- $D(x, 0)$ terminates and outputs 1 w.p. 1

D:

```
input(x,b);
if b == 1 then
    C(x);
    r = delta_rand();
    if r == 0 then
        while true then
            skip;
    else skip;
else skip;
return(1)
```

Copy of a program that returns 0 w.p. delta

PSPACE-hardness: illustrative reduction

Checked:

C almost surely terminates ⇒ *D* is (ϵ, δ)-DP

Need to check:

C doesn't almost surely terminate ⇒ *D* is **not** (ϵ, δ)-DP

If **C doesn't terminate w.p.** p on x:

- $D(x, 1)$ doesn't terminate w.p. $p + (1 - p)\delta > \delta$
- $D(x, 1)$ terminates and outputs 1 w.p. $< 1 - \delta$
- $D(x, 0)$ doesn't terminate w.p. 0
- $D(x, 0)$ terminates and outputs 1 w.p. 1

D:

```
input(x,b);
if b == 1 then
    C(x);
    r = delta_rand();
    if r == 0 then
        while true then
            skip;
    else skip;
else skip;
return(1)
```

Copy of a program that returns 0 w.p. delta

Polyspace membership: algorithm for (ϵ,0)-DP

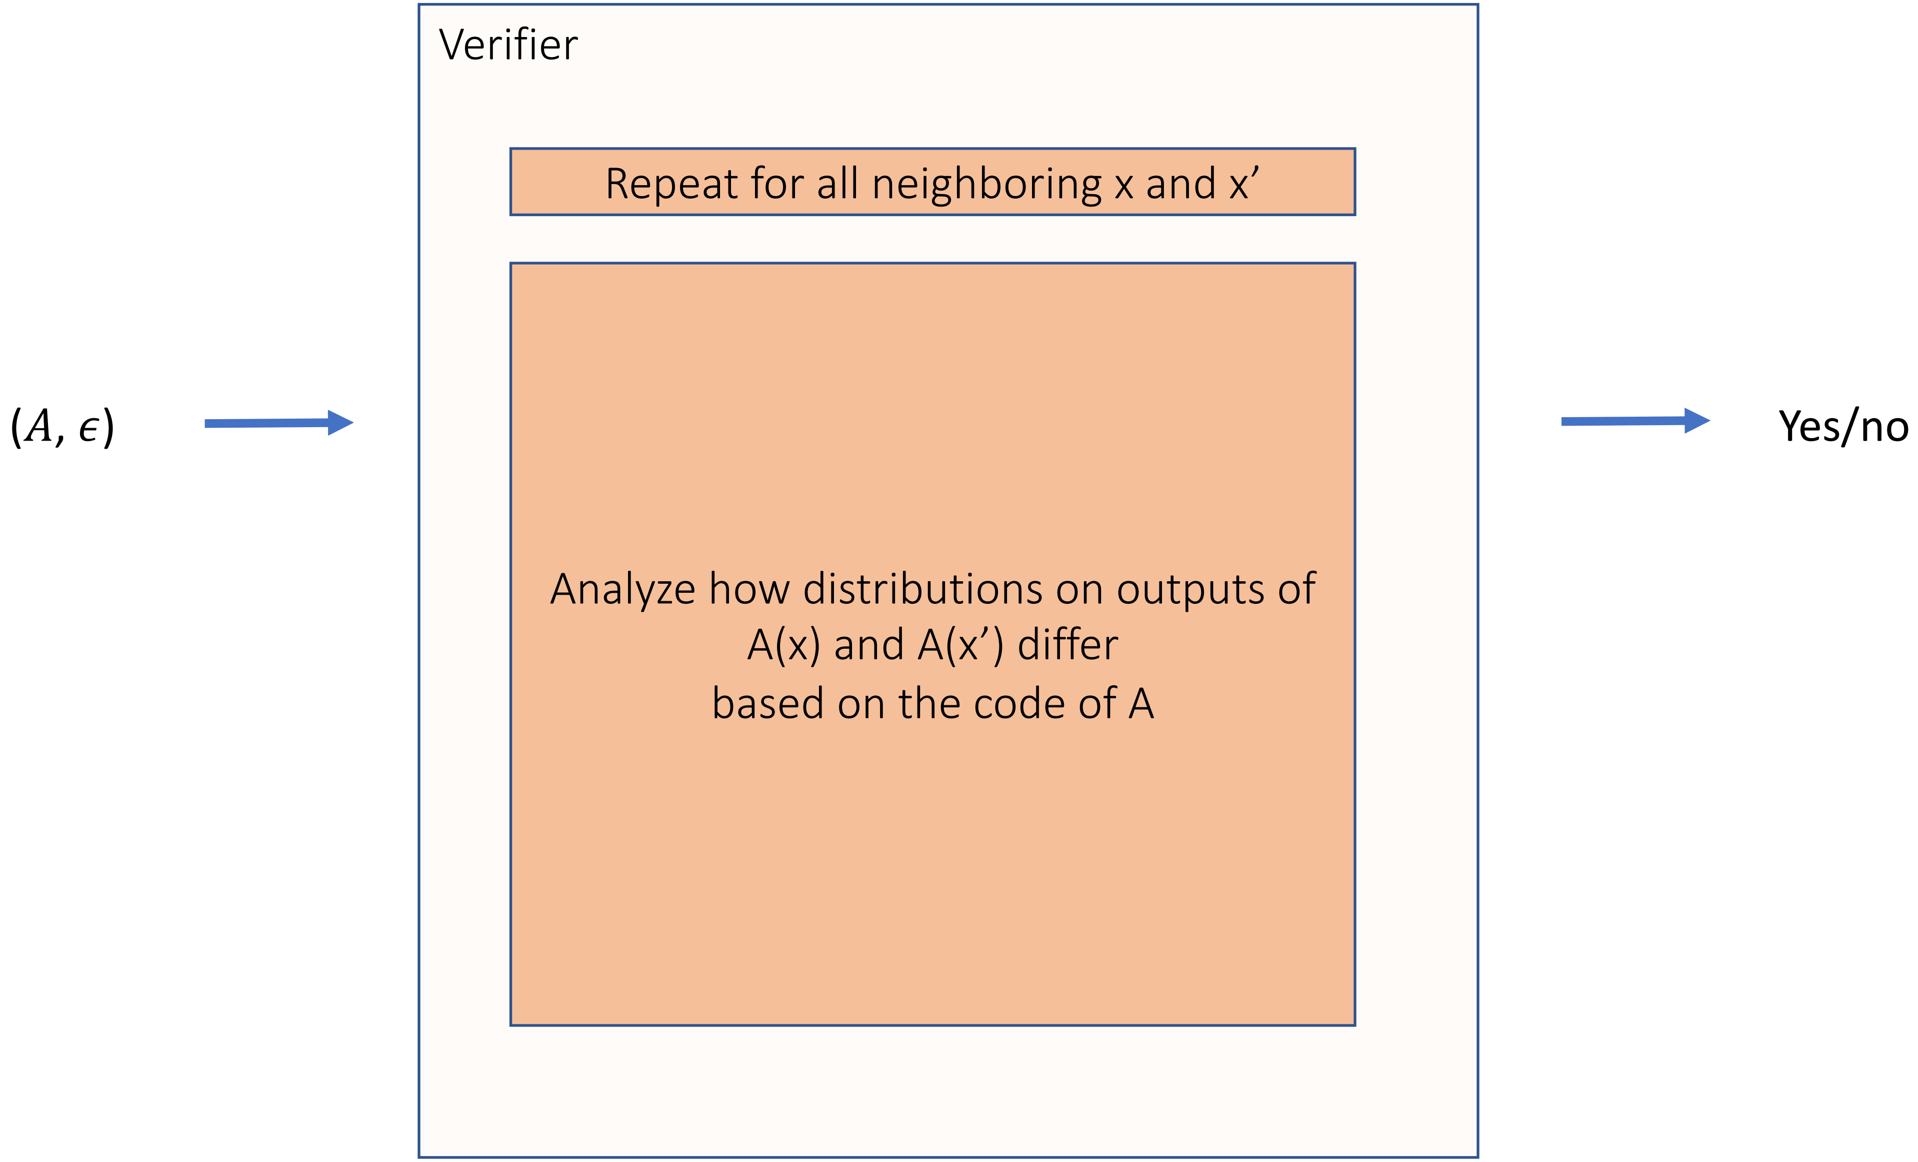

PSPACE membership: state graph

State graph depends on the input values

D(b=1):

```
1. input(b);
2. if b == 1 then
3.     r = rand();
4.     if r == 0 then
5.         while true
6.             skip;
7.     else skip;
8. else skip;
9. return(1)
```

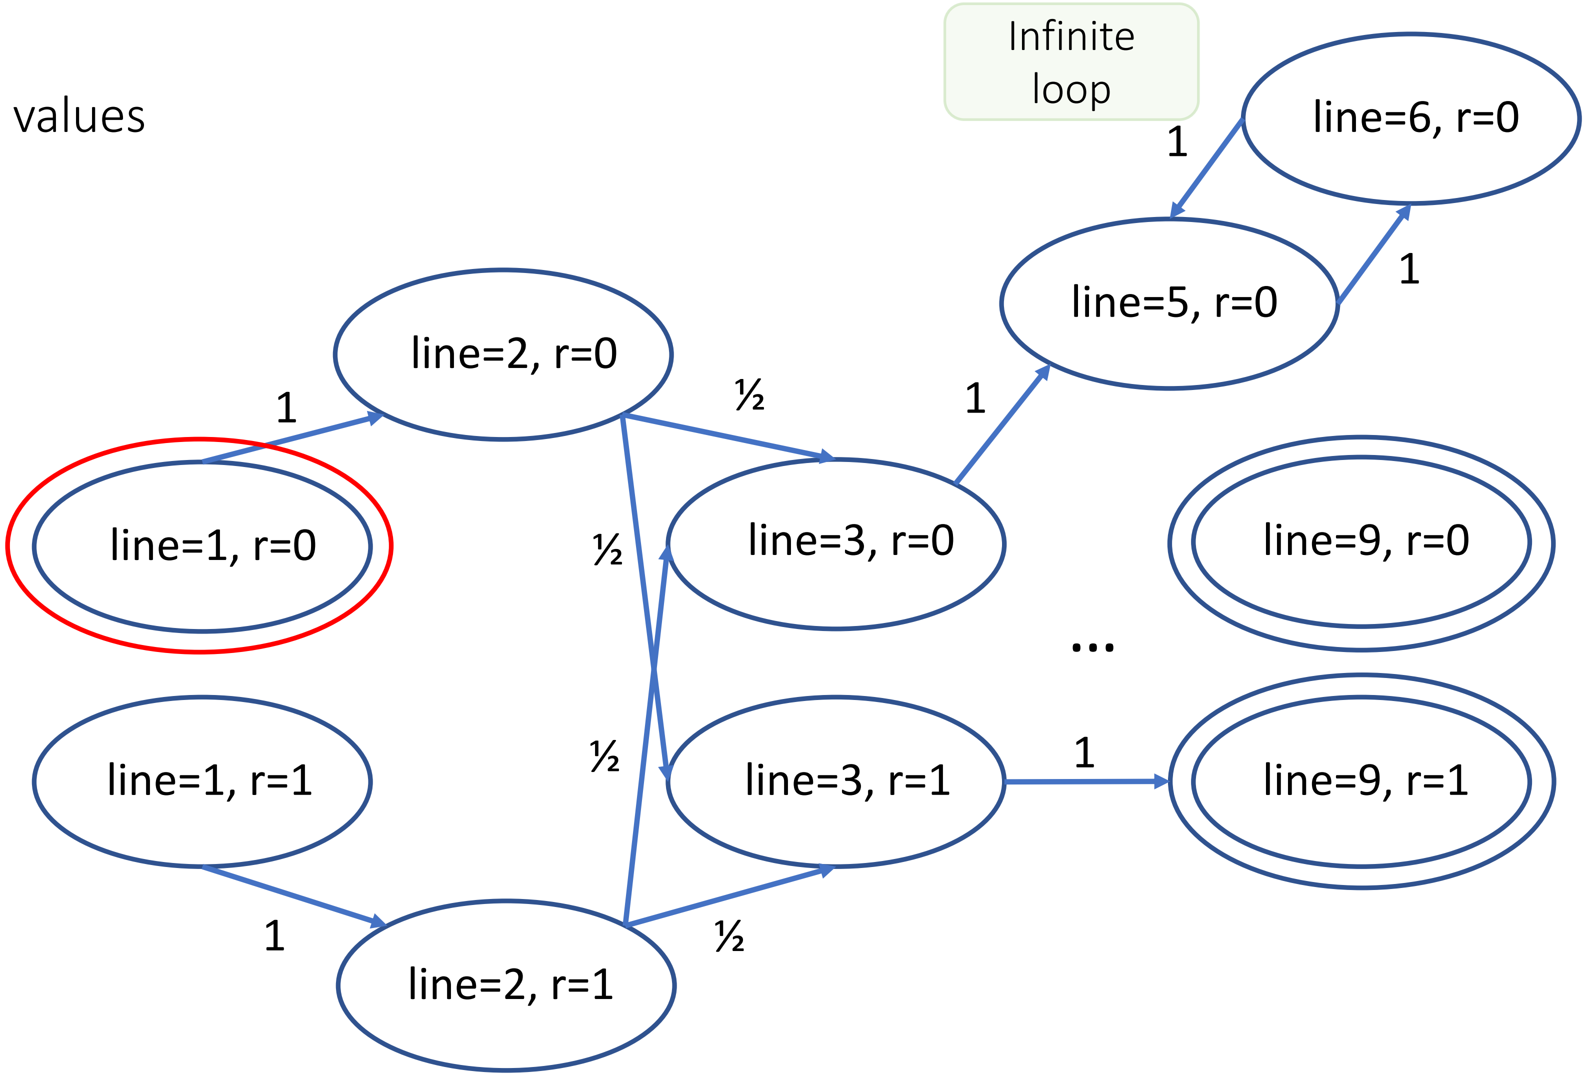

PSPACE membership: state graph

D(b=0):

```
1.  input(b);
2.  if b == 1 then
3.      r = rand();
4.      if r == 0 then
5.          while true
6.              skip;
7.      else skip;
8.  else skip;
9.  return(1)
```

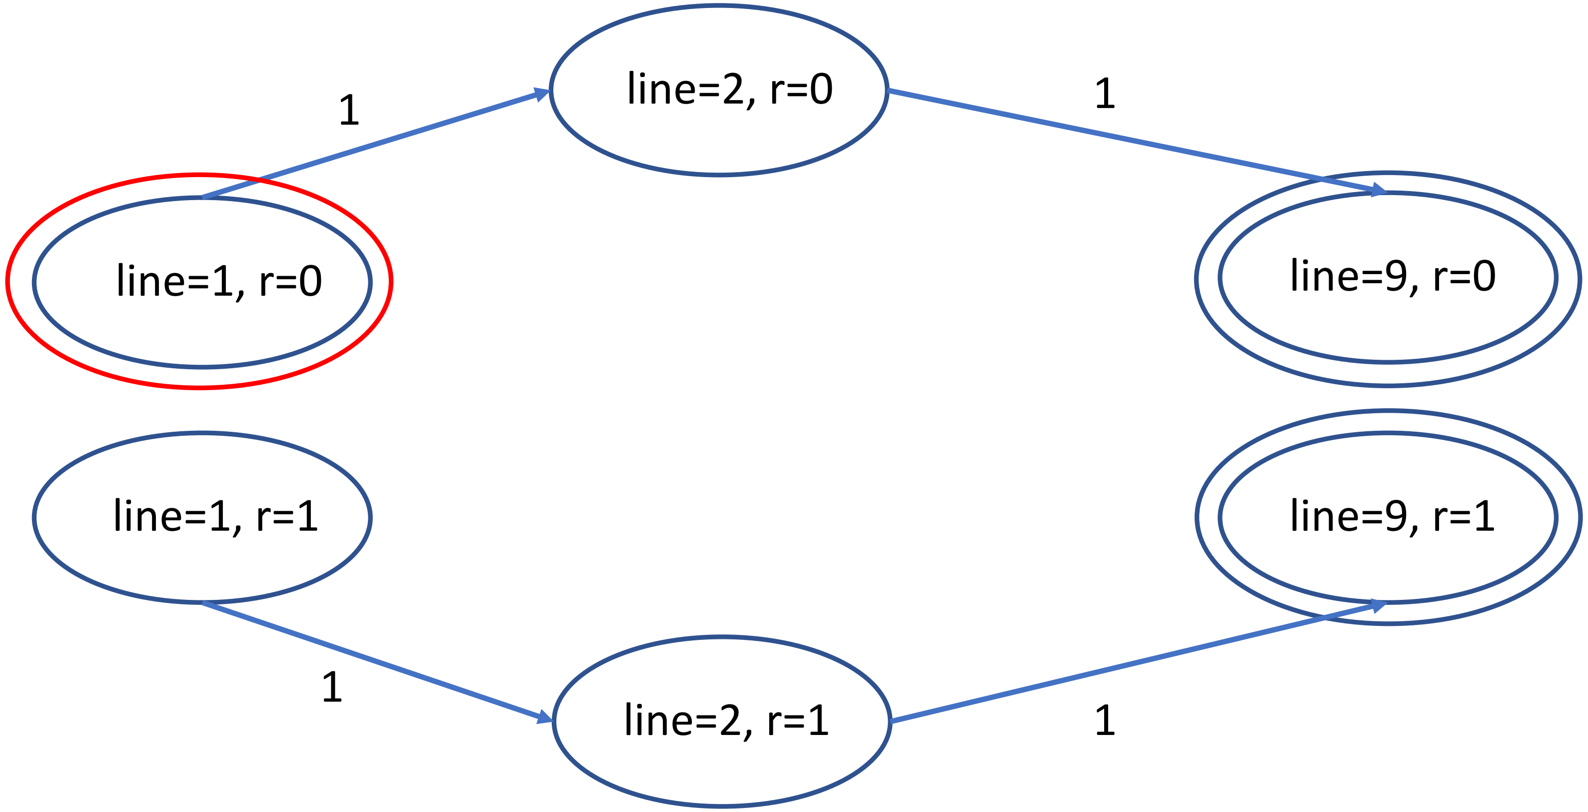

PSPACE membership: algorithm for (ϵ,0)-DP

For a program D and all neighboring inputs x, x':

- Construct the Markov chain for $D(x)$ and $D(x')$
- Compute and compare hitting probabilities

Problem: Markov chain has exp-many states ⇒ cannot store it explicitly

Need **space-efficient** algorithm for computing hitting probabilities with **implicit access** to the Markov chain

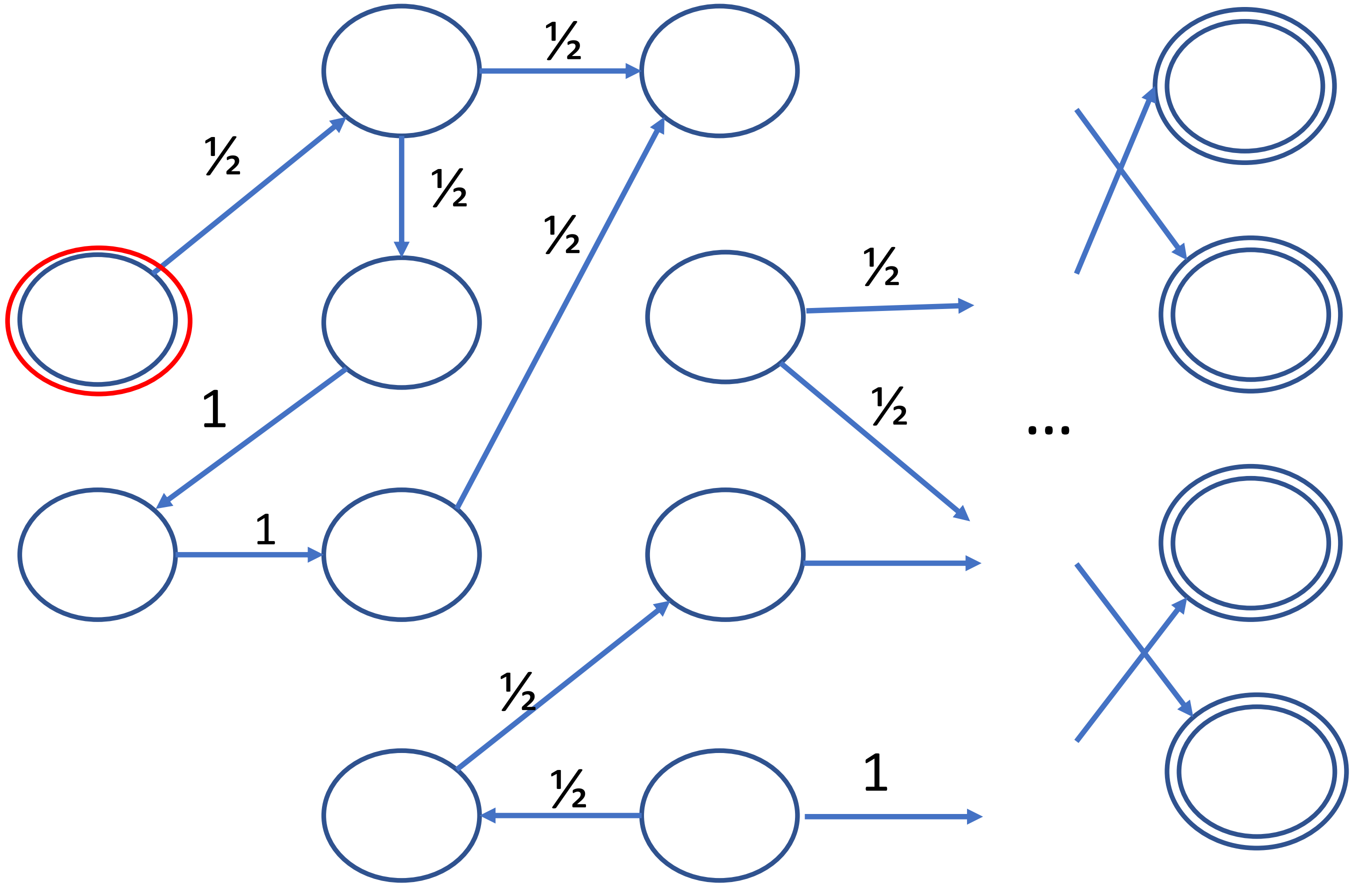

Polyspace algorithm for computing hitting probabilities in a Markov chain

Lemma [Simon'81]: M is a Markov chain with at most 2^L states

- initial distribution placing all mass on one state,
- a set of final states F each with only one self-transition,
- every non-final state has outgoing transition probability 0 or ½.

Then, there is an $O(L^6)$-space deterministic algorithm that computes the hitting probabilities of every state in F.

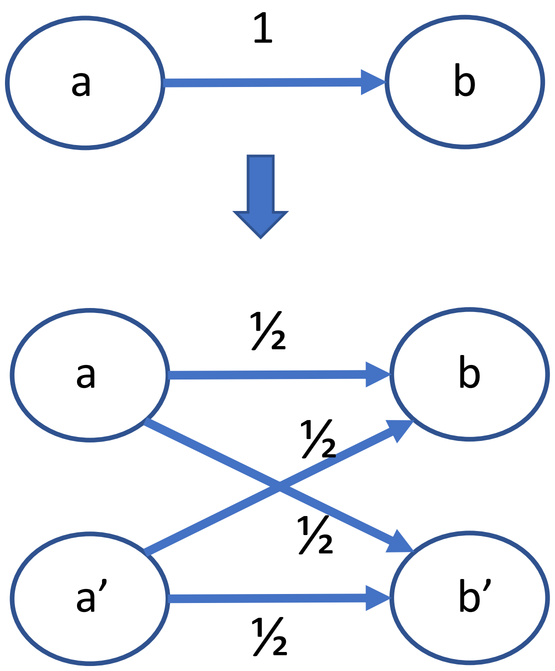

Note: to use the algorithm, we need to replace all transitions labelled by 1 in the state graph of the BPWhile program:

- Clone all states
- For each state *a* with outgoing edge w.p. 1 replace it by two edges:
  - Edge (a,b) with weight ½ to original state
  - Edge (a,b') with weight ½ to the clone-state b' of b

PSPACE membership: exponentially long numbers

For a program D

for all neighboring inputs x, x':

- Construct the Markov chain for $D(x)$ and $D(x')$
- Compare hitting probabilities of the final states with the same values

Problem:

Markov chain has exp-many states
$\Rightarrow$ hitting probabilities can be as small as $\frac{1}{2^{exp}}$

$\Rightarrow$ numbers are exponentially long

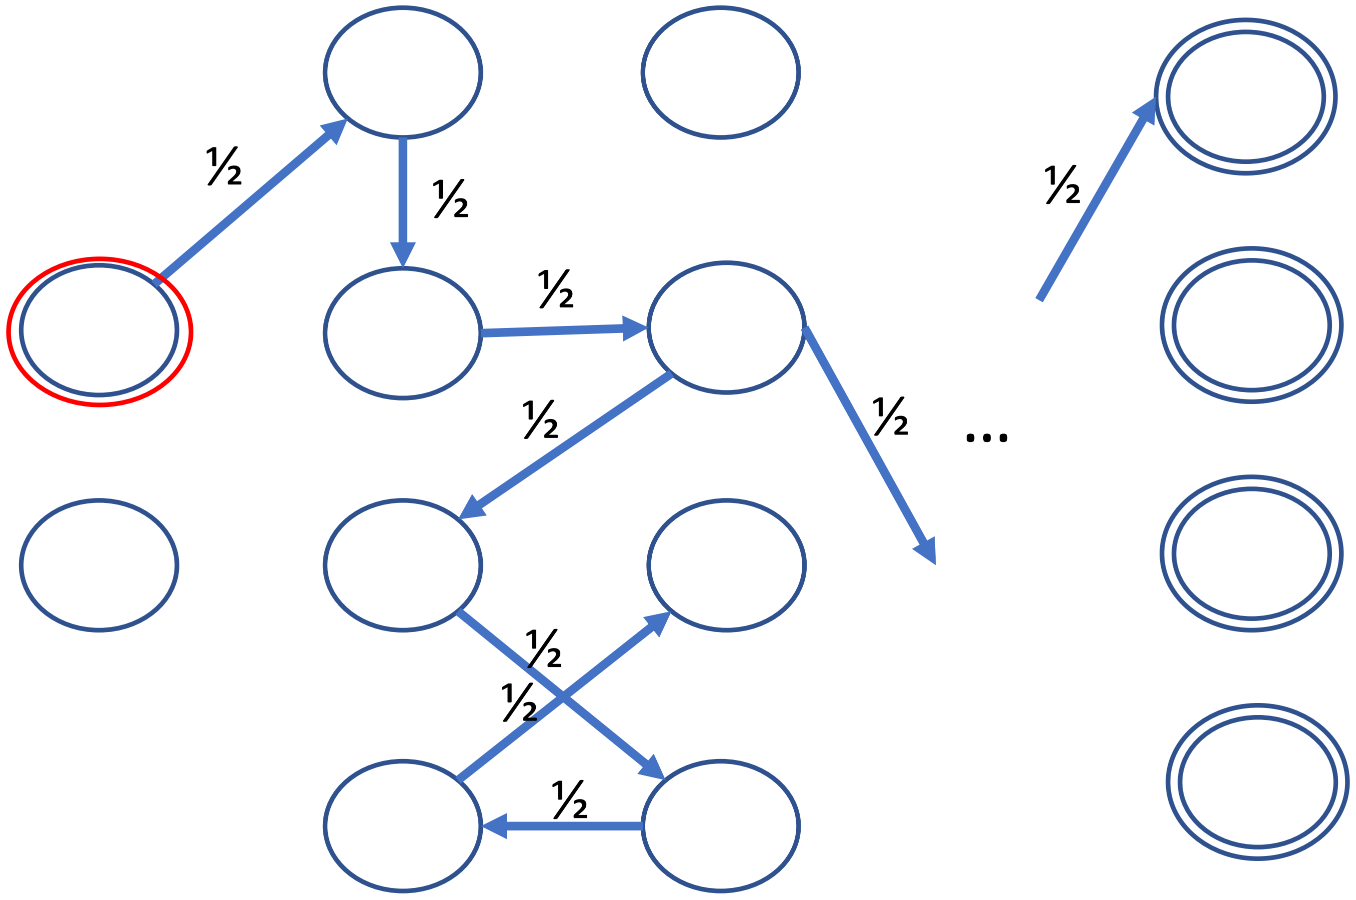

Operations with exponentially long numbers

	1	**0**	**1**	**1**	**1**	**…**	**1**	**0**	**0**	**1**
+	**0**	**0**	**1**	**0**	**1**	**…**	**1**	**1**	**1**	**1**
	1	**1**	**1**	**0**	**1**	**…**	**1**	**0**	**0**	**0**

Uniform family of log-depth circuits:

- One logspace algorithm provides implicit access to the circuits
- Each circuit has log-depth and poly size

Lemma:

Uniform families of log-depth circuits exist for:

- Comparison
- Addition
- Multiplication by a fixed rational number
- Multiplication [Ofman'62]
- Square roots [Reif'86]

Polyspace membership algorithm for (ϵ,0)-DP

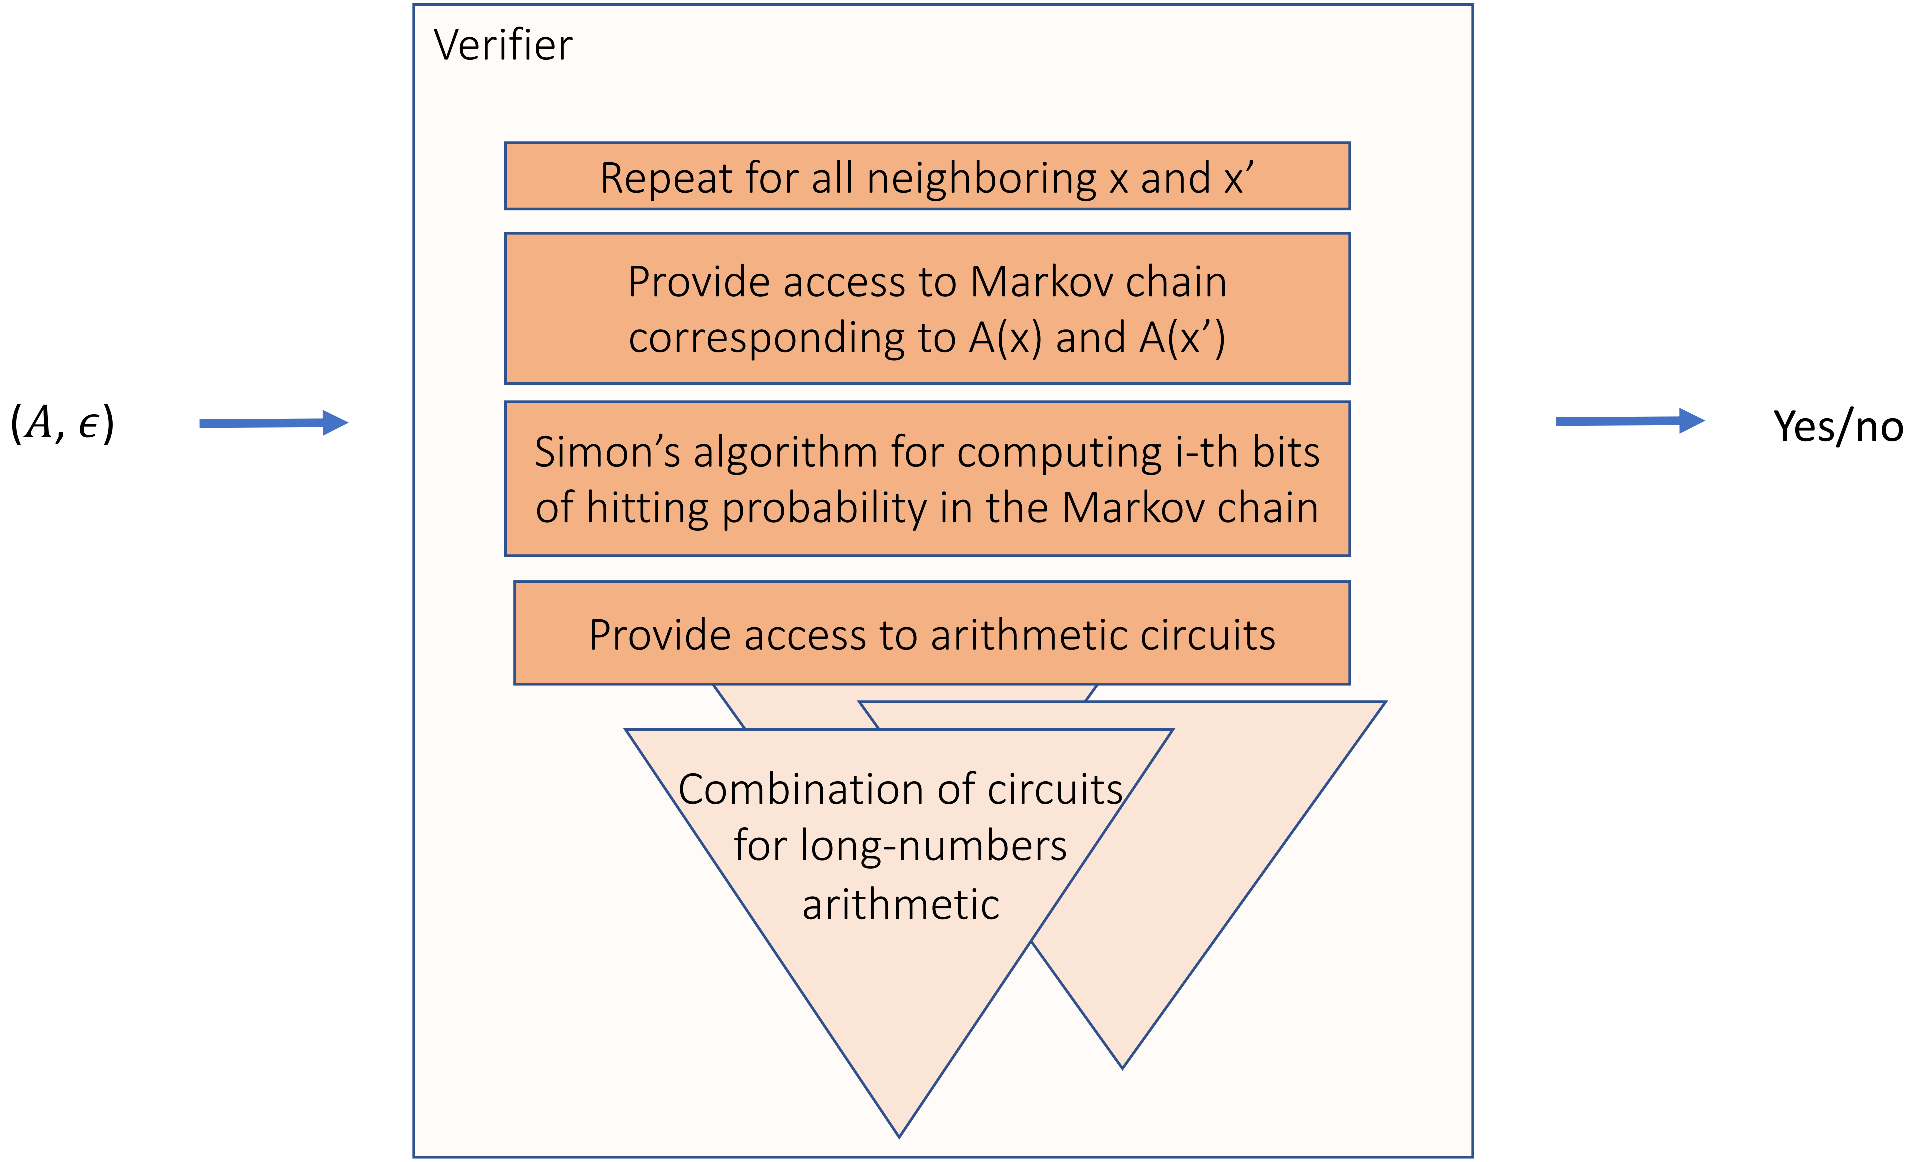

Polyspace algorithm for RDP, zCDP, TCDP

An algorithm is **RDP/zCDP/TCDP** if for a privacy parameter ρ and a fixed/any/bounded $\alpha > 1$ Rényi divergence for any neighboring inputs x, x' is at most $\rho\alpha$

[Mironov'17],[Dwork-Rothblum'16,Bun-Steinke'16], [Bun,Dwork,Rothblum,Steinke'18]

$$D_\alpha(P|Q) = \frac{1}{\alpha - 1} \log \sum \frac{p_i^\alpha}{q_i^{\alpha-1}}$$

New problems:

- To compute Rényi divergence we compute
  - Logarithms
  - Exponentiations to rational degrees
- Hence, get infinite fractions

Solution:

- Computations with a fixed precision η
- Consider gap-versions of the problem

We define Gap-RDP on (C, ρ, α, η) as follows:

- Yes-instance, if for all neighboring x, x'

$$D_\alpha(C(x)|C(x')) \le \rho\alpha$$

- No-instance, if for at least one neighboring x, x'

$$D_\alpha(C(x)|C(x')) \ge \rho\alpha + \frac{1}{2^\eta}$$

Polyspace algorithm for zCDP

Another problem: an algorithm is **zCDP** if **for all values of** α Rényi divergence for any neighboring inputs x, x' is bounded by $\rho\alpha$

Solution: showed that it is sufficient to check values of alpha from **the bounded range**:

Lemma:

C is ρ-zCDP, then for all neighboring x, x': $D_\alpha(C(x) | C(x')) \leq \rho\alpha$

C is not $(\rho + 2^{-\eta})$-zCDP => exists neighboring x, x', exists $\alpha \in (1, 1 + 2^{poly(n)}/\rho)$:

- α is a multiple of $2^{-\eta}$
- $D_\alpha(C(x) | C(x')) \geq \rho\alpha + 2^{-\eta-1}$

Polyspace membership algorithm for Gap-zCDP

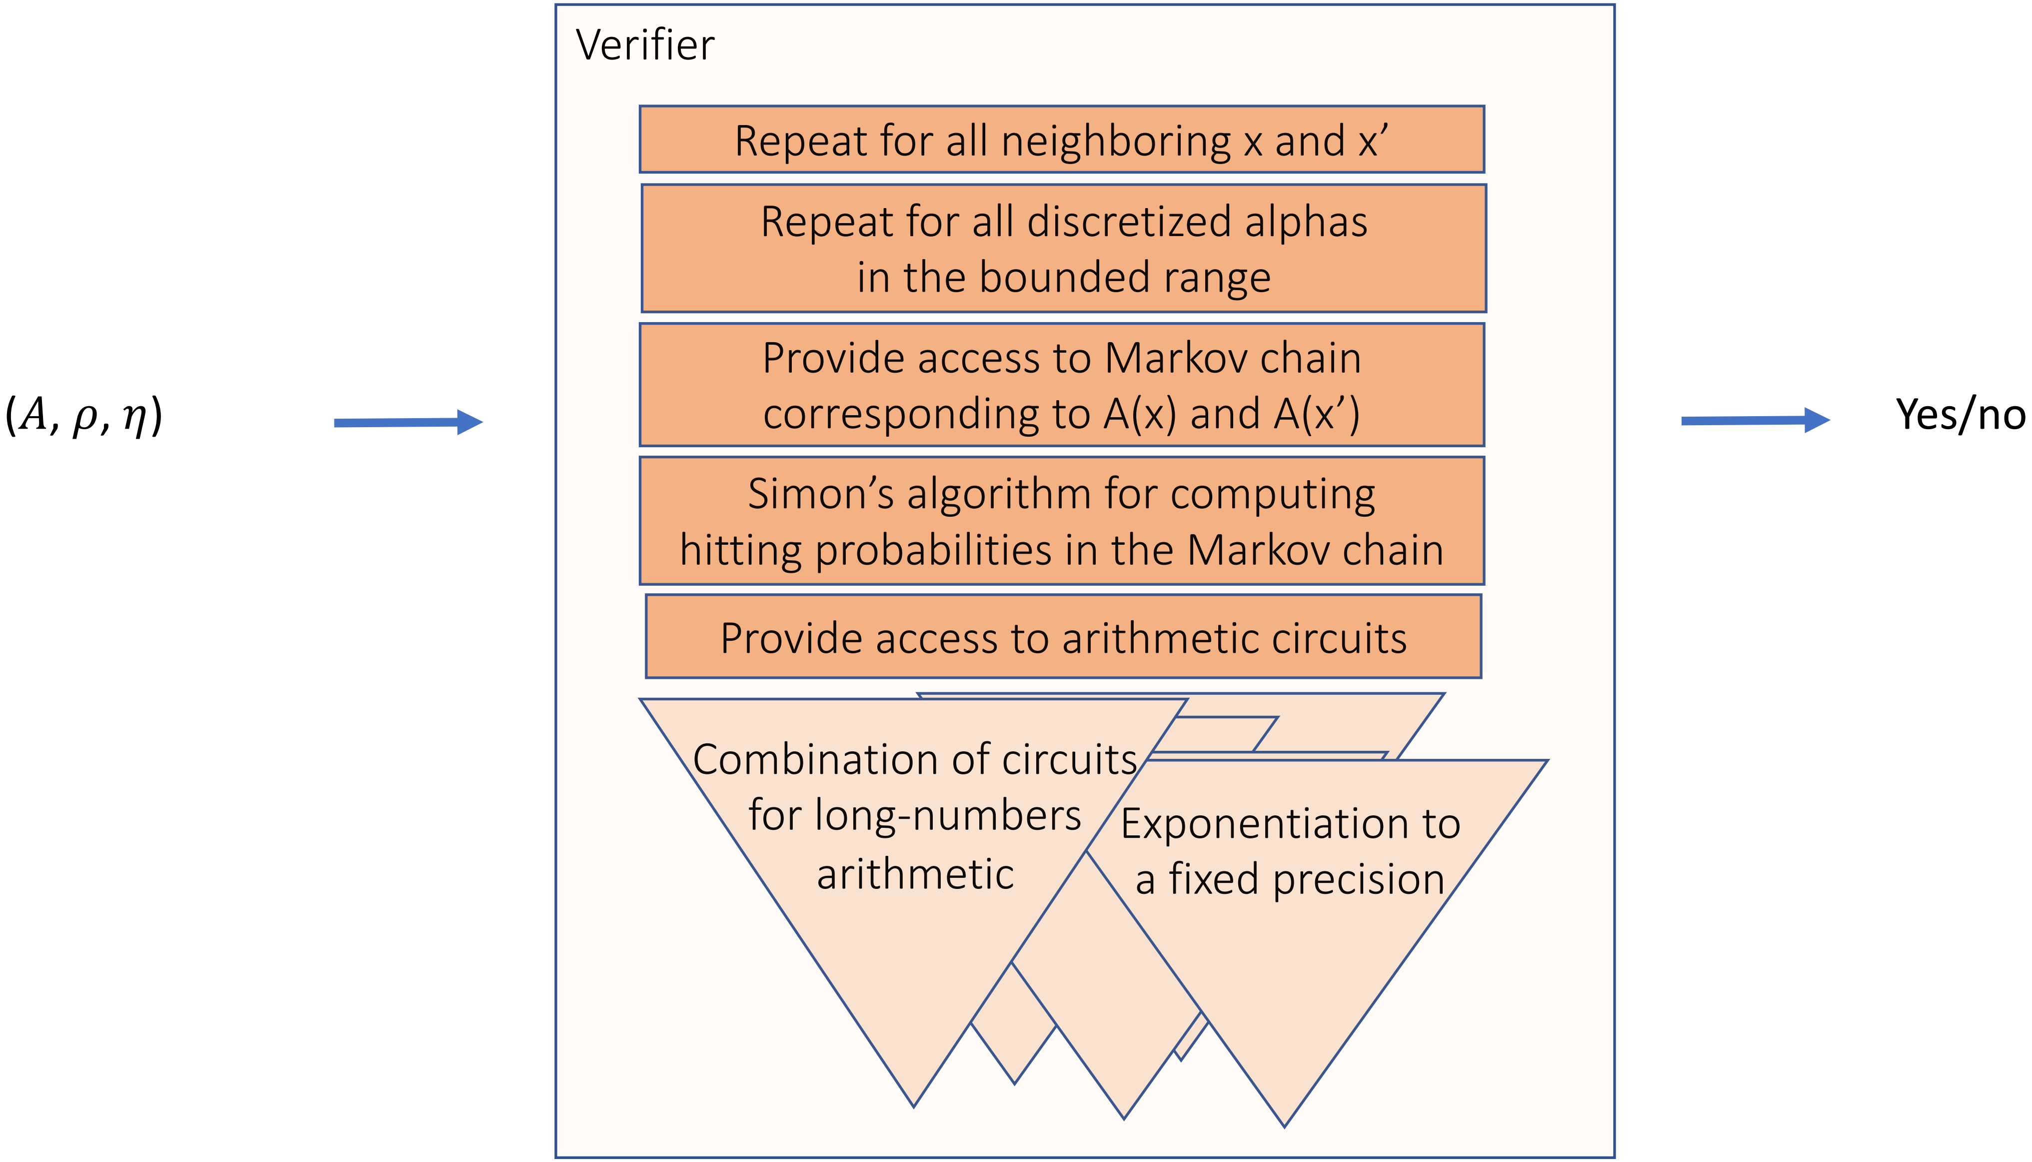

Results and future work

- We showed PSPACE-completeness for the problems of checking:
  - Pure-DP
  - Approximate-DP
  - Gap-RDP
  - Gap-zCDP
  - Gap-TCDP

- Possibly can extend the result to show PSPACE-completeness of verifying accuracy
- Improve the exact polynomial in the space complexity of the algorithm
  - Improved analysis and more efficient algorithms for Markov chains analysis and long-numbers arithmetic operations are needed for tighter results