

CS 594 – ADVANCED CRYPTO (SPRING 2026)

Alex Block


Lecture 21

April 08, 2026


MEMORY-HARD FUNCTIONS

MOTIVATION: PASSWORD STORAGE


MOTIVATION: PASSWORD STORAGE




MOTIVATION: PASSWORD STORAGE




MOTIVATION: PASSWORD STORAGE



MOTIVATION: PASSWORD STORAGE




MOTIVATION: PASSWORD STORAGE



Username	Salt	Hash
arblock	89d9 7803 4a3f6	946f f6e7 599f 342a 0a9b 9ae1 e5d3 dece 9091 52d6


MOTIVATION: PASSWORD STORAGE



Username	Salt	Hash
arblock	89d9 7803 4a3f6	946f f6e7 599f 342a 0a9b 9ae1 e5d3 dece 9091 52d6

$$\text{SHA1}(12345689d978034a3f6) \\ = 946ff6e7599f342a0a9b9ae1e5d3dece909152d6$$

MOTIVATION: PASSWORD STORAGE



(arblock, 123456)

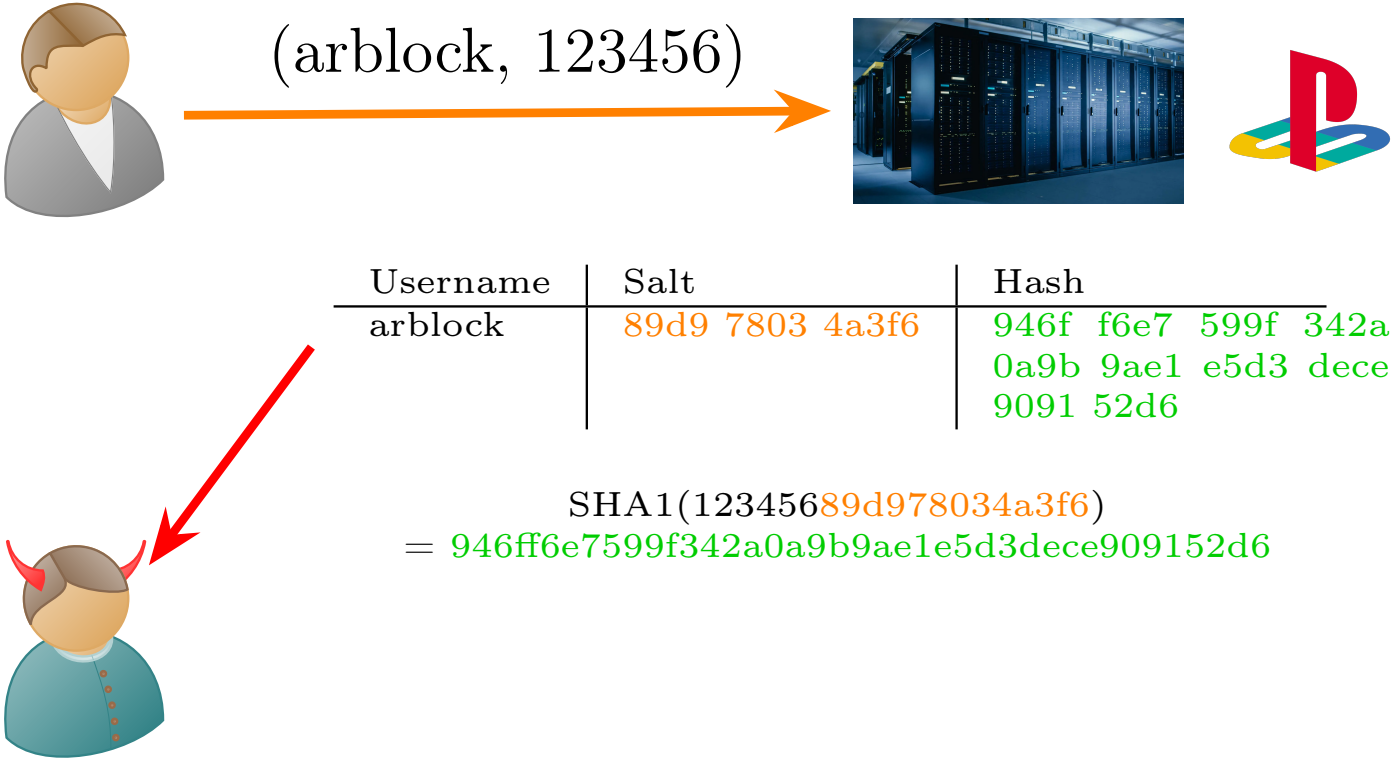


Username	Salt	Hash
arblock	89d9 7803 4a3f6	946f f6e7 599f 342a 0a9b 9ae1 e5d3 dece 9091 52d6


$$\text{SHA1}(12345689d978034a3f6) \\ = 946ff6e7599f342a0a9b9ae1e5d3dece909152d6$$




MOTIVATION: PASSWORD STORAGE



MOTIVATION: PASSWORD STORAGE



Username	Salt	Hash
arblock	89d9 7803 4a3f6	946f f6e7 599f 342a 0a9b 9ae1 e5d3 dece 9091 52d6




$$\text{SHA1}(12345689d978034a3f6) = 946ff6e7599f342a0a9b9ae1e5d3dece909152d6$$



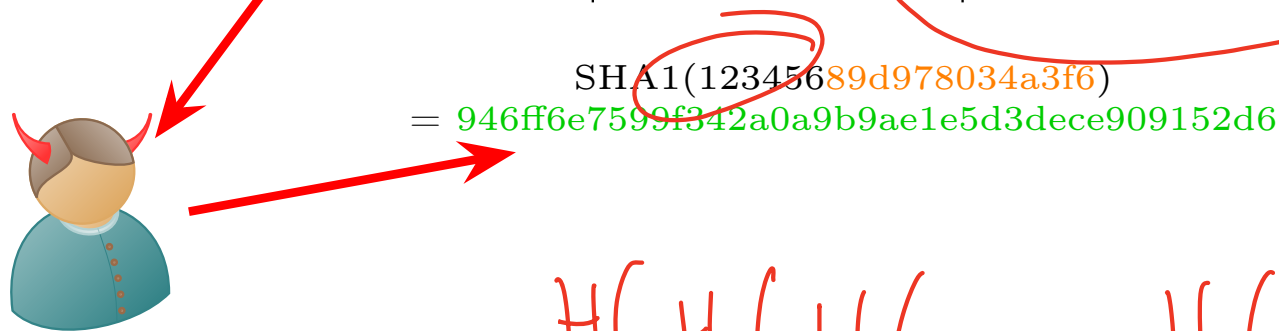
+



MOTIVATION: PASSWORD STORAGE



Username	Salt	Hash
arblock	89d9 7803 4a3f6	946f f6e7 599f 342a 0a9b 9ae1 e5d3 dece 9091 52d6



$H(H(H(\dots H(\text{username} \parallel \text{salt}) \dots)))$



+



MOTIVATION: OFFLINE ATTACKS ARE A COMMON PROBLEM

MOTIVATION: OFFLINE ATTACKS ARE A COMMON PROBLEM

- Password breaches at major companies have affected **billions** of user accounts.

MOTIVATION: OFFLINE ATTACKS ARE A COMMON PROBLEM

- Password breaches at major companies have affected **billions** of user accounts.

LastPass 

SONY

ebay

ASHLEY
MADISON®
Life is short. Have an affair.®

Linked in


Dropbox

AdultFriendFinder®

Zappos
.com 
the web's most popular shoe store!

rockyou

YAHOO!

 Adobe



livingsocial 

MOTIVATION: OFFLINE ATTACKS ARE A COMMON PROBLEM

- Password breaches at major companies have affected **billions** of user accounts.

LastPass 

SONY

ebay

ASHLEY
MADISON®
Life is short. Have an affair.®

Linked in


Dropbox

AdultFriendFinder®

Zappos
.com 
the web's most popular shoe store!

rockyou

YAHOO!

 Adobe



livingsocial 

<https://haveibeenpwned.com/>

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

Can we design a *moderately expensive* hash function?

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

Can we design a *moderately expensive* hash function?

- *Fast* to compute on regular hardware (e.g., desktop, laptop, or phone); and

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

Can we design a *moderately expensive* hash function?

- *Fast* to compute on regular hardware (e.g., desktop, laptop, or phone); and
- *Expensive* to compute on an ASIC?

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

Can we design a *moderately expensive* hash function?

- *Fast* to compute on regular hardware (e.g., desktop, laptop, or phone); and
- *Expensive* to compute on an ASIC?
 - *Application-specific Integrated circuits.*

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

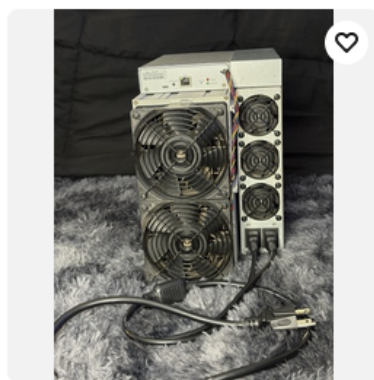
Can we design a *moderately expensive* hash function?

- *Fast* to compute on regular hardware (e.g., desktop, laptop, or phone); and
- *Expensive* to compute on an ASIC?
 - *Application-specific Integrated circuits.*
 - Became more widely available for computing hashes like SHA-3 or SHA-256 because of Bitcoin and friends.

GOAL: MODERATELY EXPENSIVE HASH FUNCTIONS

Can we design a *moderately expensive* hash function?

- *Fast* to compute on regular hardware (e.g., desktop, laptop, or phone); and
- *Expensive* to compute on an ASIC?
 - *Application-specific Integrated circuits.*
 - Became more widely available for computing hashes like SHA-3 or SHA-256 because of Bitcoin and friends.



NEW LOW PRICE

Bitmain Antminer S19 90TH/s Perfect Condition With PSU + Power Cables
New (Other)

\$162.00

25 bids · 17m left (Today 05:14 PM)

Free delivery

Located in United States

trailvault 97.1% positive (79)

MEMORY-HARD FUNCTIONS

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost

Space cost

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Space cost

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Space cost



MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Need 2× the time?
Leave computer on!

Space cost



MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Need $2\times$ the time?
Leave computer on!

Space cost



Need $2\times$ the space?
Download more RAM!

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Need 2× the time?
Leave computer on!

Space cost



Need 2× the space?
Download more RAM!

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Need 2× the time?
Leave computer on!

Space cost



Need 2× the space?
Download more RAM!

Goal

MEMORY-HARD FUNCTIONS

- Intuition: force computational cost to be dominated by *memory*.
 - CPU time is *cheap*; adding more RAM is *expensive*.

Time cost



Need 2× the time?
Leave computer on!

Space cost



Need 2× the space?
Download more RAM!

Goal

Force attacker to use *large amounts of memory* for duration of computation, even on customized hardware.

MEMORY-HARD FUNCTIONS: TWO FLAVORS

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

Data-independent MHF (iMHF)

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Data-independent MHF (iMHF)

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Data-independent MHF (iMHF)

- The structure and output of a iMHF f is *independent* of *any* input x being evaluated.

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Data-independent MHF (iMHF)

- The structure and output of a iMHF f is *independent* of *any* input x being evaluated.

Pros

- Best “memory-hard” guarantees

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHF's come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Data-independent MHF (iMHF)

- The structure and output of a iMHF f is *independent* of *any* input x being evaluated.

Pros

- Best “memory-hard” guarantees

Cons

- Susceptible to side-channel attacks!

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Pros

- Best “memory-hard” guarantees

Cons

- Susceptible to side-channel attacks!

Data-independent MHF (iMHF)

- The structure and output of a iMHF f is *independent* of *any* input x being evaluated.

Pros

- Highly resistant to side-channel attacks

MEMORY-HARD FUNCTIONS: TWO FLAVORS

- MHFs come in two flavors:
 - data-*dependent*; and
 - data-*independent*

Data-dependent MHF (dMHF)

- The structure and output of a dMHF f *depends* on the input x being evaluated.

Pros

- Best “memory-hard” guarantees

Cons

- Susceptible to side-channel attacks!

Data-independent MHF (iMHF)

- The structure and output of a iMHF f is *independent* of *any* input x being evaluated.

Pros

- Highly resistant to side-channel attacks

Cons

- Provably worse “memory-hardness”

MHFS FROM DIRECTED GRAPHS

MHFS FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.

↳ "Fancy" hash-chaining

MHFS FROM DIRECTED GRAPHS


- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.

MHFS FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.


MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.



MHFs FROM DIRECTED GRAPHS

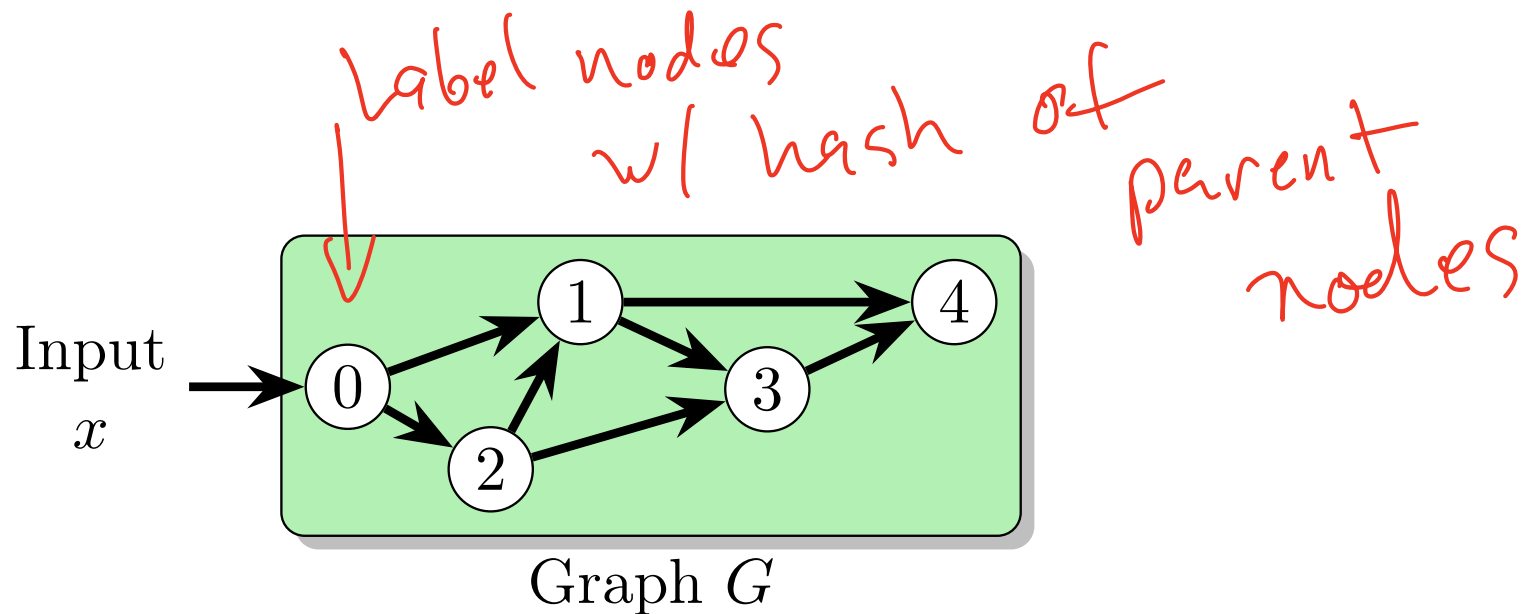
- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.



Graph G

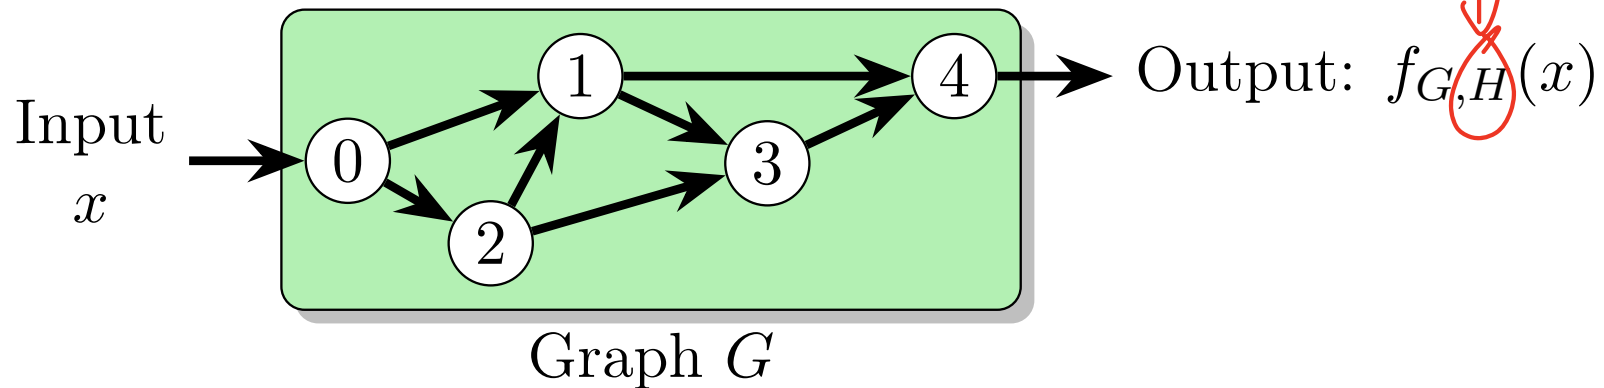
MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.



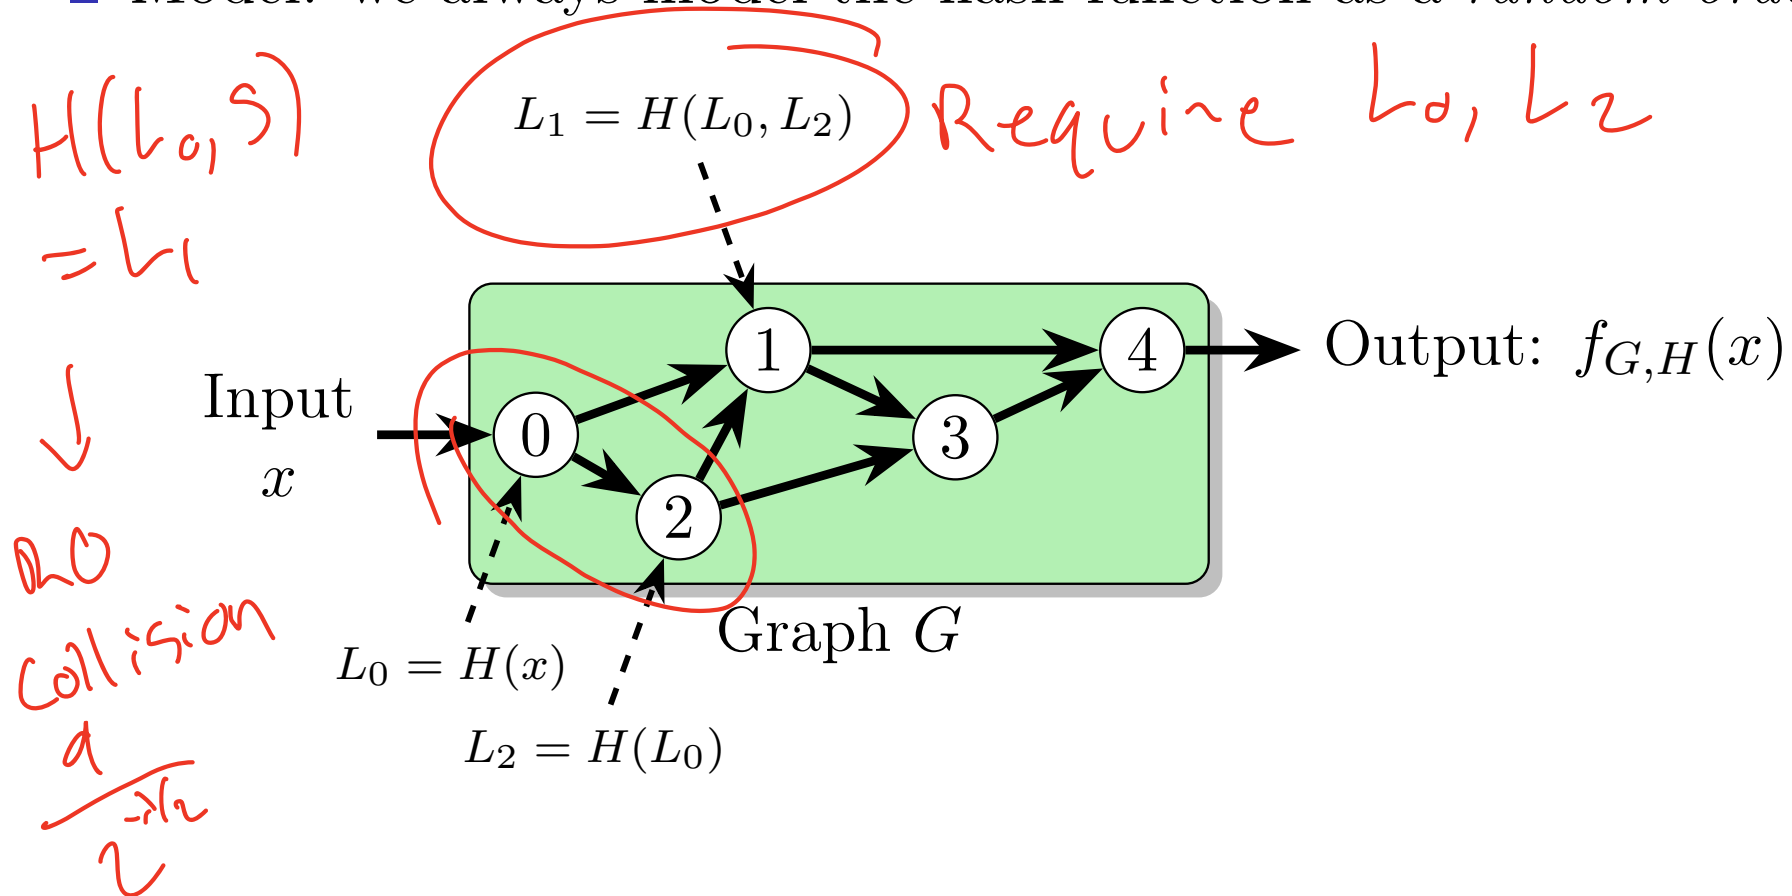
MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.




MHFs FROM DIRECTED GRAPHS

- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.

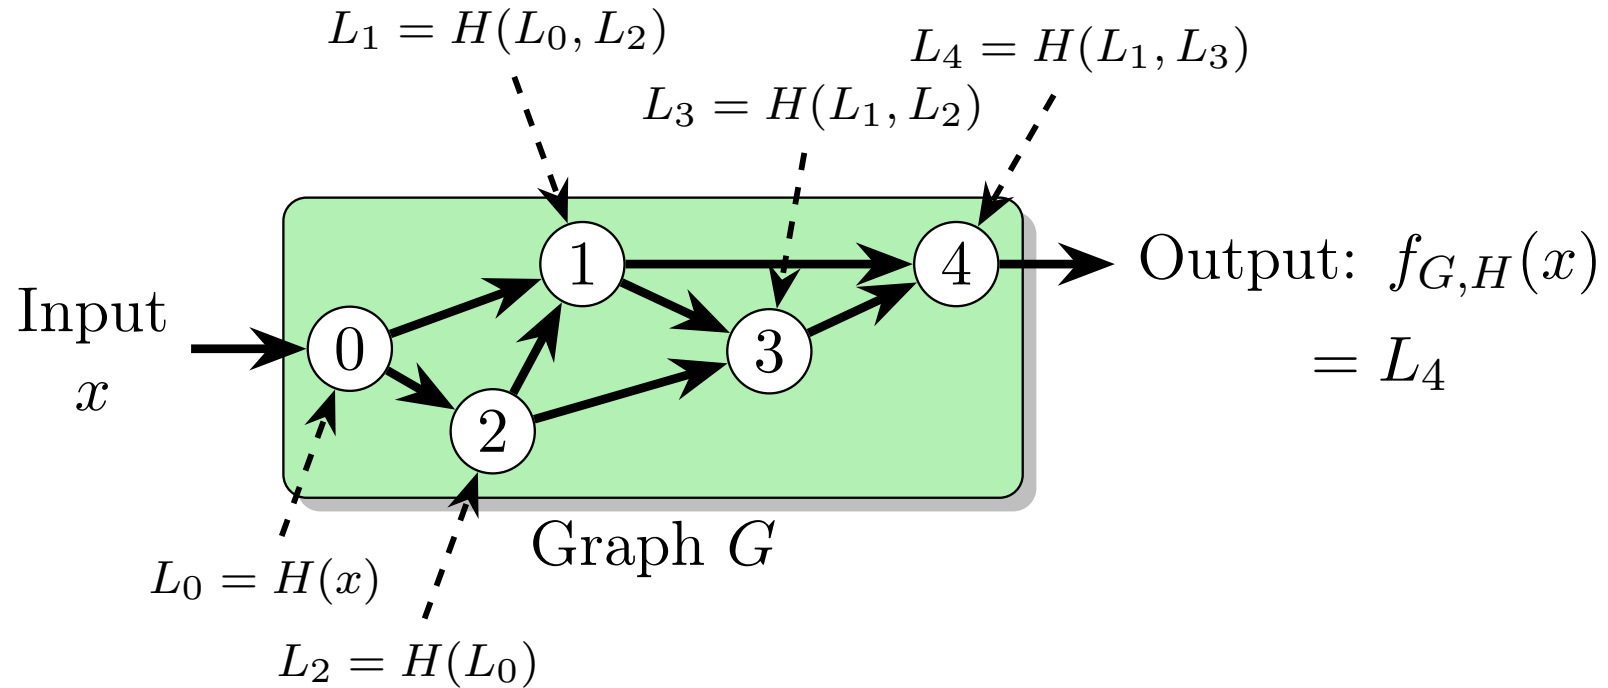


MHFs FROM DIRECTED GRAPHS


- Key idea of basically all MHFs: force hash to be computed as labels of a *directed acyclic graph*.
- Idea: use DAG to encode data dependencies to force large amounts of memory to be stored.
- Model: we always model the hash function as a *random oracle*.



MHFs FROM DIRECTED GRAPHS

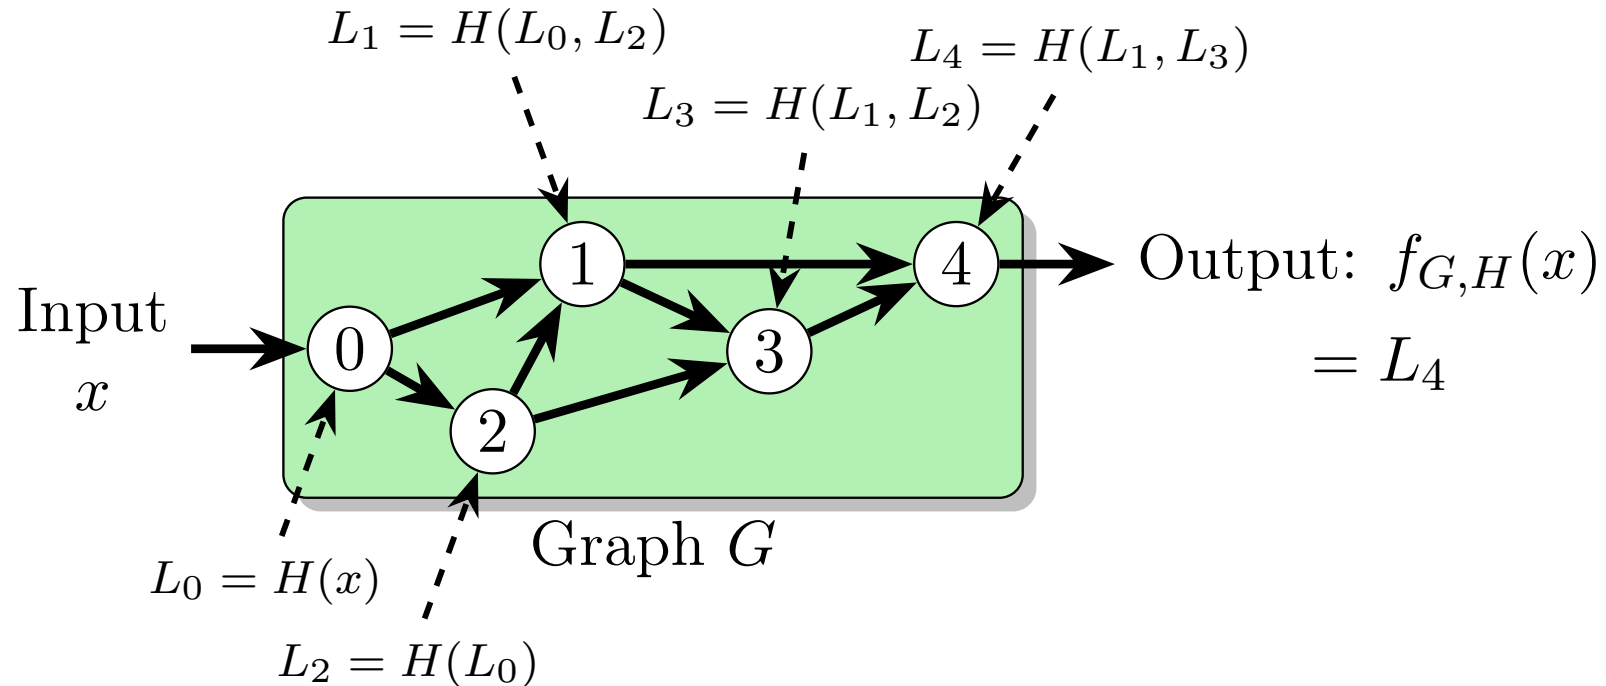


MHFs FROM DIRECTED GRAPHS




- $H: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ is a random oracle.

MHFs FROM DIRECTED GRAPHS




- $H: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ is a random oracle.
- DAG G encodes data dependencies.

MHFs FROM DIRECTED GRAPHS



- $H: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ is a random oracle.
- DAG G encodes data dependencies.
 - Maximum in-degree $\delta = O(1)$.

MHFs FROM DIRECTED GRAPHS



- $H: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ is a random oracle.
- DAG G encodes data dependencies.
 - Maximum in-degree $\delta = O(1)$.
 - $|V| = N = 2^n$ nodes/vertices (not in this example).

DMHFs vs. IMHFs

DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.

DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.

DMHFs vs. iMHFs


- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.

DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: `sCrypt`


DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt




DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt



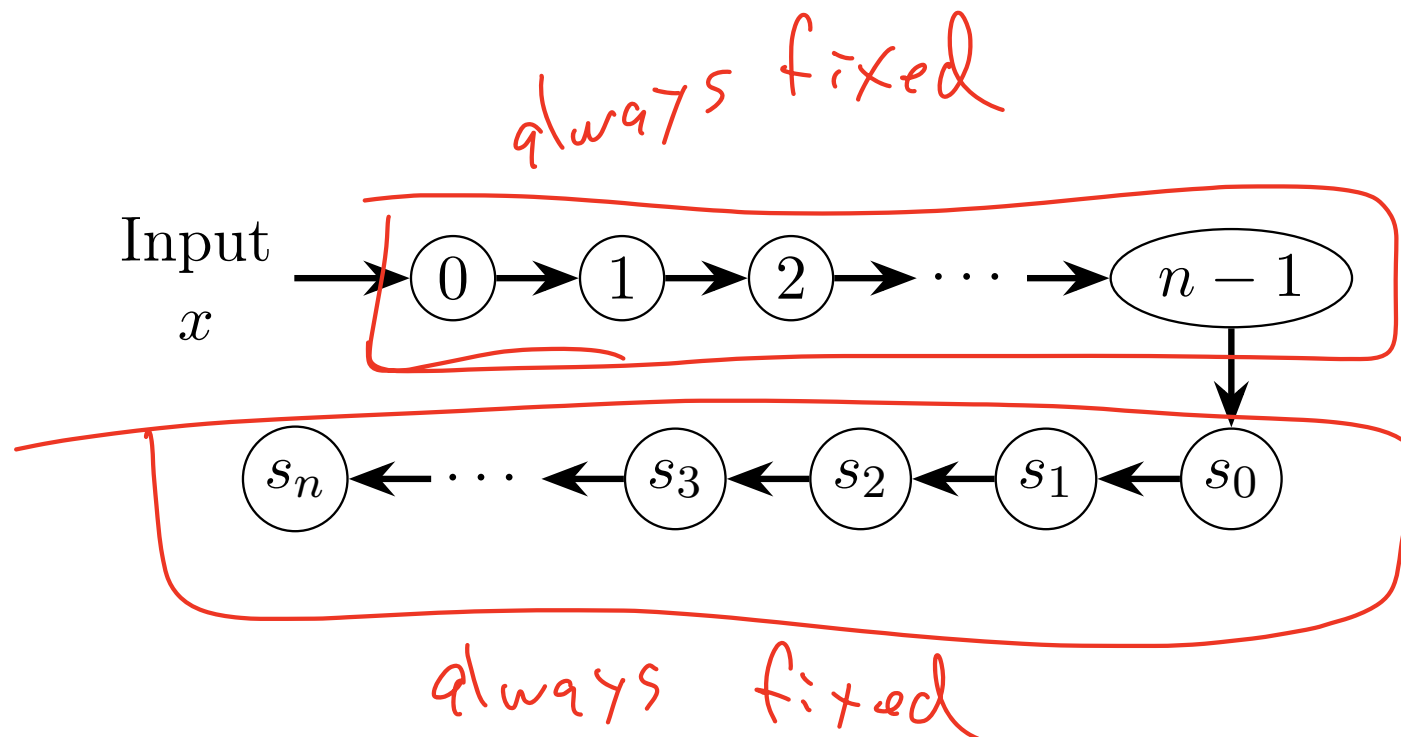
DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt




DMHFs vs. IMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt



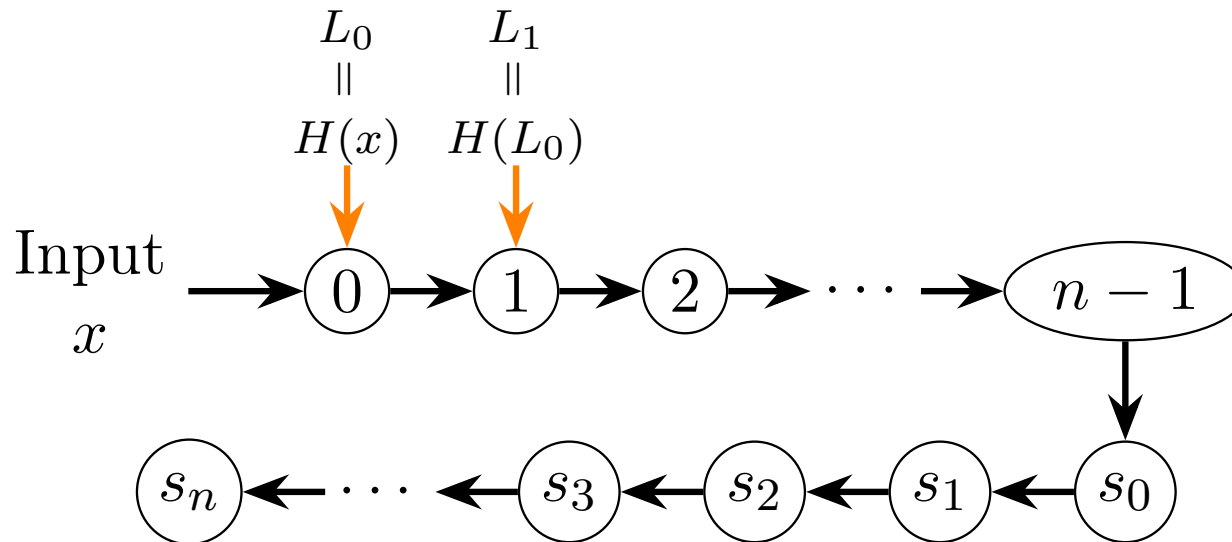
DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt



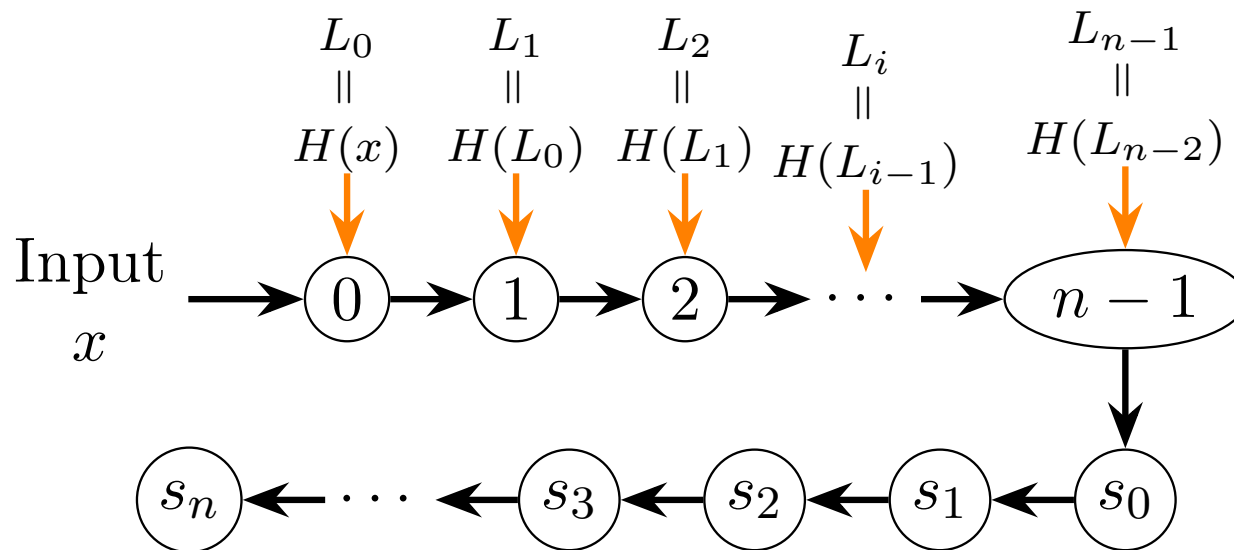
DMHFs vs. iMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt



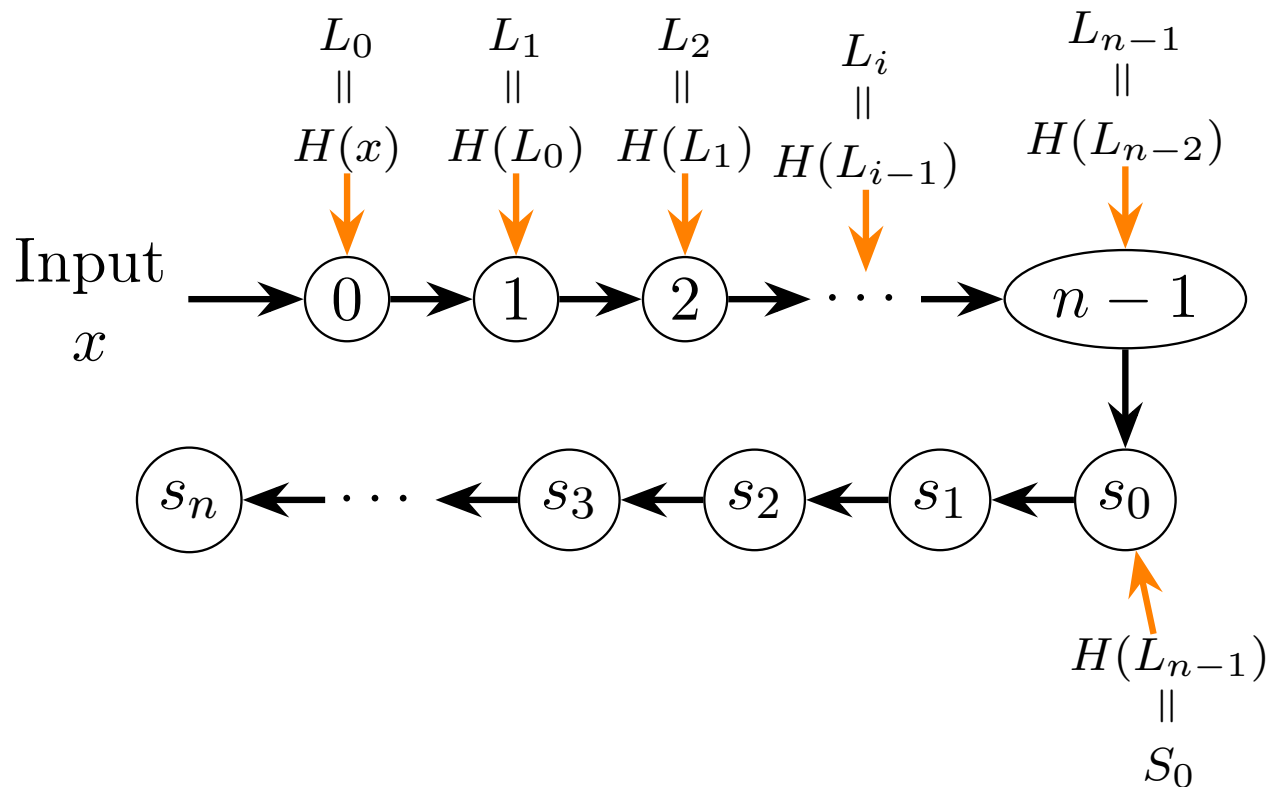
DMHFs vs. IMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dmHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt



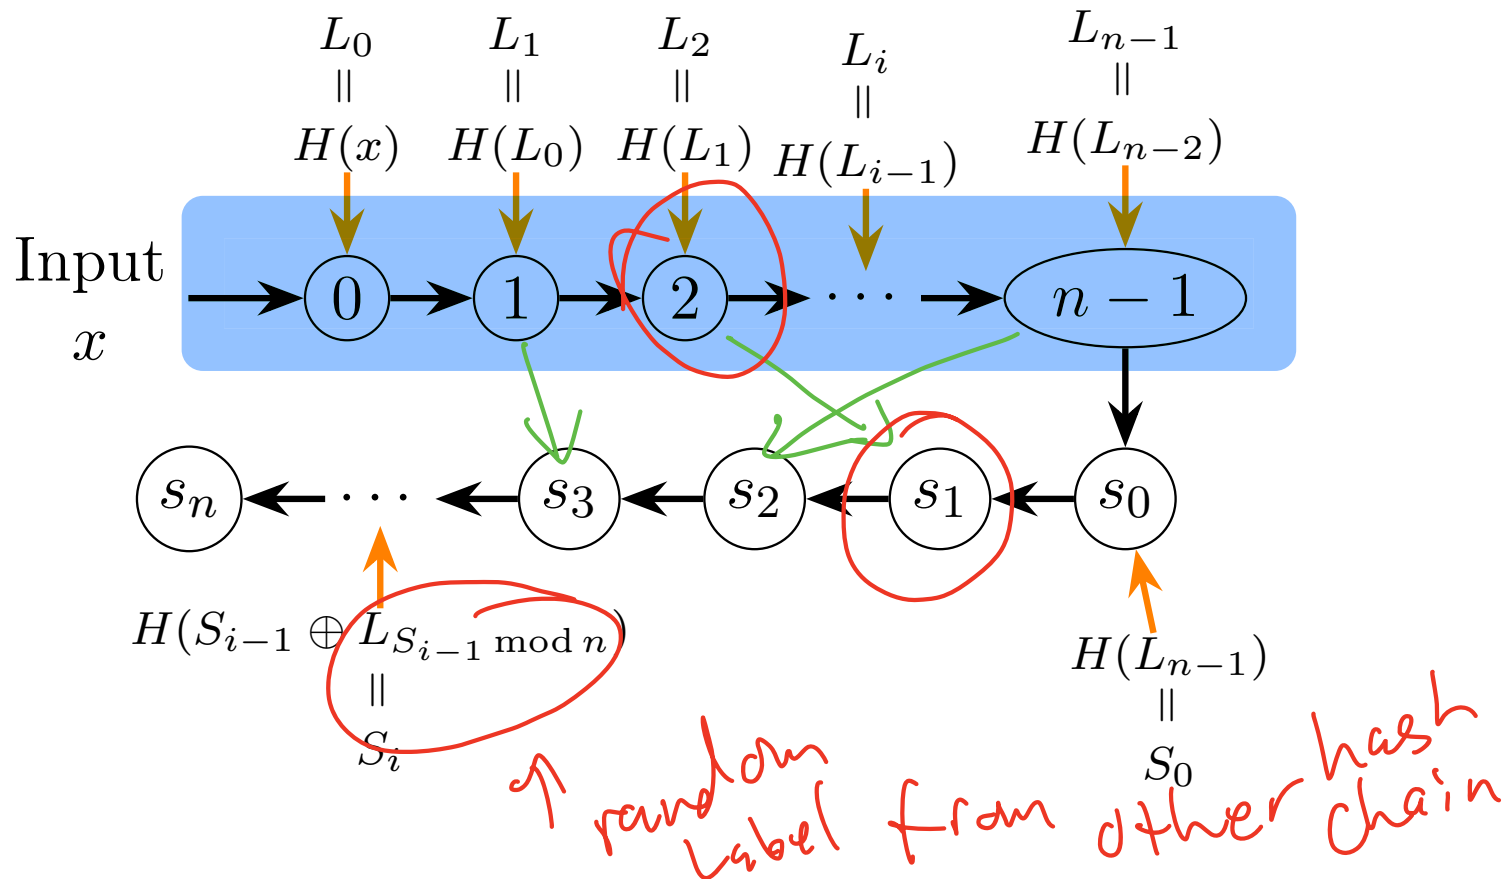
DMHFs vs. IMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dMHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt




DMHFs vs. IMHFs

- For an iMHF, the graph G is randomly sampled *independently* of any input.
 - I.e., randomly sample a graph before defining the function.
- For dmHFs, the graph G is computed *on-the-fly* using some function of the hash function/RO and the input.
 - Key example: sCrypt




EVALUATING A MHF

- Idea: *pebble the graph*.



EVALUATING A MHF

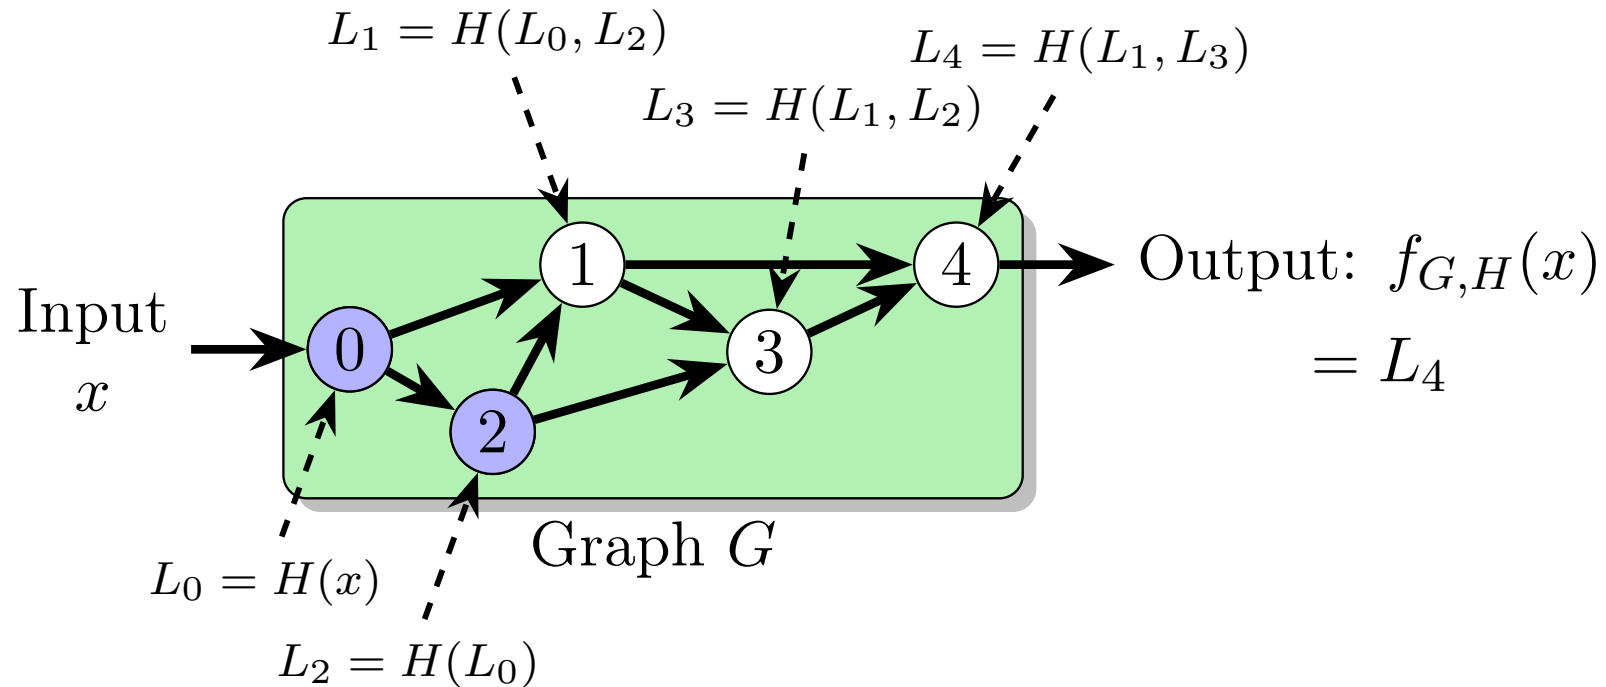
- Idea: *pebble the graph*.



- $P_0 = \{L_0\}$.

EVALUATING A MHF


- Idea: *pebble the graph*.



- $P_0 = \{L_0\}$.
- $P_1 = \{L_0, L_2\}$.

EVALUATING A MHF


- Idea: *pebble the graph*.



- $P_0 = \{L_0\}$.
- $P_1 = \{L_0, L_2\}$.
- $P_2 = \{L_2, L_1\}$.

EVALUATING A MHF

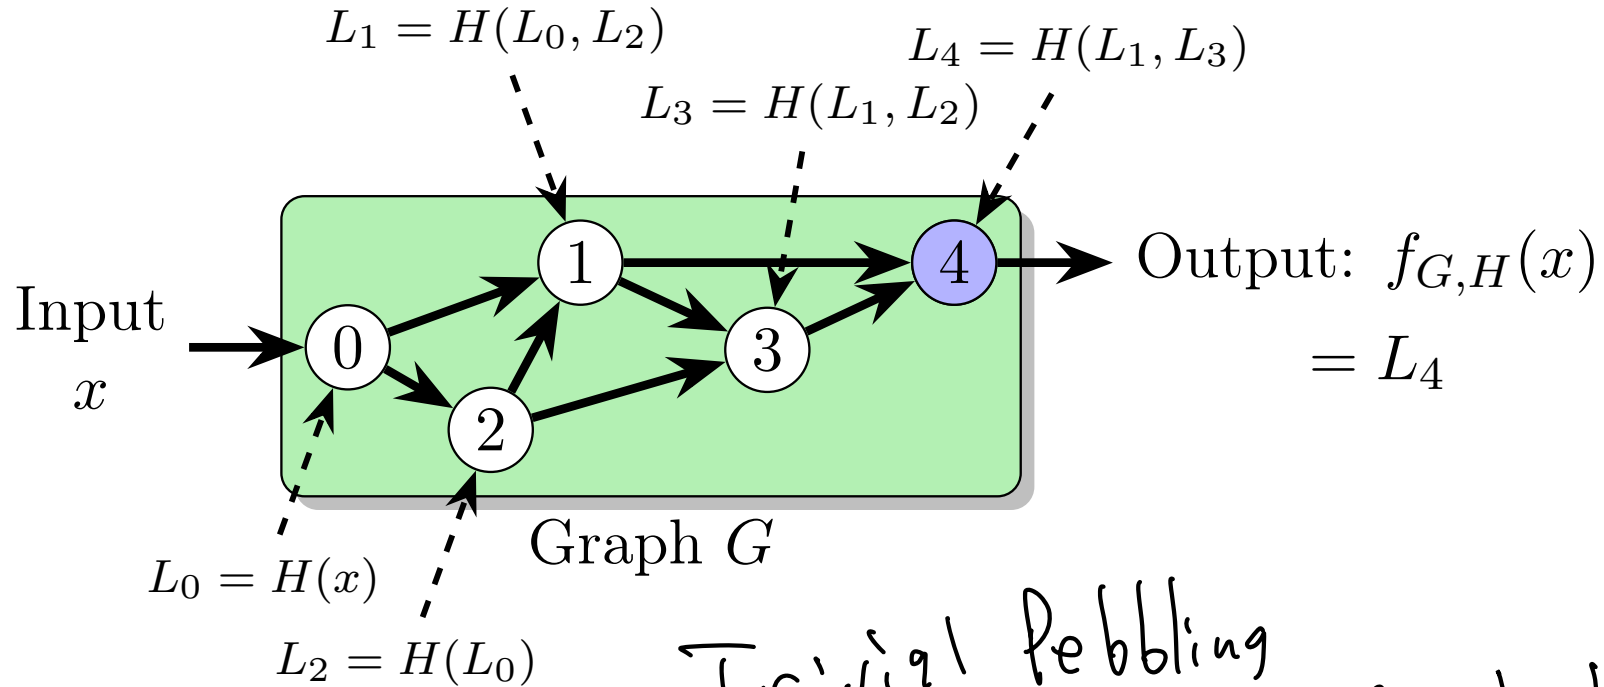
- Idea: *pebble the graph*.



- $P_0 = \{L_0\}$.
- $P_1 = \{L_0, L_2\}$.
- $P_2 = \{L_2, L_1\}$.
- $P_3 = \{L_1, L_3\}$.

EVALUATING A MHF

- Idea: *pebble the graph*.



Trivial Pebbling

$P_0 = \{L_0\}$ $P_3 = \{L_0, L_1, L_2, L_3\}$

$P_1 = \{L_1, L_2\}$ $P_4 = \{L_0, \dots, L_4\}$

$P_2 = \{L_0, L_2, L_1\}$

- P
- $P_0 = \{L_0\}$.
 - $P_1 = \{L_0, L_2\}$.
 - $P_2 = \{L_2, L_1\}$.
 - $P_3 = \{L_1, L_3\}$.
 - $P_4 = \{L_4\}$.

HOW TO MEASURE MEMORY-HARDNESS?

HOW TO MEASURE MEMORY-HARDNESS?

- Natural question: what does it actually mean to be *memory-hard*?

HOW TO MEASURE MEMORY-HARDNESS?

- Natural question: what does it actually mean to be *memory-hard*?
- There are several possible definitions of “memory-hardness.”

HOW TO MEASURE MEMORY-HARDNESS?

- Natural question: what does it actually mean to be *memory-hard*?
- There are several possible definitions of “memory-hardness.”
- Intuitively, we want to capture the idea that an adversary must lock up large amounts of memory for the *duration* of a computation.

HOW TO MEASURE MEMORY-HARDNESS?

- Natural question: what does it actually mean to be *memory-hard*?
- There are several possible definitions of “memory-hardness.”
- Intuitively, we want to capture the idea that an adversary must lock up large amounts of memory for the *duration* of a computation.
- Idea: translate memory costs to the cost of *pebbling* the underlying DAG of a MHF.

ST COMPLEXITY: ATTEMPT 1

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.

- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(\overset{\substack{\downarrow \text{\# of pebbling states}}}{T_{\mathbf{P}}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(T_{\mathbf{P}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

- \mathbf{P} is a pebbling strategy.

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(T_{\mathbf{P}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

- \mathbf{P} is a pebbling strategy.
- $T_{\mathbf{P}}$ is the time to compute \mathbf{P} .

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(T_{\mathbf{P}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

- \mathbf{P} is a pebbling strategy. # of states
- $T_{\mathbf{P}}$ is the time to compute \mathbf{P} .
- $|\mathbf{P}_i|$ is the size of the i^{th} pebbling state.

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(T_{\mathbf{P}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

- \mathbf{P} is a pebbling strategy.
 - $T_{\mathbf{P}}$ is the time to compute \mathbf{P} .
 - $|\mathbf{P}_i|$ is the size of the i^{th} pebbling state.
- ST-Complexity has a rich theory and is studied in many works.

ST COMPLEXITY: ATTEMPT 1

- First natural attempt: *space-time complexity*.
- For a DAG G , we define the *ST-Complexity* of G as

$$\text{ST}(G) = \min_{\mathbf{P}} \left(T_{\mathbf{P}} \times \max_{i \leq T_{\mathbf{P}}} (|\mathbf{P}_i|) \right)$$

- \mathbf{P} is a pebbling strategy.
 - $T_{\mathbf{P}}$ is the time to compute \mathbf{P} .
 - $|\mathbf{P}_i|$ is the size of the i^{th} pebbling state.
- ST-Complexity has a rich theory and is studied in many works.

ST-Complexity is not appropriate for
password hashing!

AMORTIZATION AND PARALLELISM

AMORTIZATION AND PARALLELISM


- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation.*

AMORTIZATION AND PARALLELISM

- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation*.
 - Password hashing is a very parallel computation!


AMORTIZATION AND PARALLELISM

- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation*.
 - Password hashing is a very parallel computation!



AMORTIZATION AND PARALLELISM

- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation*.
 - Password hashing is a very parallel computation!




$$ST_1 = S_1 \times T_1 \approx S_3 \times T_3 = ST_3$$

↑
cost of computing
 f once

↑
cost of computing
 f three times

AMORTIZATION AND PARALLELISM

- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation*.
 - Password hashing is a very parallel computation!



$$ST_1 = S_1 \times T_1 \approx S_3 \times T_3 = ST_3$$

↑
cost of computing f once


↑
cost of computing f three times

Theorem 1

There exists a function f_n (making n calls to a random oracle) such that the ST-Complexity of \sqrt{n} parallel calls to f_n is $O(ST(f))$.

AMORTIZATION AND PARALLELISM

- **Issue:** ST-Complexity can scale poorly in the number of function evaluations *for a parallel computation*.
 - Password hashing is a very parallel computation!



$$ST_1 = S_1 \times T_1 \approx S_3 \times T_3 = ST_3$$

↑ cost of computing f once

↑ cost of computing f three times

Theorem 1

There exists a function f_n (making n calls to a random oracle) such that the ST-Complexity of \sqrt{n} parallel calls to f_n is $O(ST(f))$.

For measuring memory-hardness, we would like an *amortized* memory cost (amortized area-time complexity)

CUMULATIVE (MEMORY) COMPLEXITY

CUMULATIVE (MEMORY) COMPLEXITY

- The notion of *cumulative (memory) complexity* (CC) was introduced to approximate the amortized area-time complexity of pebbling a DAG G .

CUMULATIVE (MEMORY) COMPLEXITY

- The notion of *cumulative (memory) complexity* (CC) was introduced to approximate the amortized area-time complexity of pebbling a DAG G .

$$\text{CC}(G) = \min_{\mathbf{P}} \sum_{i=1}^{T_{\mathbf{P}}} |\mathbf{P}_i|$$

CUMULATIVE (MEMORY) COMPLEXITY

- The notion of *cumulative (memory) complexity* (CC) was introduced to approximate the amortized area-time complexity of pebbling a DAG G .

$$\text{CC}(G) = \min_{\mathbf{P}} \sum_{i=1}^{T_{\mathbf{P}}} |\mathbf{P}_i|$$

- $|\mathbf{P}_i|$ is the memory used at step i of the computation.

CUMULATIVE (MEMORY) COMPLEXITY

- The notion of *cumulative (memory) complexity* (CC) was introduced to approximate the amortized area-time complexity of pebbling a DAG G .

$$\text{CC}(G) = \min_{\mathbf{P}} \sum_{i=1}^{T_{\mathbf{P}}} |\mathbf{P}_i|$$

- $|\mathbf{P}_i|$ is the memory used at step i of the computation.
- Guessing two passwords now doubles an attacker's cost!

CUMULATIVE (MEMORY) COMPLEXITY

- The notion of *cumulative (memory) complexity* (CC) was introduced to approximate the amortized area-time complexity of pebbling a DAG G .

$$\text{CC}(G) = \min_{\mathbf{P}} \sum_{i=1}^{T_{\mathbf{P}}} |\mathbf{P}_i|$$

- $|\mathbf{P}_i|$ is the memory used at step i of the computation.
- Guessing two passwords now doubles an attacker's cost!

$$\text{CC}(G, G) = 2 \cdot \text{CC}(G)$$

**NEXT TIME: CC, THE PARALLEL ROM, AND
DEPTH-ROBUSTNESS**