On the t-wise Independence of Block Ciphers

Tianren Liu

Peking University

Angelos Pelecanos

UC Berkeley

Stefano Tessaro

University of Washington

Vinod Vaikuntanathan

MIT

Lucas Gretta

UC Berkeley

William He

CMU

t-wise independence of block ciphers

- Block ciphers: (for this talk) practical encryption schemes.
  - e.g. Advanced Encryption Standard (AES).

- Symmetric-key encryption scheme
  - Users share a secret key sk.
  - Encrypt message using Enc, and decrypt with Dec.

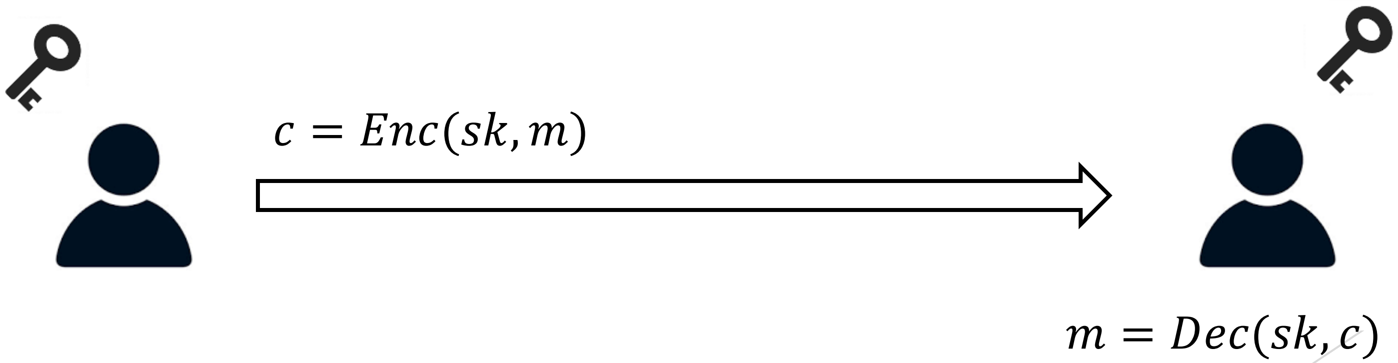

t-wise independence of block ciphers

- Block ciphers: (for this talk) practical encryption schemes.
  - e.g. Advanced Encryption Standard (AES).
- Information about AES (taken from Wikipedia)
  - Specification for the encryption of electronic data established by the U.S. National Institute of Standards and Technology (NIST) in 2001.
  - Selected by NIST after a five-year standardization process in which fifteen competing designs were presented and evaluated.
  - AES became effective as a U.S. federal government standard on May 26, 2002.
  - AES is the first (and only) publicly accessible cipher approved by the NSA for top secret information.

t-wise independence of block ciphers

- **Summary**. AES is very important,
  - used everywhere, all the time.
- We trust AES so much,
- We must have proved it is secure, right?

What does it mean to be secure?

- Cryptographers like to prove security via **reductions**.

- LWE: Solve a noisy linear system modulo a number.
- DDH: Given g^a, g^b, the element g^{ab} looks like a random element of $\mathbb{Z}_q$.

- **Goal**. Encryption scheme $\mathcal{S}$ is secure.
  - Need. Mathematical problem $\mathcal{P}$ we believe is hard.
- Proof by contradiction:
  - Assume there exists adversary $\mathcal{A}$ that breaks the security of $\mathcal{S}$.
  - Use $\mathcal{A}$ to also solve problem $\mathcal{P}$, contradiction!

Encryption scheme $\mathcal{S}$ ⟸ break

solve ⟹ Mathematical problem $\mathcal{P}$

What does it mean to be secure?

- **Crucial**. If $\mathcal{A}$ breaks the security of $\mathcal{S}$, then it can solve $\mathcal{P}$.
  - Means that $\mathcal{S}$ and $\mathcal{P}$ share some structure.

- To prove AES is secure via a reduction, we need
  - Hard mathematical problem
  - That is similar to AES.

- LWE: Solve a noisy linear system modulo a number.
- DDH: Given g^a, g^b, the element g^{ab} looks like a random element of $\mathbb{Z}_q$.

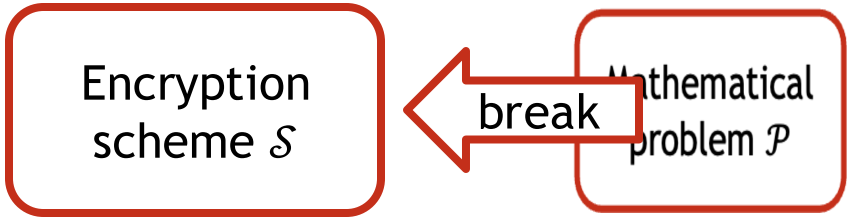

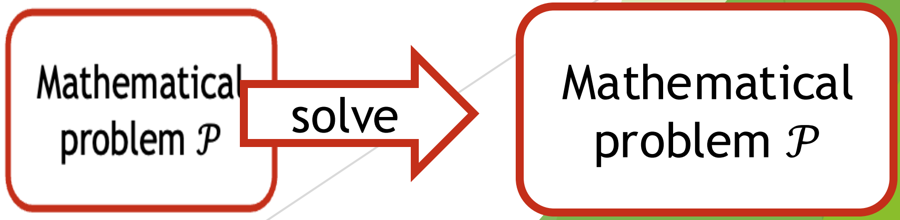

Let’s take a look at AES

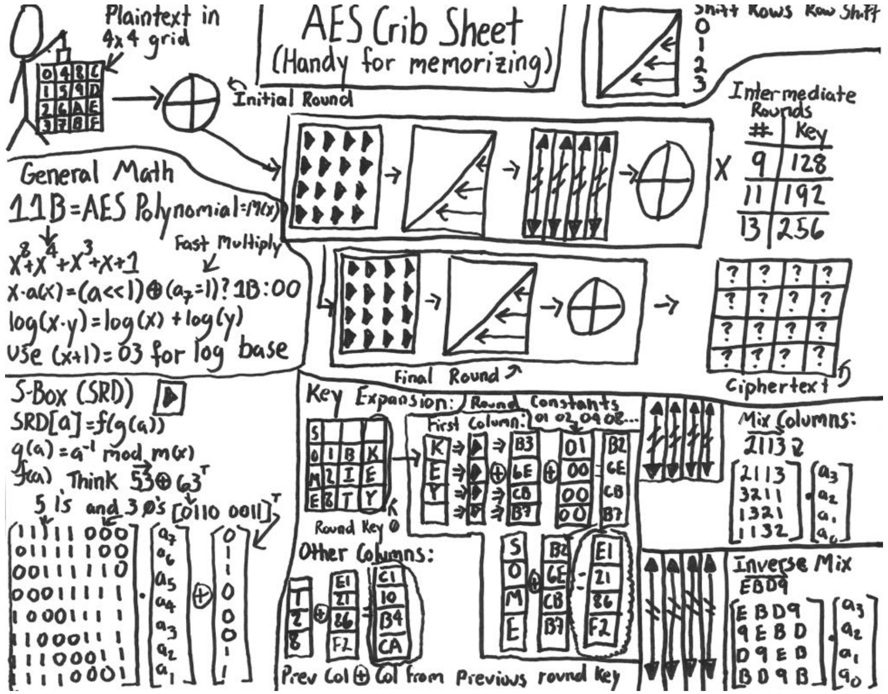

moserware.com

- No known math problems come to mind…
- Let’s try to get a theory-friendly description first.

Let’s take a look at AES

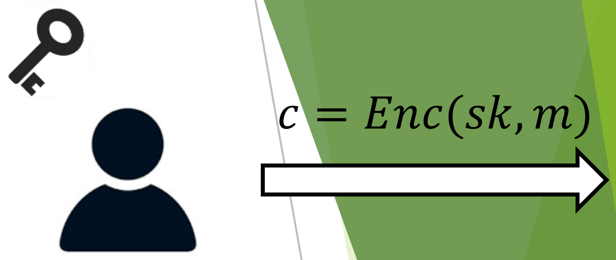

- AES describes the $Enc(\cdot)$ procedure.
  - Takes the secret key and the message as inputs.
- The encryption happens in rounds.
  - The secret key will have as many parts as rounds: $sk = (key_0, key_1, \dots, key_r)$.
  - In each round, the message is modified a little bit.
- **Philosophy**. Many simple modifications “scramble” the message to a ciphertext that is indistinguishable from random

Warm-up:
Key-Alternating Cipher (KAC)

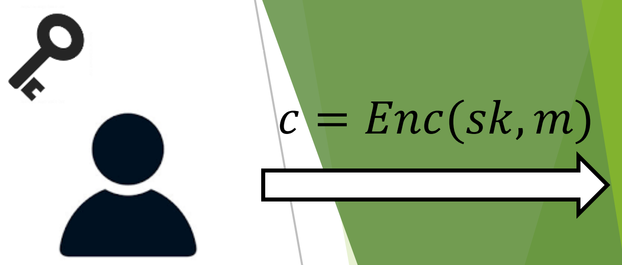

$sk = (key_0, key_1, \dots, key_r)$

Key length: $n \times$ (# rounds)

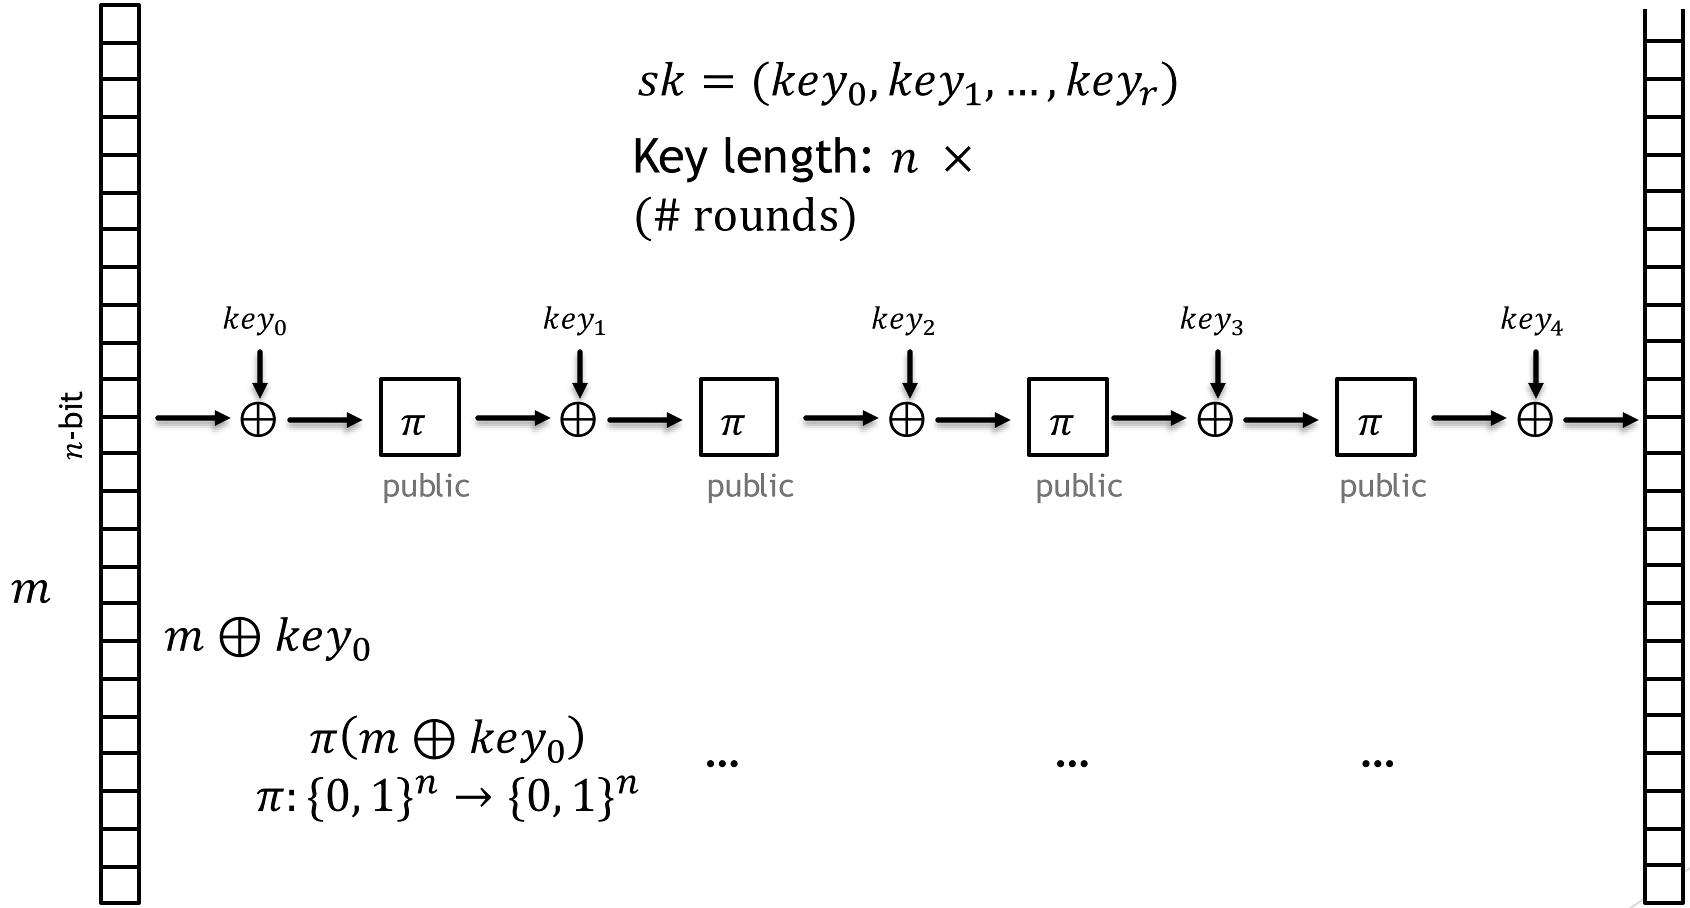

Substitution-Permutation Network (SPN)

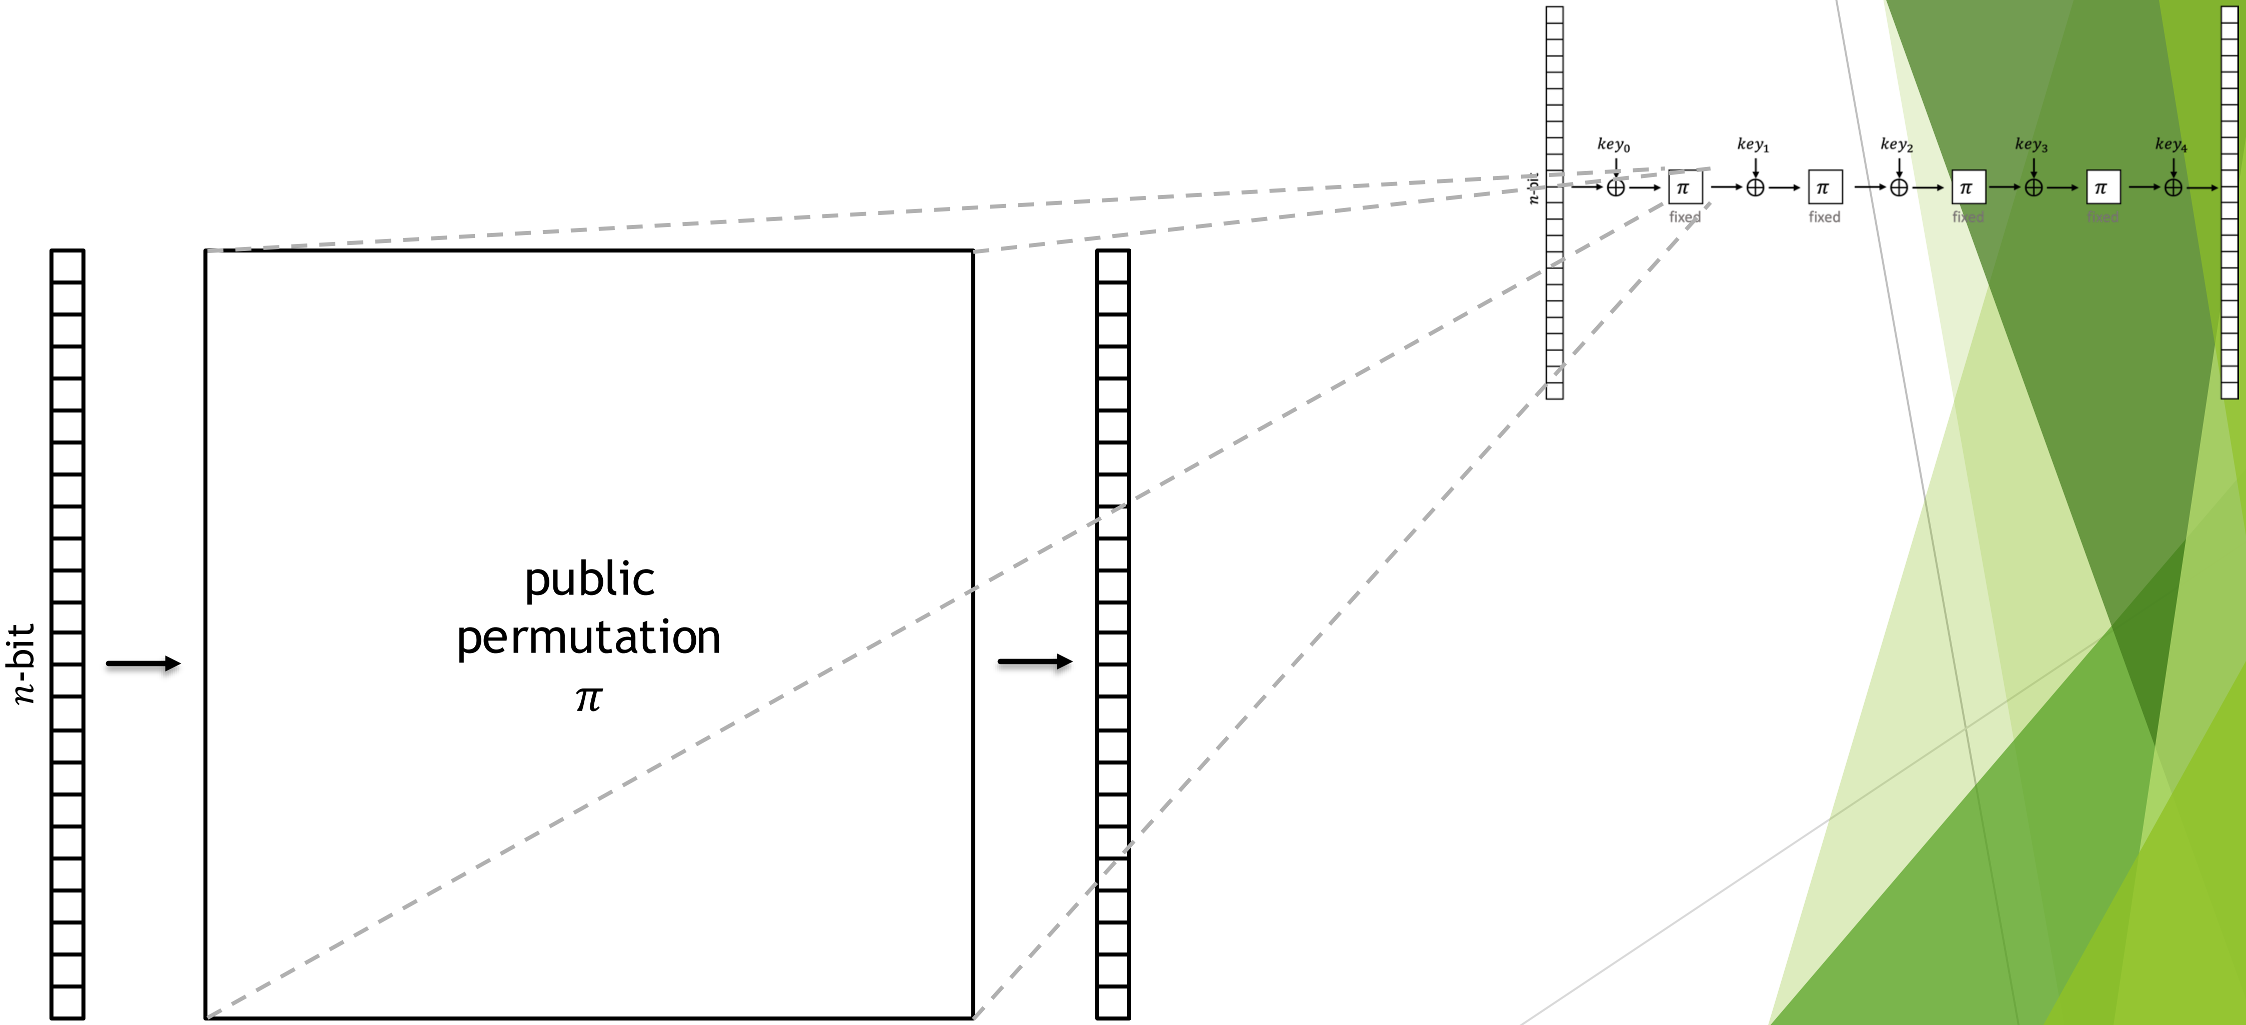

Substitution-Permutation Network (SPN)

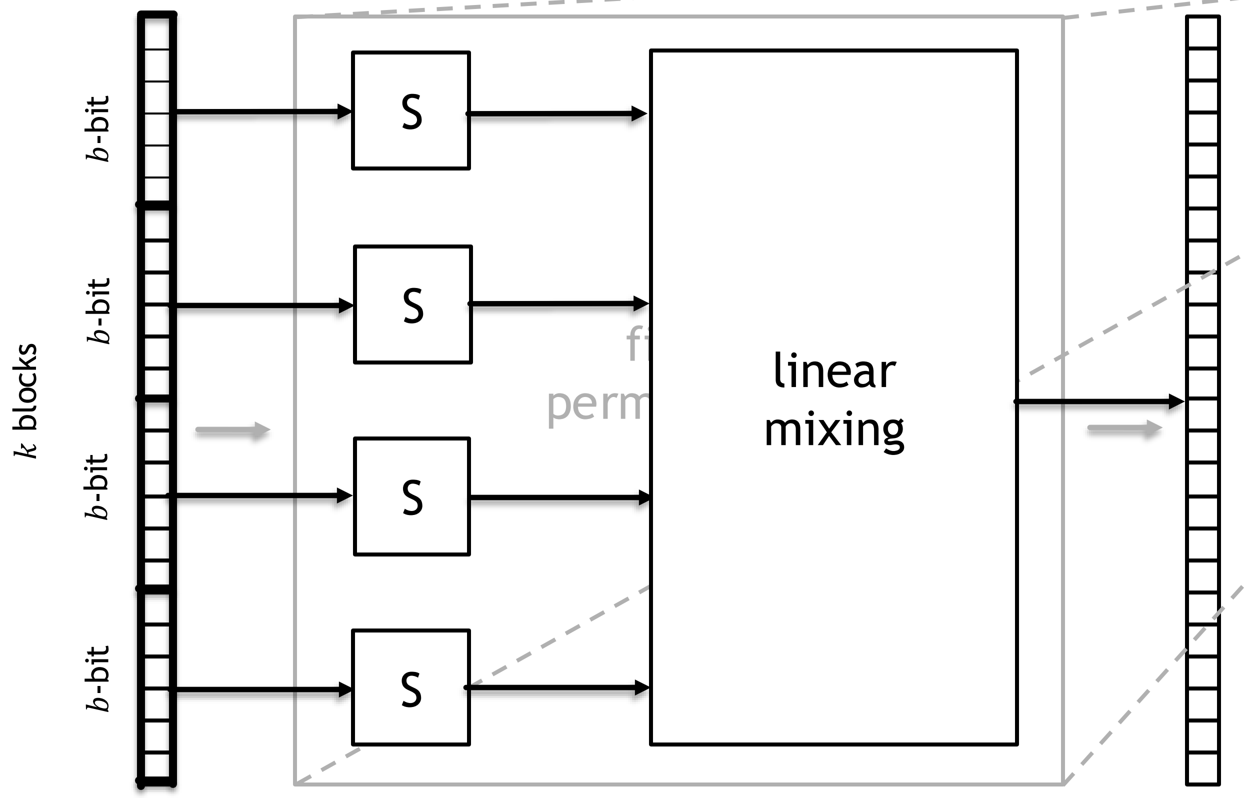

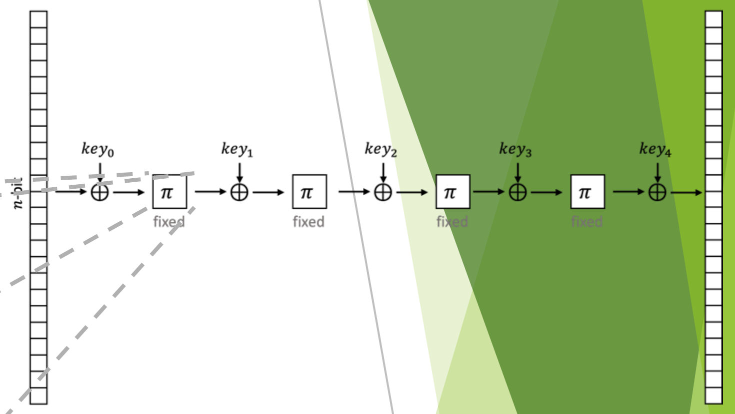

- $S: \{0,1\}^b \rightarrow \{0,1\}^b$
- Linear mixing is a matrix $\mathbb{F}_{2^b}^{k \times k}$

Substitution-Permutation Network (SPN)

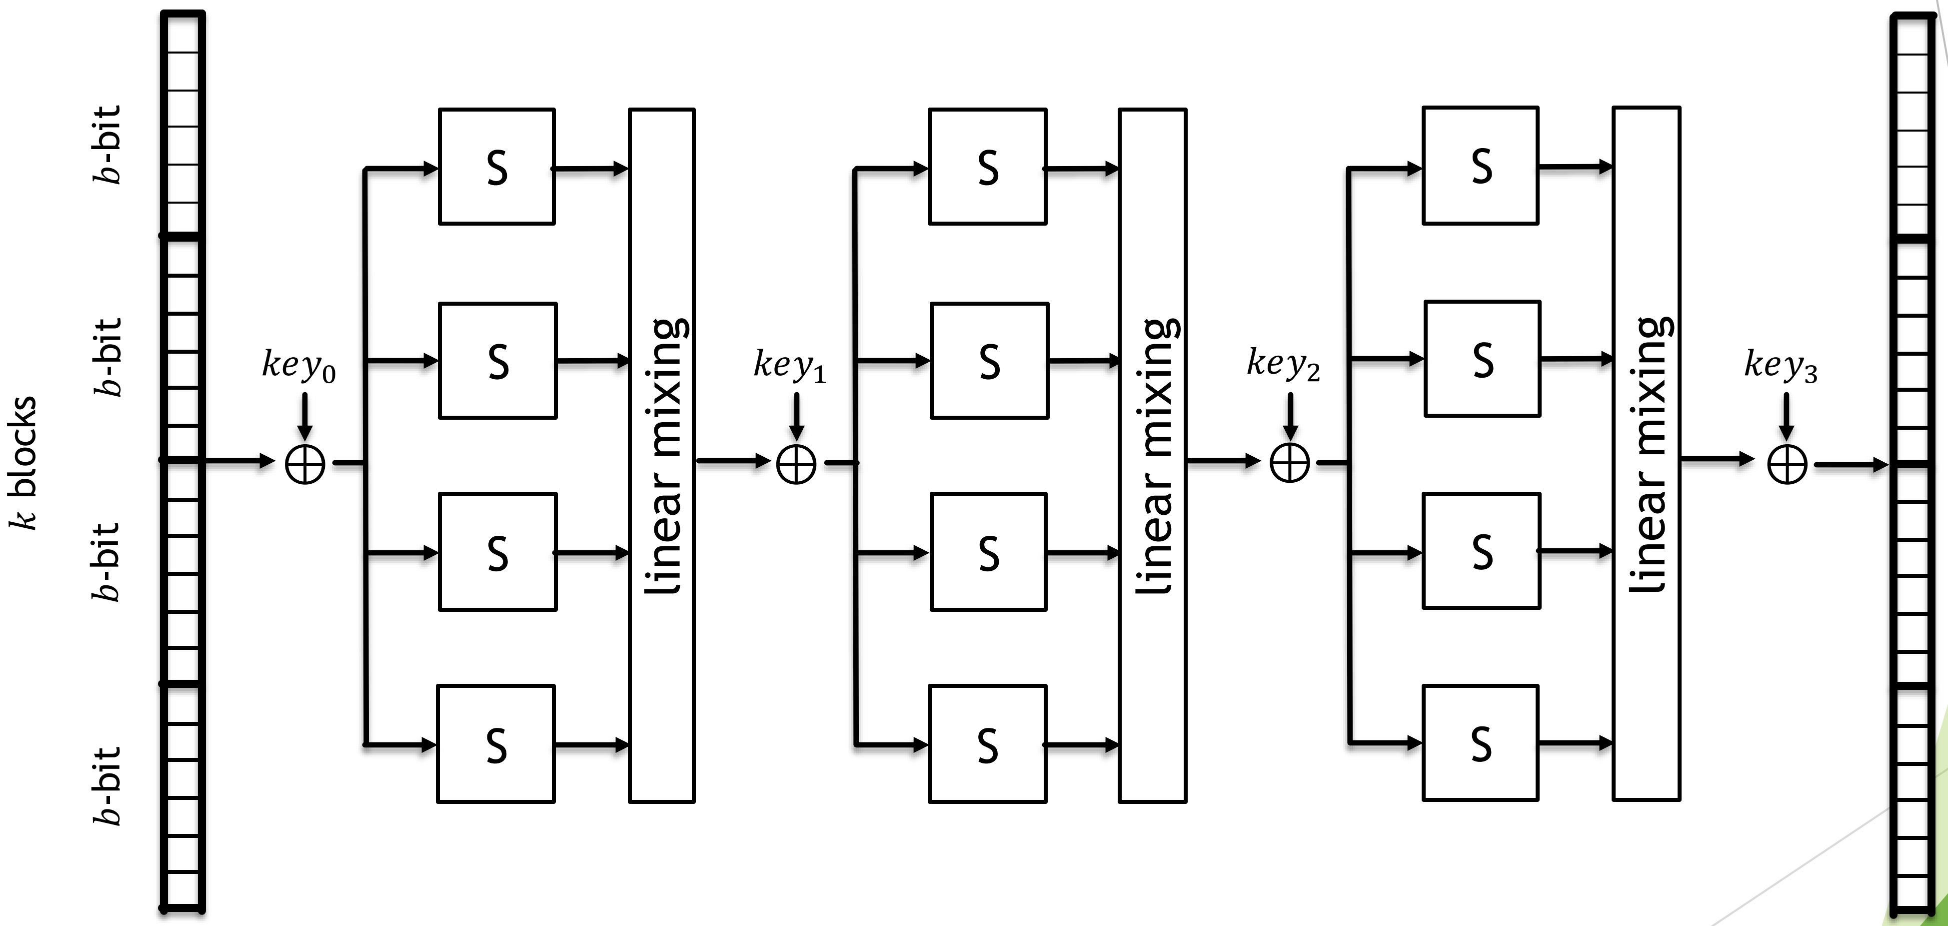

Substitution-Permutation Network (SPN)

SPN Parameters	
$n = kb$	Input length
b	S-box input length
k	Number of blocks
Number of rounds	
S-box (public perm. over $\{0,1\}^b$)	
Linear mixing ($k \times k$ matrix)	

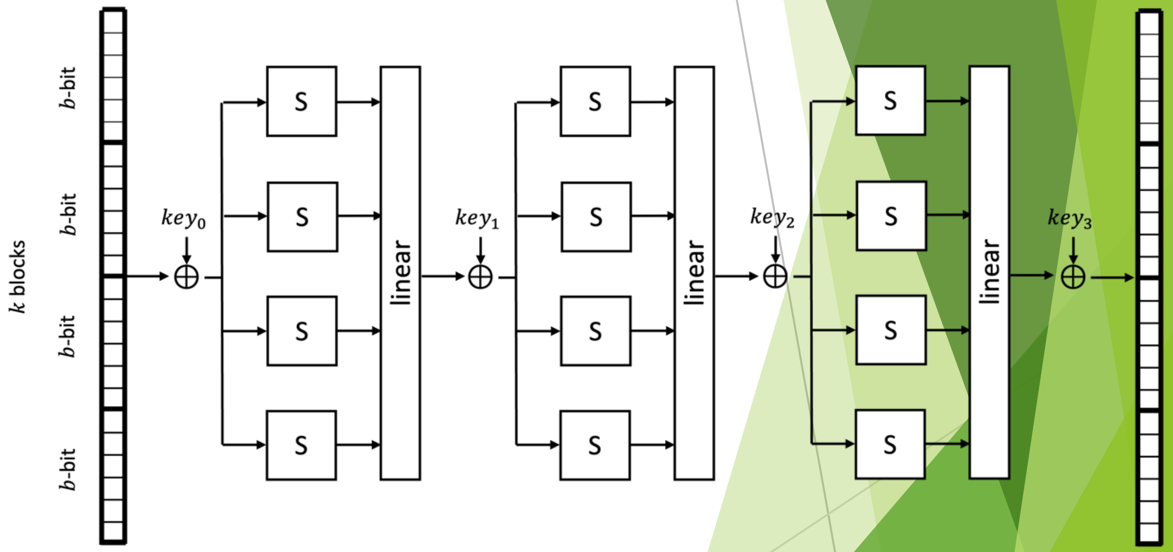

AES as an SPN

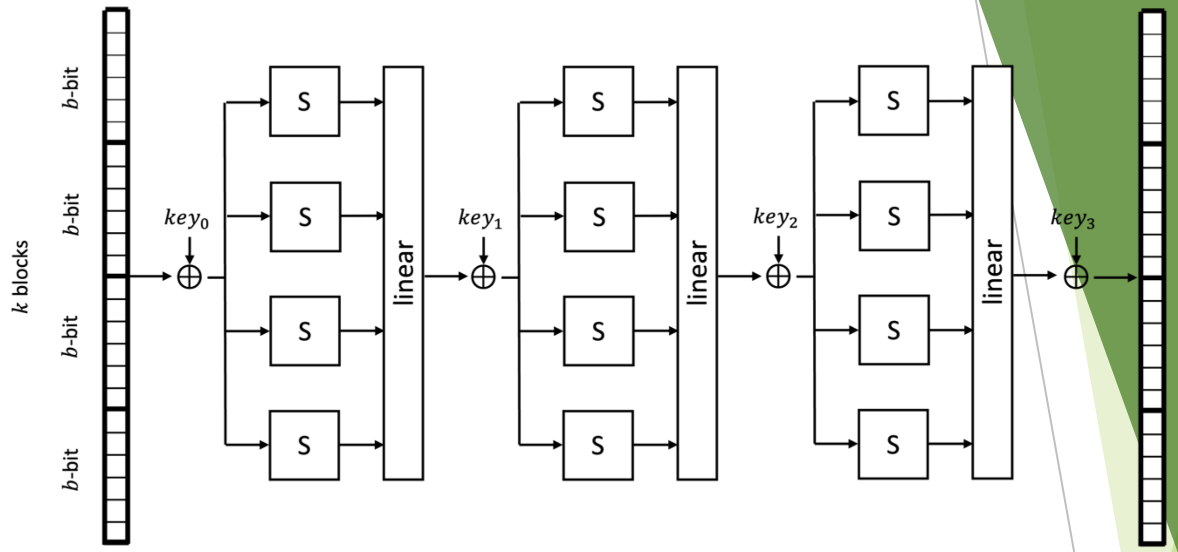

SPN Parameters		AES
$n = kb$	Input length	128
b	S-box input length	8
k	Number of blocks	16
Number of rounds		10 or 12 or 14
S-box		INV over $\mathbb{F}_{2^8}$ $x \to x^{-1}$
Linear mixing		ShiftRows & MixColumns

What about a reduction?

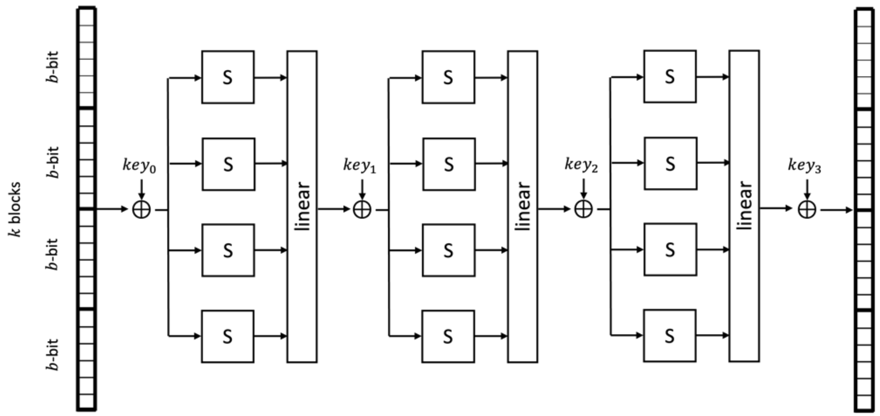

- Still no known math problems come to mind…
- What if we cannot prove security by a reduction?
- There is another way…

Security against specific attacks

- Reduction: encryption scheme is secure against **all adversaries**.
- What if we show security against **specific attacks**?

- Cryptanalysts have been very busy developing attacks against block ciphers
  - Differential attacks,
  - Linear attacks,
  - Square attacks,
  - Impossible differential,
  - Yoyo,
  - Multiple-of-8.

t-wise independence of block ciphers

AES is t-wise independent if…

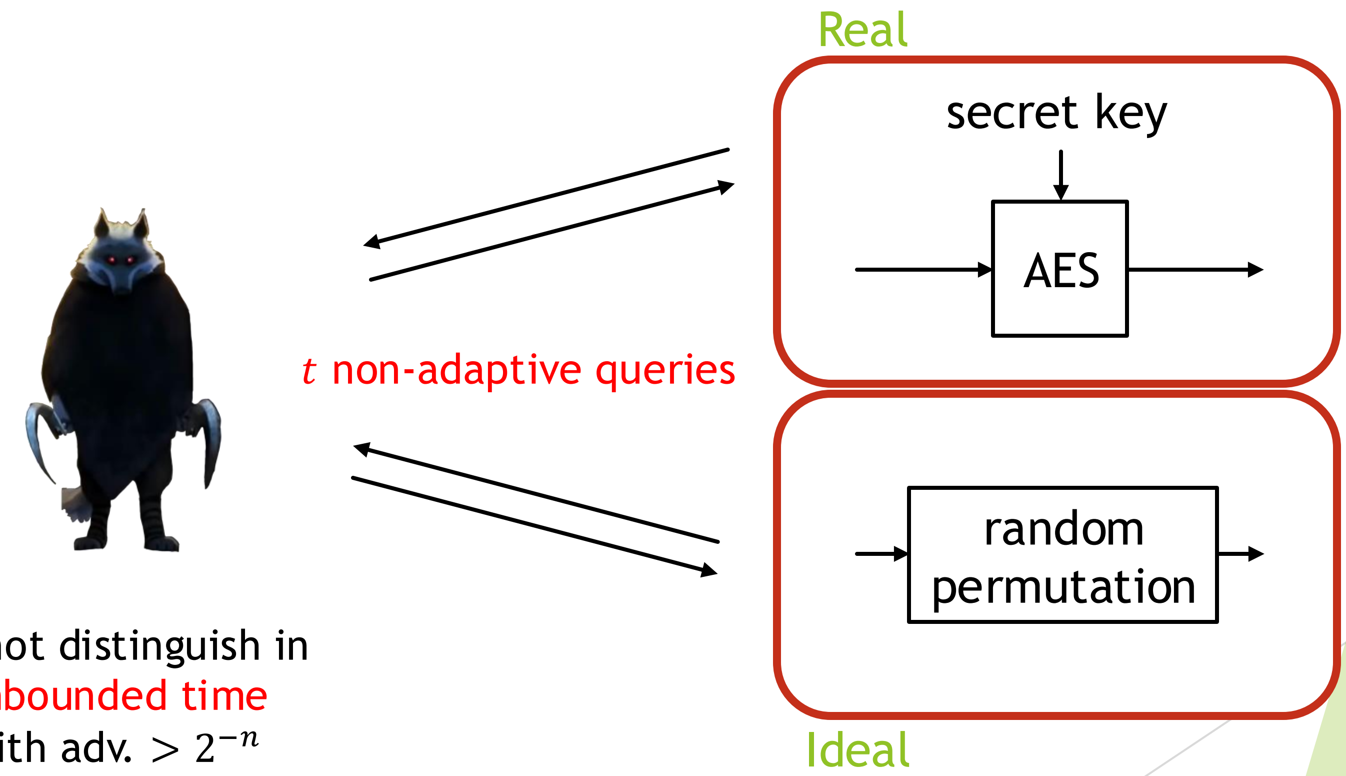

Why t-wise independence?

Definition. (ε-close to t-wise independence).For all t inputs $x_1, \ldots, x_t$

$$\text{statistical-distance}\big((y_1, \ldots, y_t), \text{uniform}\big) \leq \varepsilon$$

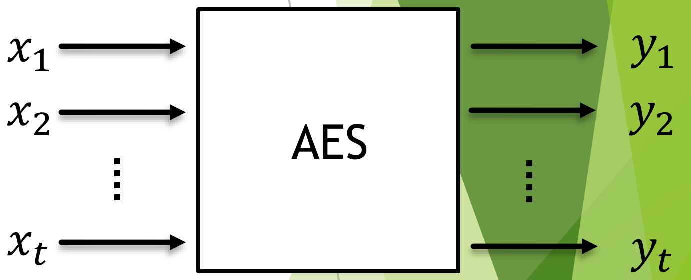

- Protects against a wide range of attacks, including
  - $t = 2$: Differential attacks [BS91], linear attacks [MY92].
  - $t = 2^d$: (truncated) degree-d differential attacks [Lai94, Knu94].
- t-wise independence is not the goal, **a lens through which to study security**:
  - Study natural constructions that are **provably** t-wise, and **plausibly** pseudorandom.

Why t-wise independence?

Definition. (ε-close to t-wise independence).For all t inputs $x_1, \ldots, x_t$

$$\text{statistical-distance}\big((y_1, \ldots, y_t), \text{uniform}\big) \leq \varepsilon$$

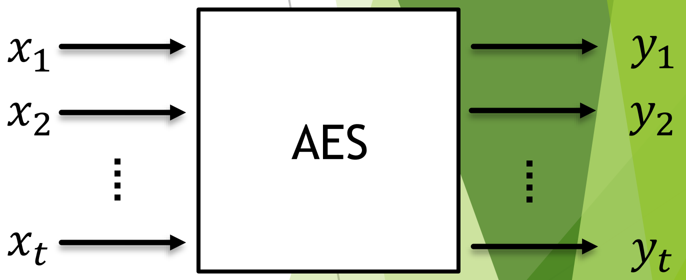

- Allows to compare block ciphers in a quantitative way
  - Say block cipher A is t-wise independent in fewer rounds than block cipher B $\Rightarrow$ block cipher A is more "secure"?
- Conjectured ([HMMR05]) that a block cipher that is 4-wise independent is also pseudorandom.

Why t-wise independence?

Definition. (ε-close to t-wise independence).For all t inputs $x_1, \ldots, x_t$

$$\text{statistical-distance}\big((y_1, \ldots, y_t), \text{uniform}\big) \leq \varepsilon$$

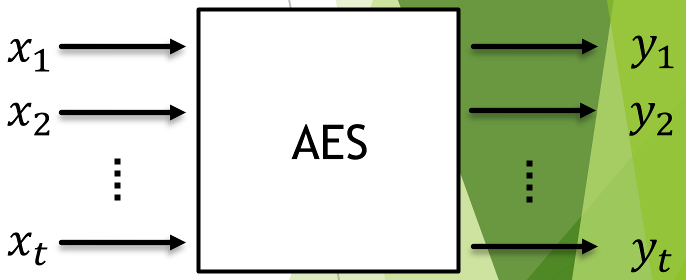

- **Feasible** (potentially unconditionally) when $|key| \geq t \cdot n$.
  - e.g., assume independent round keys.
  - i.e. We can prove things about it!

Part I. SPN results

SPN results [LTV21]

- **Theorem [LTV21].** 2-round SPN is $\approx \sqrt{\frac{2^k}{2^b}}$-close to 2-wise independent.
- **Theorem [LTV21].** 3-round SPN is $\approx \sqrt{\frac{k}{2^b}}$-close to 2-wise independent.
- Holds if linear mixing of the SPN achieves maximal branching number.
  - This is not true for AES.

Recall.
$k = 16$, $b = 8$ for AES

AES result

- **Theorem [LTV21].** 6-round AES is 0.472-close to pairwise independent.

Amplification Lemma [MPR07]

$\mathcal{F}$ is ϵ-close to t-wise independent
$\Rightarrow \mathcal{F} \circ \mathcal{F}$ is $2\epsilon^2$-close to t-wise independent.

- **Corollary.** $6r$-round AES is $(0.472^r \cdot 2^{r-1})$-close to pairwise.
  - To achieve 2^{-128} security, we set $r \approx 1500$.
  - 9000-round AES is 2-wise independent!

Part II. SPN* Results

Idealized model: SPN* [BV06]

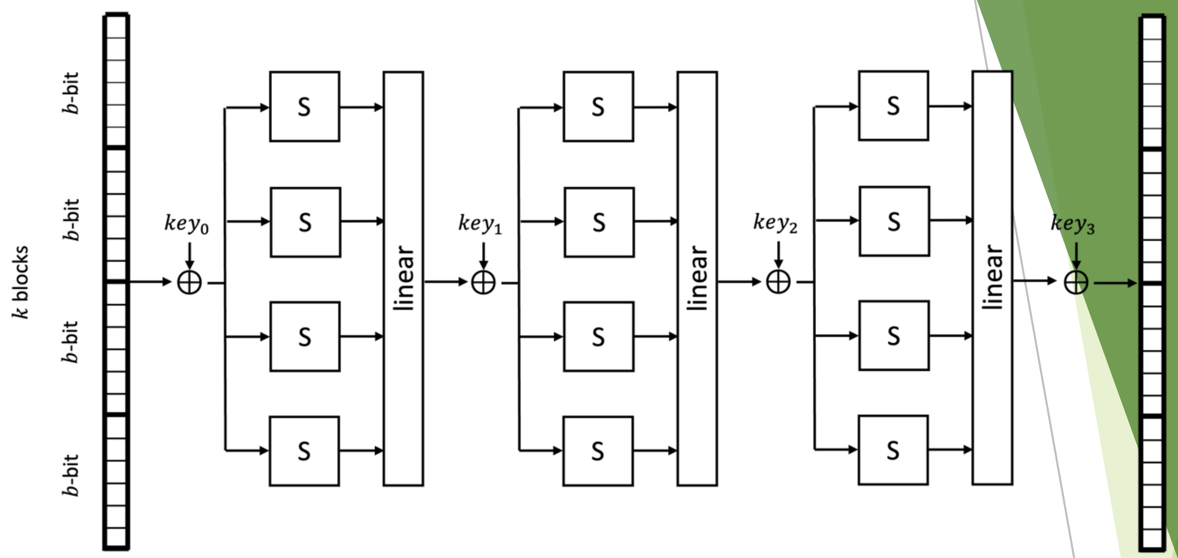

SPN Parameters		AES	SPN*
$n = kb$	Input length	128	n
b	S-box input length	8	b
k	Number of blocks	16	k
Number of rounds		10 or 12 or 14	r
S-box		INV over $\mathbb{F}_{2^8}$ $x \to x^{-1}$	random permutation over $\mathbb{F}_{2^8}$
Linear mixing		ShiftRows & MixColumns	Linear over $\mathbb{F}_{2^8}$

Random S-box Substitution-Permutation Network (SPN*)

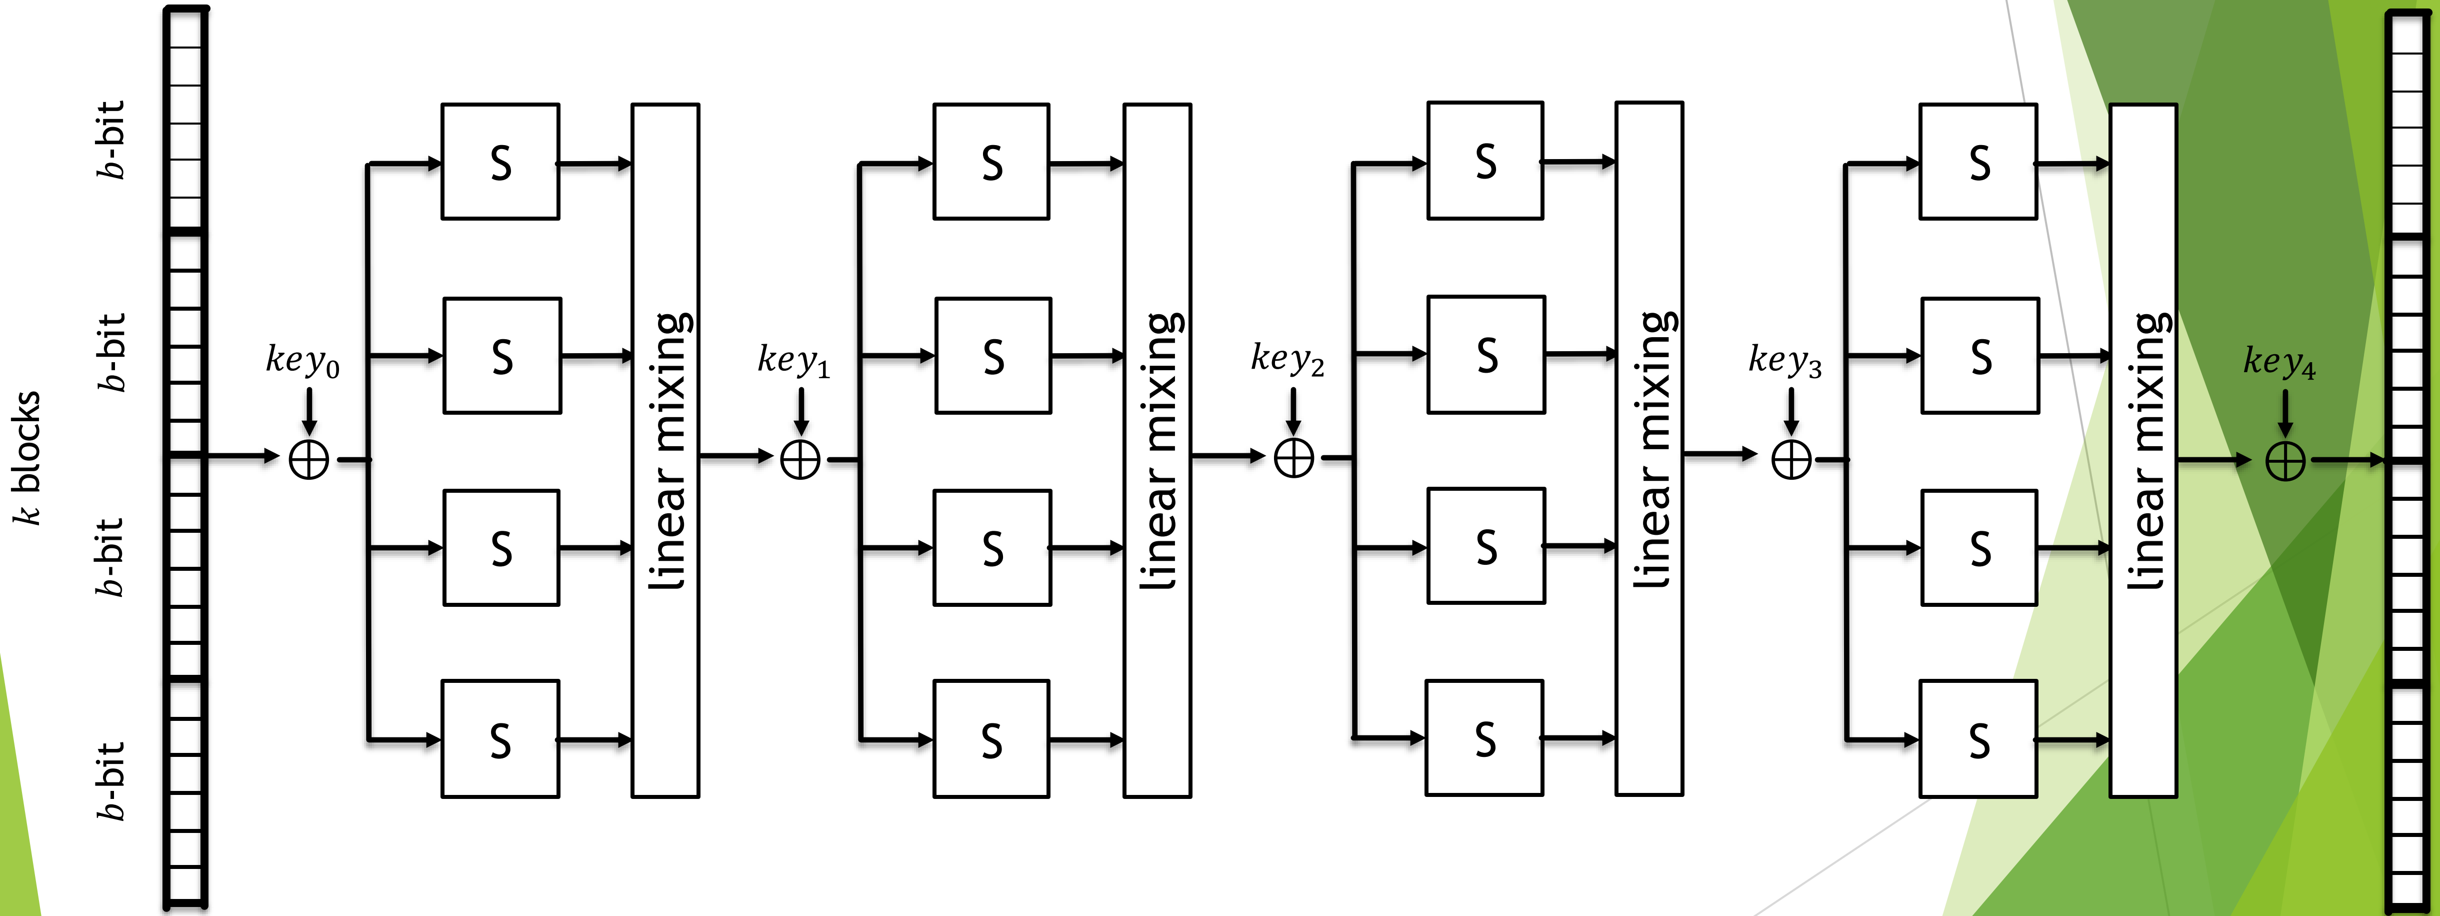

Random S-box Substitution-Permutation Network (SPN*)

The random S-boxes are now the key!

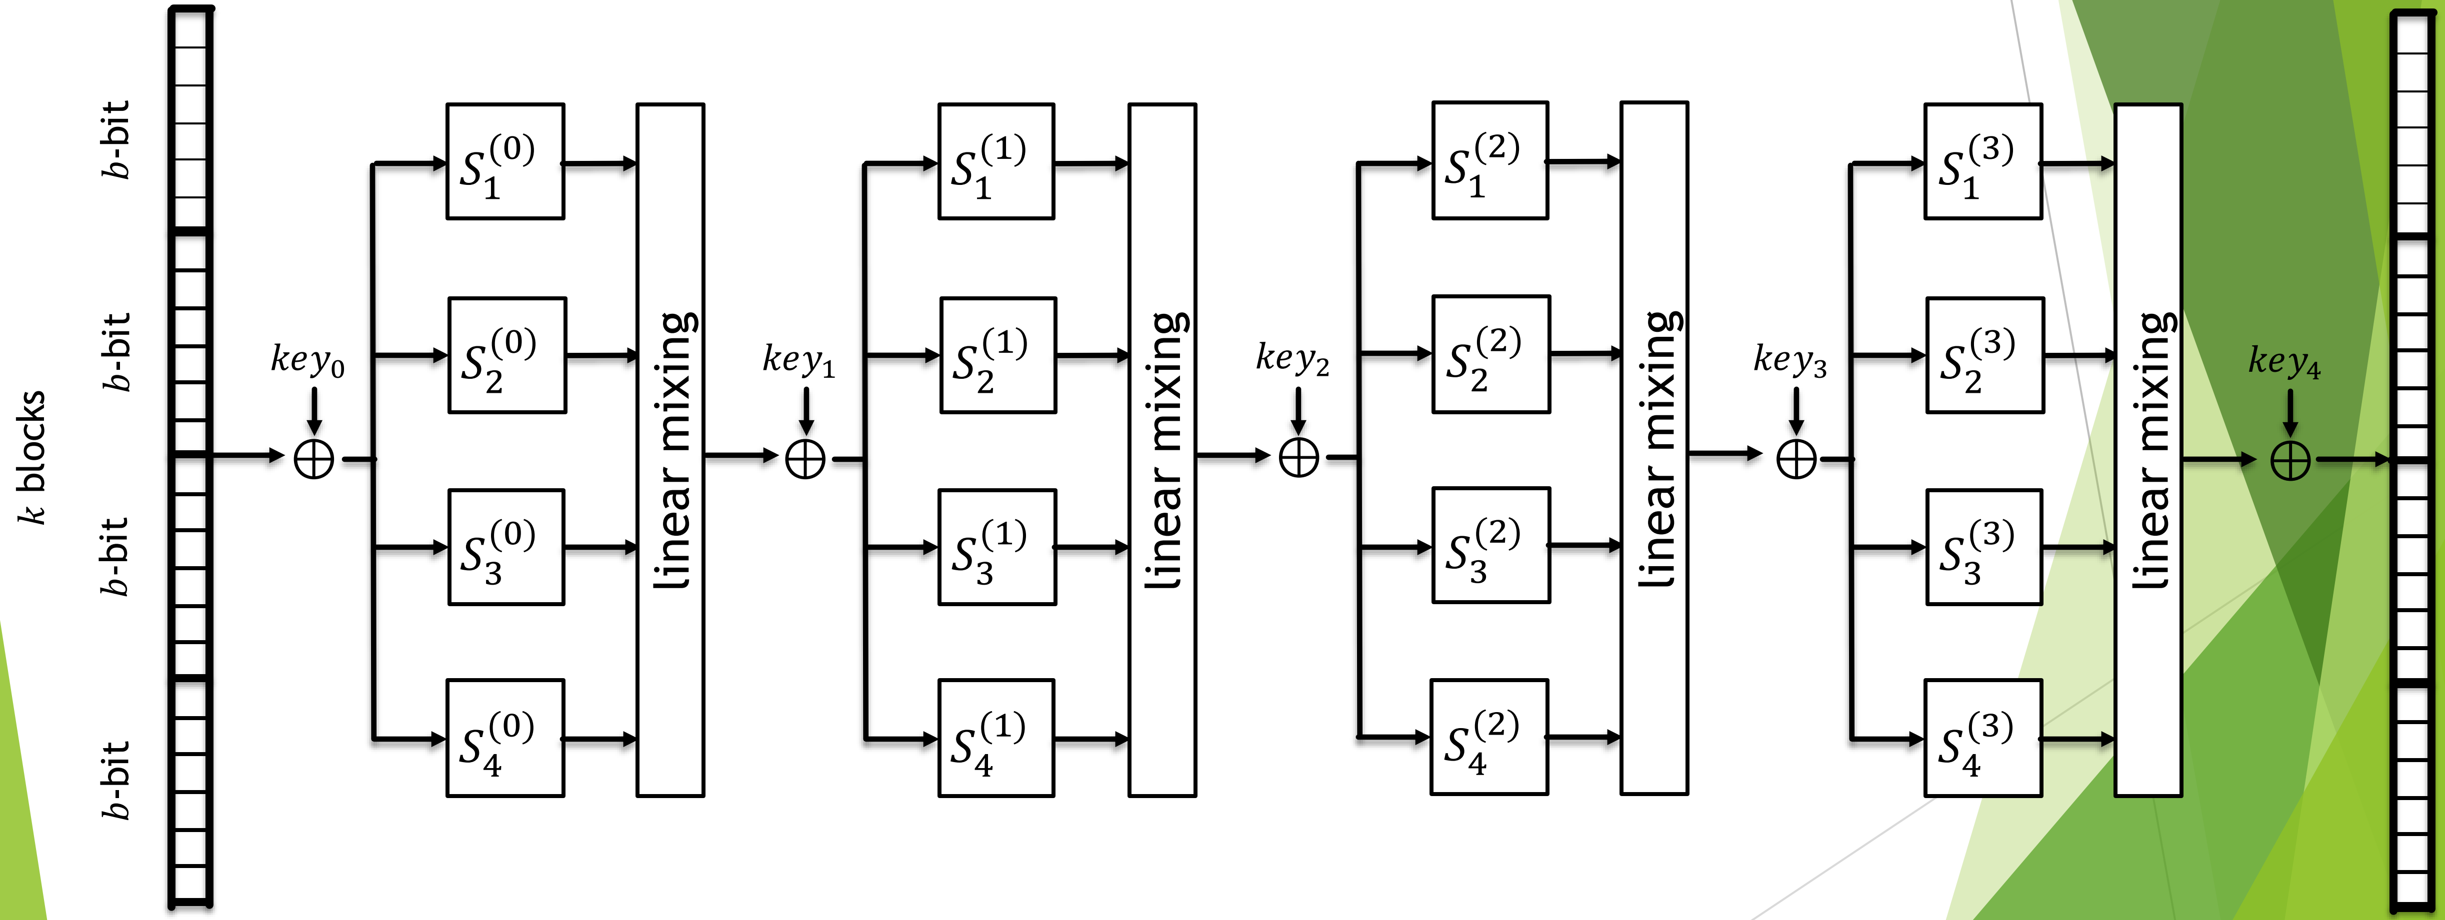

Random S-box Substitution-Permutation Network (SPN*)

The random S-boxes are now the key!

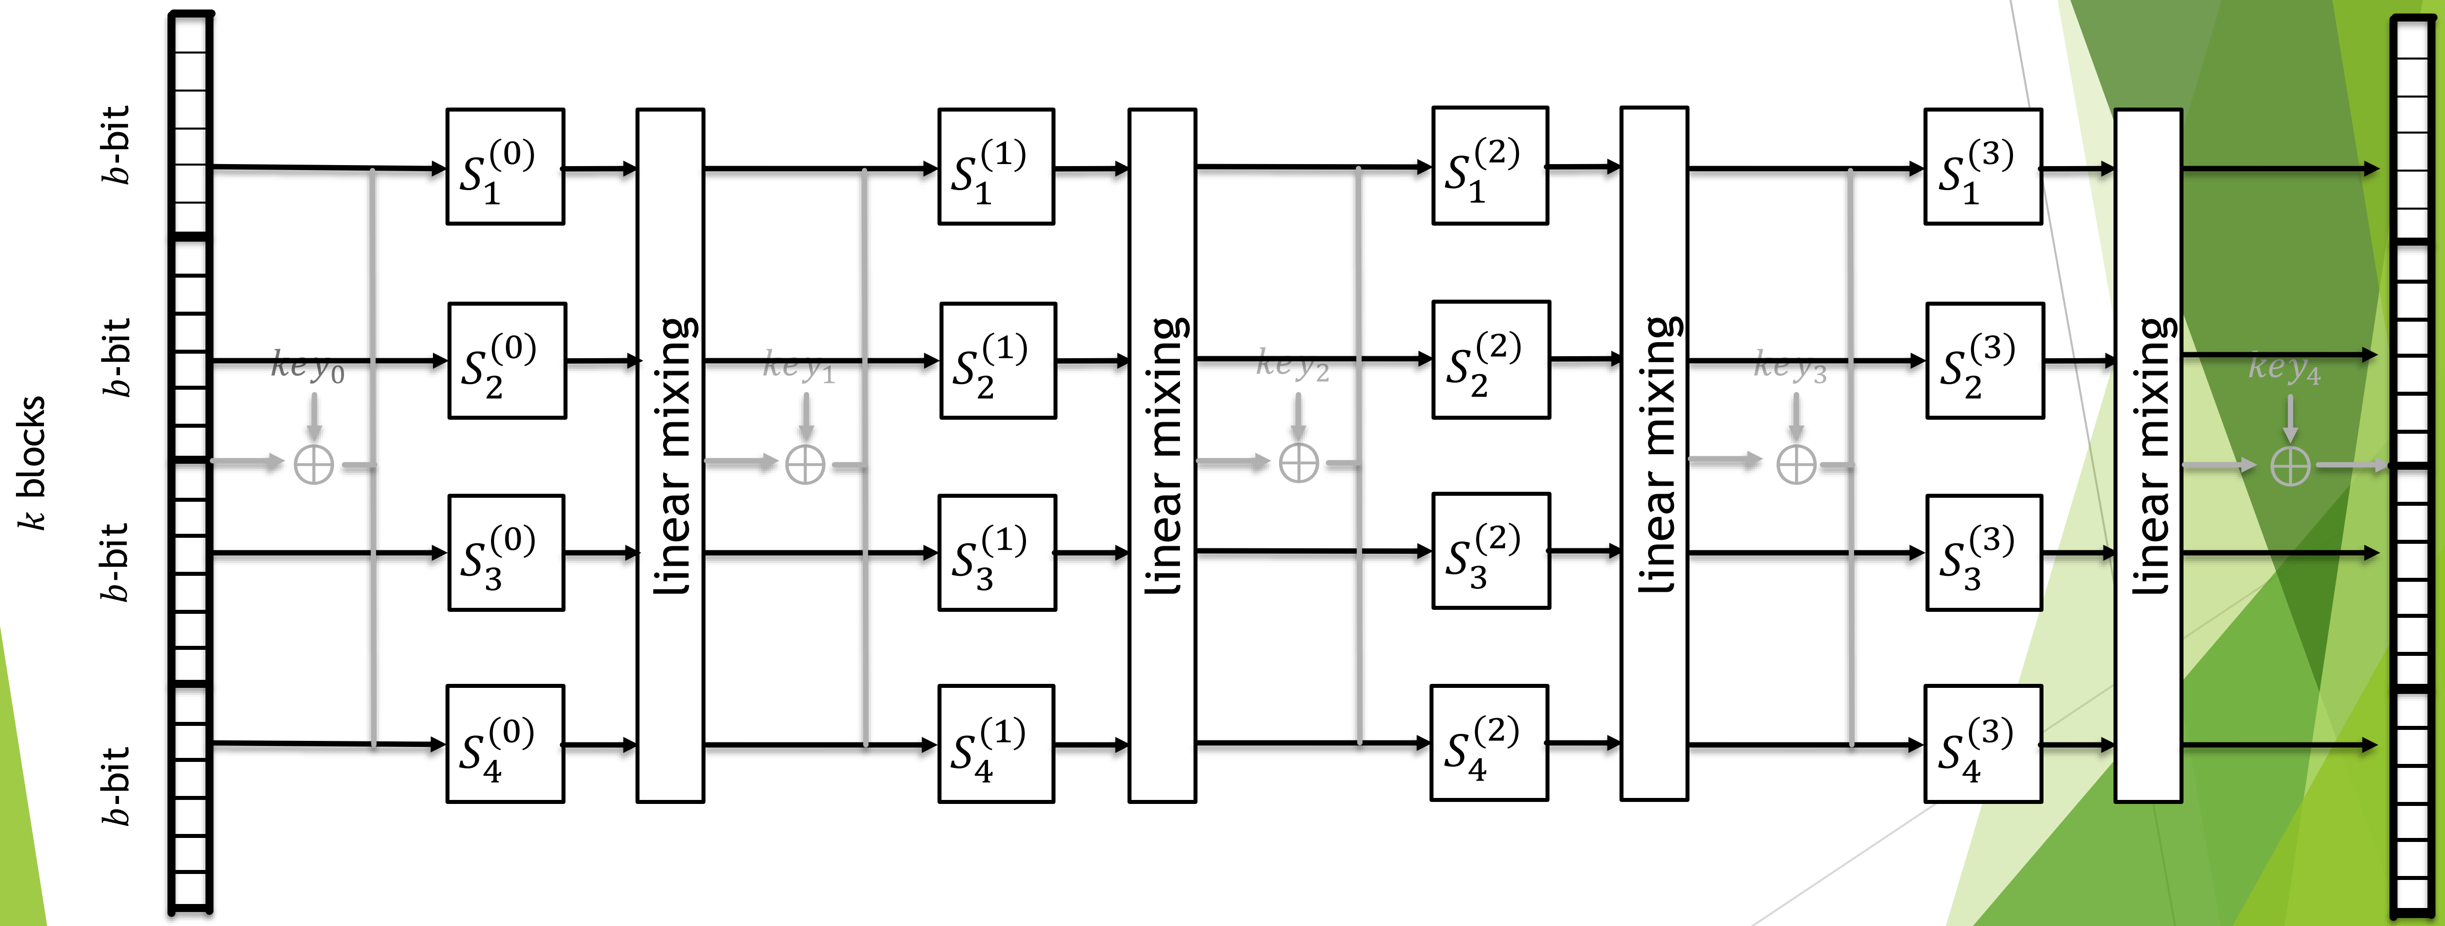

Usefulness of random S-box model

1. Block ciphers with random S-boxes already exist (GOST, Kufu).

2. Random S-box results can be translated to non-random S-box results.
   - Called *censored SPN*.

Lemma [LPTV23].

$$\underbrace{ARK \circ INV \circ \cdots \circ INV \circ ARK}_{O\left(b^2 \cdot 2^b \cdot \log\frac{1}{\varepsilon}\right)} \approx_\varepsilon \text{ random permutation over } \mathbb{F}_{2^b}$$

Usefulness of random S-box model

Lemma [LPTV23].

$$\underbrace{ARK \circ INV \circ \cdots \circ ARK \circ INV}_{O\left(b^2 \cdot 2^b \cdot \log \frac{1}{\varepsilon}\right)} \approx_\varepsilon \text{ random permutation over } \mathbb{F}_{2^b}$$

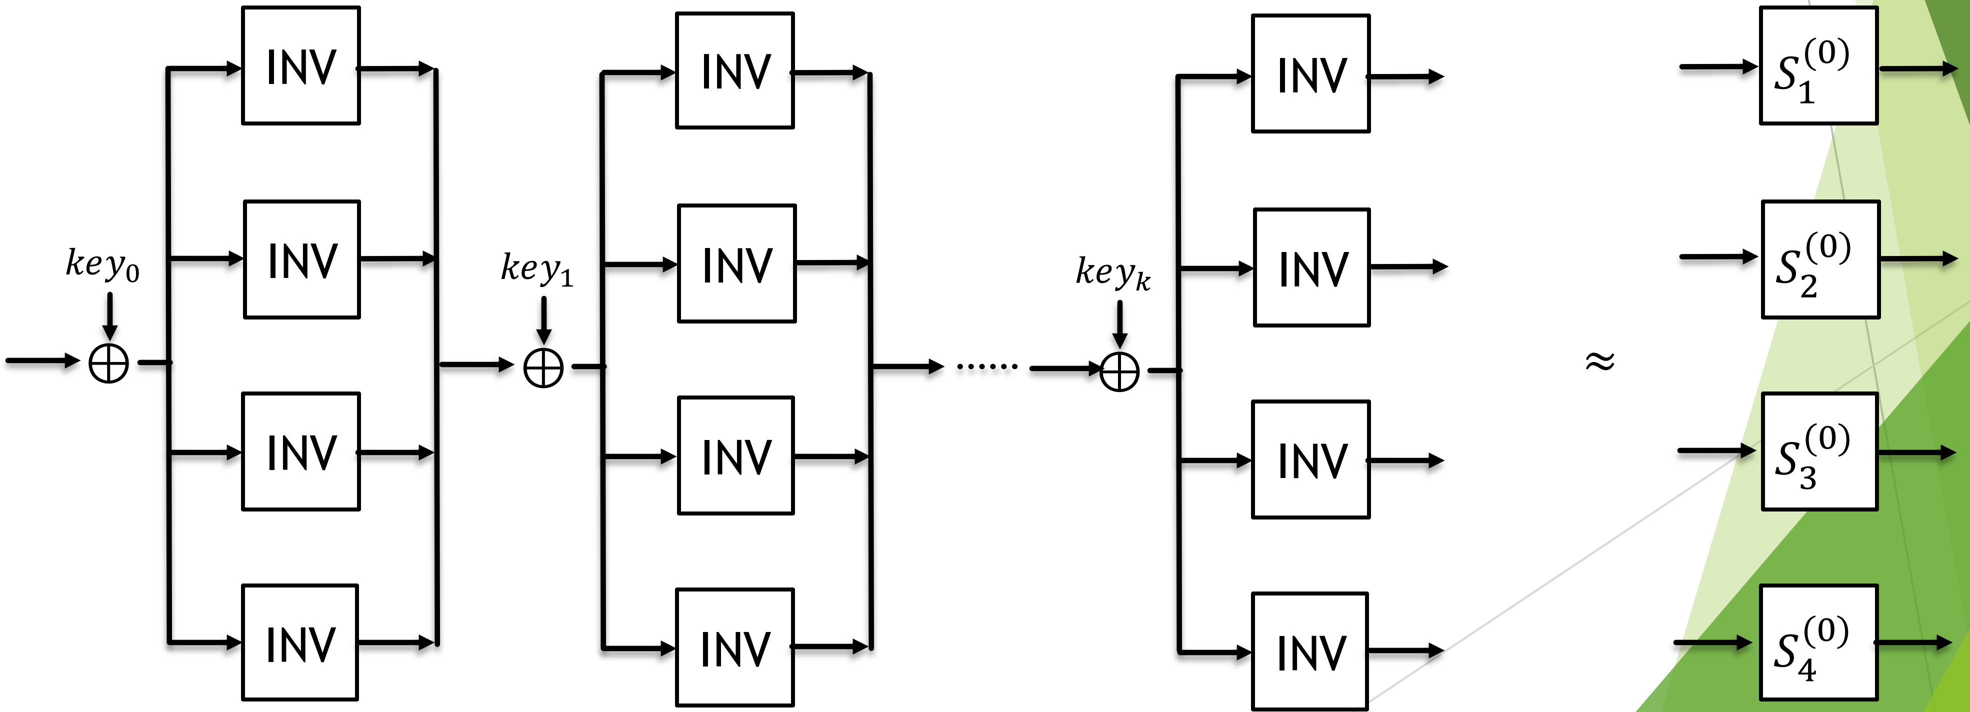

Usefulness of random S-box model

3. Ideal S-boxes allow us to identify desirable mixing properties.
   - Our proofs rely on the **maximal branching number** of the mixing matrix.
   - Our tight bounds explain how parameters affect convergence:
     - Number of blocks k increases, the SPN* converges faster.

SPN* Results

Theorem [LPTV23]. r rounds of SPN* suffice to reach ε-close to t-wise independence

Rounds r	ε-Closeness	t
2	$2^{-\Omega(kb)}$	$O(1)$
2	2^{-b}	$2^{\left(0.499-\frac{1}{4k}\right)b}$
$O(k)$	$2^{-\Omega(kb)}$	$2^{\left(0.499-\frac{1}{4k}\right)b}$
$O(\log t)$	$2^{-\Omega(kb)}$	$2^{0.499b}$

Limitation
t goes up to $\approx \sqrt{2^b} \ll 2^{kb}$

Constant t:
2 rounds suffice

Large t:
$\min\{O(k), O(\log t)\}$ suffice

AES*

- Consider the random S-box version of AES.
- The keys will be the S-boxes with random permutations over $\mathbb{F}_{2^8}$.

Theorem [LPTV23]. 7-round AES* is 2^{-128}-close to pairwise independent.

  - A lot of progress was done before by [BV06].
- Above result is tight: numerically verified.
- Can simulate AES* using censored AES:

Theorem [LPTV23]. 192-round censored AES is 2^{-128}-close to pairwise independent.

Censored SPN

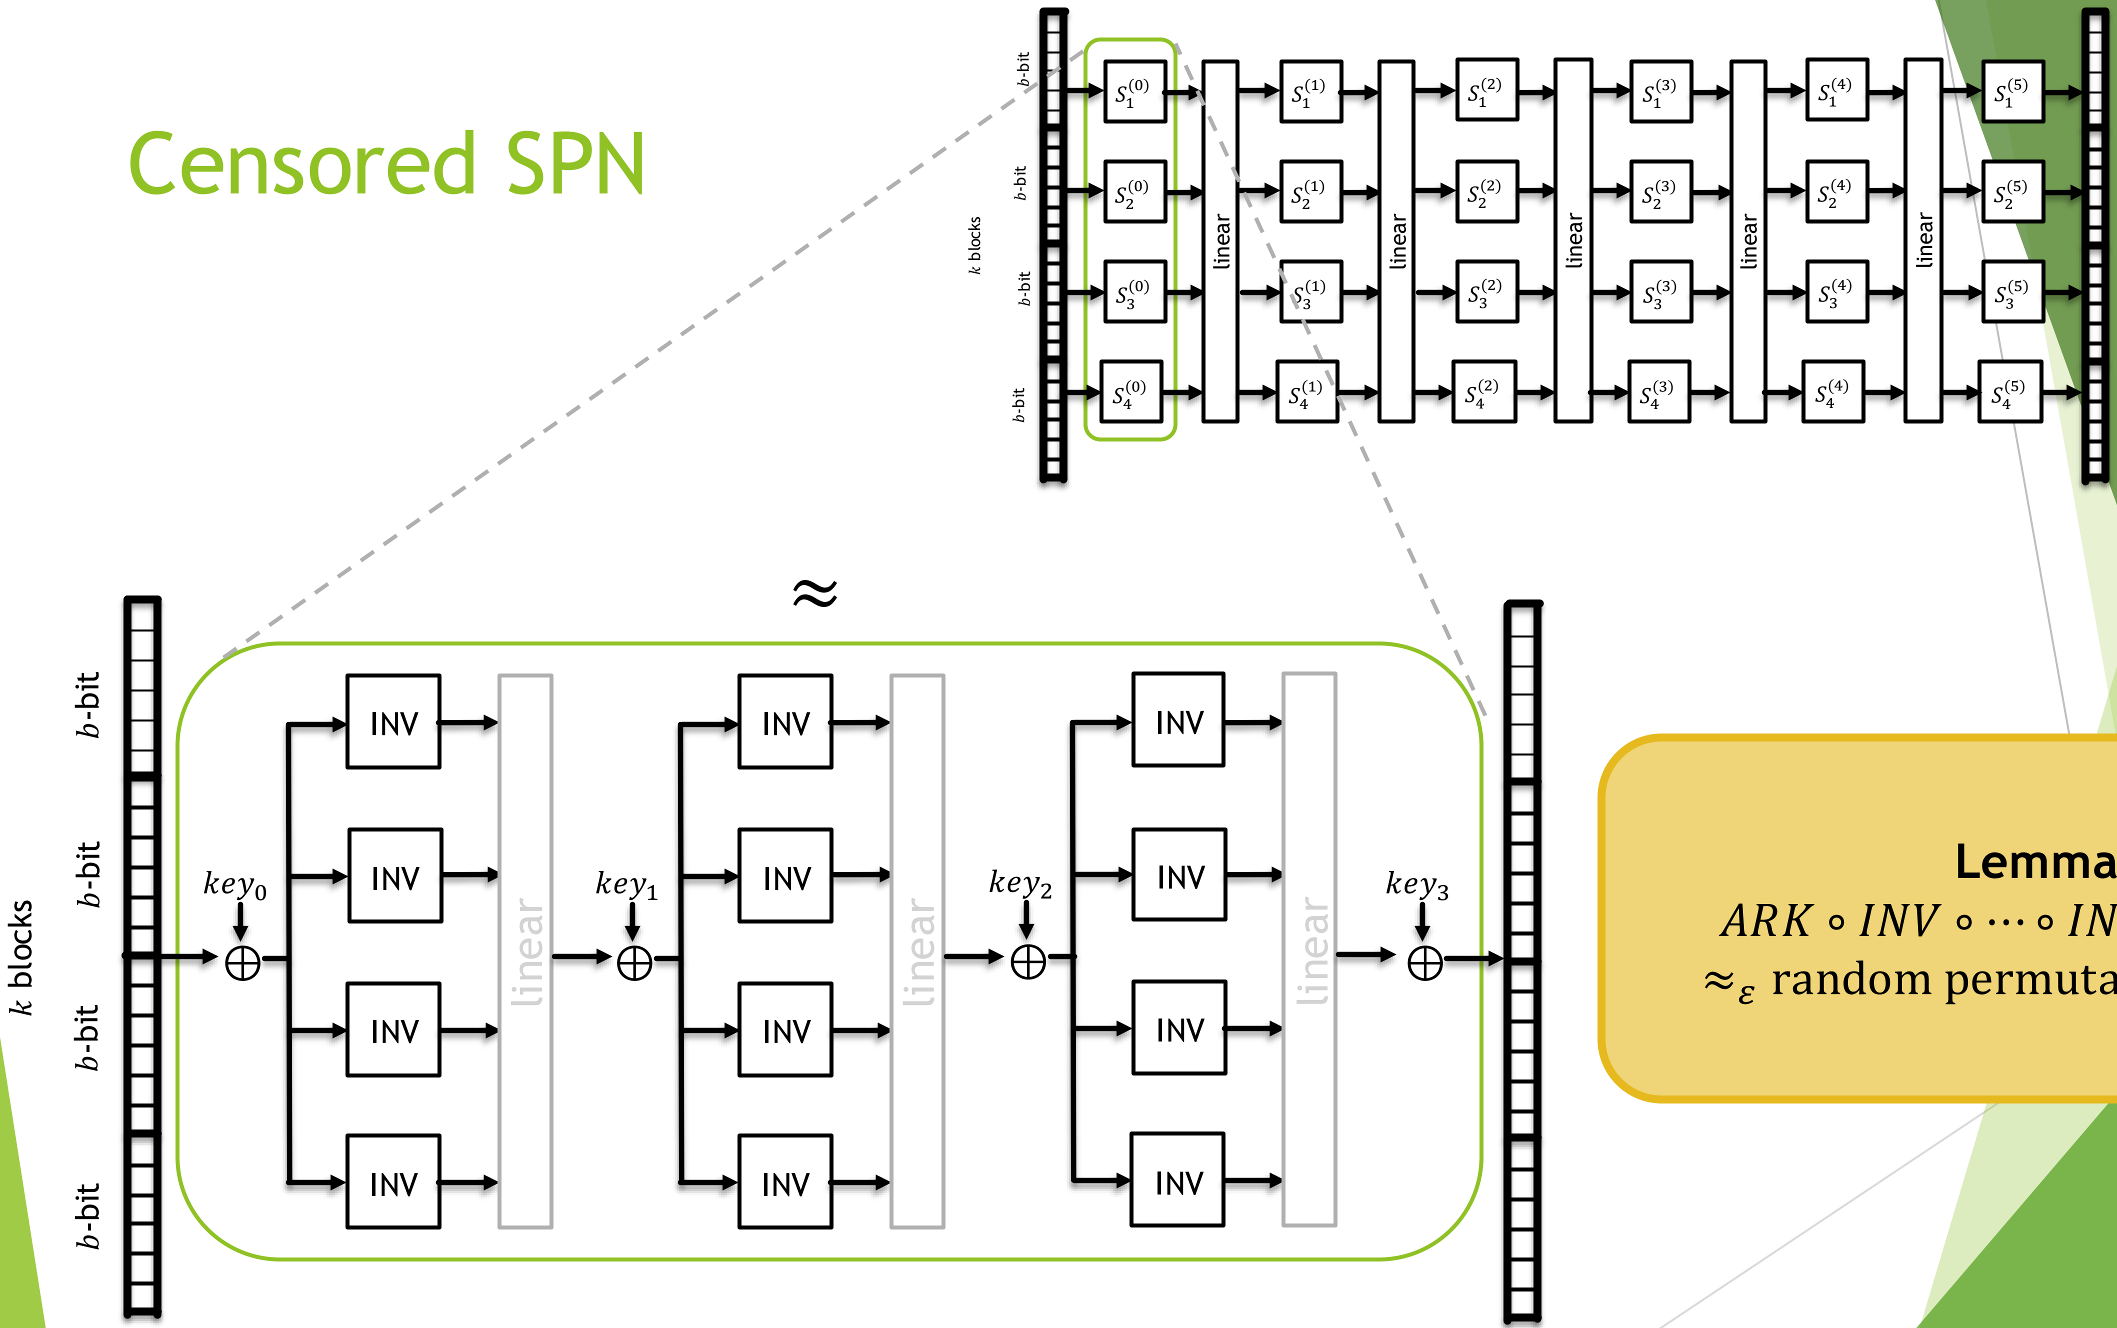

Lemma.

$ARK \circ INV \circ \cdots \circ INV \circ ARK \approx_\varepsilon$

$\approx_\varepsilon$ random permutation over $\mathbb{F}_{2^b}$

Censored SPN

Lemma.
$ARK \circ INV \circ \cdots \circ INV \circ ARK \approx_\varepsilon$ random permutation over $\mathbb{F}_{2^b}$

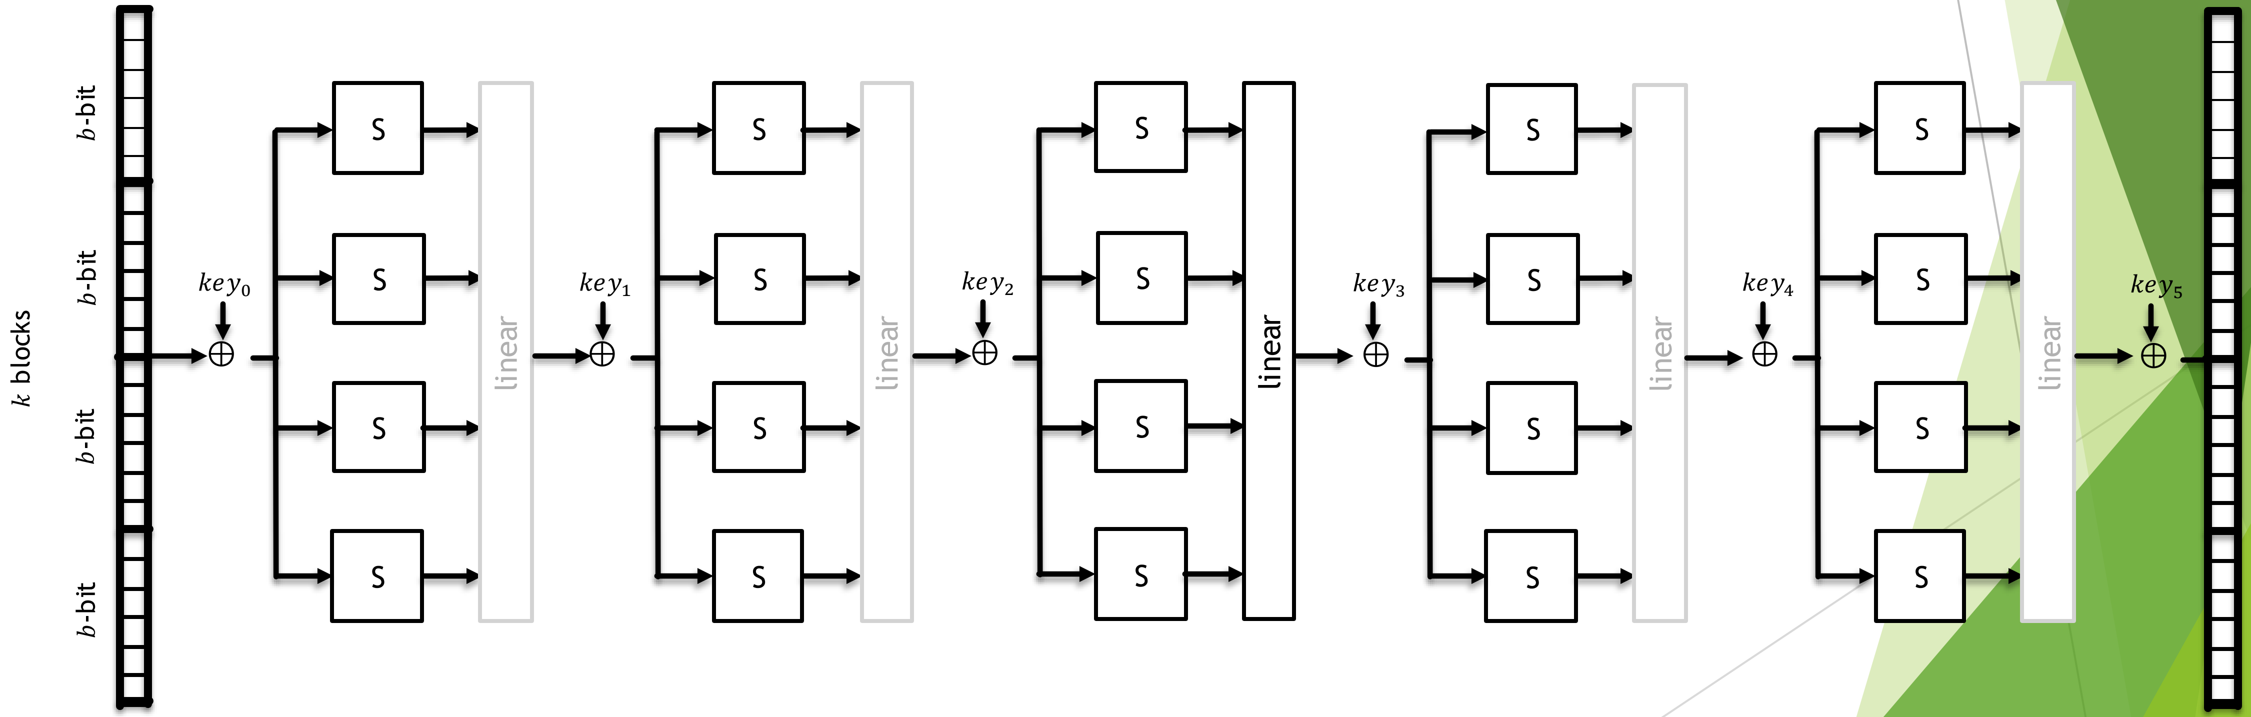

Censored AES

Theorem. 192-round censored AES is 2^{-128}-close to pairwise independent.

- Reasonable to expect that removing many mixing layers hurts security.
- **Evidence** that the true AES is pairwise independent in < 200 rounds.
- Contrast this with > 9000 rounds of AES in [LTV21]!

Part IIa. SPN* Technical Details

Layouts (2-wise)

- The random S-box destroys any correlation, except equality.
  - Let $\Delta X := X_1 \oplus X_2, \Delta Y := Y_1 \oplus Y_2$.
- Define $\mathrm{layout}(\Delta X) := \{i \mid \Delta X[i] = 0\}$.
  - Also known as "activity pattern".
  - Random S-boxes preserve $\mathrm{layout}(\Delta X) = \mathrm{layout}(\Delta Y)$.
- ΔY is a uniformly random difference from the layout.

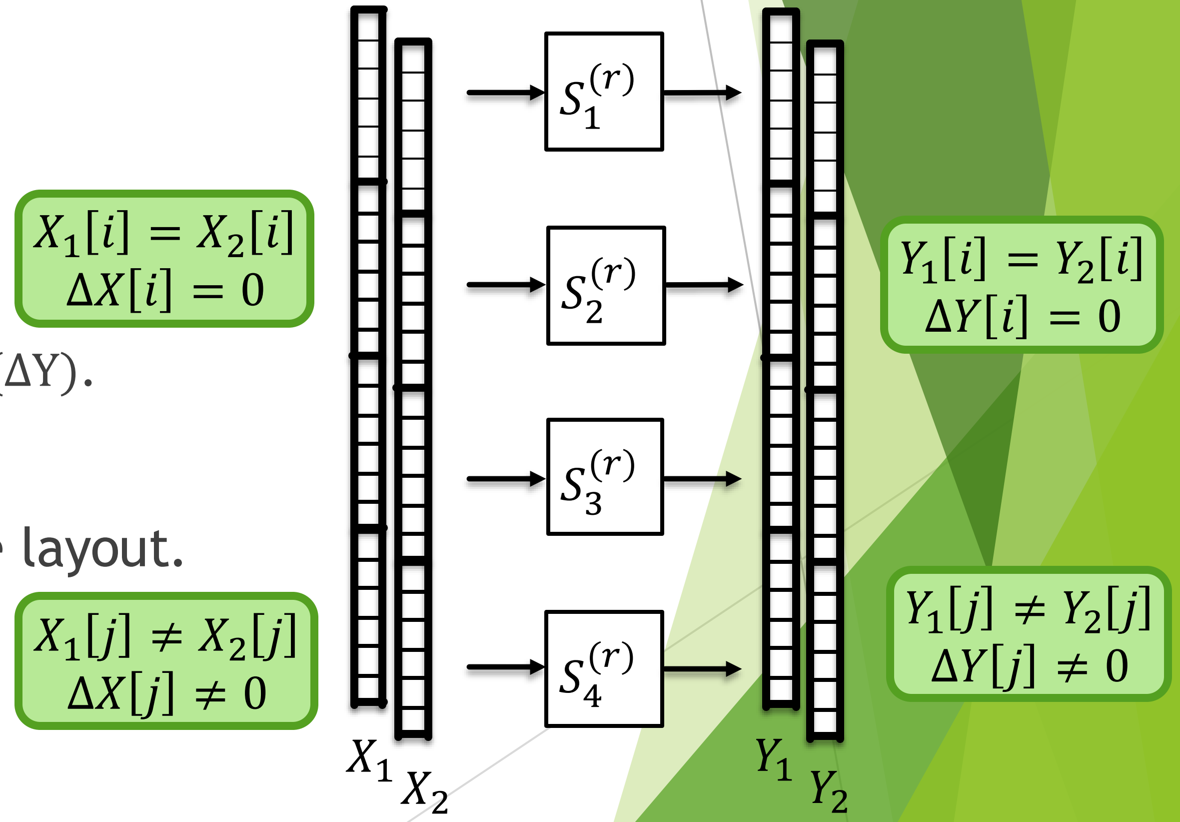

Layouts: example

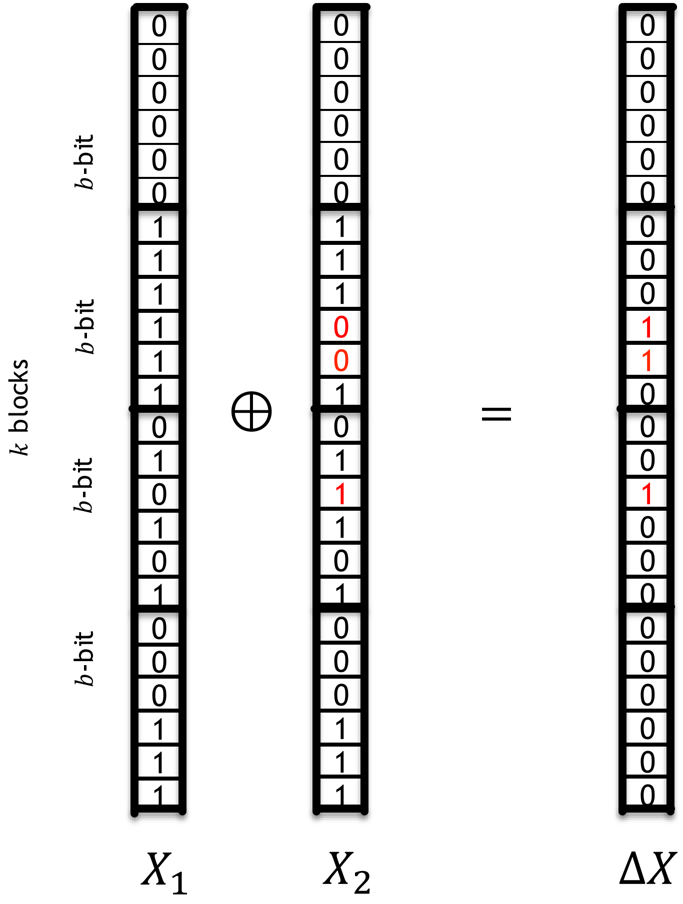

Layout
$\mathrm{layout}(\Delta X) = \{0, 3\}$

Define weight of layout
$|\mathrm{layout}(\Delta X)| = 2$

Lower weight ⇒ more distinct blocks

Layout Graph

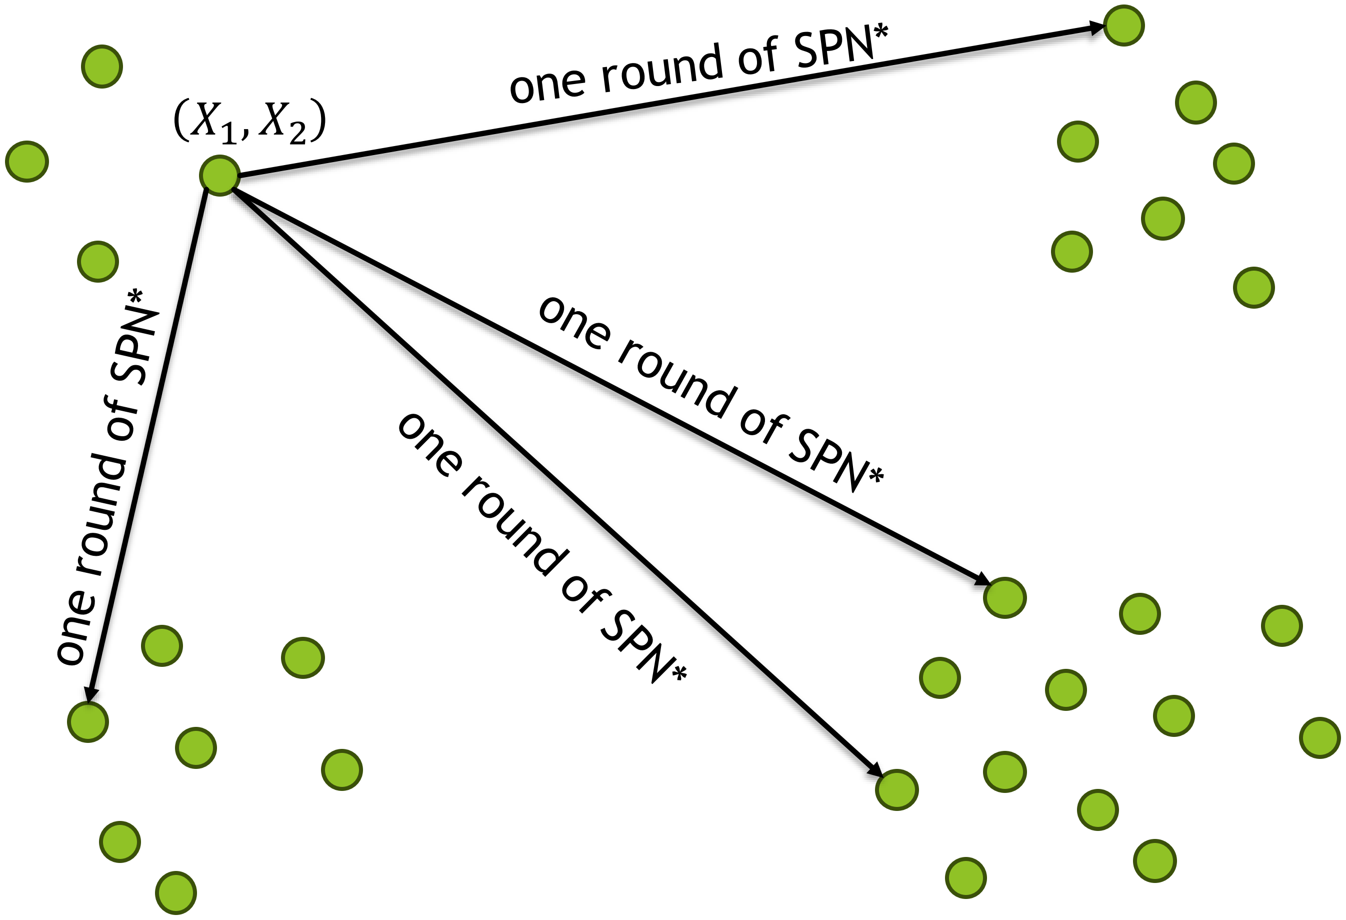

Each pair of inputs is a node

One round of SPN* will change $(X_1, X_2) \rightarrow$ another pair

Layout Graph

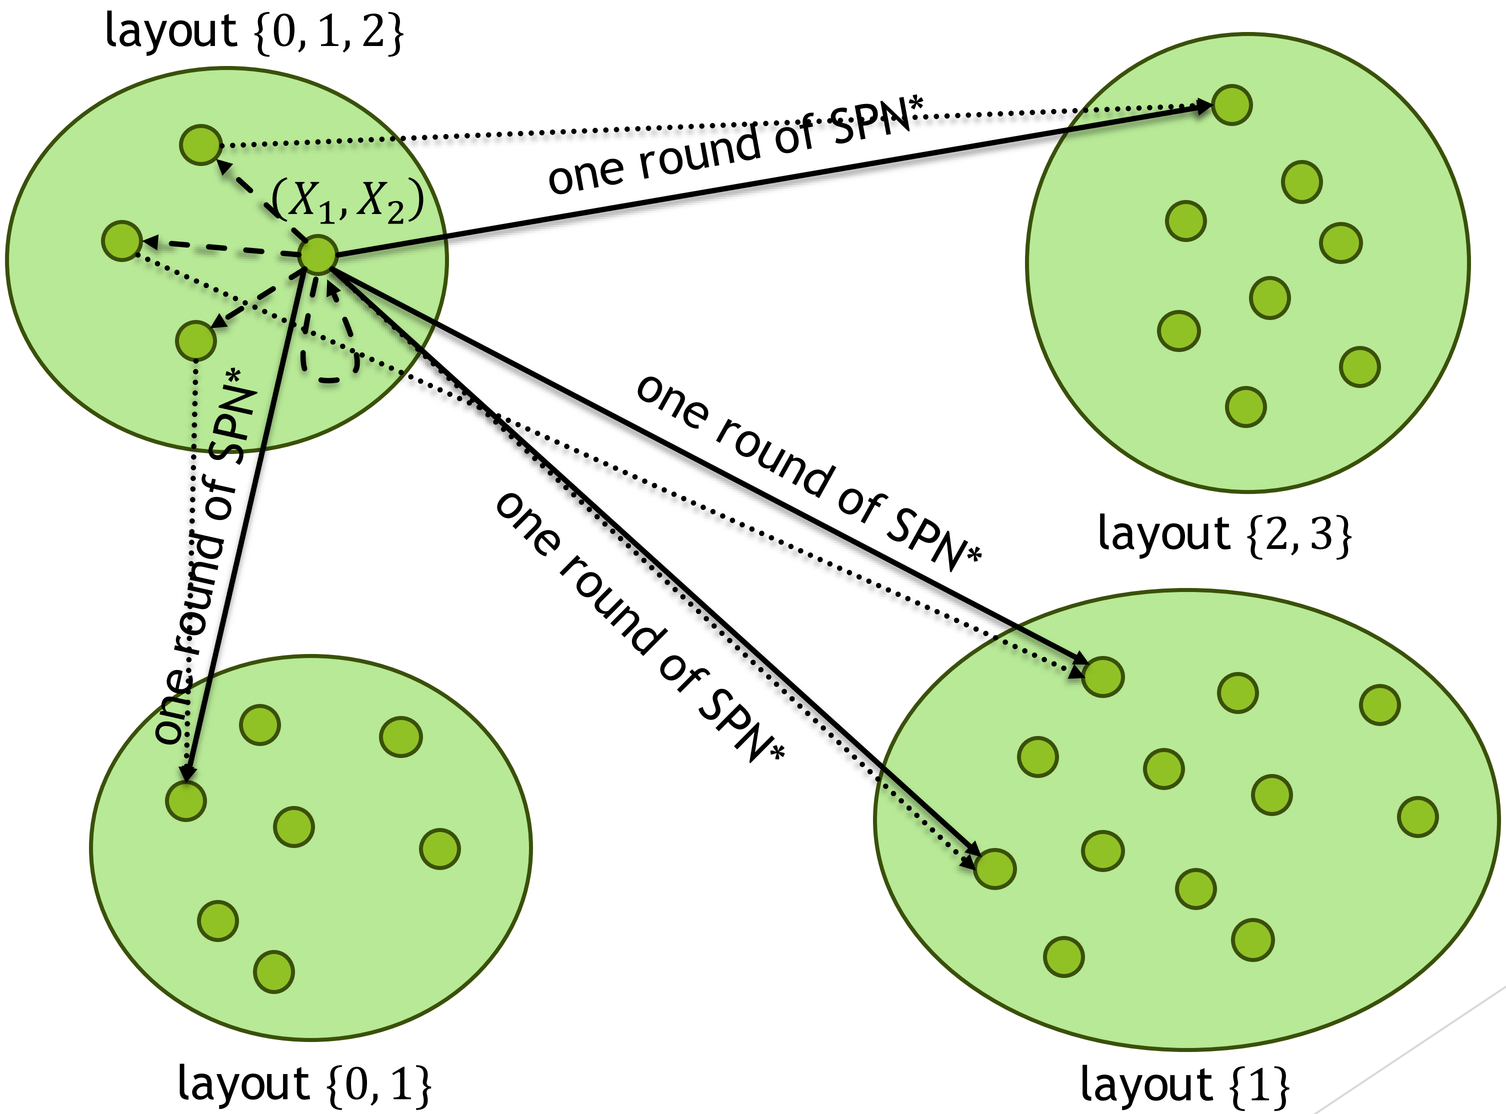

<u>SPN* round</u>
random S-boxes + linear mixing

1. Random S-boxes choose uniform node in the layout

2. Linear mixing maps each node to another node (possibly outside the layout)

<u>Crucial</u>
Random S-boxes randomize within a layout

Layout Graph

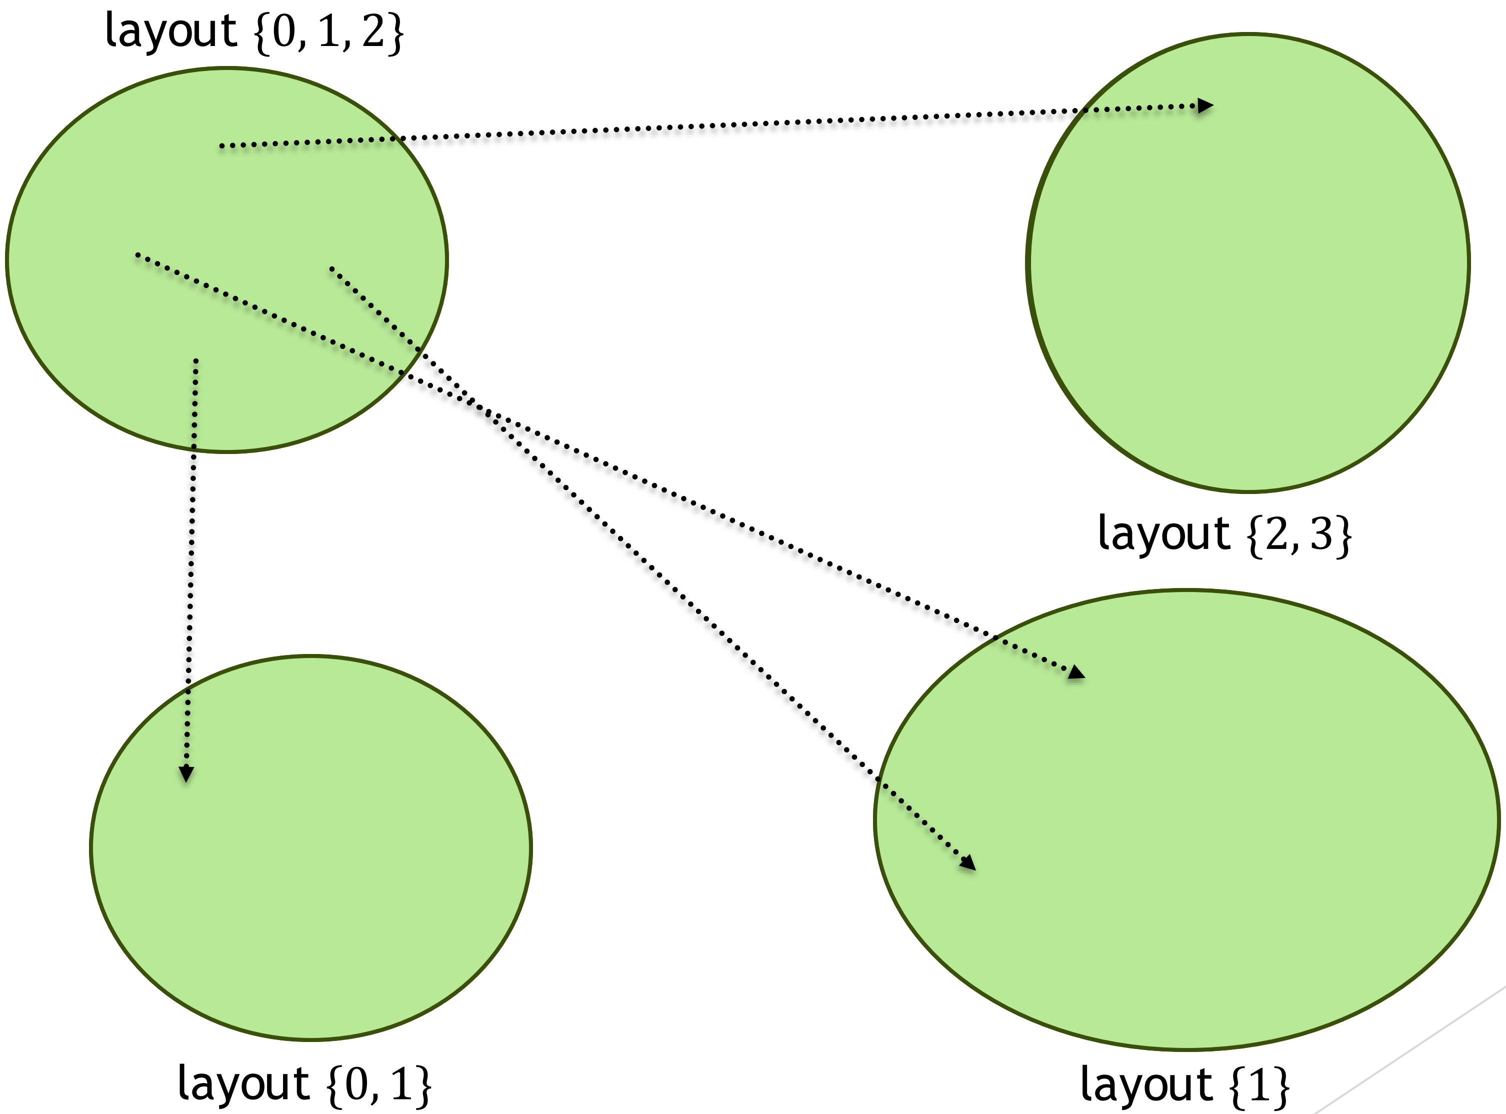

Crucial
Random S-boxes randomize within a layout

Layout Walk

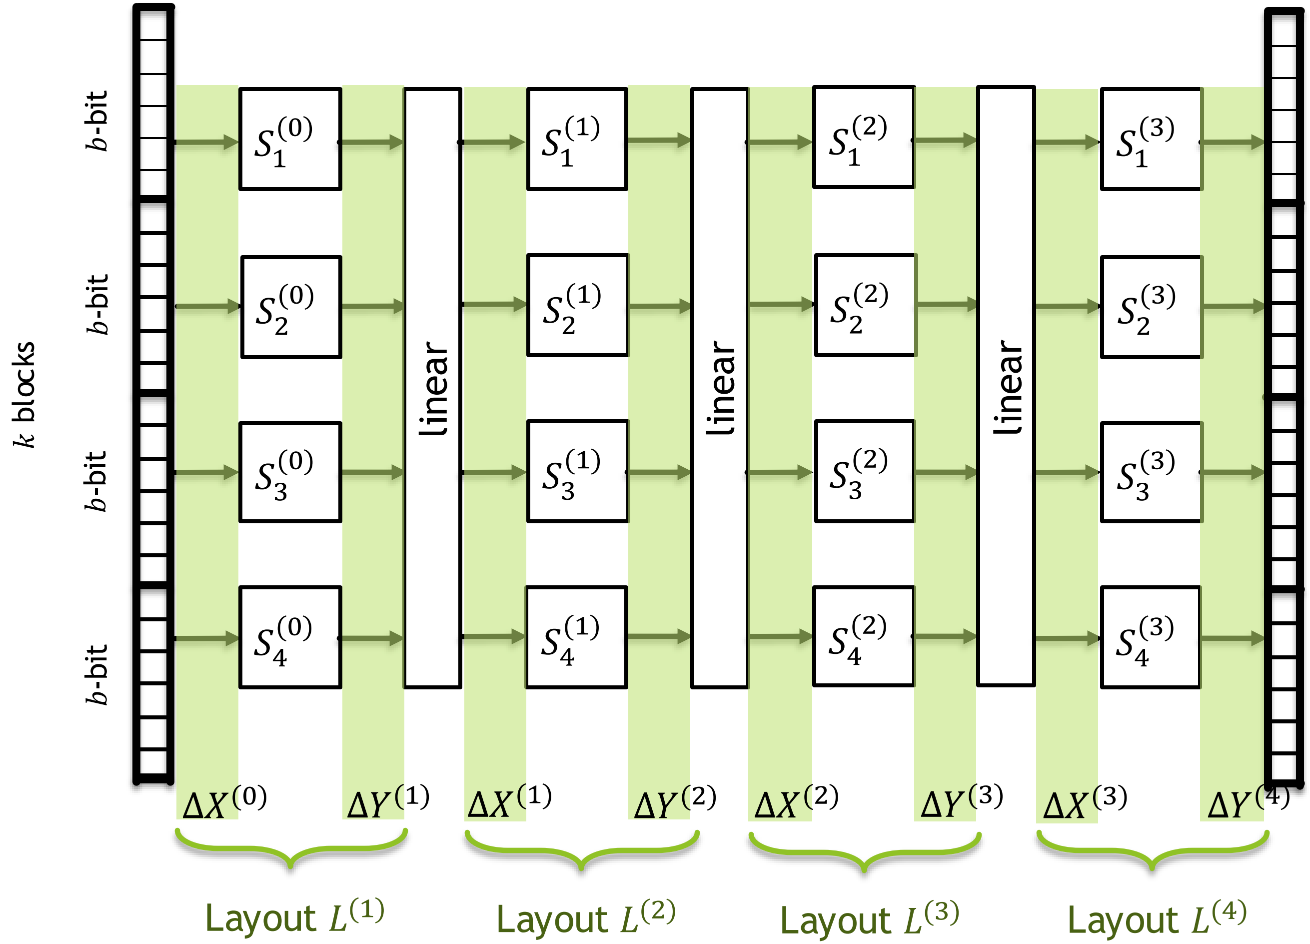

<u>Simplified Problem</u>
Only consider the layouts

<u>Define</u>
Random walk on layout graph $L^{(1)} \to L^{(2)} \to L^{(3)} \to \cdots$

<u>Suffices</u>
Bound the mixing time of the layout graph.

<u>First step</u>
Understand how mixing affects the layout.

<u>Recall</u>
$\Delta Y^{(i)}$ is sampled uniformly from $L^{(i)}$.

Layout Transition Probability

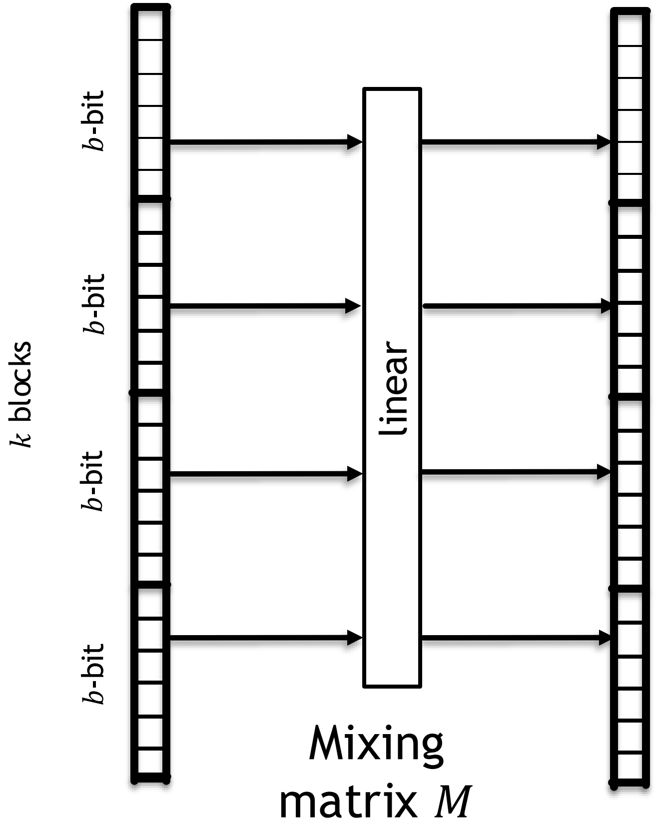

$\Delta Y \sim$ layout I

$\Delta X = M \cdot \Delta Y$
layout J?

Sample input difference
ΔY from layout I

First step
Understand how mixing affects the layout.

Define
$\mathbb{P}[\text{layout } I \rightarrow \text{layout } J] \coloneqq \mathbb{P}_{\Delta Y \text{ in } I}[M \cdot \Delta Y \text{ in } J]$

$\mathbb{P}[\Delta X \text{ in } J]$ after mixing?
- Depends on mixing matrix
- Assume maximal branch number

Layout Transition Probability

Define
$\mathbb{P}[\text{layout } I \to \text{layout } J] \coloneqq \mathbb{P}_{\Delta Y \text{ in } I}[M \cdot \Delta Y \text{ in } J]$

$$\mathbb{P}[I \to J] = \frac{\#[\Delta Y \text{ in } I \mid M \cdot \Delta Y \text{ in } J]}{\#[\Delta Y \text{ in } I]}$$

$$= \frac{\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \mathbb{I}[M \cdot \Delta Y = \Delta X]}{(2^b - 1)^{k - |I|}}$$

It holds that

- $\#[\Delta Y \text{ in } I] = \left(2^b - 1\right)^{k - |I|}$
- $\#[\Delta Y \text{ in } I \mid M \cdot \Delta Y \text{ in } J] = \sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \mathbb{I}[M \cdot \Delta Y = \Delta X]$

Goal
Bound $\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \mathbb{I}[M \cdot \Delta Y = \Delta X]$

Layout Transition Probability

<u>Goal</u>

Bound $\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \mathbb{I}[M \cdot \Delta Y = \Delta X]$

<u>Define</u>

"ΔX sat I" if $\forall i \in I, \Delta X[i] = 0$

<u>Contrast with</u>

"ΔX in I" if $\begin{cases} \forall i \in I, \Delta X[i] = 0 \\ \forall i \notin I, \Delta X[i] \neq 0 \end{cases}$

Easier to work with

solutions to system of linear equations

$$\sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \mathbb{I}[M \cdot \Delta Y = \Delta X] = \begin{cases} (2^b)^{k-|I|-|J|}, & |I| + |J| \leq k \\ 1, & \text{otherwise} \end{cases}$$

$$\sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \frac{1}{2^{bk}} = (2^b)^{k-|I|-|J|}$$

$$\sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \mathbb{I}[M \cdot \Delta Y = \Delta X] = \sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \frac{1}{2^{bk}} \pm (\text{up to } 1)$$

Layout Transition Probability

$$\sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \mathbb{I}[M \cdot \Delta Y = \Delta X] = \sum_{\substack{\Delta Y \text{ sat } I \\ \Delta X \text{ sat } J}} \frac{1}{2^{bk}} \pm (\text{up to } 1)$$

<u>Inclusion-exclusion principle</u>
relates (ΔY in I) with (ΔY sat I)

$$\sum_{\substack{\Delta Y \text{ in } I \\ \Delta X \text{ in } J}} \mathbb{I}[M \cdot \Delta Y = \Delta X] = \sum_{\substack{\Delta Y \text{ in } I \\ \Delta X \text{ in } J}} \frac{1}{2^{bk}} \pm (\text{up to } 2^k)$$

Layout Transition Probability

$$\mathbb{P}[I \to J] = \frac{\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \mathbb{1}[M \cdot \Delta Y = \Delta X]}{\sum_{\Delta Y \text{ in } I} 1}$$

$$= \frac{\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \frac{1}{2^{bk}} \pm (\text{up to } 2^k)}{\sum_{\Delta Y \text{ in } I} 1}$$

$$= \frac{\sum_{\Delta Y \text{ in } I} \sum_{\Delta X \text{ in } J} \frac{1}{2^{bk}}}{\sum_{\Delta Y \text{ in } I} 1} \pm \frac{(\text{up to } 2^k)}{\sum_{\Delta Y \text{ in } I} 1}$$

$$= \underbrace{\sum_{\Delta X \text{ in } J} \frac{1}{2^{bk}}}_{\mathbb{P}_{\Delta X}[\Delta X \text{ in } J]} \pm \underbrace{\frac{(\text{up to } 2^k)}{2^{b(k-|I|)}}}_{\frac{1}{2^{\Theta(bk)}} \text{ if } |I| \leq \frac{k}{2}}$$

Lemma. If $|I| \leq \frac{k}{2}$, distribution of layout after mixing is $\frac{1}{2^{\Theta(bk)}}$-close to stationary.

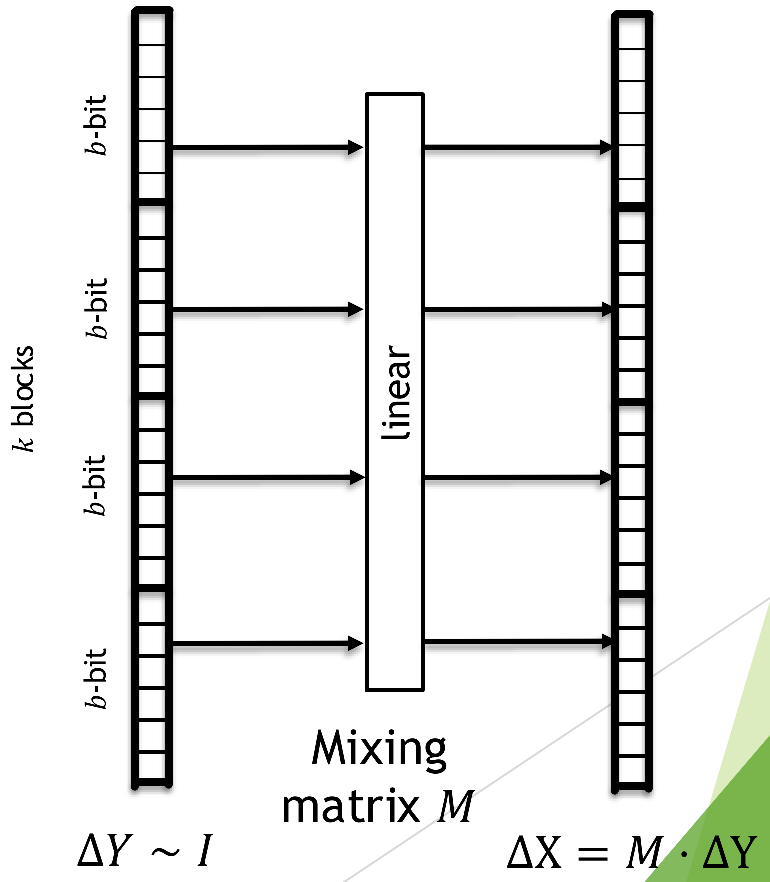

Proof Overview (2-wise)

Lemma. If $|I| \leq \frac{k}{2}$, distribution of layout after mixing is $\frac{1}{2^{\Theta(bk)}}$-close to stationary.

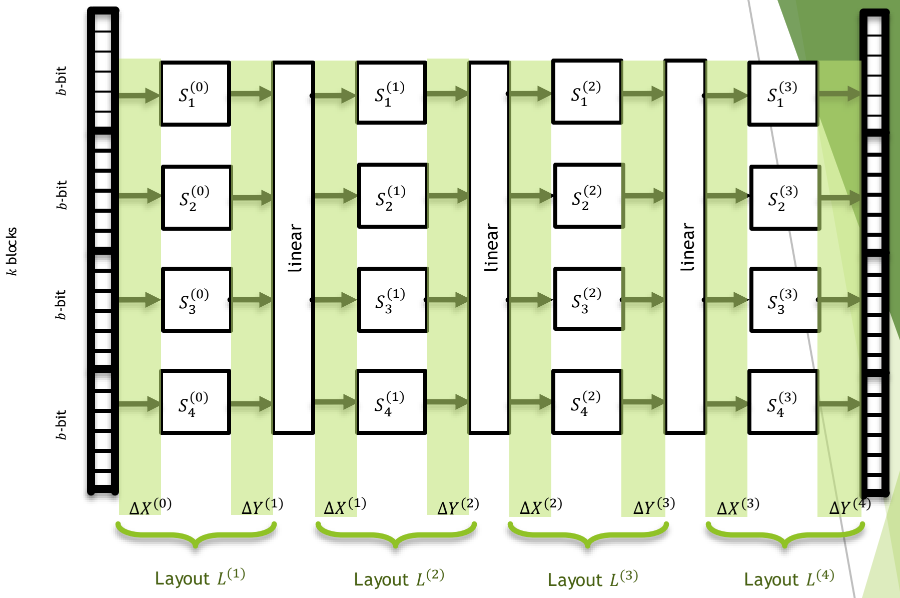

- $|L^{(1)}| \leq \frac{k}{2} \Rightarrow L^{(2)}$(actually $\Delta Y^{(2)}$) is $\frac{1}{2^{\Theta(bk)}}$-close to stationary

- $|L^{(1)}| \geq \frac{k}{2} \Rightarrow |L^{(2)}| \leq \frac{k}{2}$

  $\Rightarrow L^{(3)}$ (actually $\Delta Y^{(3)}$) is $\frac{1}{2^{\Theta(bk)}}$ −close to stationary

Proof Overview (t-wise)

- $X_a[i] = X_b[i] \Leftrightarrow Y_a[i] = Y_b[i]$.
- Generalize to t-wise layouts.
  - Remember whether $X_a[i] = X_b[i]$.

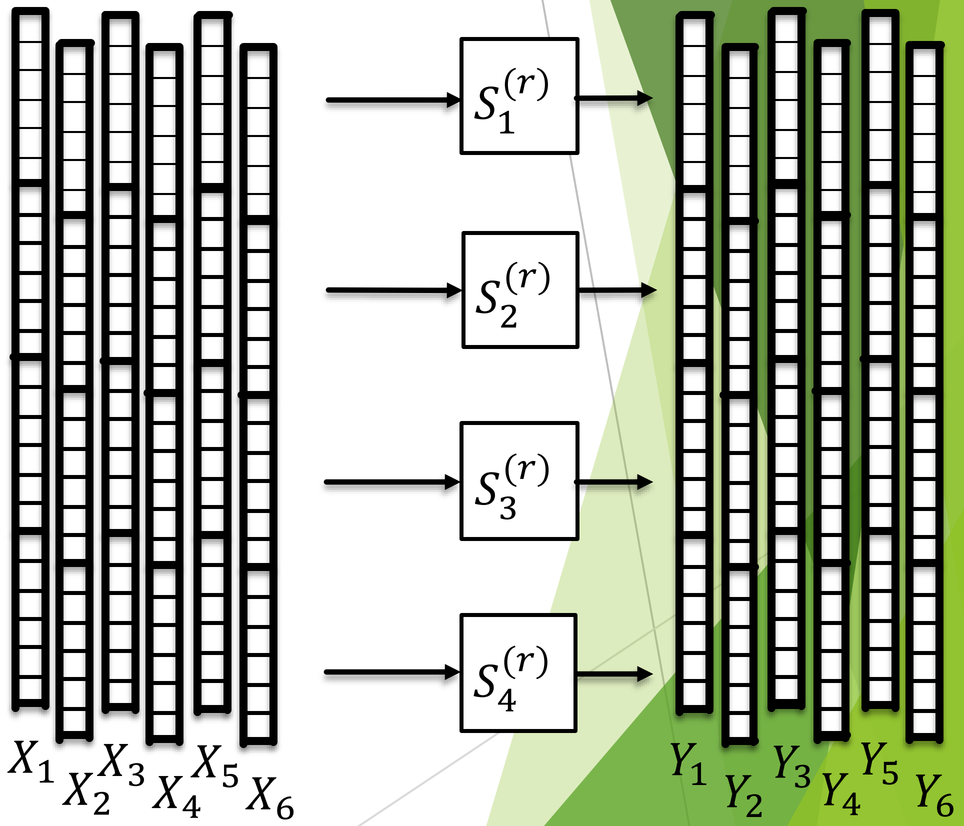

Part IIb. Censored AES Technical Details

Pairwise Independence of AES*

- We saw before a way to approximate the transition probability $\mathbb{P}[I \to J]$ up to some small error.

- Turns out that we can compute $\mathbb{P}[I \to J]$ exactly:

Lemma [BV06]. If M has maximal branch number, the layout transition probability equals

$$\mathbb{P}[I \to J] = \mathbb{P}_{\Delta Y \text{ in } I}[M \cdot \Delta Y \text{ in } J] = \sum_{i=0}^{|I|+|J|-k-1} (-1)^i \frac{\binom{k-1+i}{k-1}}{(2^b-1)^{k-|J|+i}}.$$

Pairwise Independence of AES*

Lemma [BV06]. If M has maximal branch number, the layout transition probability equals

$$\mathbb{P}[I \to J] = \mathbb{P}_{\Delta Y \text{ in } I}[M \cdot \Delta Y \text{ in } J] = \sum_{i=0}^{|I|+|J|-k-1} (-1)^i \frac{\binom{k-1+i}{k-1}}{(2^b-1)^{k-|J|+i}}.$$

- **Issue.** The AES* mixing does not have maximal branch number.
- Still possible to compute the exact adjacency matrix T of the layout graph.
  - Size $(2^{16}-1) \times (2^{16}-1) \approx 65K \times 65K$.
  - TOO LARGE!
  - Turns out T has rank $5^4 = 625$, and thus we can compute powers of T

Pairwise Independence of AES*

- Can numerically compute the exact convergence to pairwise independence.

Theorem. The 7-round AES* is 2^{-128}-close to pairwise independent.

- For censored AES, we will also need the following result:

Theorem. The 3-round AES* is $2^{-23.42}$-close to pairwise independent.

Pairwise Independence of censored AES

Theorem. The 3-round AES* is $2^{-23.42}$-close to pairwise independent.

- **Question.** How to go from random S-box to INV S-box?

Lemma.

$$\underbrace{ARK \circ INV \circ \cdots \circ ARK \circ INV}_{8 \text{ times}} \approx_{2^{-29.39}} \text{random permutation over } \mathbb{F}_{2^8}$$

Pairwise Independence of censored AES

Start with a 3-round AES*

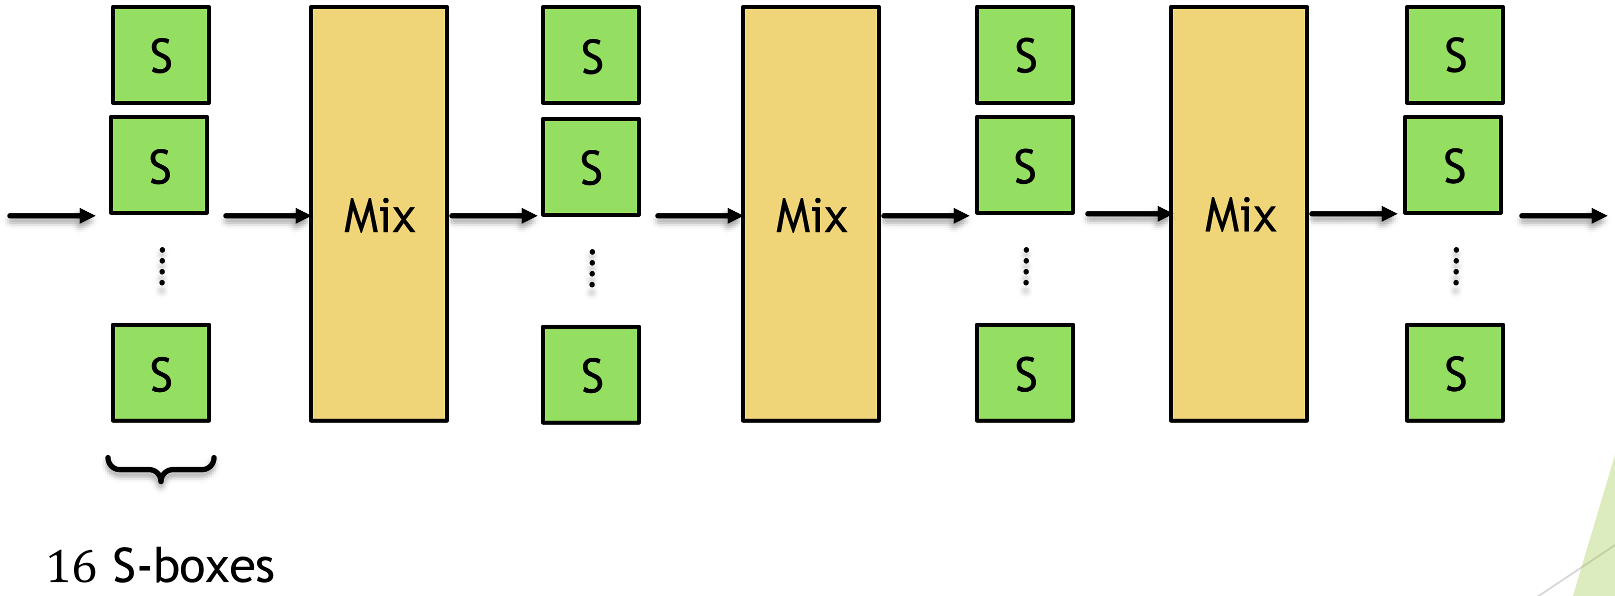

3-round AES* $\approx_{2^{-23.42}}$ pairwise independent

Pairwise Independence of censored AES

Replace each random S-box with 8 rounds of INV S-boxes

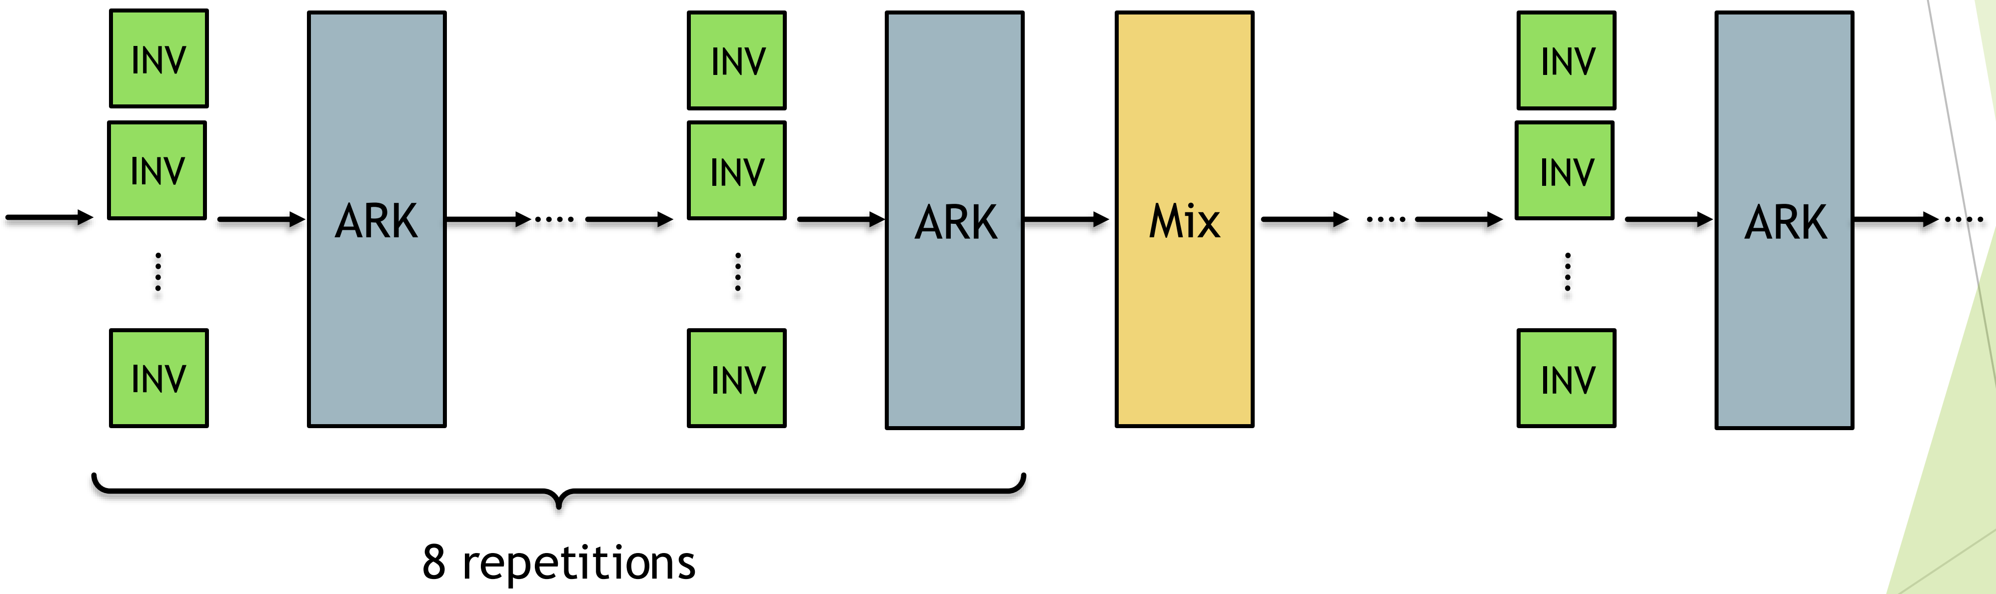

A total of $4 \cdot 8 = 32$ rounds of INV S-boxes

Pairwise Independence of censored AES

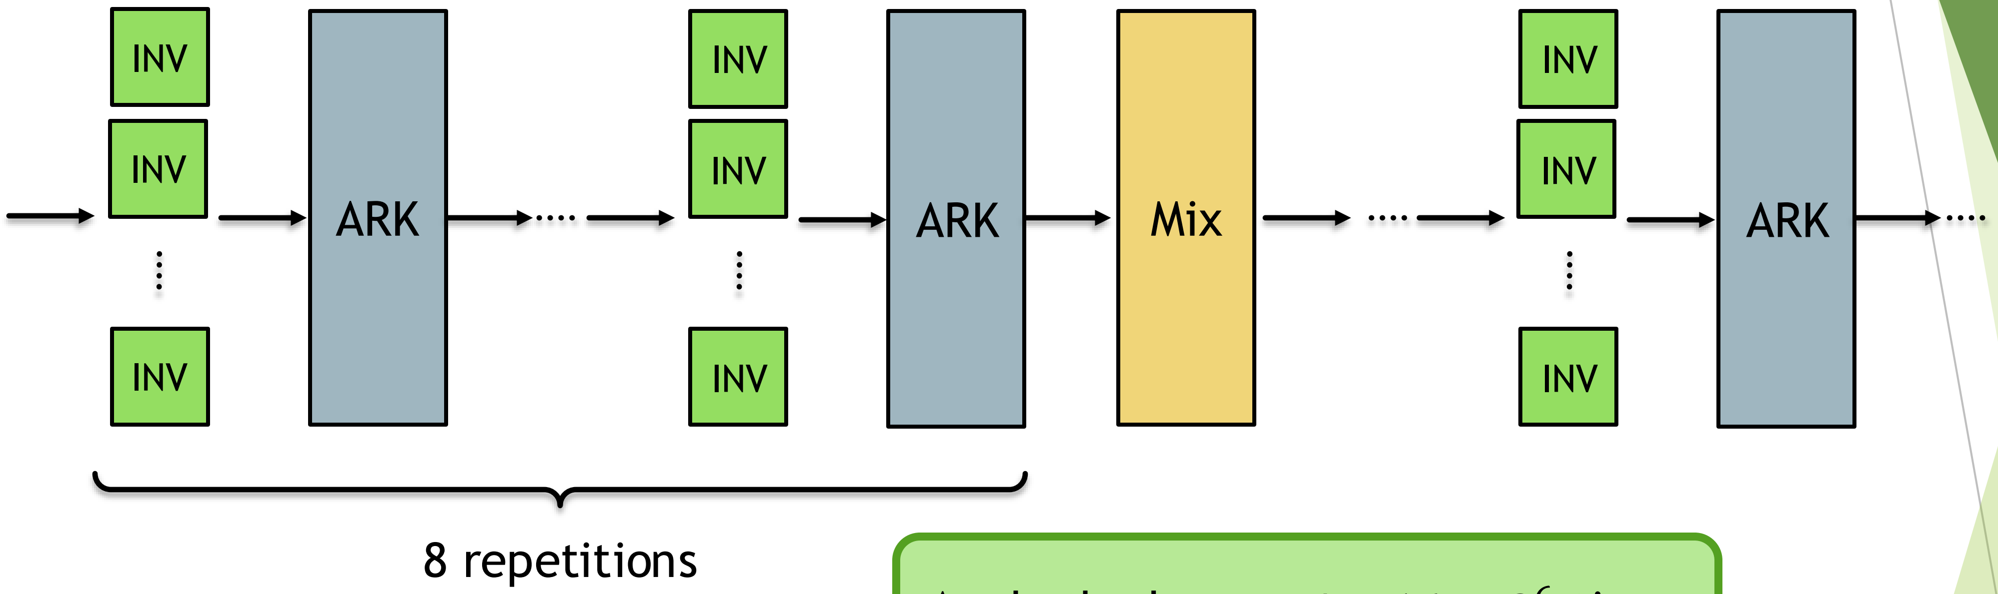

Apply the lemma $4 \times 16 = 2^6$ times

Lemma.

$$\underbrace{ARK \circ INV \circ \cdots \circ ARK \circ INV}_{8\,\text{times}} \approx_{2^{-29.39}} \text{random permutation over } \mathbb{F}_{2^8}$$

32-round censored AES $\approx_{2^{-23.39}}$ 3-round AES*

Pairwise Independence of censored AES

Note that this is censored AES!

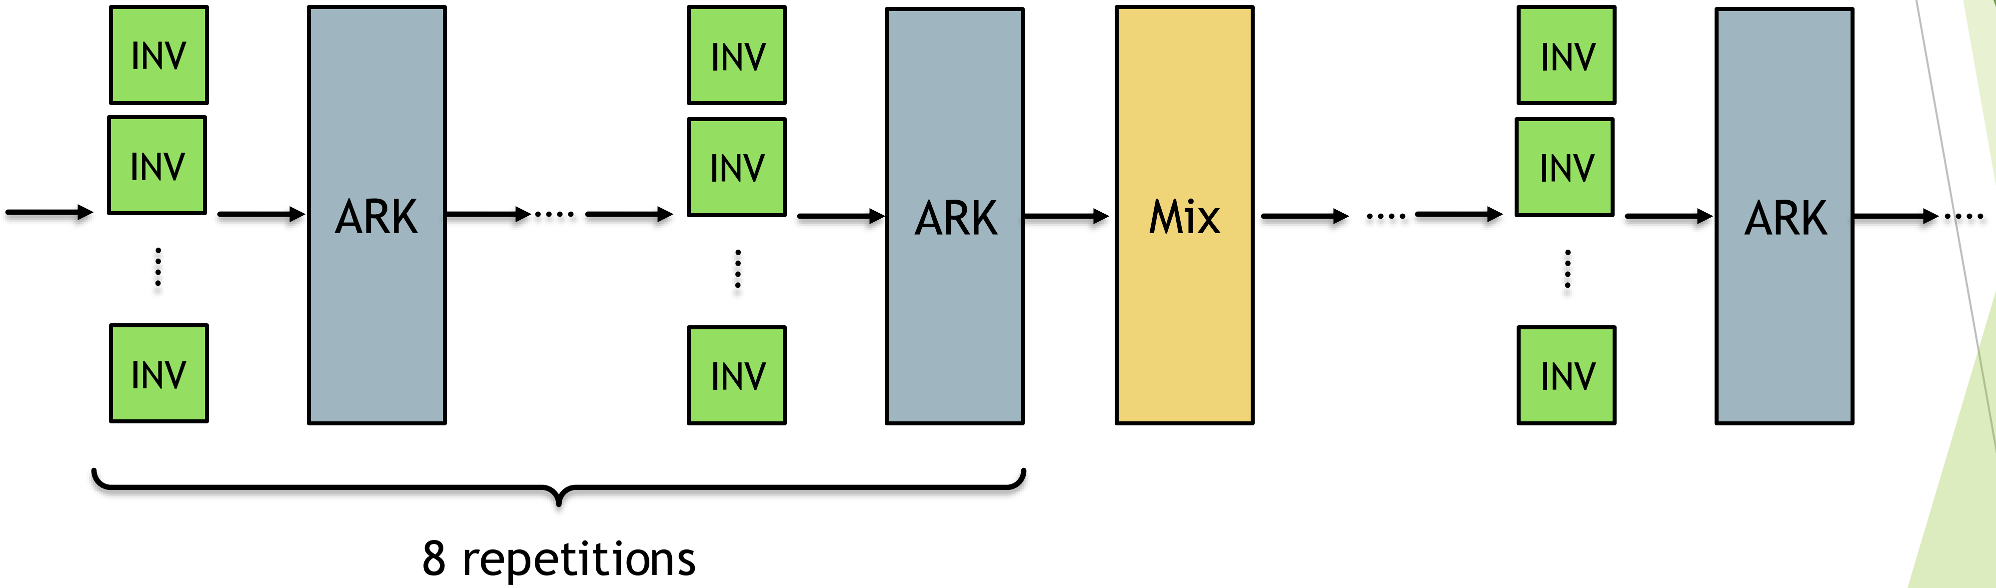

32-round censored AES $\approx_{2^{-23.39}}$ 3-round AES* $\approx_{2^{-23.42}}$ pairwise independent

Pairwise Independence of censored AES

32-round censored AES $\approx_{2^{-23.39}}$ 3-round AES* $\approx_{2^{-23.42}}$ pairwise independent

$\Rightarrow$ 32-round censored AES $\approx_{2^{-22.39}}$ pairwise independent

Amplification Lemma [MPR07]

$\mathcal{F}$ is ϵ-close to t-wise independent
$\Rightarrow \mathcal{F} \circ \mathcal{F}$ is $2\epsilon^2$-close to t-wise independent.

6 repetitions of the 32-round censored AES is $2^5 \cdot (2^{-22.39})^6 < 2^{-128}$

$\Rightarrow$ 192-round censored AES is pairwise independent!

Part III. Reversible circuits

An emerging block cipher: Reversible circuits

- Another way to a create a keyed-permutation on $\{0,1\}^n$.

n wires

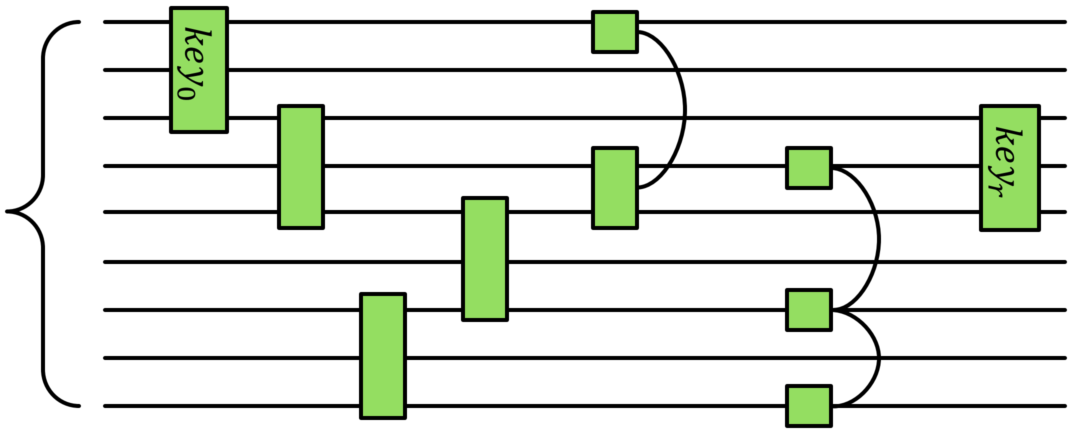

- Each “round” is now a random 3-bit gate.
- The secret key includes the gates.
  - key_0 is a random permutation of $\{0,1\}^3$ and the wires it acts on.

Reversible circuits

- We can now ask the same security question as for SPNs:

n wires

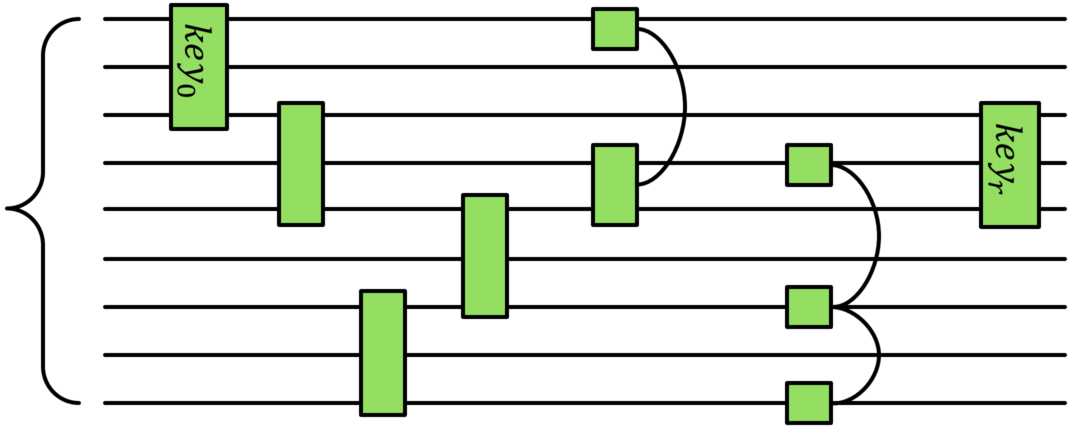

How many rounds (gates) do we need to obtain a t-wise independent permutation?

Background on reversible circuits

- Introduced by Gowers [Gow96] who wanted to study P vs NP.
  - Subsequent work by [HMMR05, BH08] shows that $\tilde{O}\left(n^2 t^2 \cdot \log \frac{1}{\varepsilon}\right)$ gates suffice.
- Quantum physicists [BHH16, HHJ21] study *random quantum circuits*.
  - Connections to quantum pseudorandomness, black holes, many-body systems...
- More than just encryption [CCMR24].
  - Reversible circuits are pseudorandom $\Rightarrow$ candidate obfuscation schemes.
  - Inspired by the thermalizing processes of statistical mechanics.

Background on reversible circuits

- He and O’Donnell [HO24] also study circuits with nearest-neighbor gates.

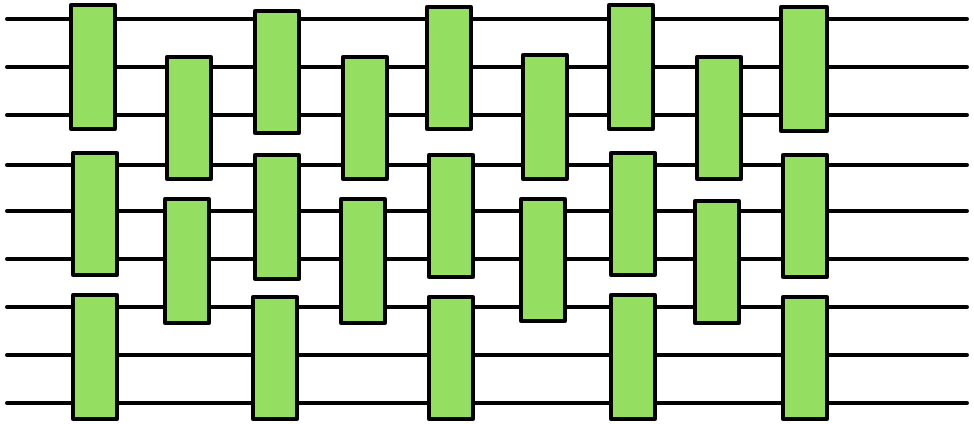

  - A more practical construction.

t-wise independence of reversible circuits

Theorem [GHP24]. For $t \leq 2^{n/50}$, a random reversible circuit with $\tilde{O}\left(nt \cdot \log\frac{1}{\varepsilon}\right)$ gates is ε-close to t-wise independent.

- Our analysis uses log-Sobolev inequalities, instead of spectral gaps.
  - Avoids the extra factors from prior work.

- **Optimal** up to polylogs for constant ε.

t-wise independence of reversible circuits

Theorem [GHP24]. For $t \leq 2^{n/50}$, a random reversible circuit with $\tilde{O}\left(nt \cdot \log\frac{1}{\varepsilon}\right)$ gates is ε-close to t-wise independent.

- Nice result for free:

Corollary. A random circuit with $L \leq 2^{n/50}$ gates cannot be compressed to less than $\frac{L}{n^3 \log n}$ gates whp.

  - Pointed to us by [CHH+24].
  - Our bounds imply *incompressibility* of random circuits.

Part IV. Open questions

Many t-wise independence questions remain…

- t-wise independence of any (non-idealized) block cipher for $t > 2$?
- Improved AES analysis?
  - Prove that real AES is at least as secure as censored AES?
- t-wise independence of SPN* beyond $t = \sqrt{2^b}$.
  - Can we push t up to $2^{\Theta(kb)}$ (or even 2^b)?

… and many block cipher questions remain!

- Are reversible circuits pseudorandom?
- Study other classes of attacks.
  - e.g., algebraic attacks via the Polynomial Calculus proof system [AL15].
- The role of key scheduling.
  - Given a secure block cipher with independent keys, what key scheduler preserves its security?

Research Program Goals

- Continue a research program put forward by [LTV21].

- Goal i
  - t-wi
  - Man

Goal

Security of practical encryption schemes from a theoretical viewpoint

- Solving these problems is an important quest.
  - Likely requires new techniques from mathematics and TCS.

Thank you!

Questions?