Differentially Private Release and Learning of Threshold Functions

Mark Bun — Harvard
Kobbi Nissim — Harvard and Ben-Gurion
Uri Stemmer — Ben-Gurion
Salil Vadhan — Harvard

Median Estimation

Data domain [T] = {1, ..., T}

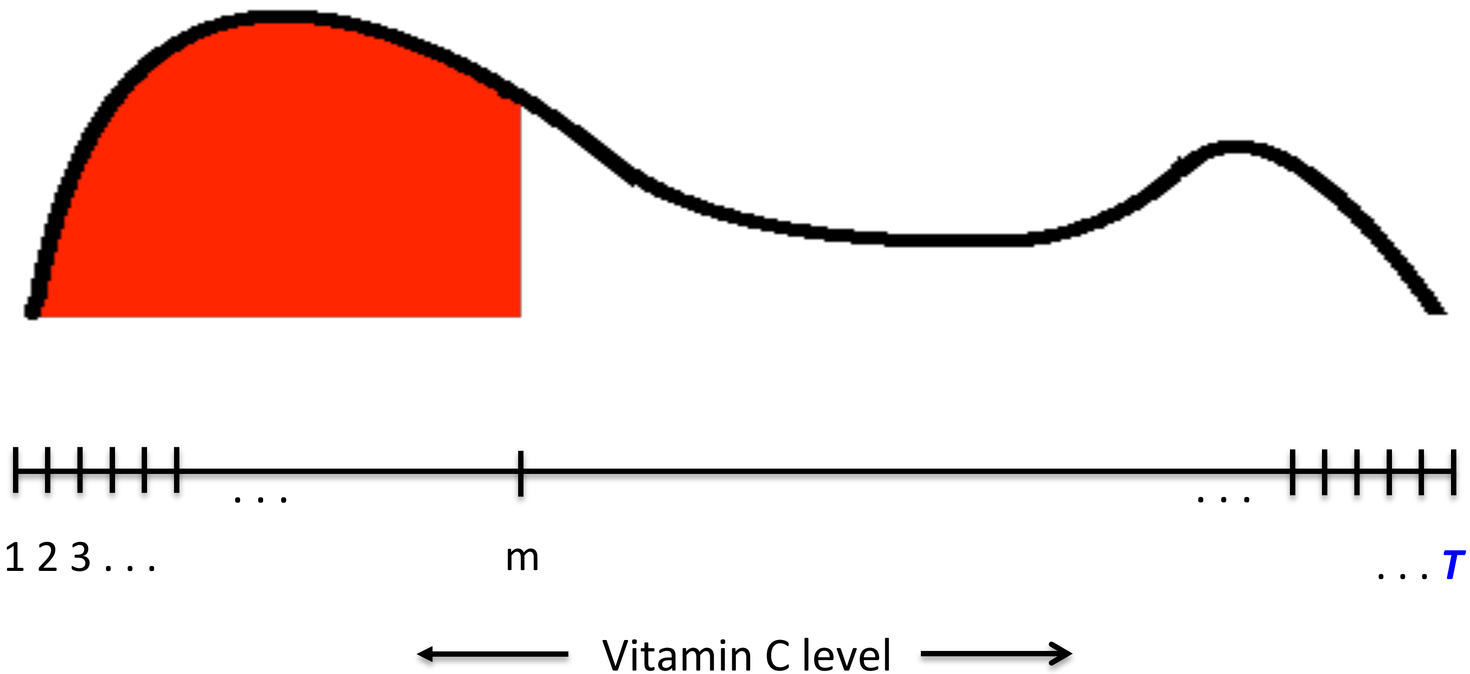

Median Estimation

Data domain [*T*] = {1, ..., *T*}

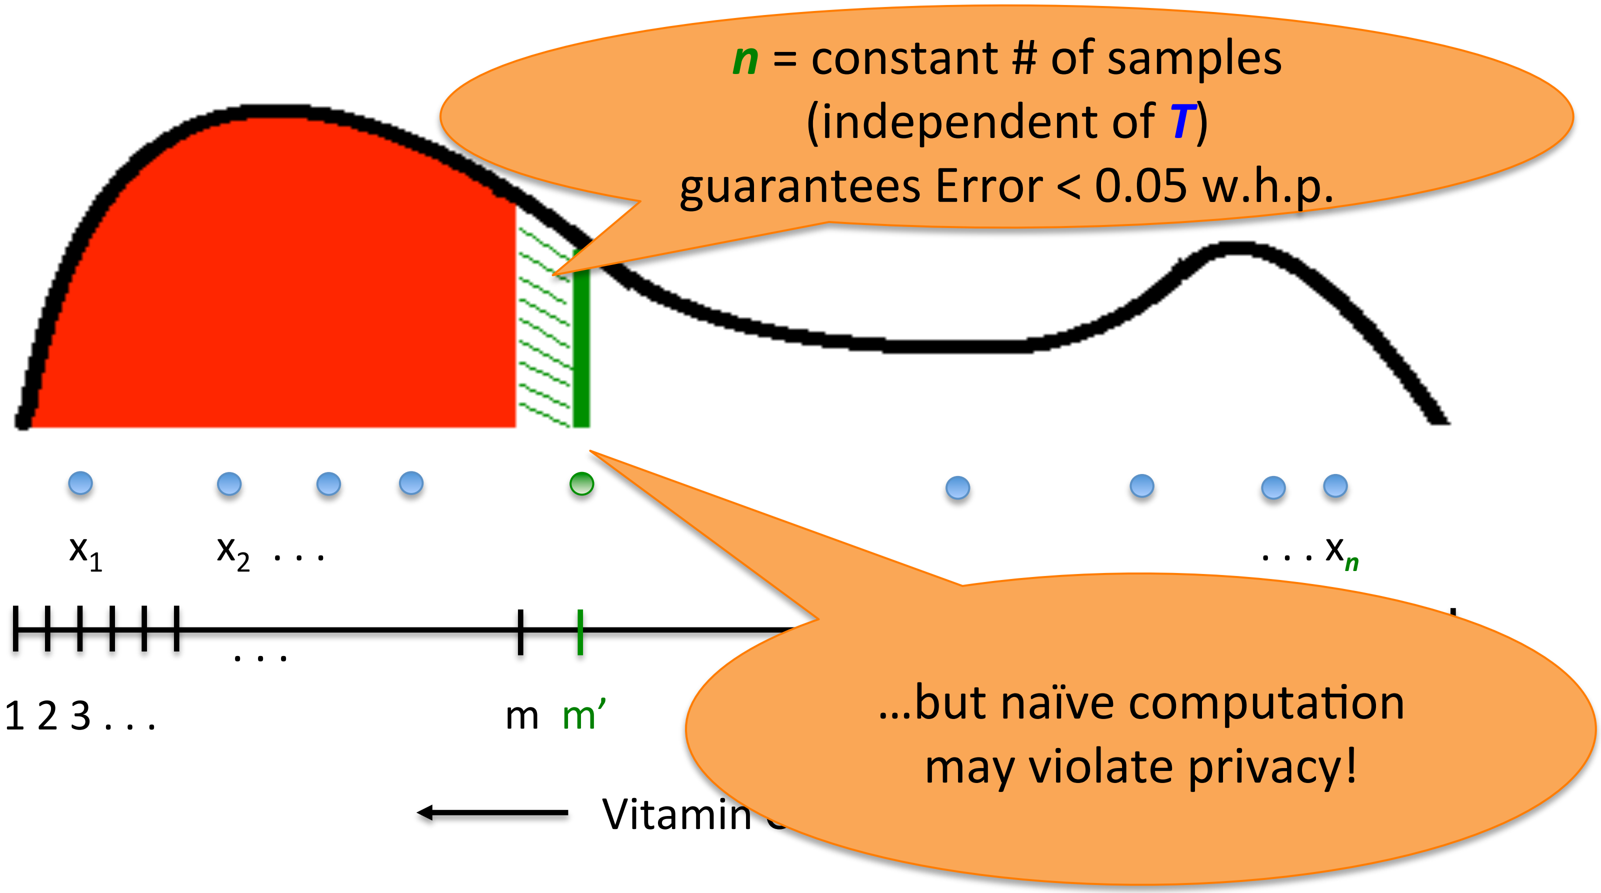

Privacy-Preserving Data Analysis

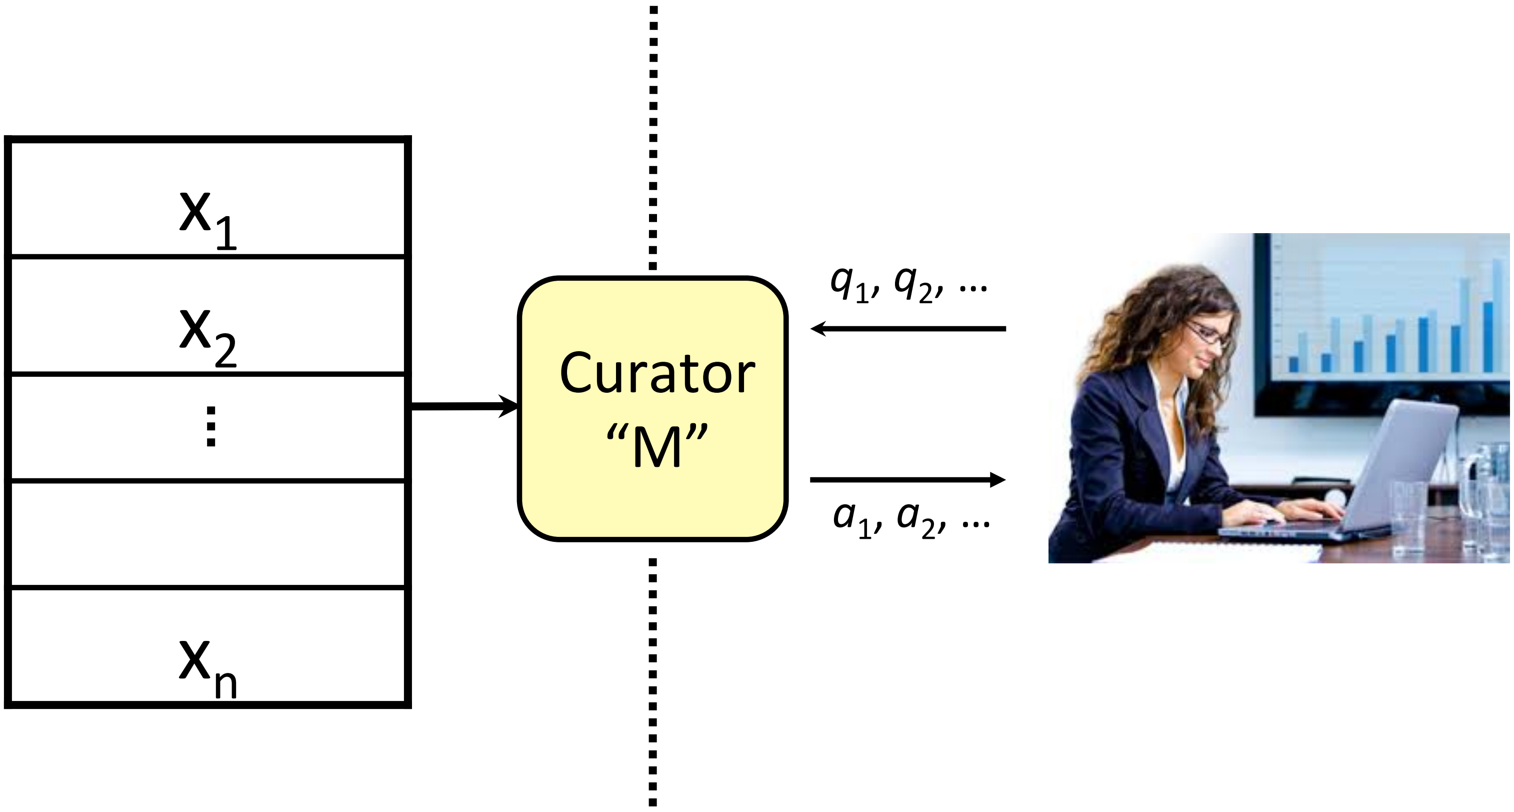

Want curators that are:

- Differentially Private
- Accurate for "Threshold" Tasks
- Sample Efficient

This Talk

- Sample complexity of threshold tasks with approx. differential privacy

- These tasks have higher sample complexity than their non-private counterparts (***n*** grows w/ ***T***)

- Network of reductions to the simpler "interior point problem"

- New combinatorial lower bound techniques
  - Distributed computing, Ramsey theory

Differential Privacy

DN03+Dwork, DN04, BDMN05,

Dwork-McSherry-Nissim-Smith06, Dwork-Kenthapadi-McSherry-Mironov-Naor06

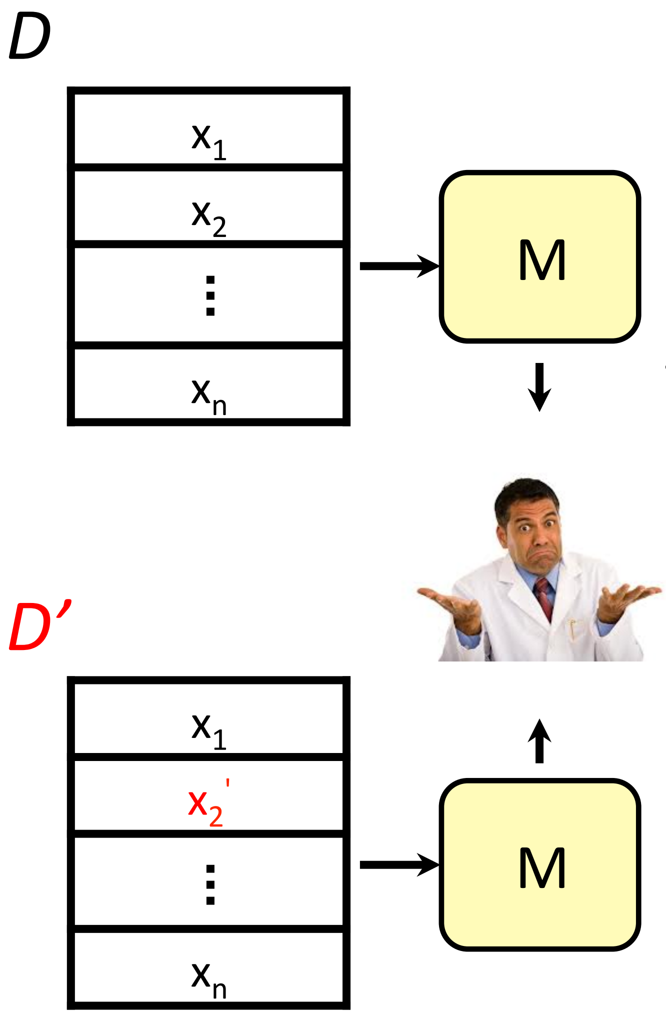

D and *D′* are **neighbors** if they differ on one row

small const., e.g. $\varepsilon = 0.1$
$e^{\varepsilon} \approx 1 + \varepsilon$

"cryptographically small"
require $\delta << 1/n$, often $\delta = \mathrm{negl}(n)$

M is **(ε,δ)-differentially private** if for all neighbors *D*, *D′* and $S \subseteq \mathrm{Range}(M)$:

$$\Pr[M(D') \in S] \leq e^{\varepsilon} \Pr[M(D) \in S] + \delta$$

◆Privacy ◆Accuracy ◆Sample Complexity

Accuracy for Approx. Medians

$\mathcal{P}$ = unknown distribution over [*T*] with median m

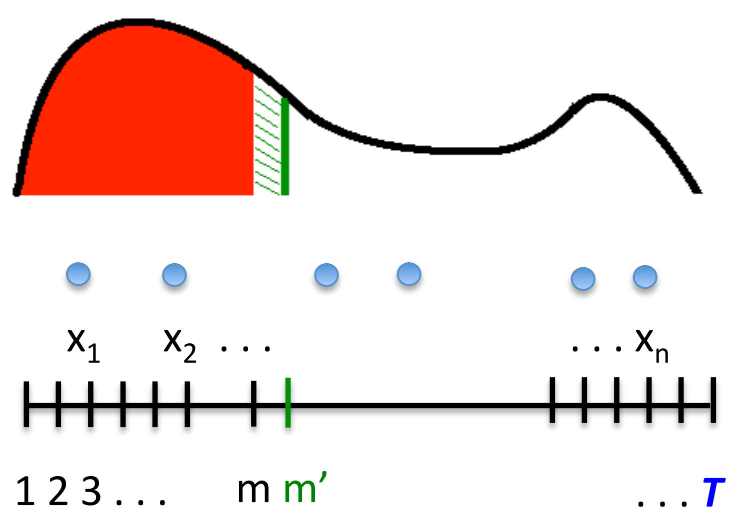

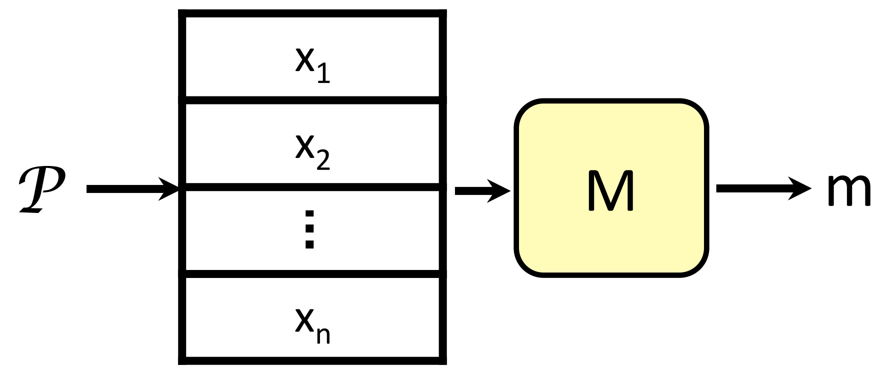

M is **accurate** if
$\Pr_{x \sim \mathcal{P}}[x \in [m, m']] < 0.05$

(w.p. 99% over sample, coins(M))

◆Privacy ◆Accuracy ◆Sample Complexity

Private Approx. Medians

[McSherry-Talwar07]

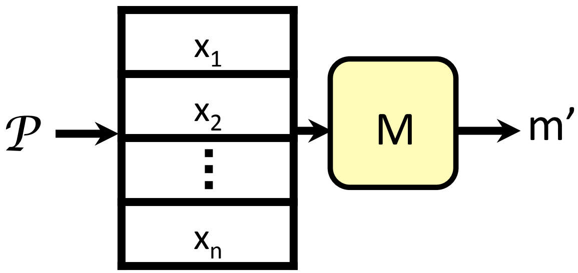

Sample m' with exponentially decaying PMF:

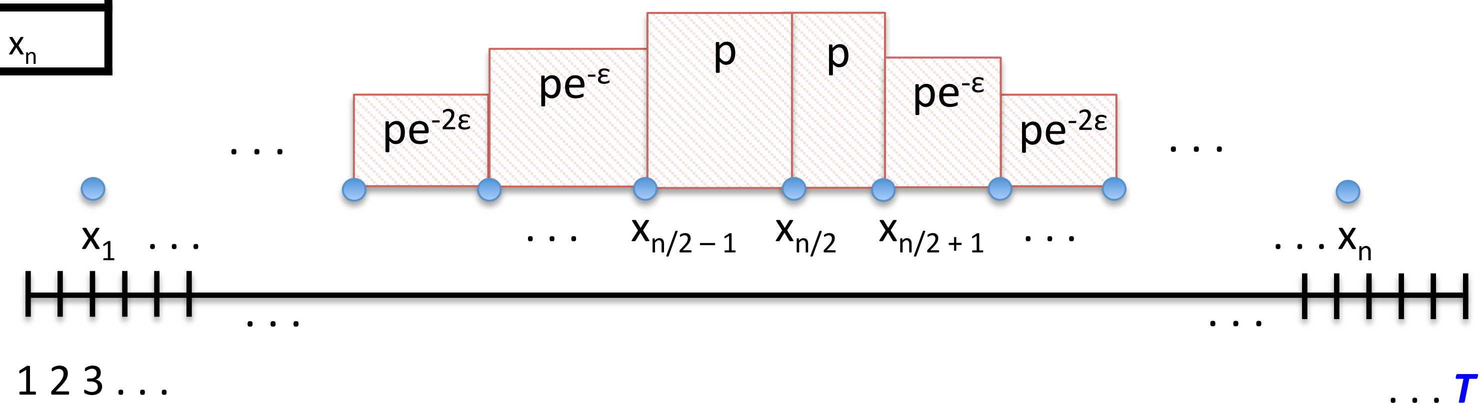

- (ε, 0)-differential privacy: Changing one person's data alters PMF by factor of e^{ε}
- Accuracy: n = O(log T) samples suffice to produce approx. median

◆Privacy ◆Accuracy ◆Sample Complexity

Accuracy for Threshold Tasks

$f_t : [\boldsymbol{T}] \rightarrow \{0,1\}$ with $f_t(x) = 1$ if $x \leq t$, $= 0$ if $x > t$

Threshold Estimation

For each $t \in [\boldsymbol{T}]$: "What fraction of dist. $\mathcal{P}$ satisfies the threshold property f_t?"

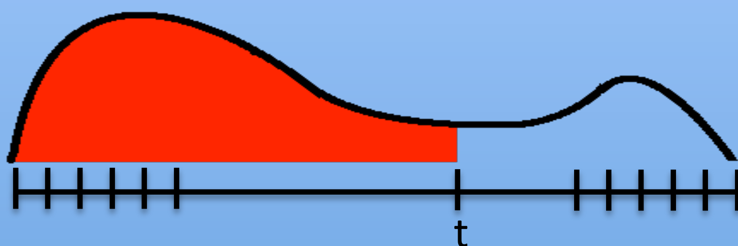

M is **accurate** if

$|a_t - f_t(\mathcal{P})| < 0.05$ for every threshold t

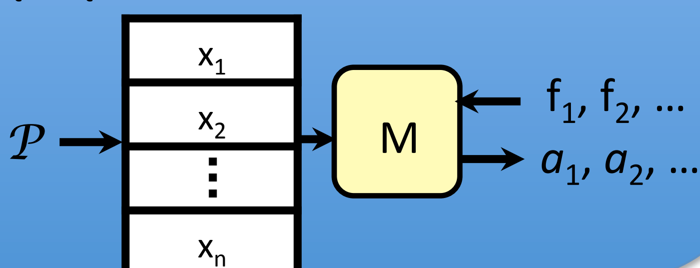

(Properly) PAC Learning Thresholds [Val84]

"What threshold function generalizes labeled examples from $\mathcal{P}$?"

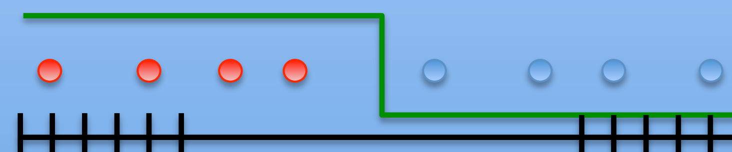

M is **accurate** if

$f_s(x) = f_t(x)$ w.p. > 0.95 over $x \sim \mathcal{P}$

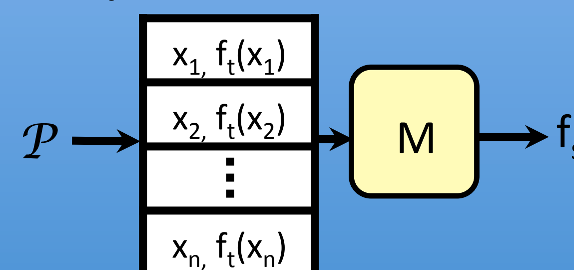

◆Privacy ◆**Accuracy** ◆Sample Complexity

Accuracy for Threshold Tasks

$f_t : [T] \rightarrow \{0,1\}$ with $f_t(x) = 1$ if $x \leq t$, $= 0$ if $x > t$

<u>Threshold Estimation</u>

= CDF Learning

"What fr… …hre…

Non-private sample complexity = O(1) [DKW56]

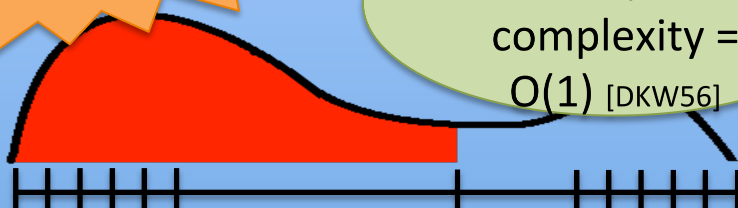

M is **accurate** if

$|a_t - f_t(\mathcal{P})| < 0.05$ for every threshold t

= Quantile Estimation

$\mathcal{P}$ → x_1, x_2, …, x_n

<u>(Properly) PAC Learning Thresholds</u> [Val84]

"What threshold fu… labeled example…

Non-private sample complexity = O(1) [Vapnik]

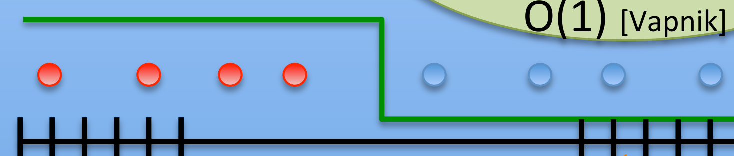

M is **accurate** if

$f_s(x) = f_t(x)$ w.p. …

= Empirical Risk Minimization

$\mathcal{P}$ → $x_n, f_t(\ldots)$

◆Privacy ◆Accuracy ◆Sample Complexity

Sample Complexity for Diff. Privacy

How big does ***n*** have to be to guarantee accuracy *and* privacy for threshold tasks?

Question: Is there an additional **price of diff. privacy** over statistical accuracy alone?

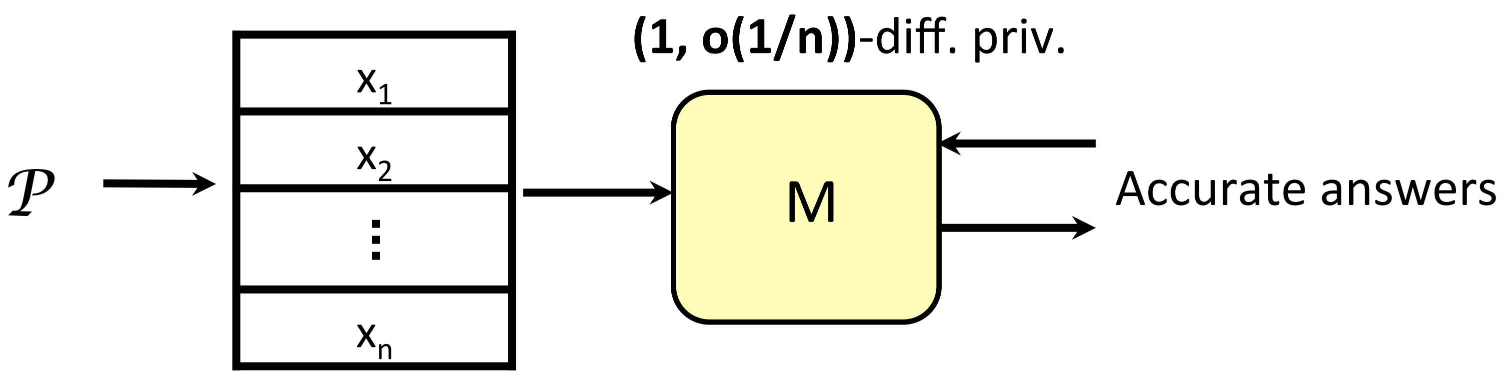

Sample Complexity for Diff. Privacy

	Approximate Medians / Releasing Thresholds	**(Proper) PAC Learning of Thresholds**
No privacy	$n = \Theta(1)$ [DKW56]	$n = \Theta(1)$ [Vapnik]
(1, o(1/n))-diff. privacy Upper bounds:	$O(T^{1/2})$ [DN03,DN04,BDMN05,DMNS06] $O(\log T)$ [BLR08,DNPR10,CSS10,DNRR15] $8^{\log^* T (1+o(1))}$ [Beimel-Nissim-Stemmer13]	$O(\log T)$ [KLNRS08] $8^{\log^* T (1+o(1))}$ [BNS13]
Lower bound: **(This work)**	$\boldsymbol{\Omega(\log^* T)}$	$\boldsymbol{\Omega(\log^* T)}$

First separation
(generalizes to higher VC-dim)

Iterated logarithm
= # of logs needed to reach 1

This work: Plus somewhat improved upper bounds

- Privacy
- Accuracy
- Sample Complexity

Lower Bounds for (ε, 0)-Diff. Priv.

Volume-based ("packing") arguments

- Tight characterization of (ε, 0)-DP [Hardt-Talwar10, Beimel-Nissim-Stemmer13a]
- Break down even for δ=negl(***n***) [De11, Beimel-Nissim-Stemmer13b]

Lower bounds via info. theory & comm. complexity

- LBs for two-party privacy problems
  [McGregor-Mironov-Pitassi-Reingold-Talwar-Vadhan10]
- Characterization of (ε, 0)-DP learning [Feldman-Xiao14]

Lower Bounds for (ε, δ)-Diff. Priv.

Reconstruction attacks [Dinur-Nissim03]

- Connection to sparse recovery [Dwork-McSherry-Talwar07]
- Combinatorial (hereditary) discrepancy [Muthukrishnan-Nikolov12, Nikolov-Talwar-Zhang13, Nikolov15]

Probabilistic fingerprinting codes [Boneh-Shaw95, Tardos03]

- LBs for contingency tables [B.-Ullman-Vadhan14, Steinke-Ullman15]
- LBs for convex optimization, PCA [Bassily-Smith-Thakurta14, Dwork-Talwar-Thakurta-Zhang14]

Lower Bounds for (ε, δ)-Diff. Priv.

Prior techniques for (ε, δ)-DP exploit **high dimensionality** of concepts/data

This work: Lower bounds for (ε, δ)-DP even for **simple** concepts (i.e. VC-dimension = 1)

Techniques

- Equivalence between threshold tasks and the "Interior Point Problem"

- New upper and lower bounds for solving IPP with approx. differential privacy

$$\log^* T \leq n \leq 2^{\log^* T}$$

Interior Point Problem

- Input: Database $D = (x_1, \ldots, x_n) \in [T]^n$
- Output: Any $p \in [T]$ with $\min_i x_i \leq p \leq \max_i x_i$

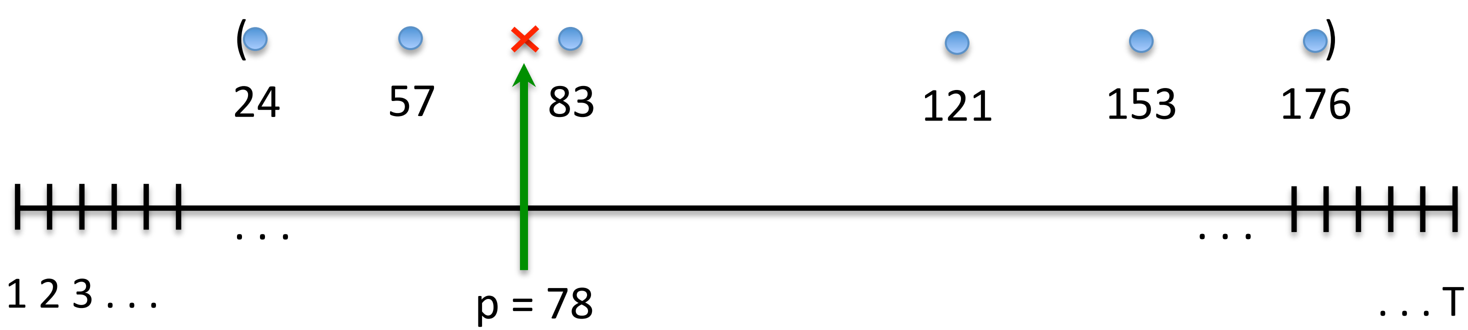

Want (ε, δ)-diff. privacy + success w.p. 2/3

General Reductions

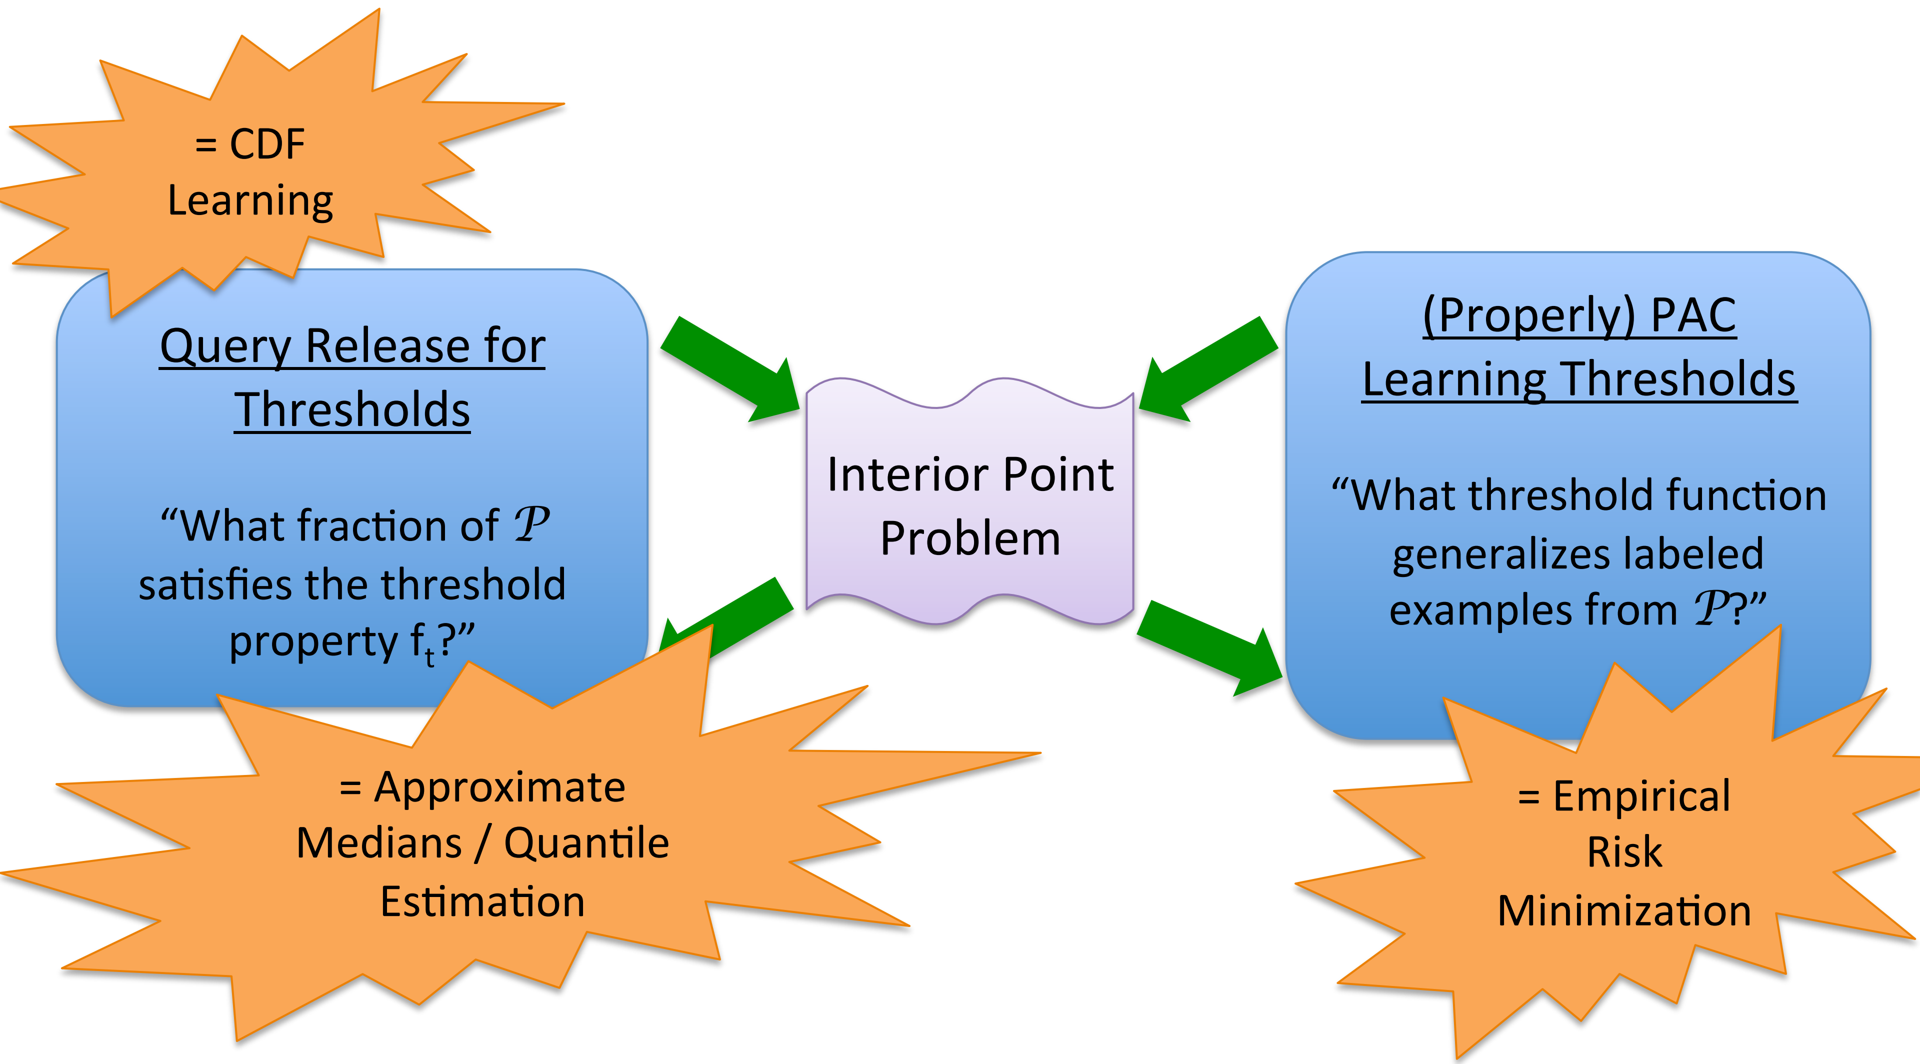

Results for Interior Point

- <u>Lower bound:</u> Sample complexity of IPP is

$$n \geq \Omega(\log^* T)$$

- <u>Upper bound:</u> Sample complexity of IPP is

$$n \leq 2^{\log^* T (1+o(1))}$$

  - ➢ Simpler algorithm inspired by lower bound construction
  - ➢ Better dependence on error in applications

A Detour in Distributed Computing

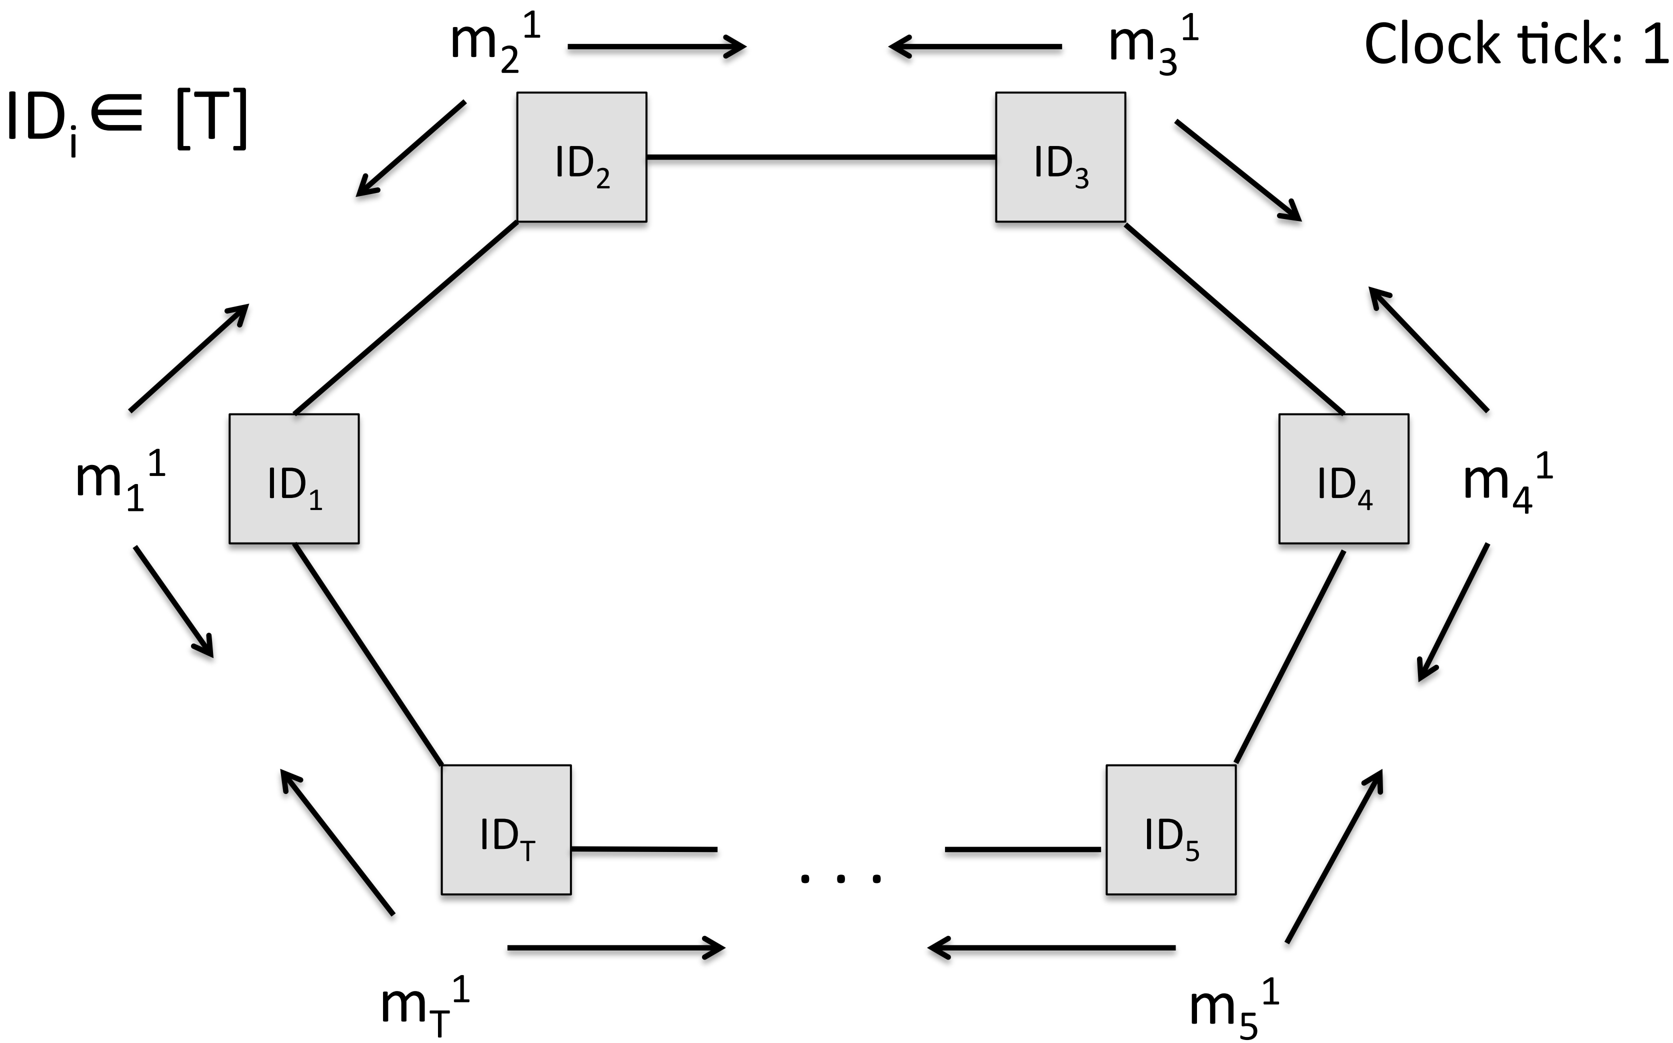

A Detour in Distributed Computing

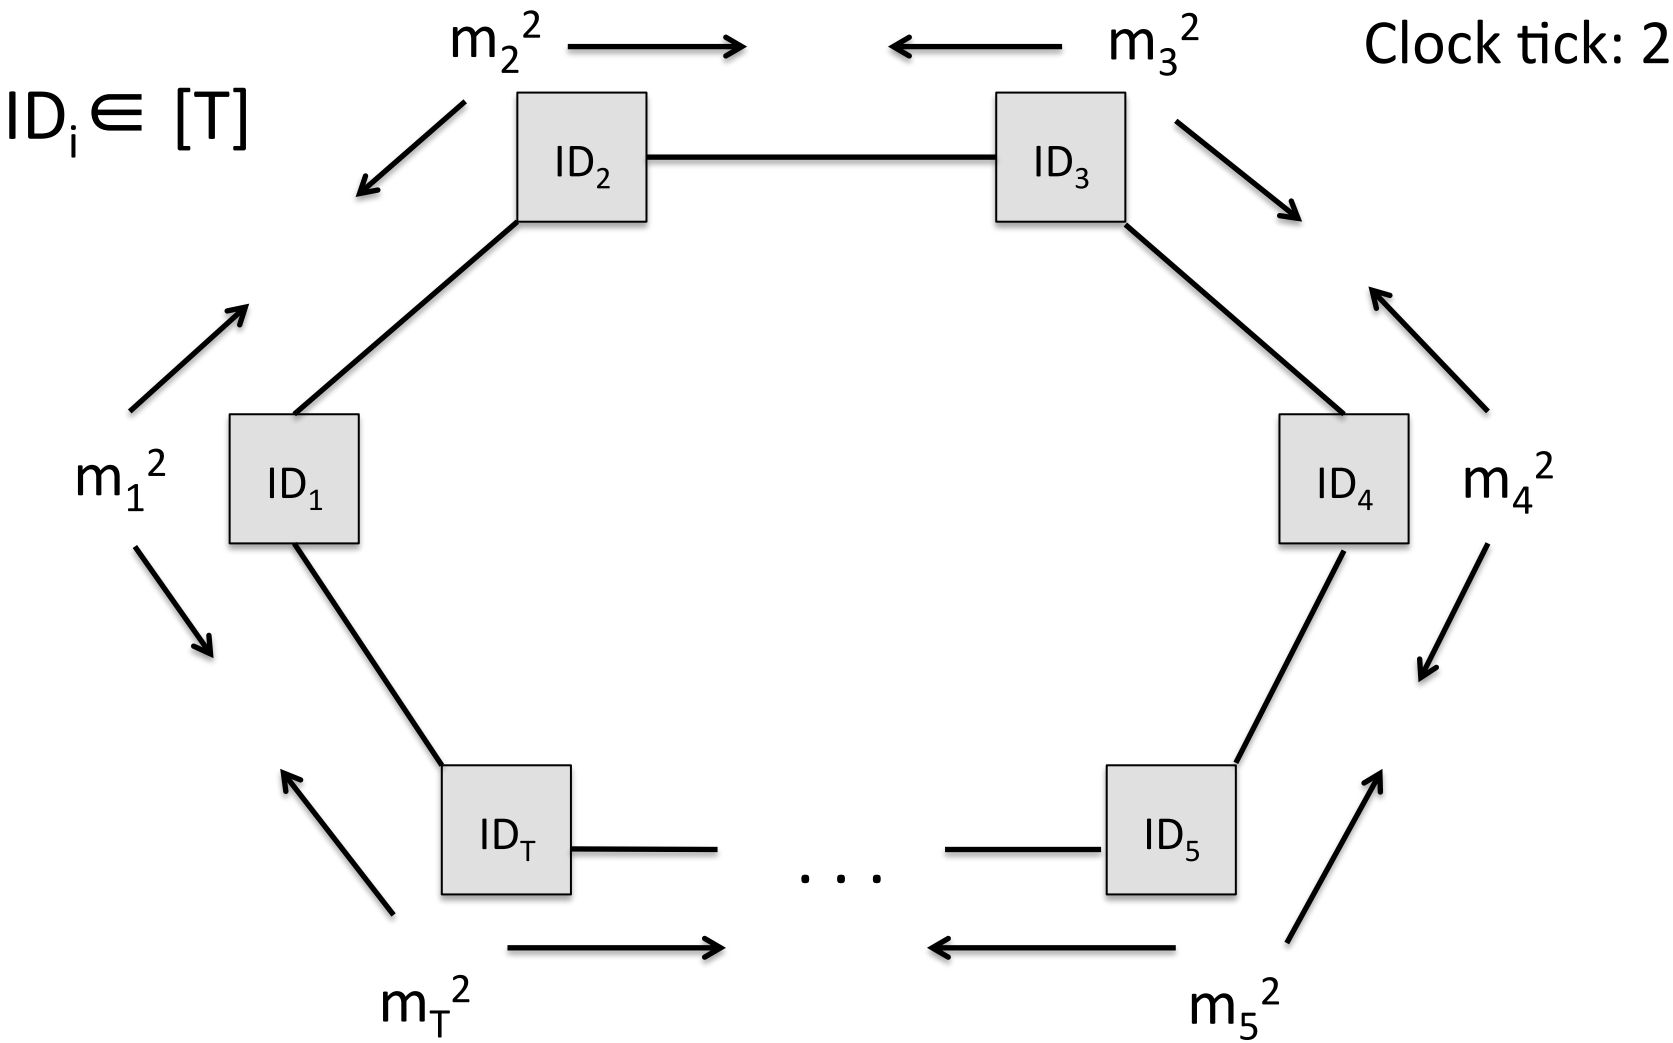

A Detour in Distributed Computing

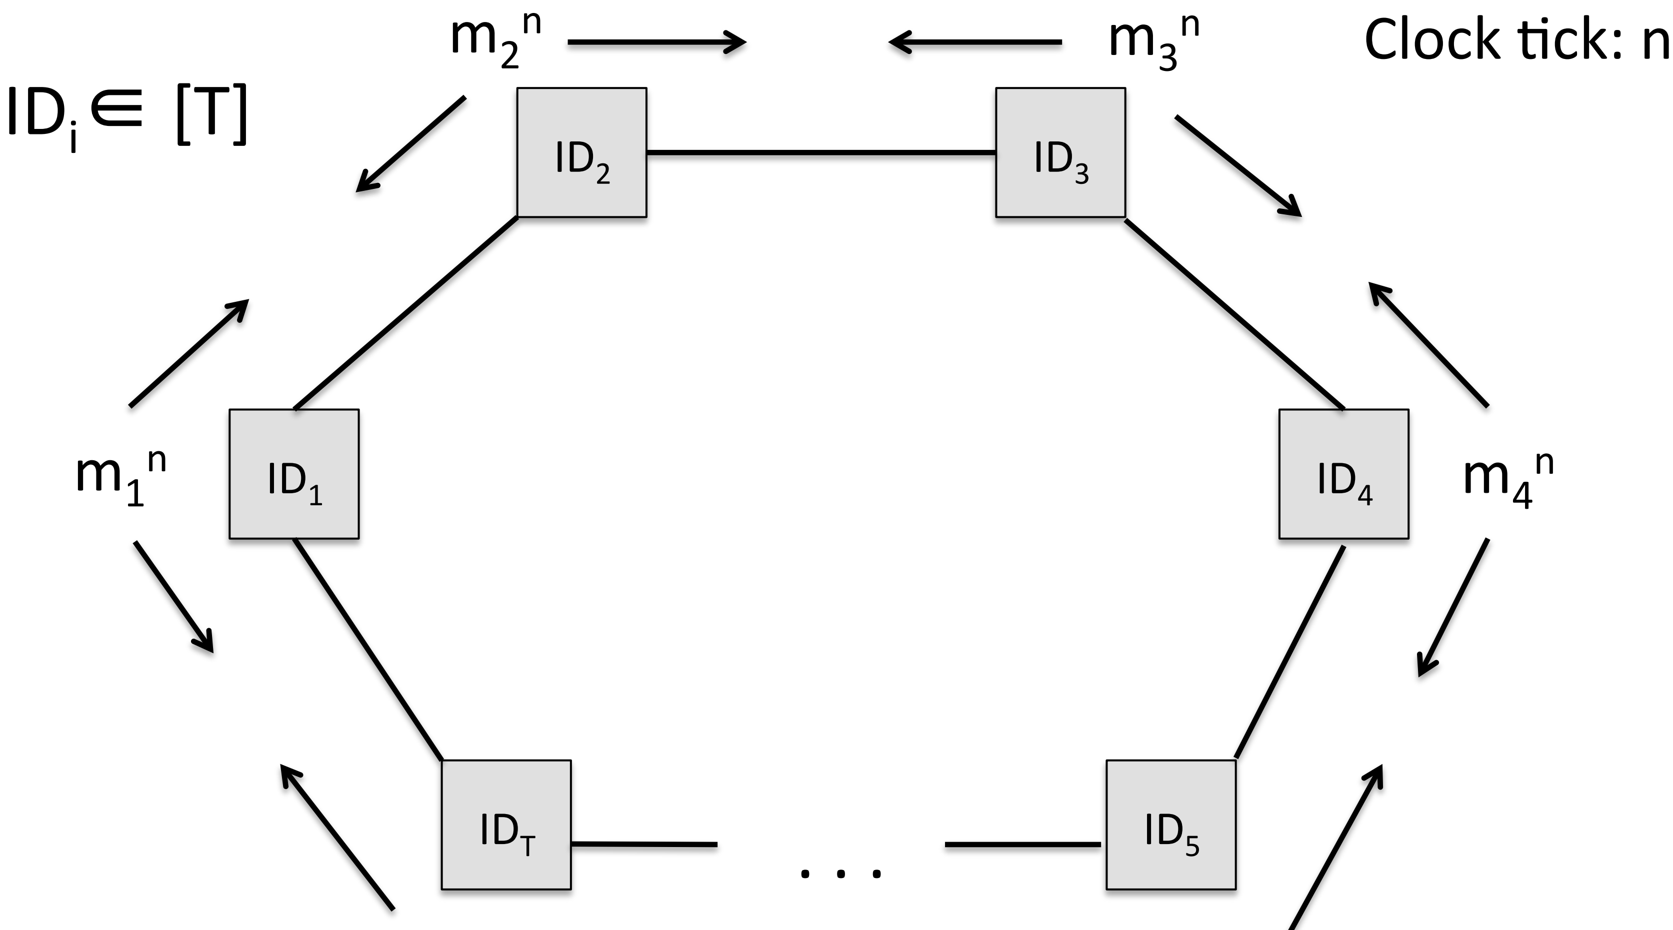

<u>Question:</u> How many clock ticks are needed for the processors to agree on a 3-coloring?

A Detour in Distributed Computing

<u>Question:</u> How many clock ticks are needed for the processors to agree on a 3-coloring?

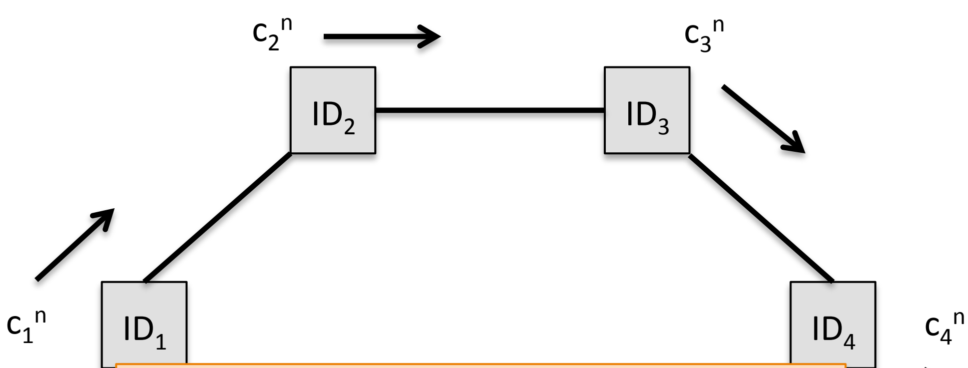

$|c_i^{n+1}| \approx \log |c_i^n|$

$\Rightarrow$Reach constant # colors after n = O(log* T) rounds

[Cole-Vishkin86]:
O(log* T) rounds suffices

$c_i^1 = ID_i \in [T]$

<u>After round n:</u>
j = first index where c_i^n disagrees w/ c_{i-1}^n

$c_i^{n+1} = j \;||\; (c_i^n)_j$

A Detour in Distributed Computing

<u>Question:</u> How many clock ticks are needed for the processors to agree on a 3-coloring?

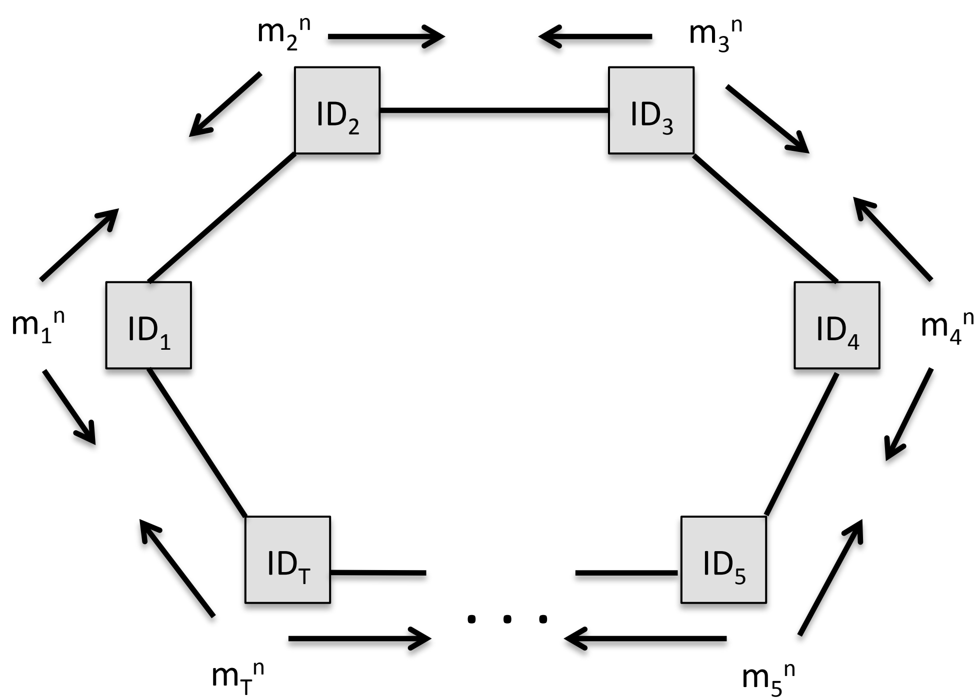

[Linial92]:
$\Omega(\log^* T)$ rounds required

<u>Key observation:</u>
Processor i's information after round n is $(ID_{i-n}, ID_{i-n+1}, \ldots, ID_{i+n})$

<=> Existence of a coloring
$C: \binom{[T]}{2n+1} \to \{1,2,3\}$

Choose Your Own Adventure

Two proofs of the interior point lower bound:

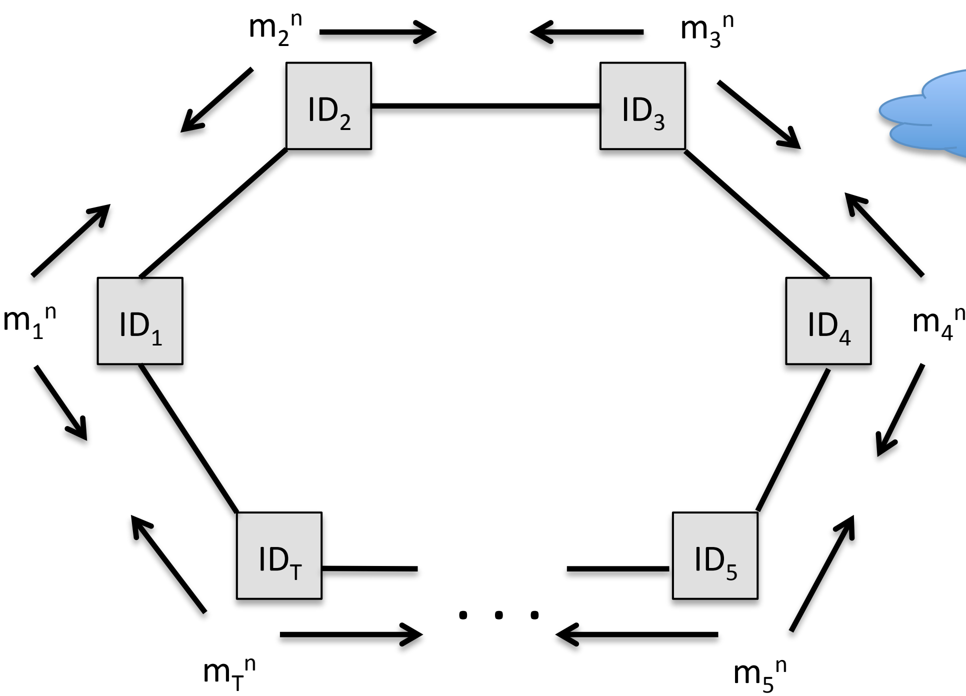

[Cole-Vishkin86]:

O(log* T) rounds suffices

Thanks to Avi Wigderson

IPP Proof 1

[Linial92]:

Ω(log* T) rounds required

IPP Proof 2

Choose Your Own Adventure

Two proofs of the interior point lower bound:

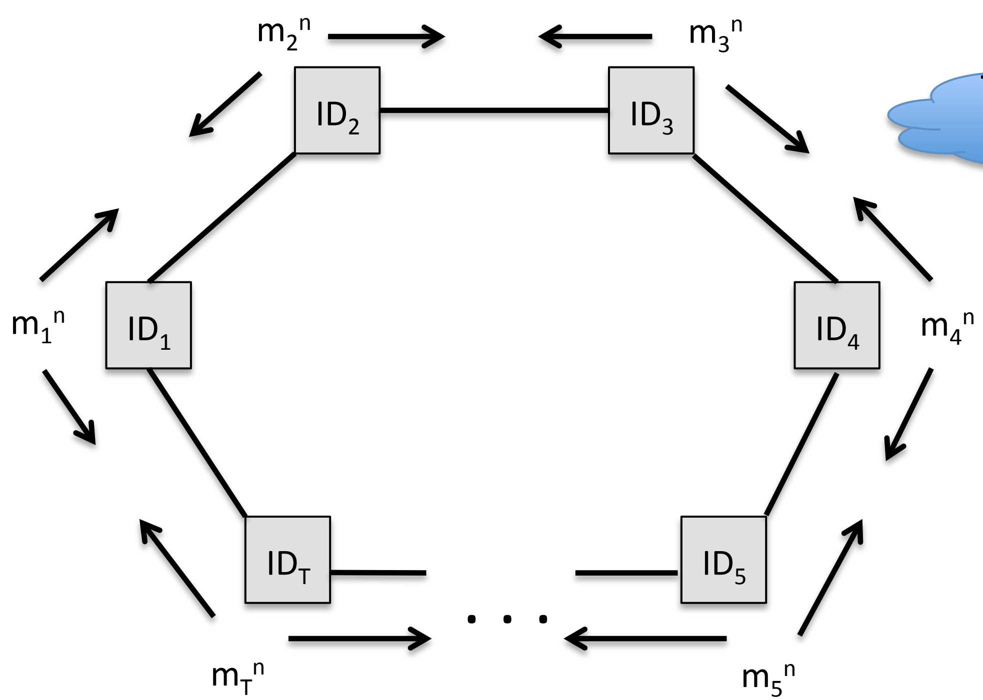

[Cole-Vishkin86]:
O(log* T) rounds suffices

Thanks to Avi Wigderson

IPP Proof 1

go to slide 37 ↵

[Linial92]:
Ω(log* T) rounds required

IPP Proof 2

go to slide 27 ↵

Ramsey's Theorem

"Sufficiently large objects must necessarily contain a given structure" --Wikipedia

or: How to get really big numbers to appear in your proofs

Baby version: "Theorem on Friends and Strangers"

Any group of 6 people contains either:

3 mutual friends or
3 mutual strangers

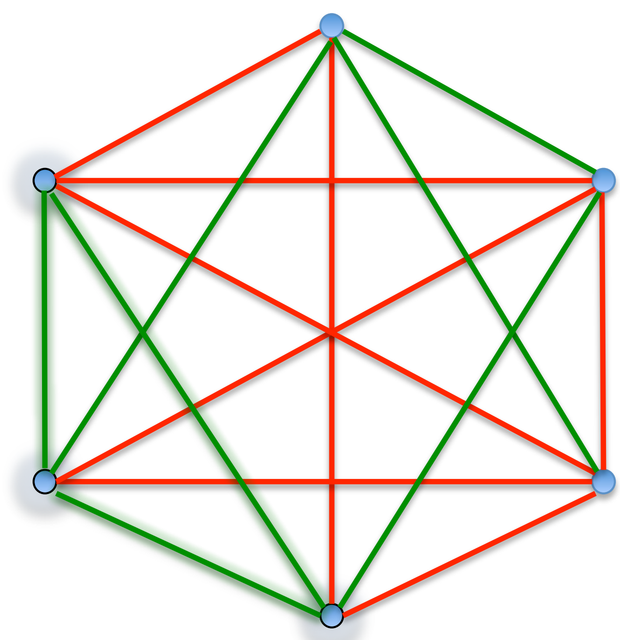

Ramsey's Theorem

- Ground set $[T]$
- Coloring function $C : \binom{[T]}{n} \rightarrow [K]$

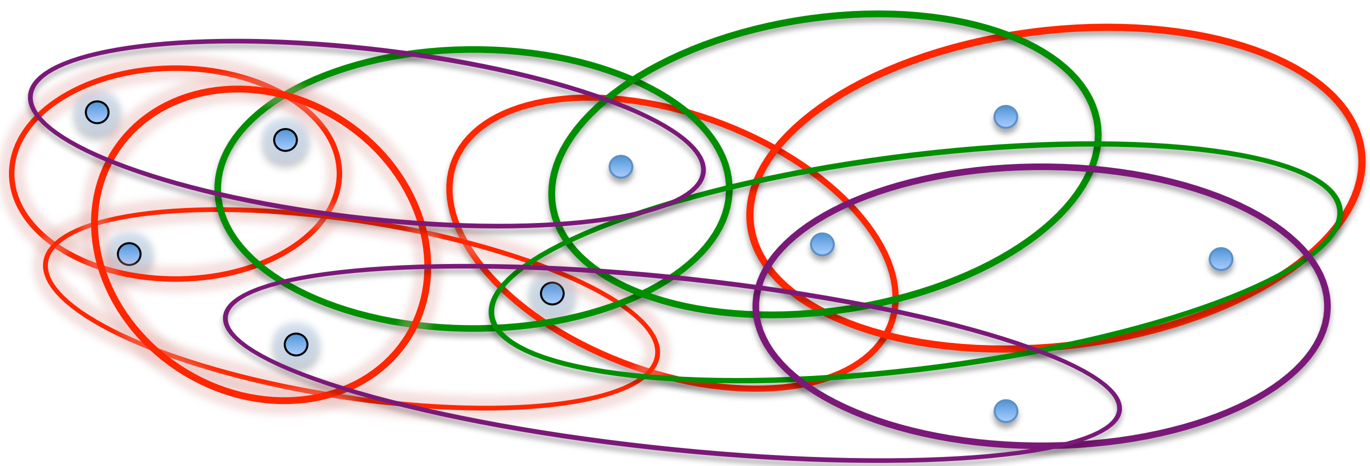

Thm: For $T > R(n, m, K)$, there exists a monochromatic S of size m (i.e. C is constant on $\binom{S}{n}$)

Ramsey vs. Interior Point

Ramsey's Thm		Interior Point Problem
Ground set	$[\boldsymbol{T}]$	Data domain
Hyperedge	$\{x_1, \ldots, x_{\boldsymbol{n}}\} \in \binom{[\boldsymbol{T}]}{\boldsymbol{n}}$	Database
Coloring function	$C_M \leftrightarrow M$	DP Mechanism

Claim: M solves IPP $\Rightarrow$ $\exists$ coloring $C_M : \binom{[\boldsymbol{T}]}{\boldsymbol{n}} \to [\boldsymbol{n}]$

with no size-($3\boldsymbol{n}$) monochromatic set

$\therefore$ By Ramsey, $\boldsymbol{T} < R(\boldsymbol{n}, m=3\boldsymbol{n}, K=\boldsymbol{n})$ $\quad (=\text{tower}(\boldsymbol{n}))$

[Erdős-Rado52]

Defining a Coloring

Write $D \in \binom{[T]}{n}$ as $\{x_1 < x_2 \ldots < x_n\}$

Define $C_M(D) = \text{argmax}_k \Pr[M(D) \in [x_k, x_{k+1})]$

E.g. $C_M(D) = 3$

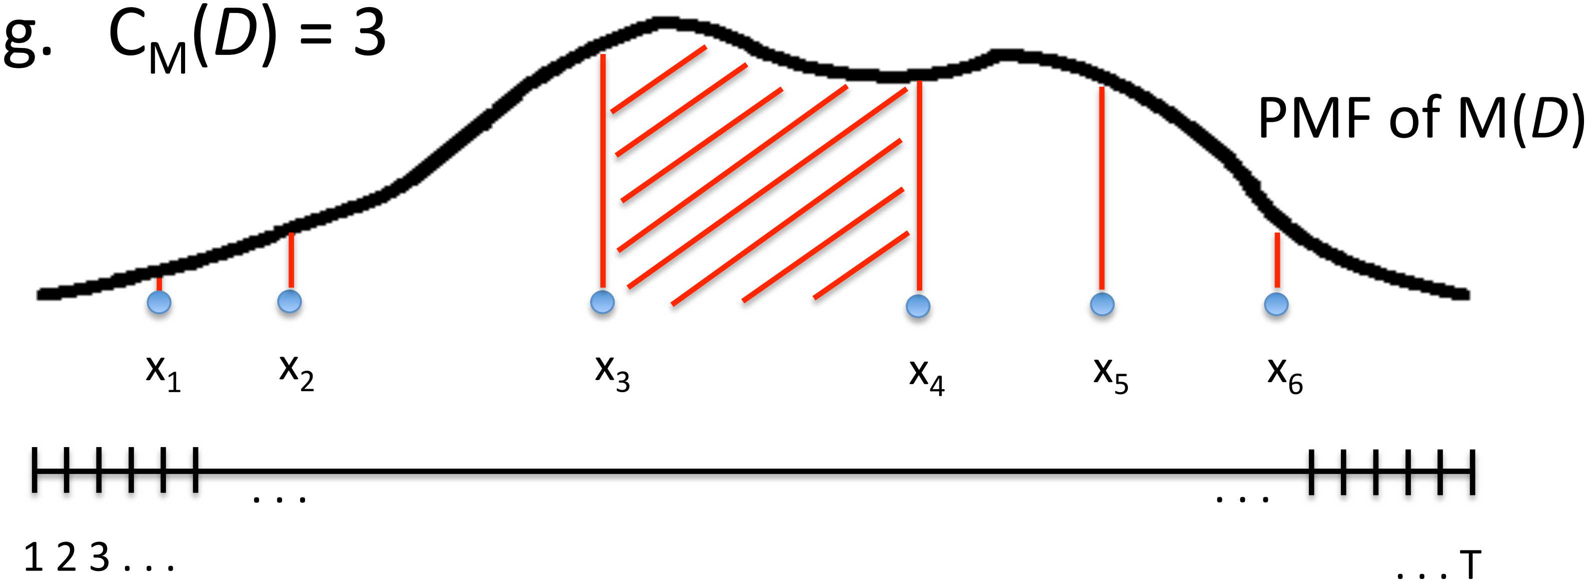

Let $S = \{x_0 < x_1 < \ldots < x_{m+1}\}$ (recall $m \approx 3\boldsymbol{n}$)

Suppose (for contradiction): $C_M(D) = k \quad \forall D \in \binom{[T]}{n}$

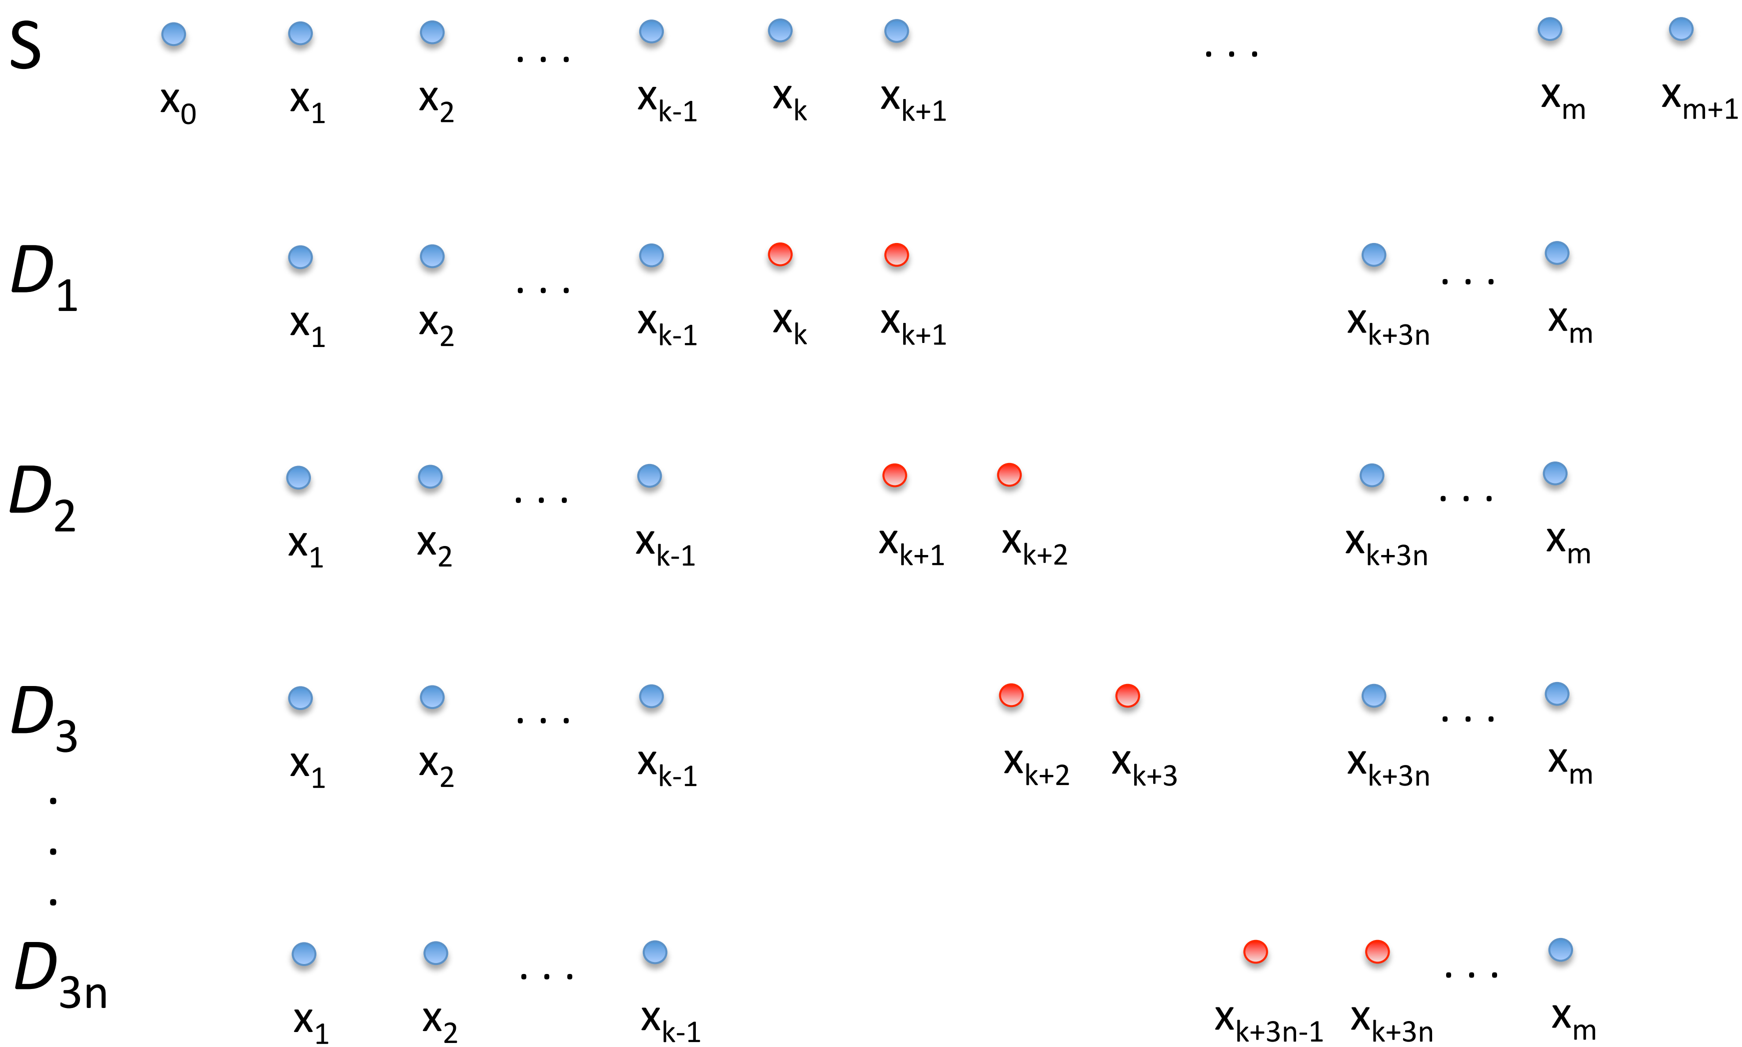

Let $S = \{x_0 < x_1 < \dots < x_{m+1}\}$ (recall $m \approx 3n$)

Suppose (for contradiction): $C_M(D) = k \quad \forall D \in \binom{[T]}{n}$

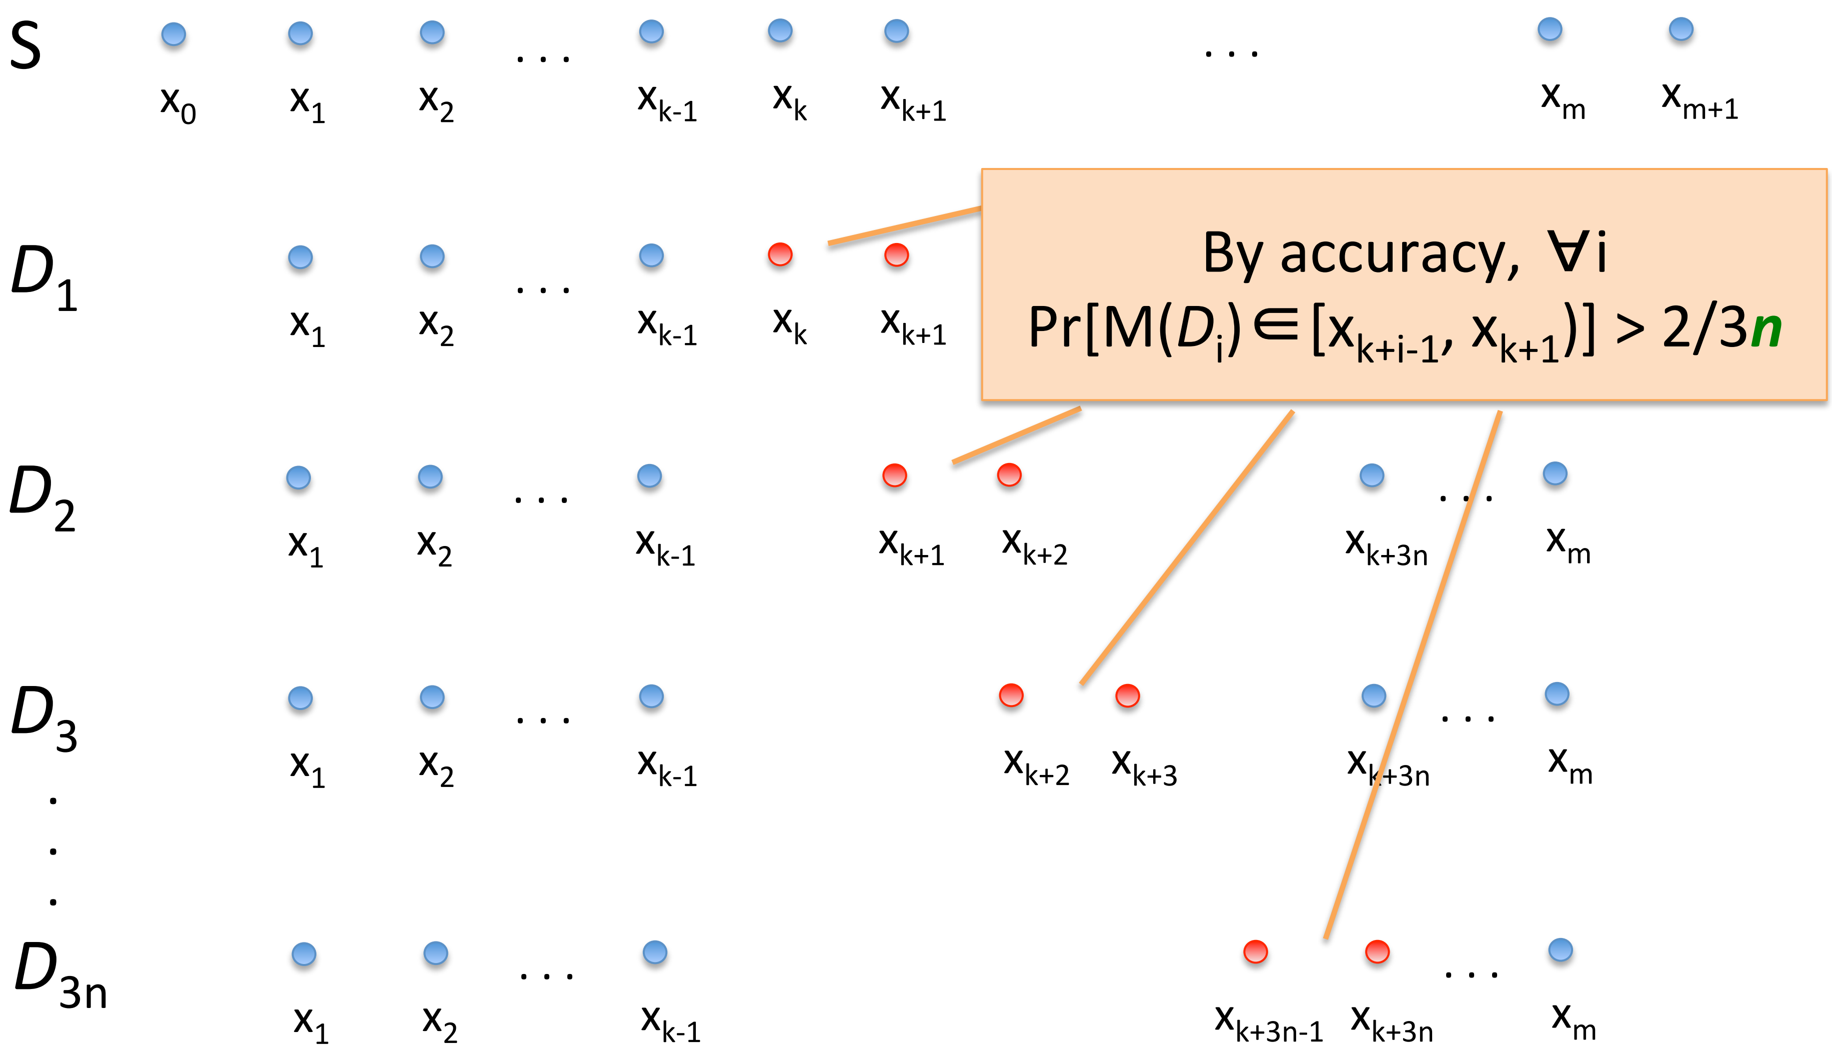

By privacy, $\forall i$

$\Pr[M(D^*) \in [x_{k+i-1}, x_{k+1})] \;>\; e^{-2\varepsilon}(2/3\boldsymbol{n}) - 2\delta \;>\; 1/3\boldsymbol{n}$

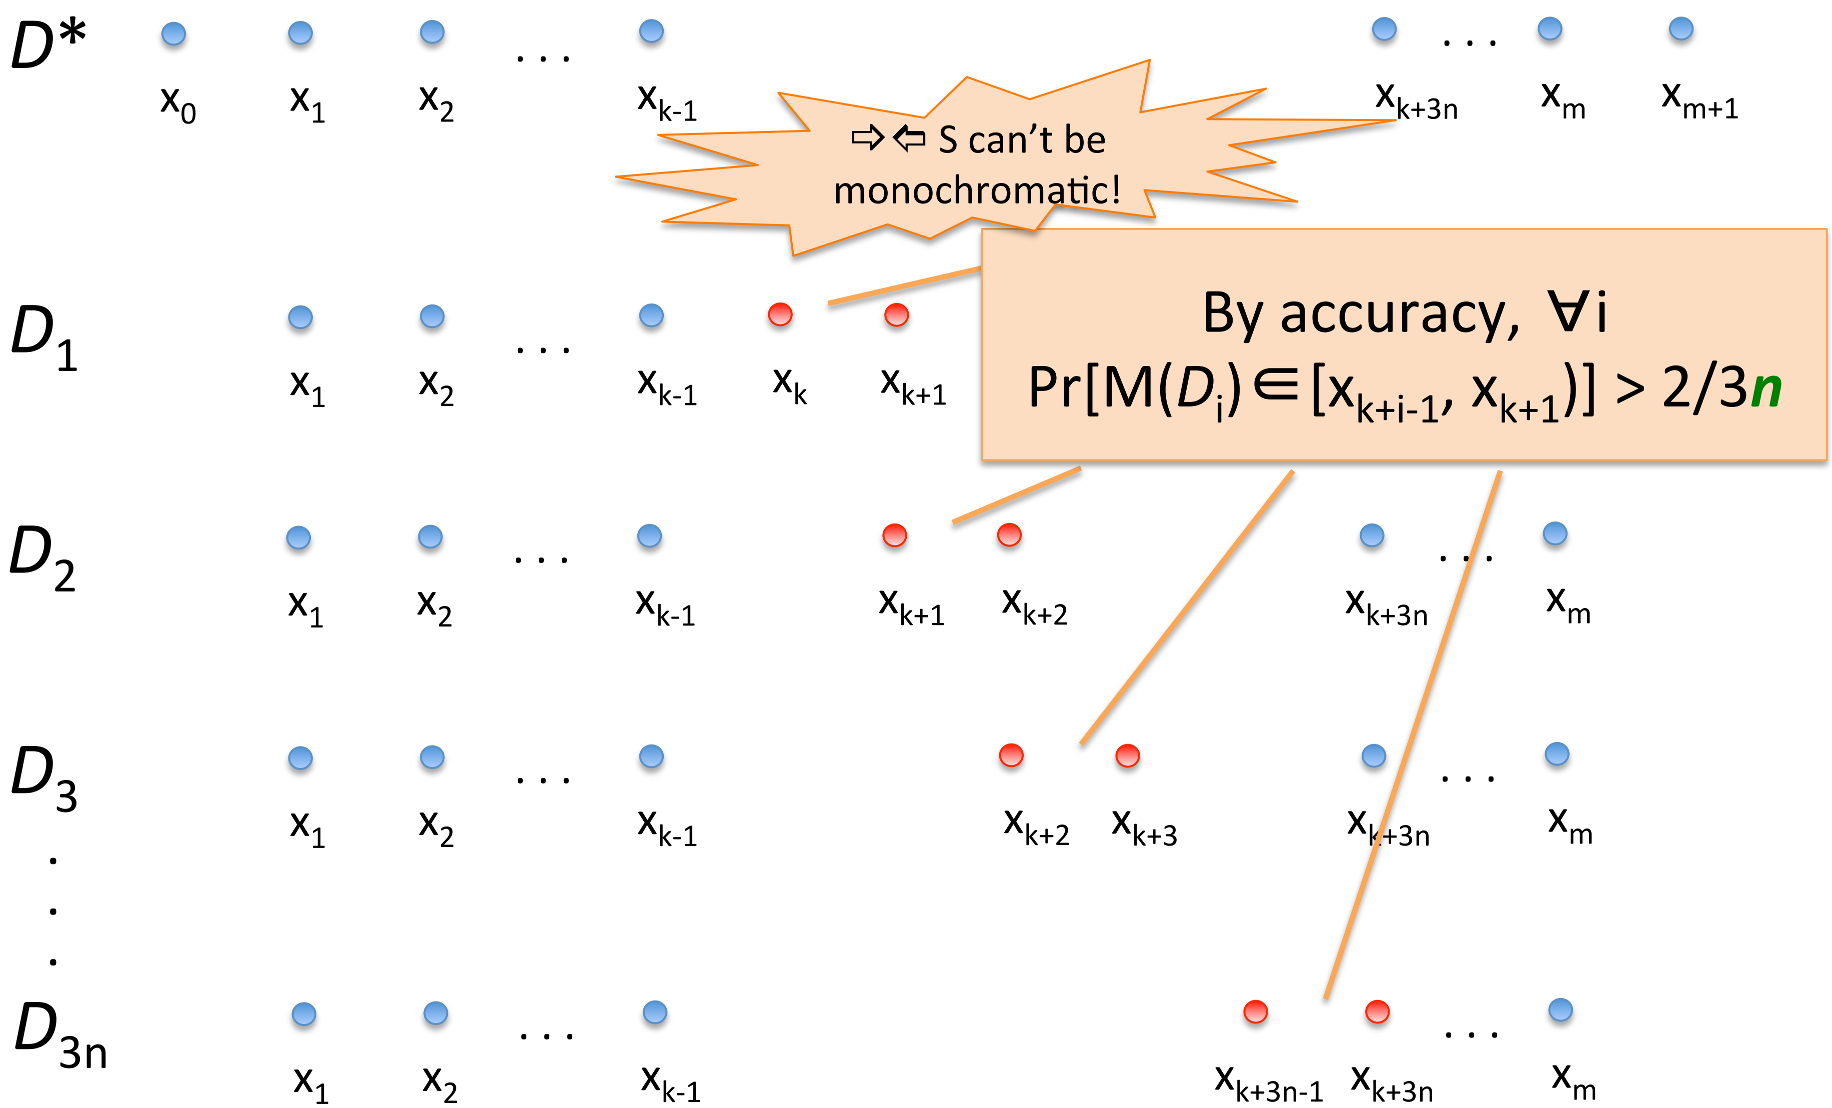

Proof Recap

M privately solves IPP on $[\boldsymbol{T}]$

$\Rightarrow$ $\exists$ coloring $C_M : \binom{[\boldsymbol{T}]}{\boldsymbol{n}} \rightarrow [\boldsymbol{n}]$ with no size-$(3\boldsymbol{n})$ monochromatic set

$\Rightarrow$ $\boldsymbol{T} < R(\boldsymbol{n}, m=3\boldsymbol{n}, K=\boldsymbol{n}) = \text{tower}(\boldsymbol{n})$ by Ramsey

$\Rightarrow$ $\boldsymbol{n} > \Omega(\log^* \boldsymbol{T})$

Conclusions

- Diff. privacy-preserving **reductions** between threshold tasks
- Price of (ε, δ)-diff. privacy for **simple statistics**
- Open questions:
  - Combinatorial characterization of sample complexity?
    [e.g. HT10, Har11, NTZ13, BNS13]
  - Sample complexity of *improper* PAC learning?
    [e.g. BKN10, FX14]

Thank you!

SUPPLEMENTARY CONTENT

Interior Point Lower Bound

- Recursively construct hard distributions $\mathcal{P}_n$ on domain size $T(n) \approx \text{tower}(n)$ $\Rightarrow$ $n \geq \log^* T$
- Base case: For $n = 1$, set $T(1) = 2$

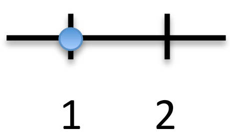

Output 1 w.p. ≥ 2/3

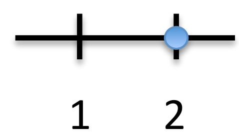

Output 1 w.p. $\geq \frac{(2/3)-\delta}{e^{\varepsilon}} > \frac{1}{3}$

- Inductive case:

Suppose M solves IPP on $\mathcal{P}_{n+1}$ over domain $[T(n+1)]$

$\Rightarrow$ construct M' for IPP on $\mathcal{P}_n$ over $[T(n)]$

Interior Point Lower Bound

To sample D_{n+1} from $\mathcal{P}_{n+1}$:

1. Sample $D_n = (x_1, \ldots, x_n)$ from $\mathcal{P}_n$
2. Sample $y_0 \in [b^{T(n)}]$ at random
3. For i = 1, …, n, sample y_i that agrees with y_0 up to base b-“digit” x_i

cf. Cole-Vishkin86

D_n =

x_1
x_2
x_3
⋮
x_n

→

D_{n+1} =

y_0
y_1
y_2
y_3
⋮
y_n

Interior Point Lower Bound

To sample D_{n+1} from $\mathcal{P}_{n+1}$:

1. Sample $D_n = (x_1, \ldots, x_n)$ from $\mathcal{P}_n$
2. Sample $y_0 \in [b^{T(n)}]$ at random
3. For i = 1, …, n, sample y_i that agrees with y_0 up to base b-“digit” x_i

cf. Cole-Vishkin86

E.g. $b = 10$

$D_n =$

$x_1 = 3$
$x_2 = 5$
$x_3 = 3$
$\vdots$
$x_n = 4$

$D_{n+1} =$

$y_0 = 8675309$
$y_1 = 8674812$
$y_2 = 8675365$
$y_3 = 8671863$
$\vdots$
$y_n = 8675052$

Interior Point Lower Bound

$\mathcal{P}_{n+1}$ is a hard distribution by contradiction:

M solves IPP on $\mathcal{P}_{n+1}$

$\Rightarrow$ M' solves IPP on $\mathcal{P}_n$

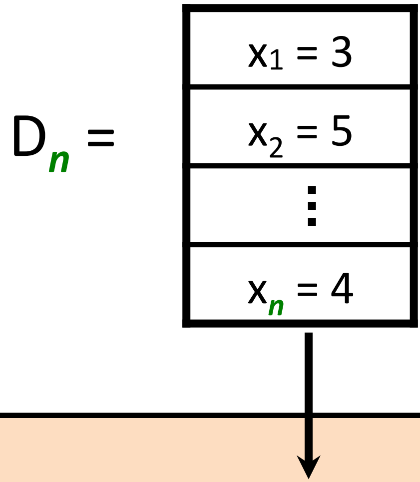

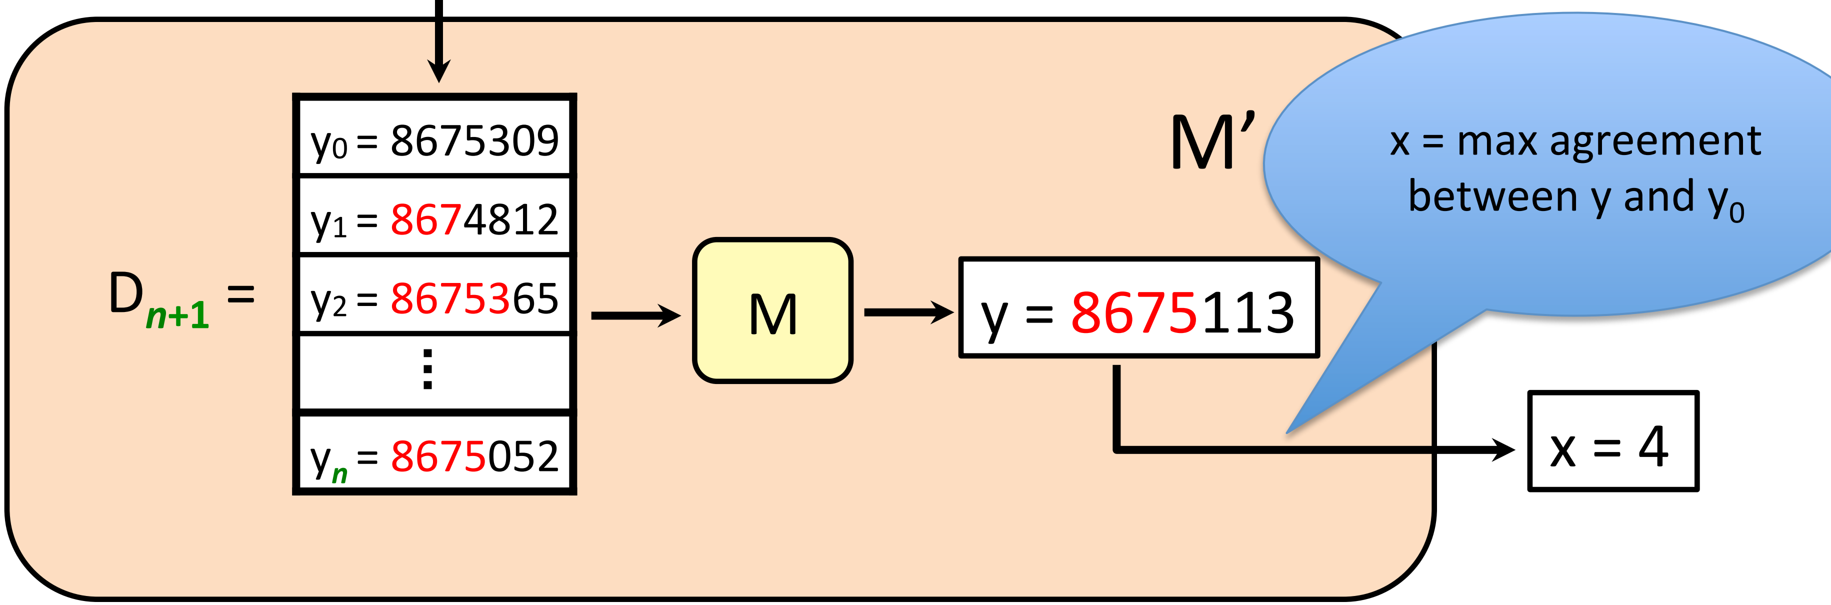

Interior Point Lower Bound

If M succeeds,

- $x \geq \min_i x_i$ by construction
- $x \leq \max_i x_i$ by privacy (whp)

$\Rightarrow$ M' succeeds

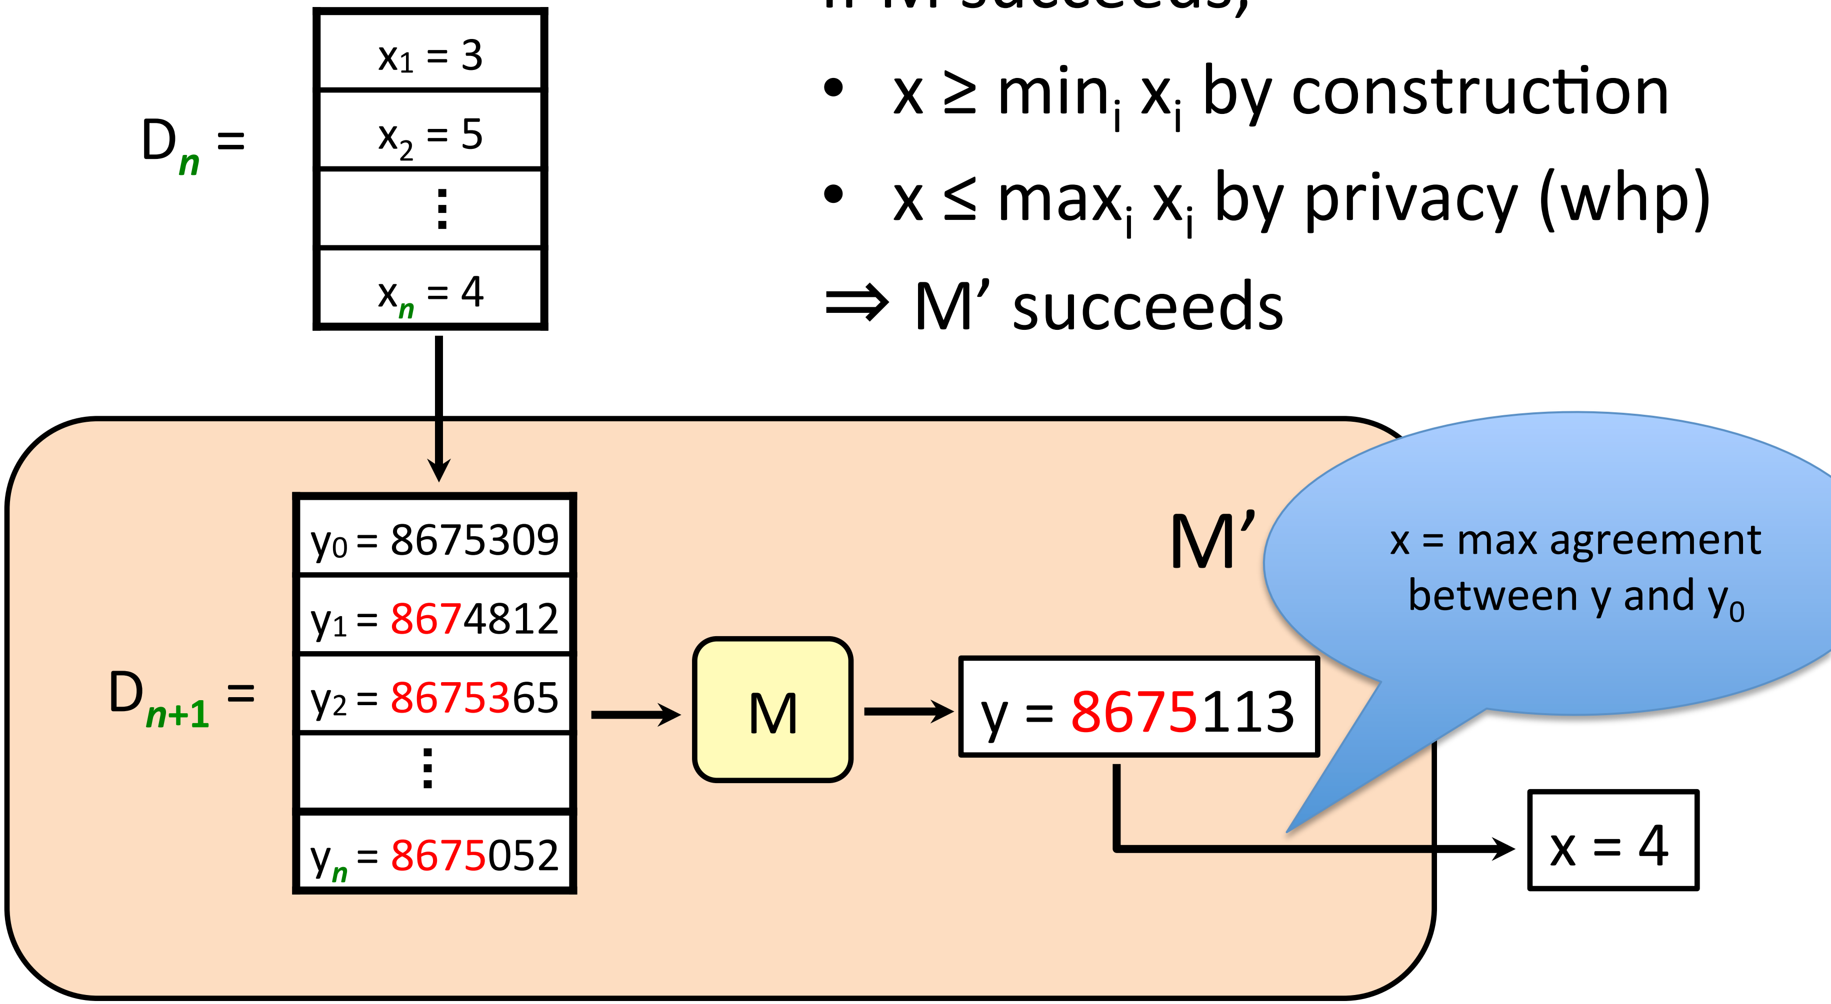

Interior Point Lower Bound

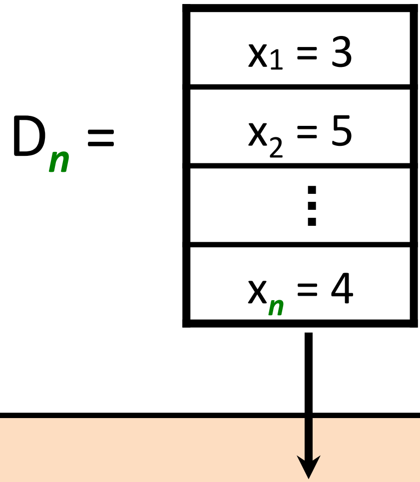

If M succeeds,

- $x \geq \min_i x_i$ by construction
- <u>$x \leq \max_i x_i$ by privacy (whp)</u>

$\Rightarrow$ M' succeeds

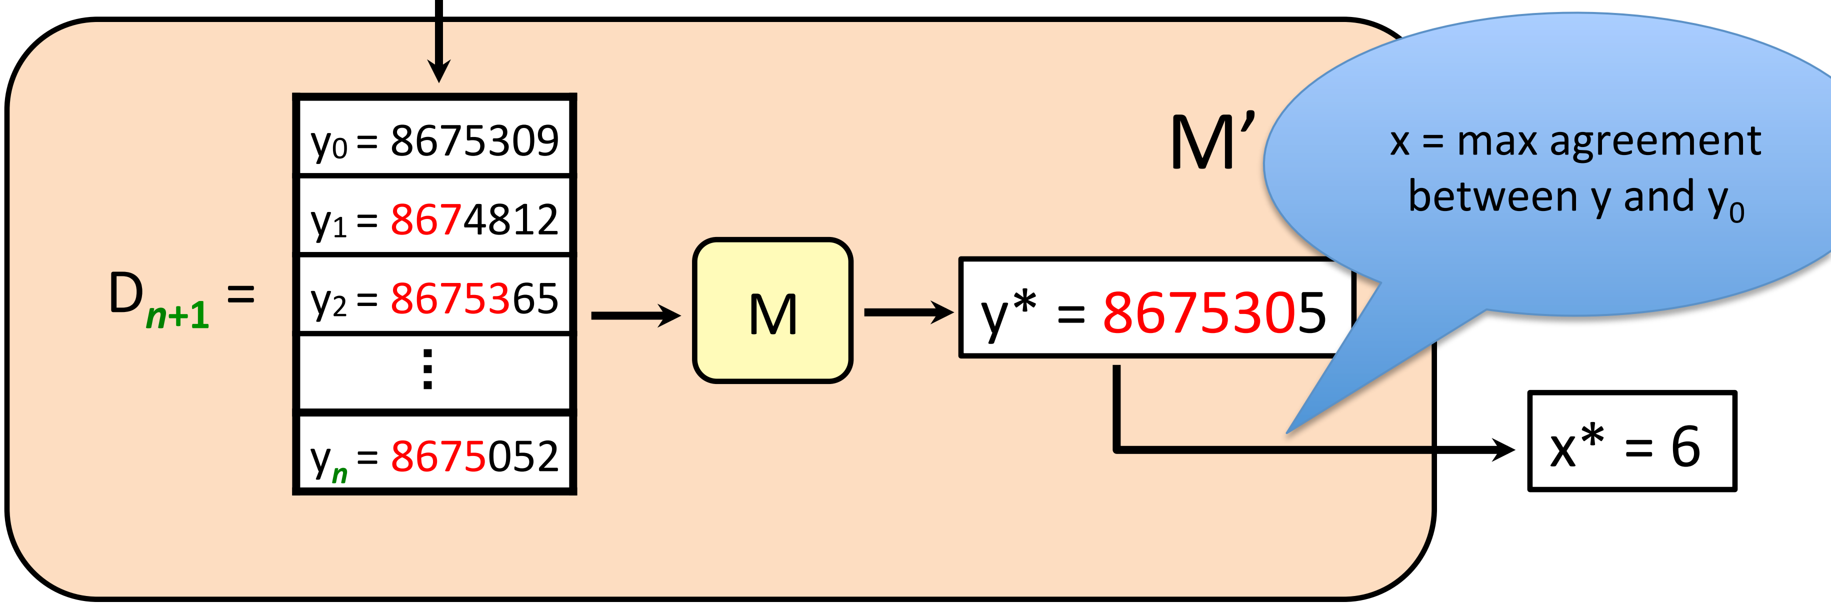

Interior Point Lower Bound

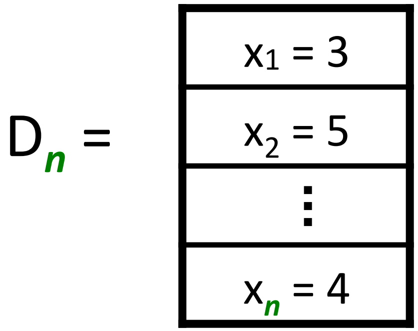

If M succeeds,

- $x \geq \min_i x_i$ by construction
- $x \leq \max_i x_i$ by privacy (whp)

$\Rightarrow$ M' succeeds

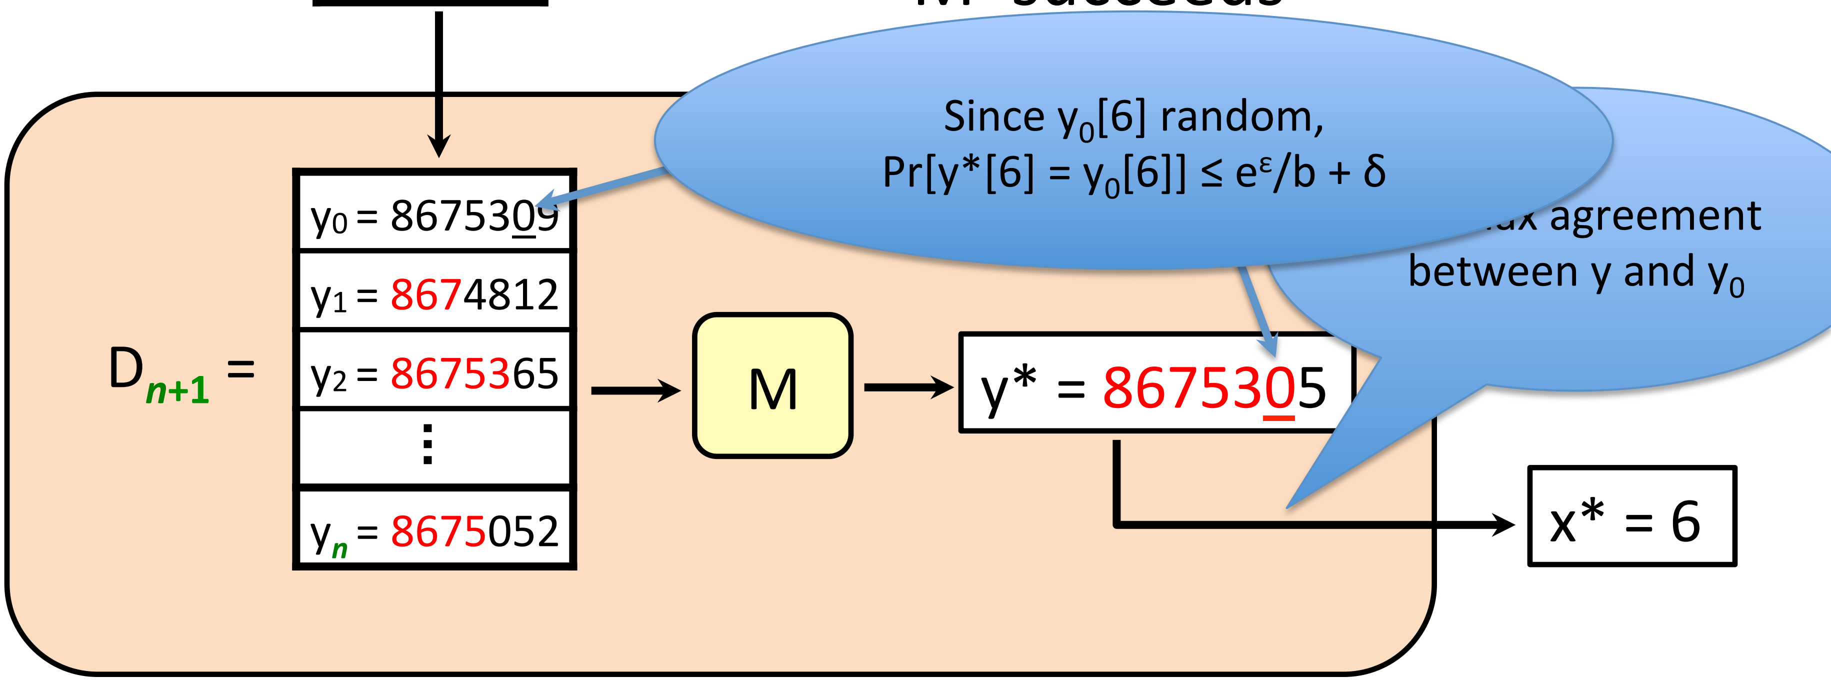

Interior Point Lower Bound

- Recursively construct hard distributions $\mathcal{P}_n$ on domain size $T(n) \approx \text{tower}(n)$ $\Rightarrow$ $n \geq \log^* T$
- Base case: For $n = 1$, set $T(1) = 2$

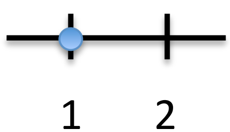

Output 1 w.p. ≥ 2/3

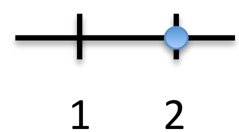

Output 1 w.p. $\geq \frac{(2/3)-\delta}{e^{\varepsilon}} > \frac{1}{3}$

- Inductive case:

  Suppose M solves IPP on $\mathcal{P}_{n+1}$ over domain $[T(n+1)]$

  => construct M' for IPP on $\mathcal{P}_n$ over $[T(n)]$