

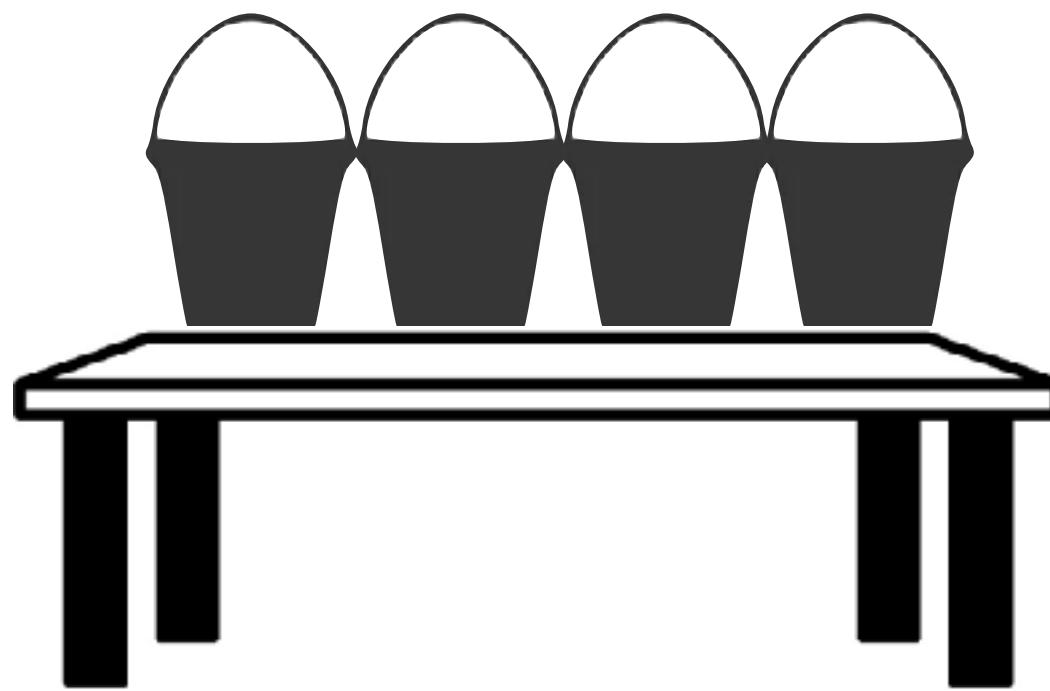
Practical Perfect Hashing

(Minimal & Dynamic)

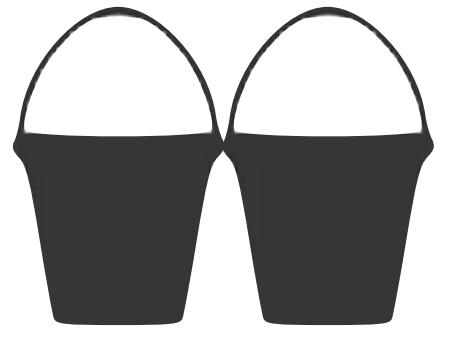
The background of the slide features a dark blue gradient. Overlaid on this are several glowing, translucent geometric shapes. These include several nested, perspective-viewed triangles that create a sense of depth. In the center, there is a hexagonal pattern composed of many thin, glowing lines that radiate outwards, resembling a complex network or a stylized sunburst. The overall effect is futuristic and mathematical.

Niv Dayan - CSC2525 Research Topics in Database Management

Hash Tables



Hash Tables



**Maps keys to
random buckets**



Resolves collisions



**Expected constant
time operations**

Many DB applications



Hash Join



In-memory index



Key-value Stores

Collision Resolution



Chaining



Open addressing
(Linear probing, etc)



Cuckoo hashing

Collision Resolution **relies on storing full keys**



Chaining



Open addressing
(Linear probing, etc)

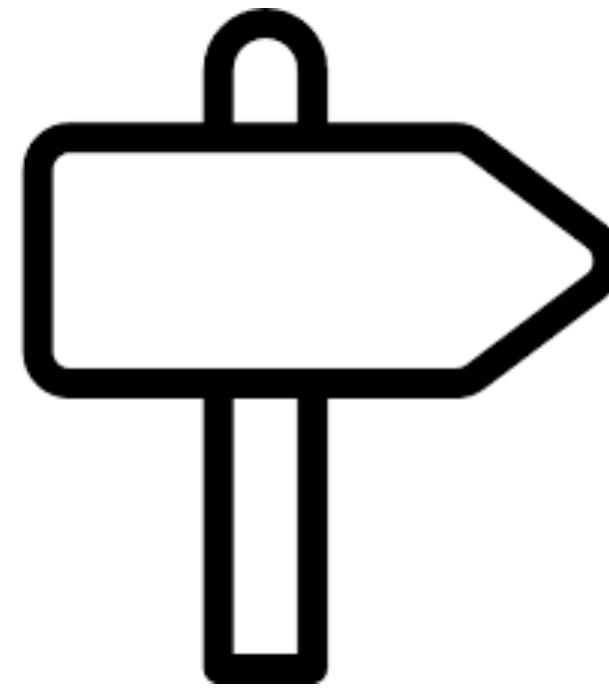


Cuckoo hashing

Problems storing full keys



space
Keys may be large

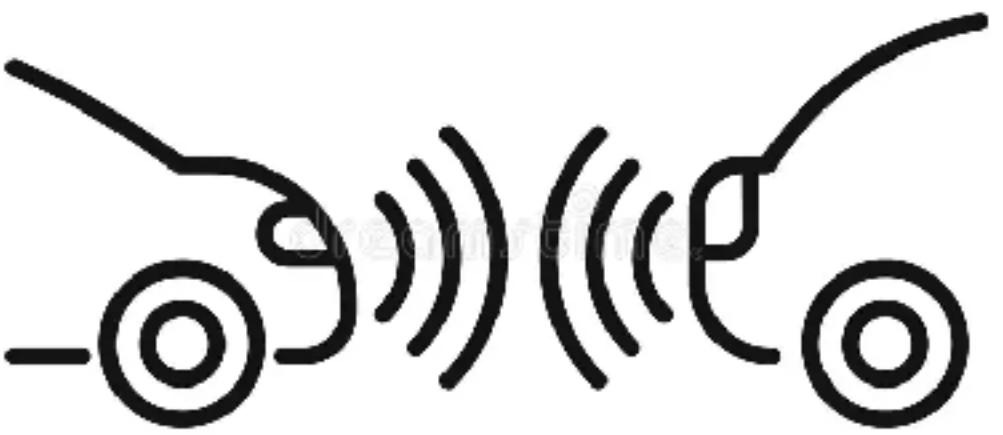


Requires indirection
if keys are var-length

Perfect Hashing



**Does not store
the keys**



**Collision-free
queries from key
to payload**

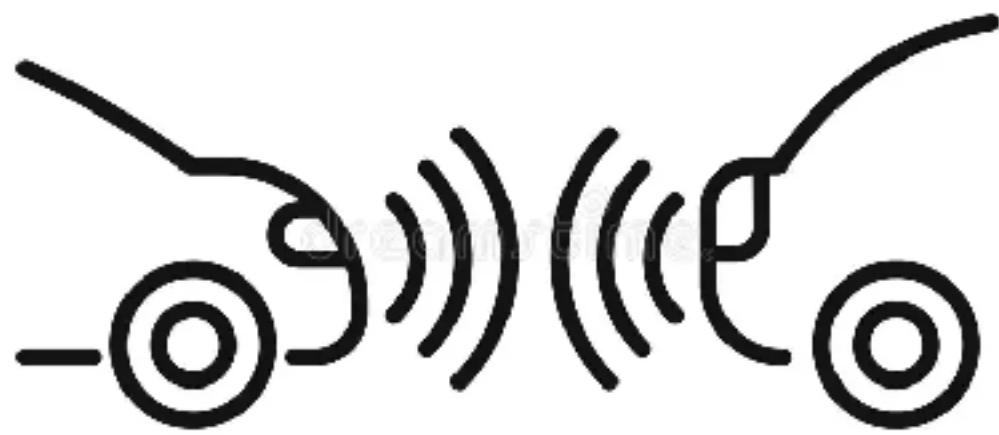


**Higher construction/
insertion overheads to
resolve collisions**

Perfect Hashing



Does not store
the keys



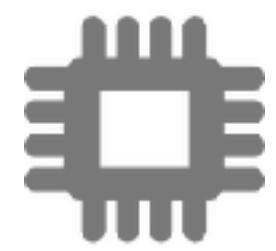
Collision-free
queries from key
to payload



Higher construction/
insertion overheads to
resolve collisions

Only effective when we query existing keys! Why?

Application Example: Key-value Stores

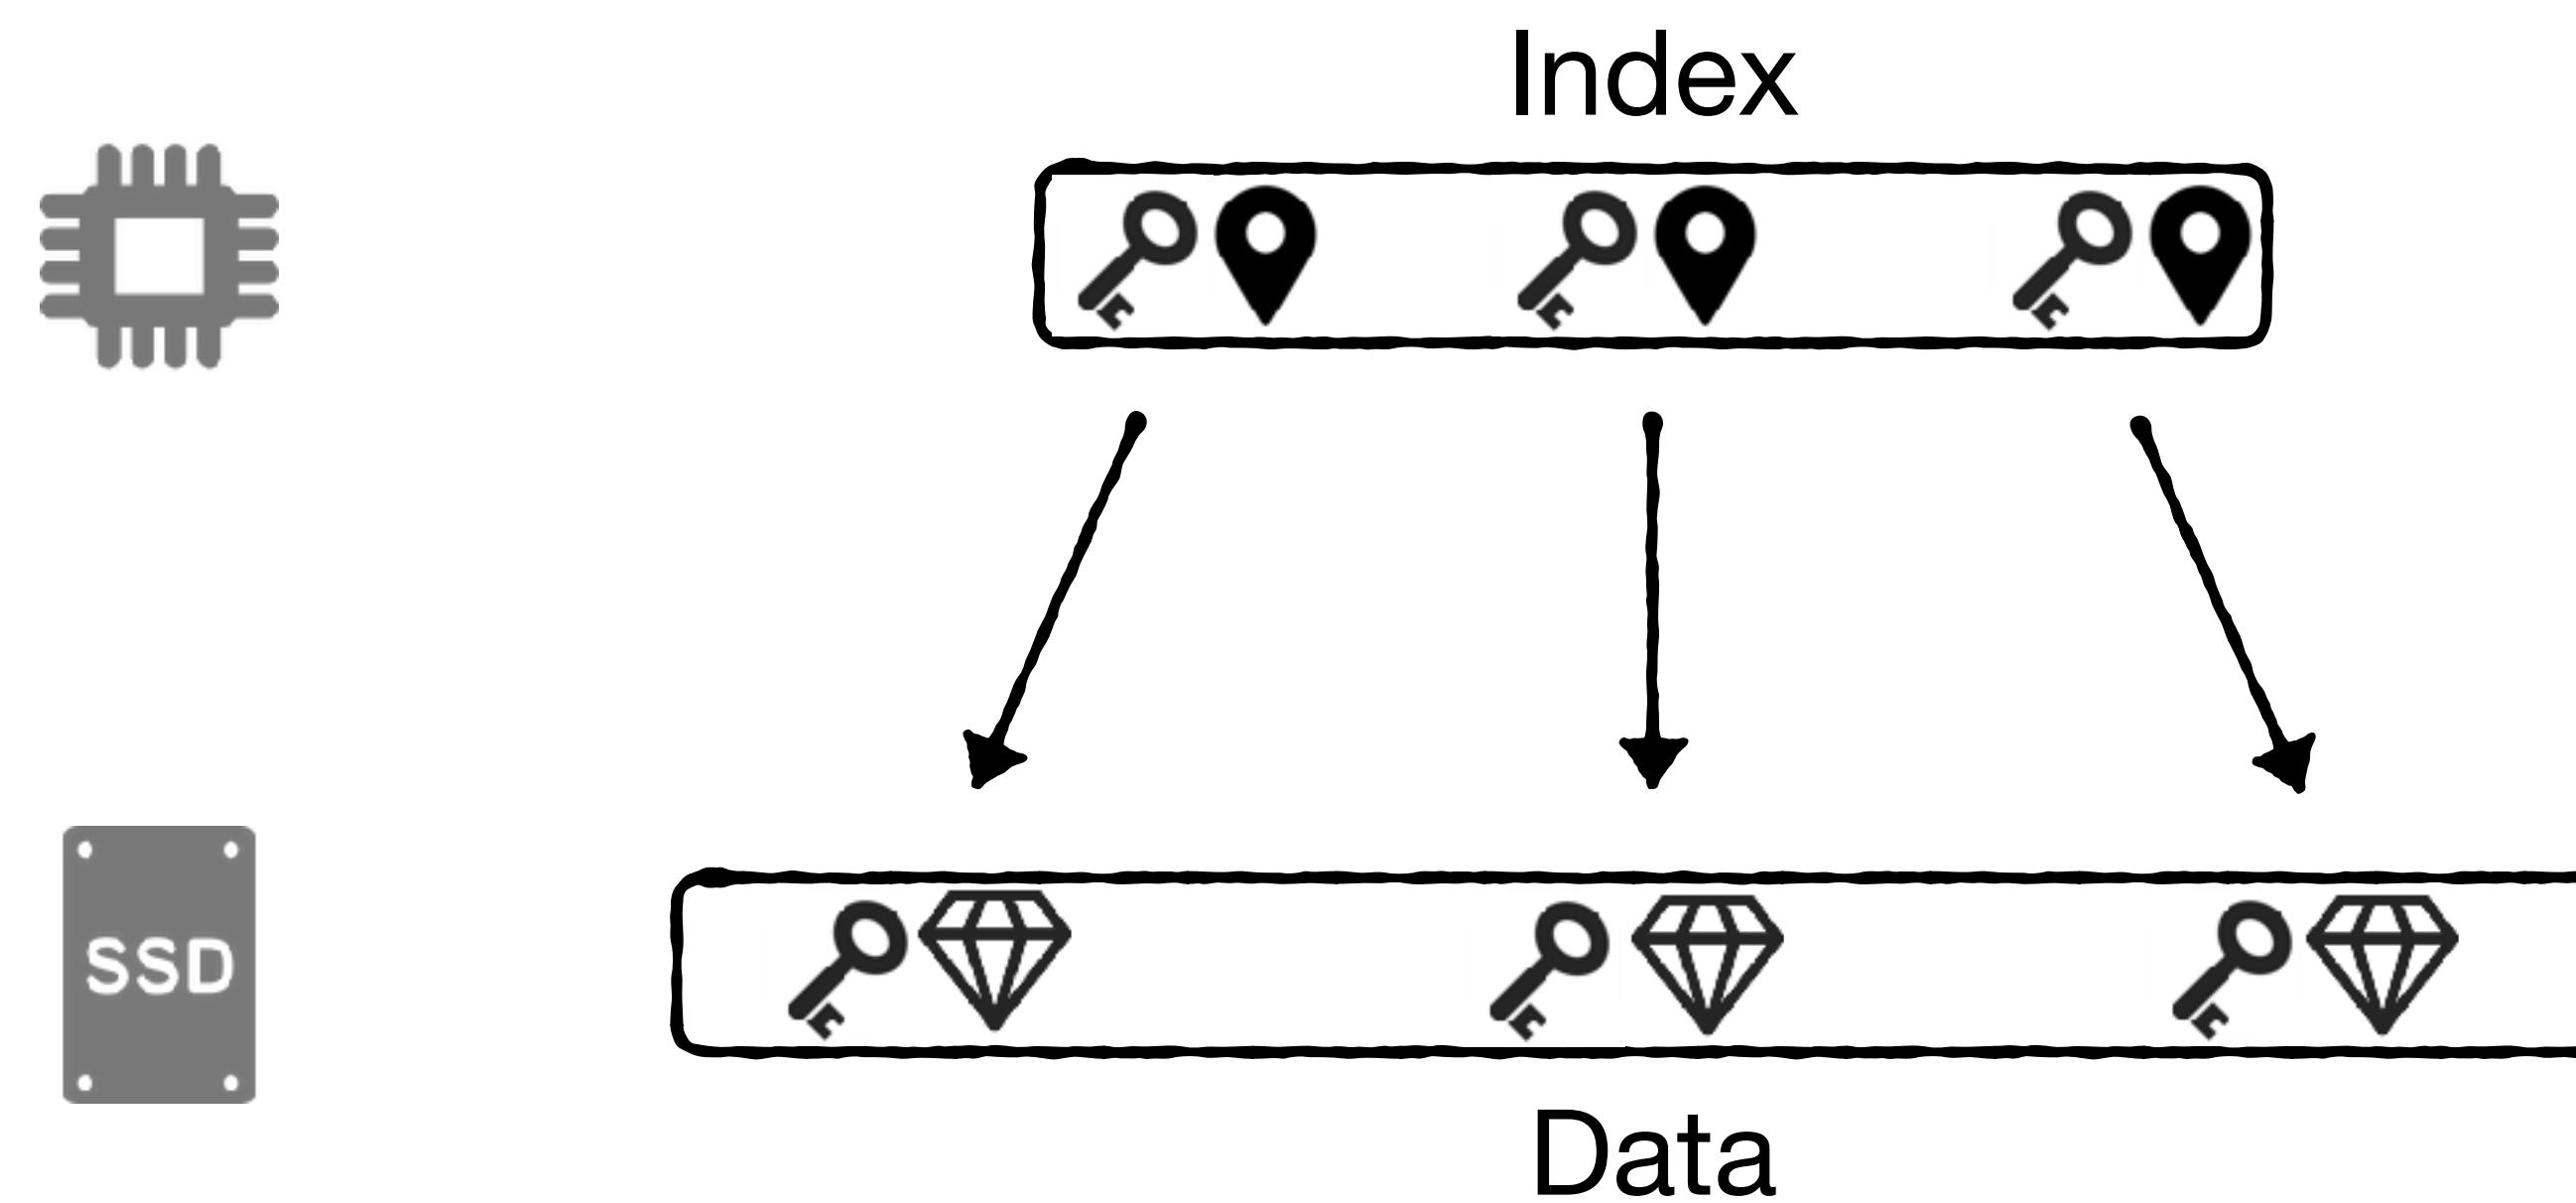


Index

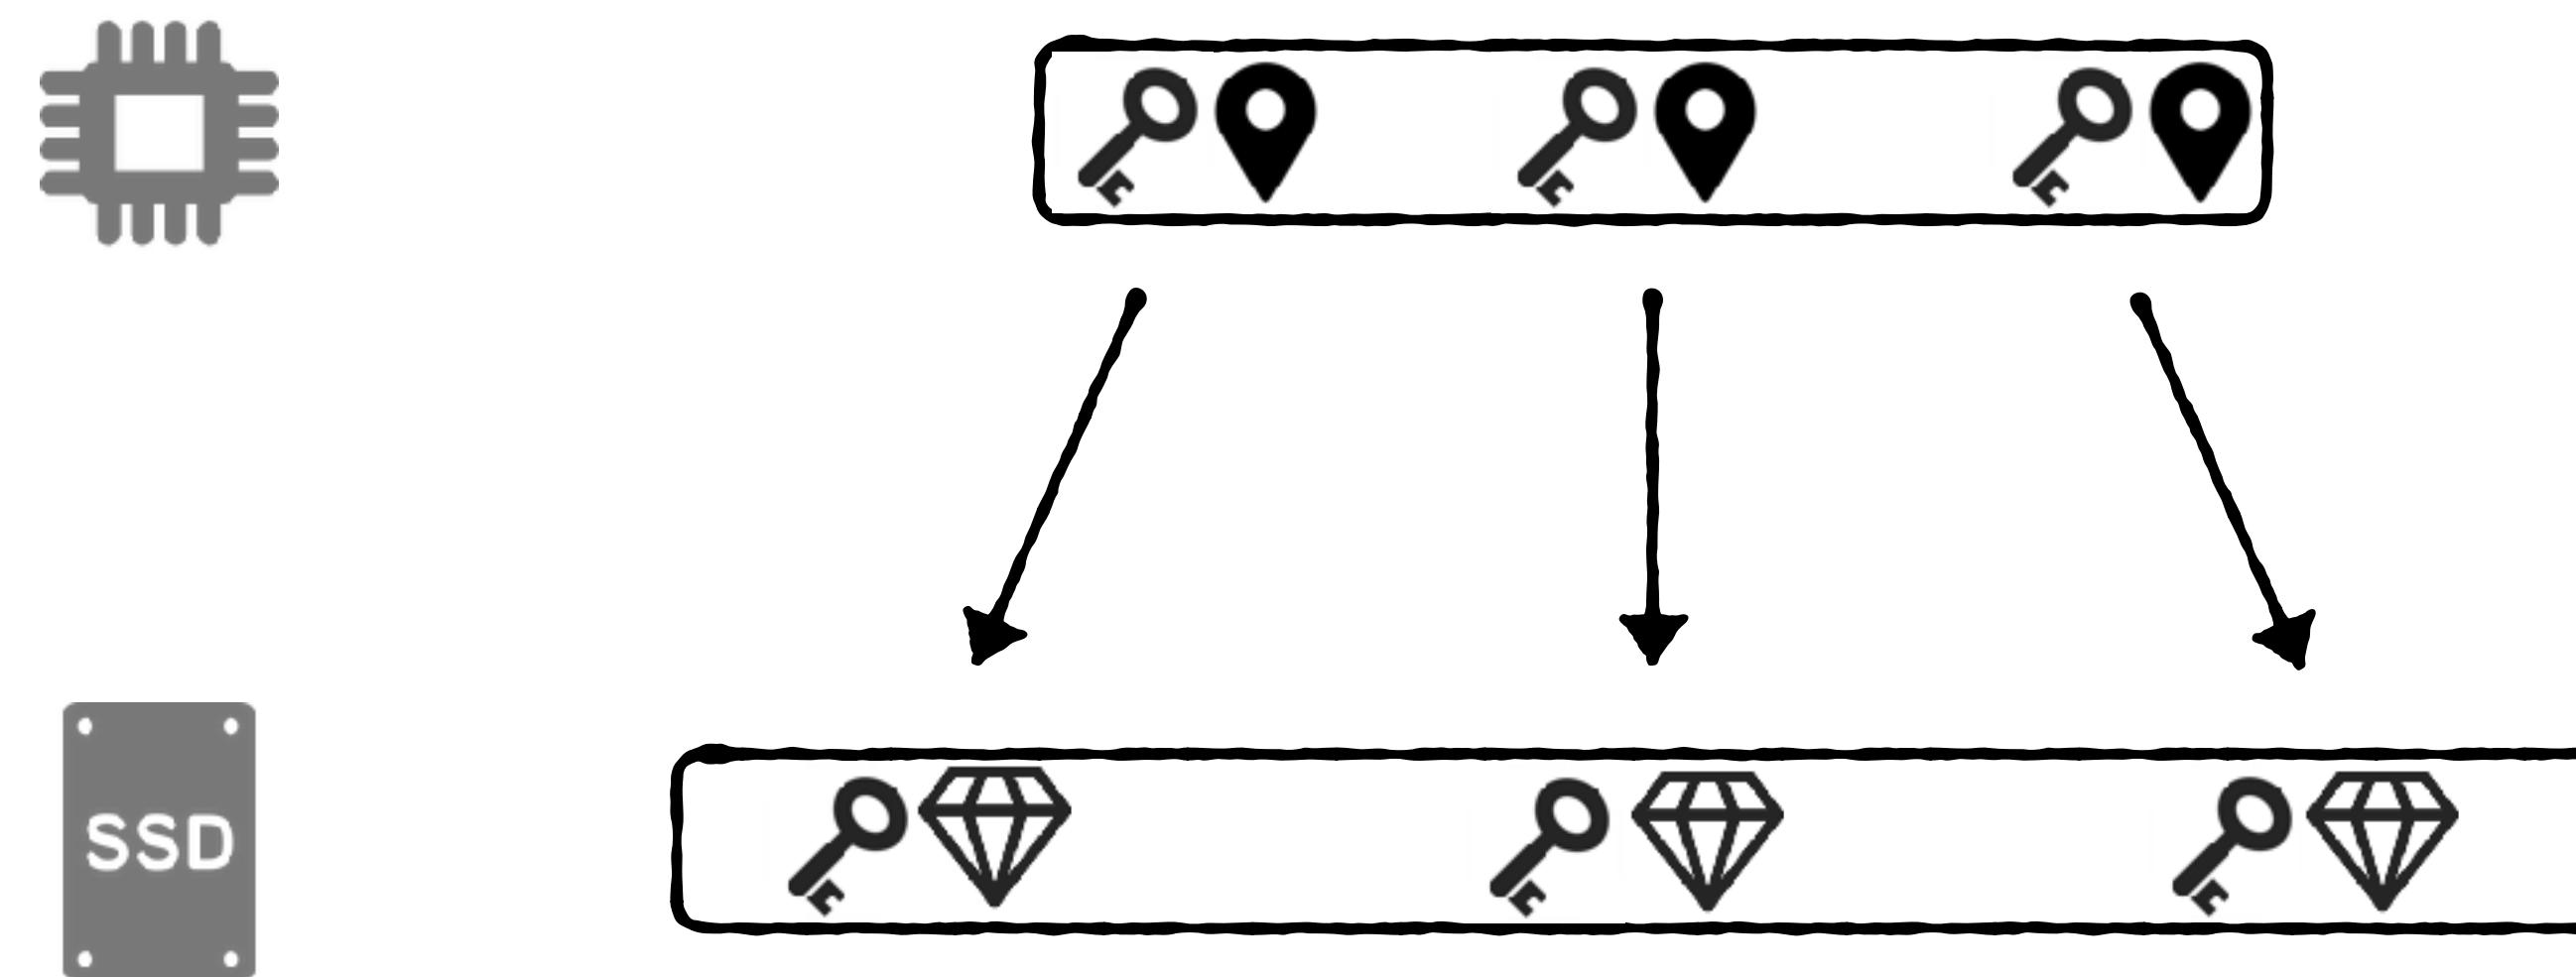


Data

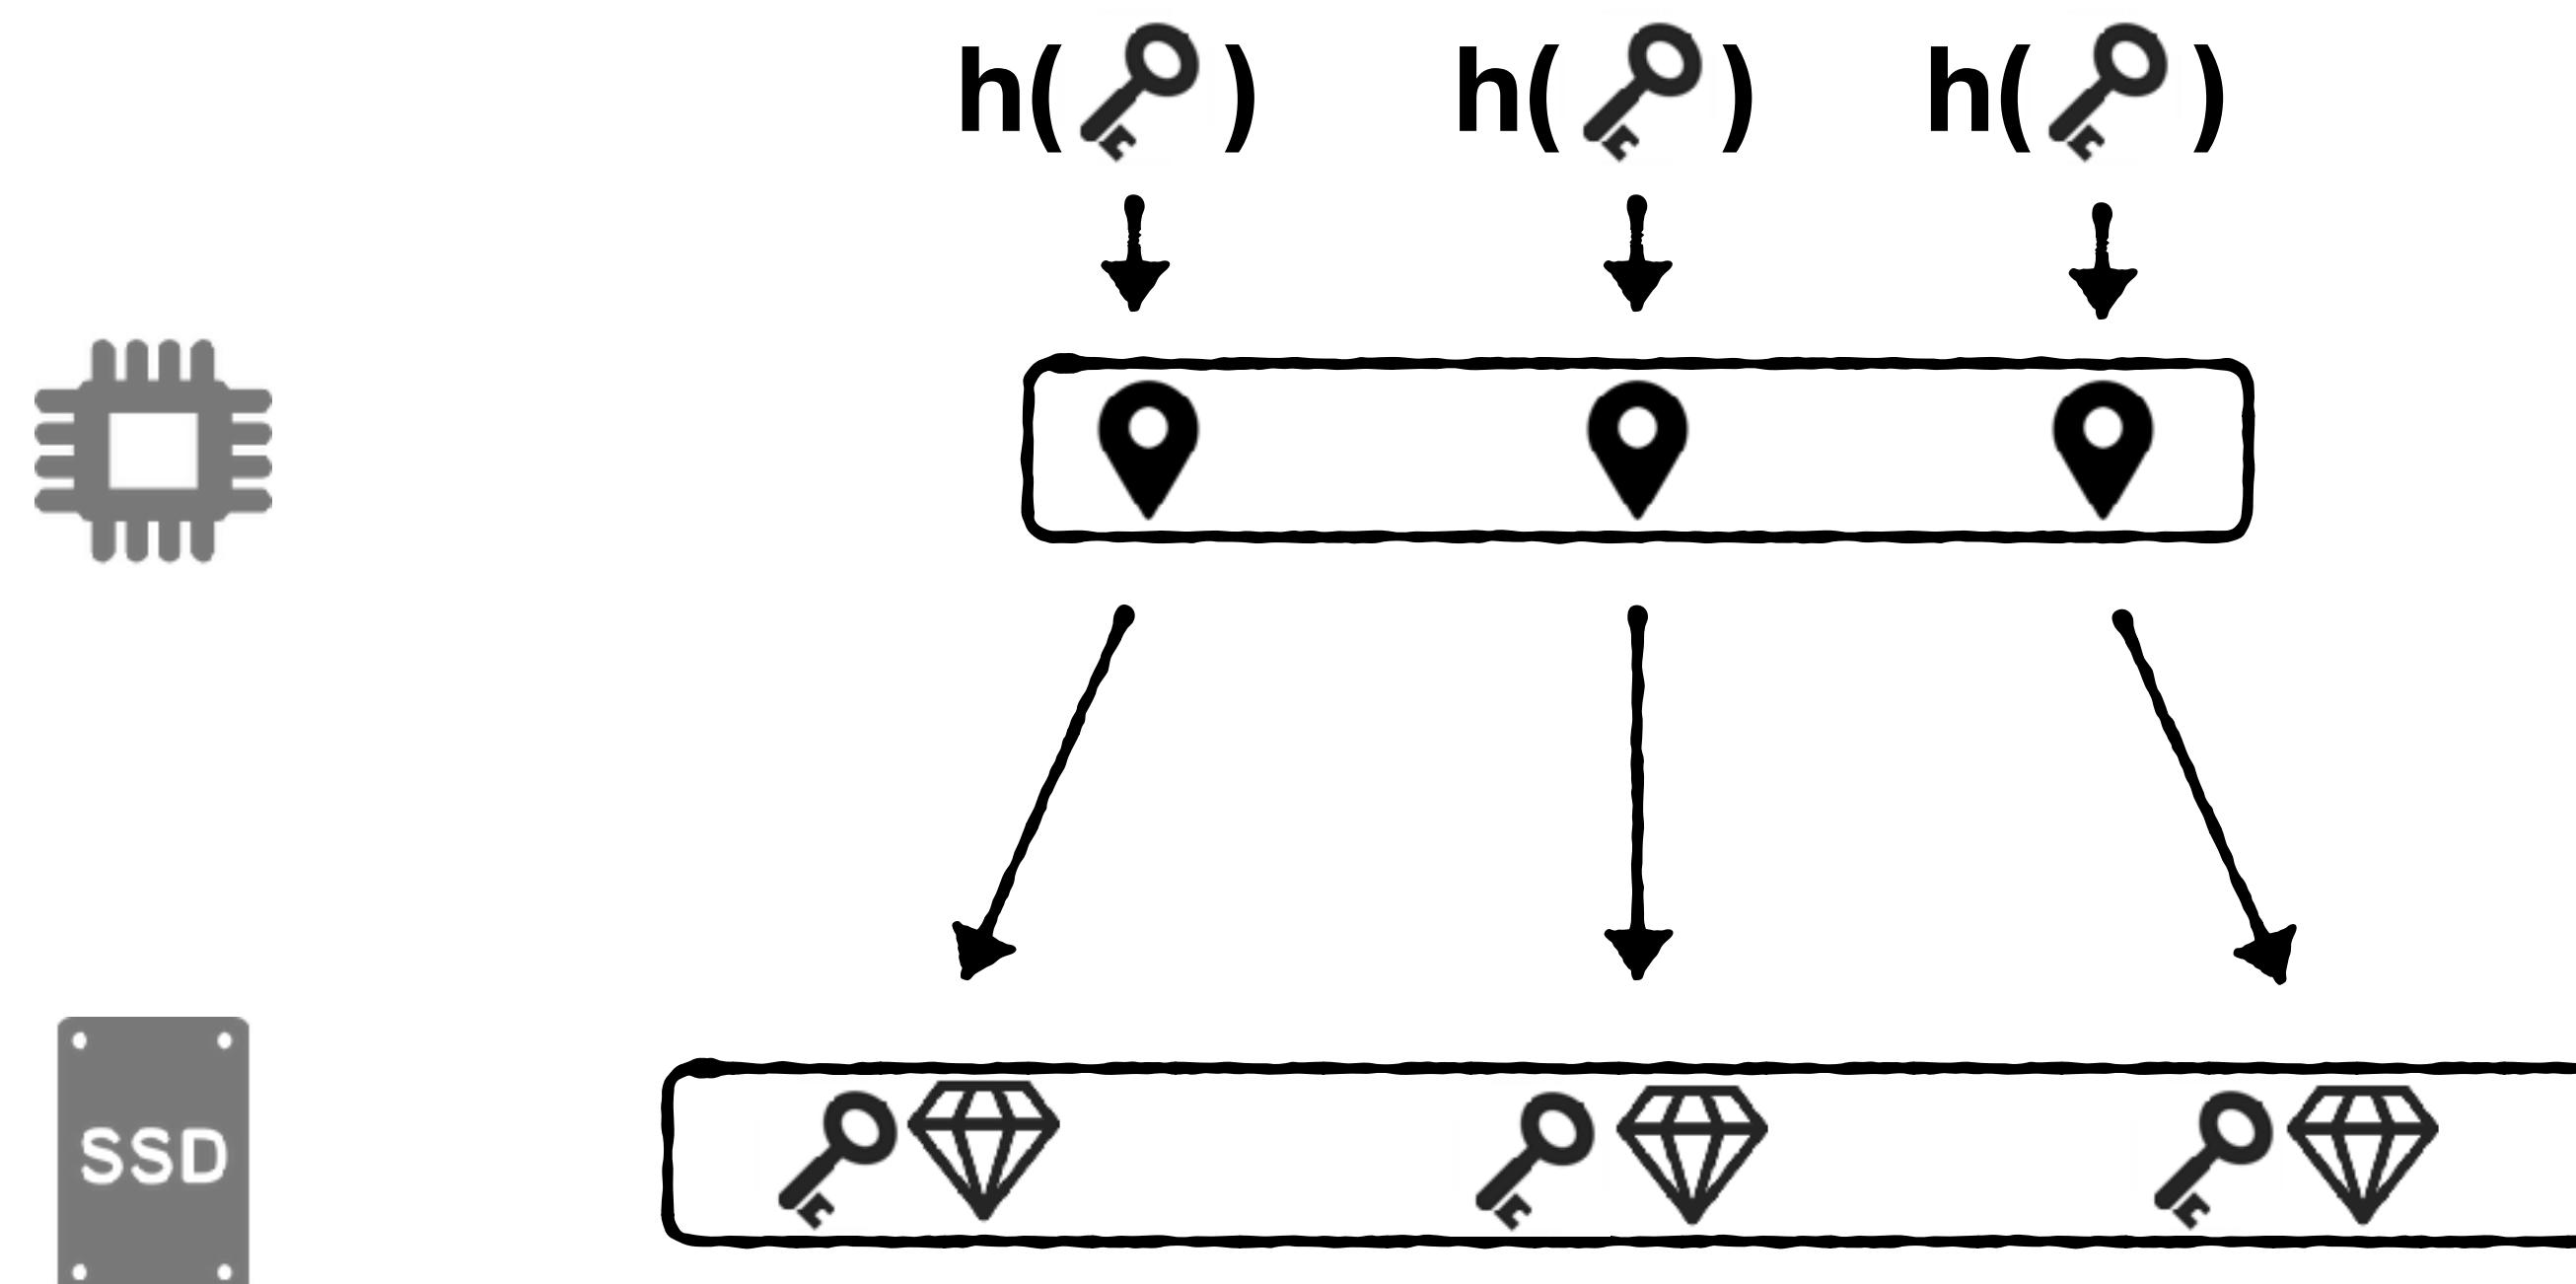
Key-value Stores



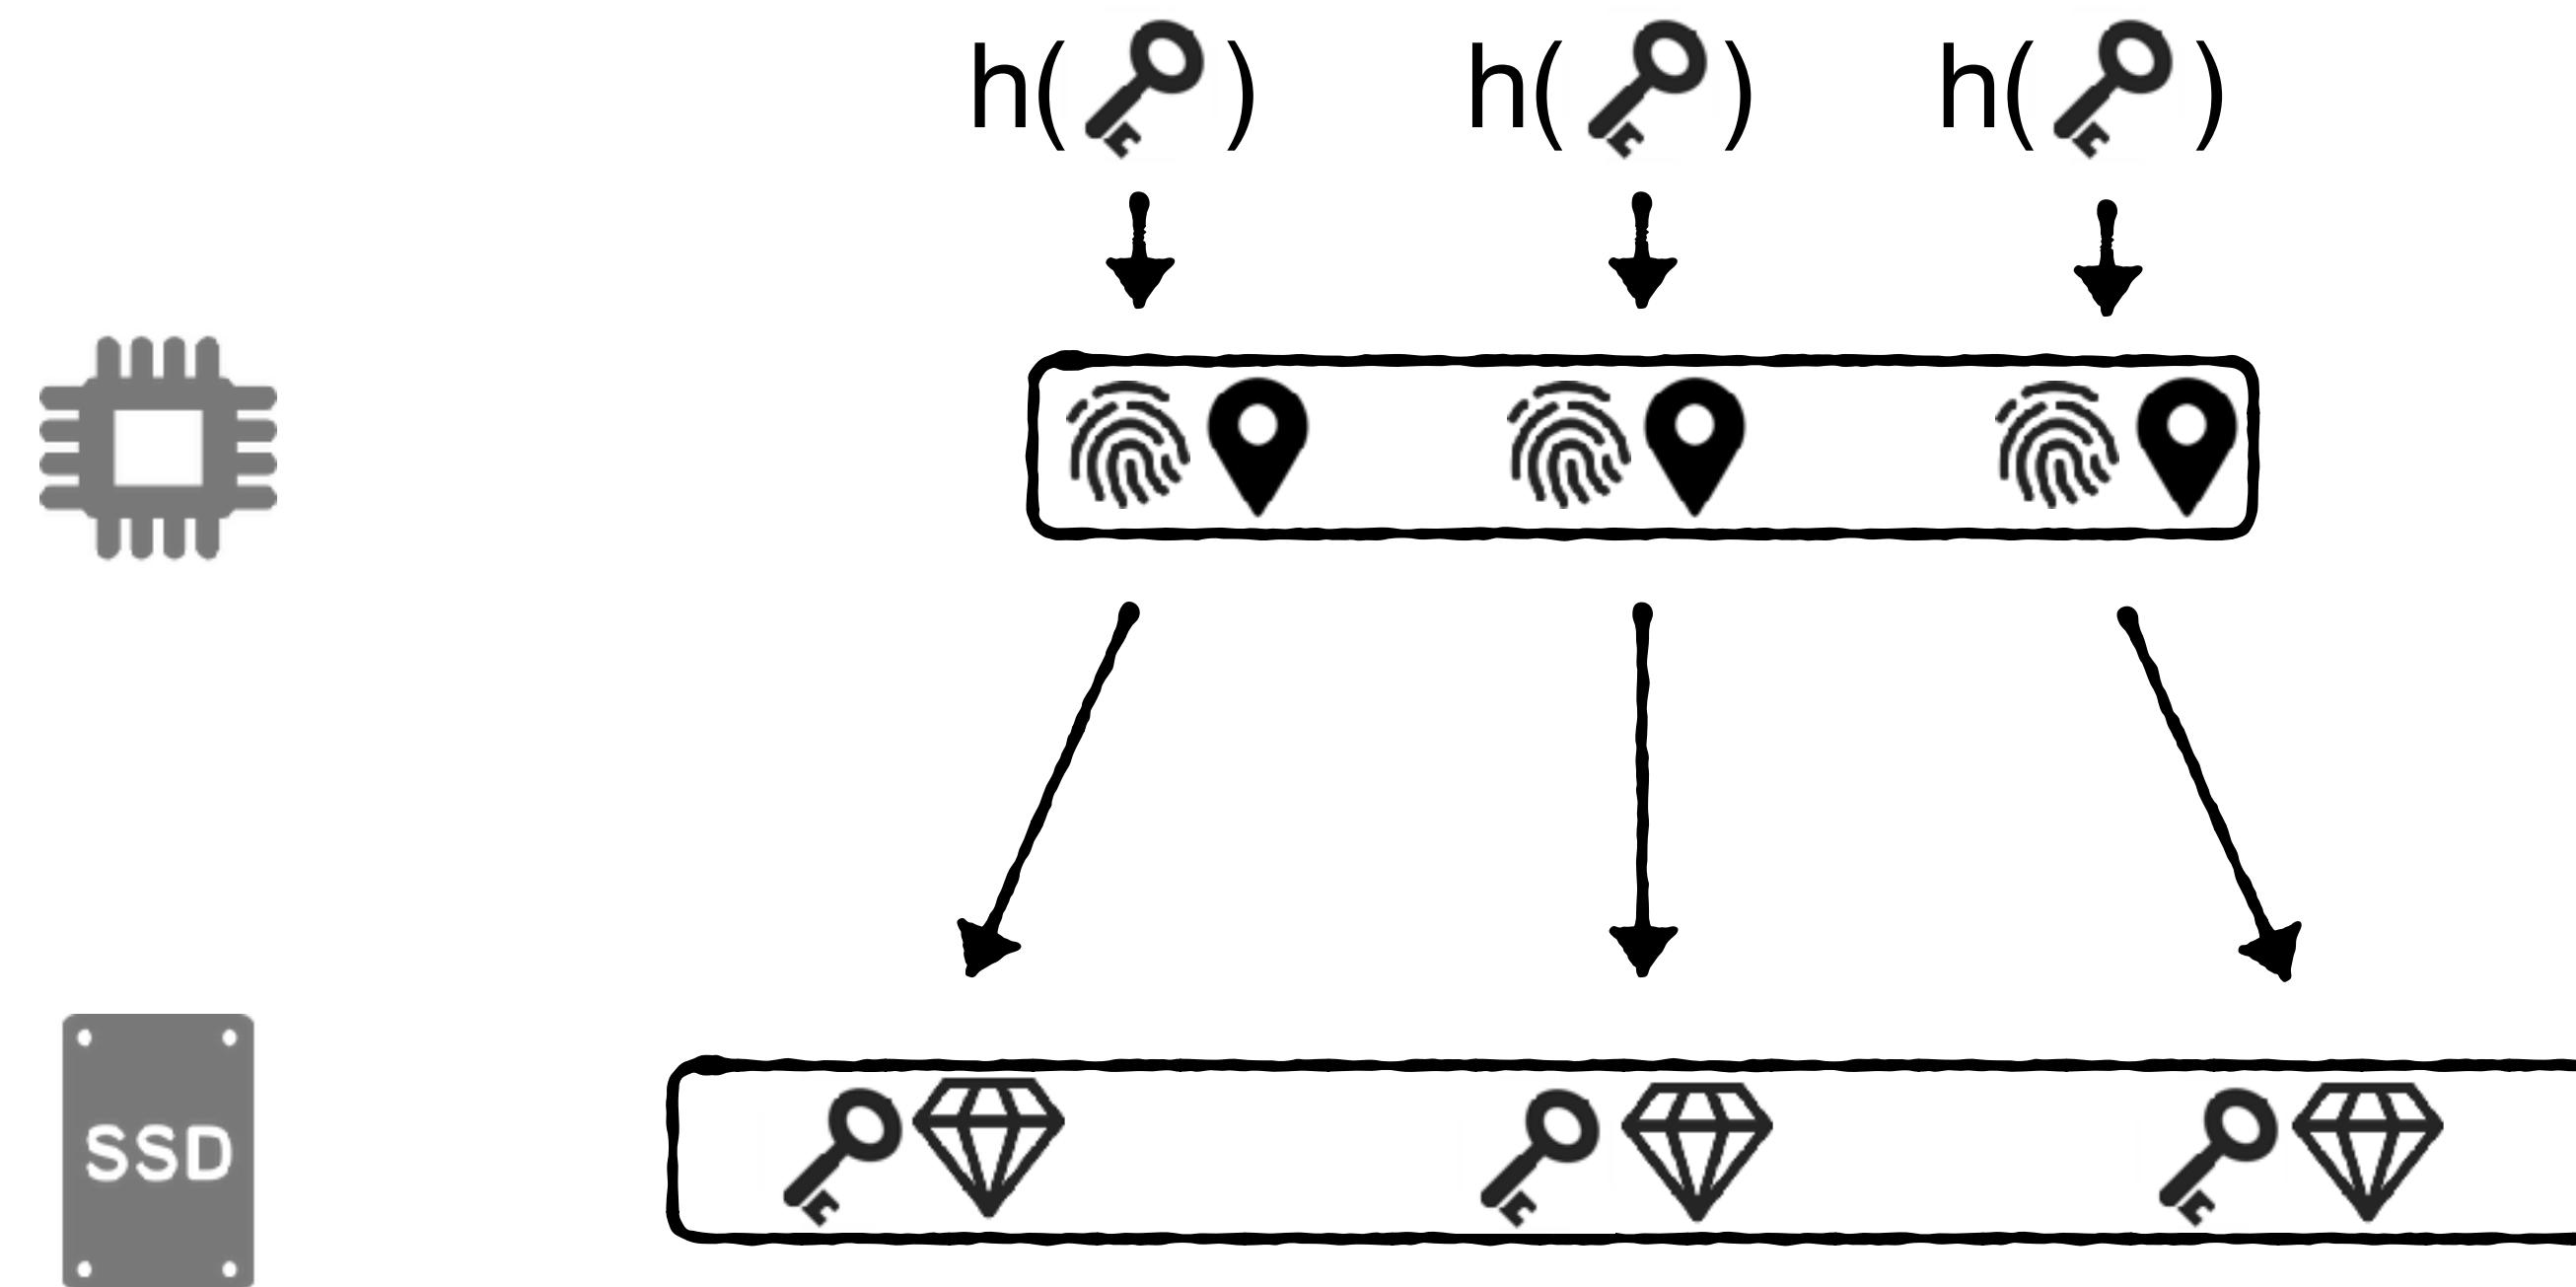
Index can be hash table containing keys (e.g., bitcask)



can we get away with not storing keys?

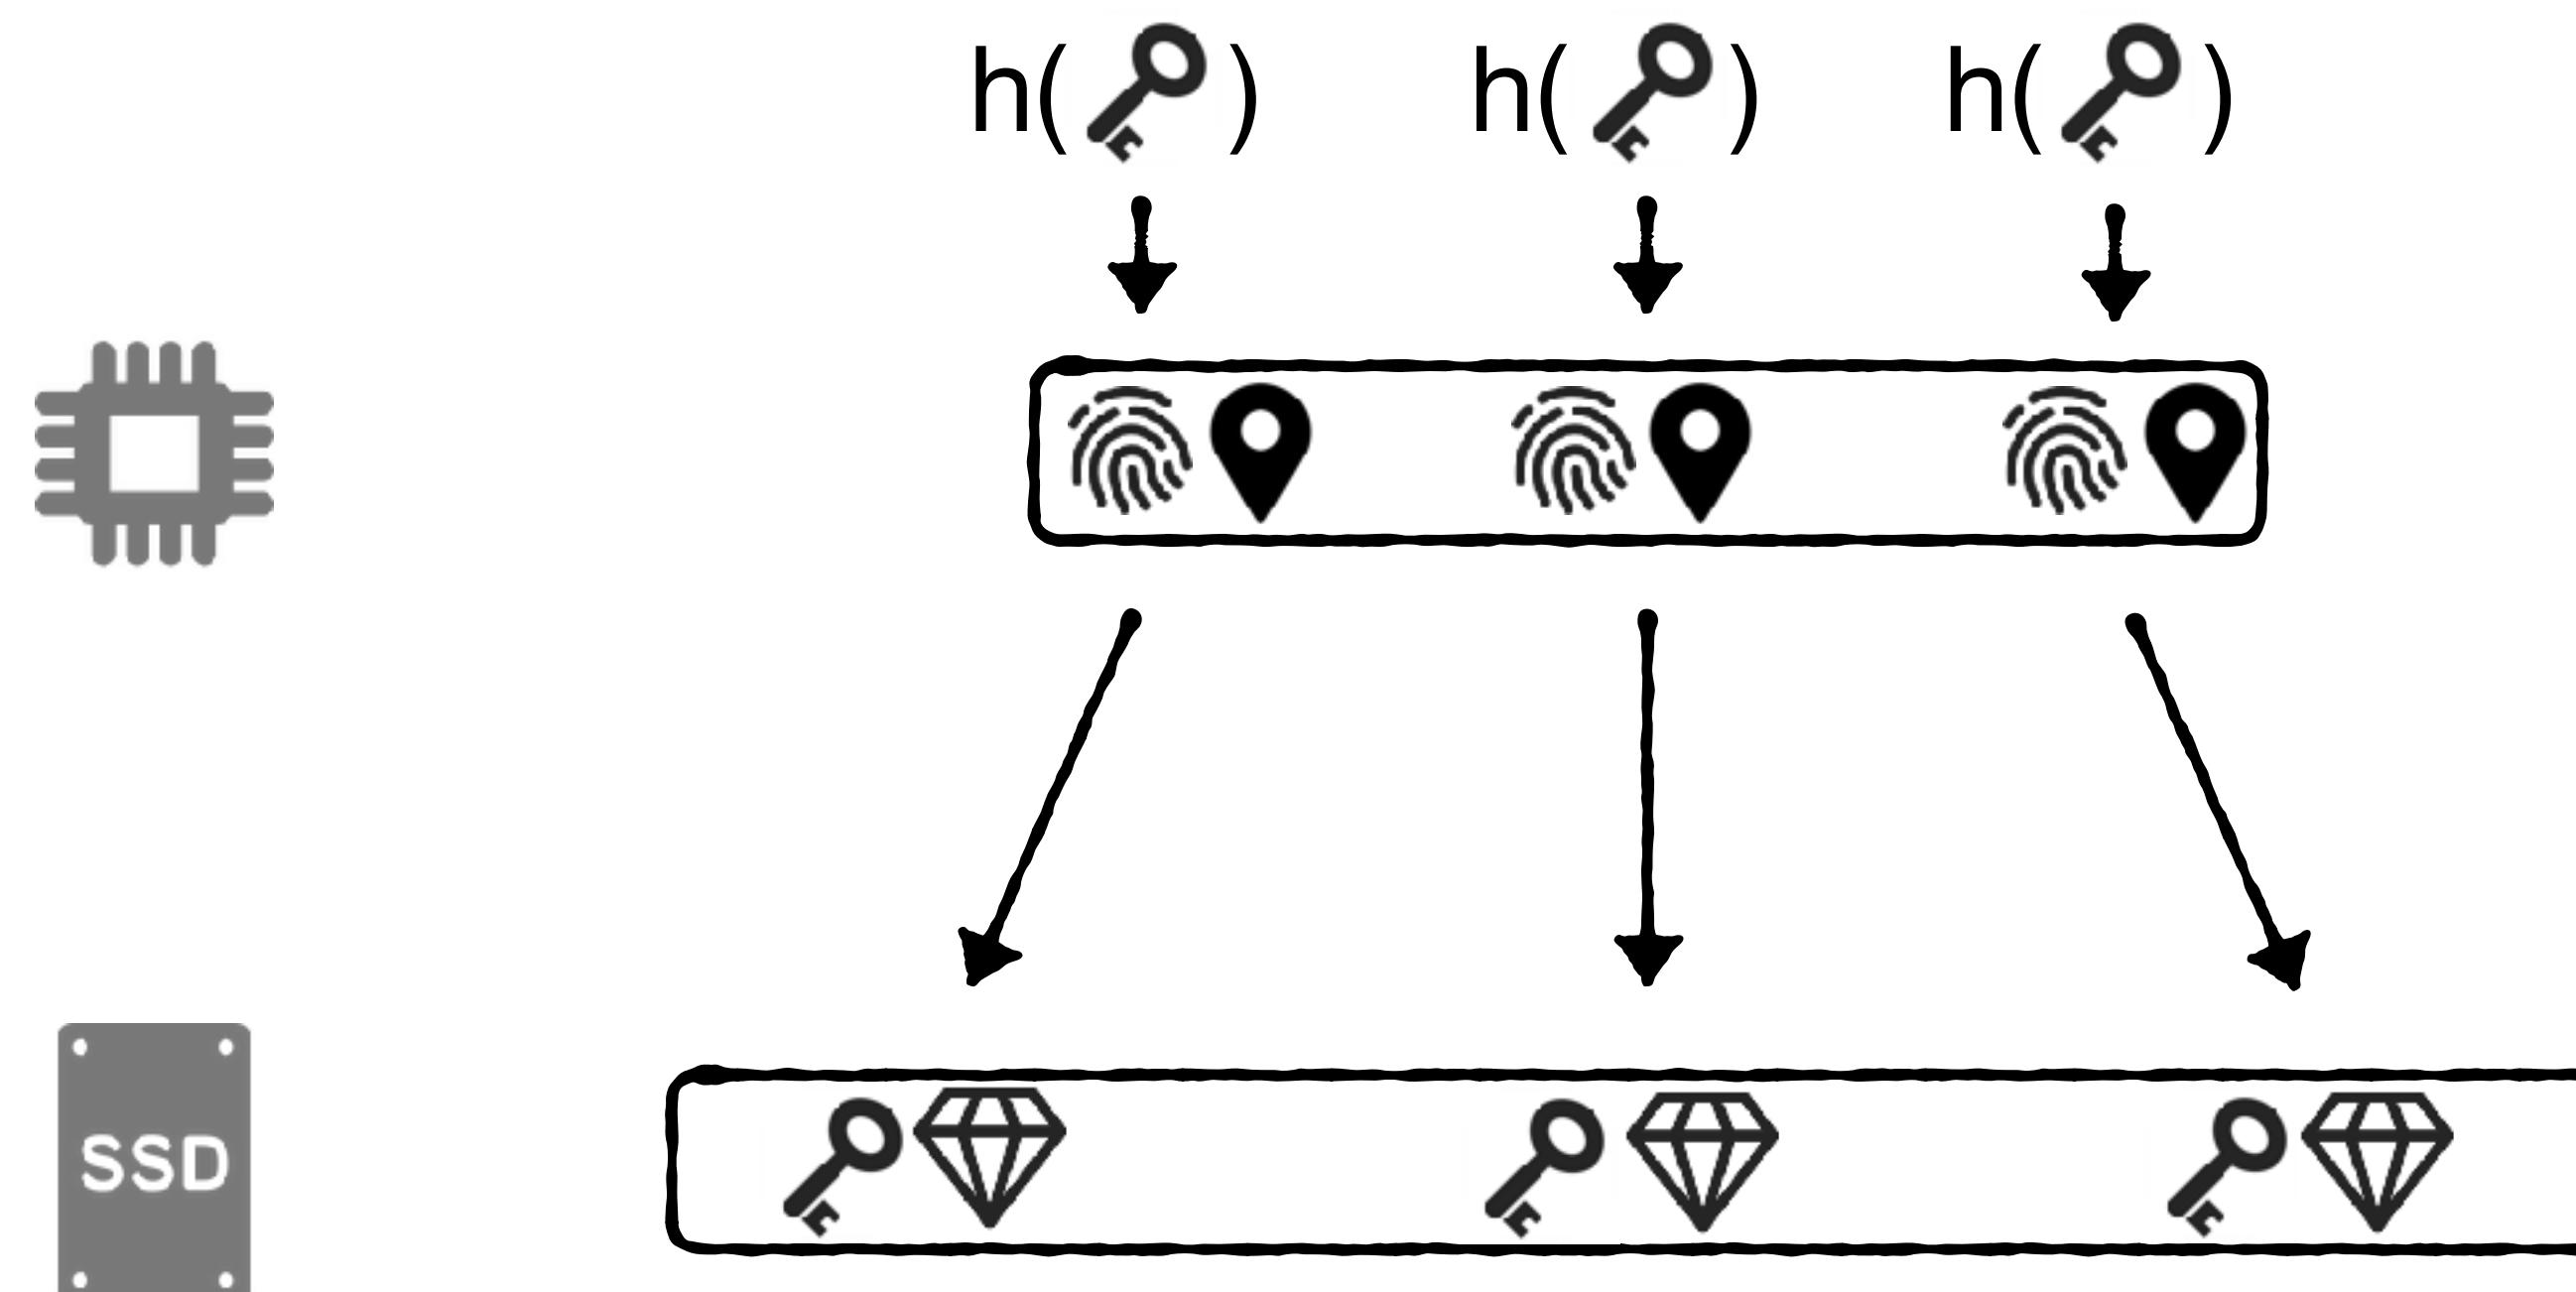


Use a filter (e.g., quotient filter)



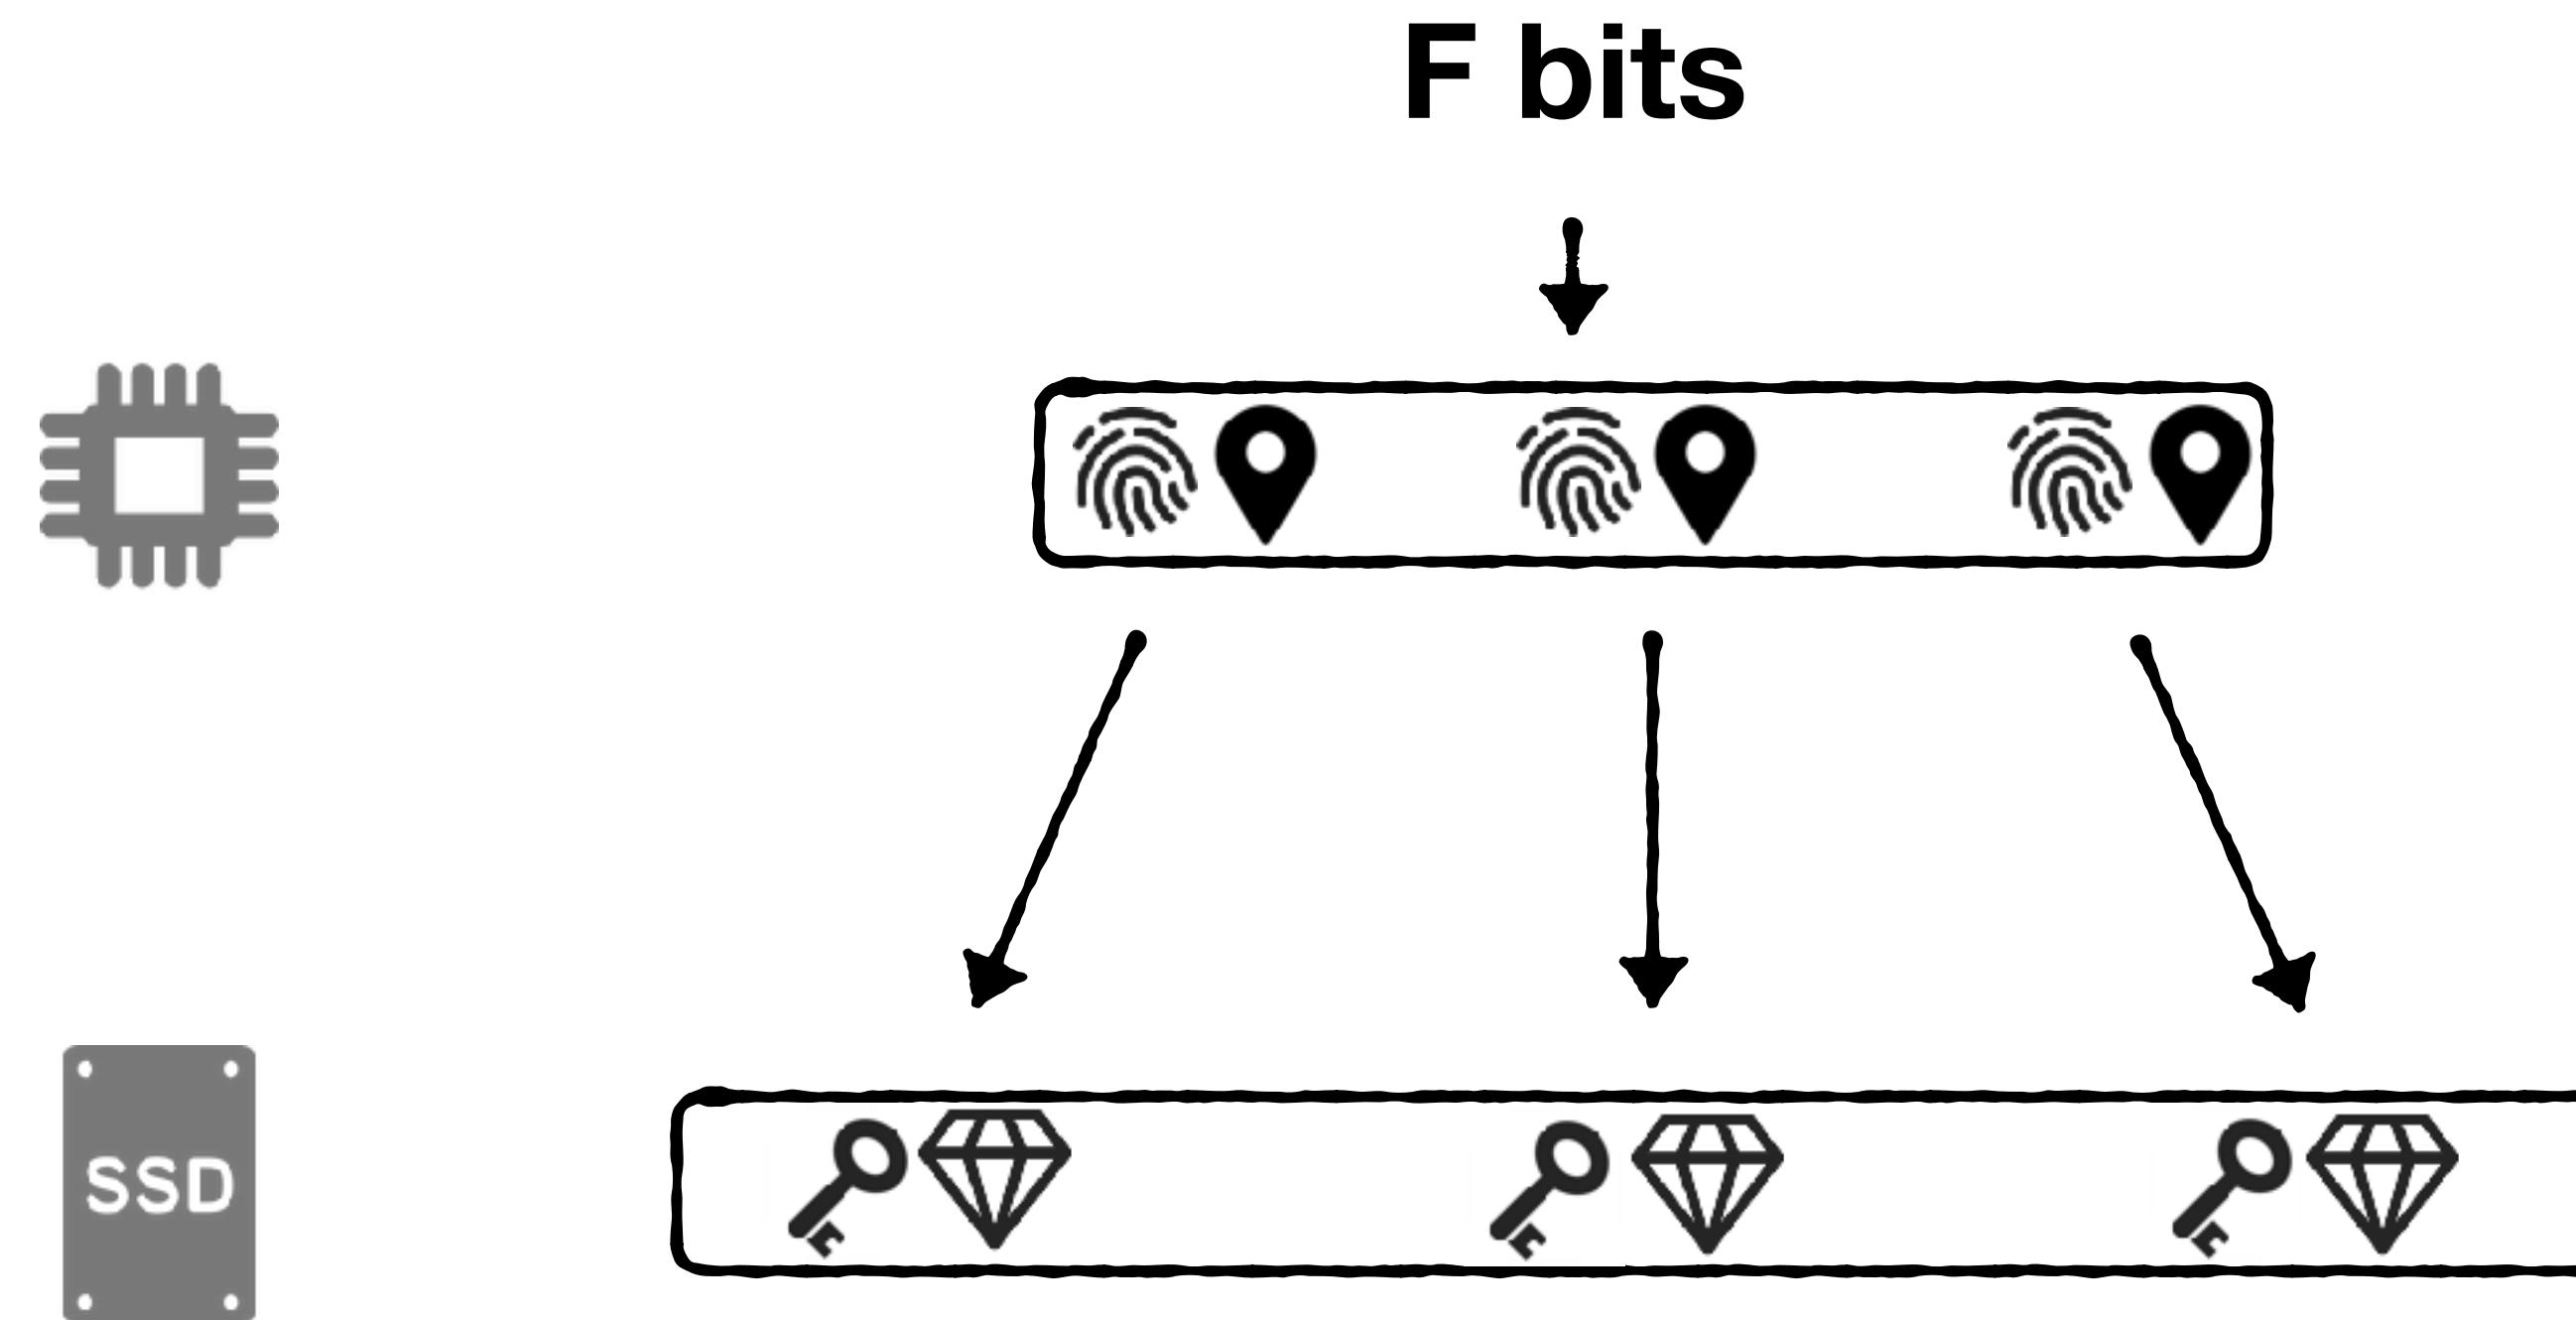
Use a filter (e.g., quotient filter)

Replace keys with fingerprints to save space



Query I/O costs

non-existing key?
existing key?

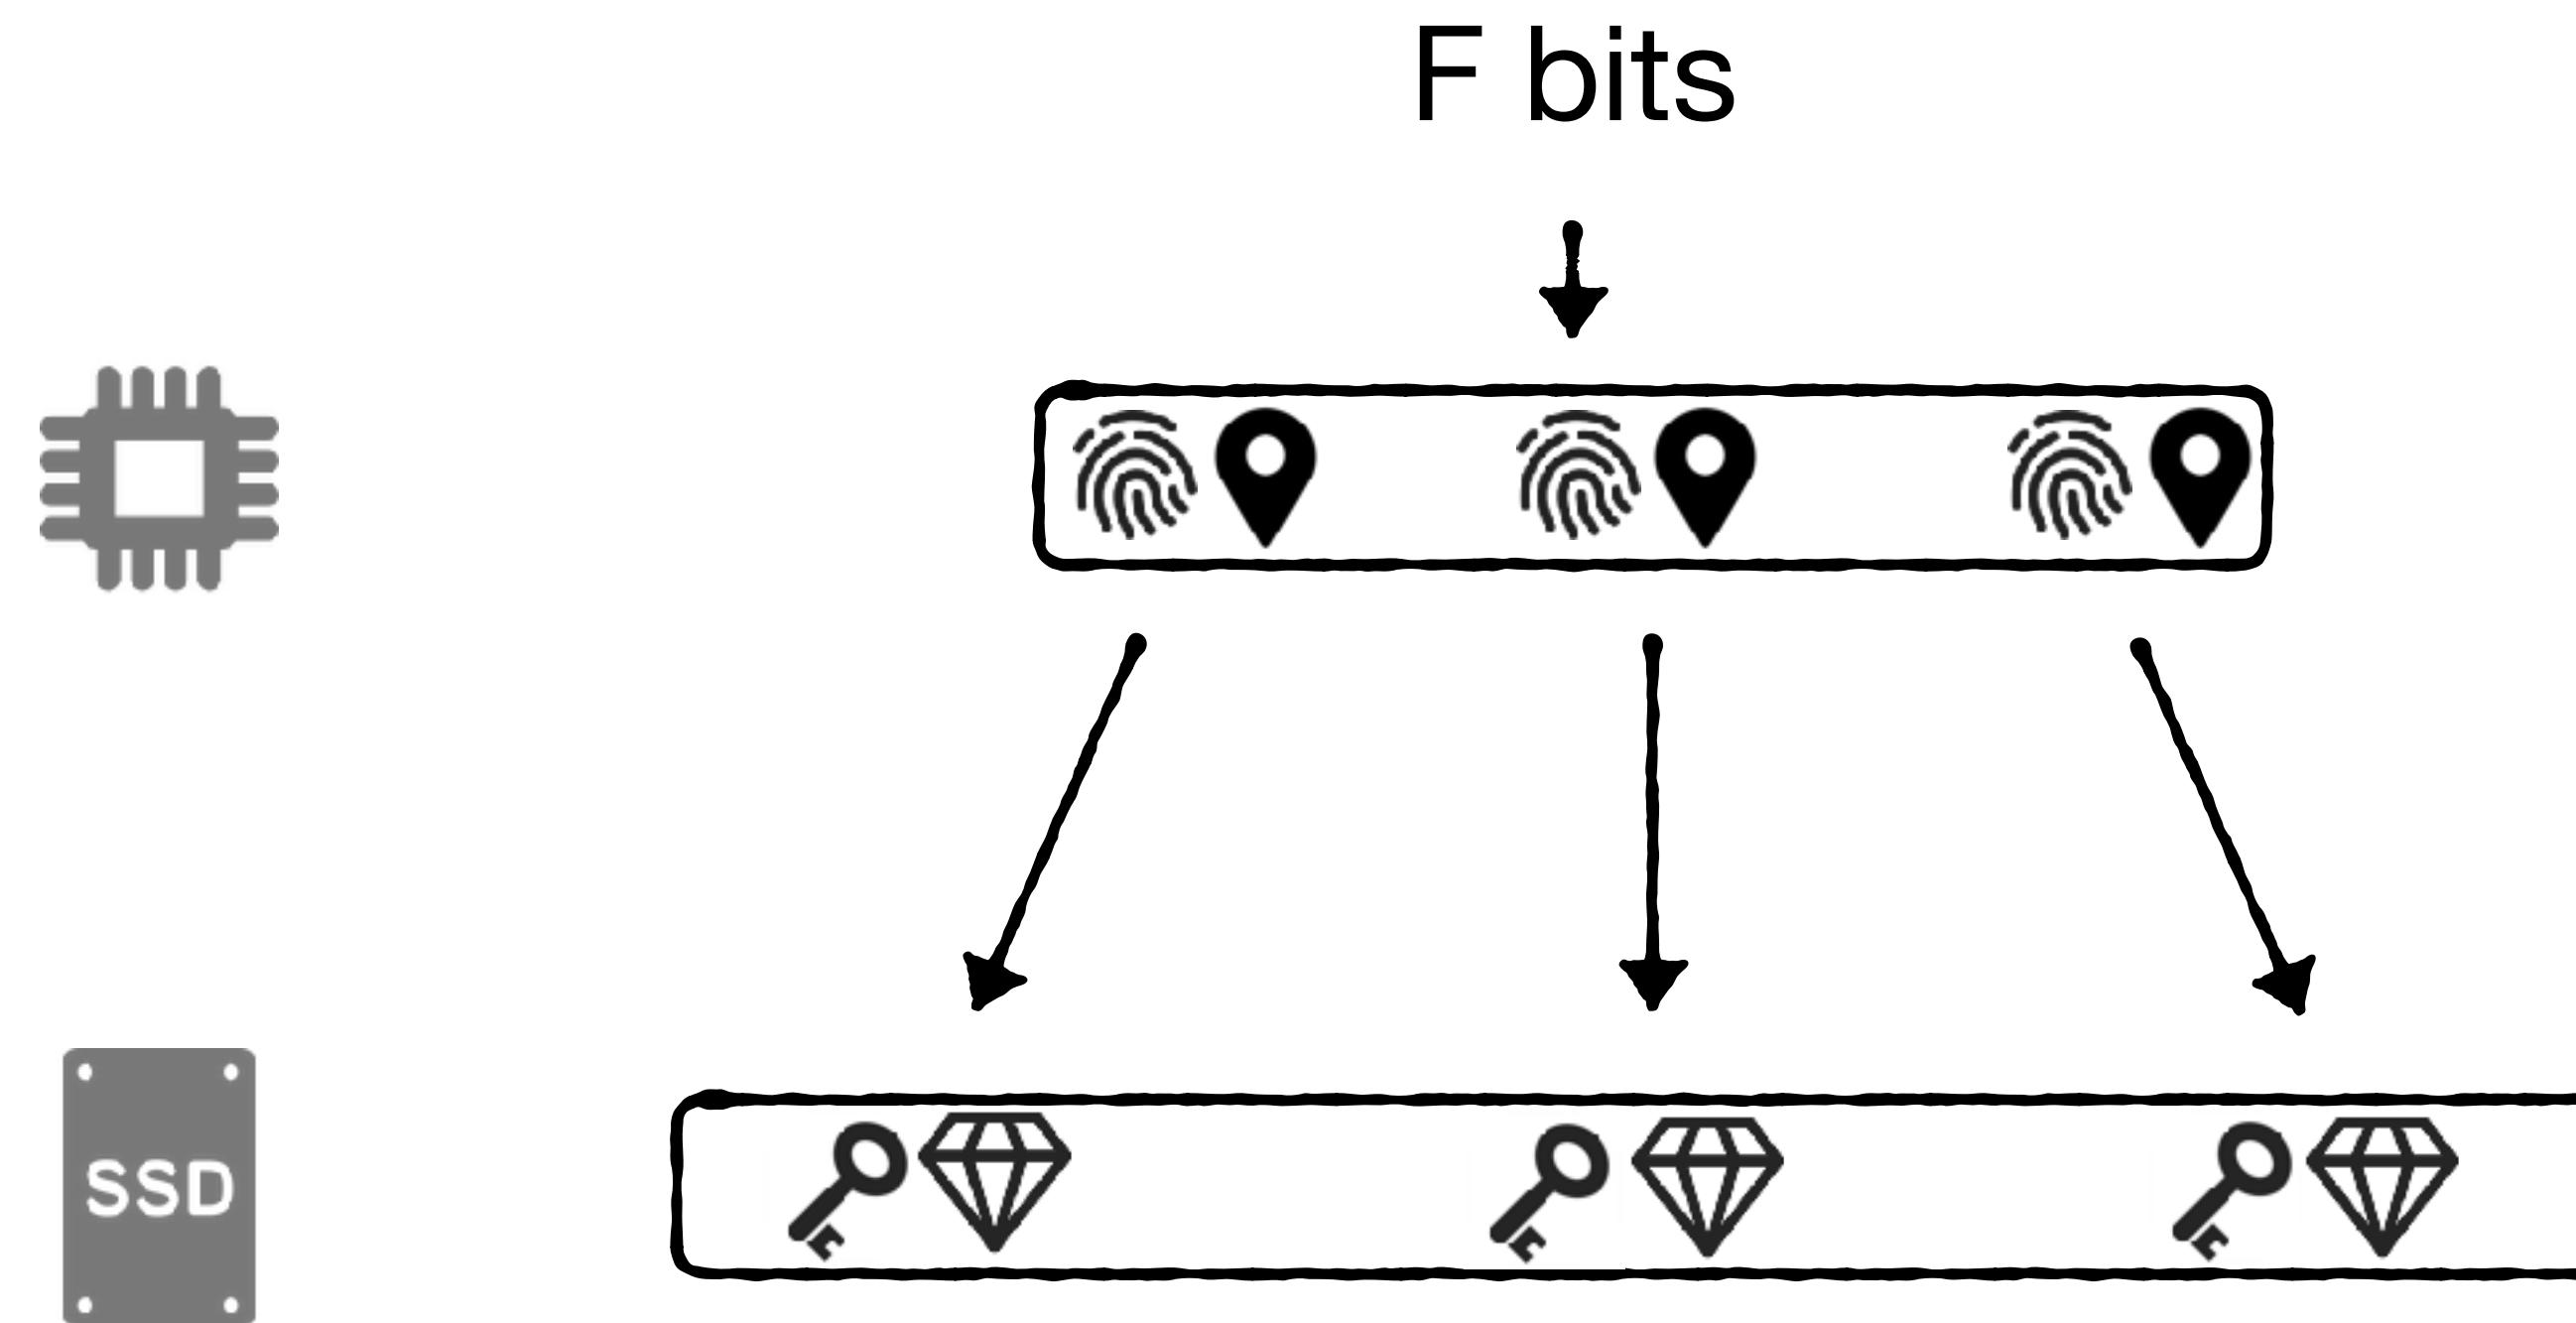


Query I/O costs

non-existing key?

2^F

existing key?



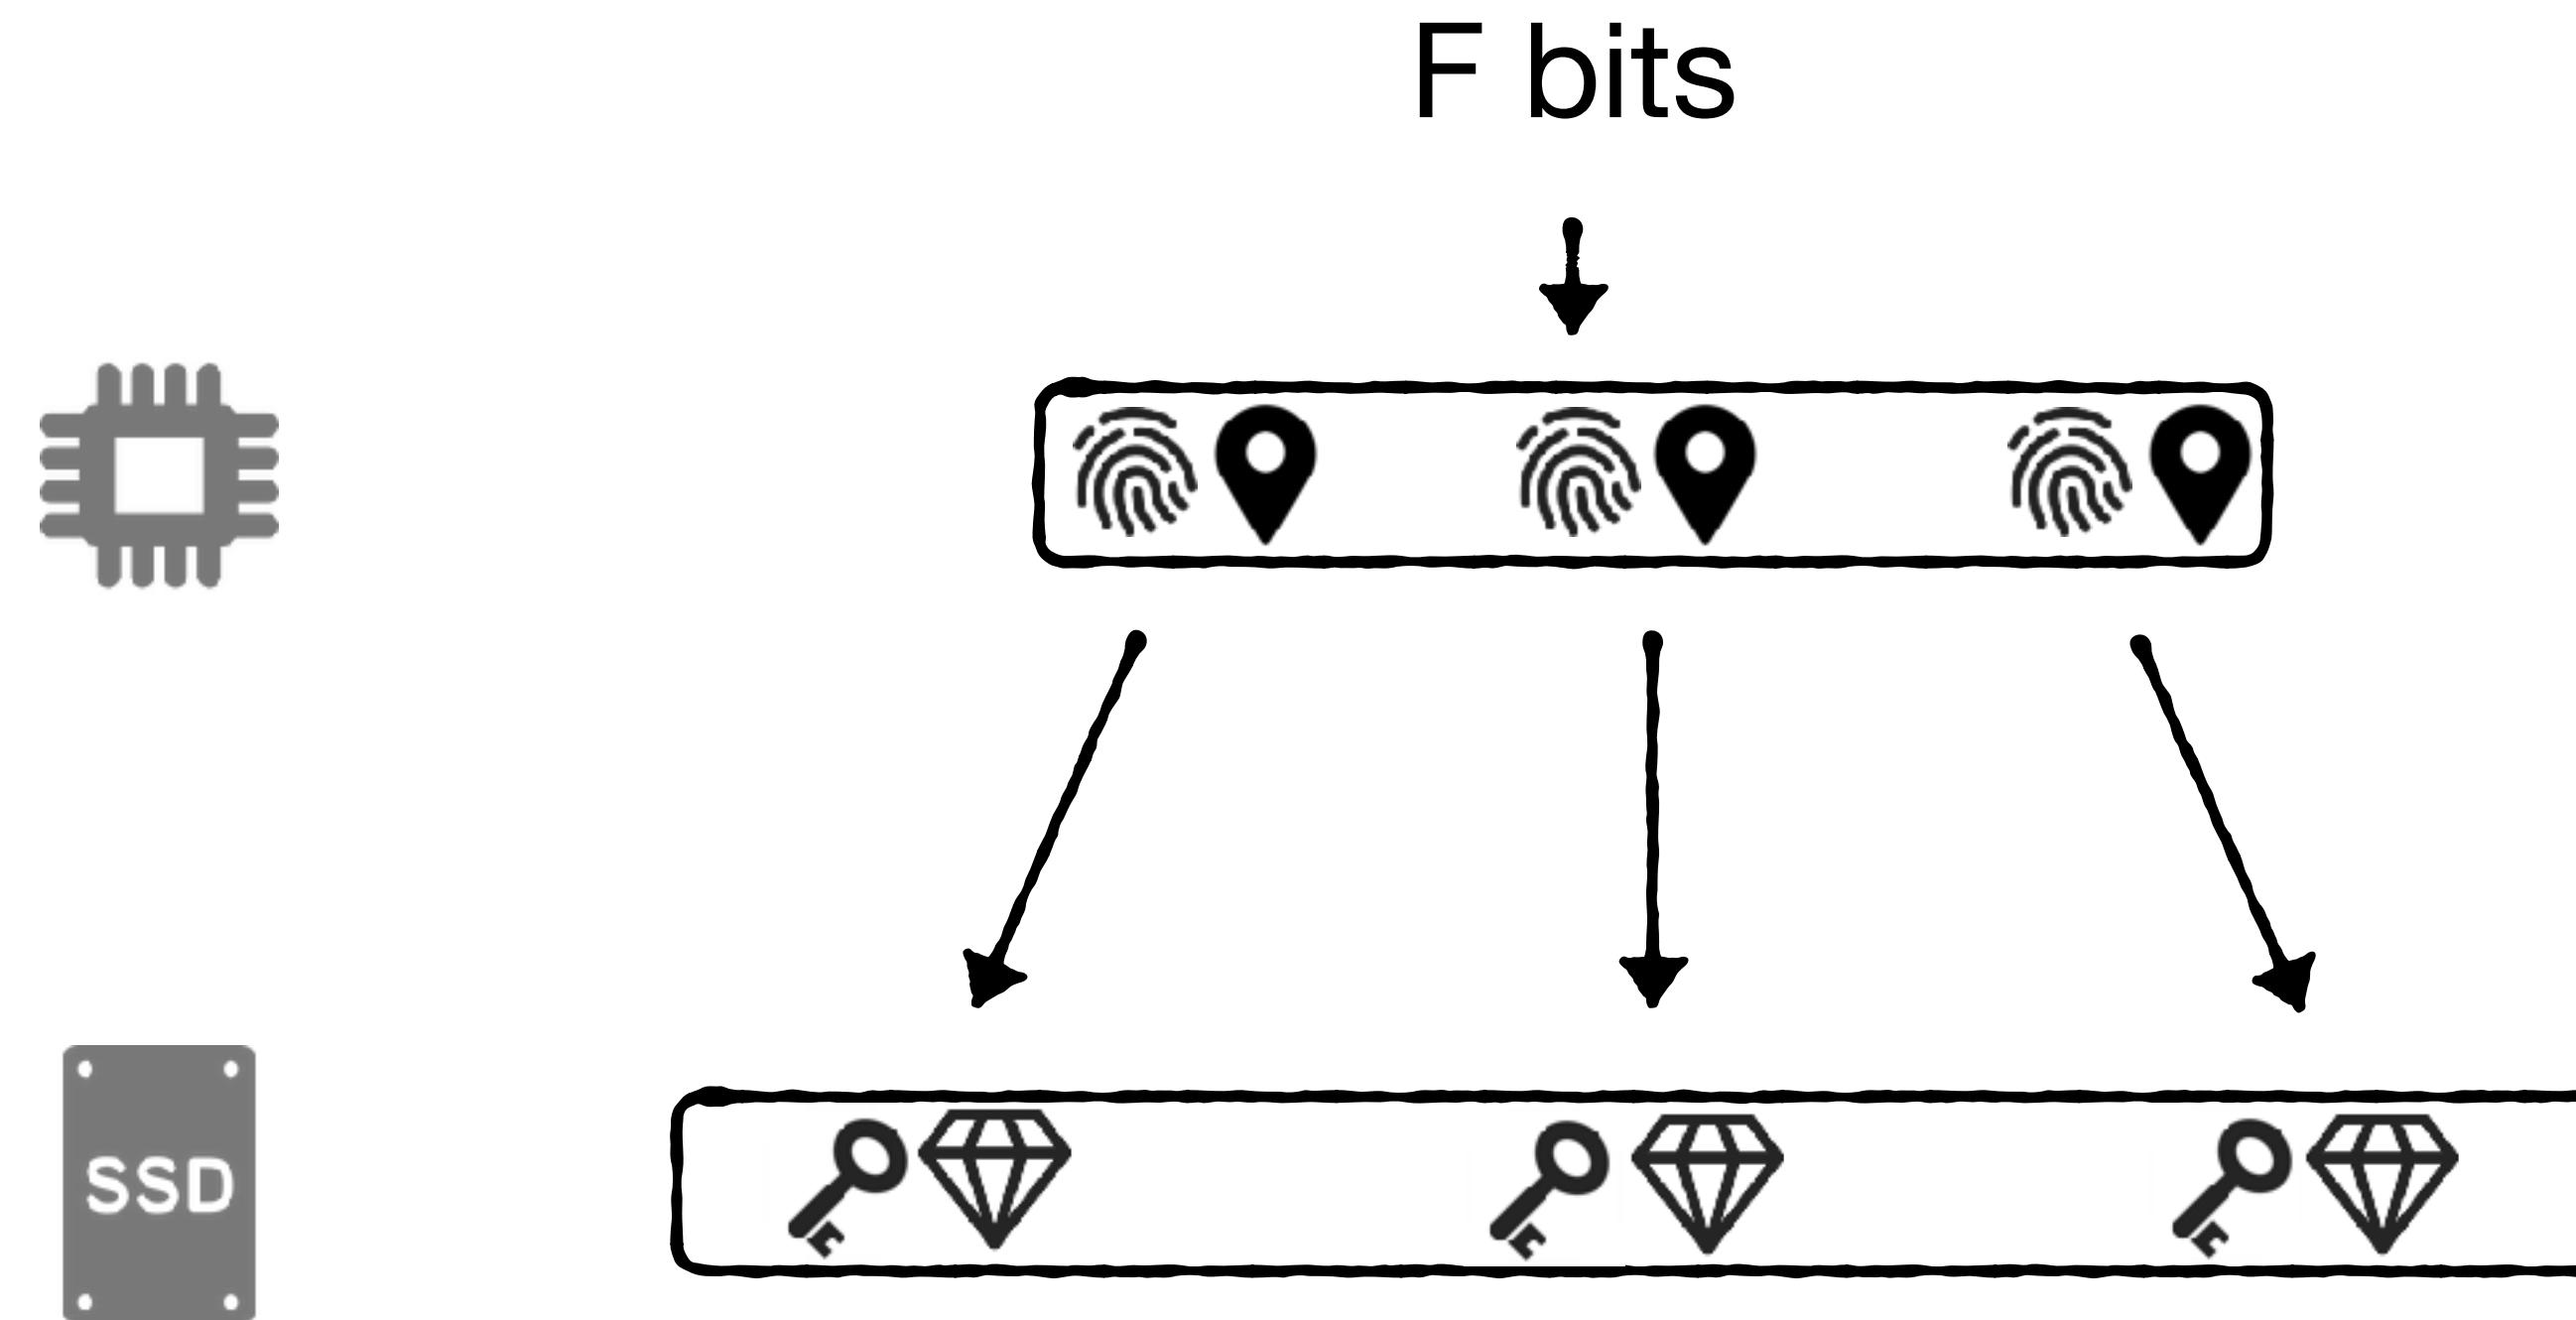
Query I/O costs

non-existing key?

2^{-F}

existing key?

$1+2^{-F}$



Query I/O costs

non-existing key?

2^{-F}

existing key?

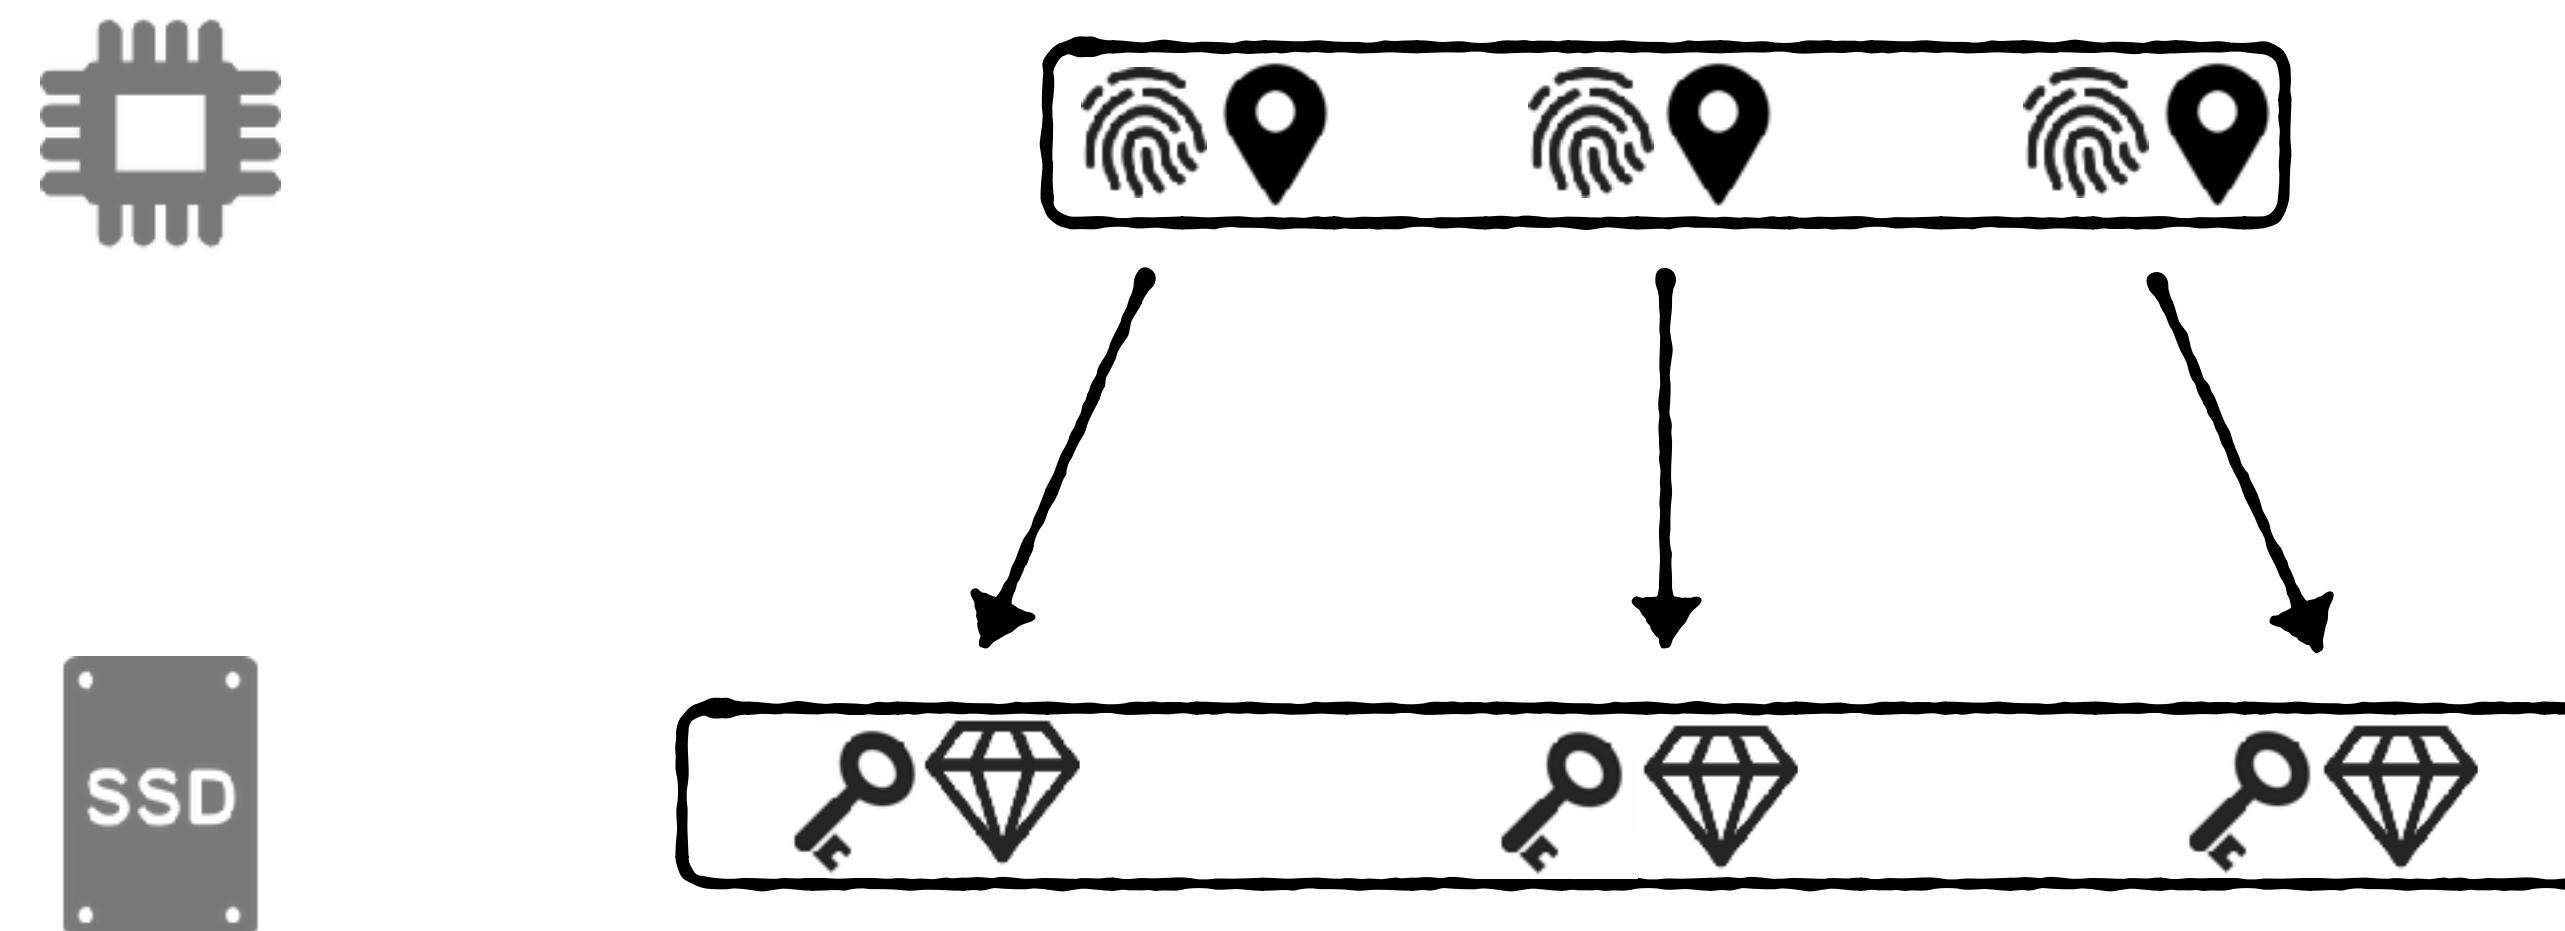
$1+2^{-F}$



Let's focus on queries to existing keys

(1) more common

(2) minimize latency for useful work



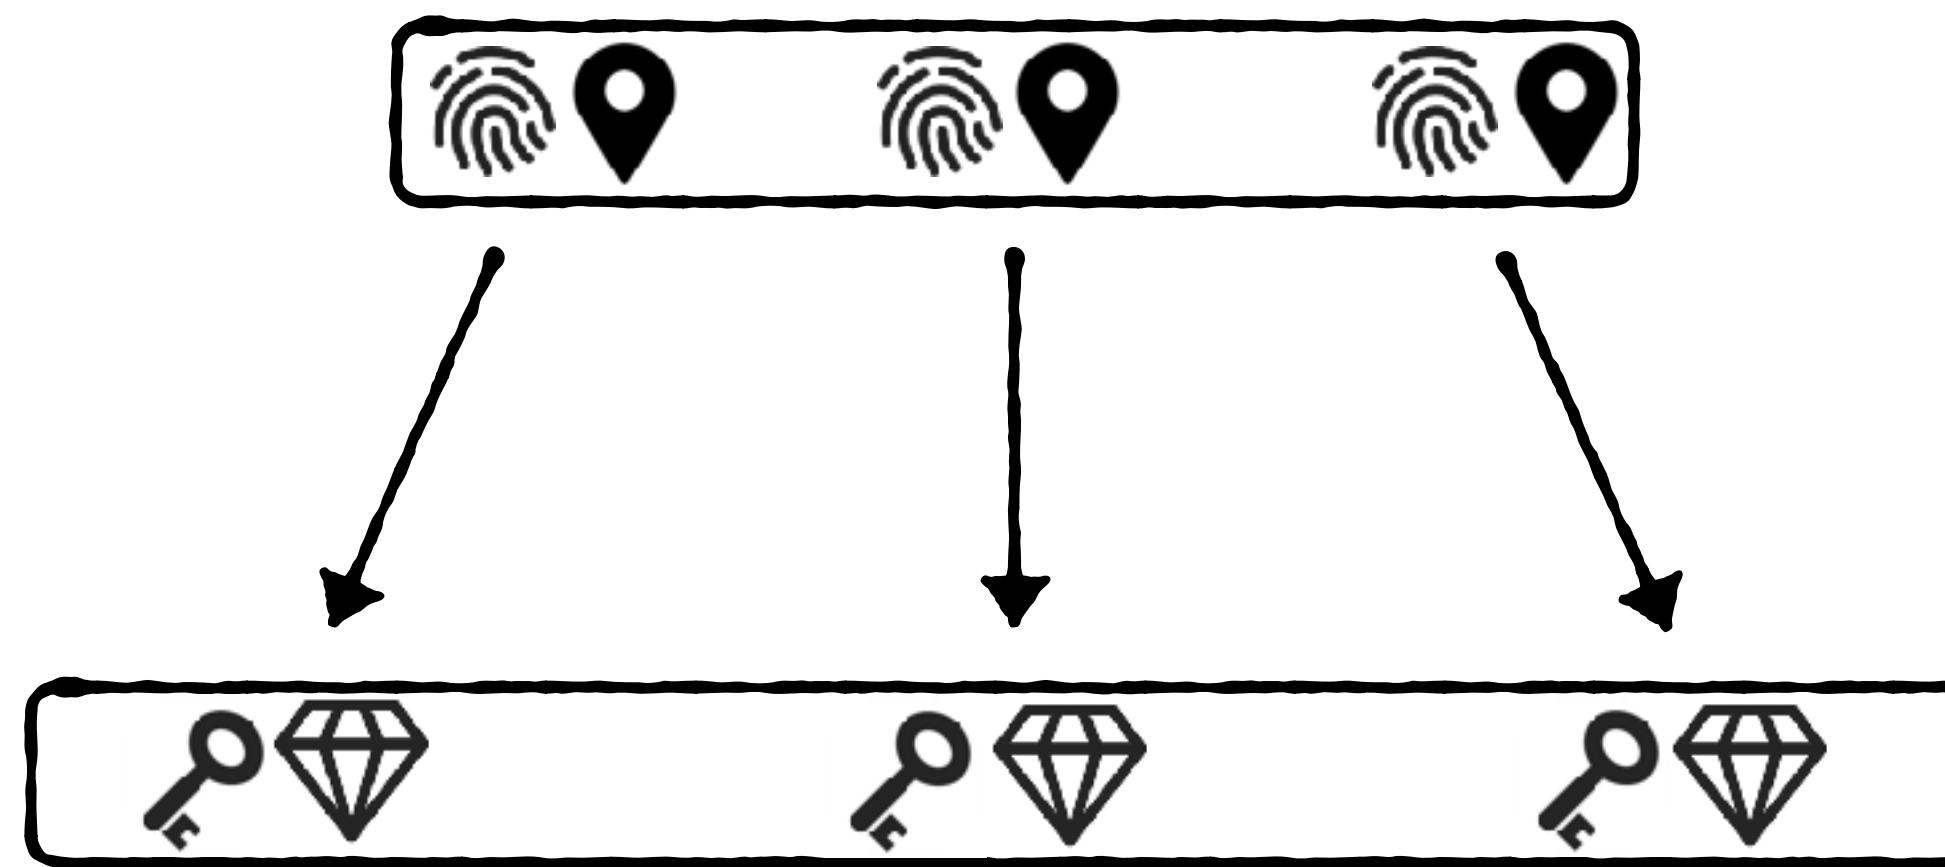
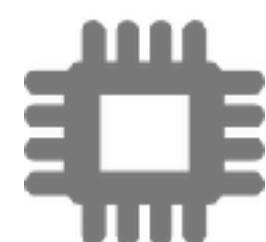
Query I/O costs

existing key?

1+2-F



Due to fingerprint collisions



Query I/O costs

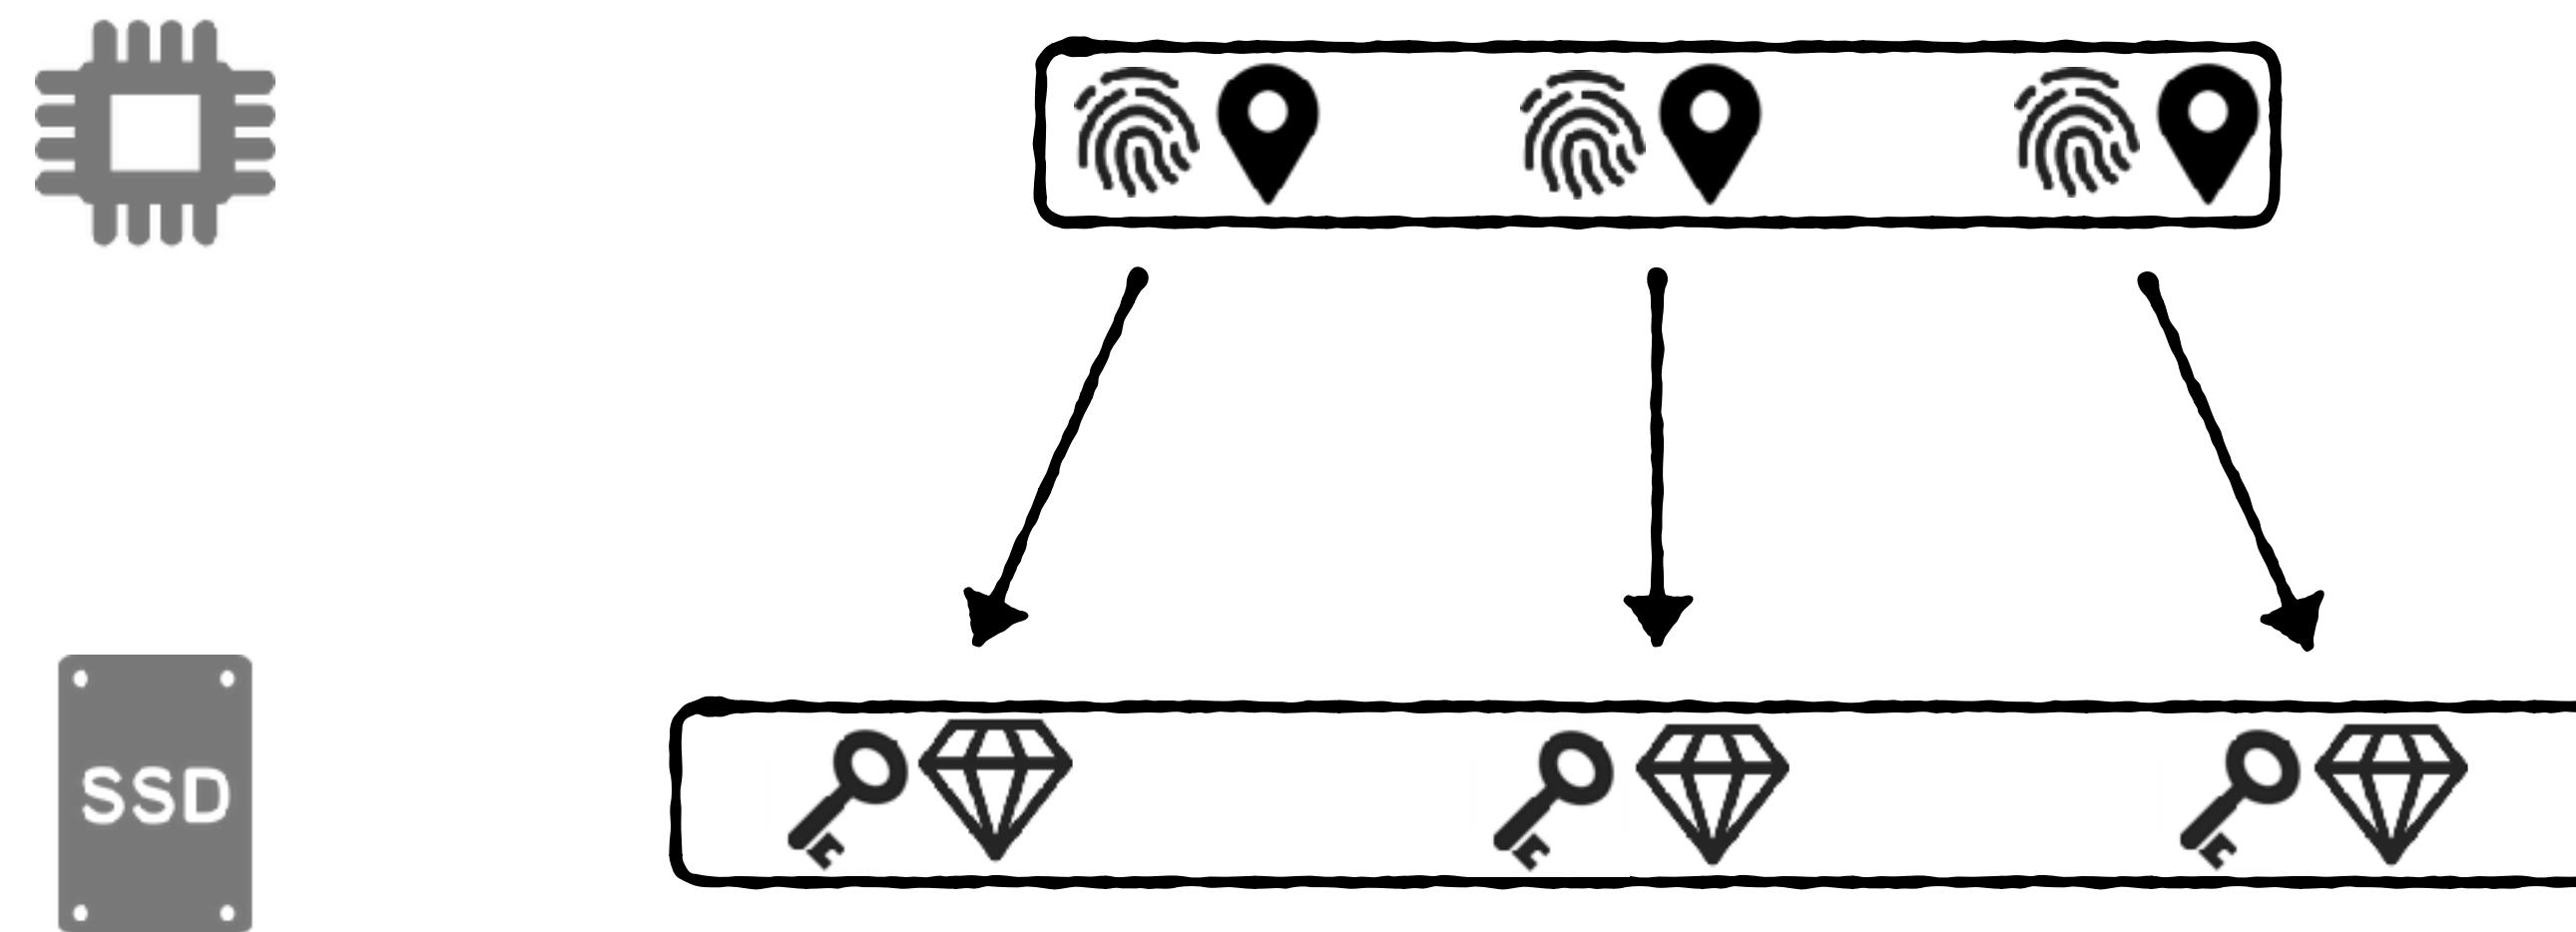
existing key?

$1+2^{-F}$



Due to fingerprint collisions

Can we reduce by increasing F



Query I/O costs

existing key?

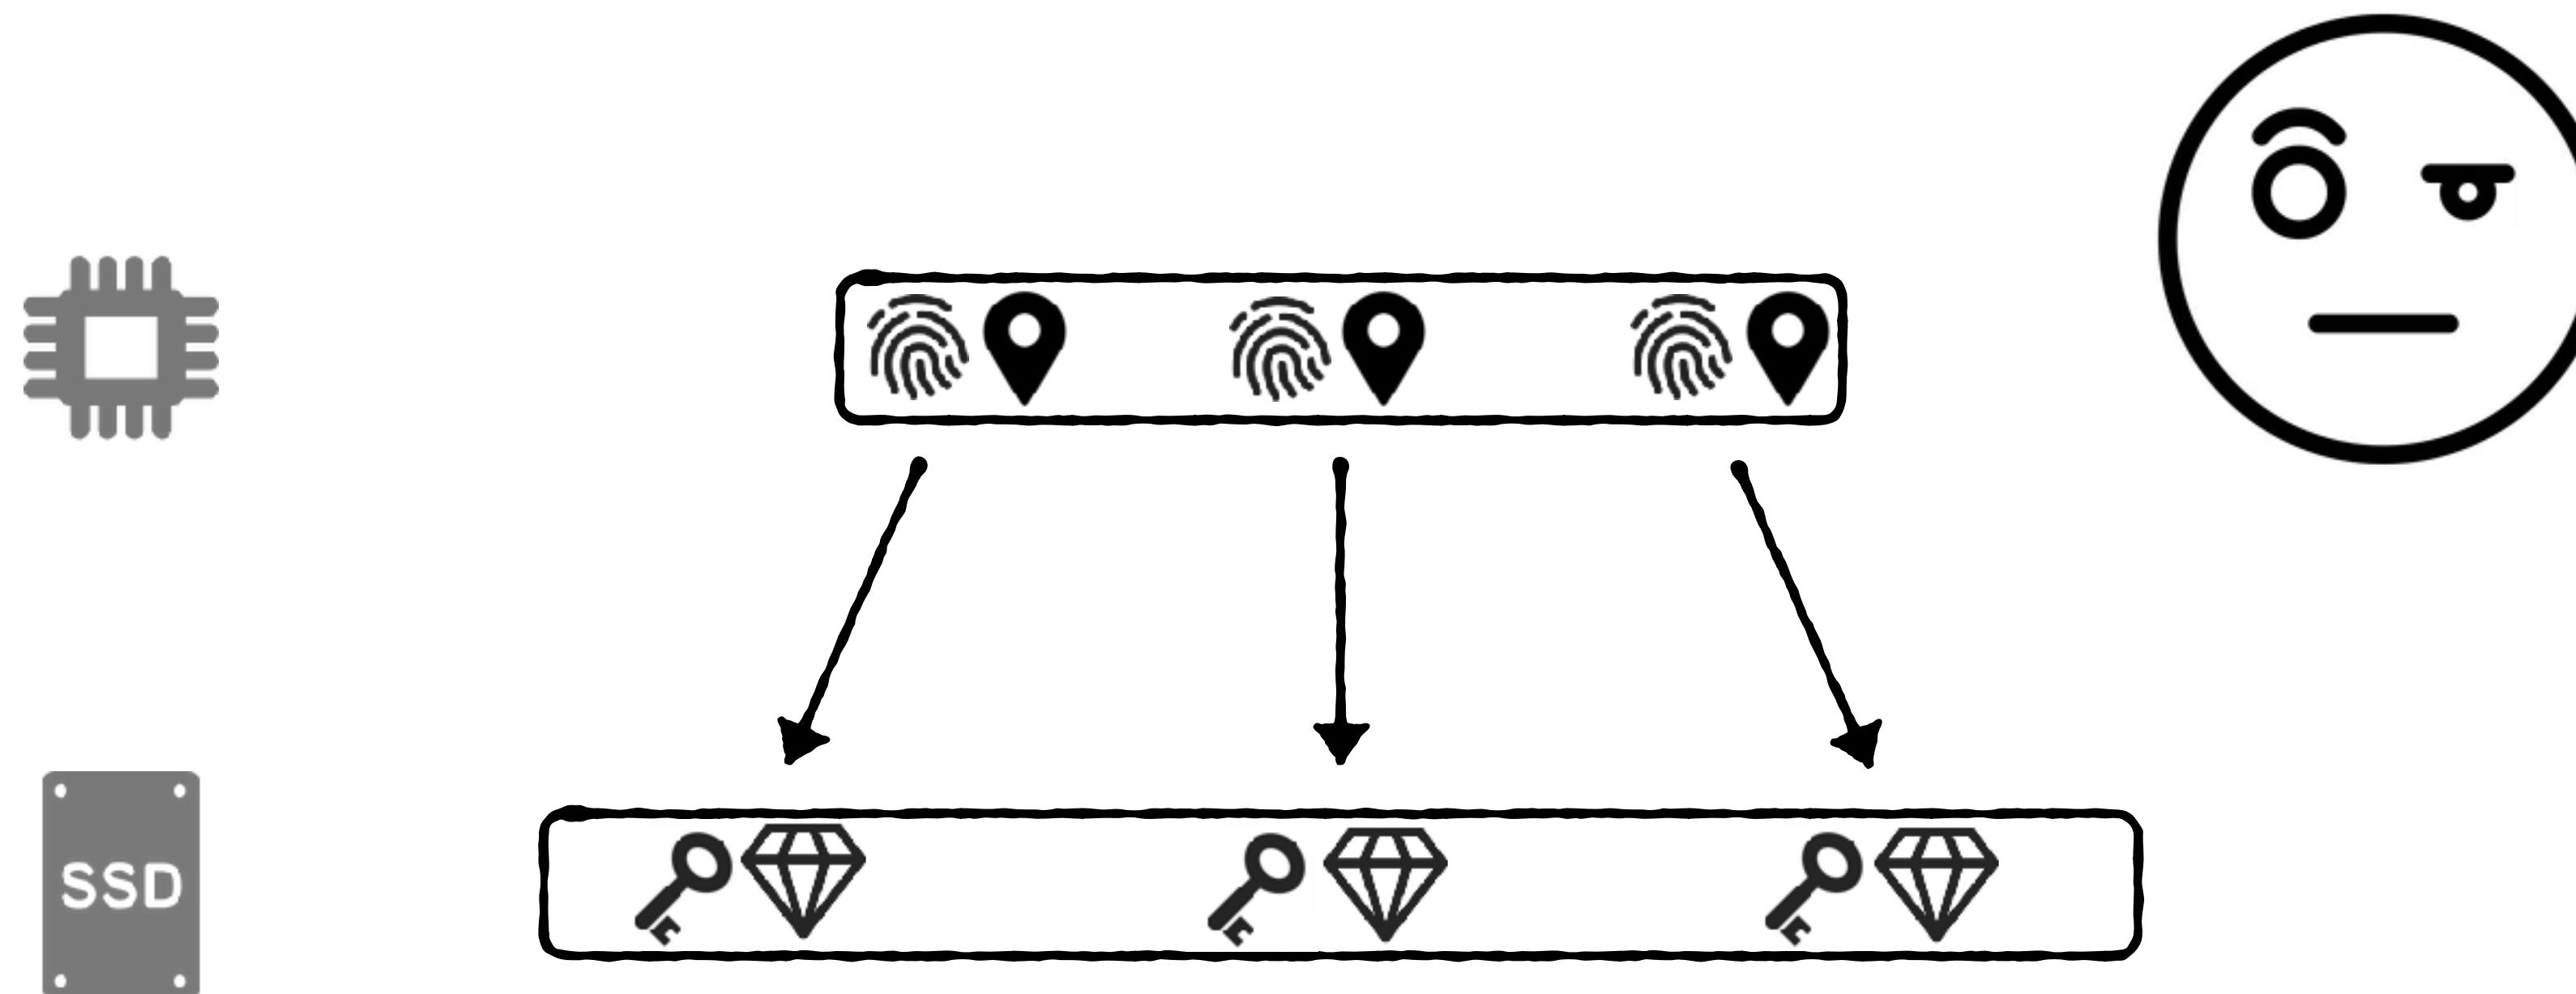
$1+2^{-F}$



Due to fingerprint collisions

Can we reduce by increasing F

Is there a better way?



Perfect Hashing



Minimal

space-efficient
static data



Dynamic

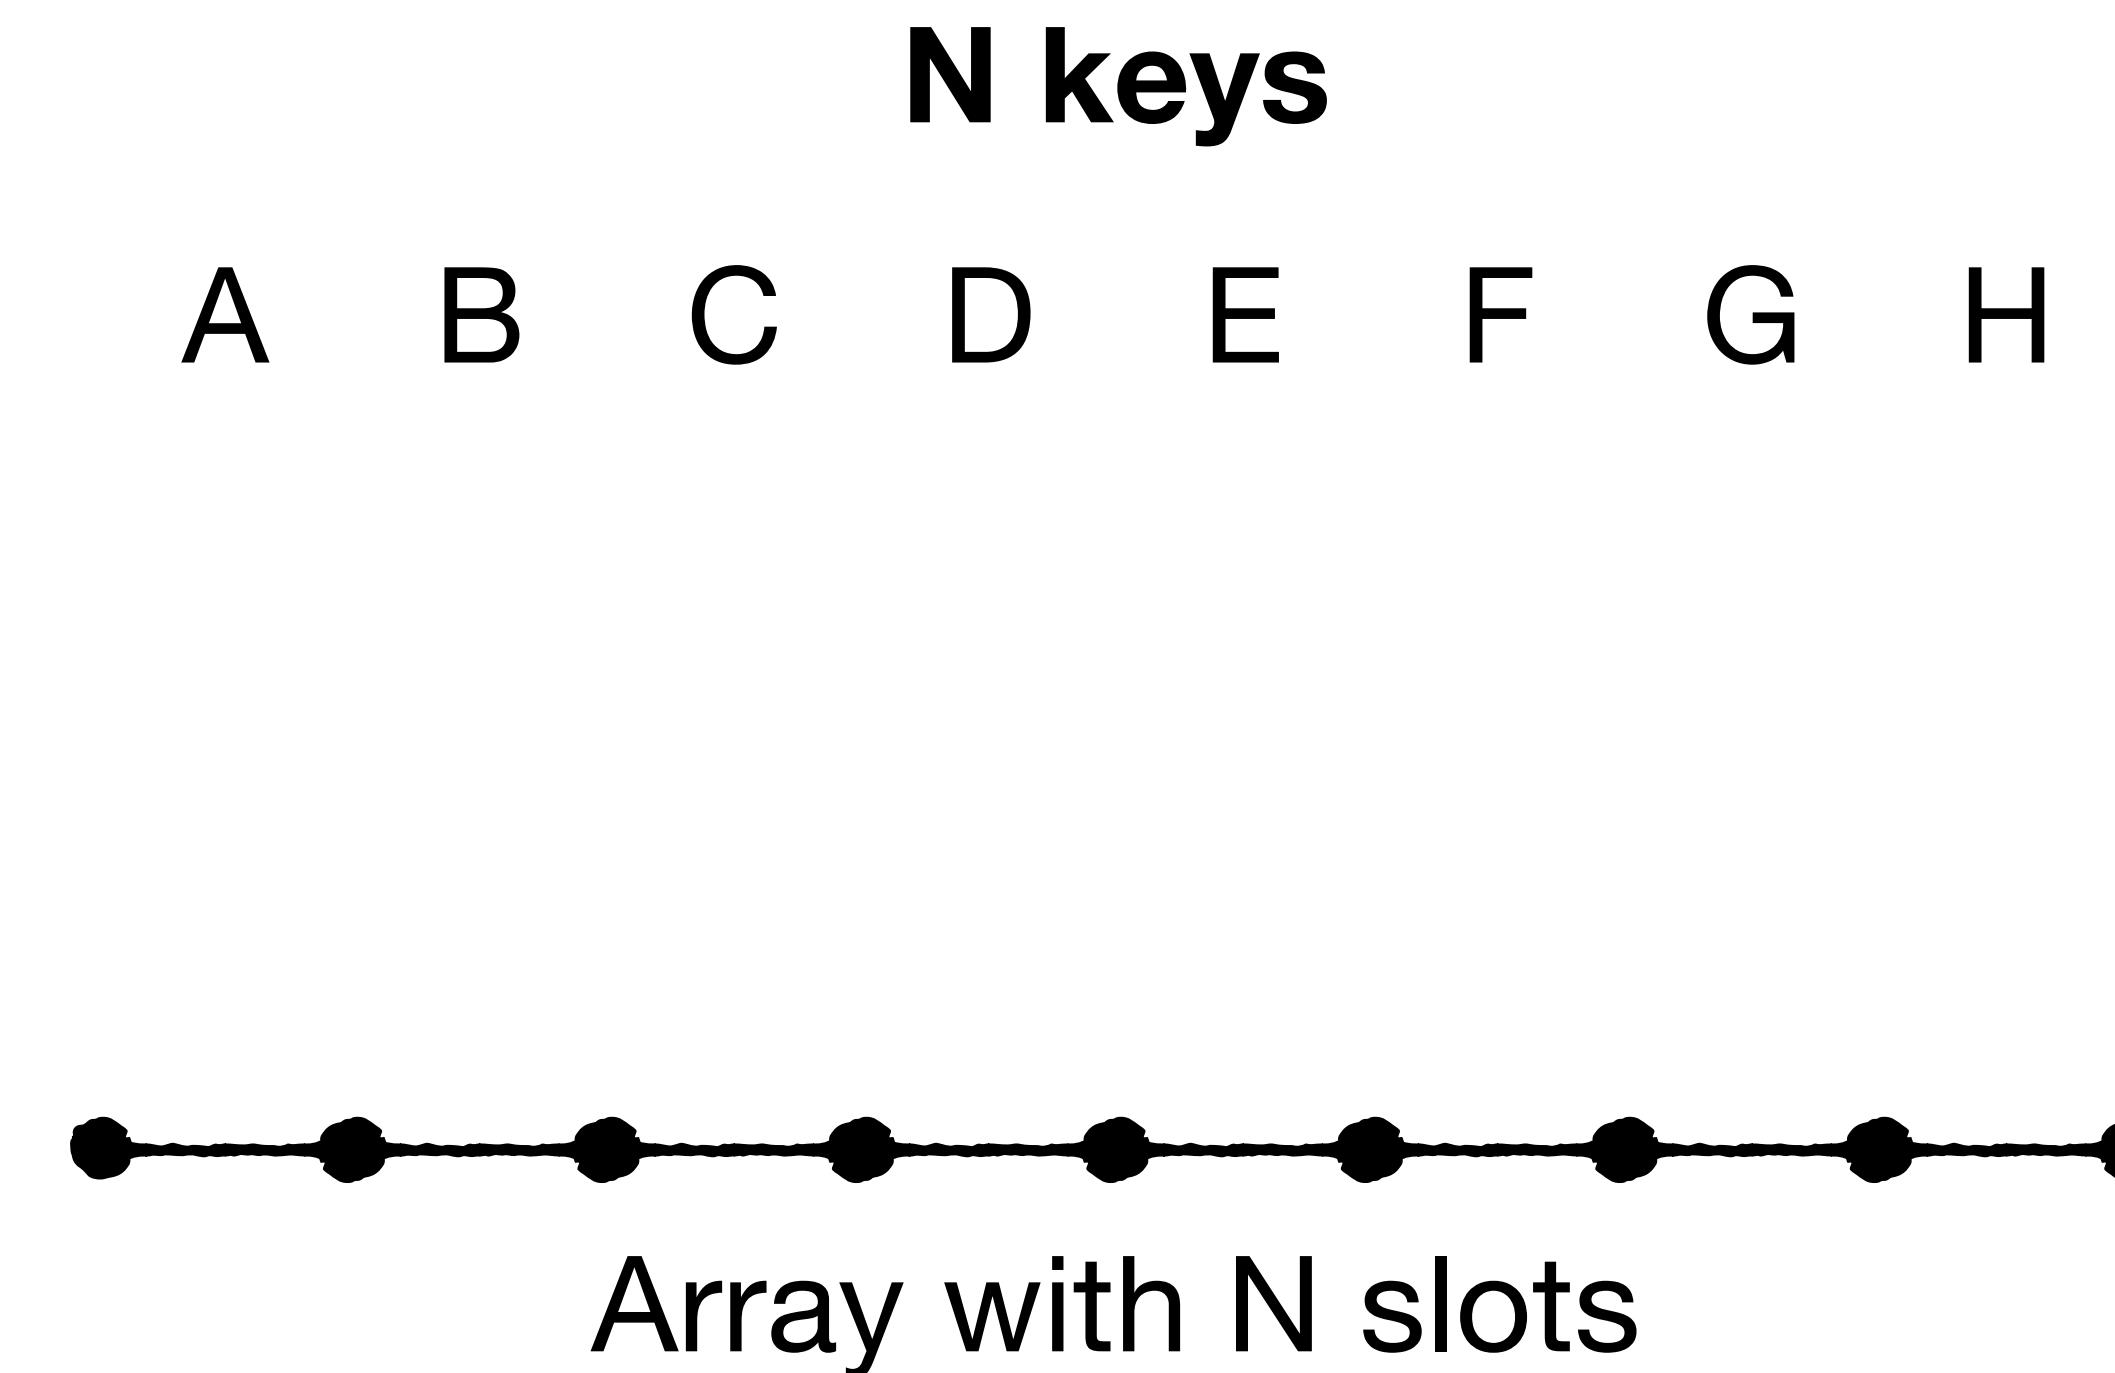
more space
supports insertions

Minimal Perfect Hashing

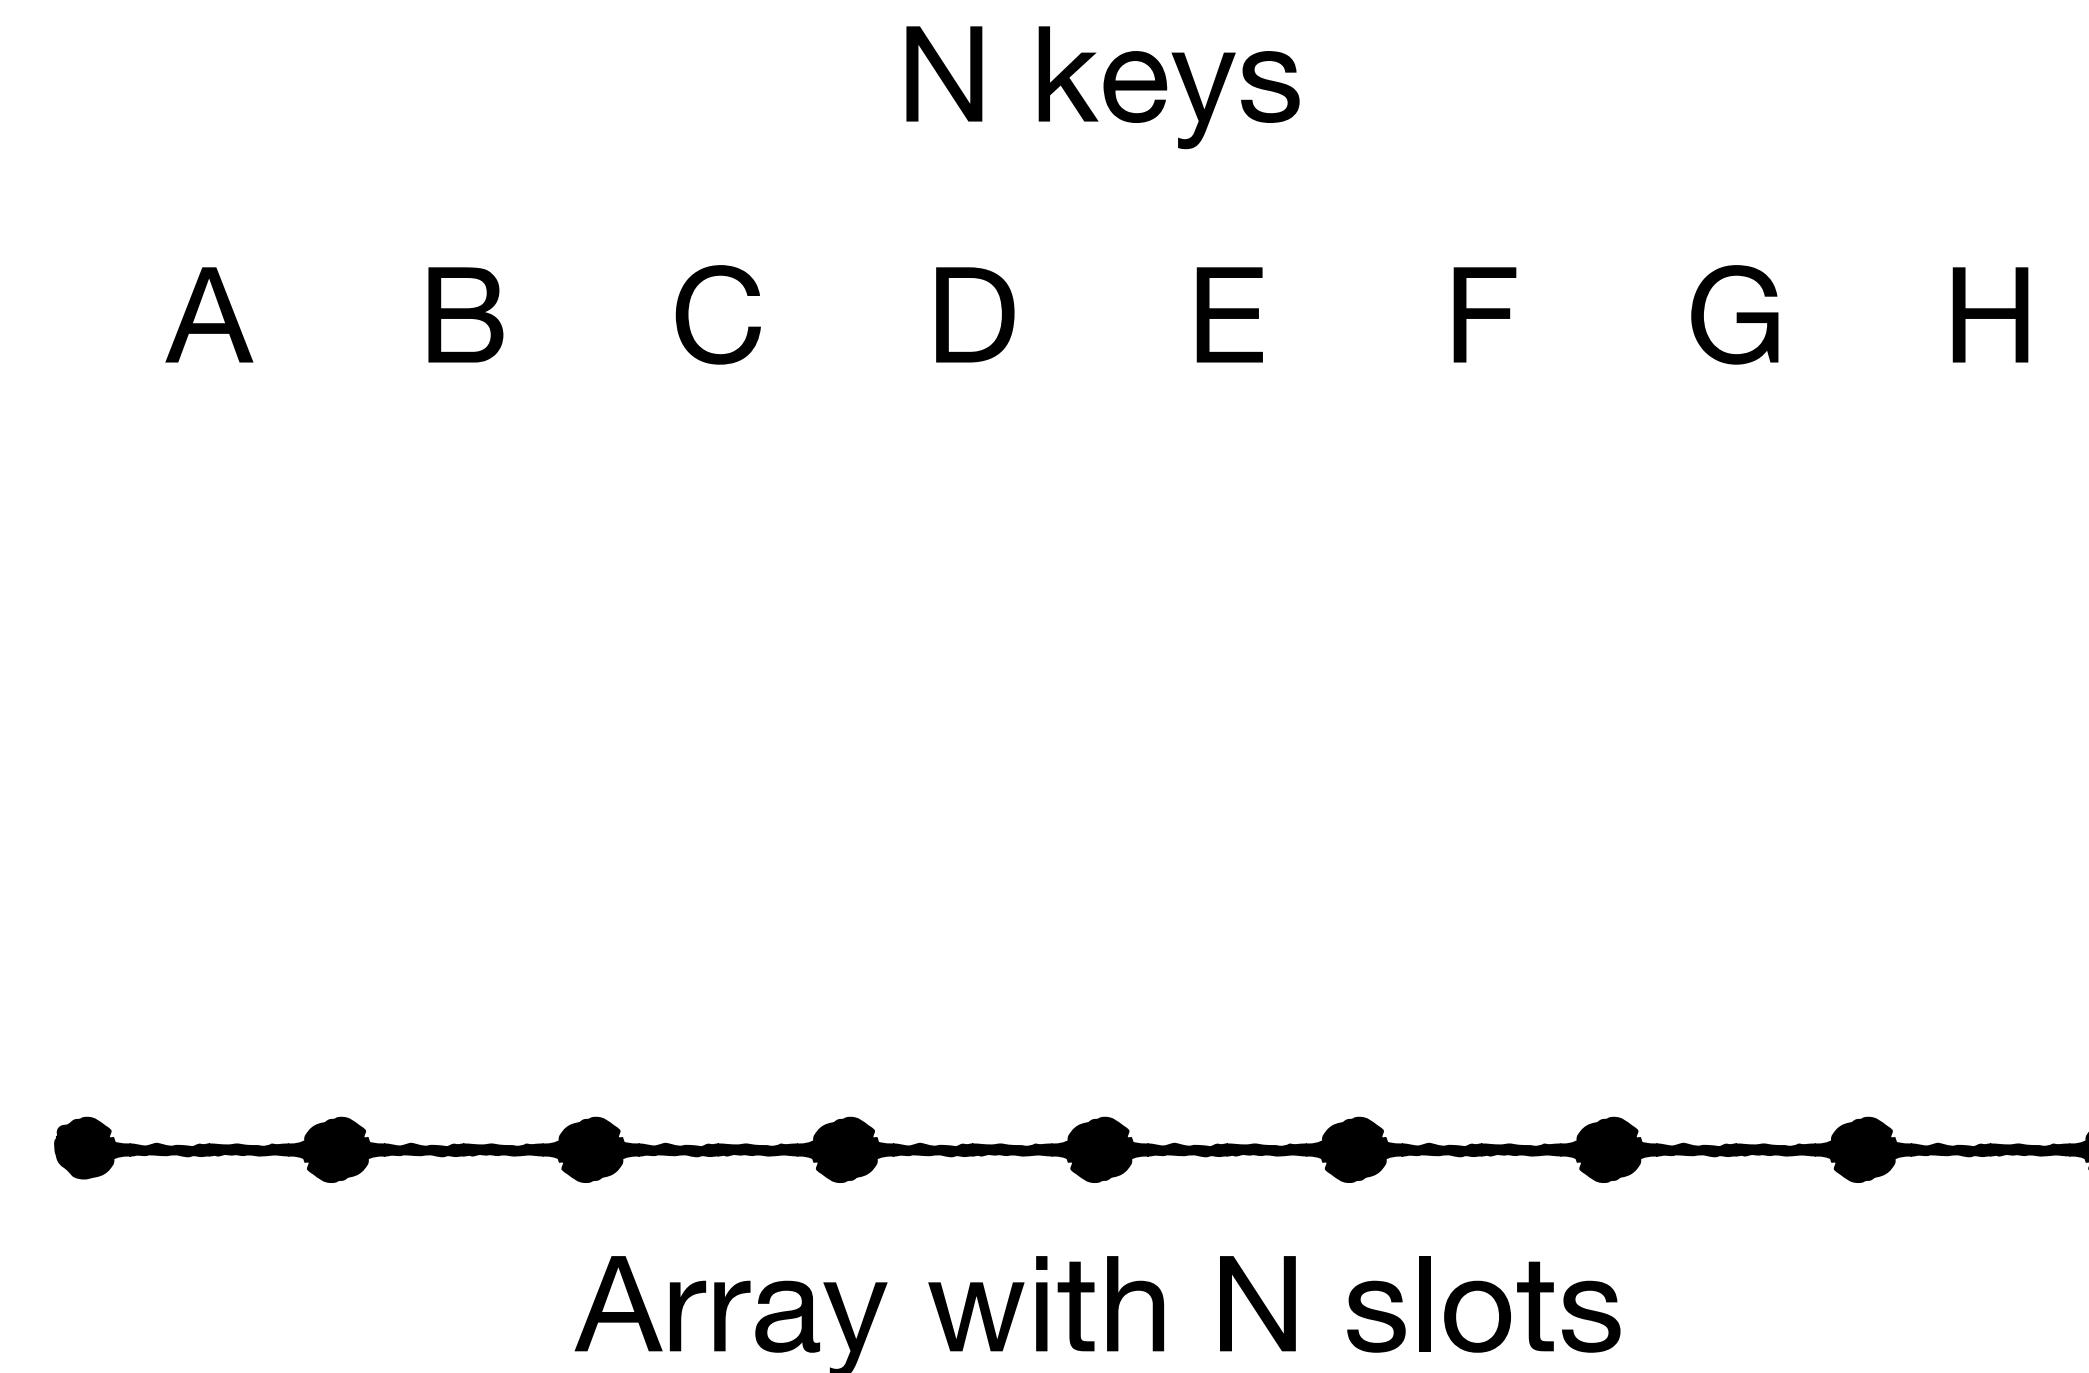


Array with N slots

Minimal Perfect Hashing

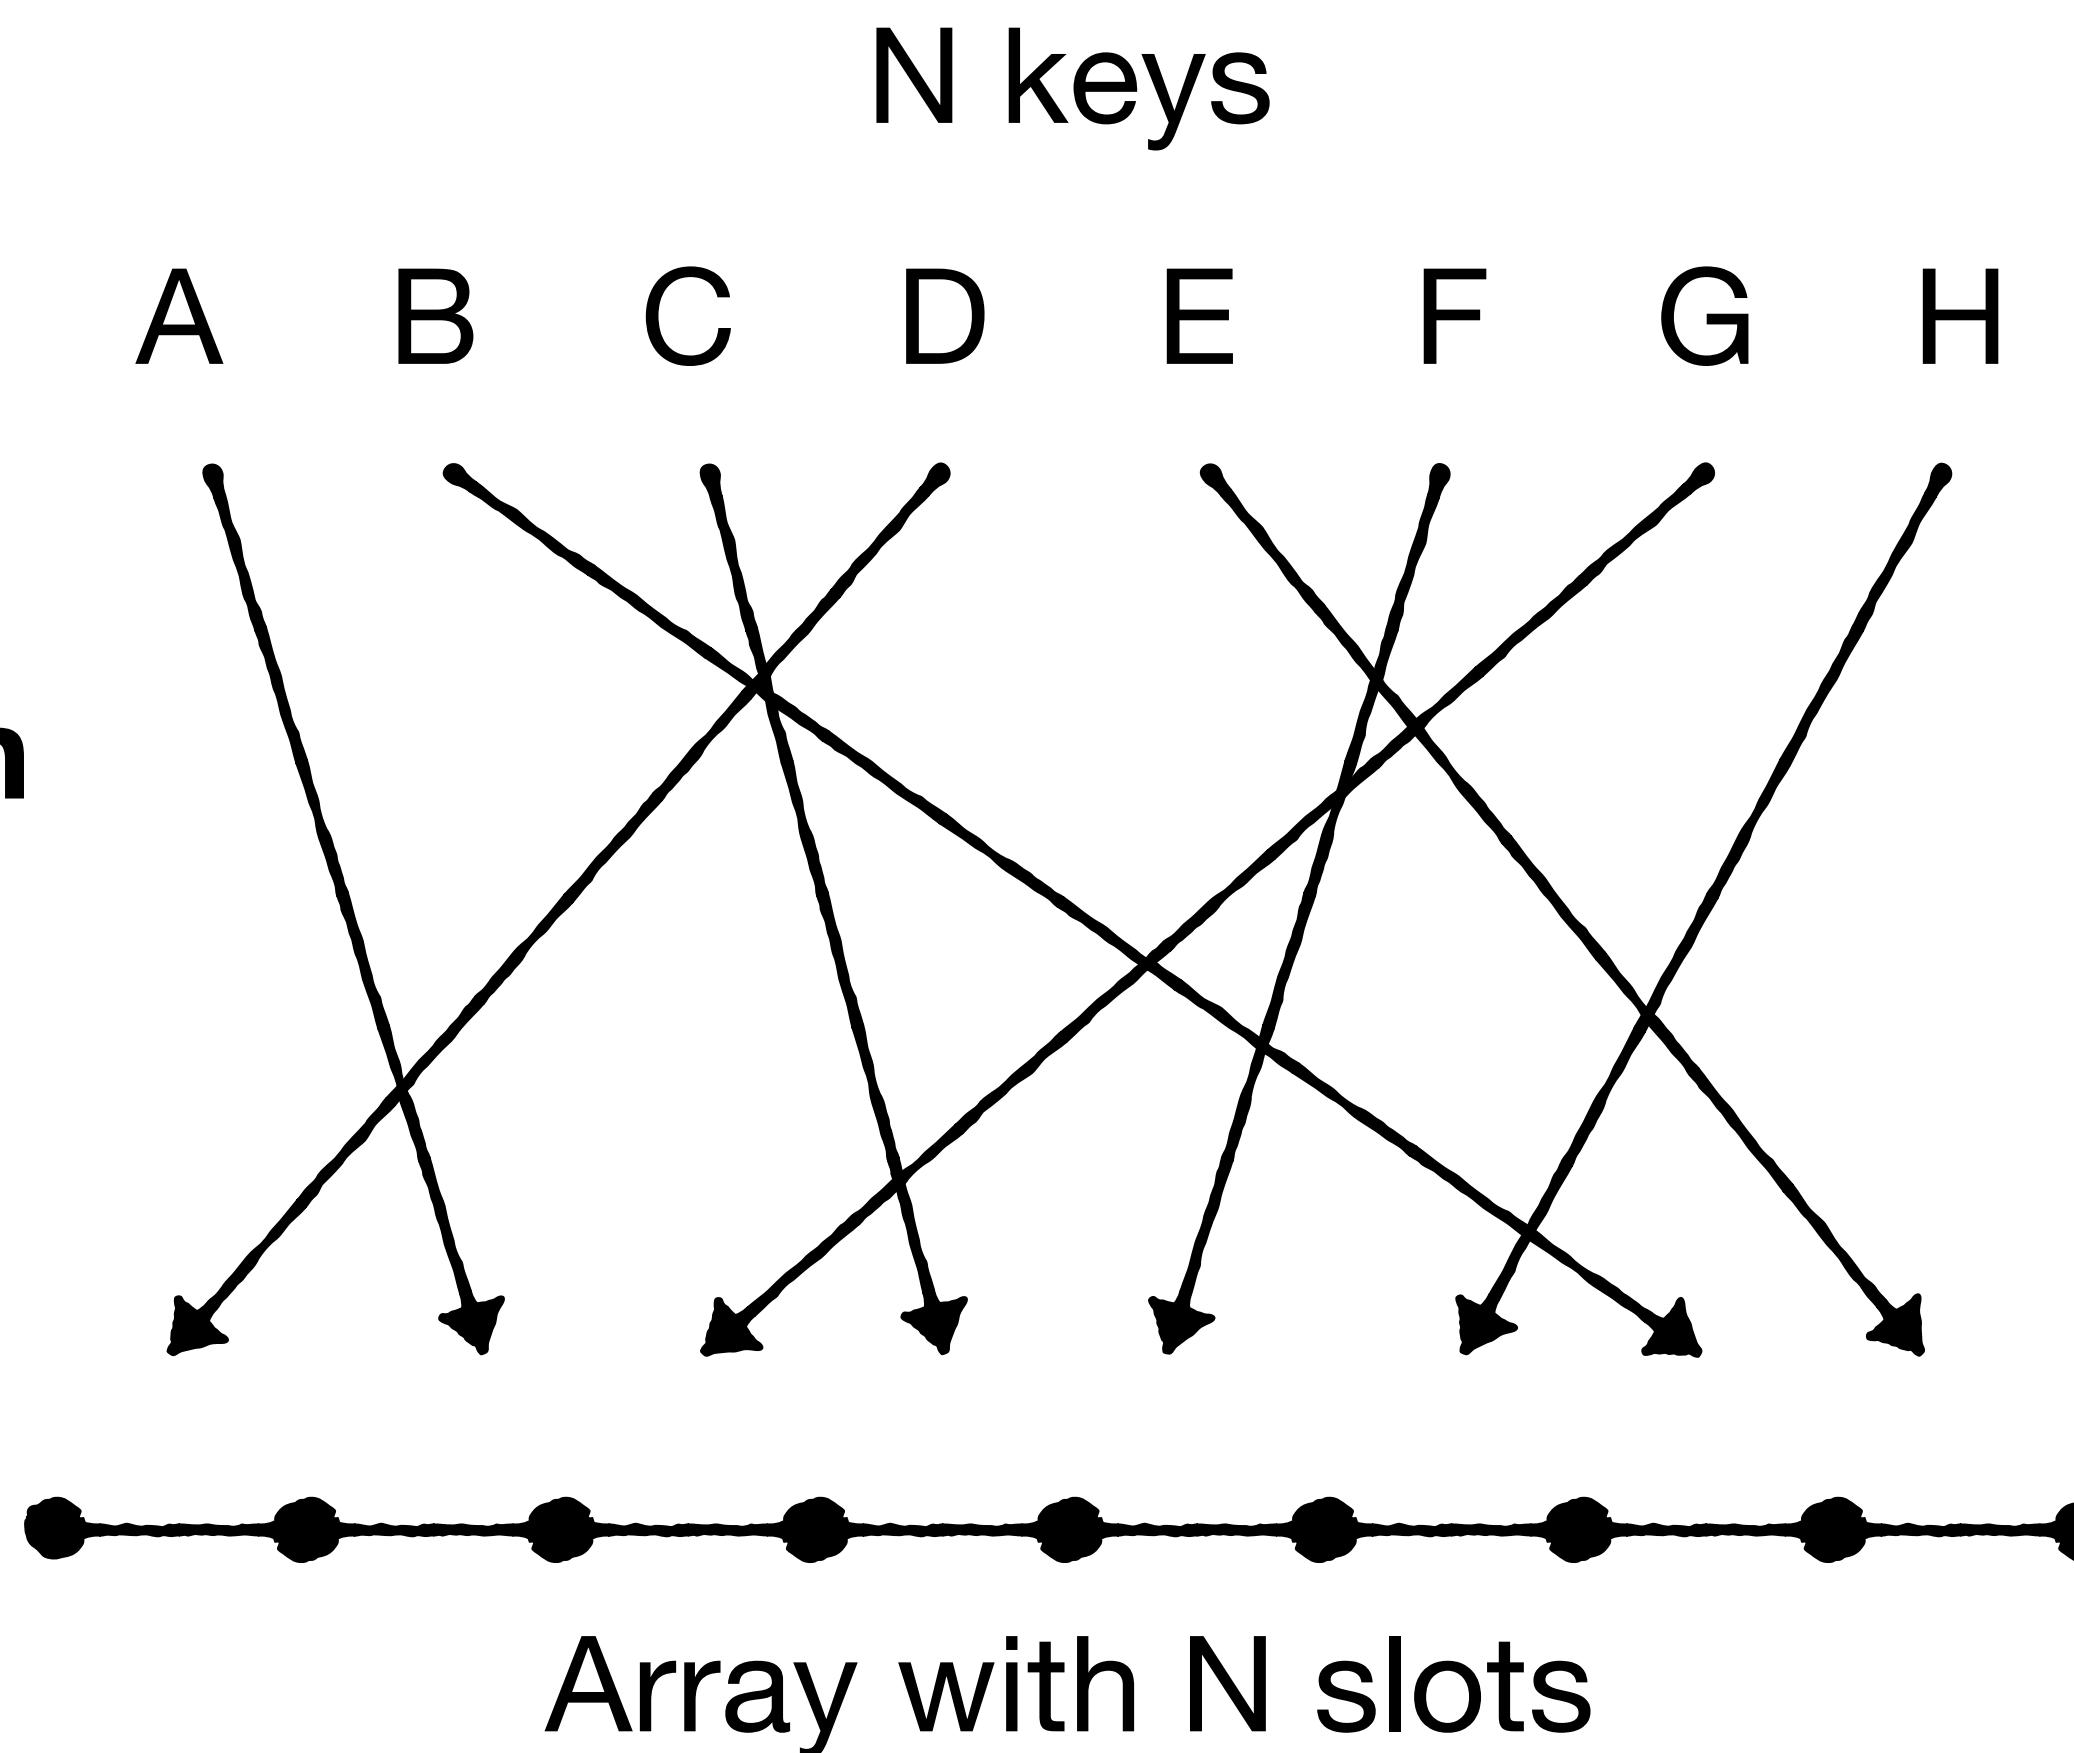


Minimal Perfect Hashing



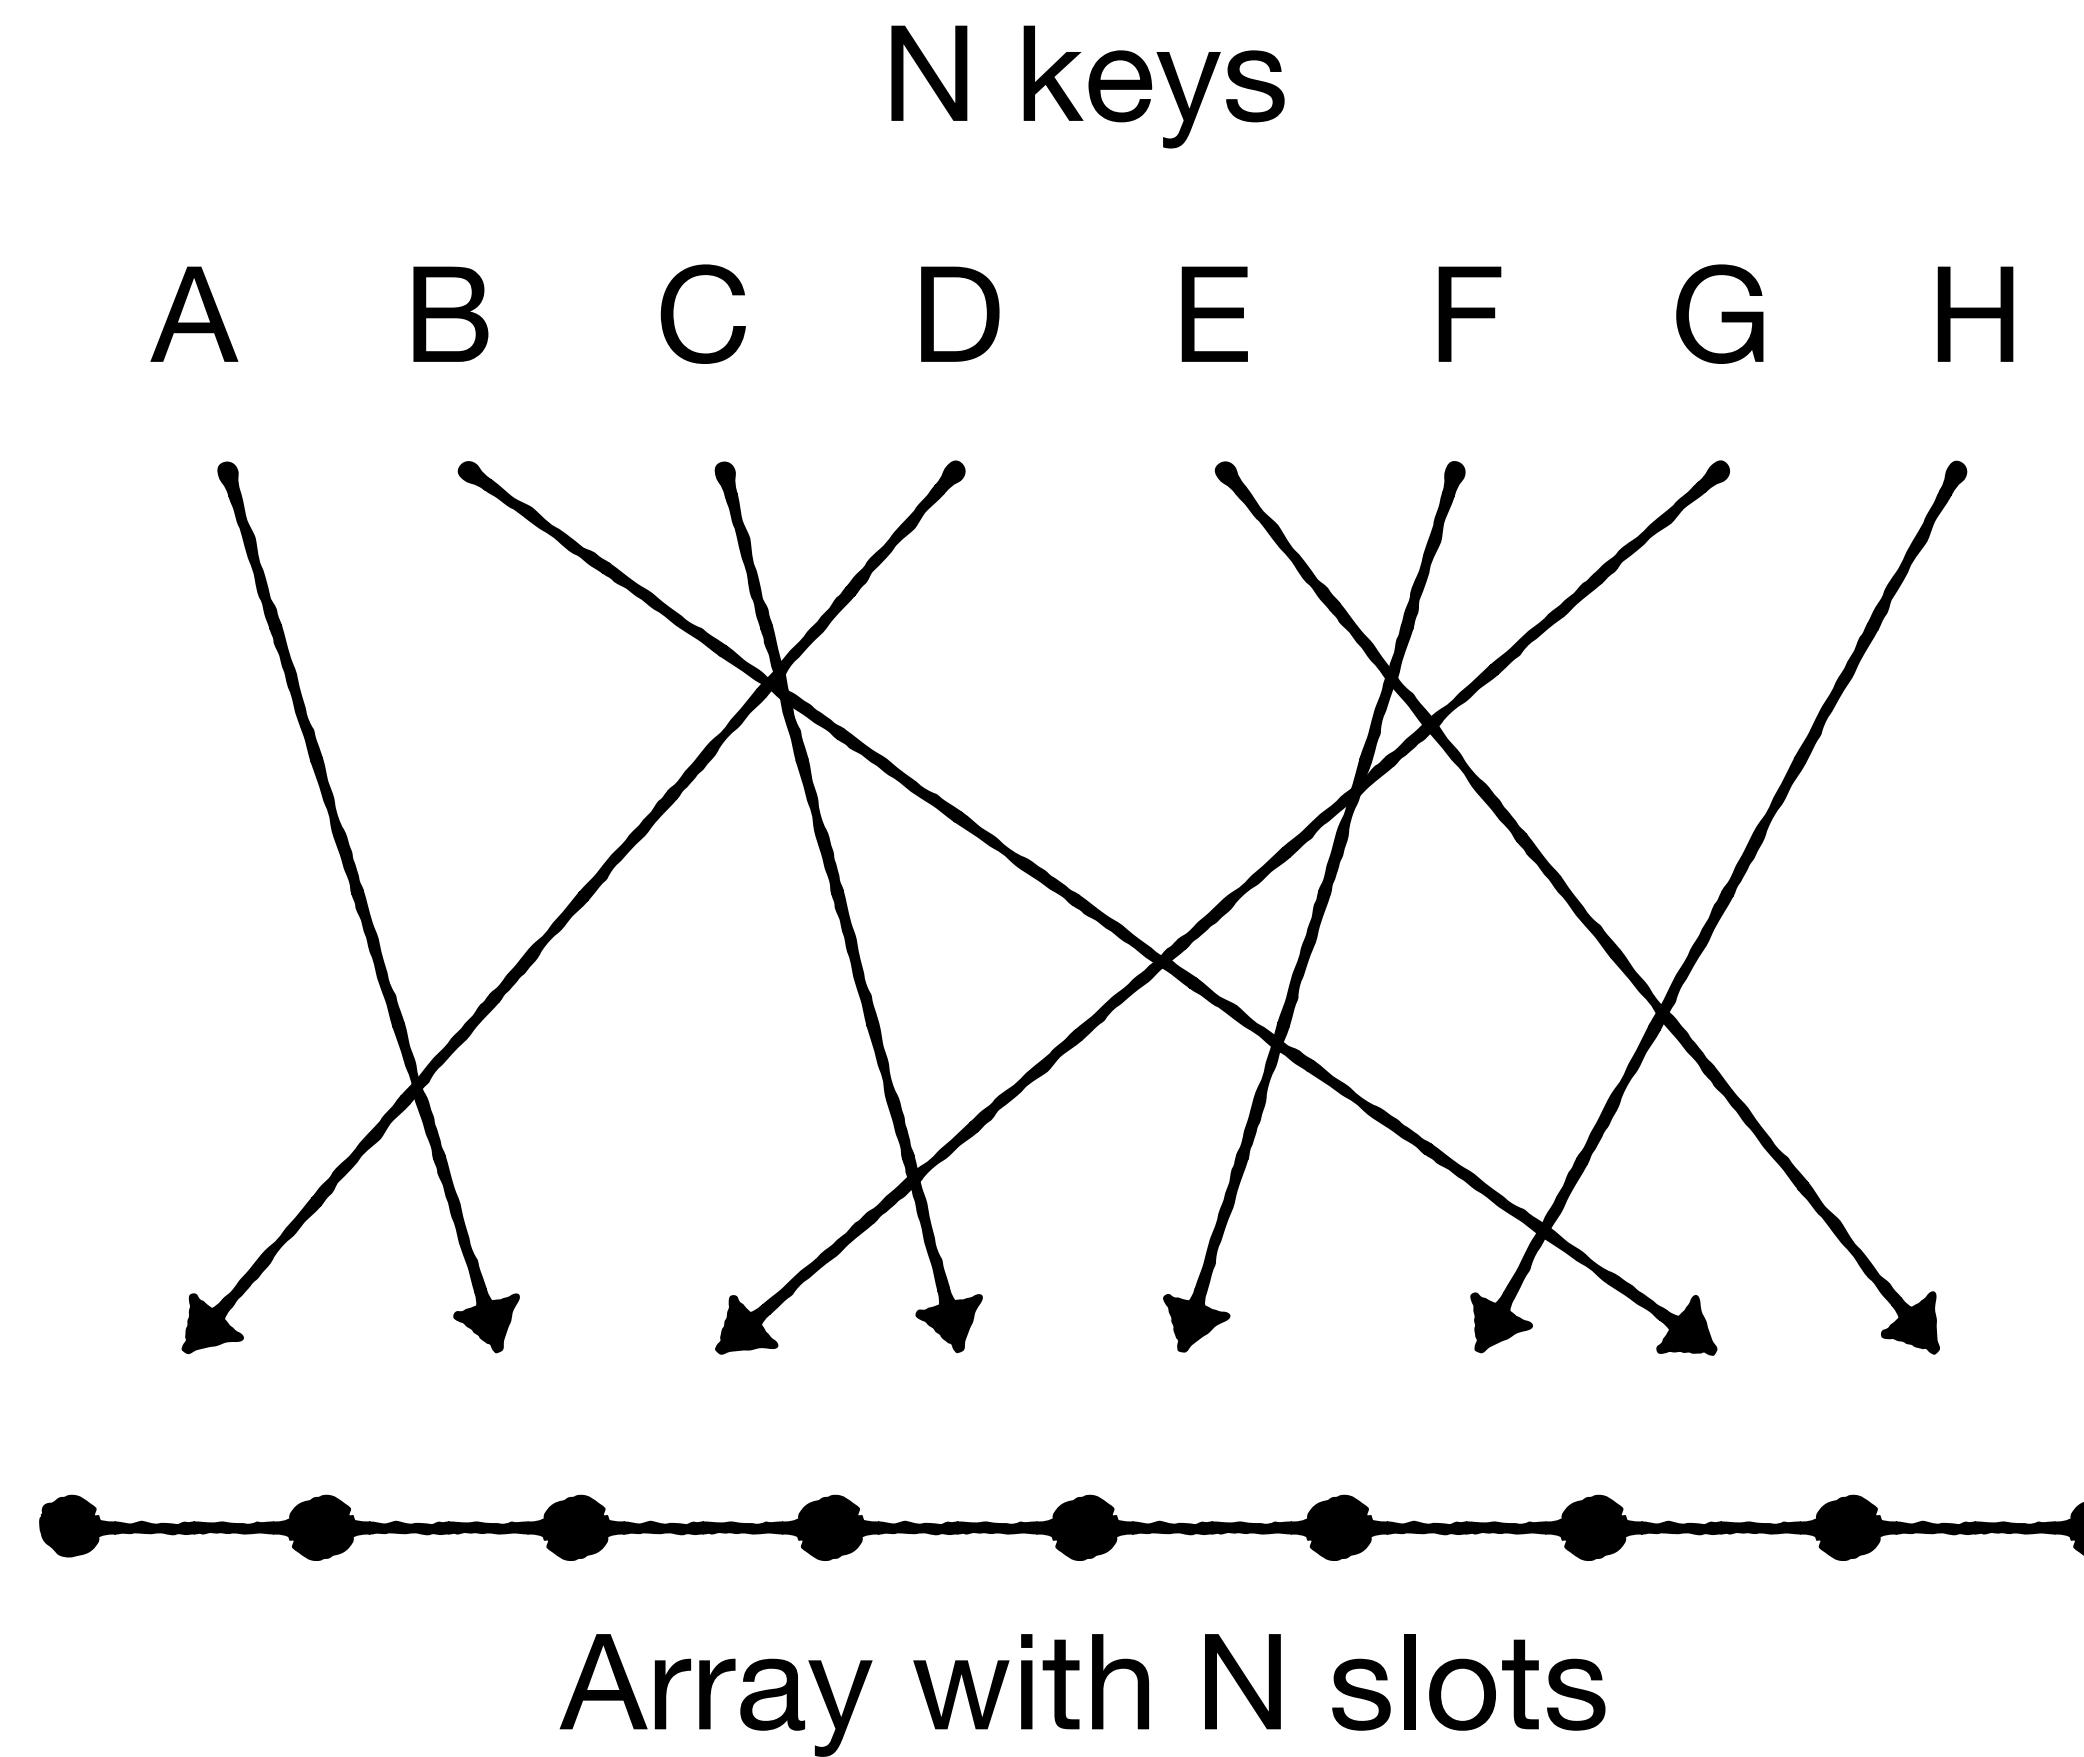
100% load factor! No extra capacity as with normal hash tables

**Goal: Establish bijection
(one-to-one mapping)**

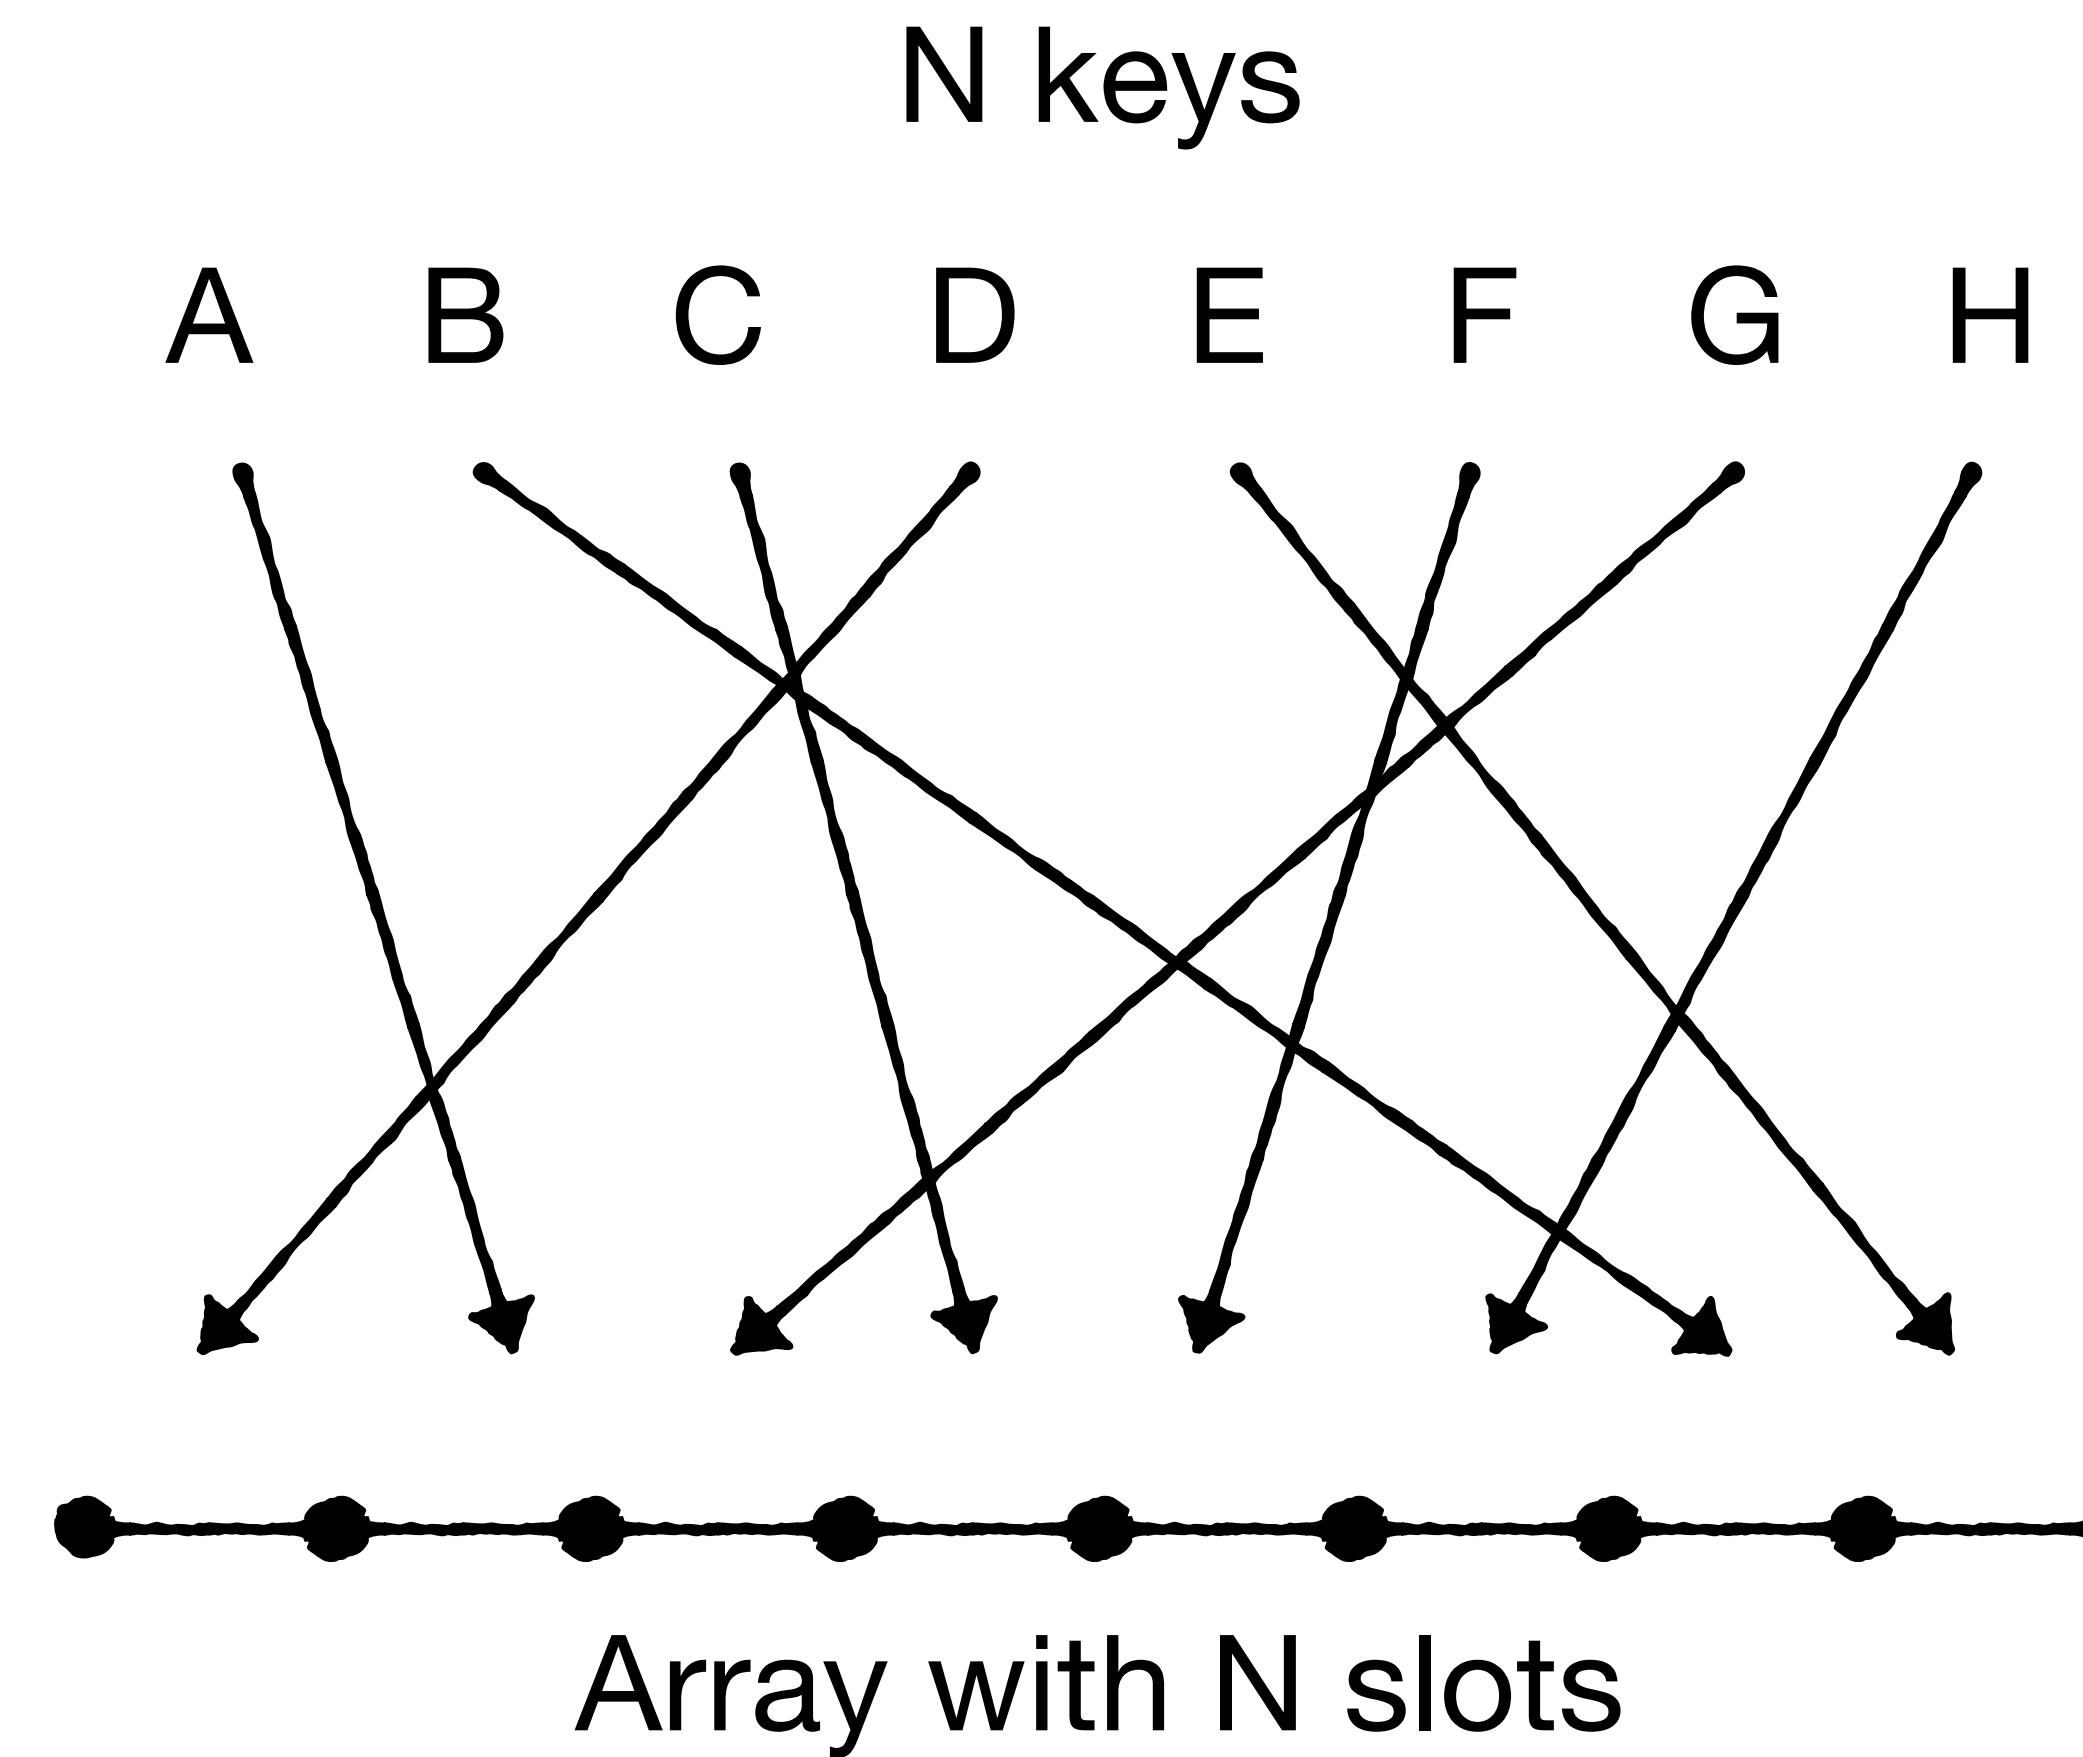


Goal: Establish bijection
(one-to-one mapping)

Collision-free



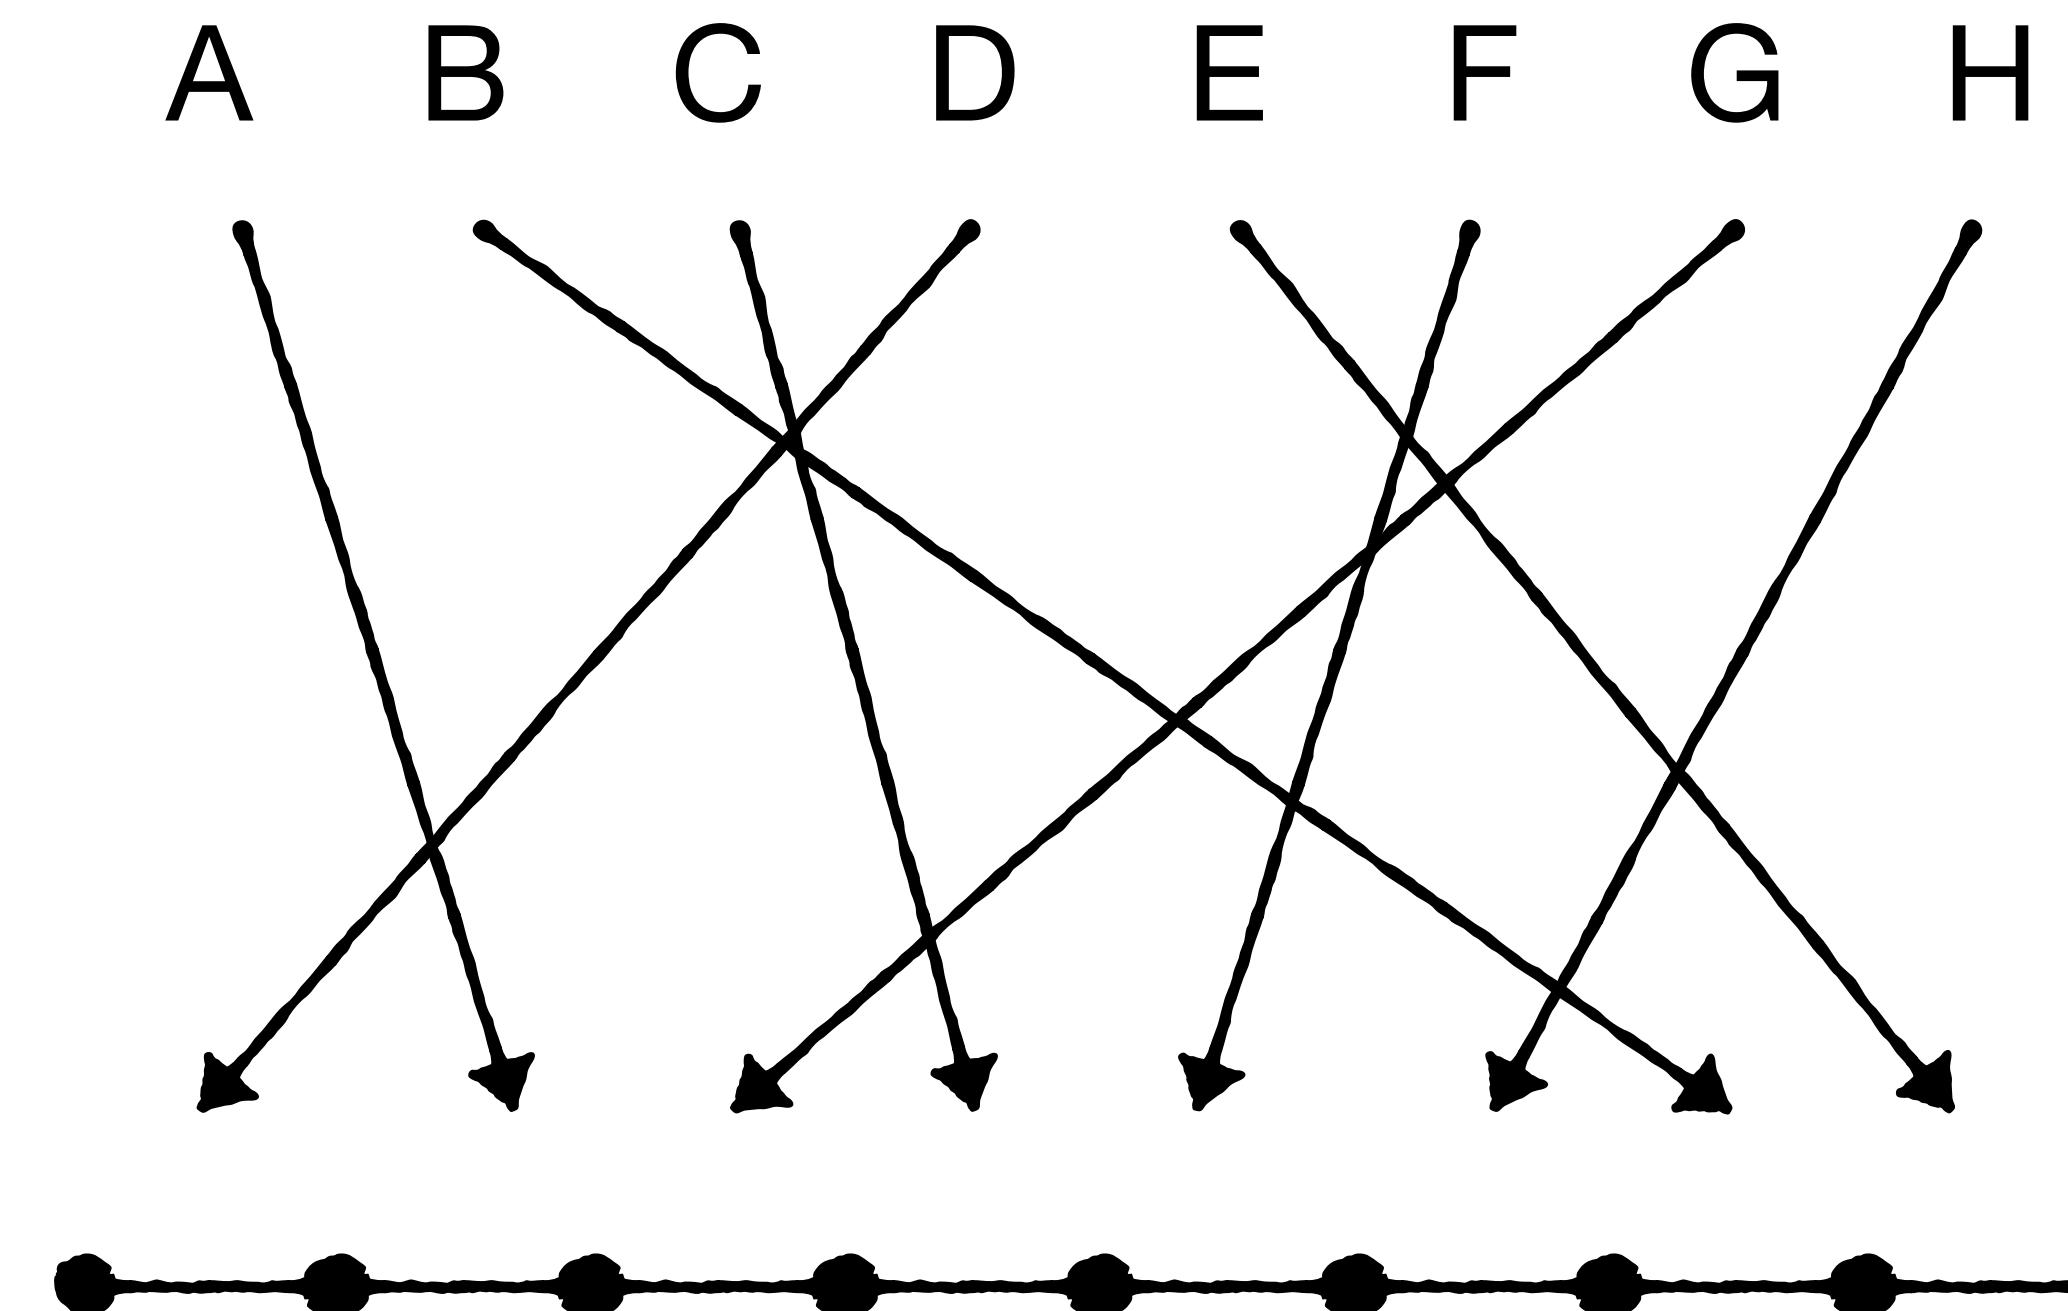
What's the probability of general hash function creating bijection?



What's the probability of general hash function creating bijection?

possible bijections (permutations)?

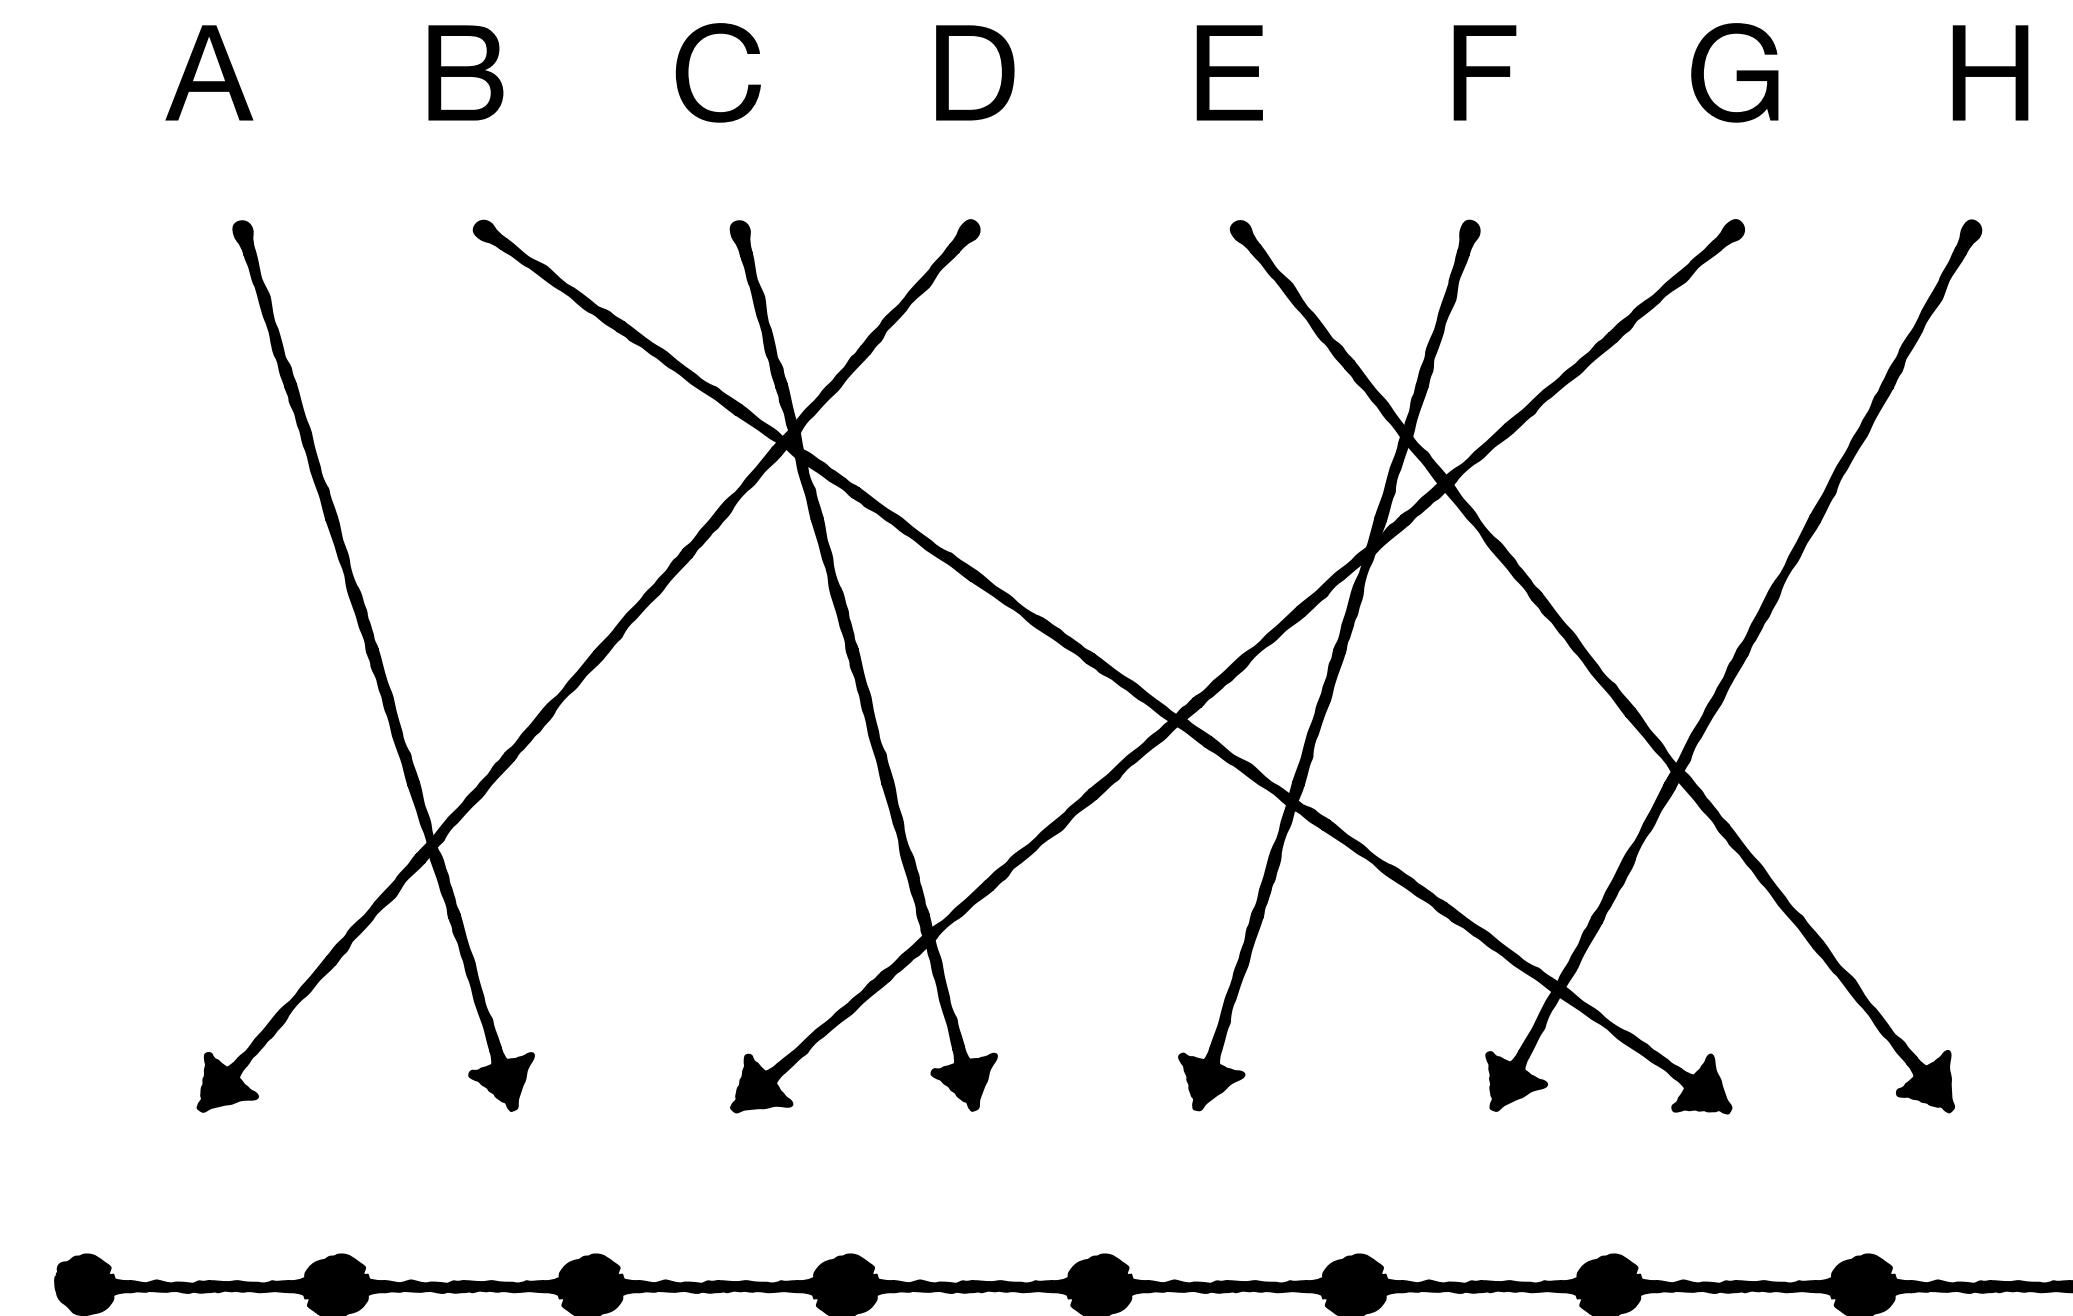
possible assignments?



What's the probability of general hash function creating bijection?

possible bijections (permutations)? $N!$

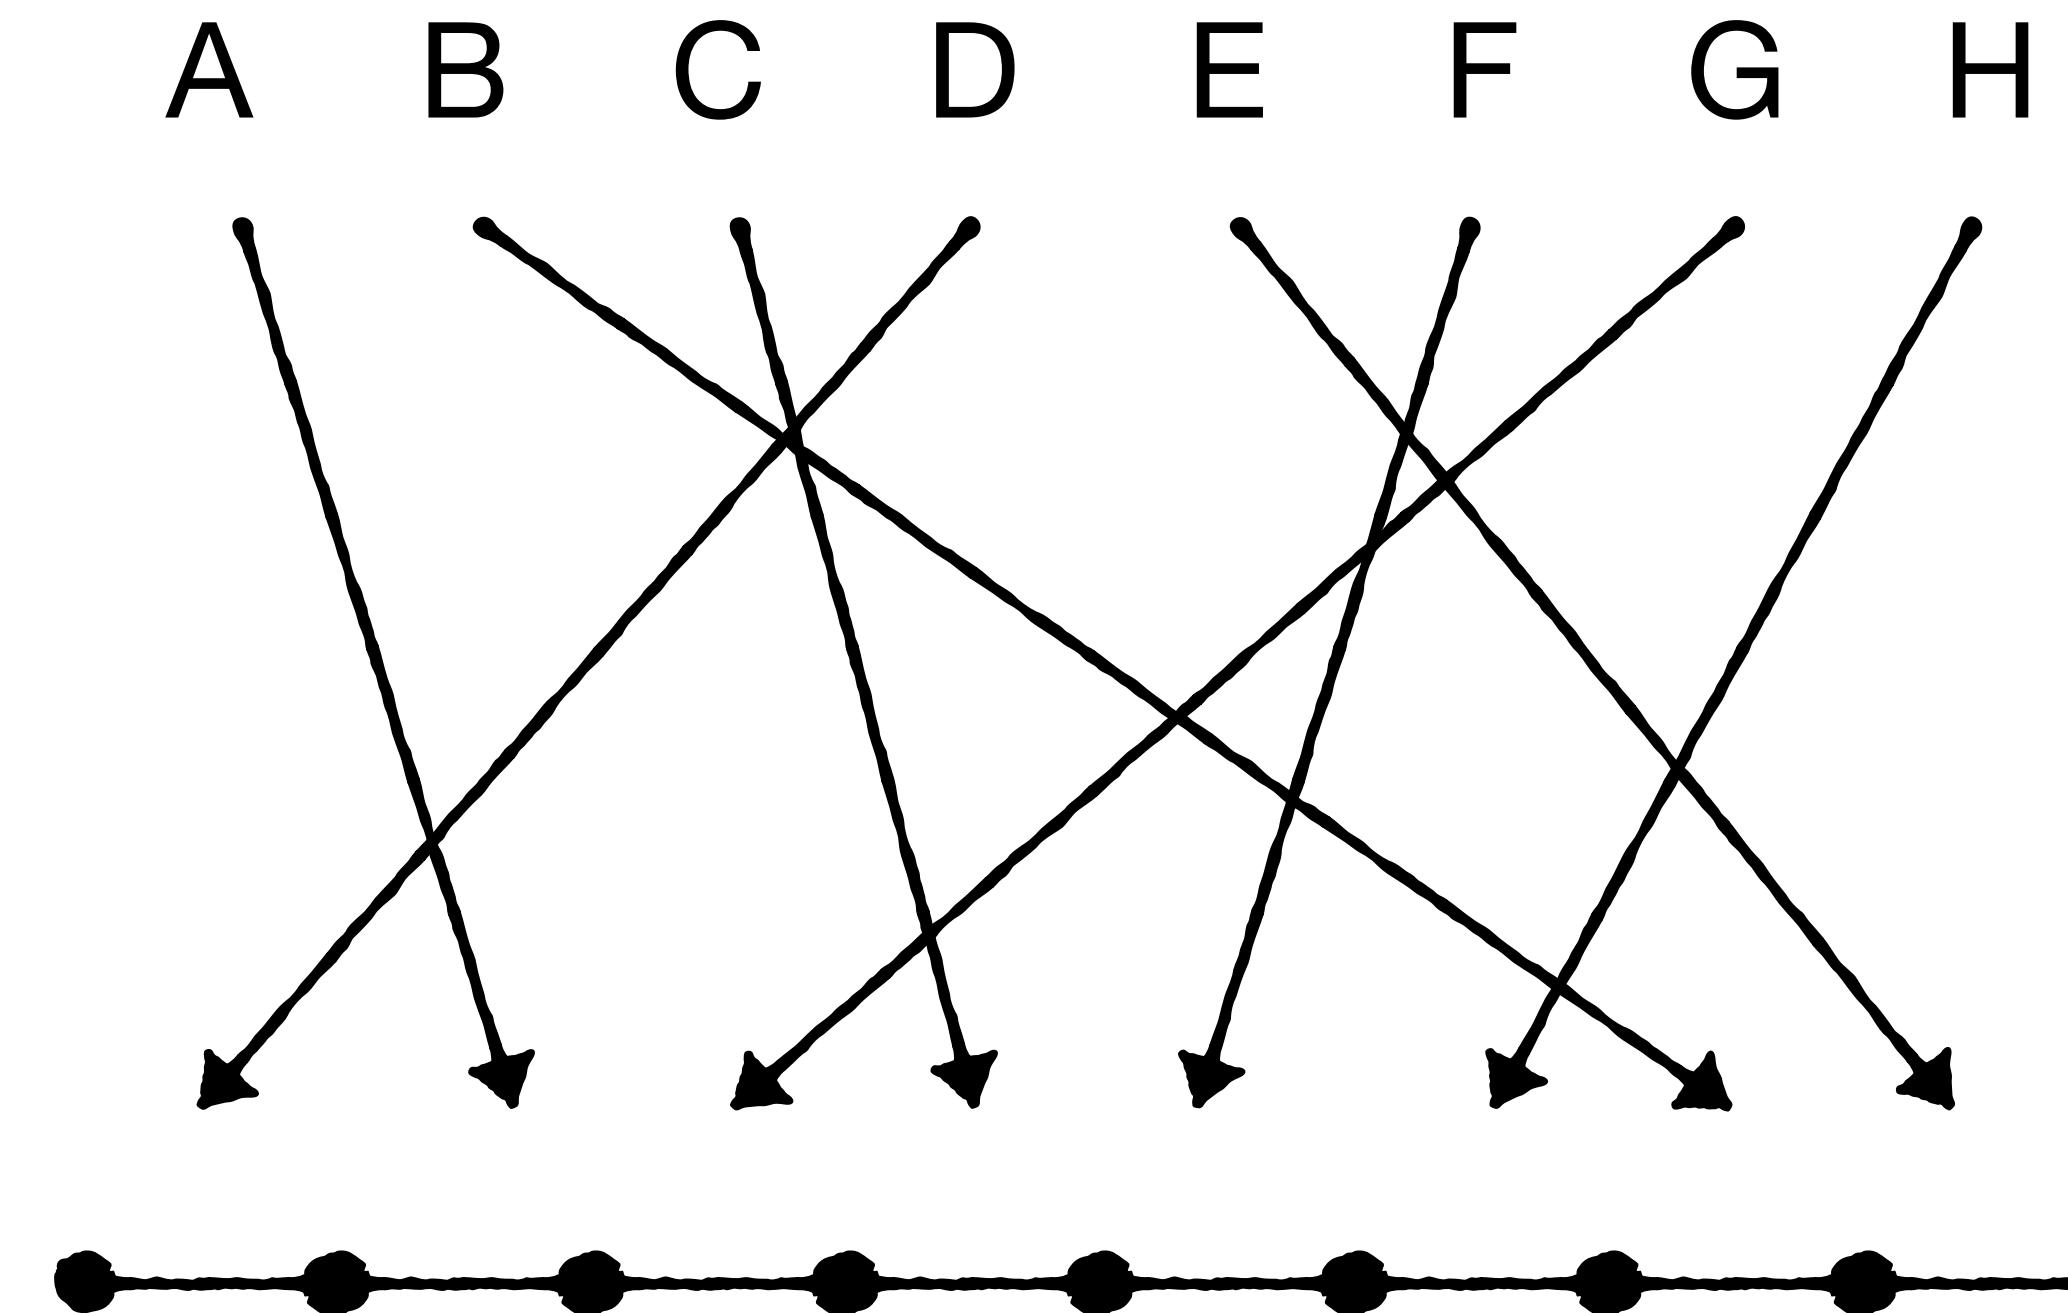
possible assignments?



What's the probability of general hash function creating bijection?

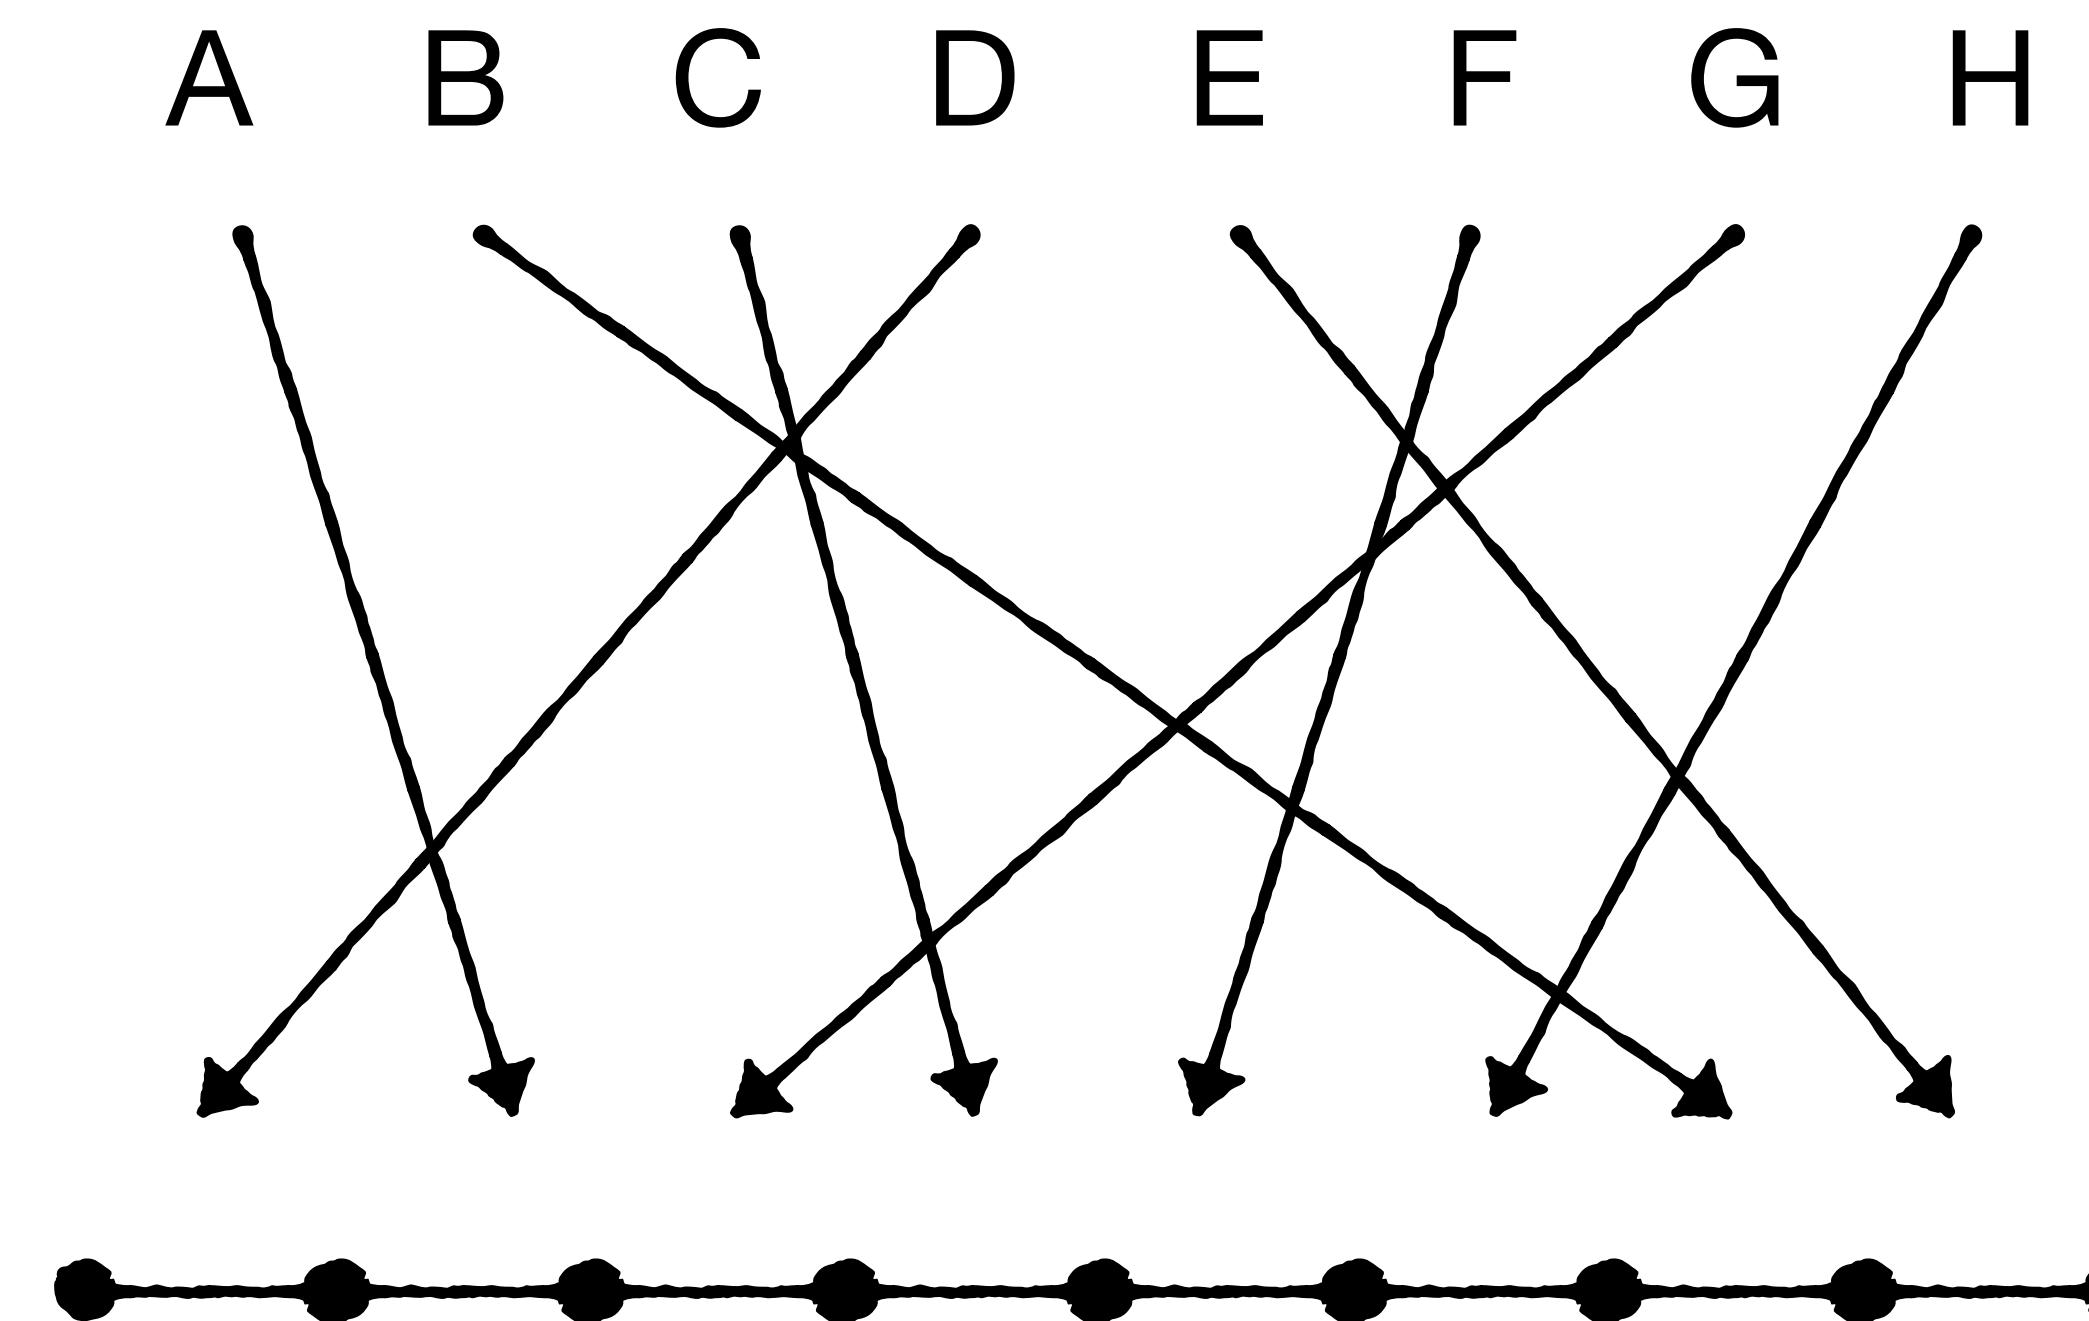
possible bijections (permutations)? $N!$

possible assignments? N^N



What's the probability of general hash function creating bijection?

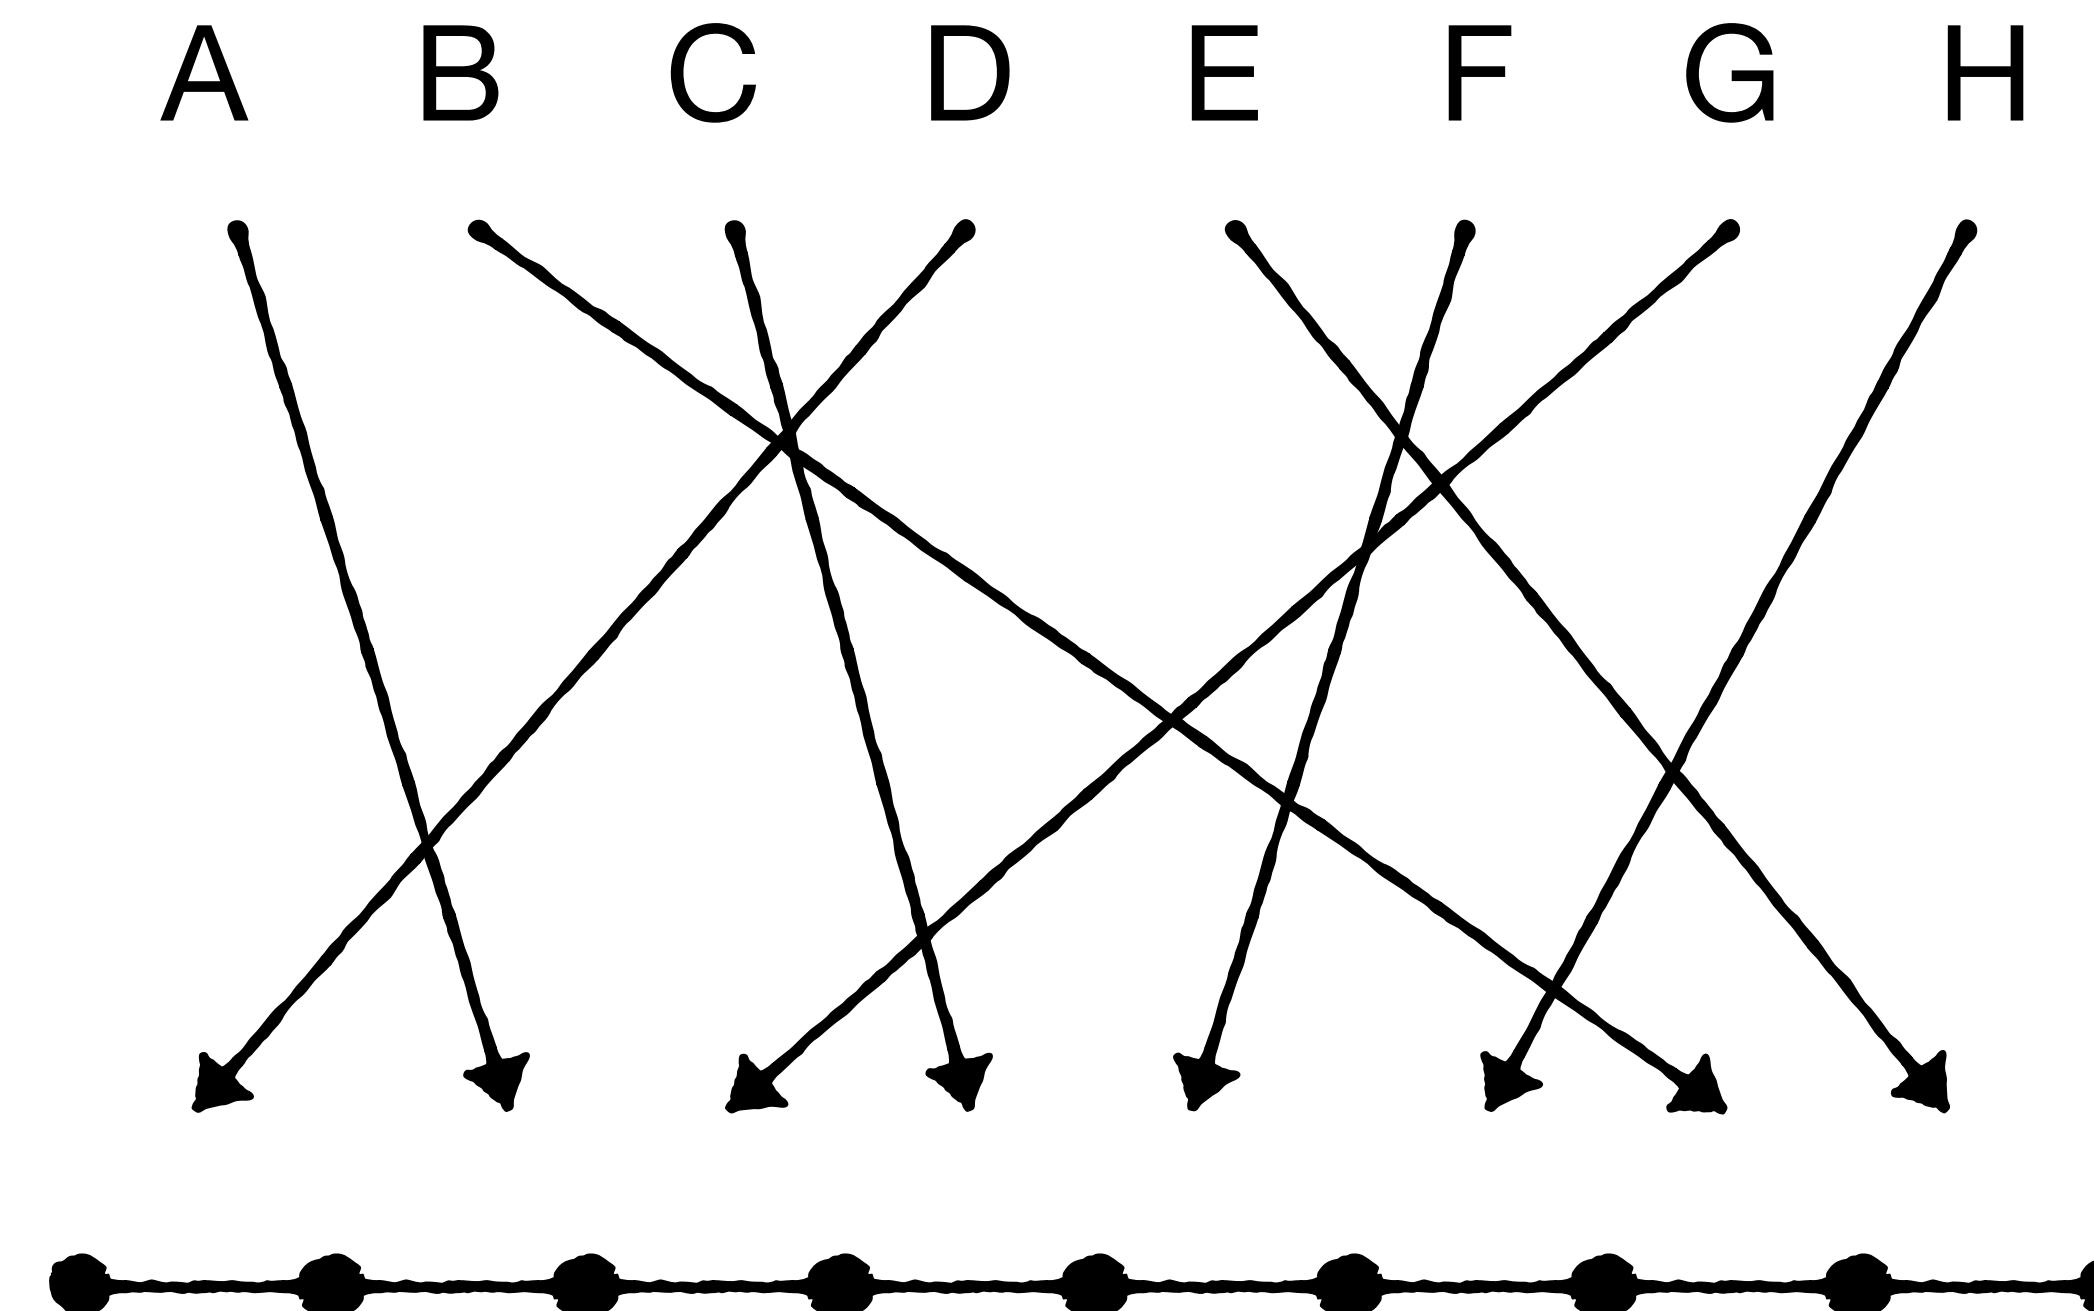
$$\frac{N!}{N^N}$$



What's the probability of general hash function creating bijection?

$$\frac{N!}{N^N} \approx \sqrt{2 \pi N} \cdot e^{-N}$$

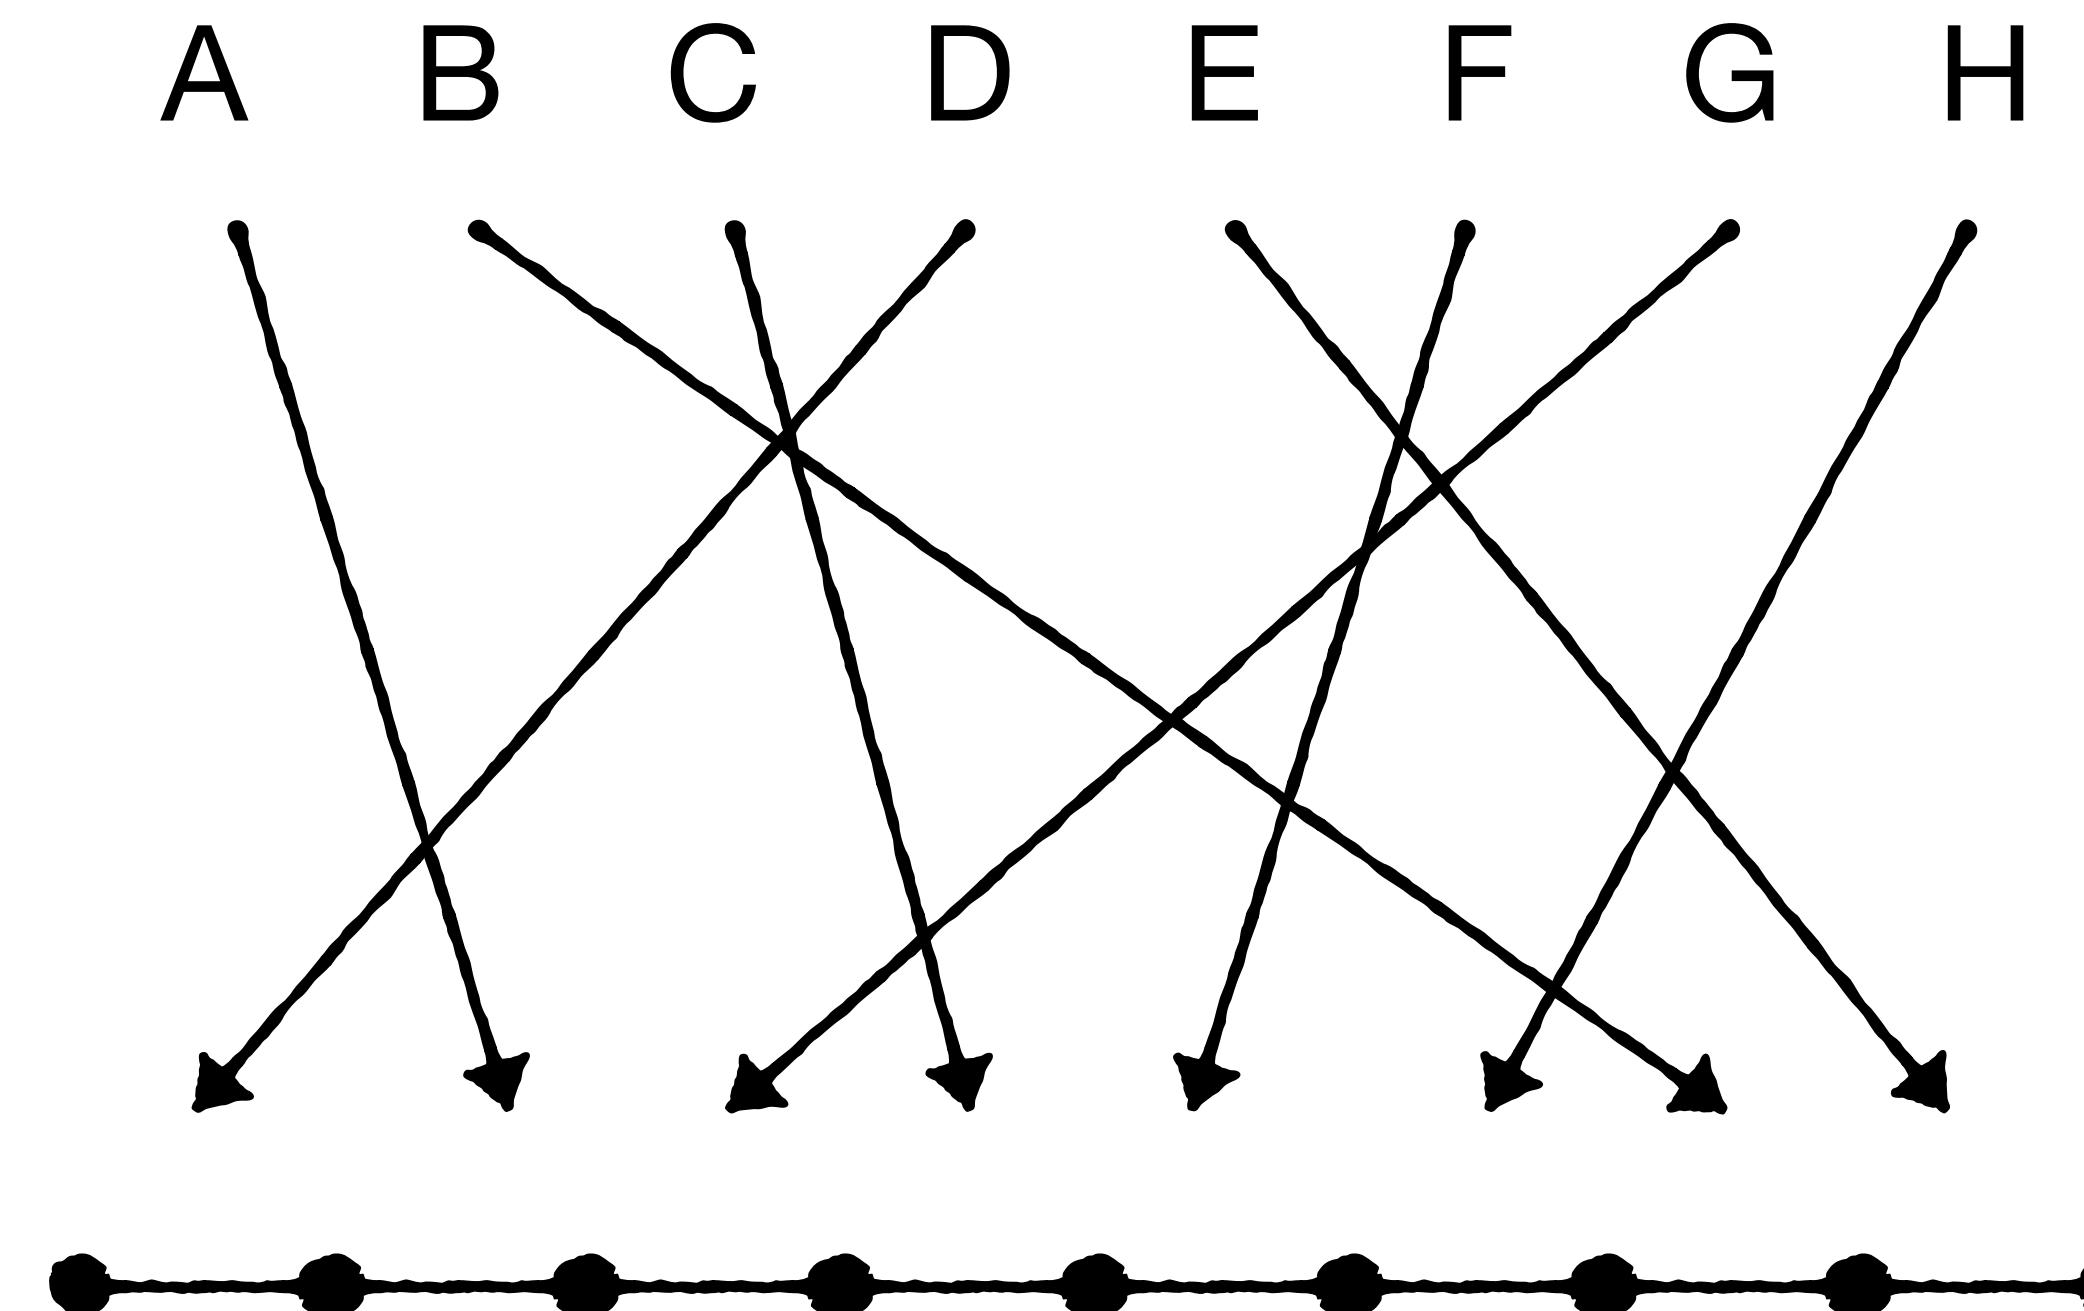
By **Stirling's approximation**

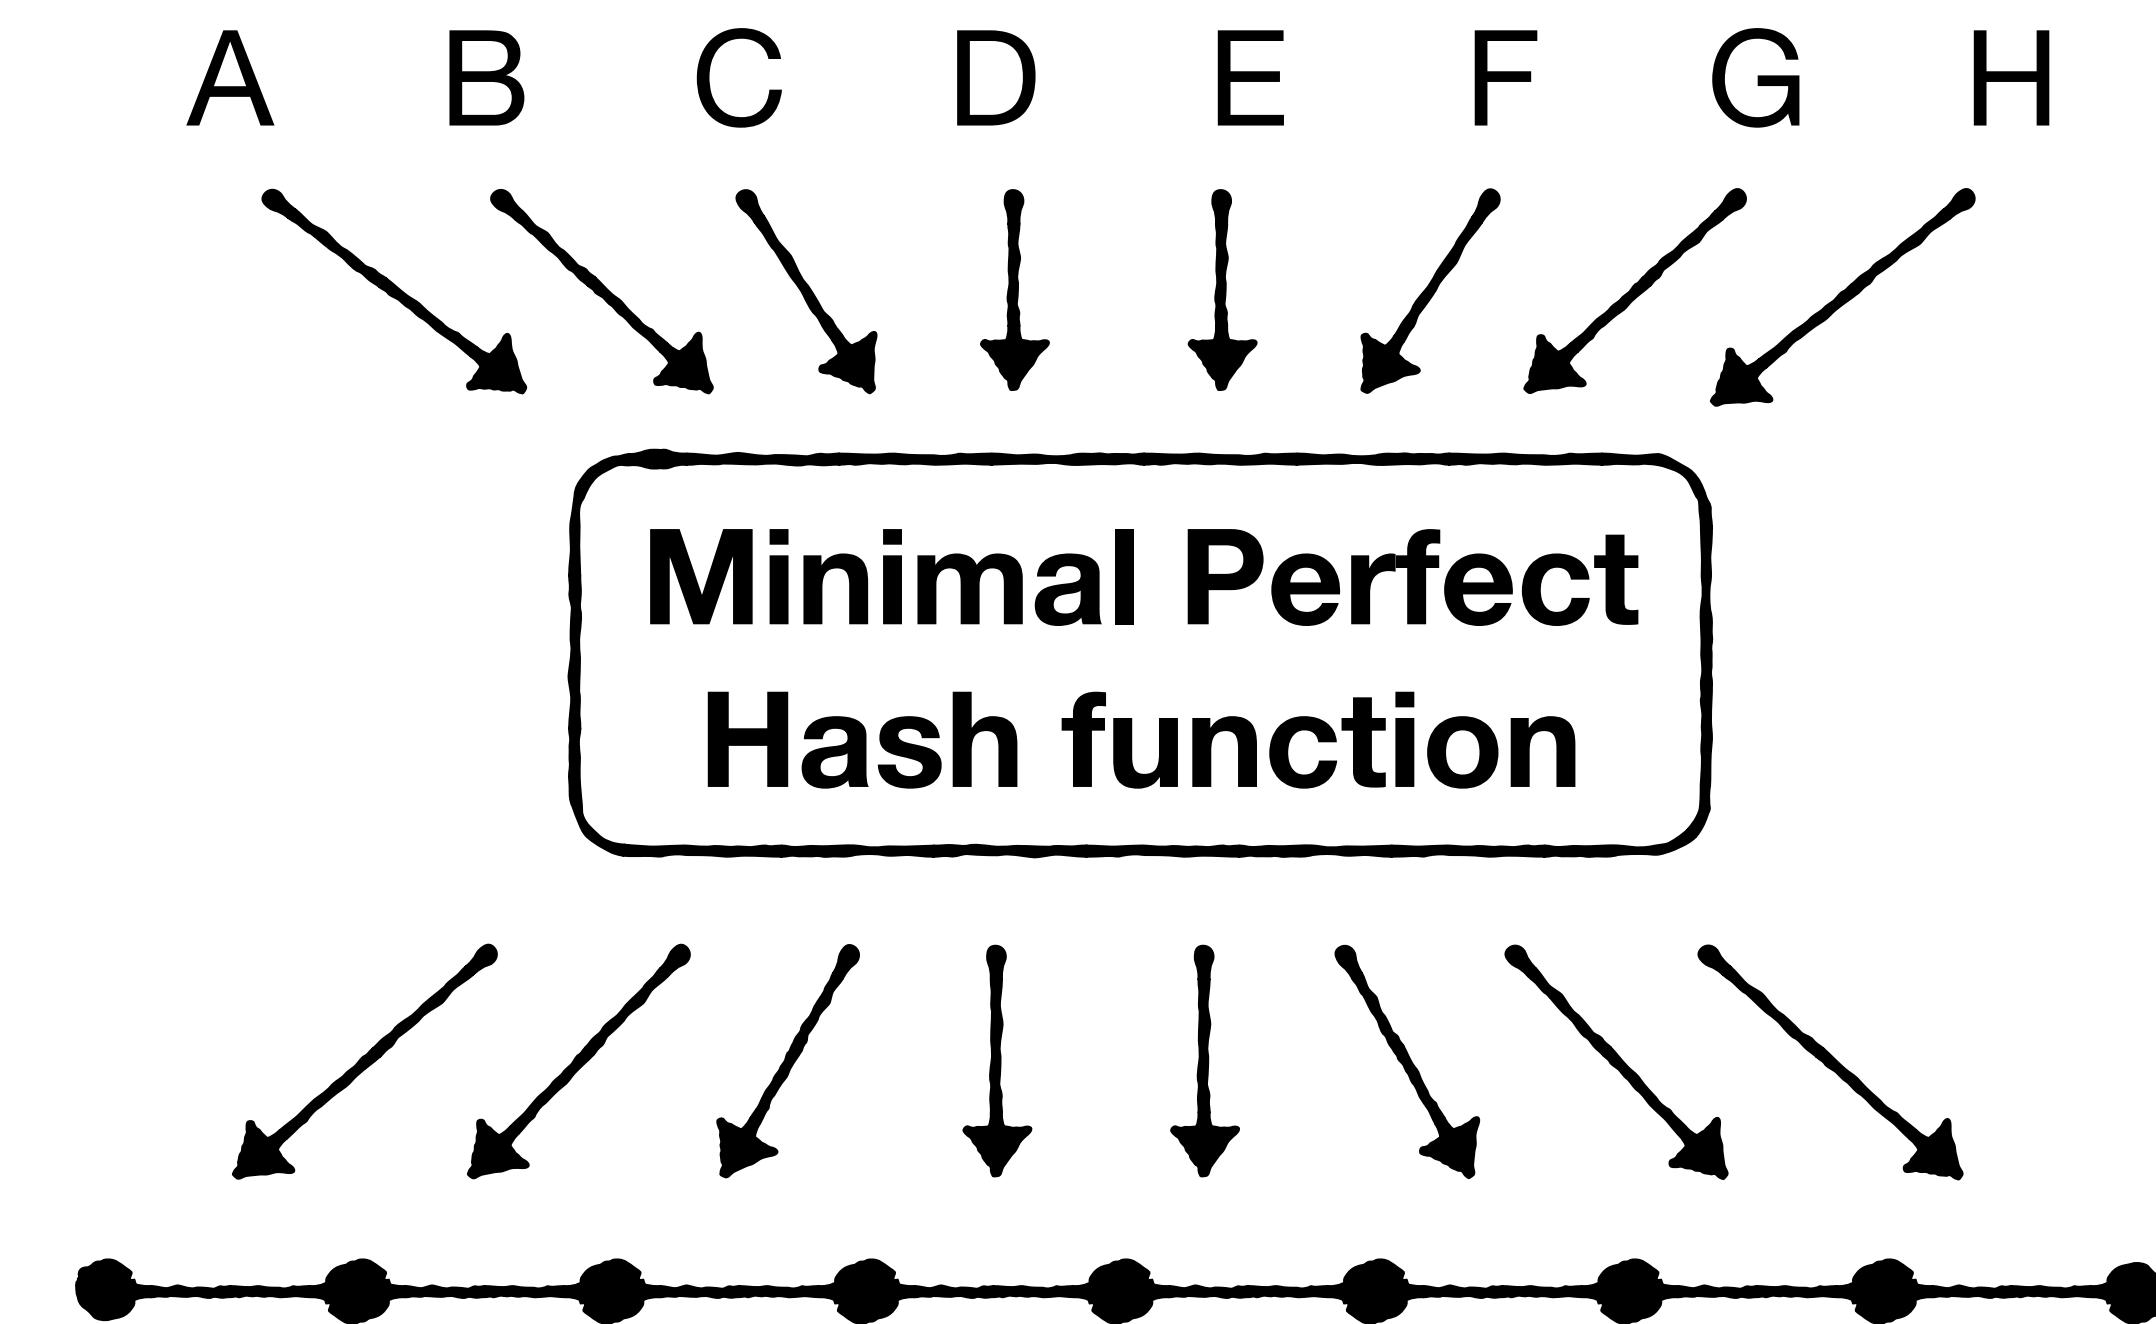


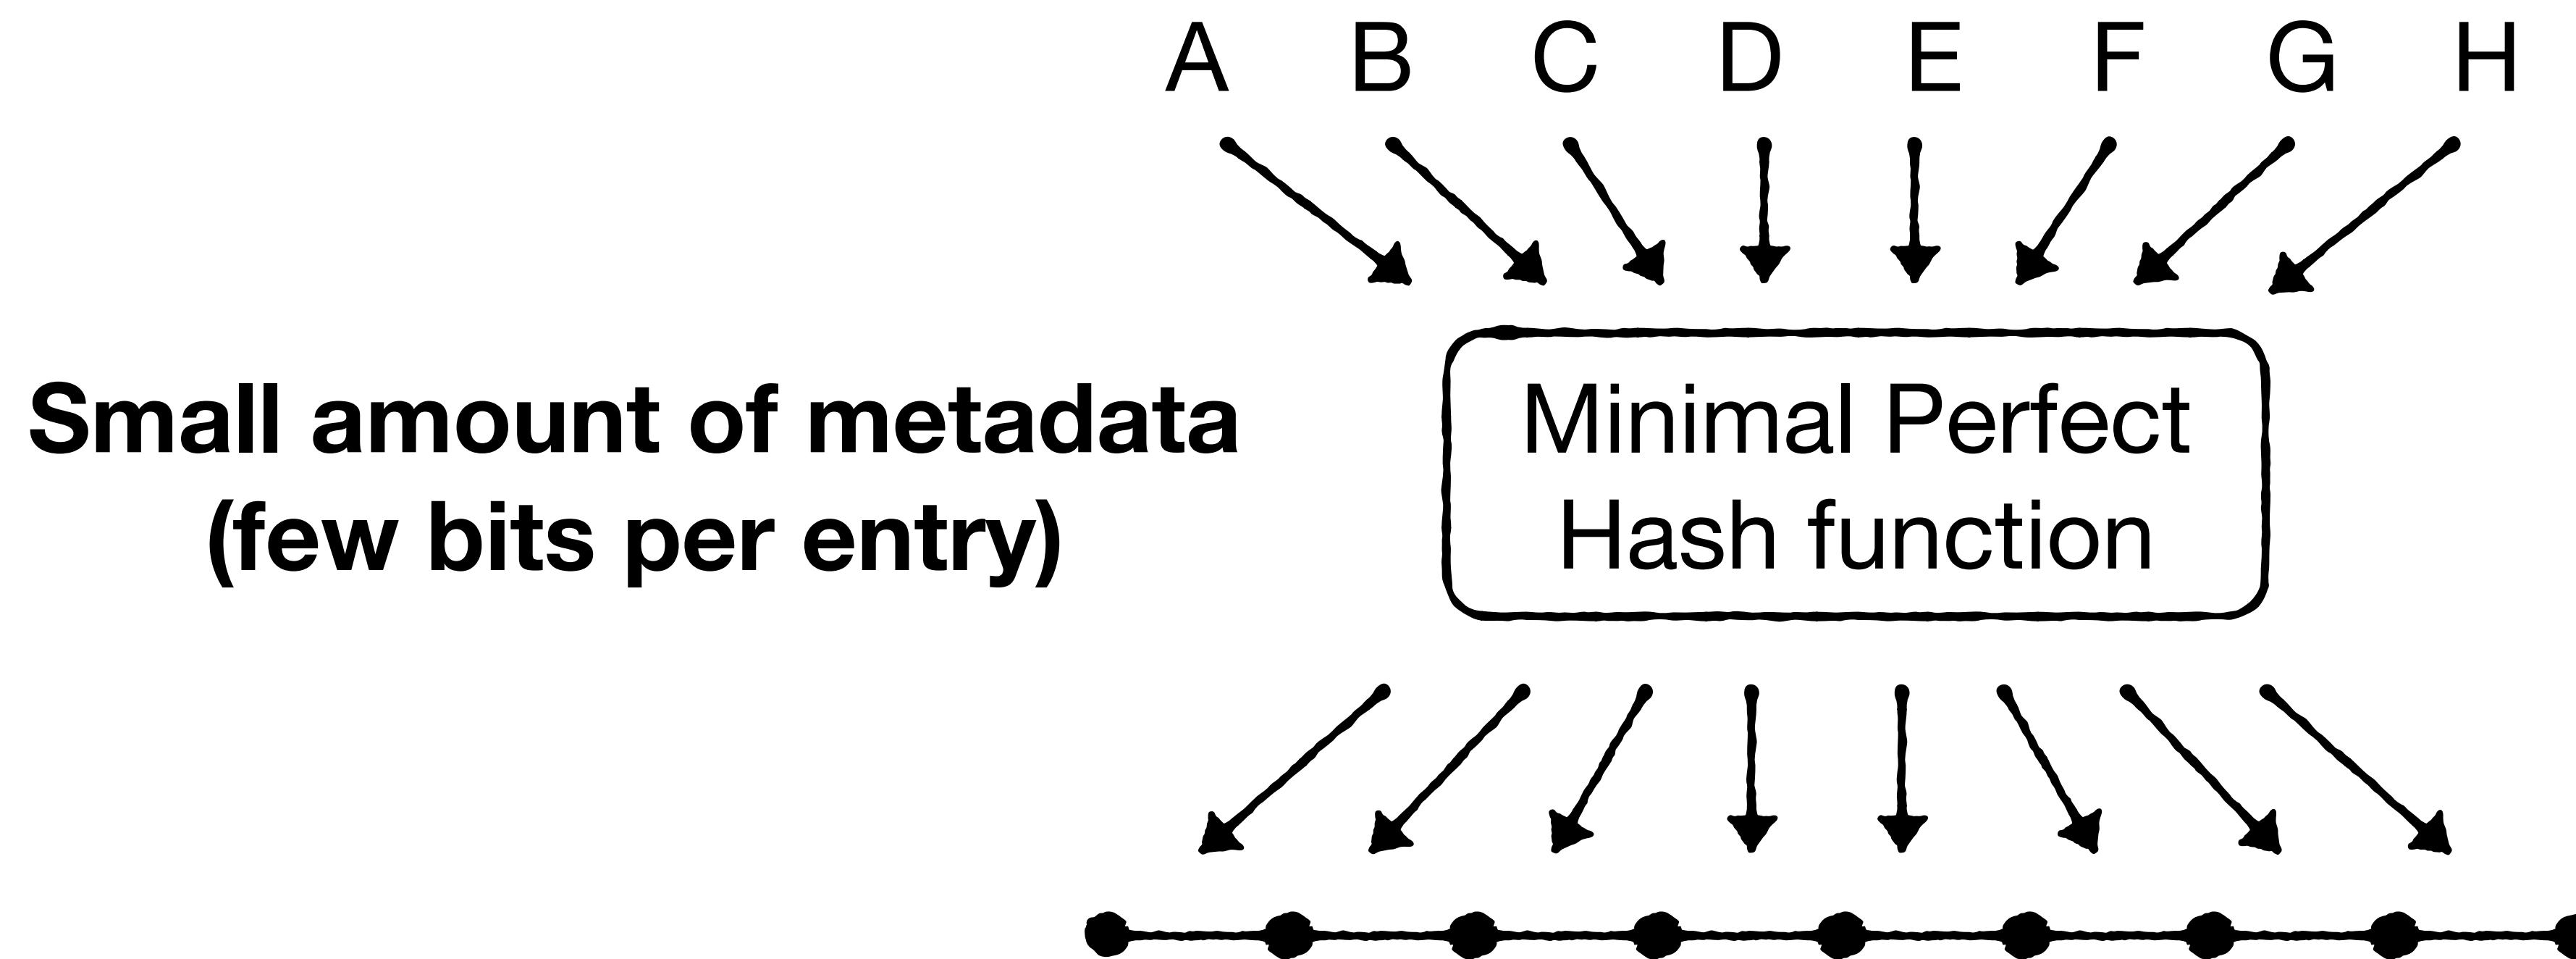
What's the probability of general hash function creating bijection?

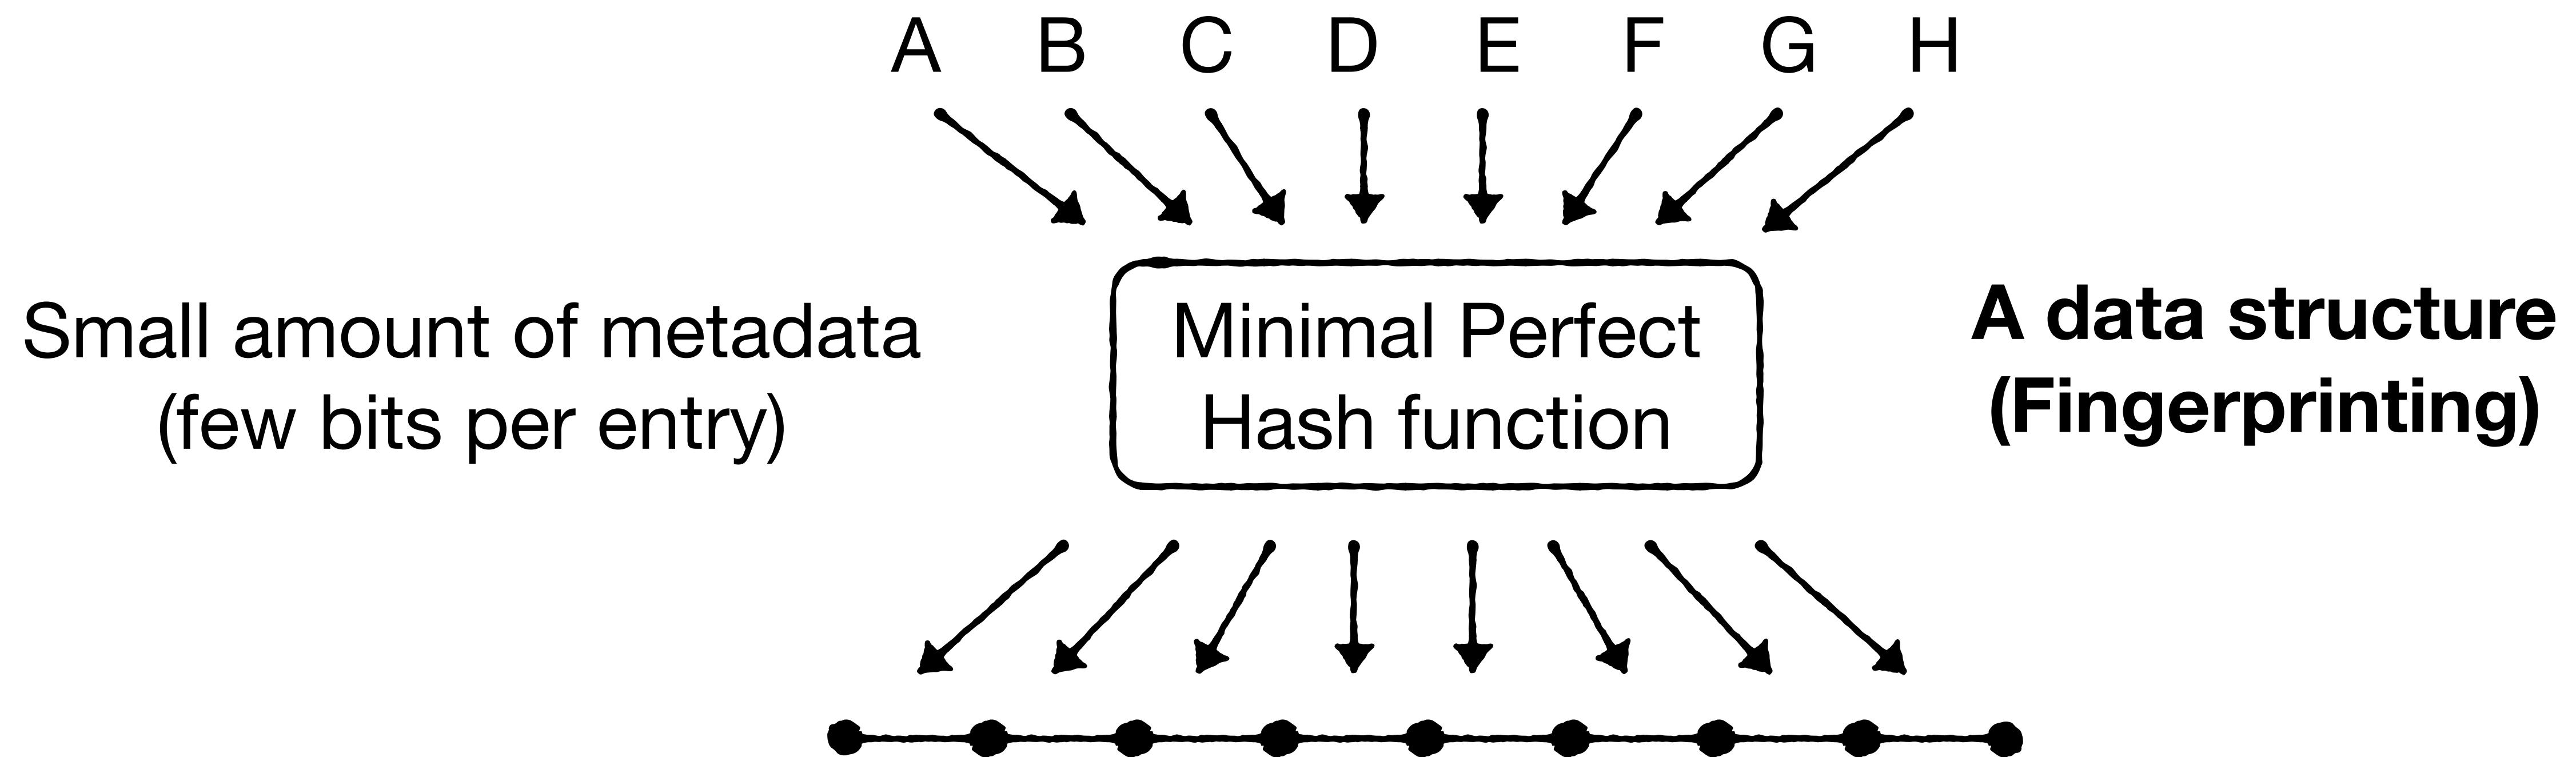
$$\lim_{N \rightarrow \infty} \approx \sqrt{2 \pi N} \cdot e^{-N} = 0$$

What can we do instead?

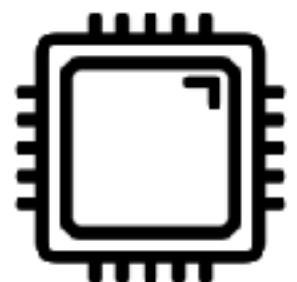








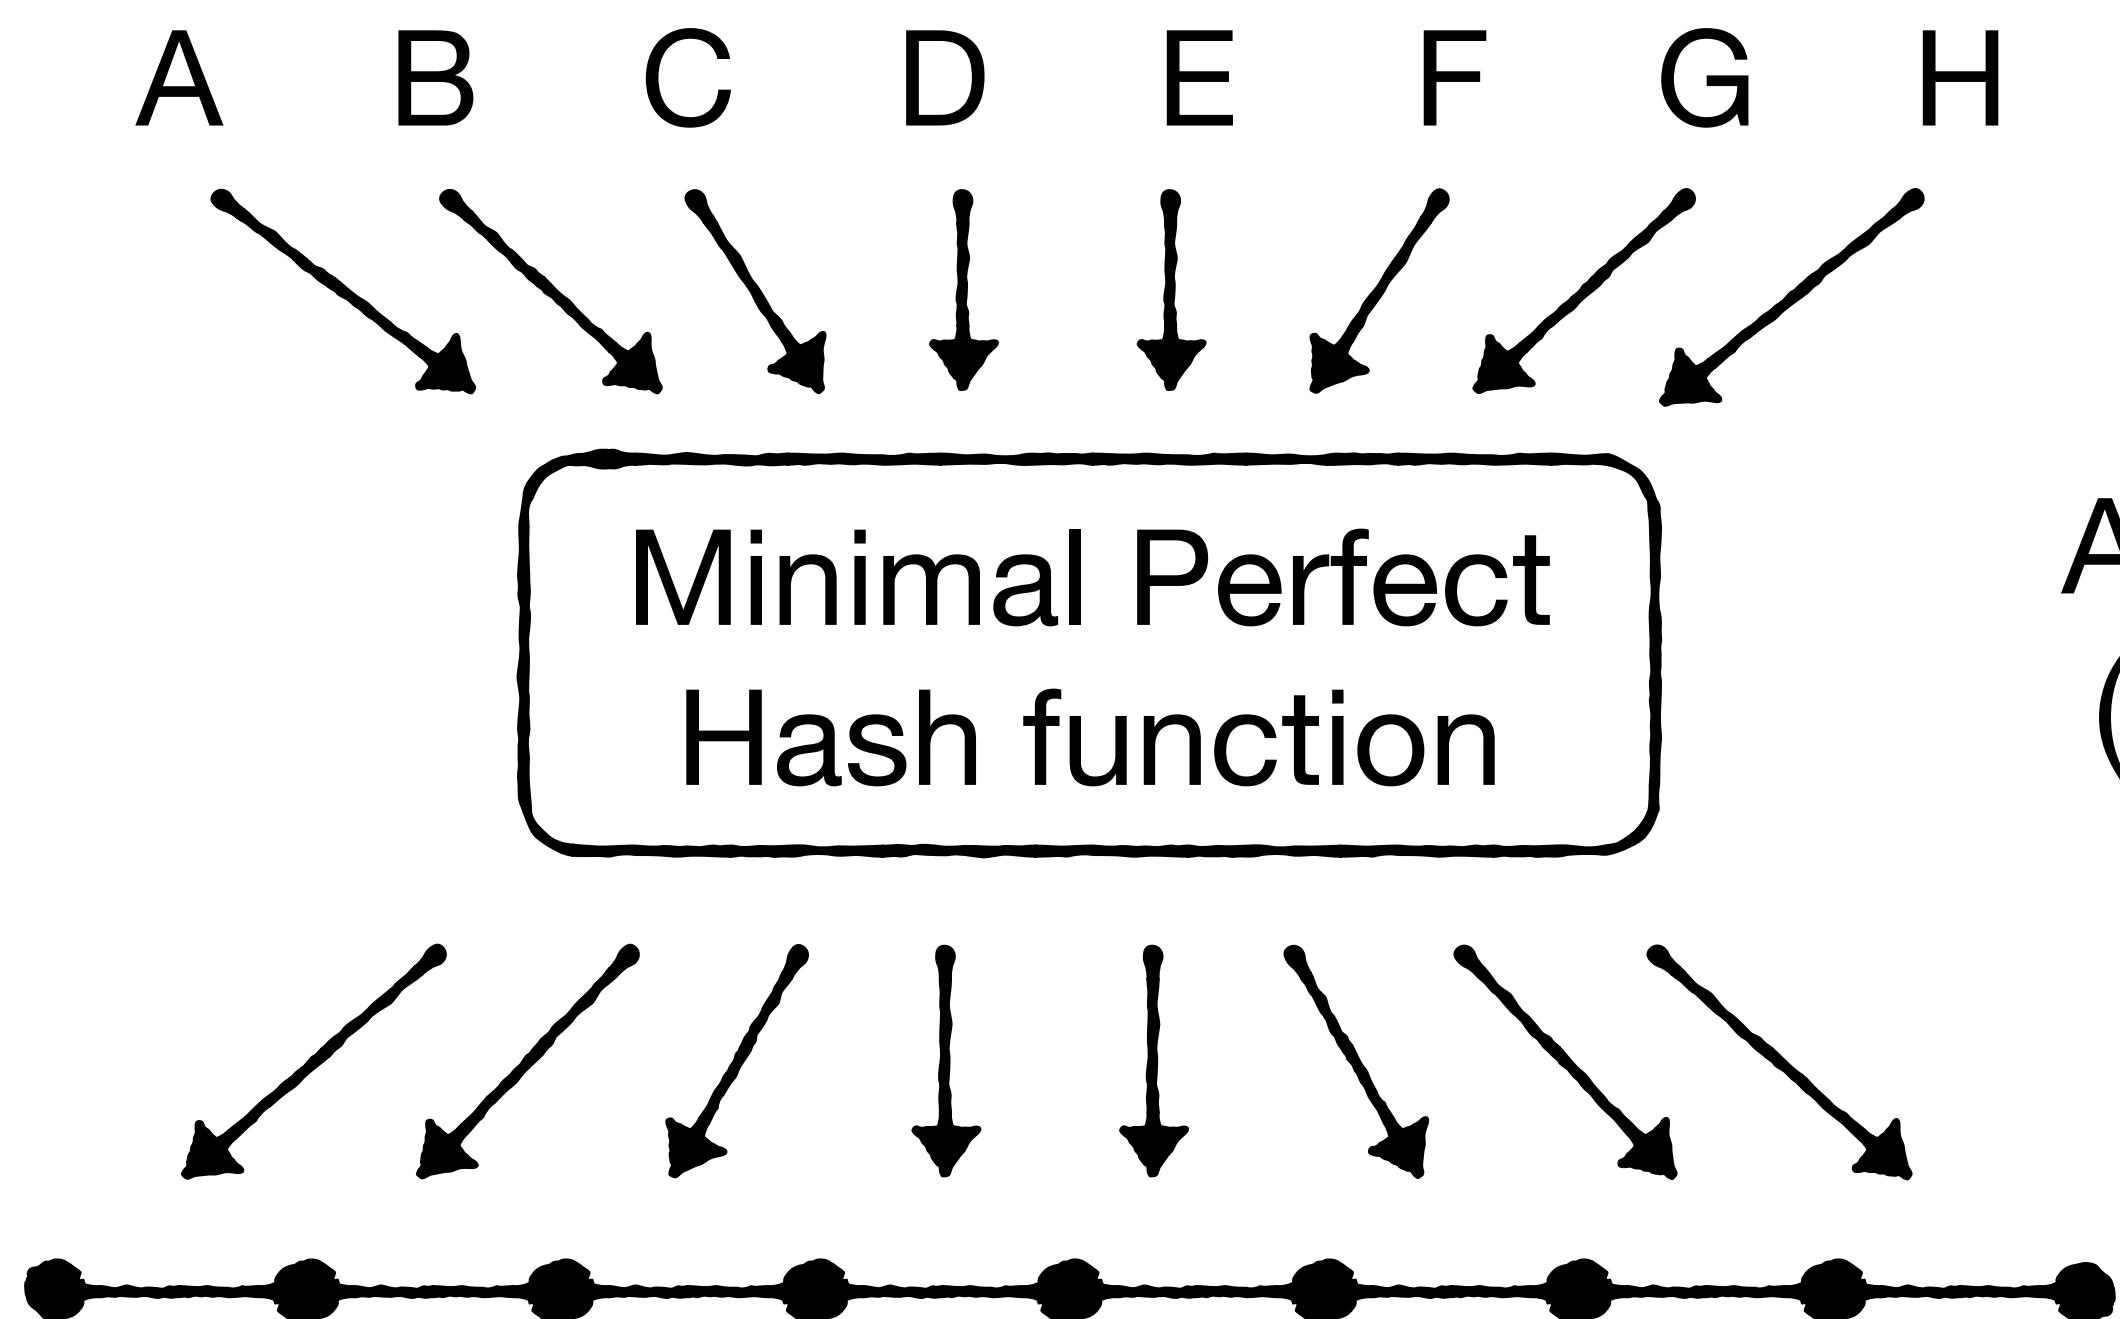
**Memory
(Bits / entry)**



Query cost



Construction time



**A data structure
(Fingerprinting)**

Fingerprinting

Fingerprinting-based Minimal Perfect Hashing Revisited. JEA 2023.

Piotr Beling.

Retrieval and Perfect Hashing using Fingerprinting. JEA 2014.

Ingo Müller, Peter Sanders, Robert Schulze & Wei Zhou.

Fast and Scalable Minimal Perfect Hashing for Massive Key Sets. SEA 2017.

Antoine Limasset, Guillaume Rizk, Rayan Chikhi, Pierre Peterlongo.

Meraculous: de novo genome assembly with short paired-end reads. PloS one 2011.

Jarrod A. Chapman ,Isaac Ho, Sirisha Sunkara, Shujun Luo, Gary P. Schroth, Daniel S. Rokhsar

Perfect Hashing for Network Applications. ISIT 2006.

Yi Lu, Balaji Prabhakar, Flavio Bonomi.

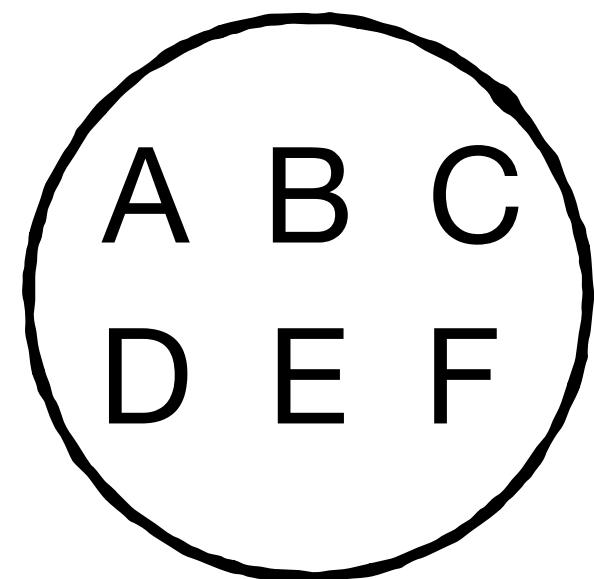
Fingerprinting

Fast and Scalable Minimal Perfect Hashing for Massive Key Sets. SEA 2017.

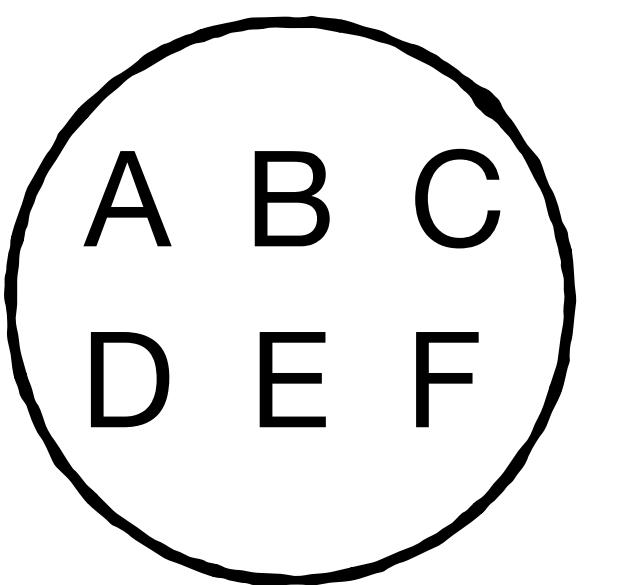
Antoine Limasset, Guillaume Rizk, Rayan Chikhi, Pierre Peterlongo.

Accessible & Experimental

Keys

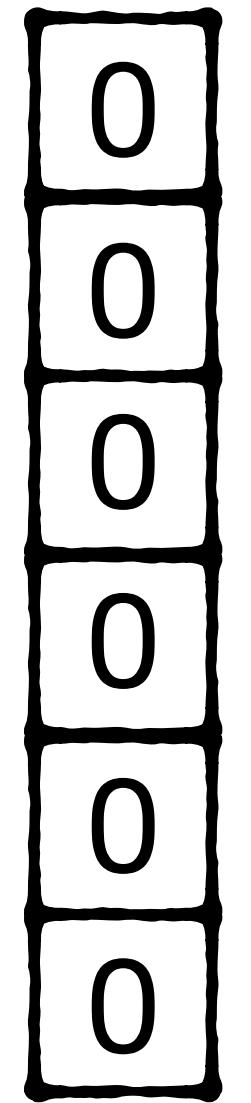


Keys

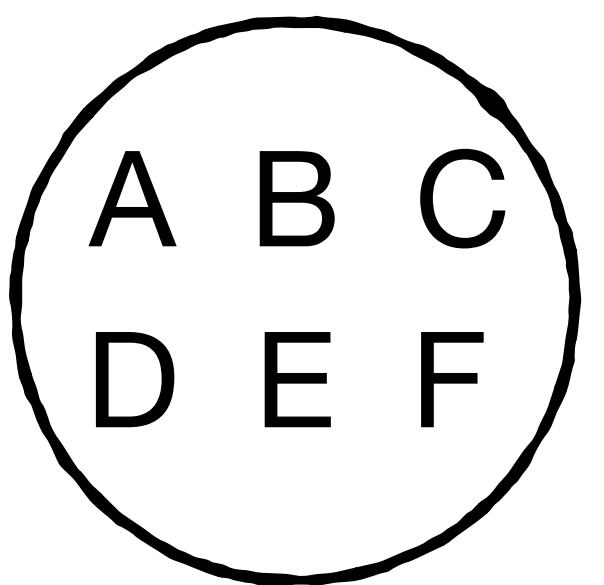


Hash₁

Bitmap B₁

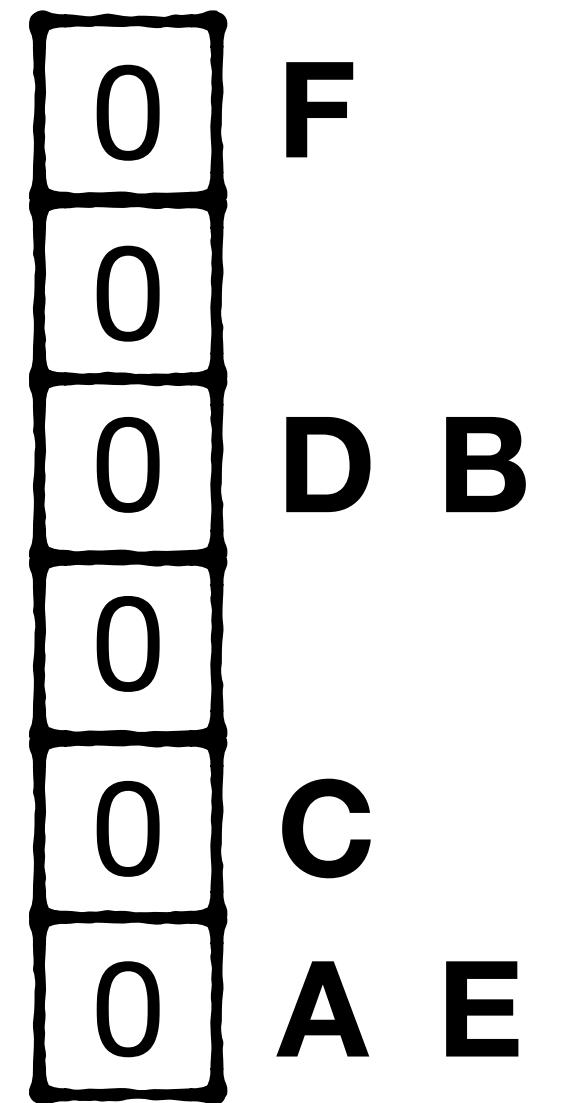


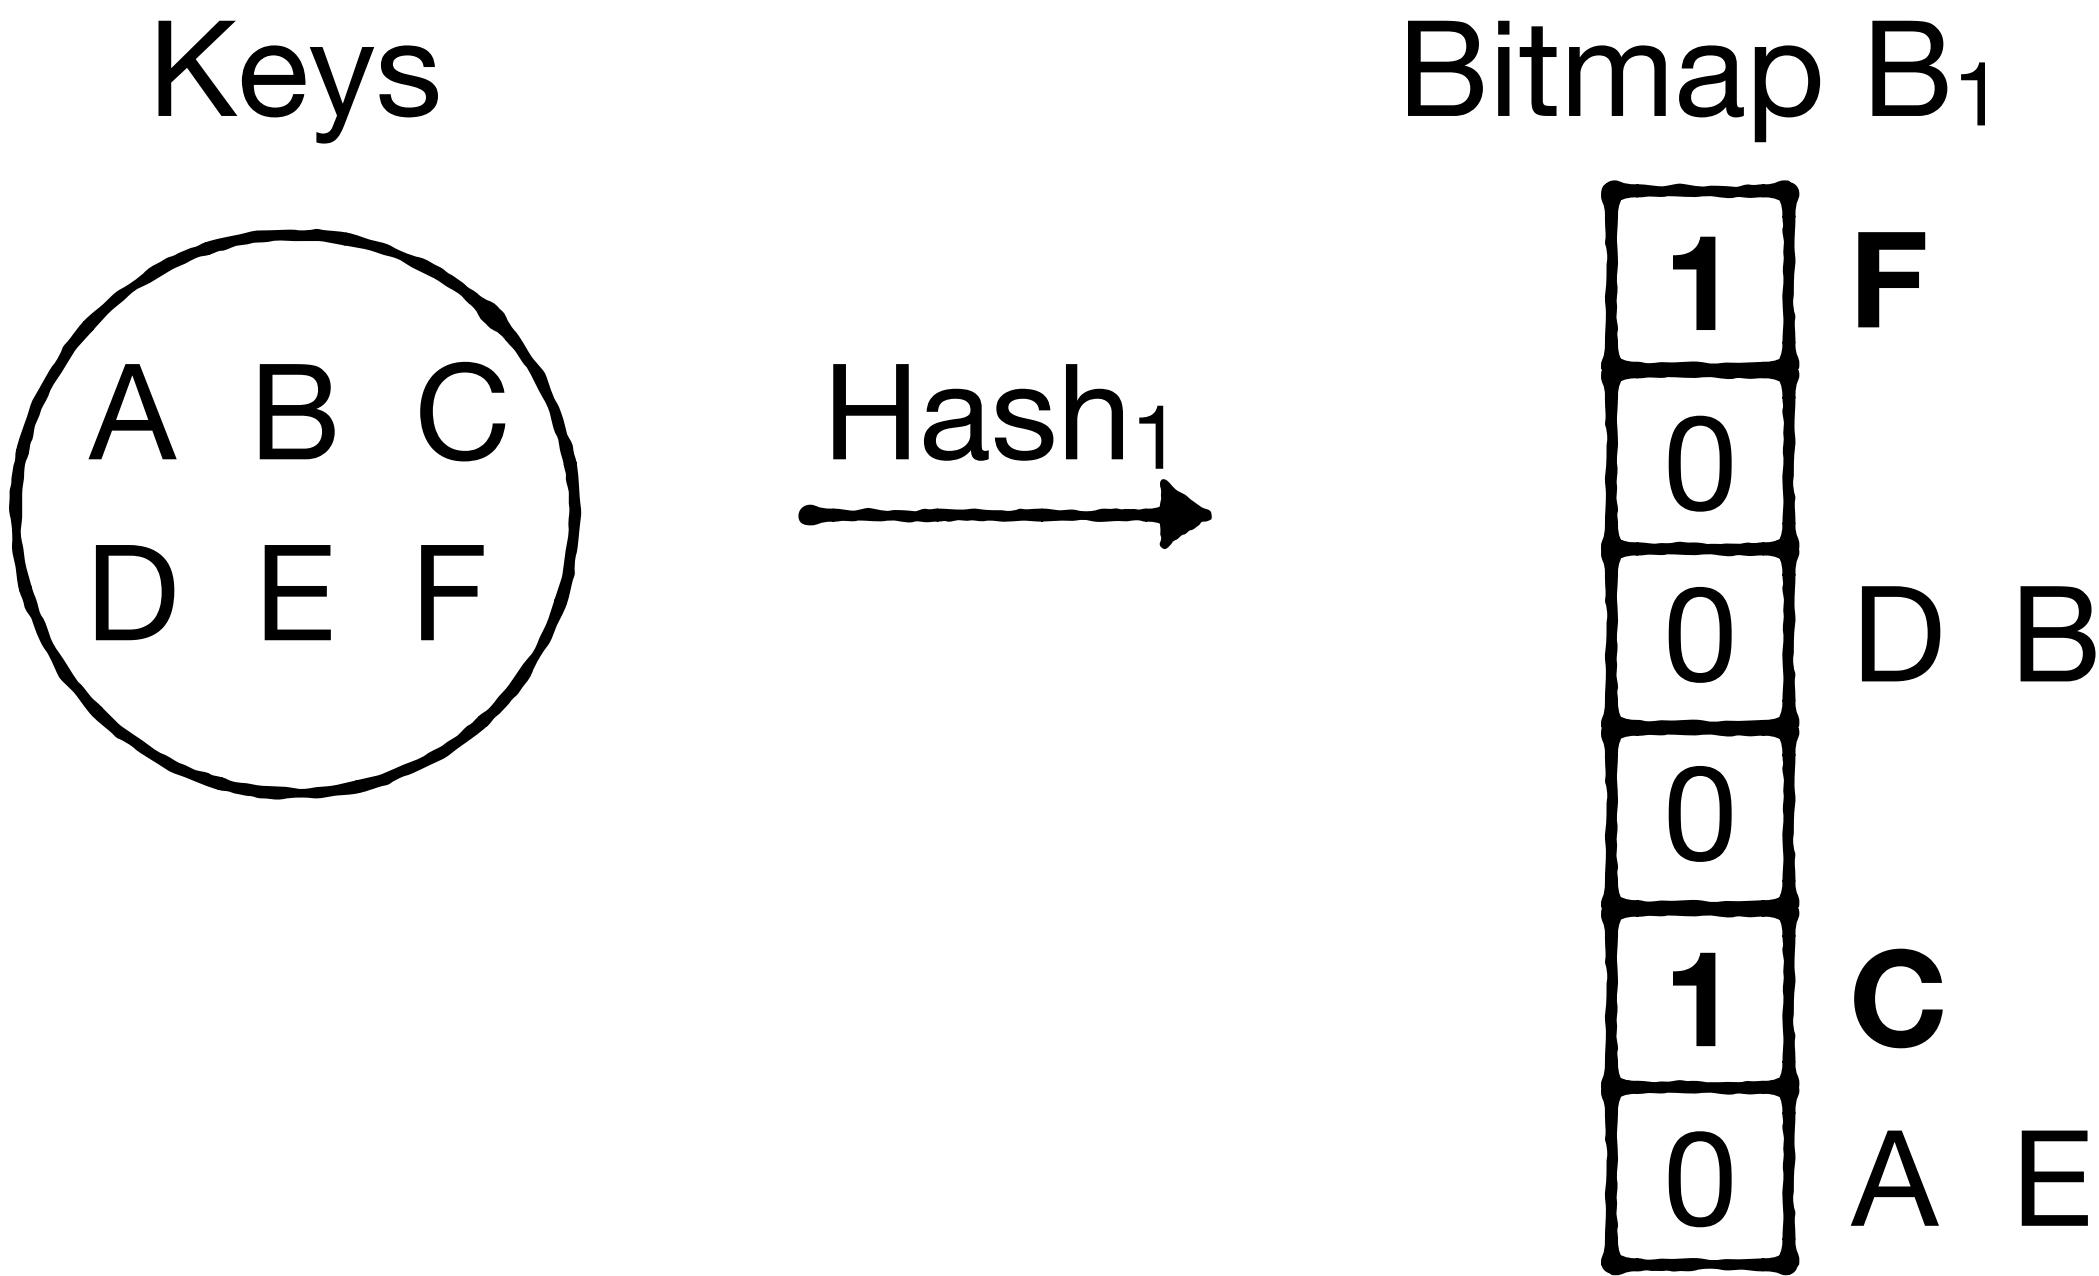
Keys



Hash₁ →

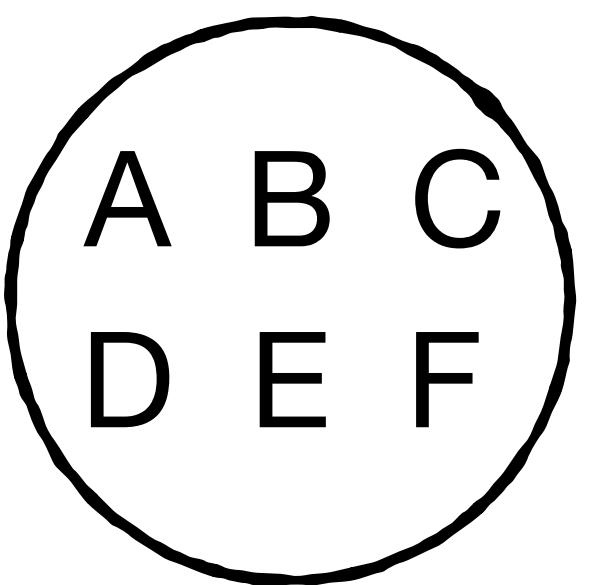
Bitmap B₁





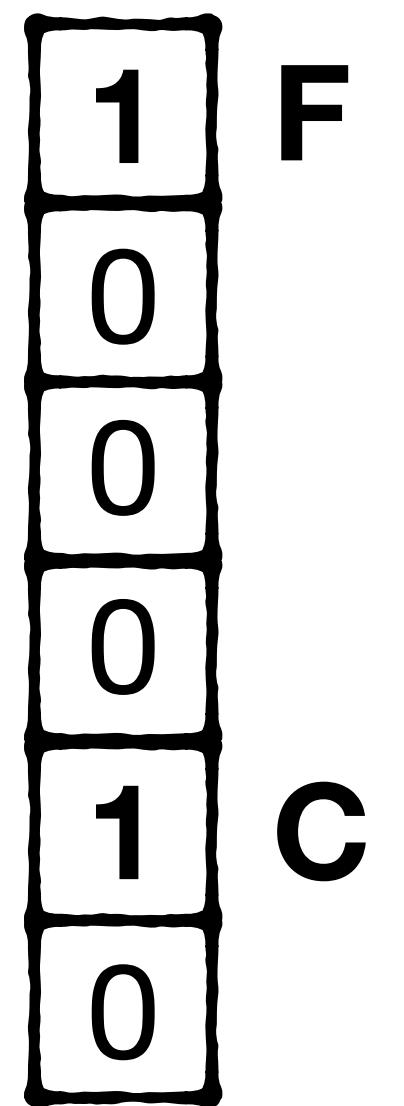
Set to 1 only if one entry mapped to this bit

Keys

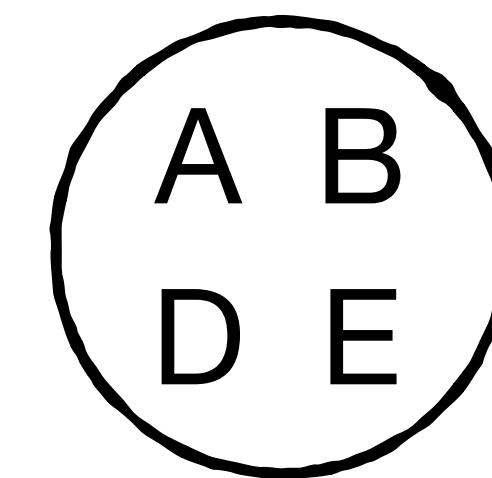


Hash₁ →

Bitmap B₁

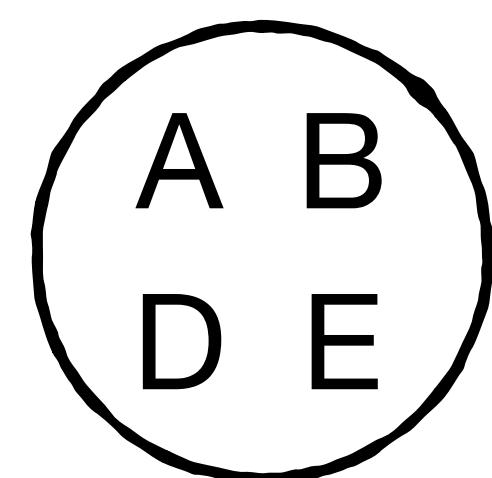


Continue
recursively



Bitmap B₁

1	F
0	
0	
0	
1	C
0	



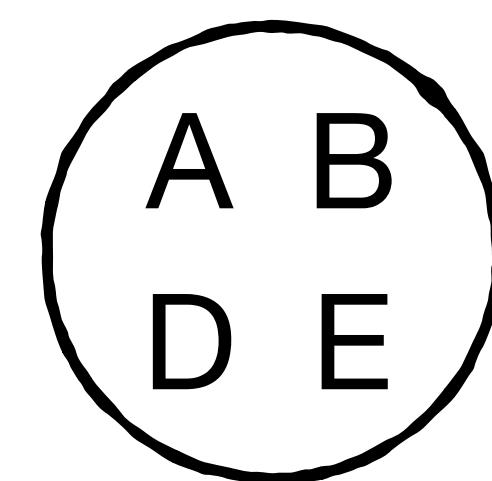
Bitmap B₂

Hash₂ →

0
0
0
0

Bitmap B₁

1	F
0	
0	
0	
1	C
0	



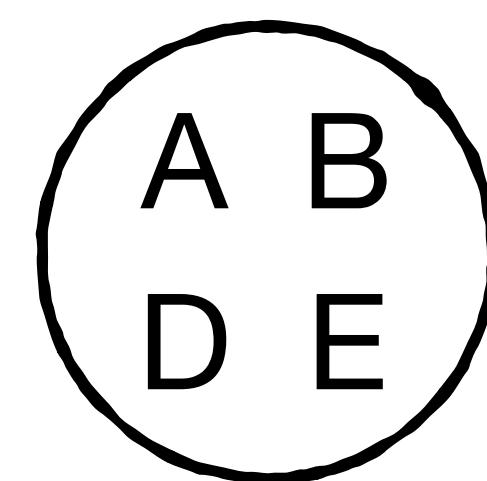
Bitmap B₂

0	E
0	B D
0	A
0	

Hash₂

Bitmap B₁

1	F
0	
0	
0	
1	C
0	



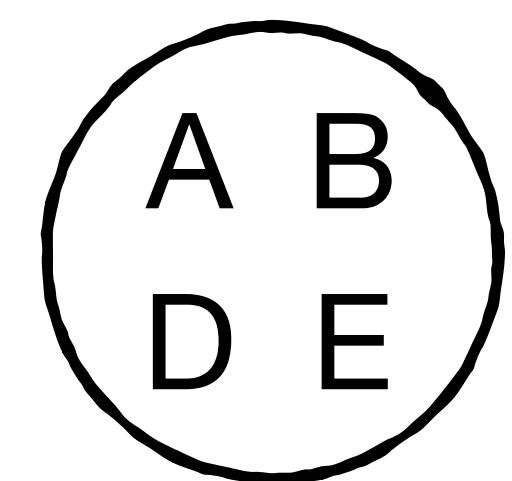
Bitmap B₂

1	E
0	B D
1	A
0	

Hash₂ →

Bitmap B₁

1	F
0	
0	
0	
1	C
0	

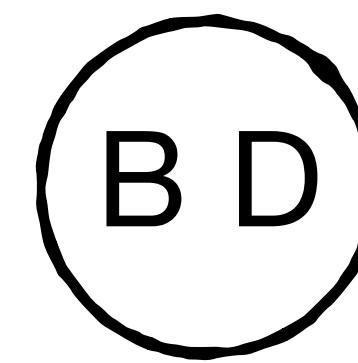


Bitmap B₂

1	E
0	
1	A
0	

Hash₂ →

Continue



Bitmap B₁

1	F
0	
0	
0	
1	C
0	

Bitmap B₂

1	E
0	
1	A
0	

Continue

B D

Bitmap B_1

1
0
0
0
1
0

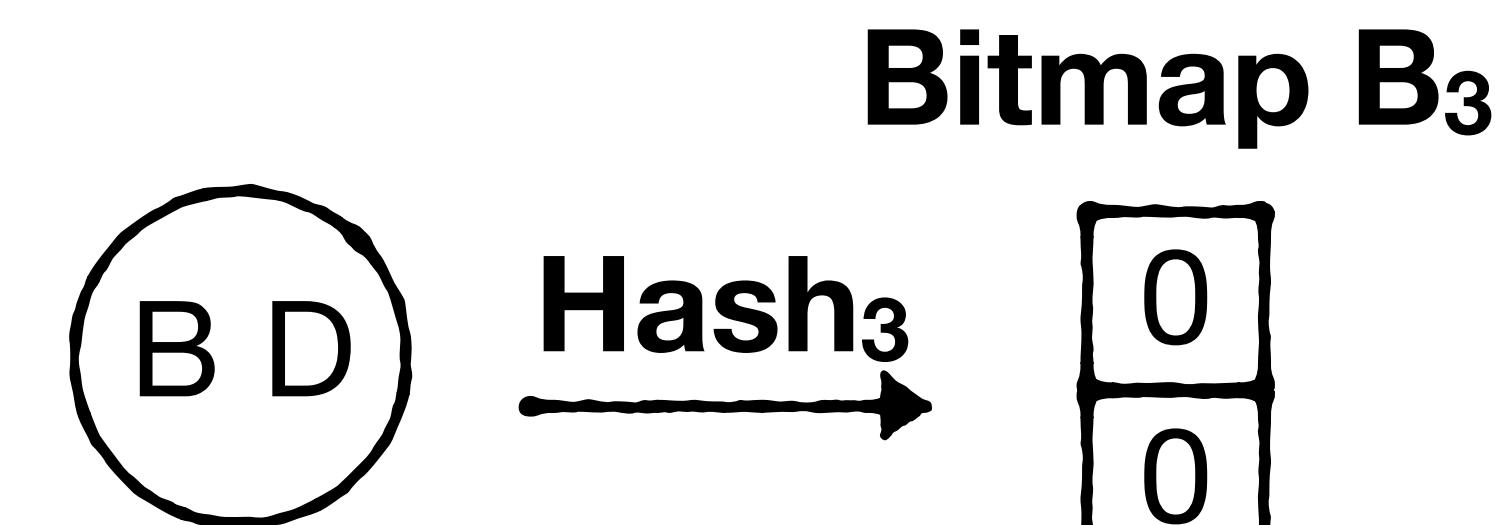
F

Bitmap B_2

1
0
1
0

E

A



Bitmap B_3

0
0

Bitmap B_1

1
0
0
0
1
0

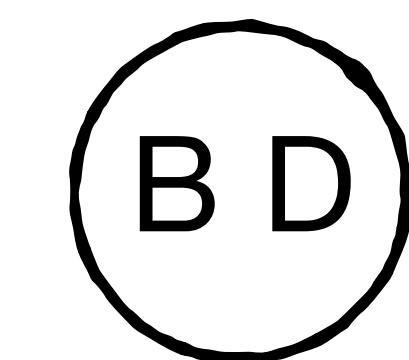
F

Bitmap B_2

1
0
1
0

E

A



Bitmap B_3

0
0

B

D

$\xrightarrow{\text{Hash}_3}$

Bitmap B_1

1
0
0
0
1
0

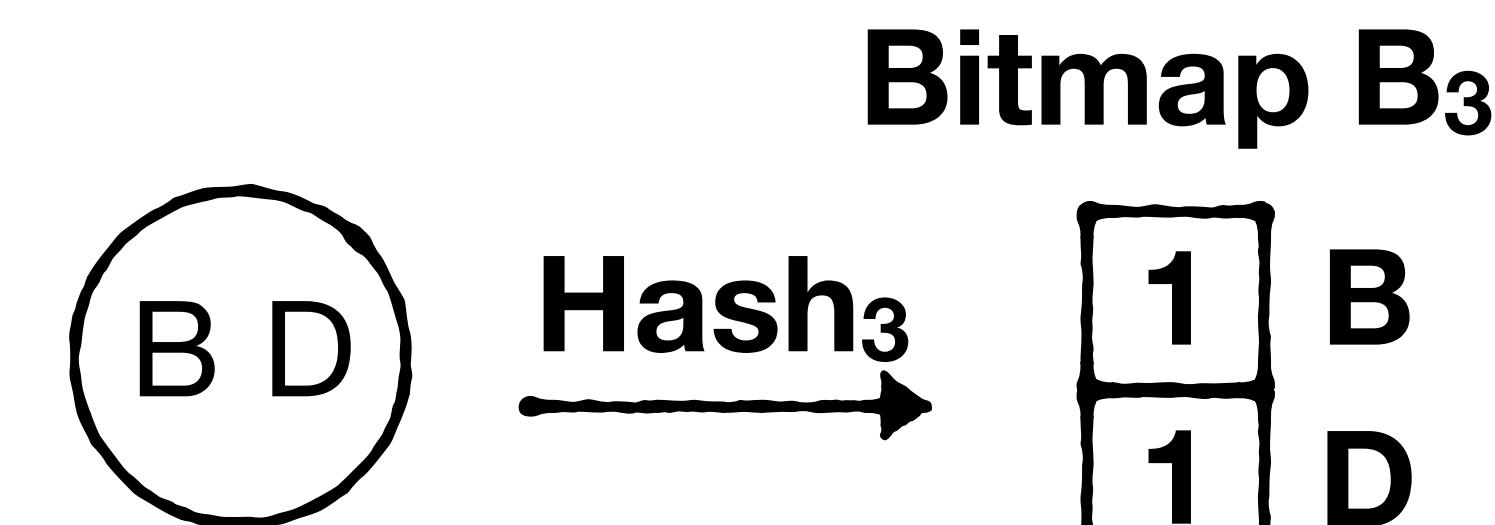
F

Bitmap B_2

1
0
1
0

E

A



Bitmap B_3

1
1

B

D

Bitmap B₁

1	F
0	
0	
0	
1	C
0	

Bitmap B₂

1	E
0	
1	A
0	

Bitmap B₃

1	B
1	D

Concatenate bitmaps

Bitmap B₁

1	F
0	
0	
0	
1	C
0	

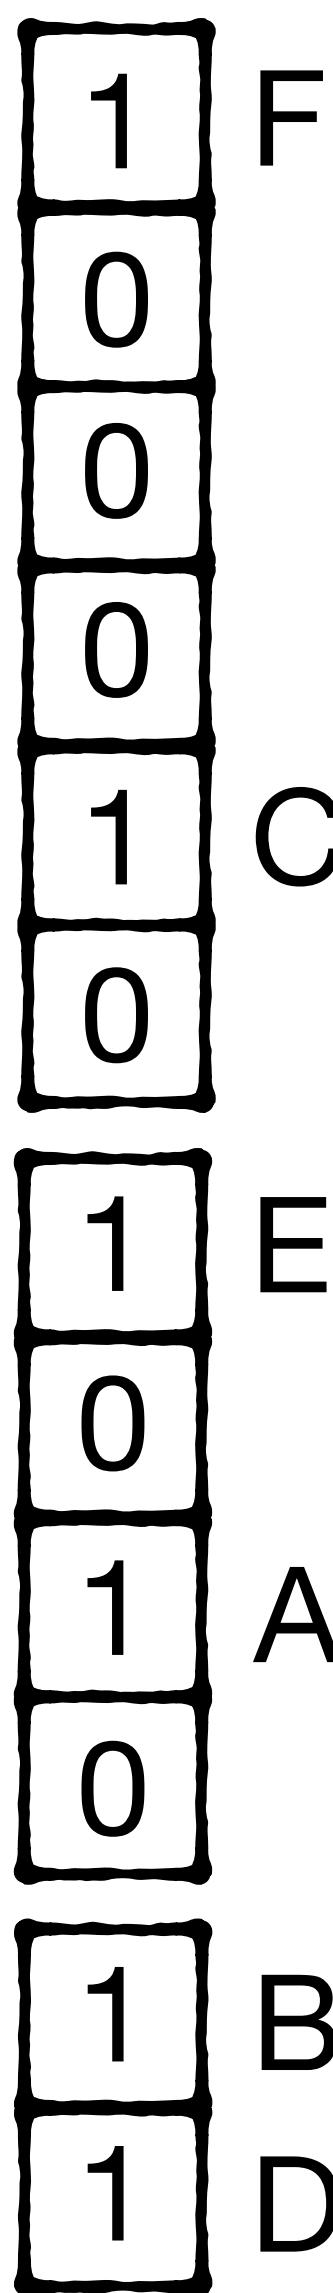
Bitmap B₂

1	E
0	
1	A
0	

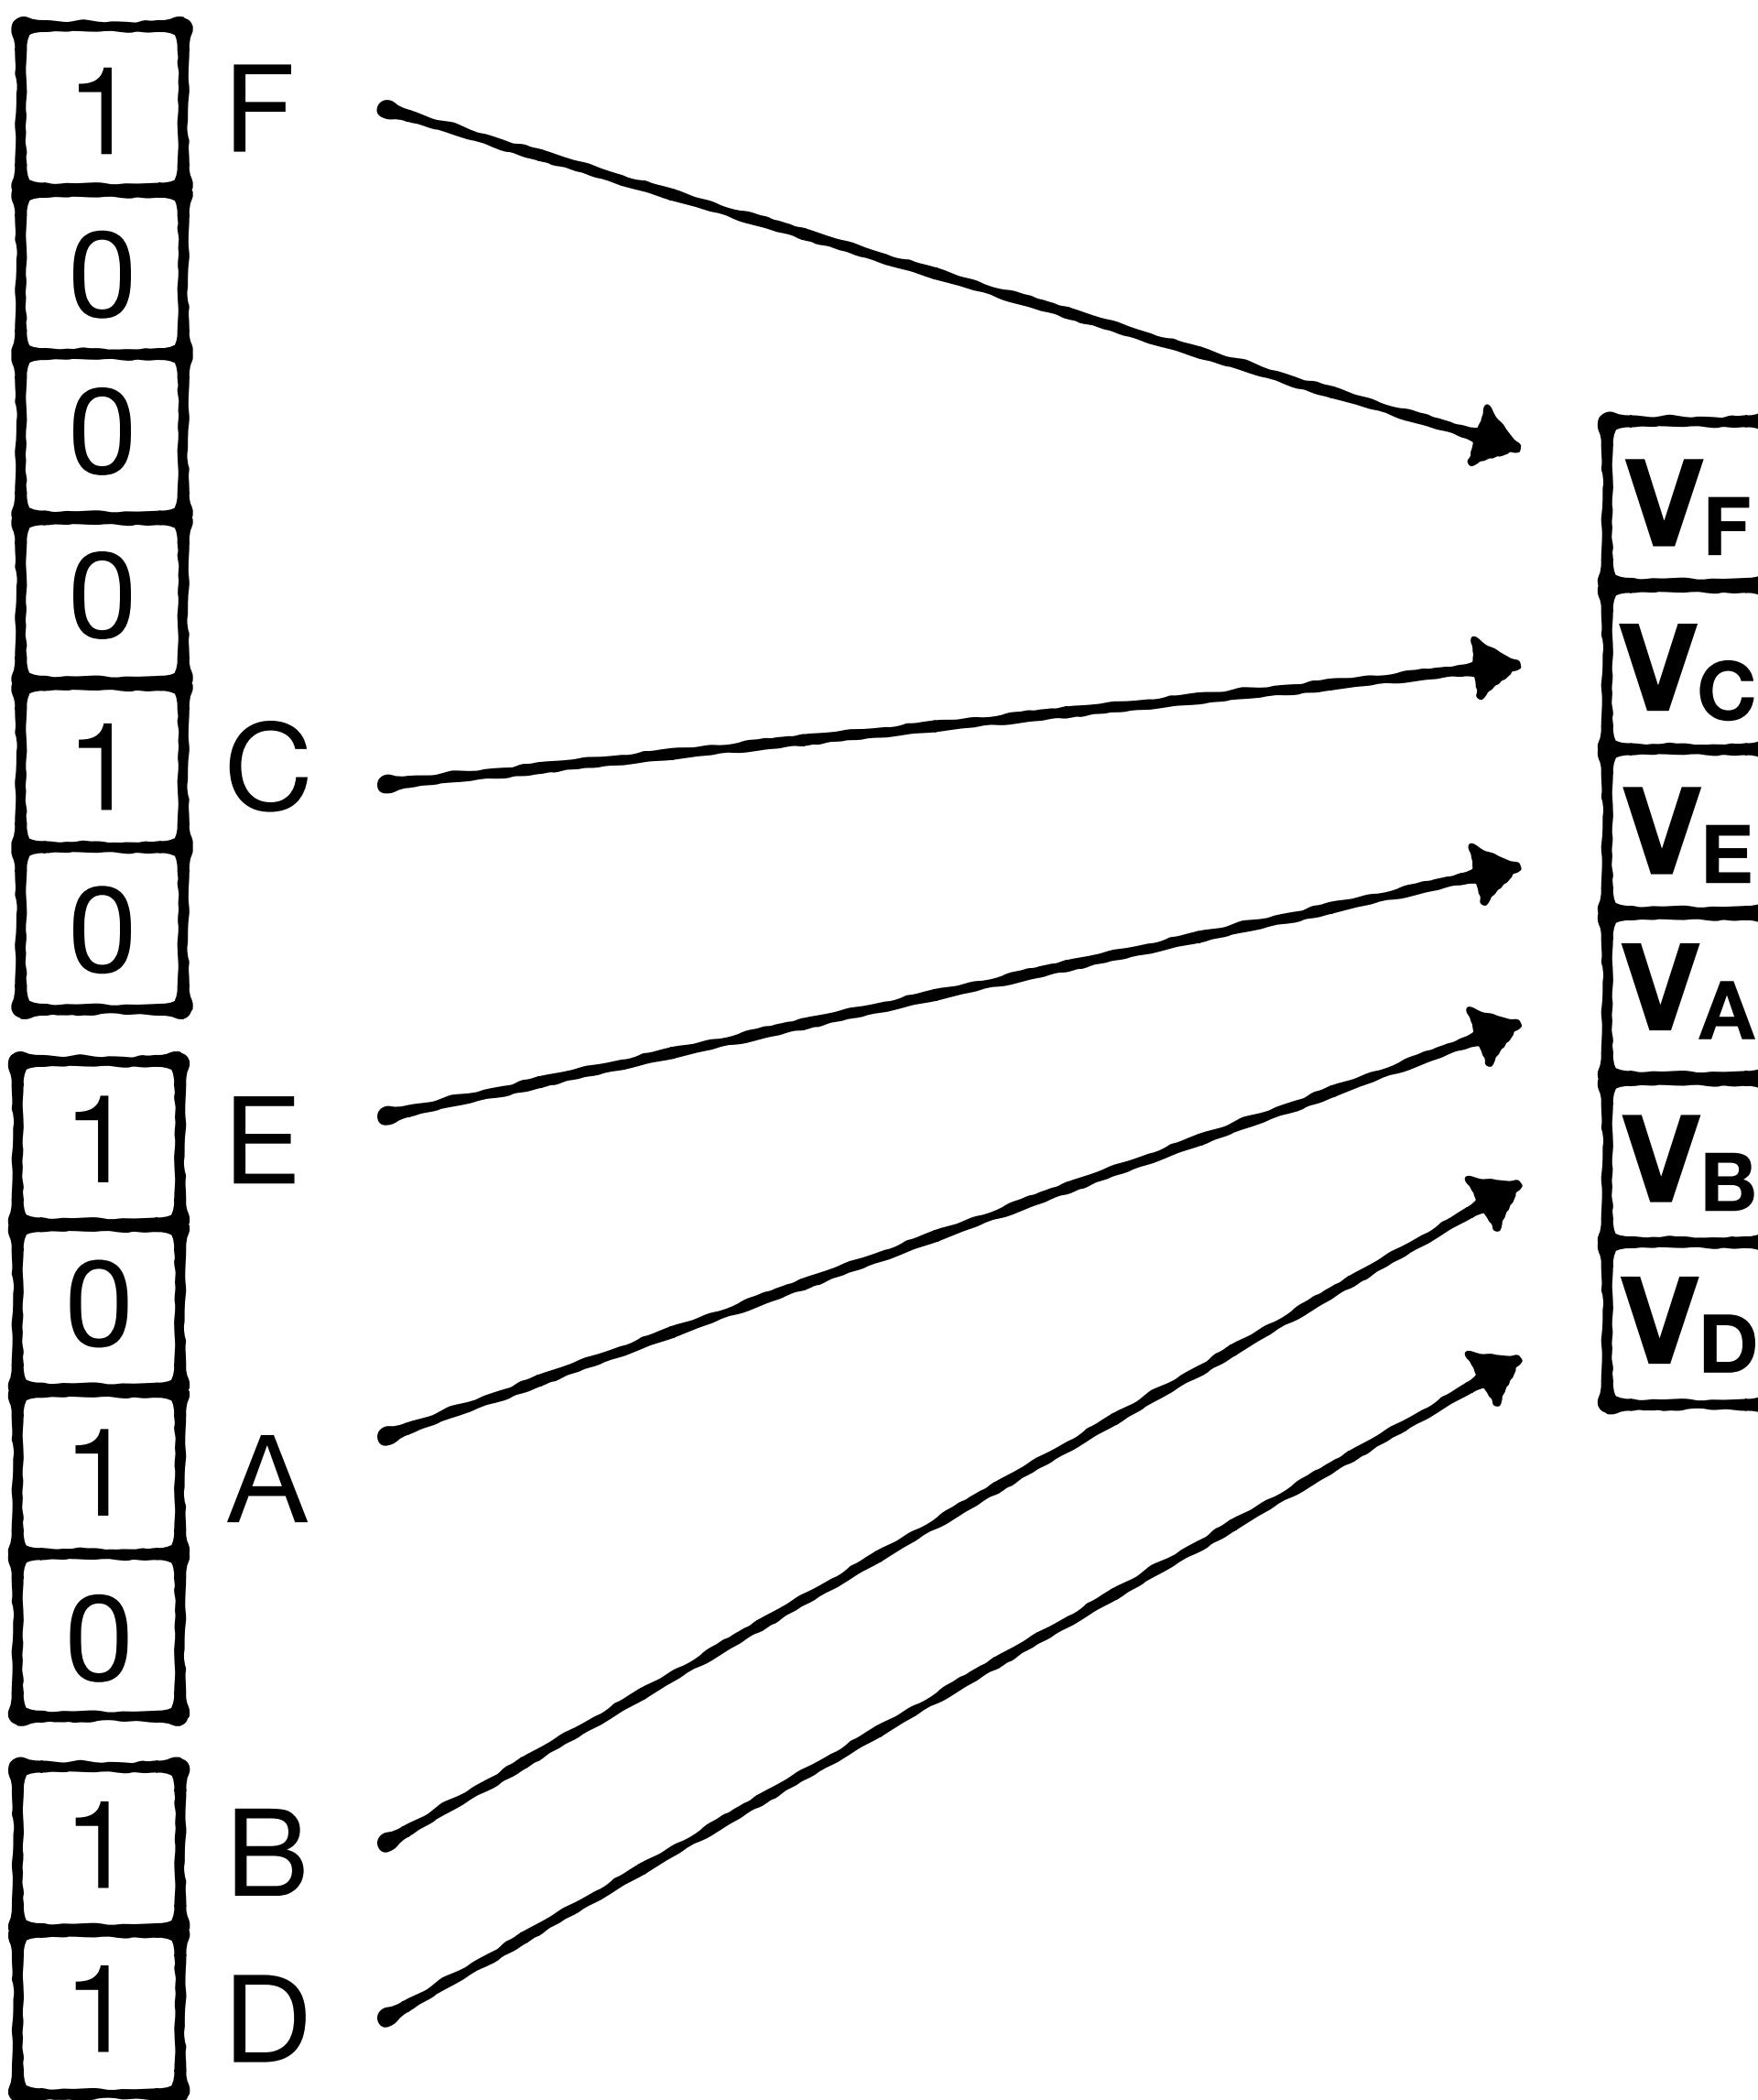
Bitmap B₃

1	B
1	D

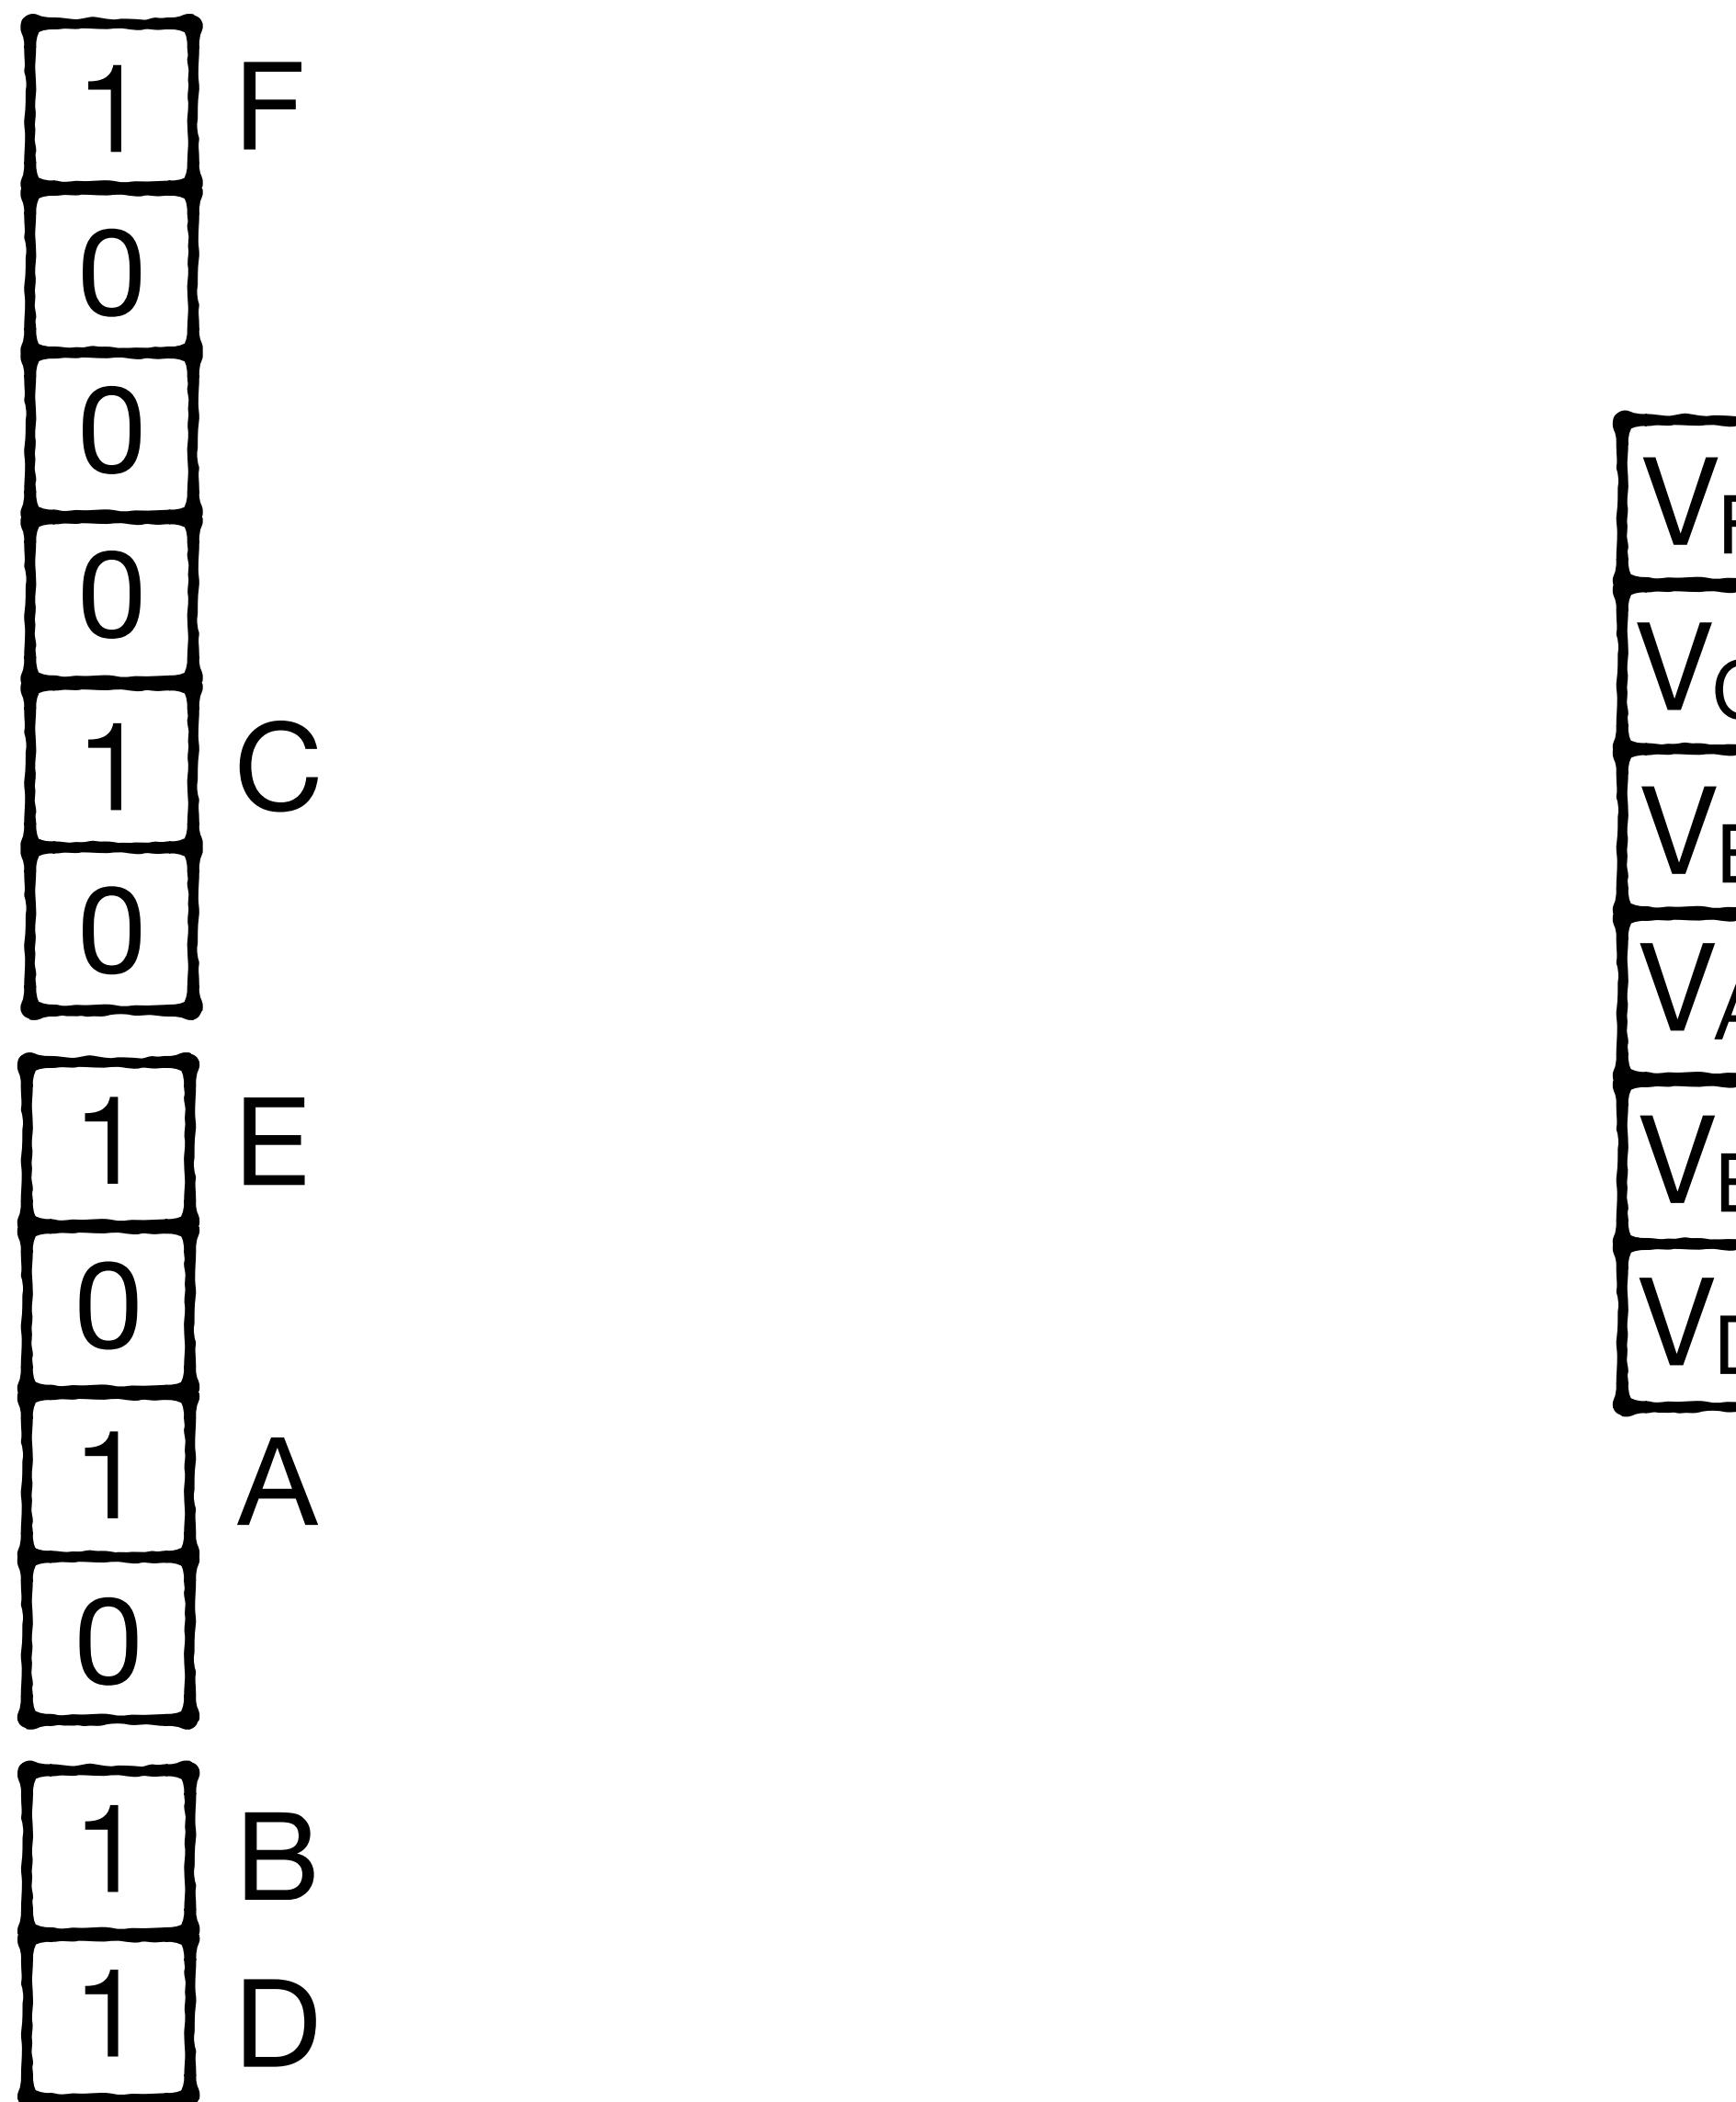
Concatenate bitmaps



Bijection is now established :)

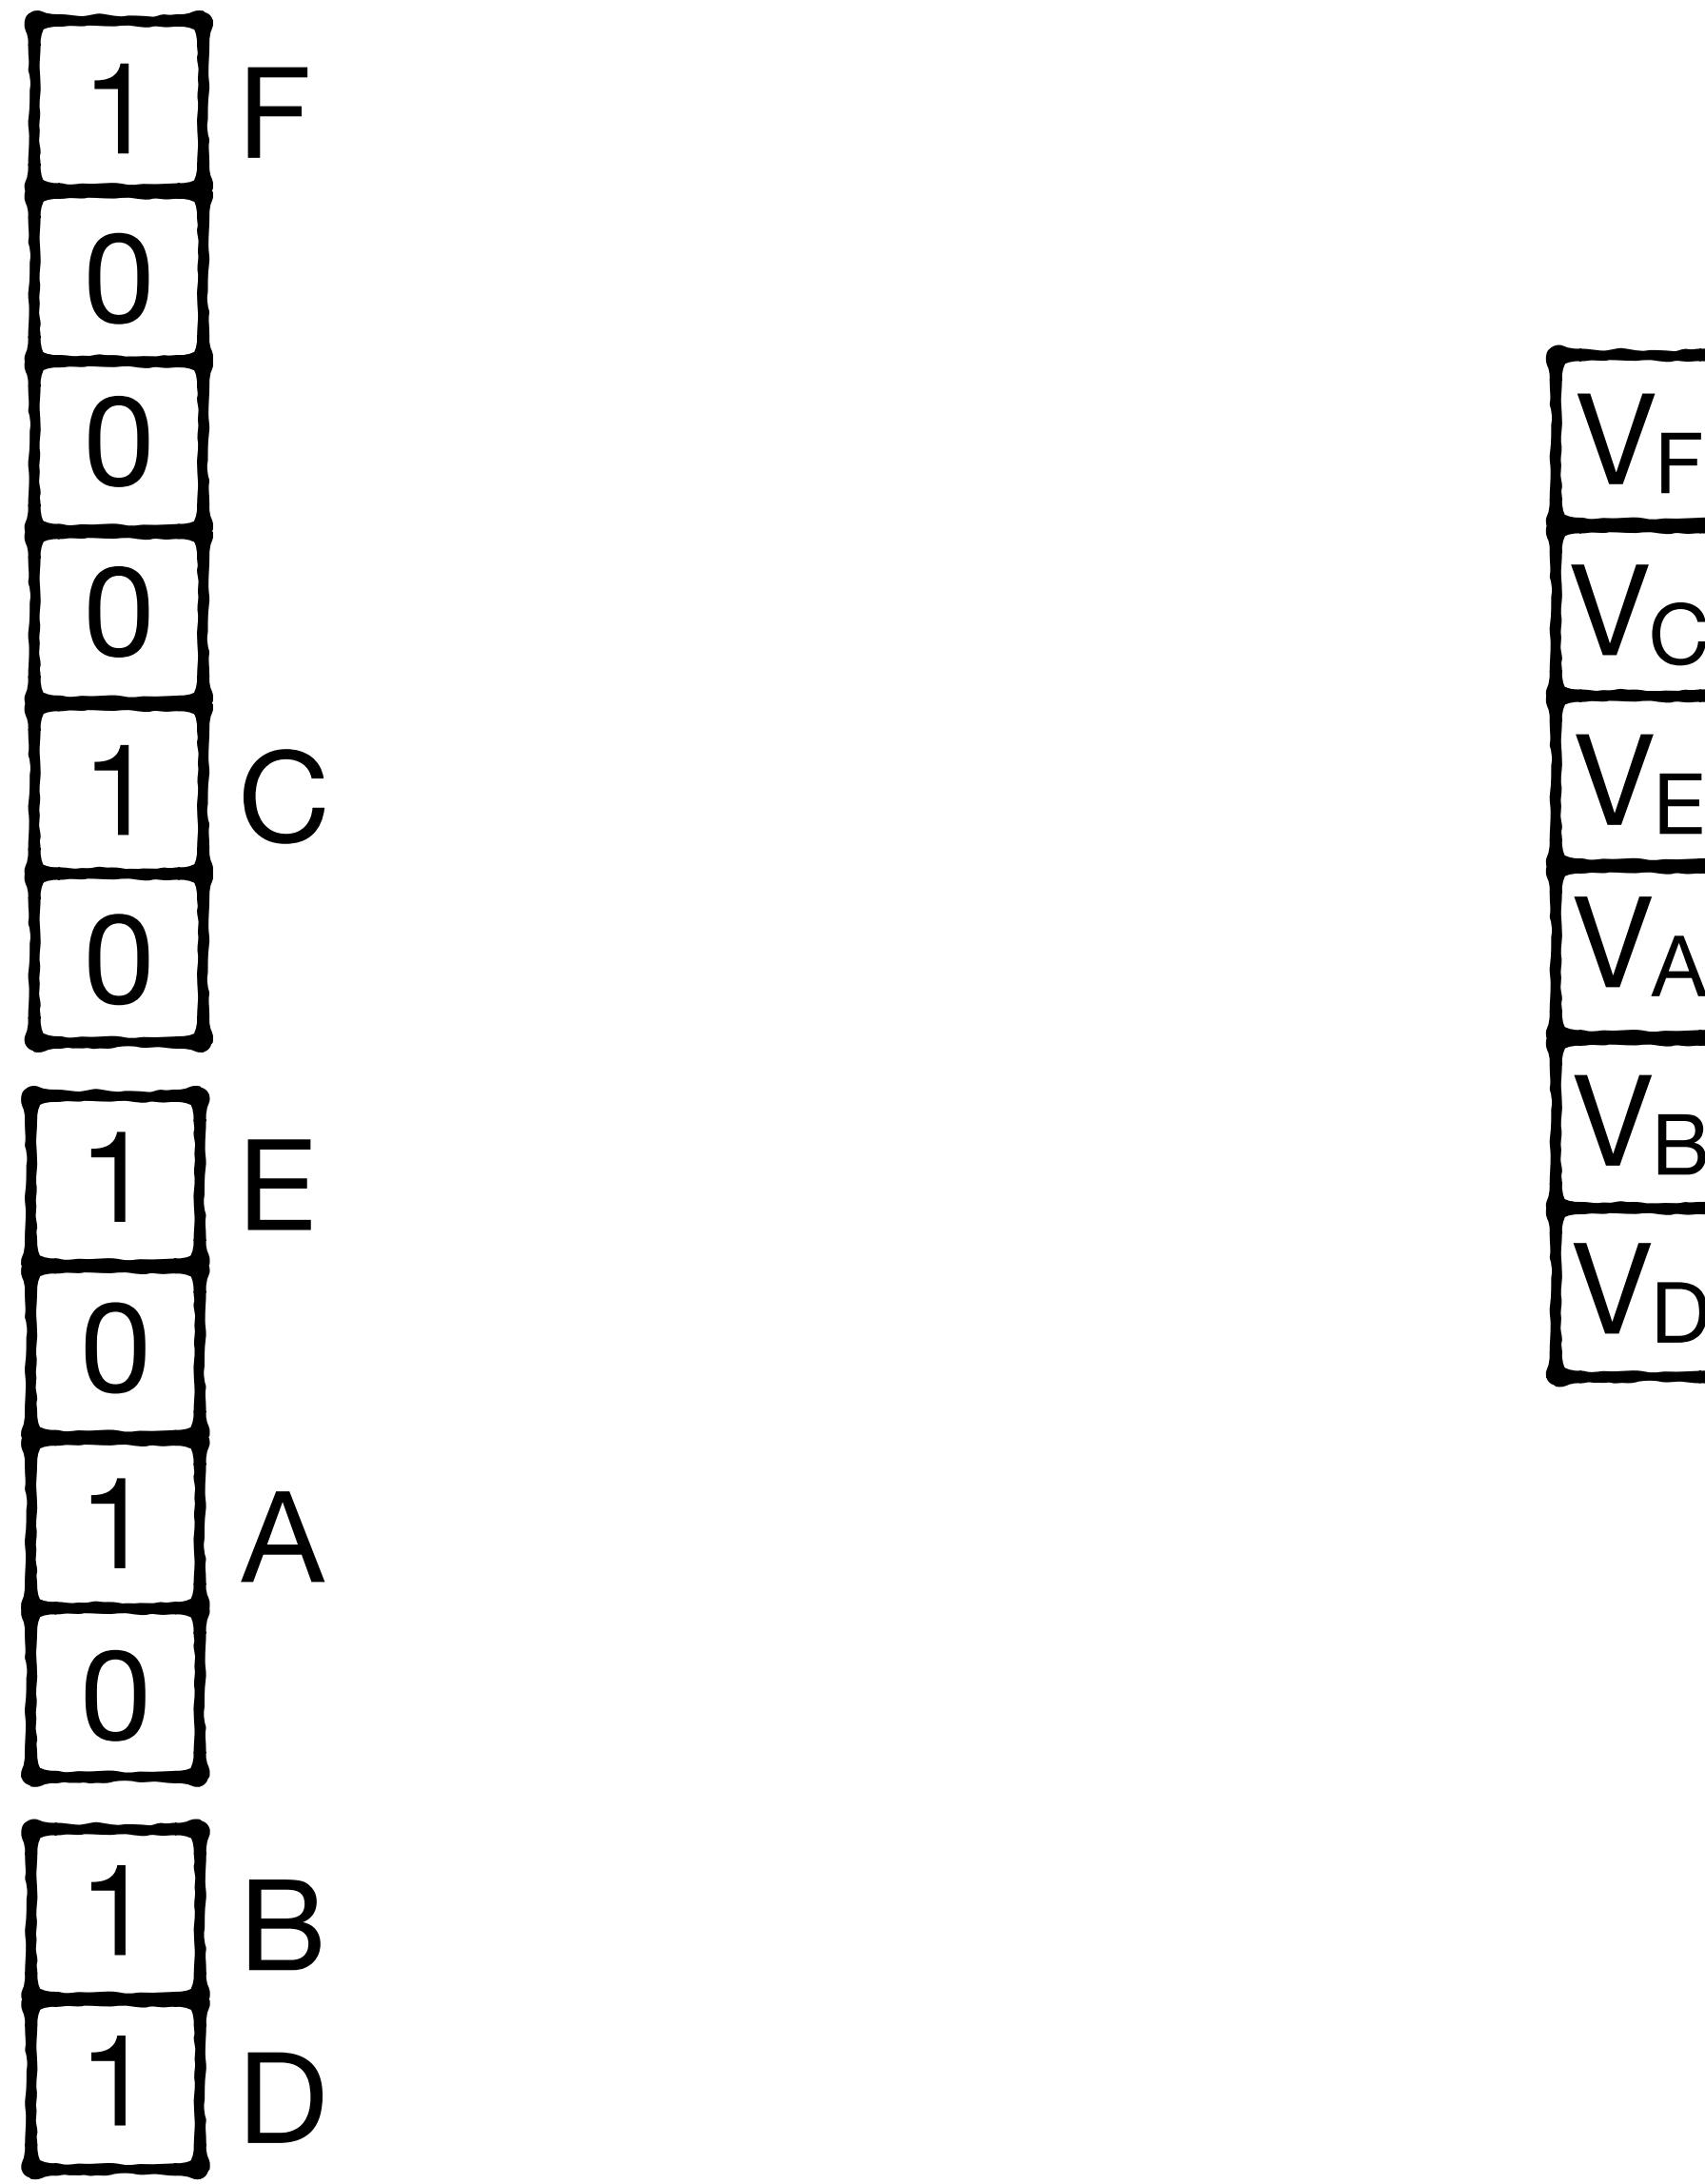


Bijection is now established :)



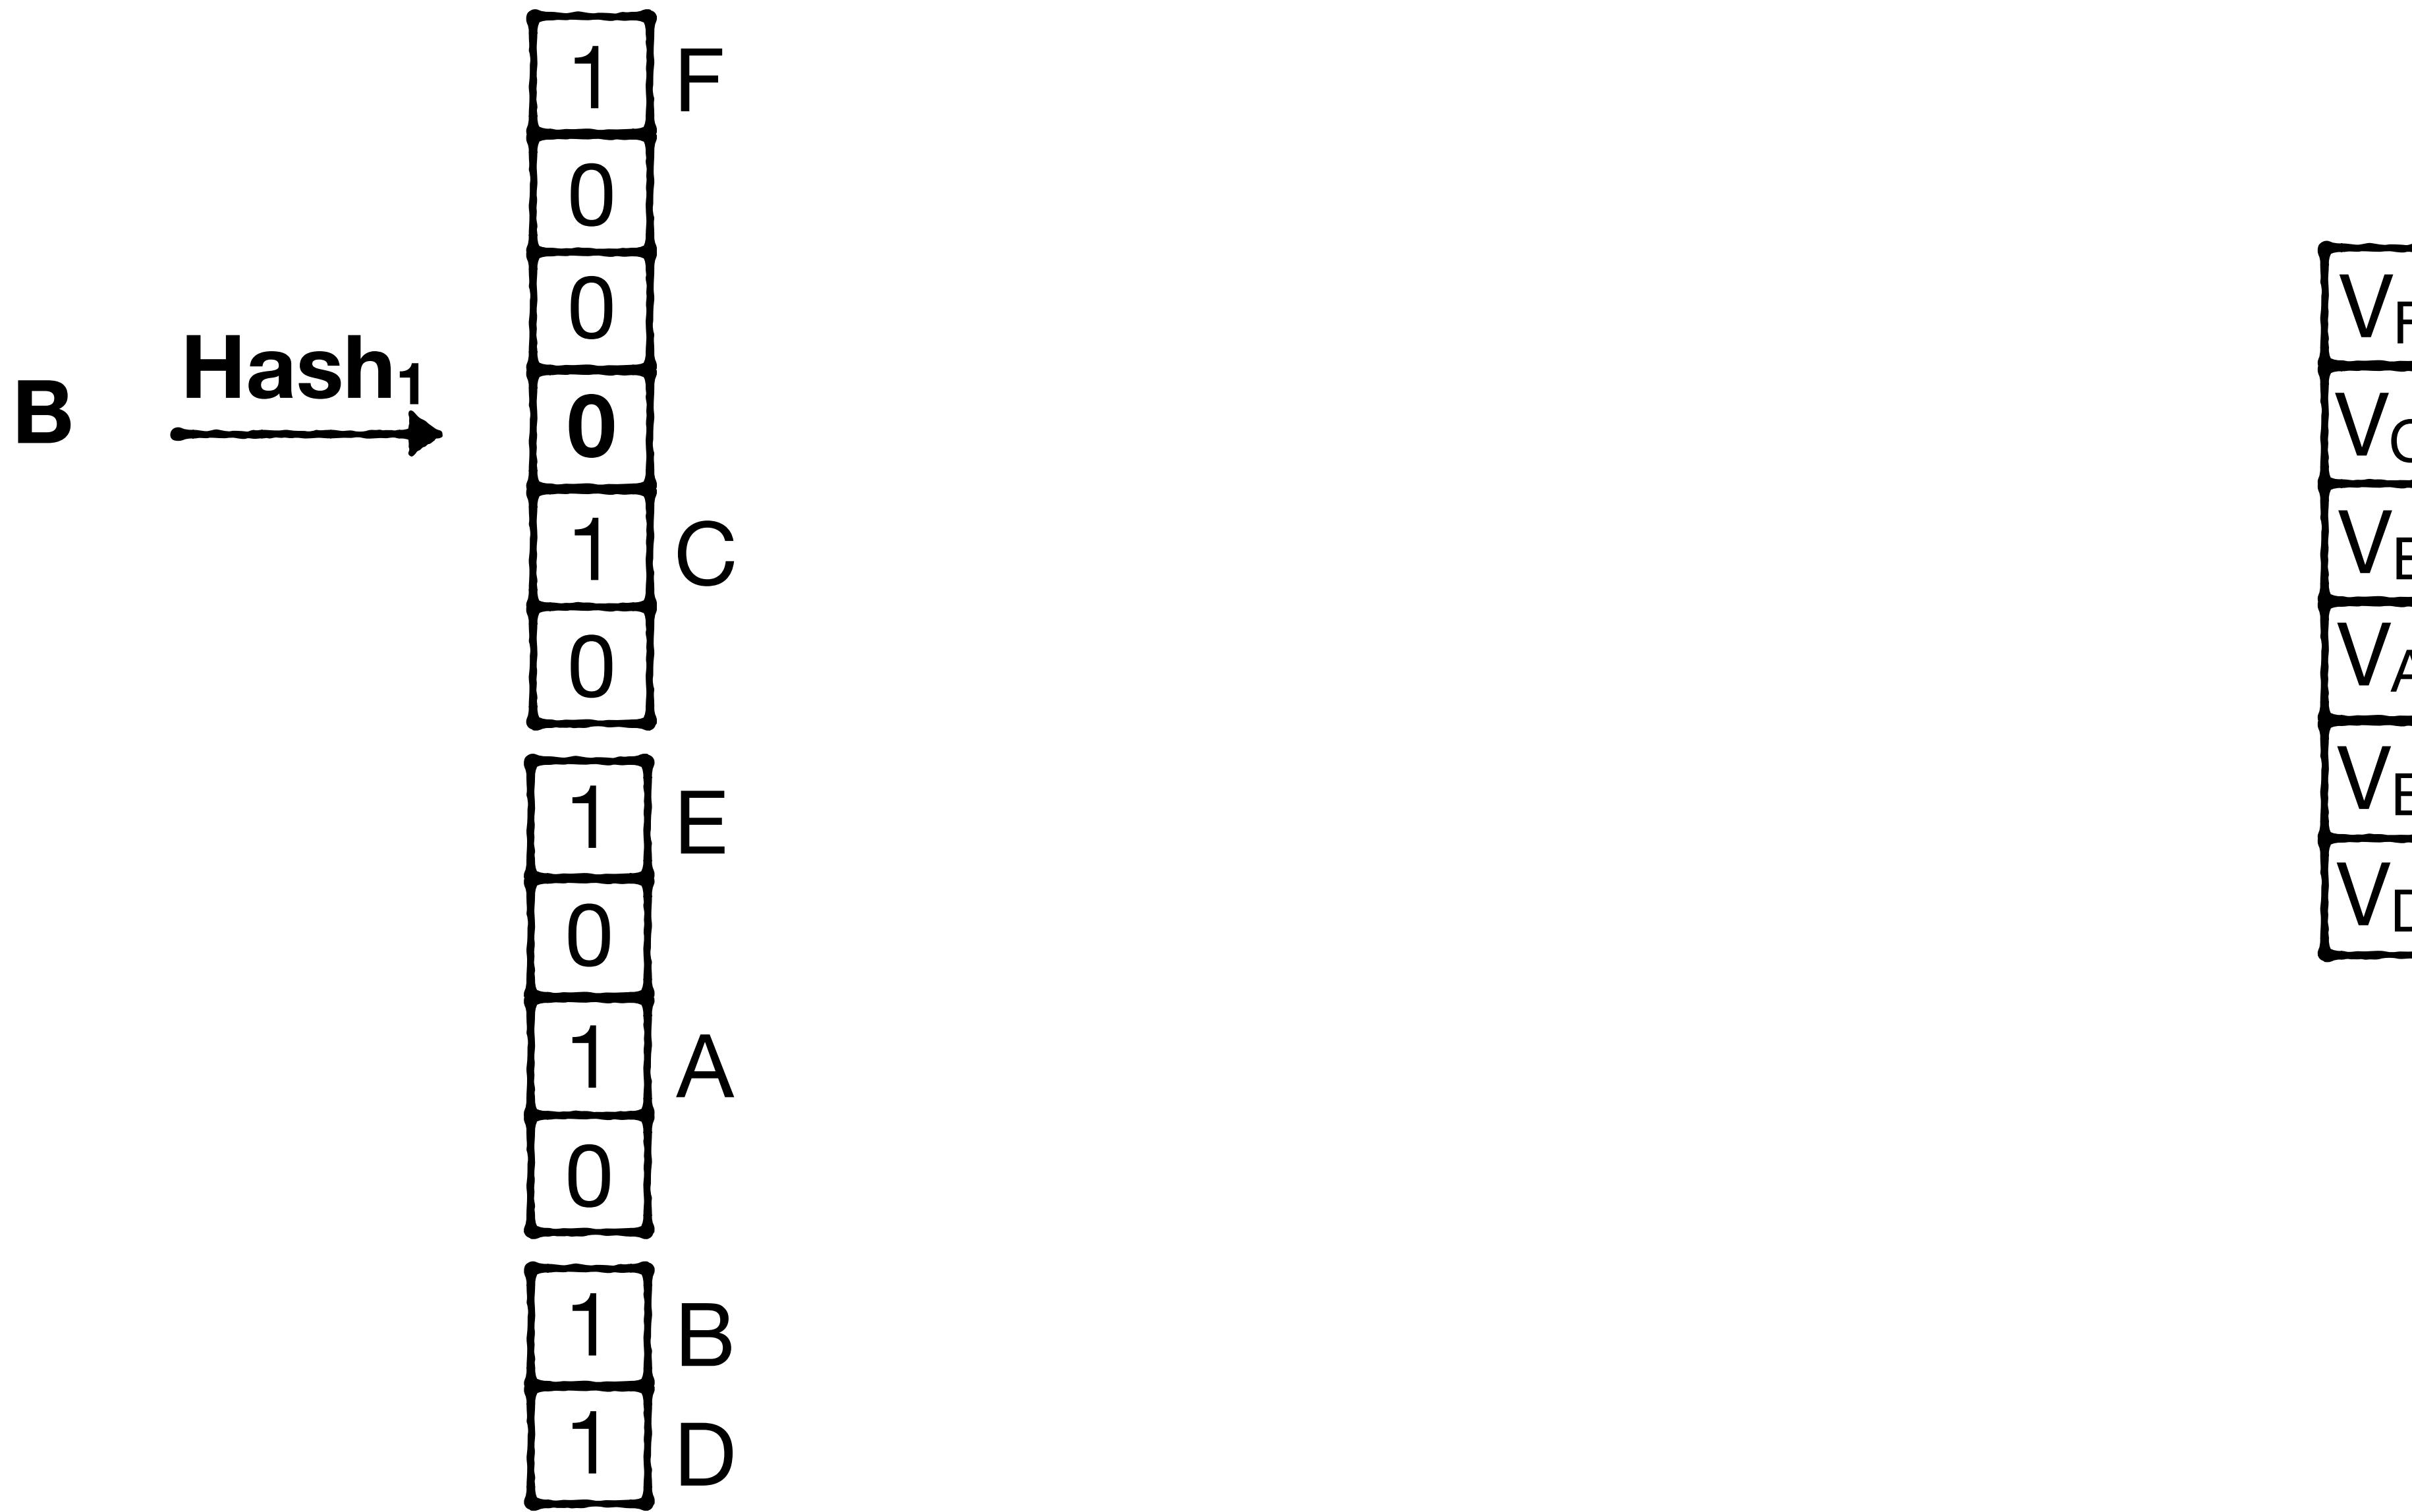
**Array stores a value/pointer
associated with each key**

How to query?



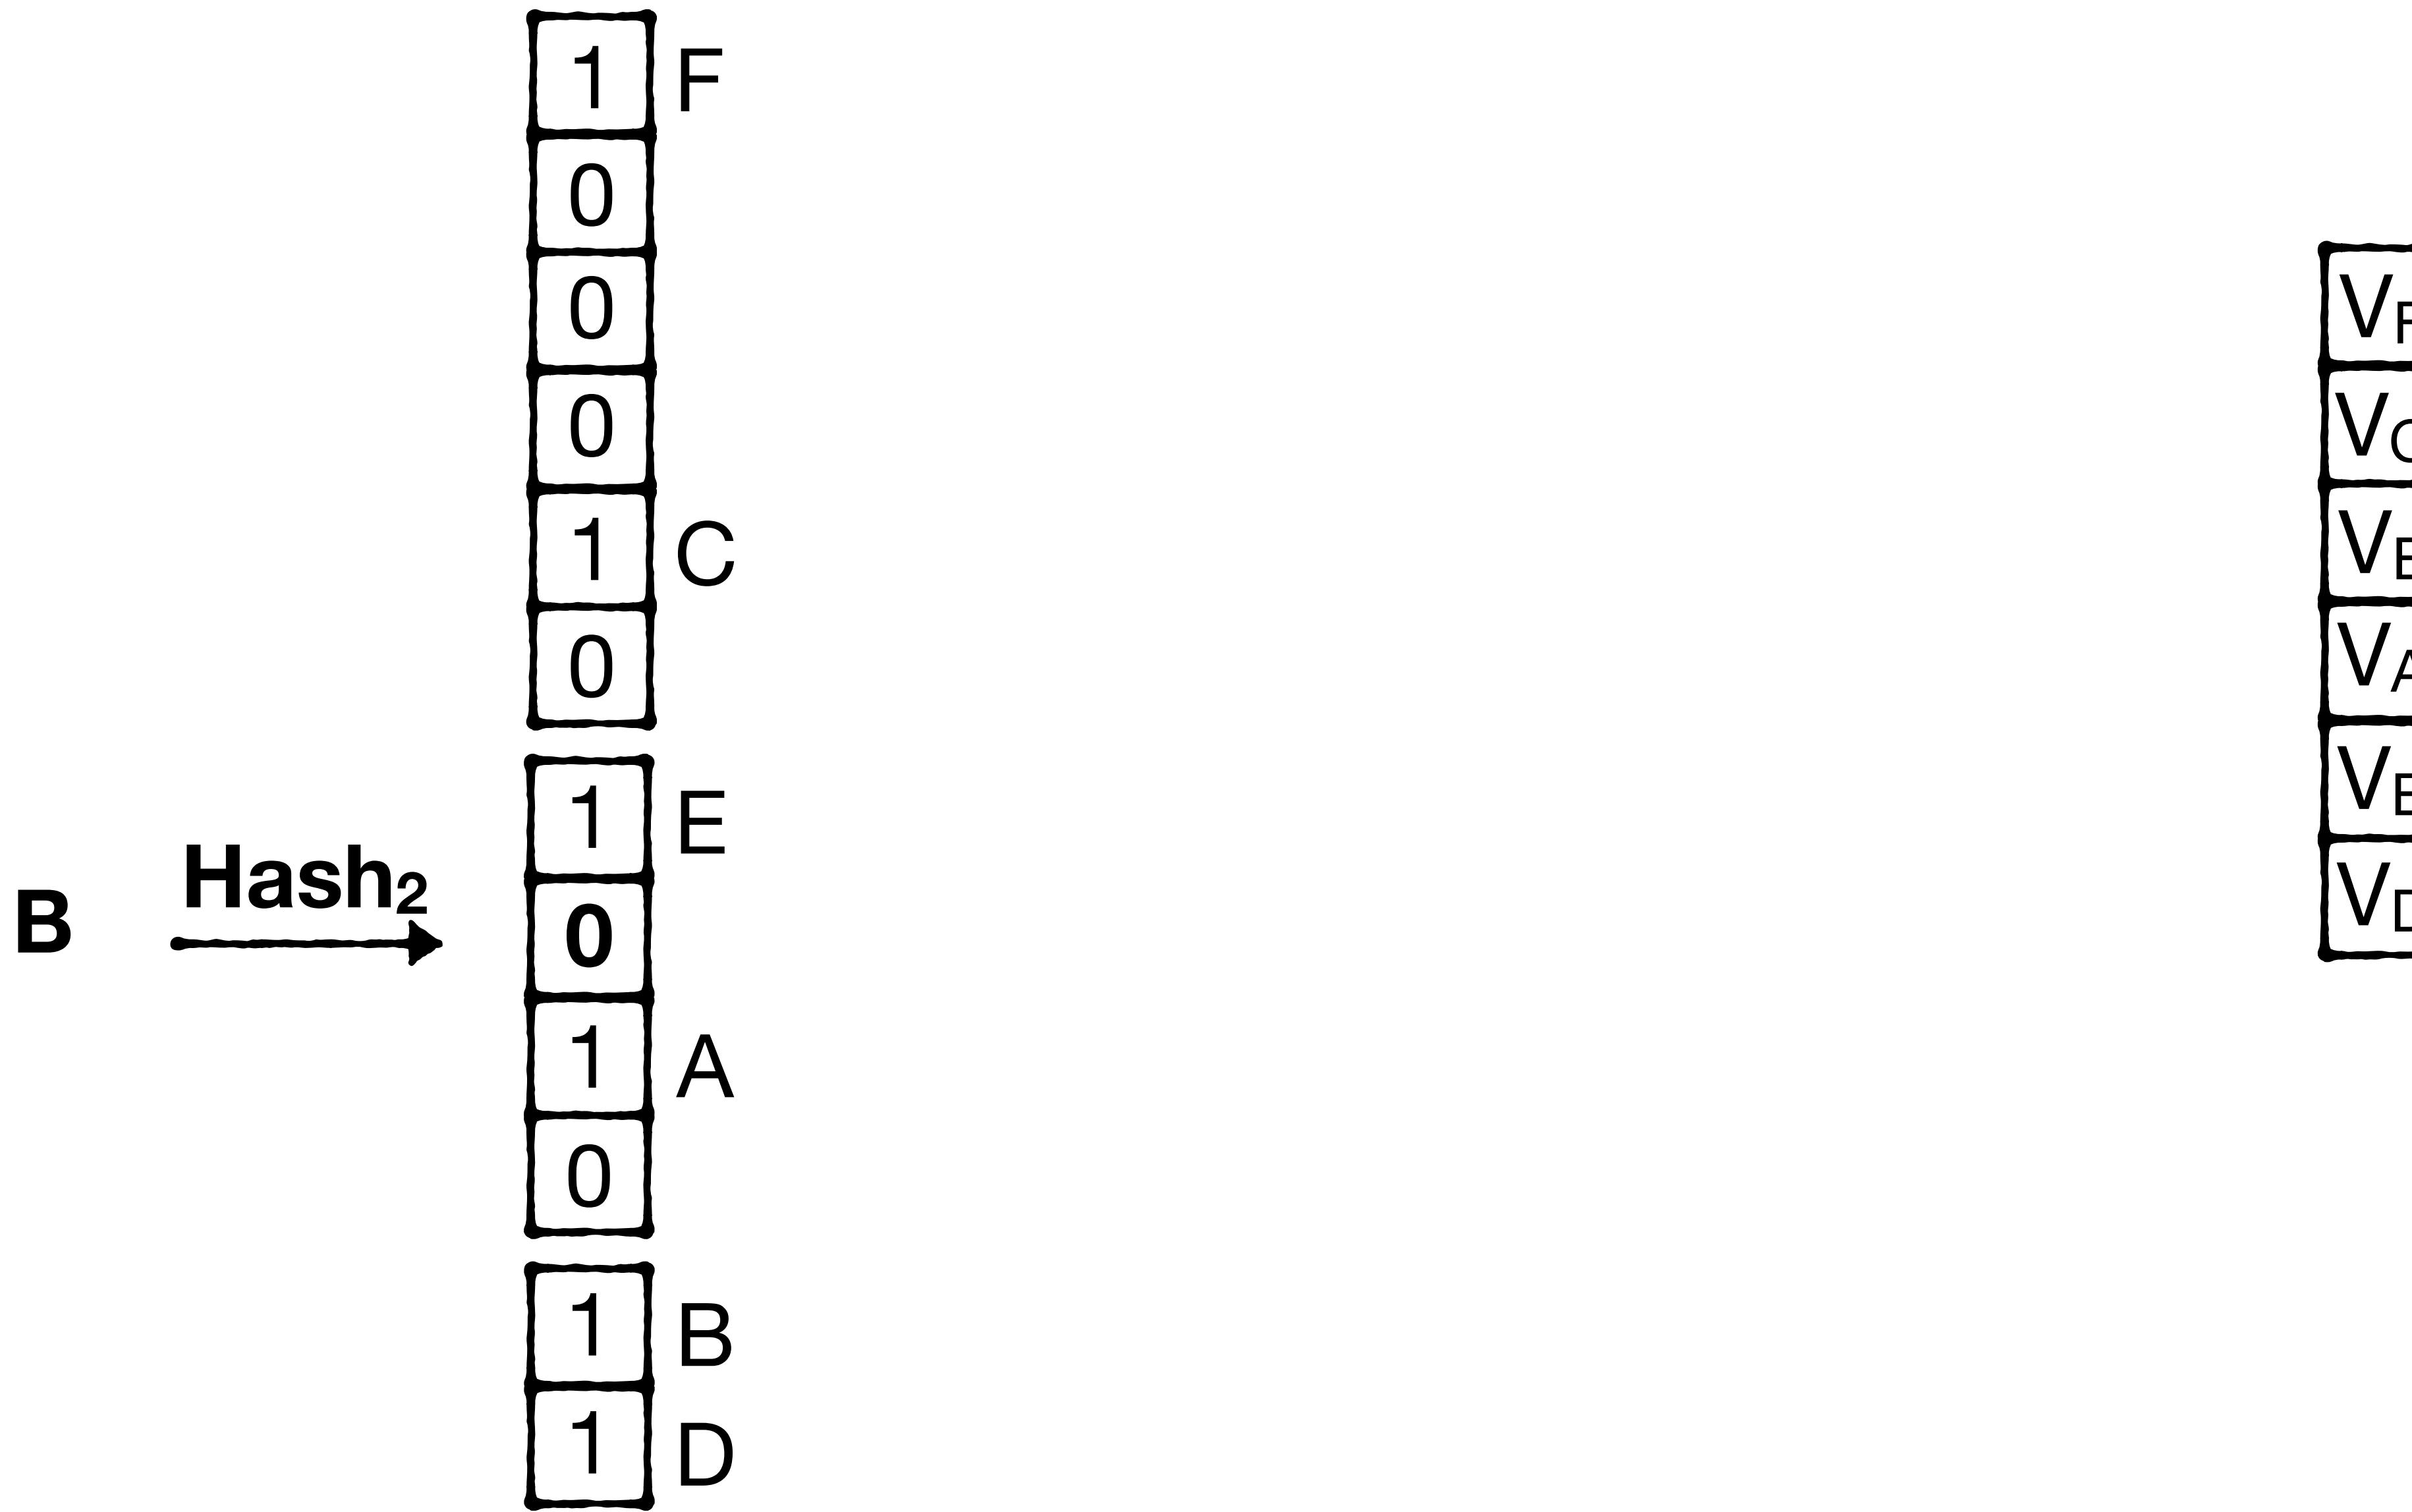
How to query?

Check bitmaps until finding 1



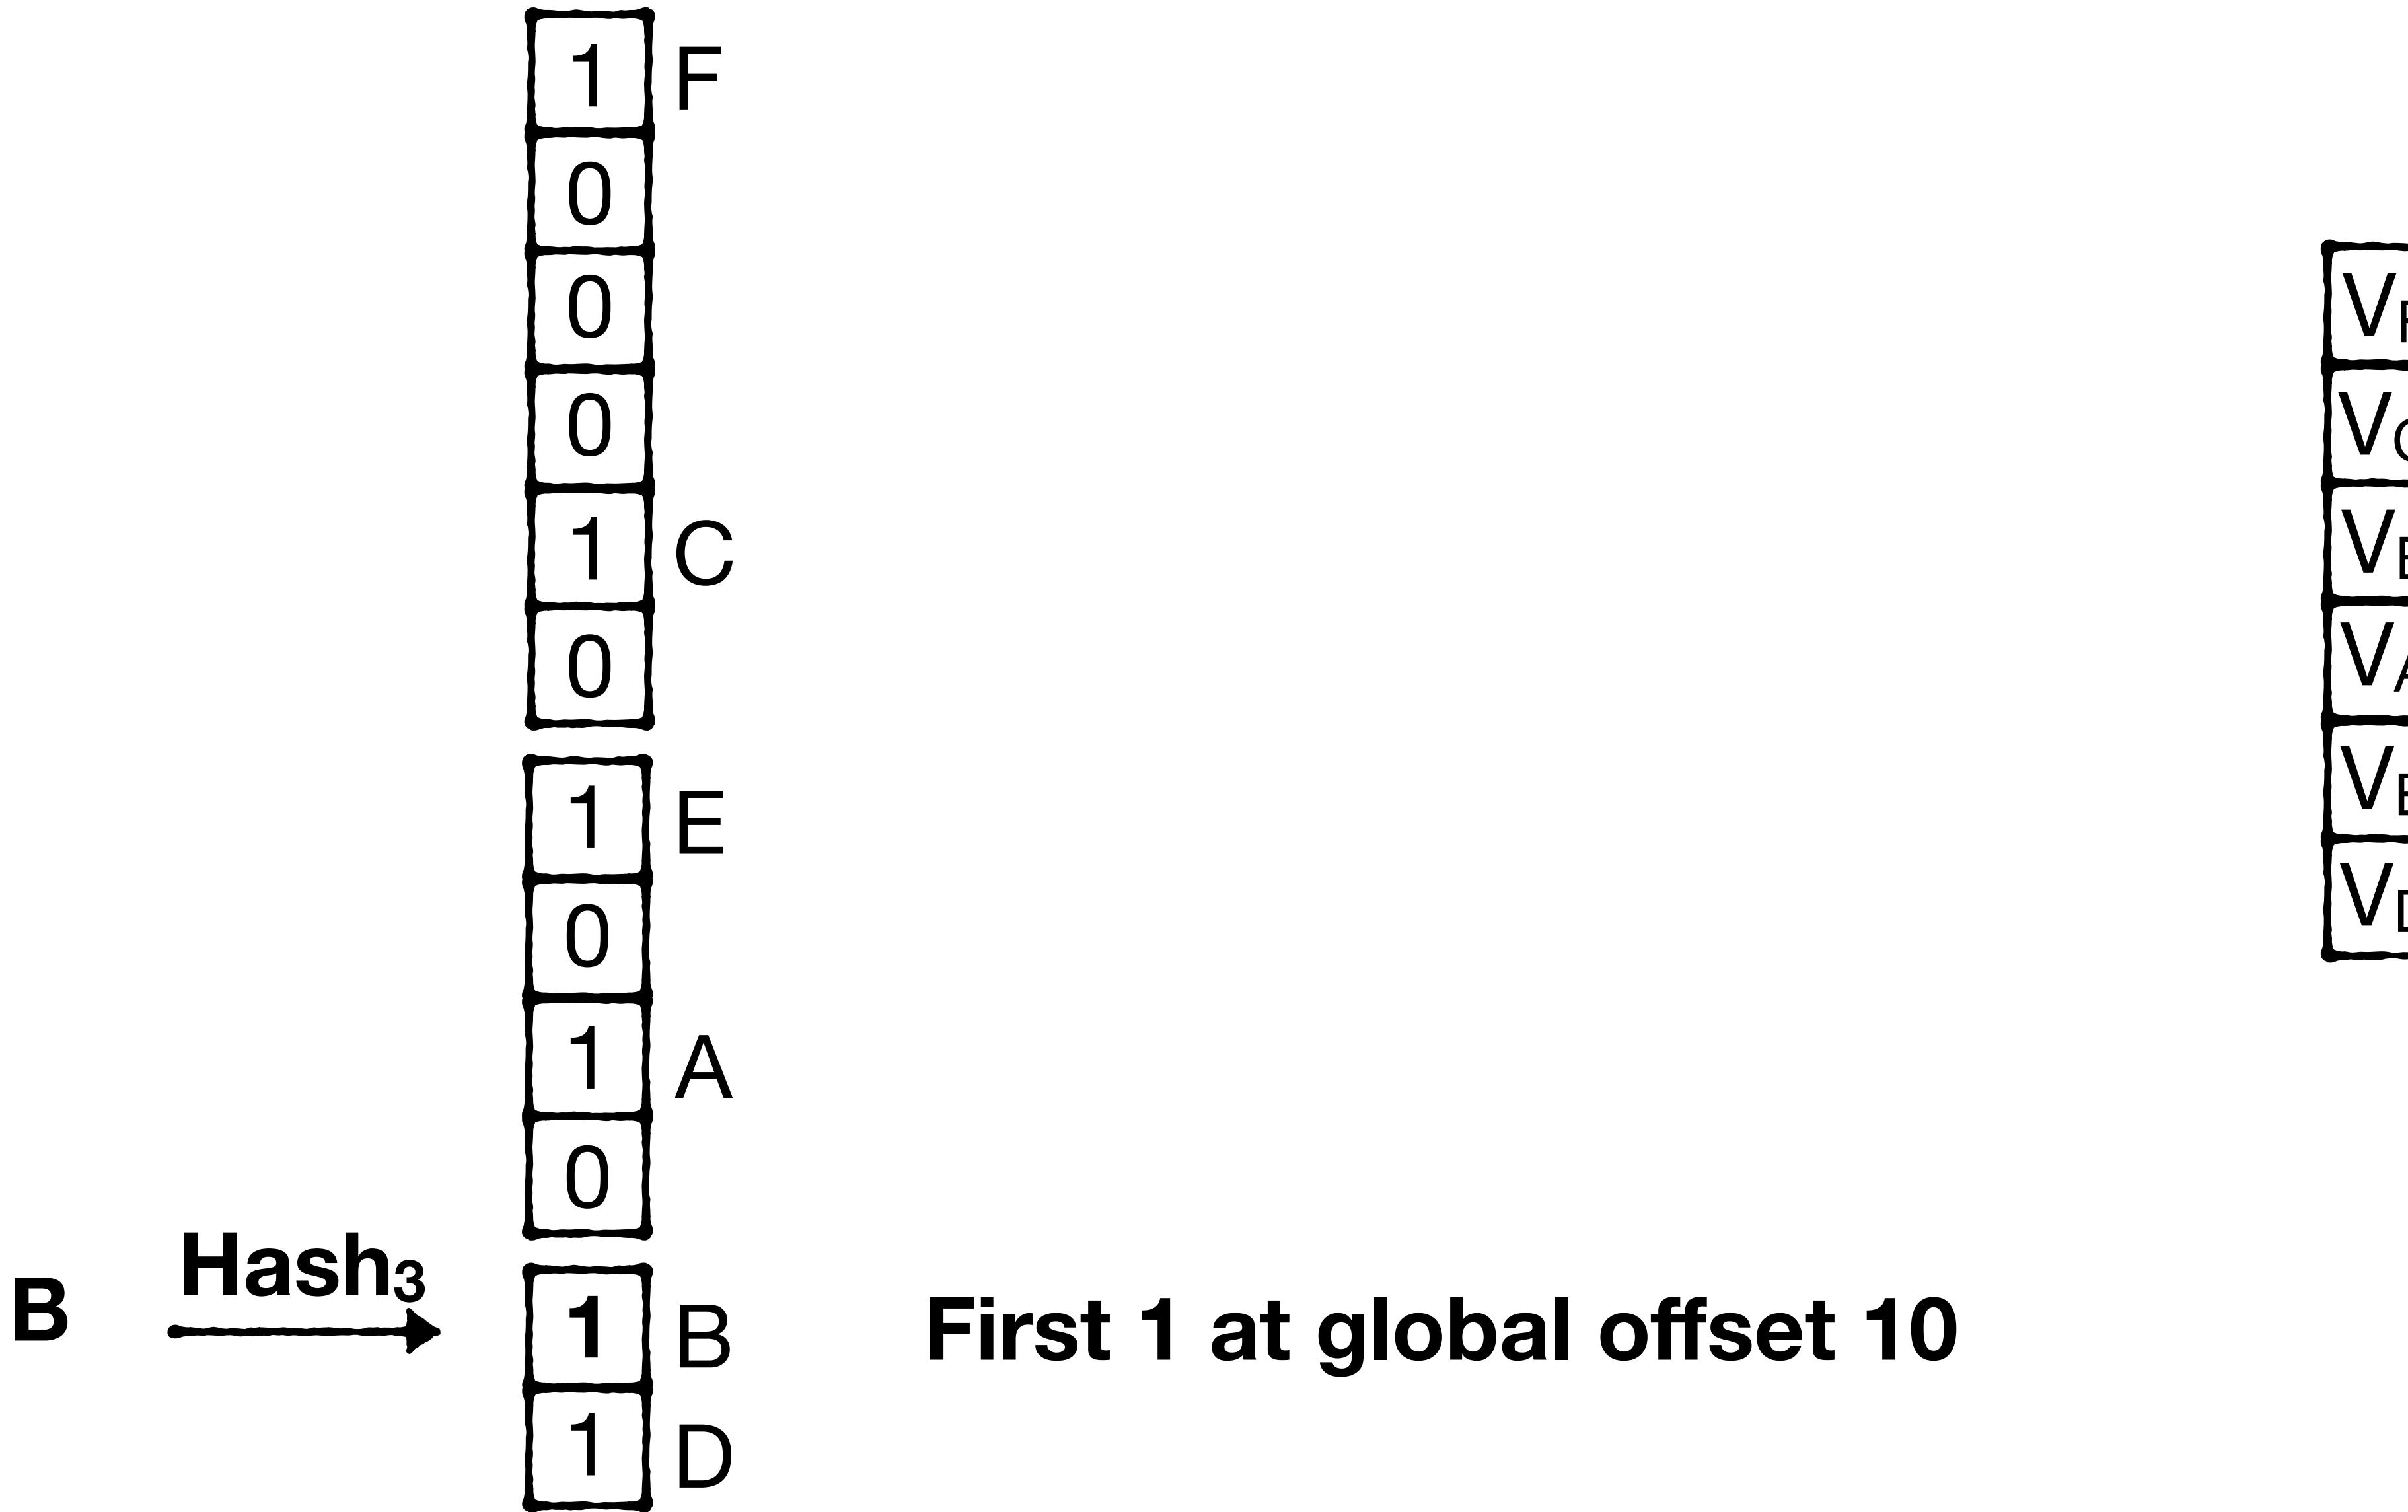
How to query?

Check bitmaps until finding 1



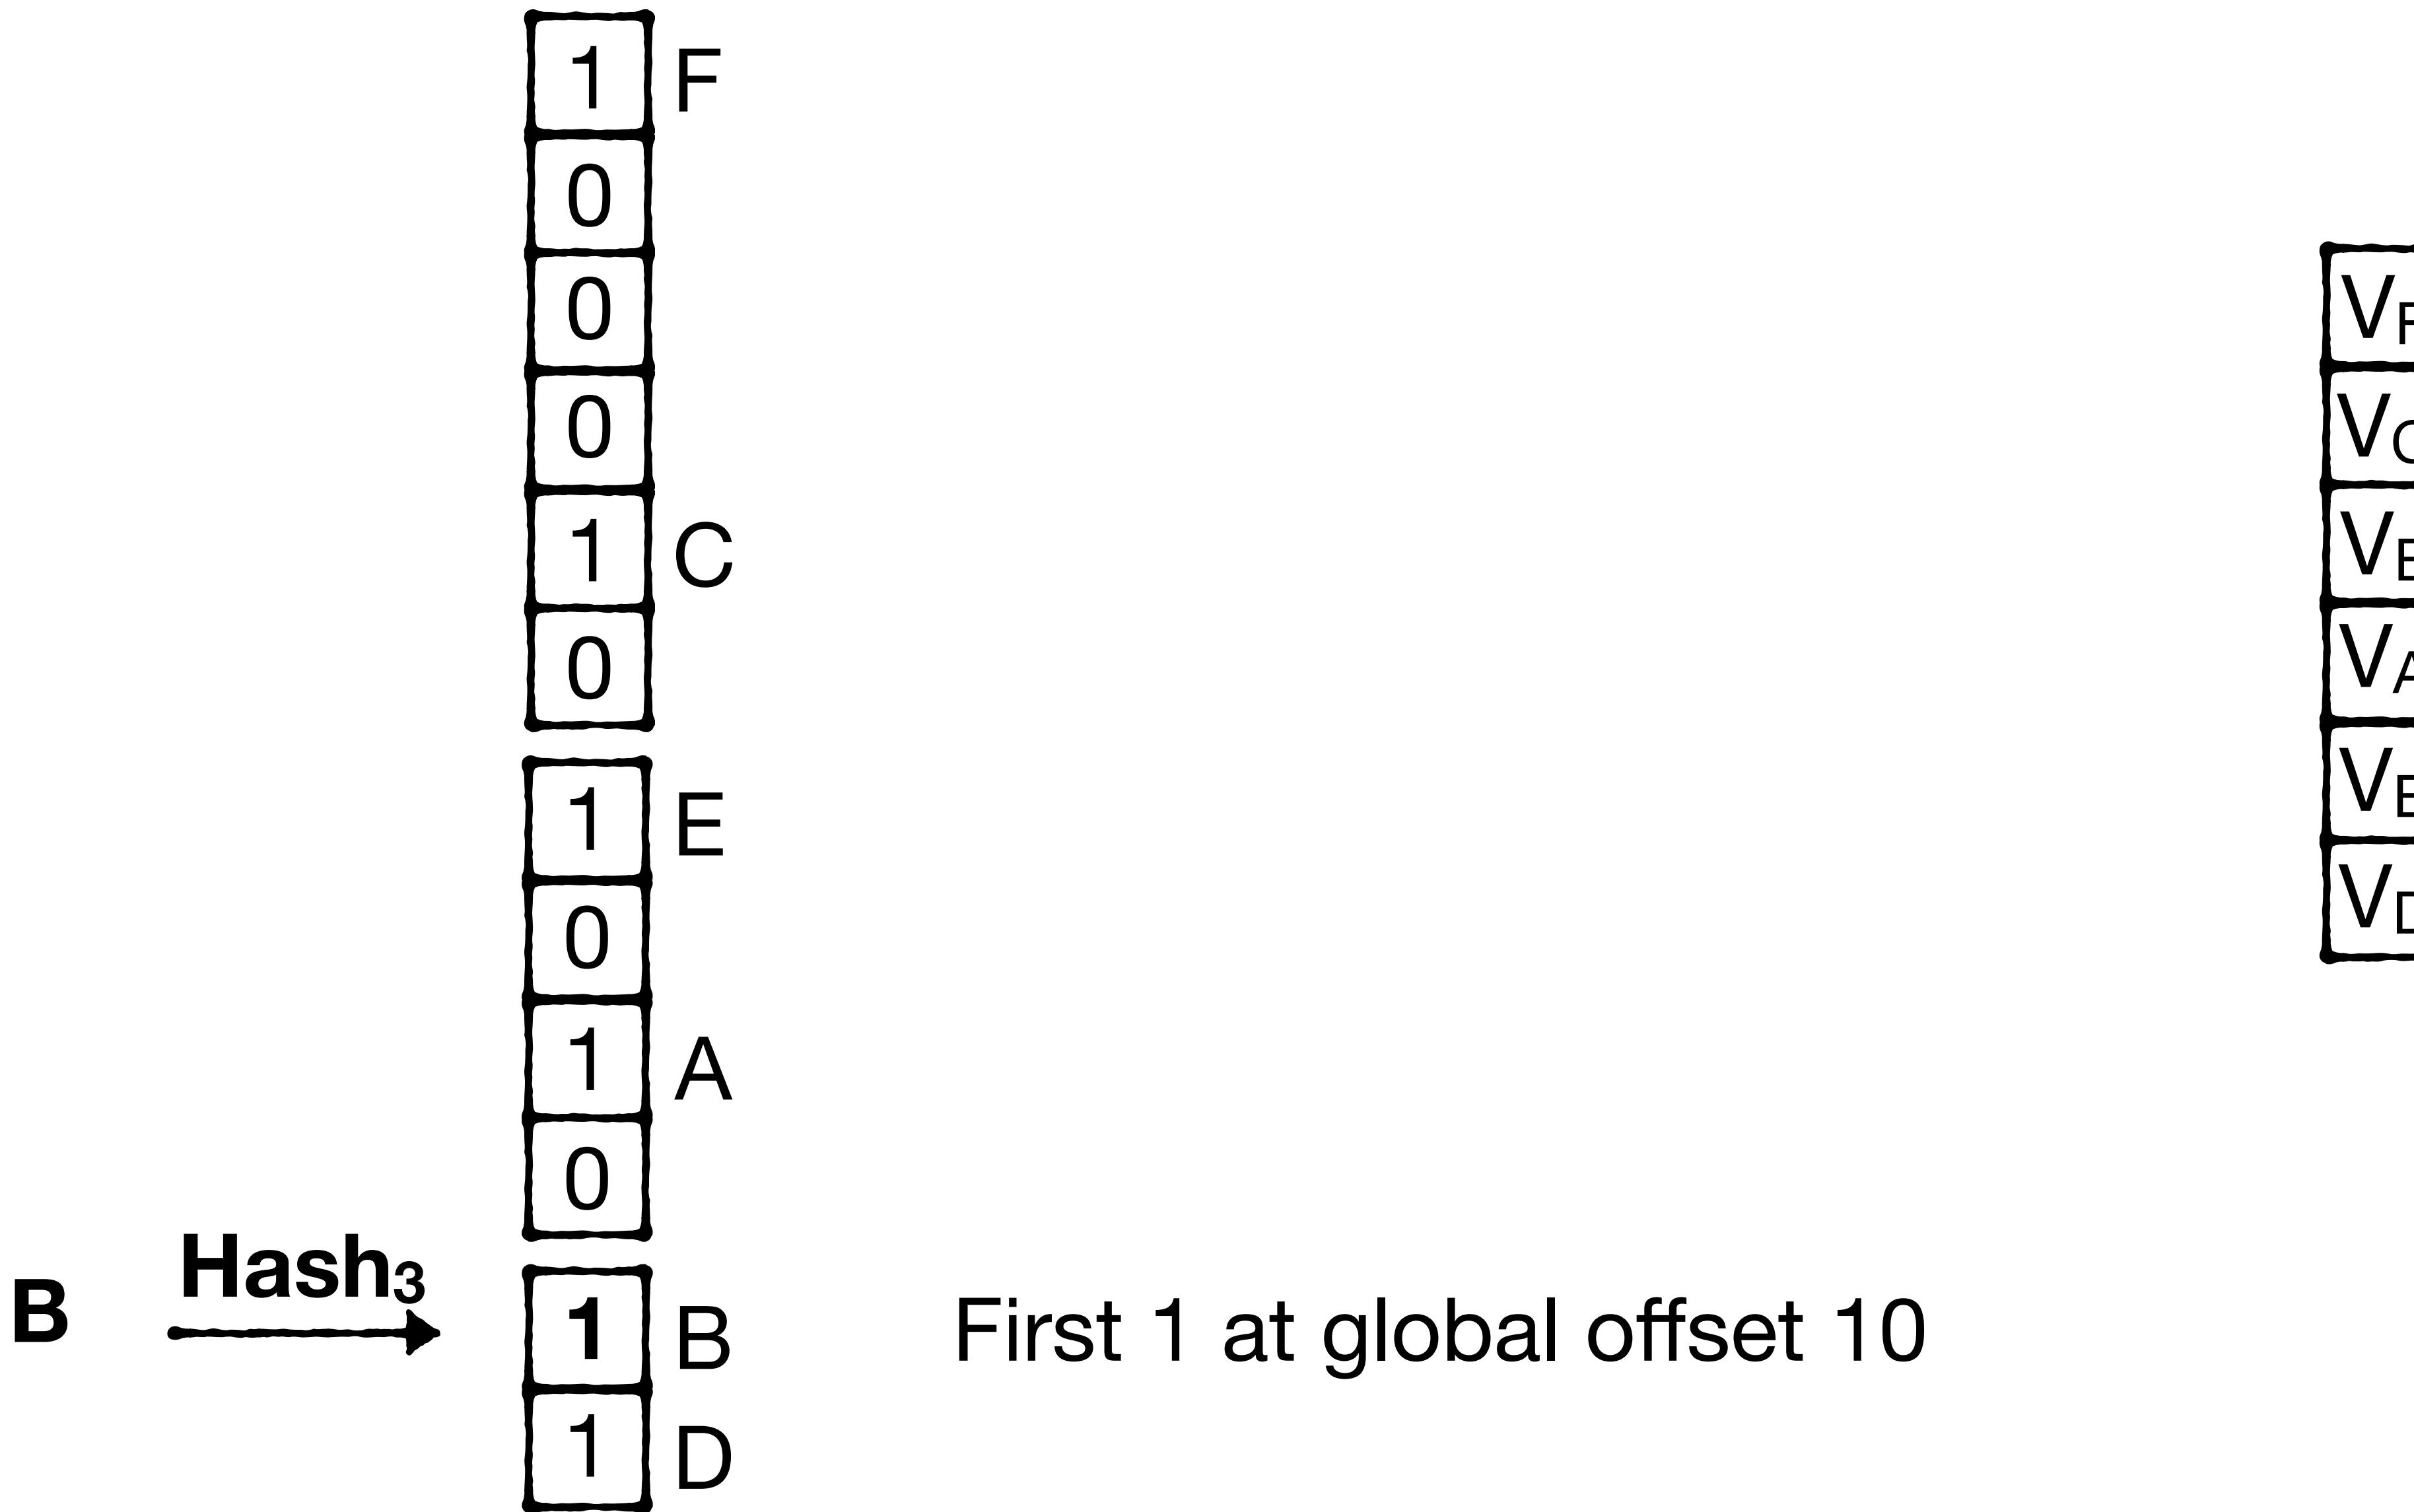
How to query?

Check bitmaps until finding 1



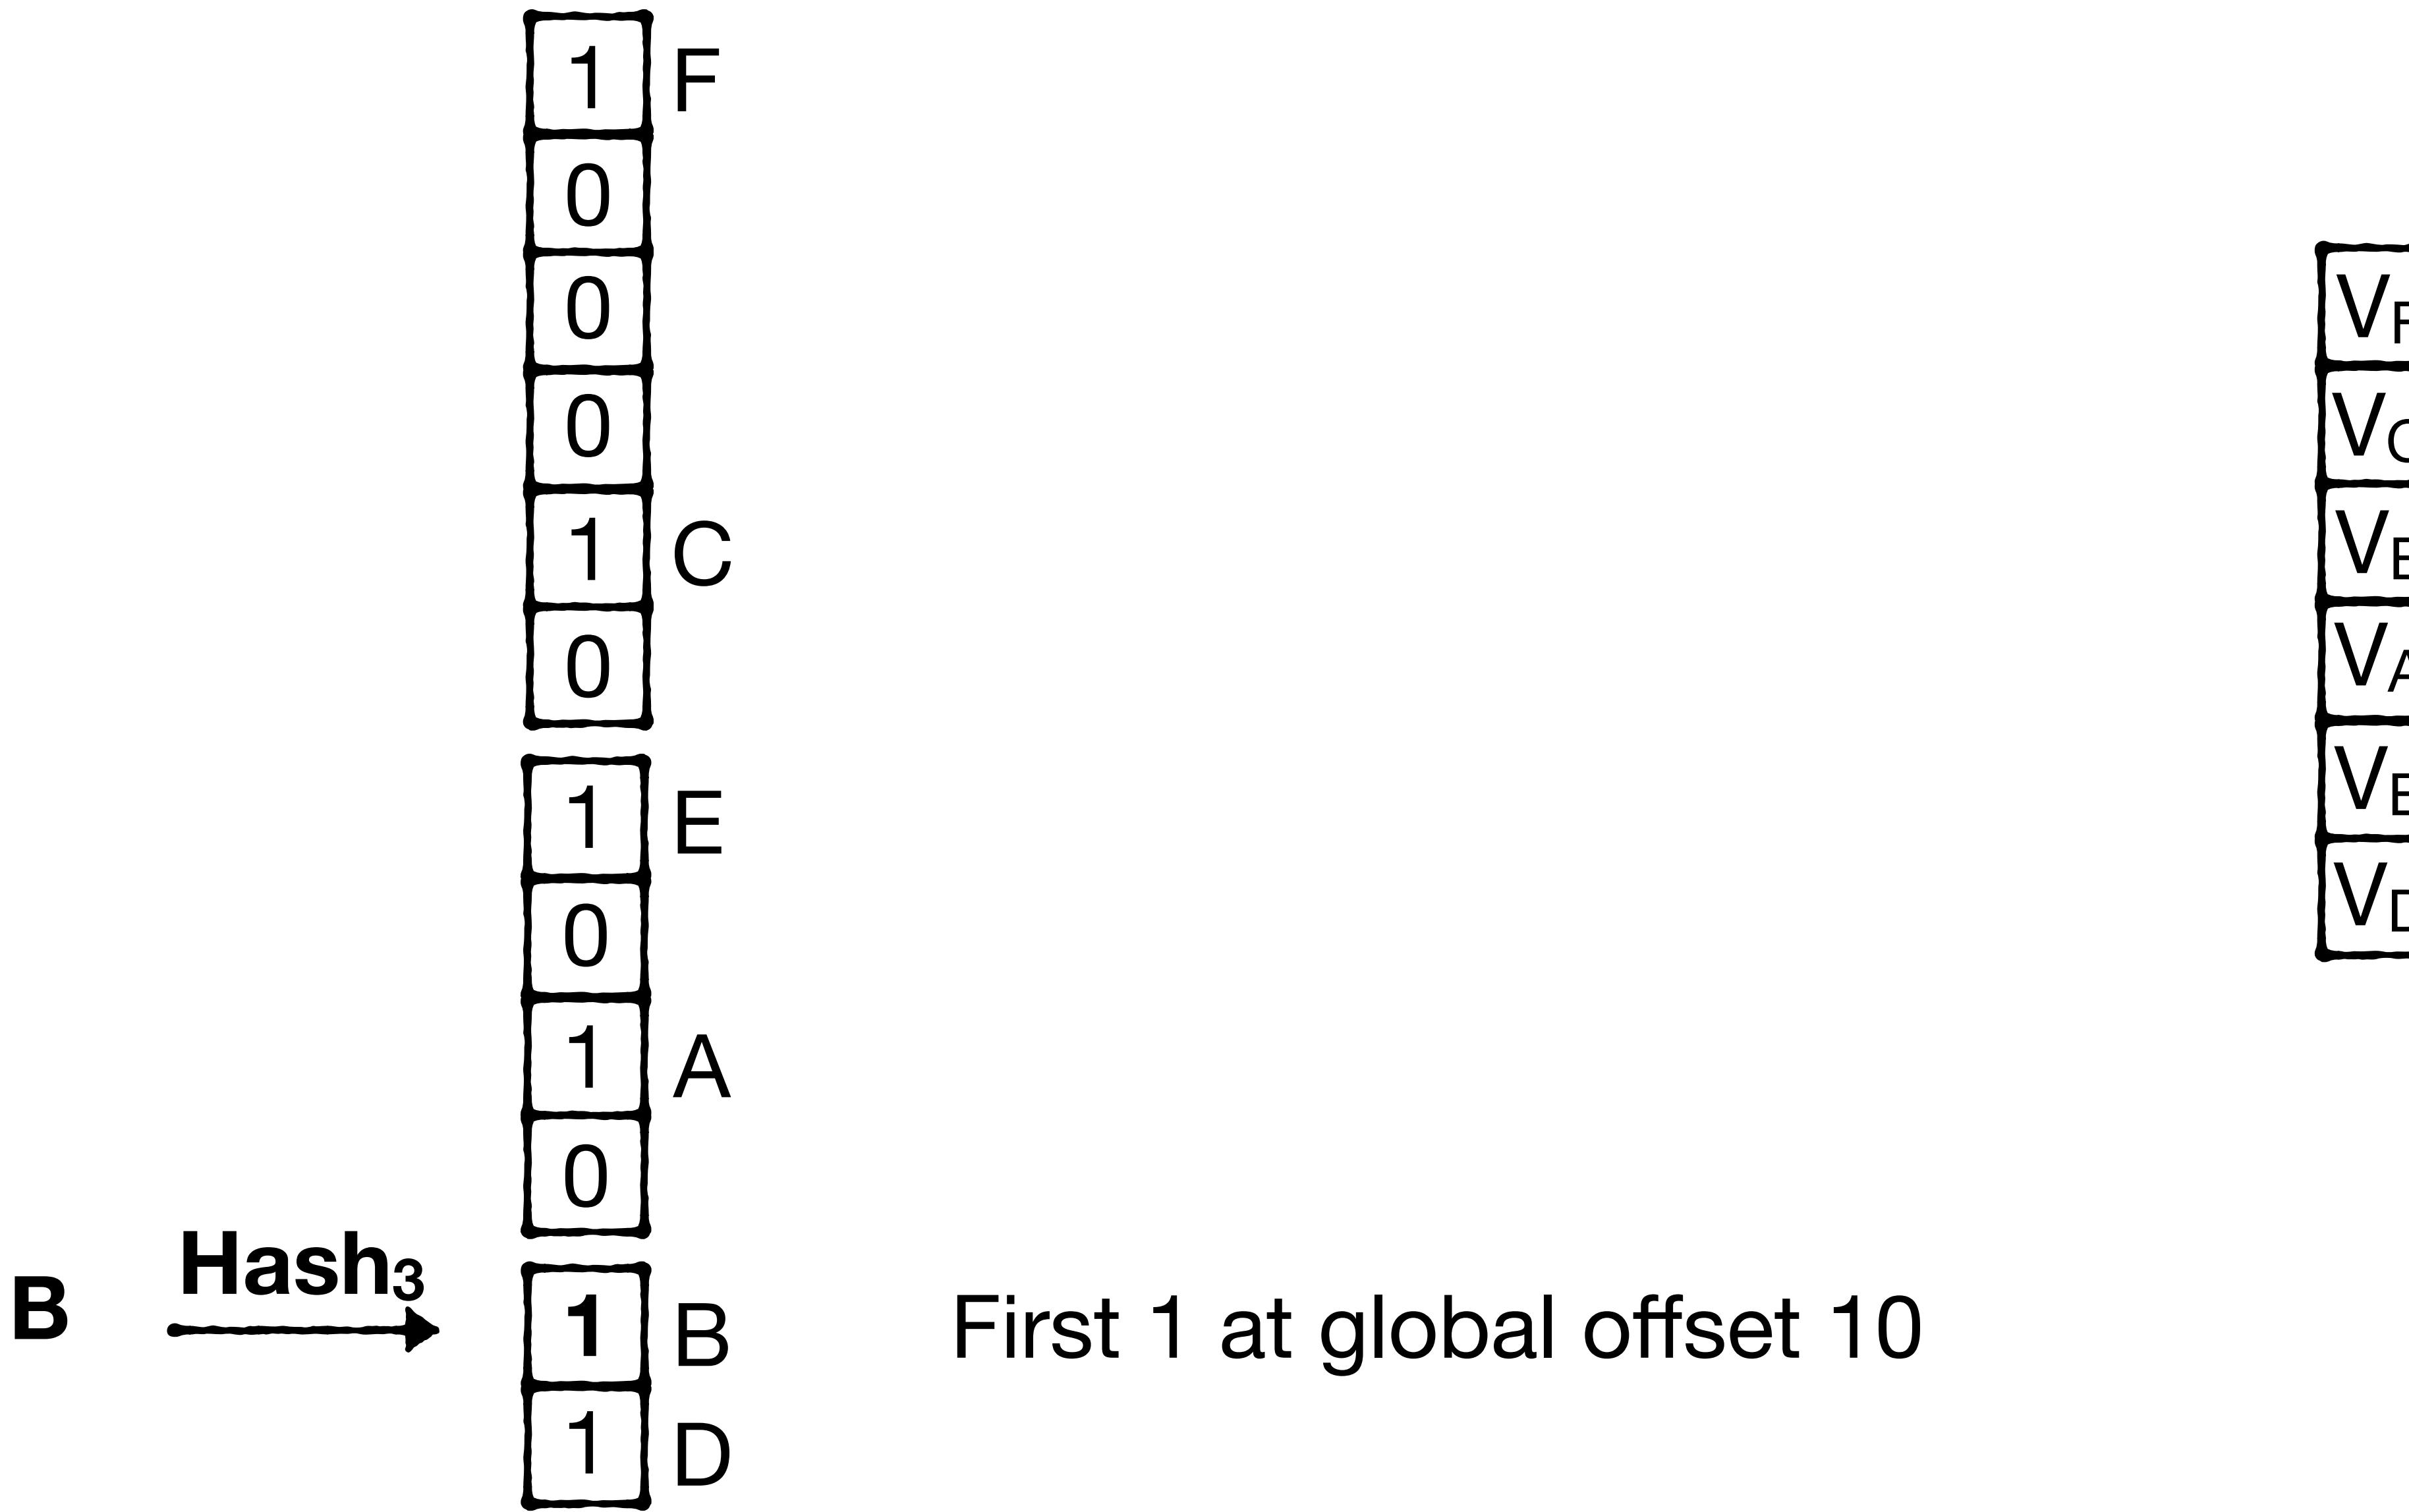
How to query?

Check bitmaps until finding 1
Next? :)



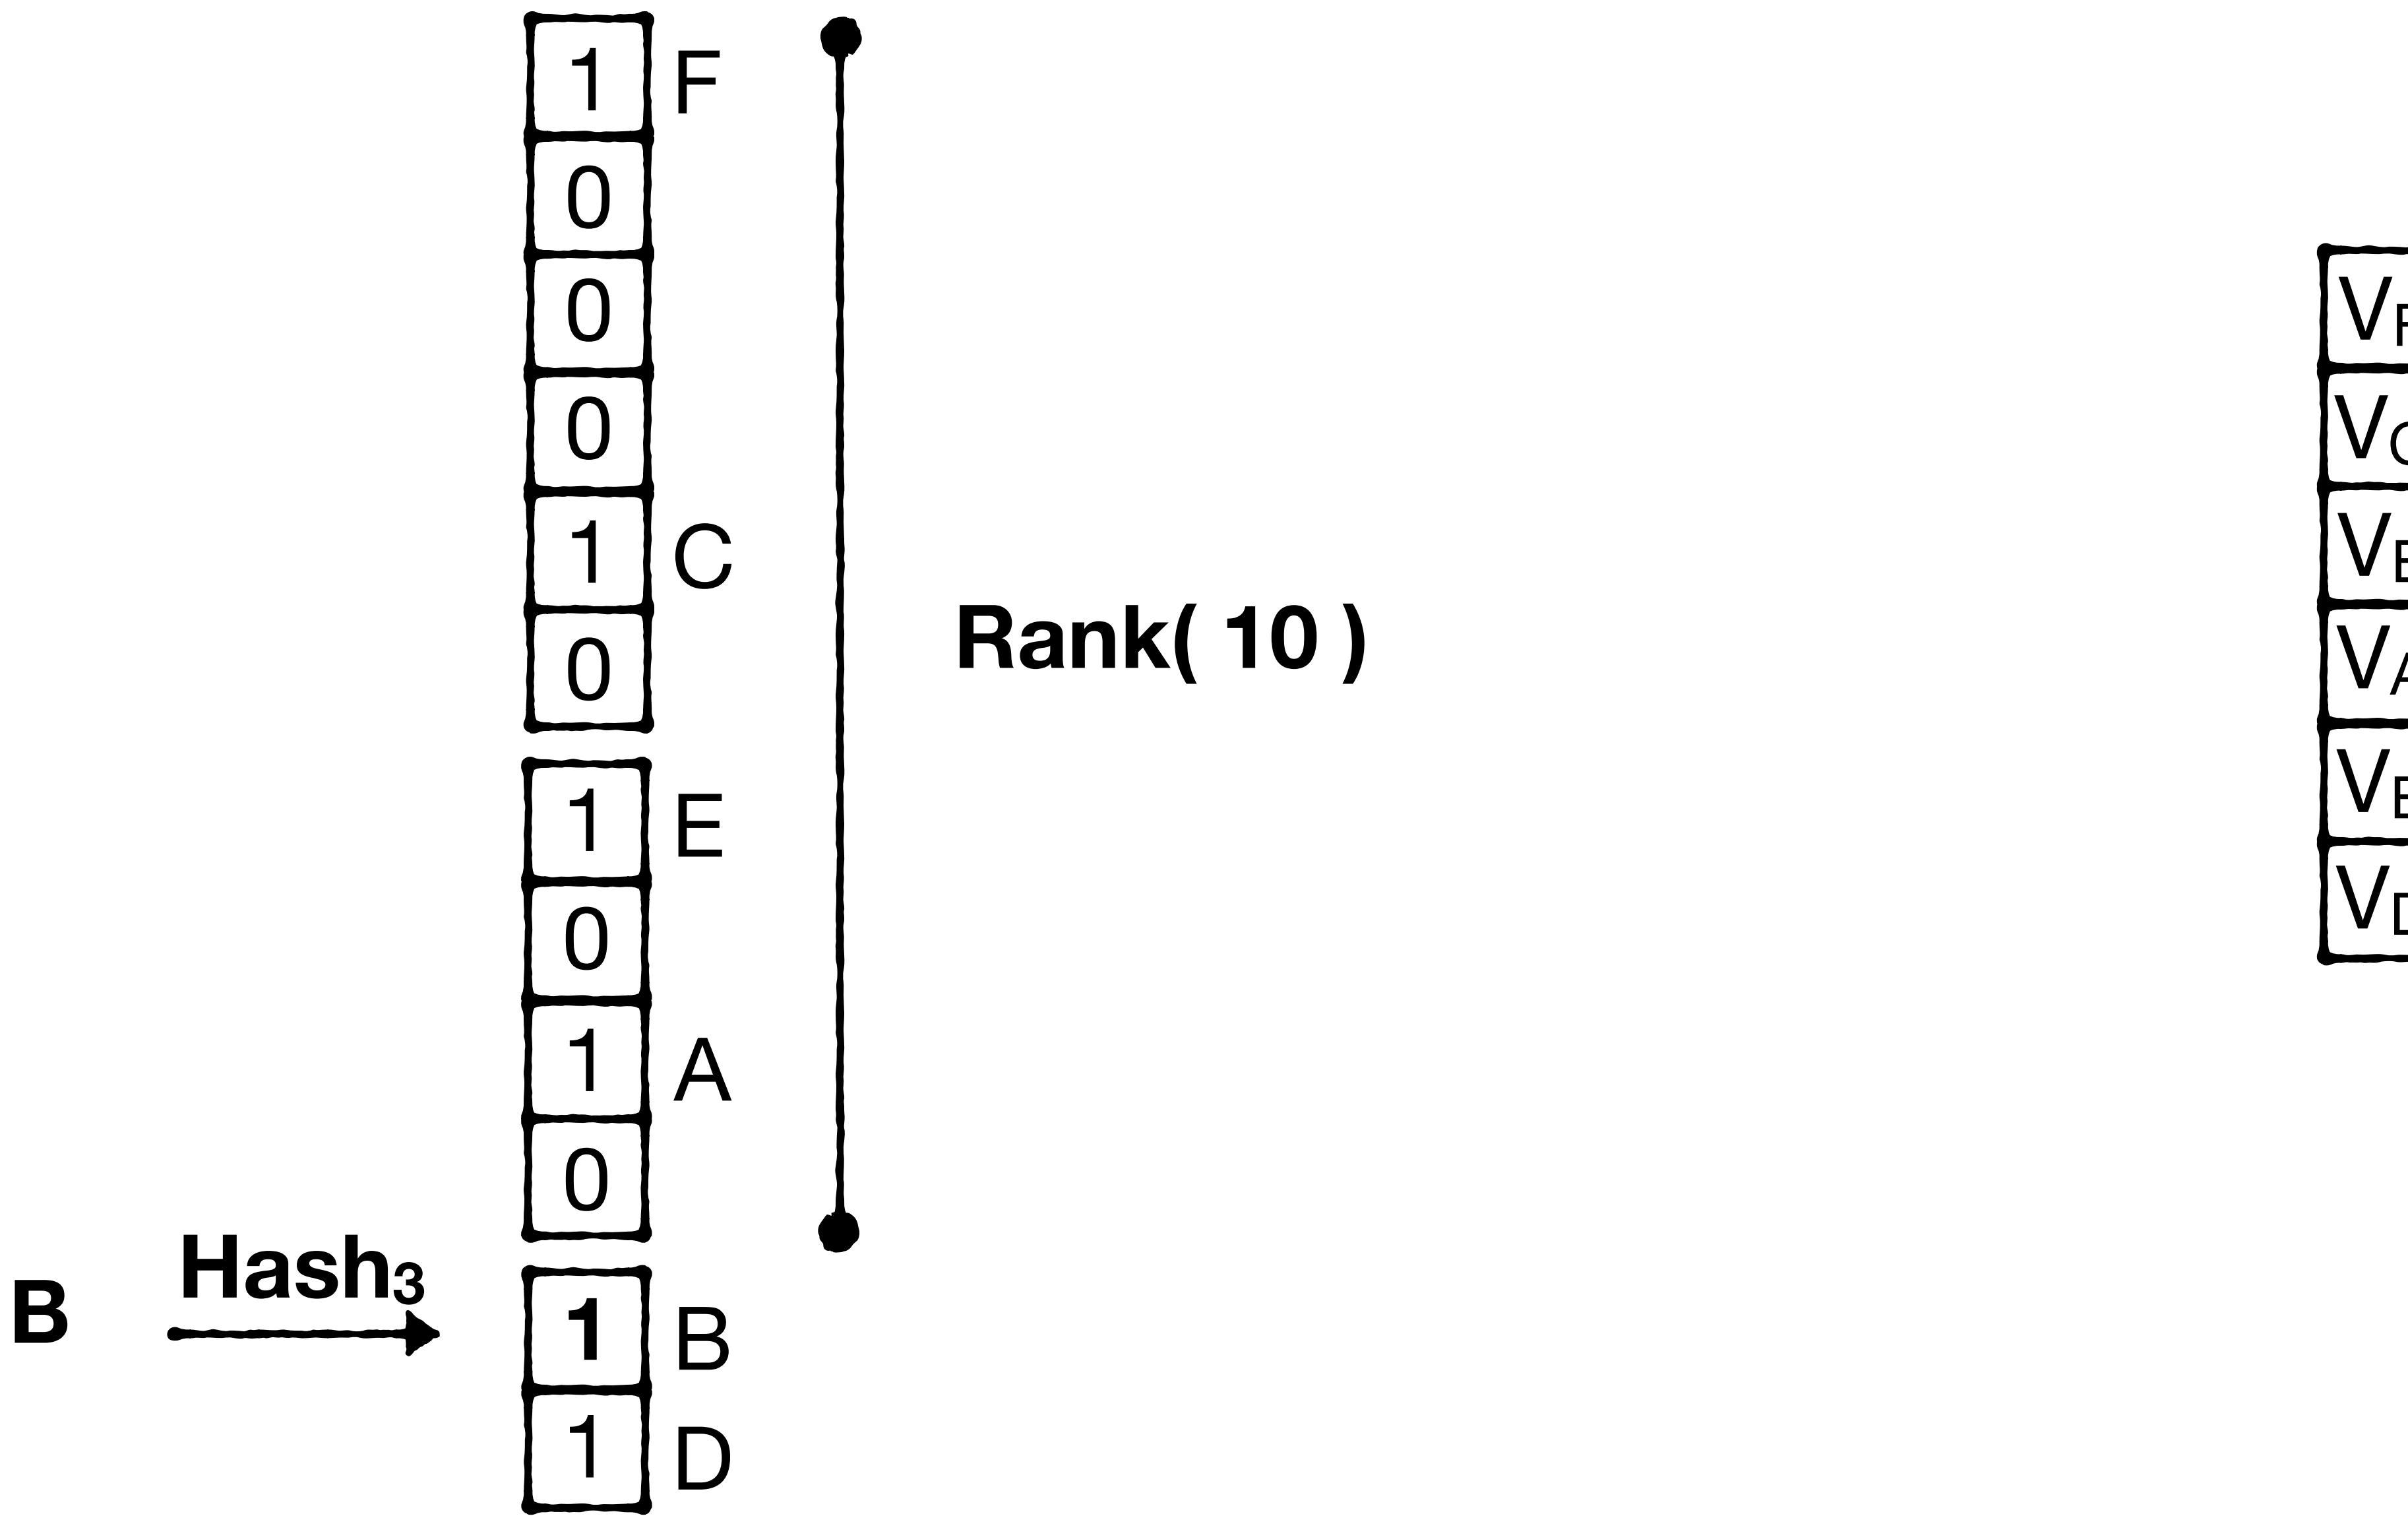
How to query?

Check bitmaps until finding 1
Rank on offset of first 1



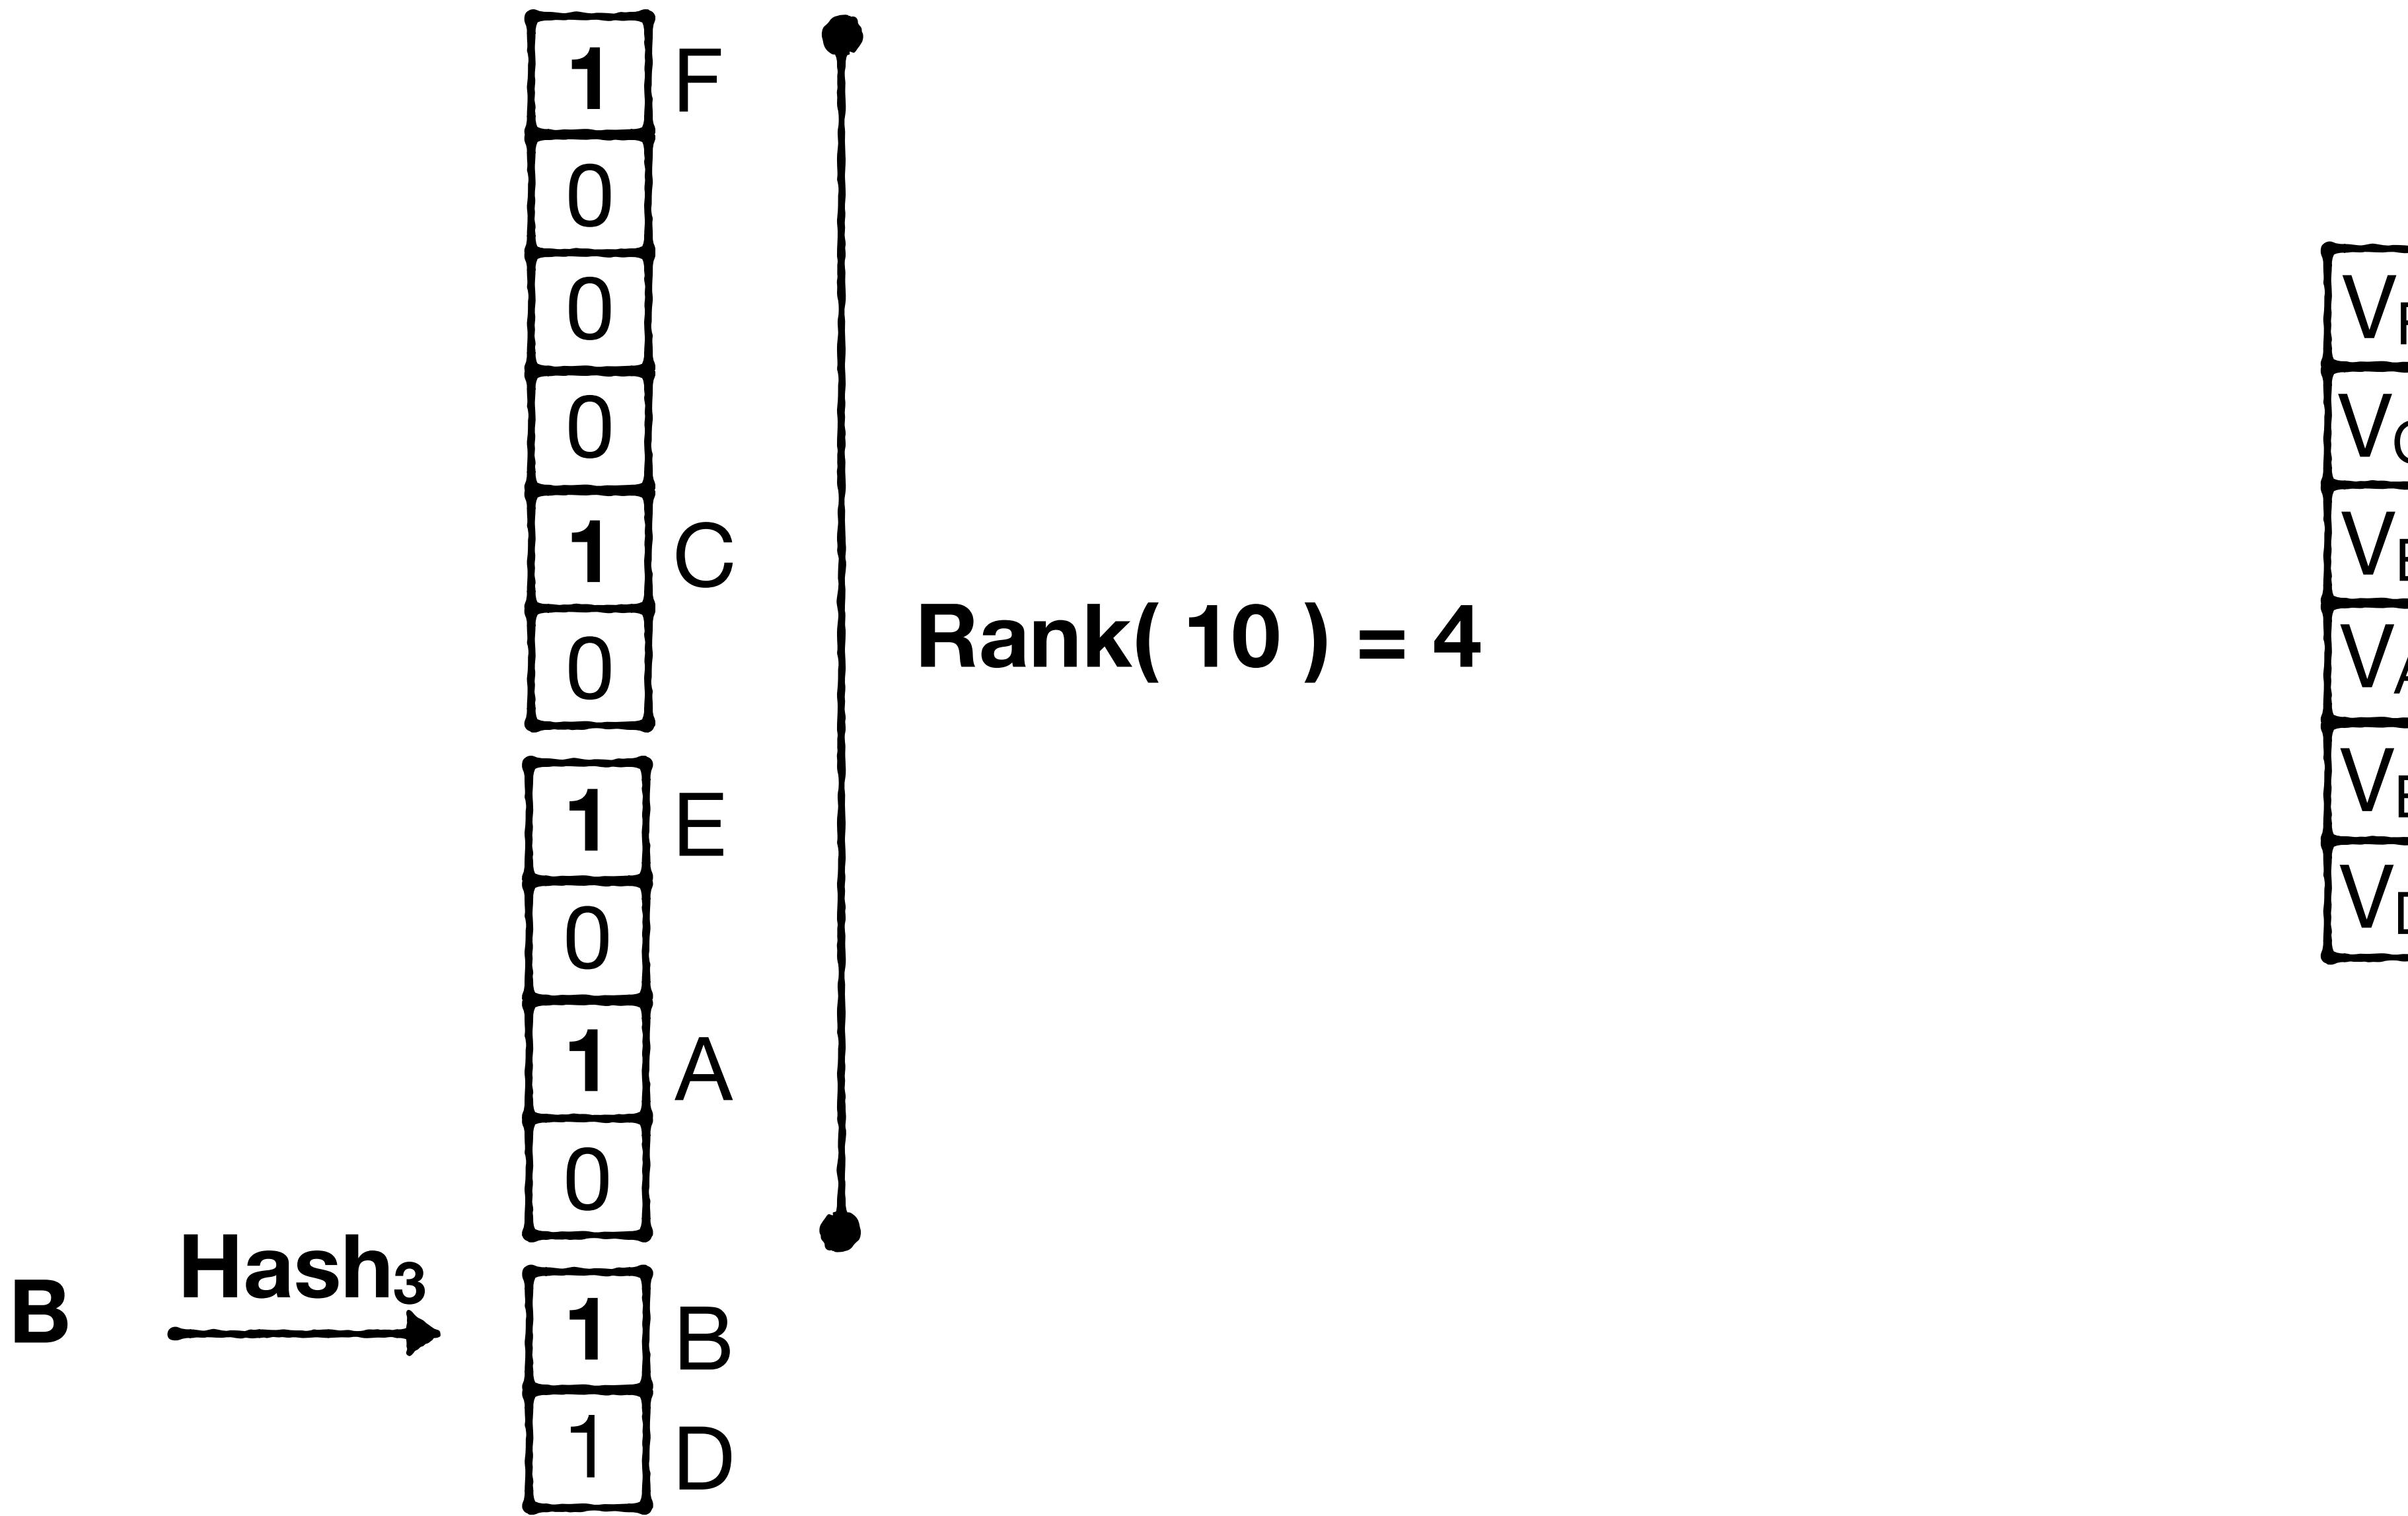
How to query?

Check bitmaps until finding 1
Rank on offset of first 1



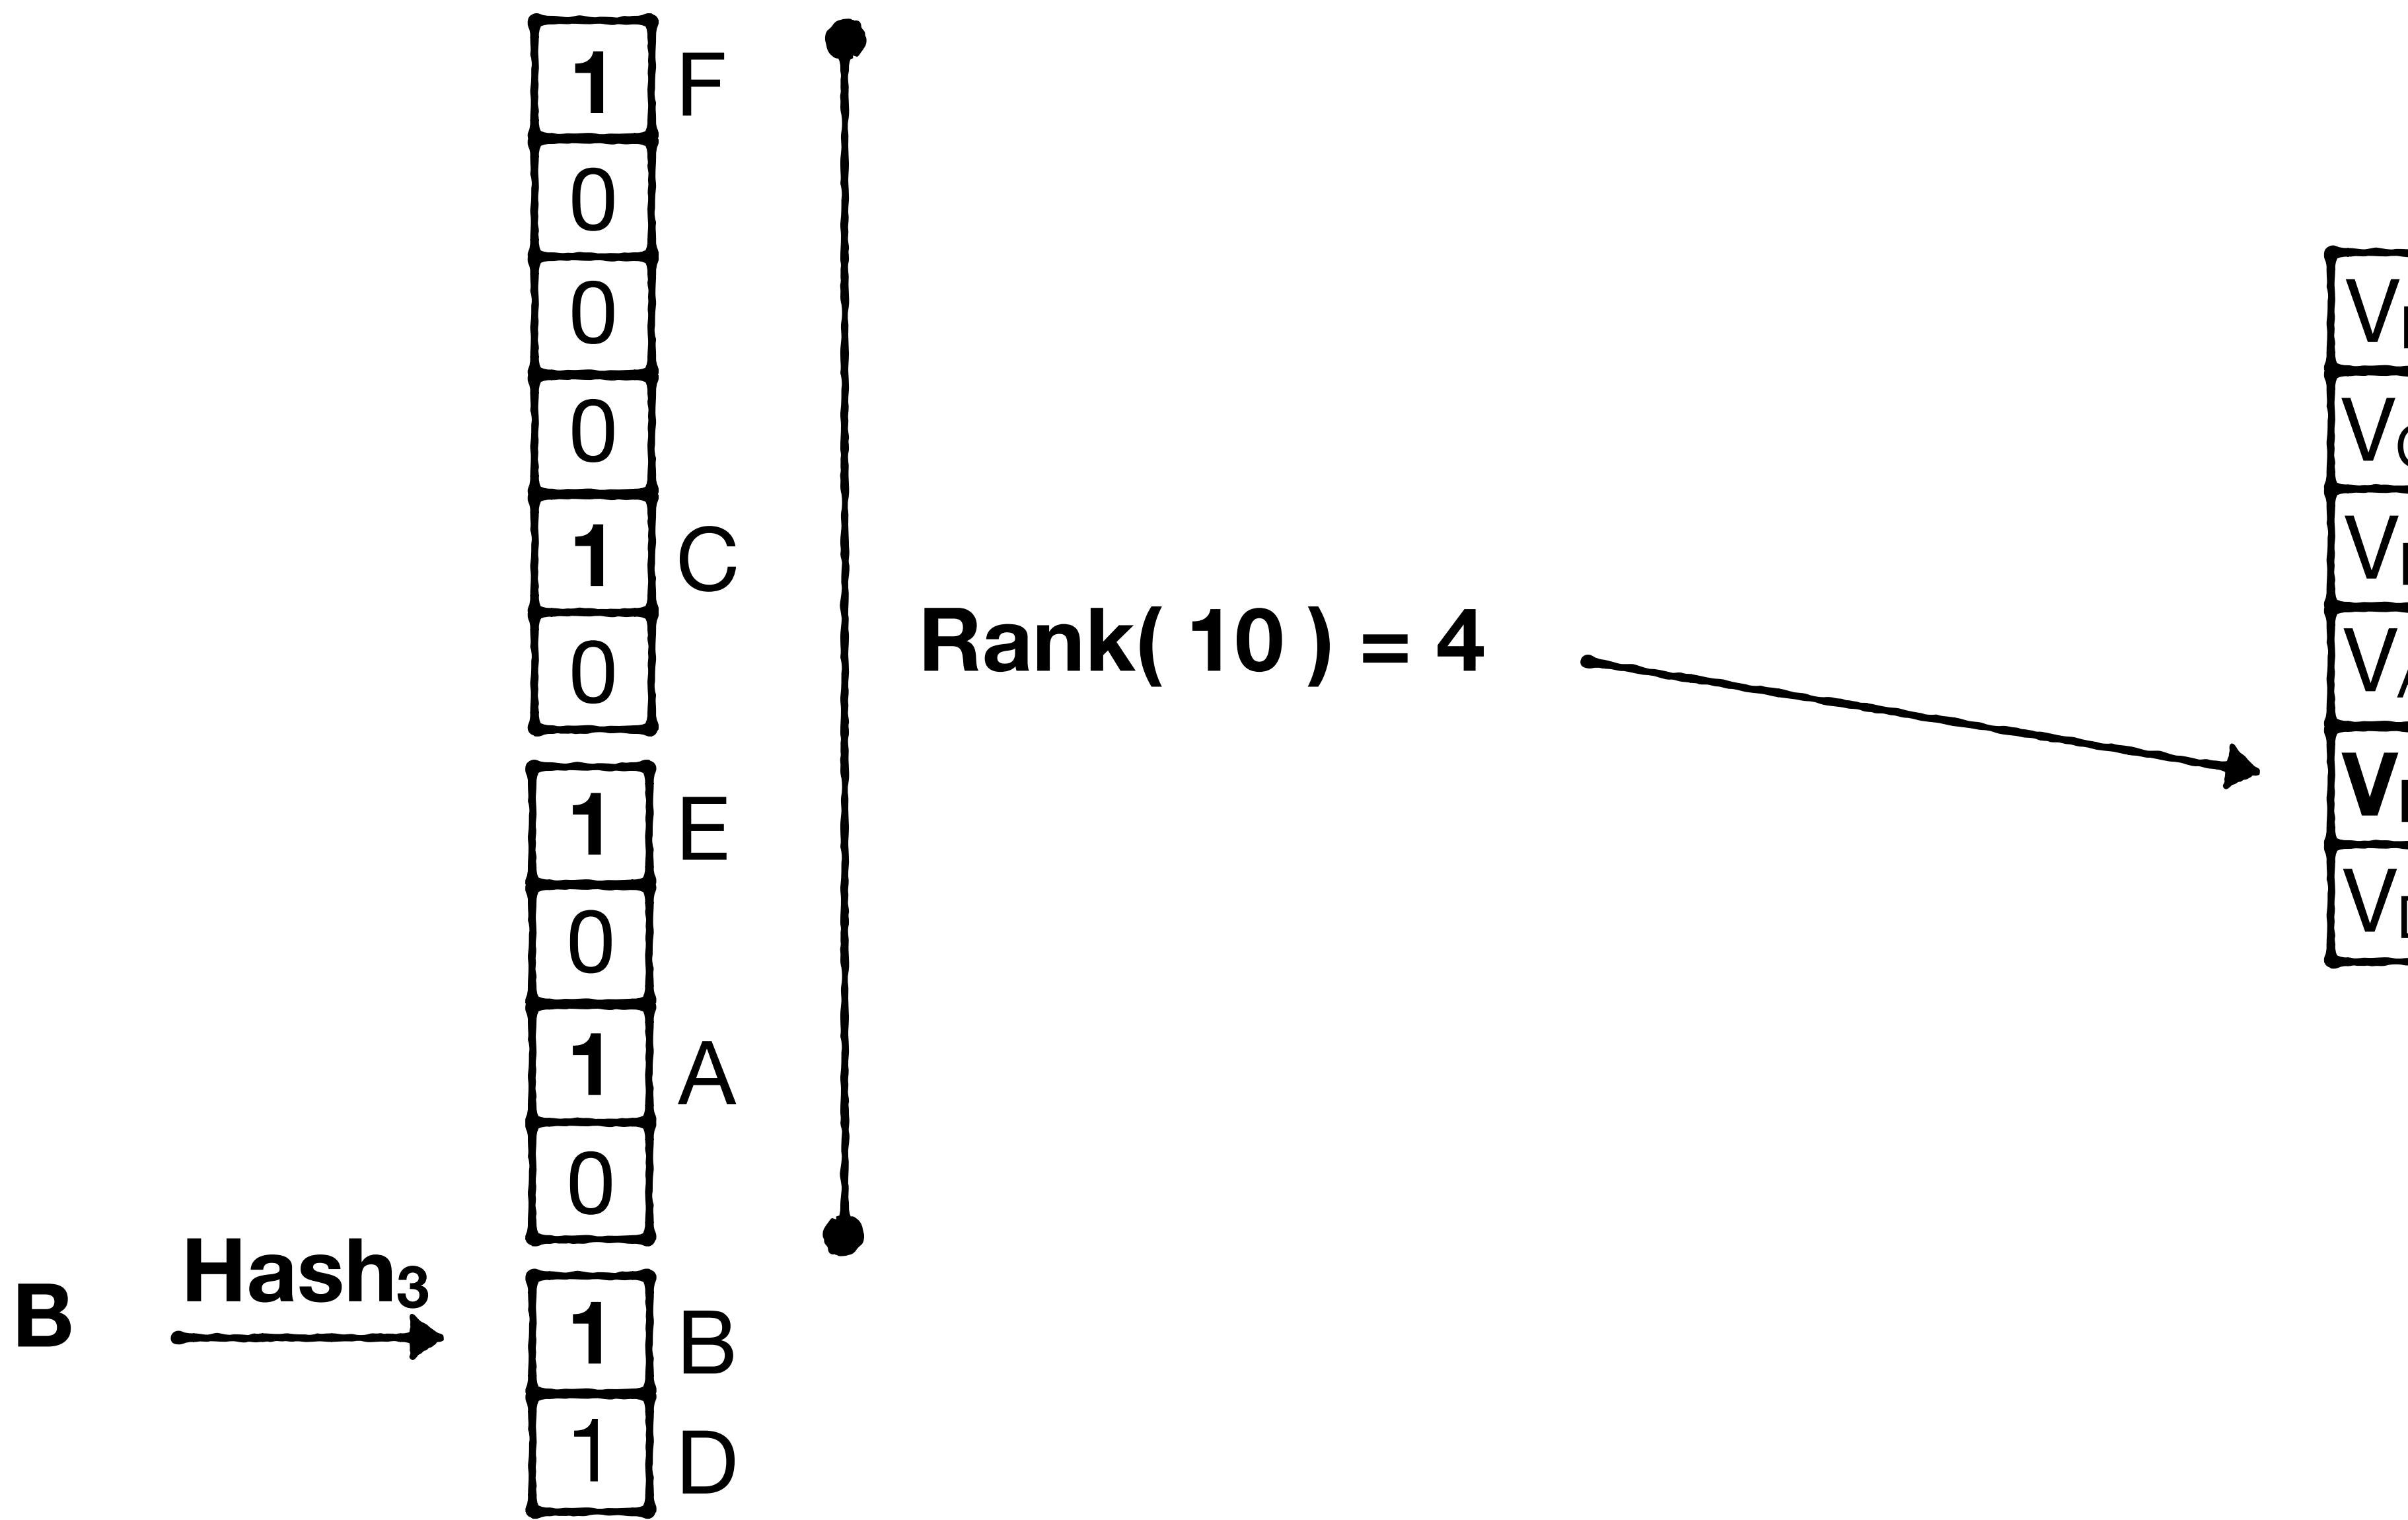
How to query?

Check bitmaps until finding 1
Rank on offset of first 1



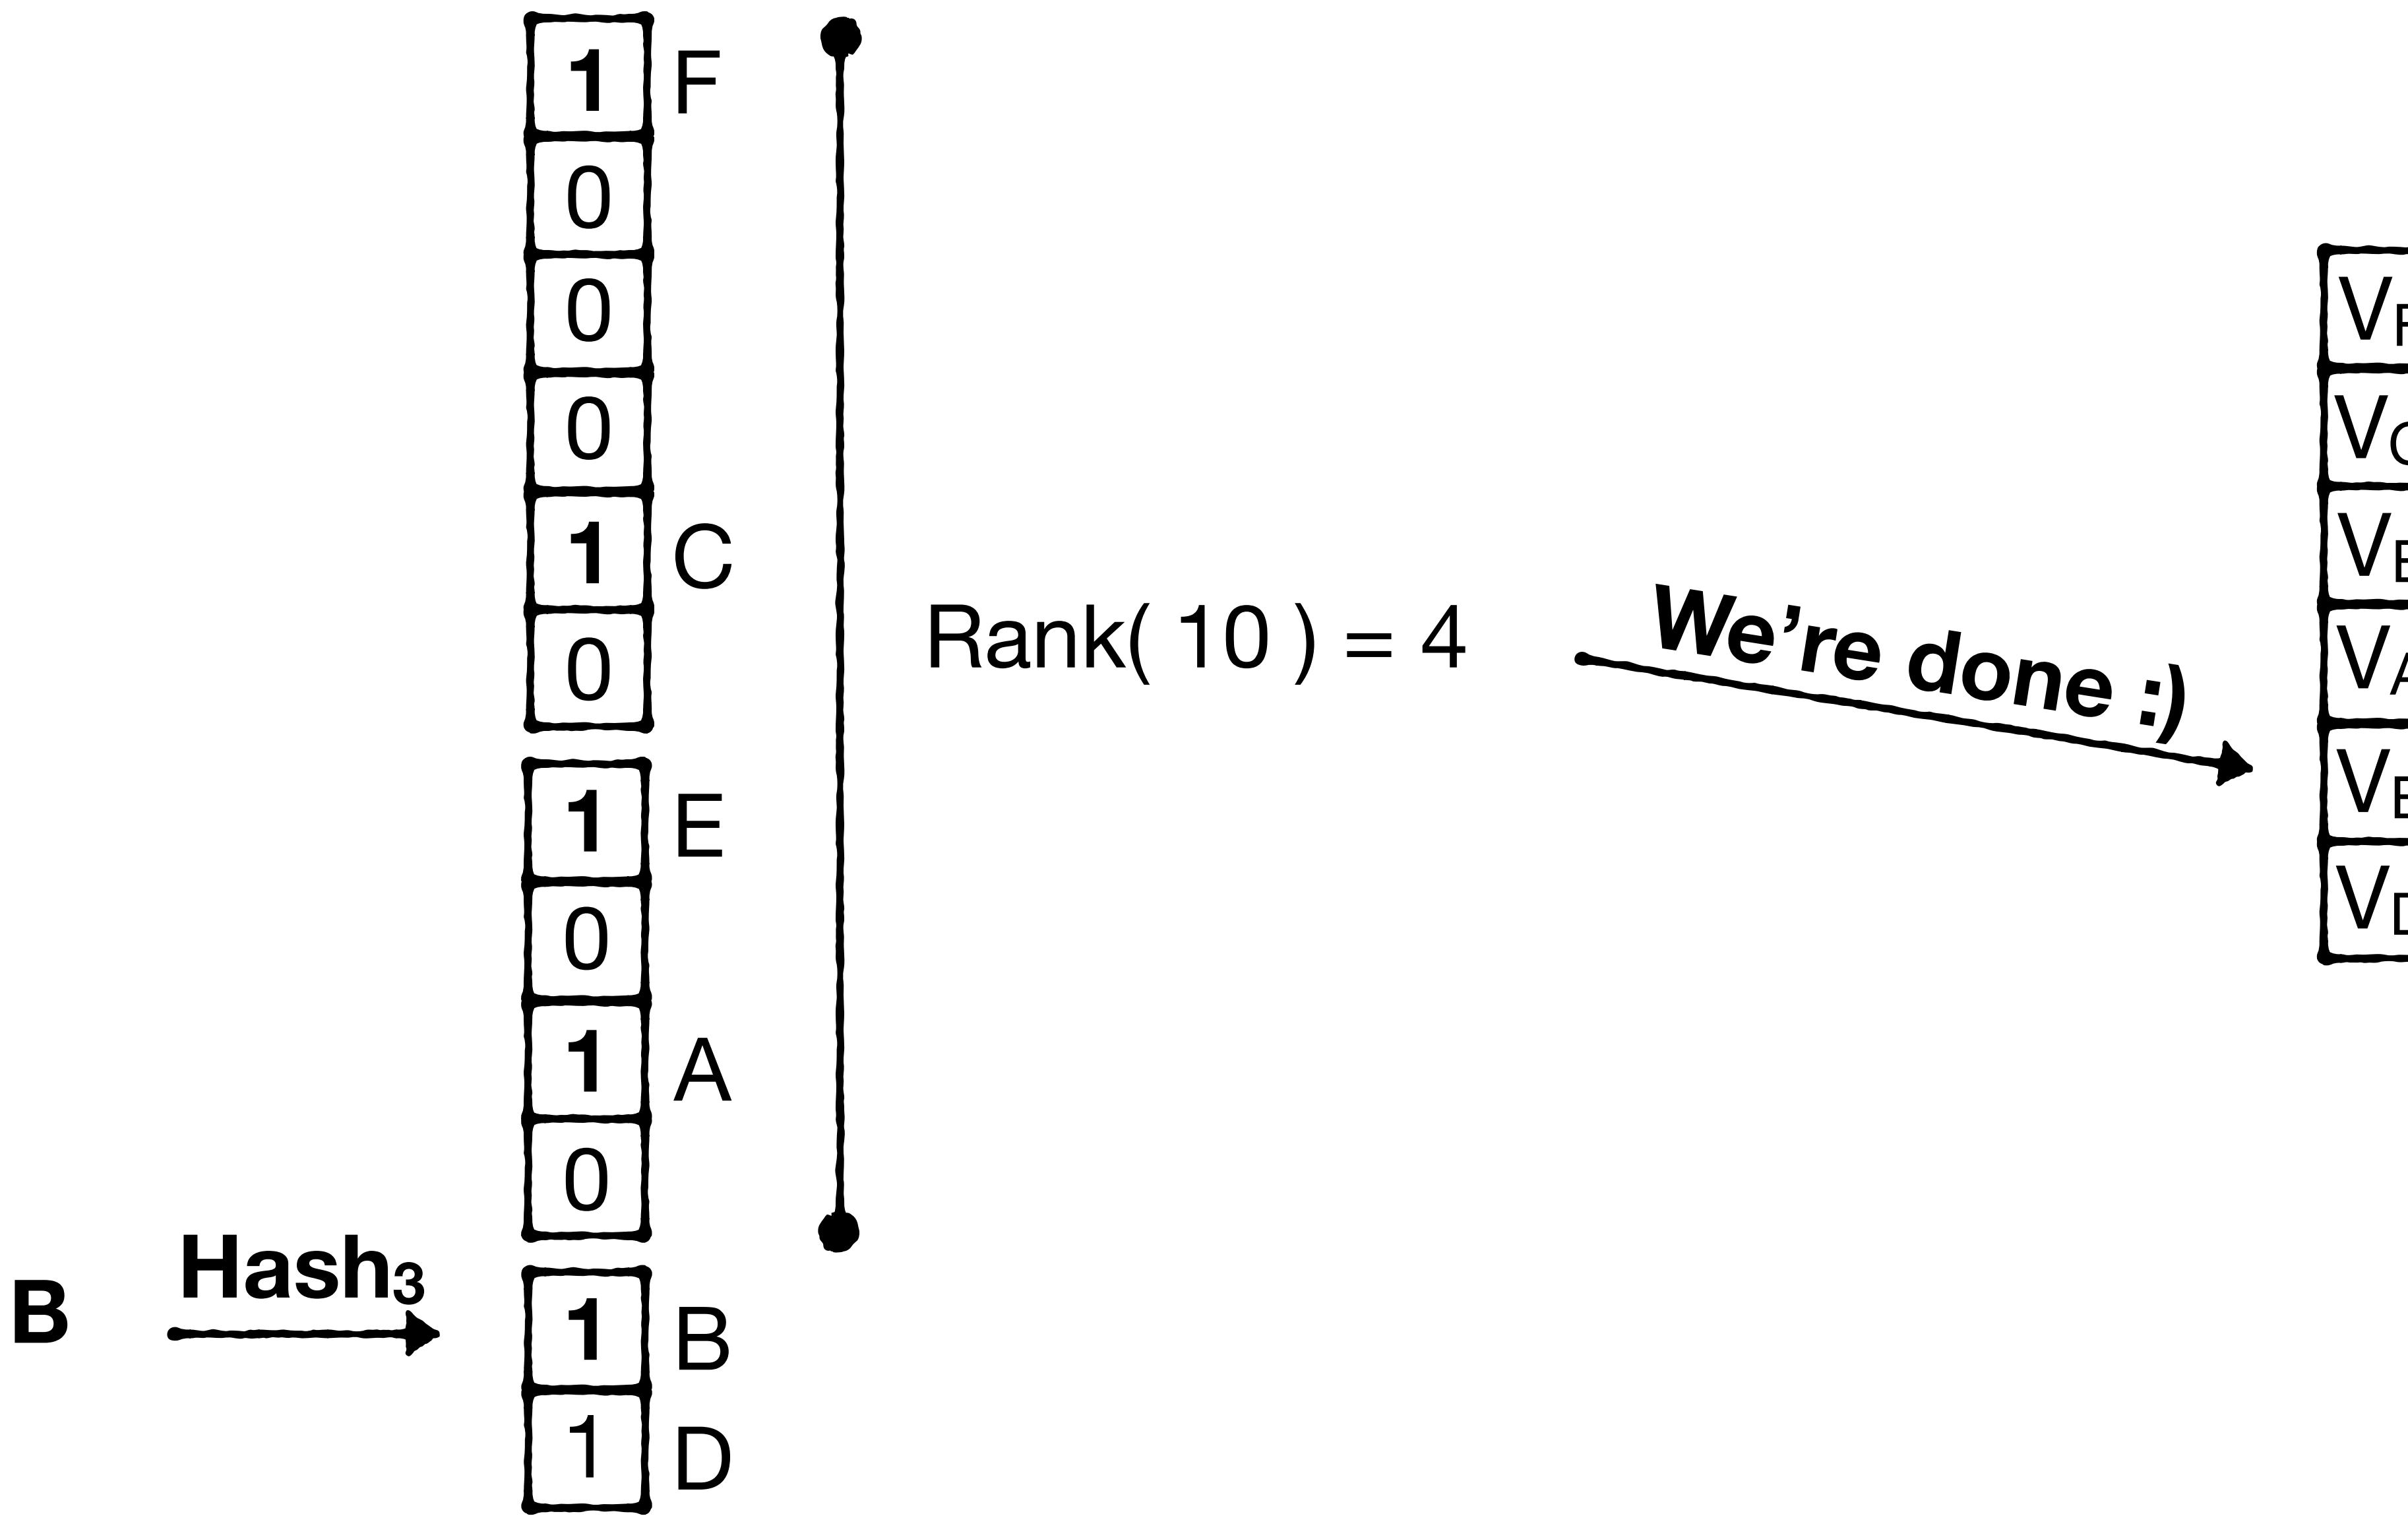
How to query?

Check bitmaps until finding 1
Rank on offset of first 1



How to query?

Check bitmaps until finding 1
Rank on offset of first 1



How to run rank quickly?

1
0
0
0
1
0
1
0
1
0
1
1

Rank(0) = 0

1
0
0
0
0
1
0
1
0
1
1
0
1
1

Rank(4) = 1

How to run rank quickly?

Rank(8) = 3

Rank prefix sums as we saw 2 weeks ago

Rank array

0
1
3

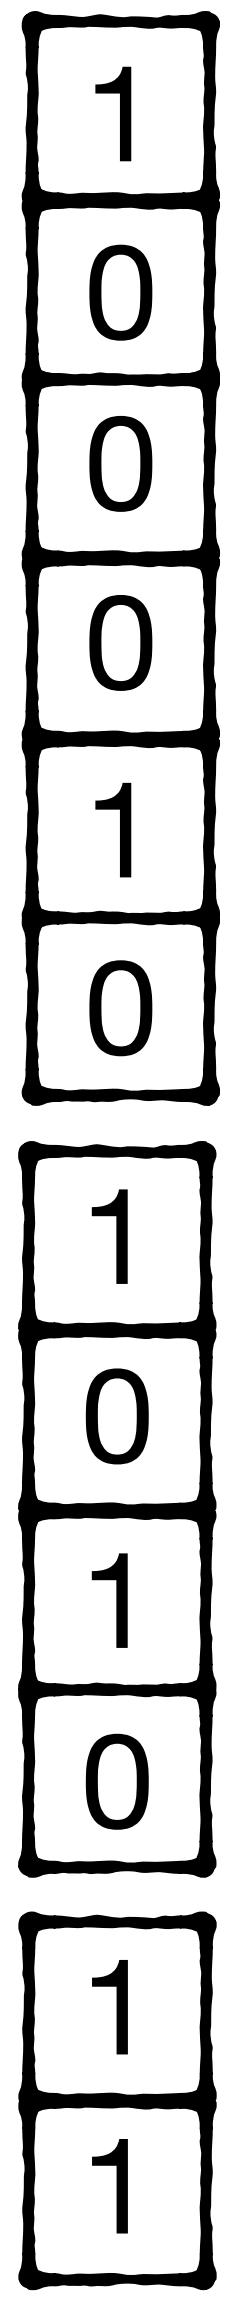
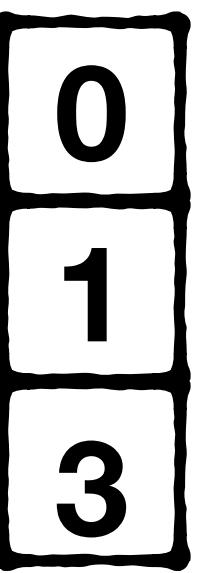
1
0
0
0
1
1
0
1
0
1
1

How to run rank quickly?

Rank prefix sums as we saw 2 weeks ago

Rank(10)

Rank array



Rank(10) $\xrightarrow{\lfloor 10/4 \rfloor = 2}$ Rank array

0
1
3

1
0
0
0
1
1
0
1
0
1
1
0
1
1

Rank(10) $\xrightarrow{\lfloor 10/4 \rfloor = 2}$ **Rank array**

0
1
3

1
0
0
0
1
1
0
1
0
1
1
0
1
0
1
1
1

Count additional 1s

Rank(10) $\xrightarrow{\lfloor 10/4 \rfloor = 2}$ Rank array

popcount(B & (2ⁱ - 1))

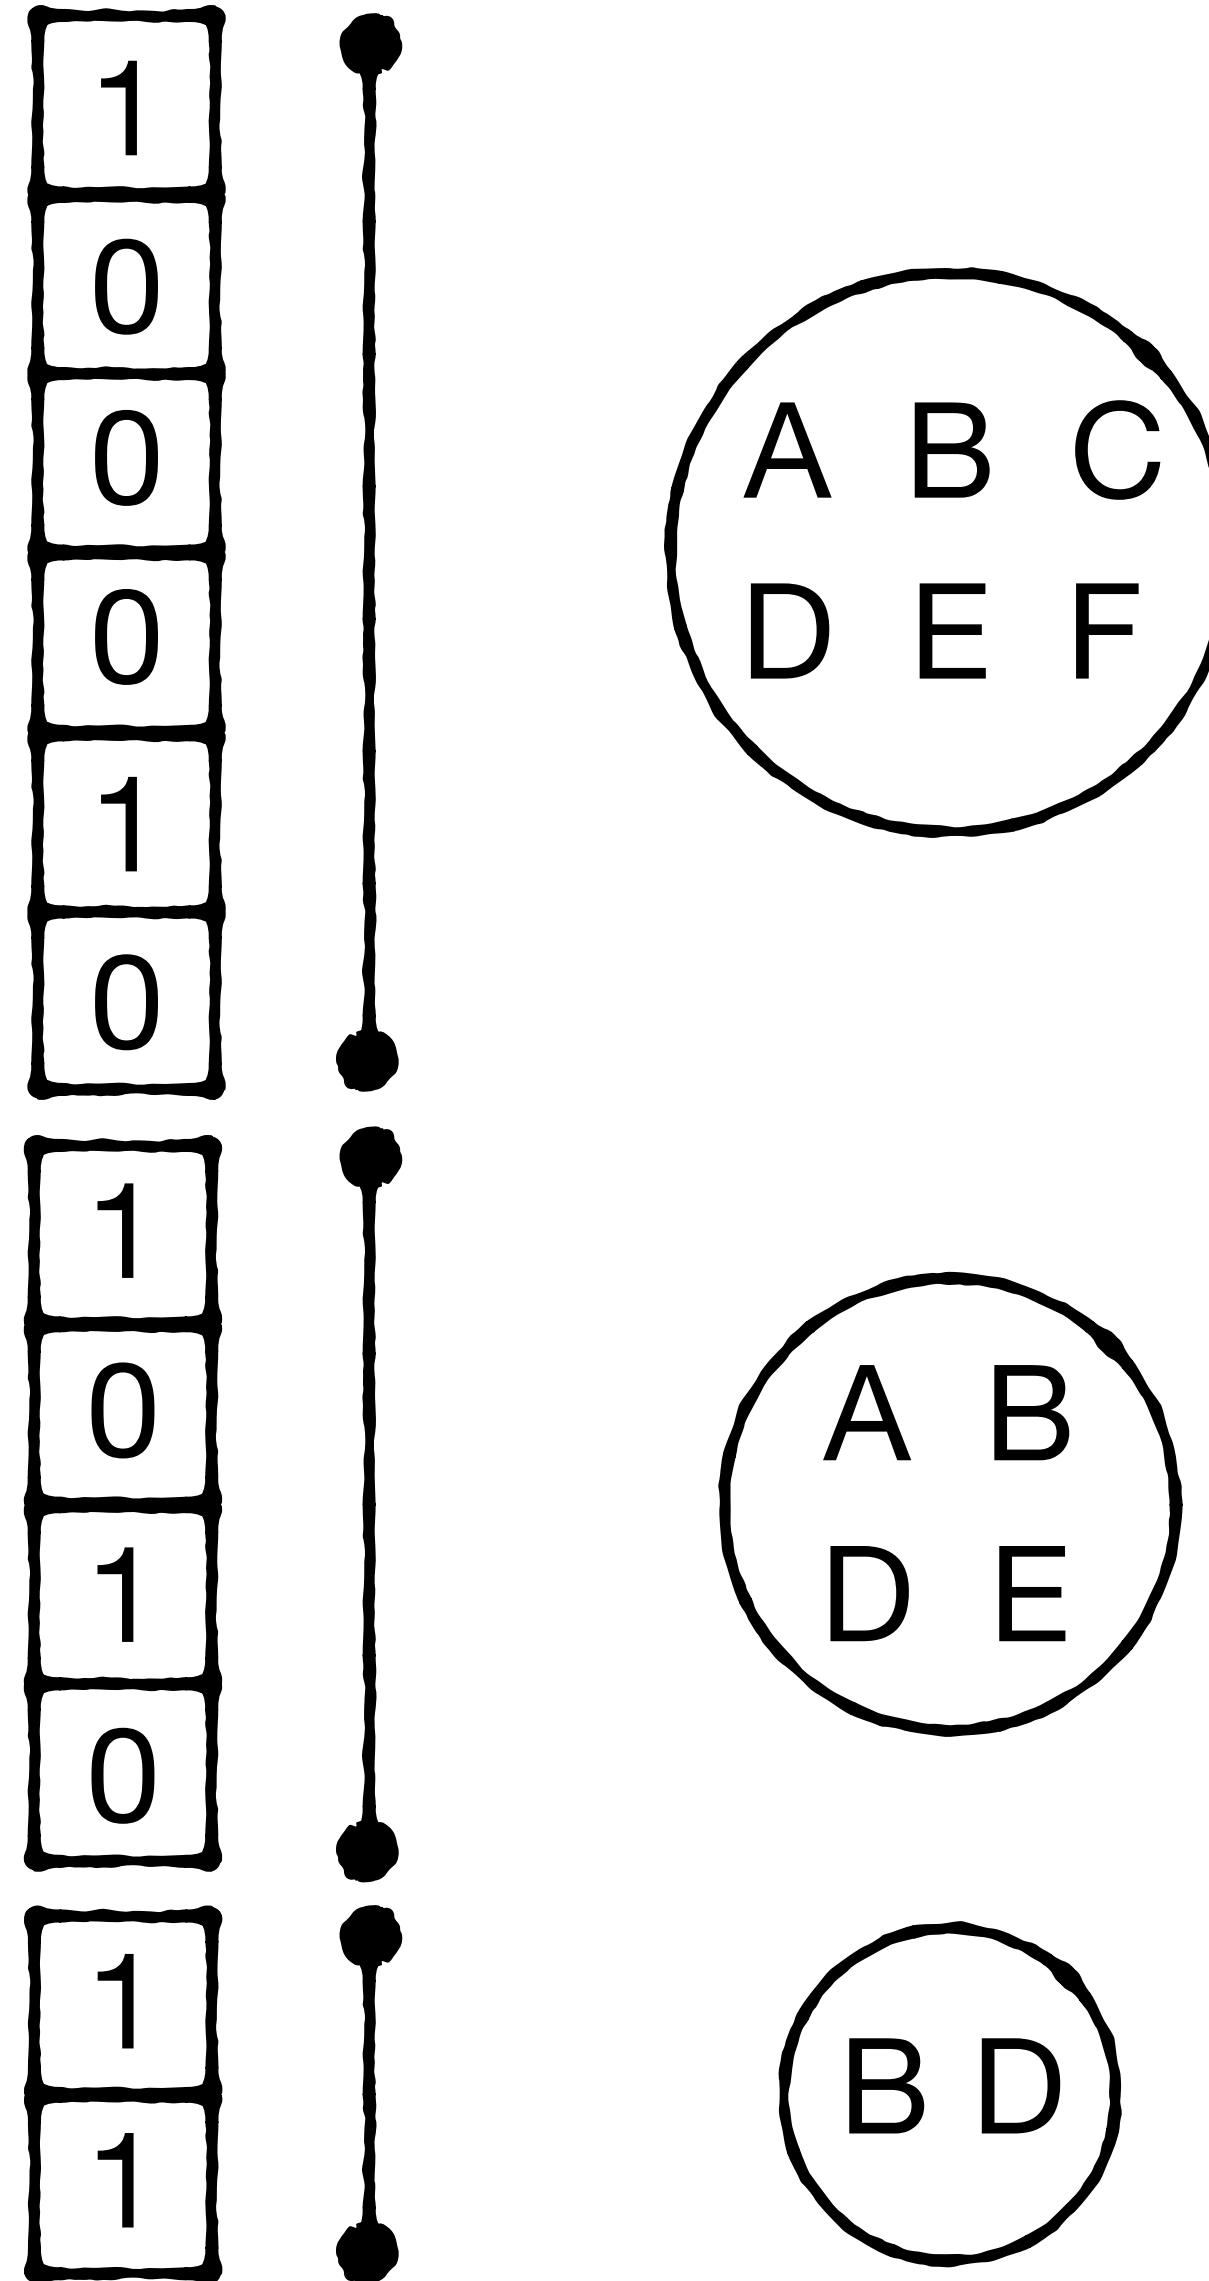
0
1
3

1
0
0
0
1
1
0
1
1
0
1
1

1
0
0
0
1
0
1
1
0
1
0
1
1

Total bitmap size?

Total bitmap size?



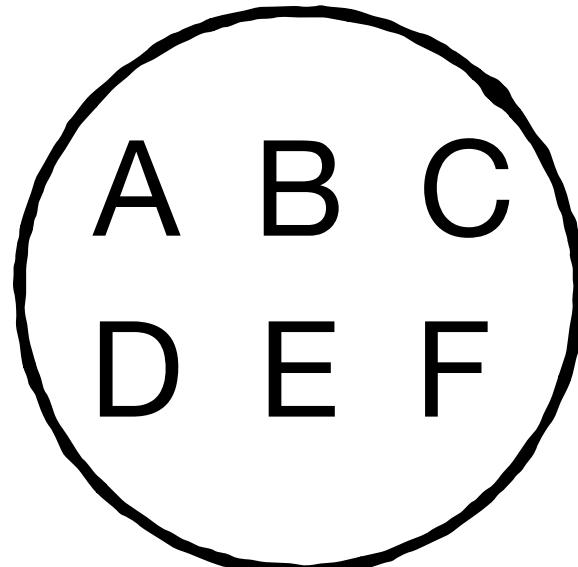
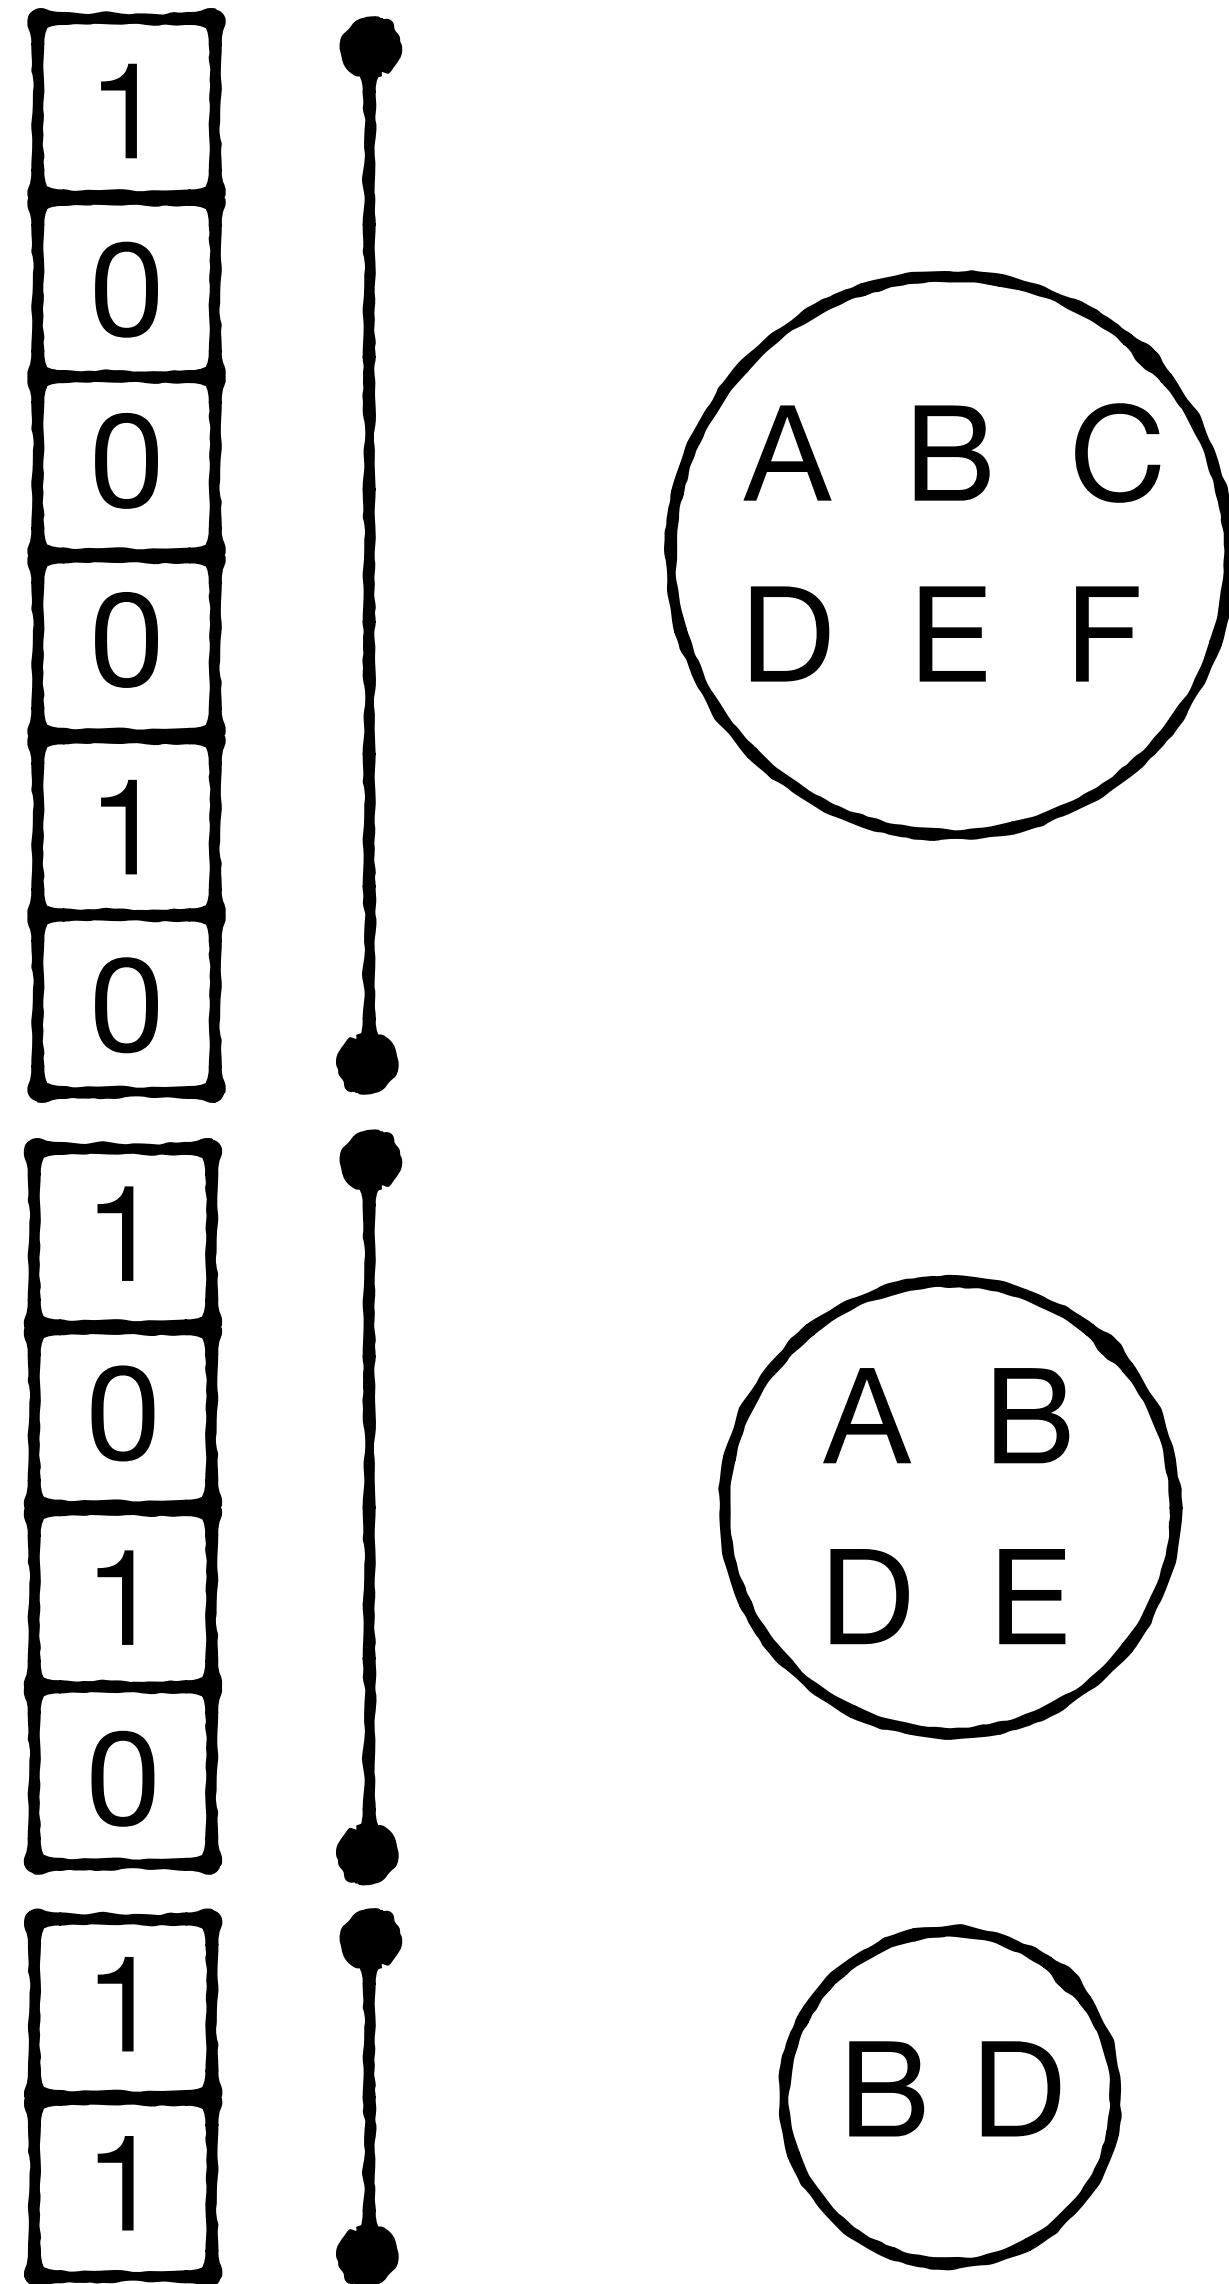
Tried 6 entries. 2 succeeded

Tried 4 entries. 2 succeeded

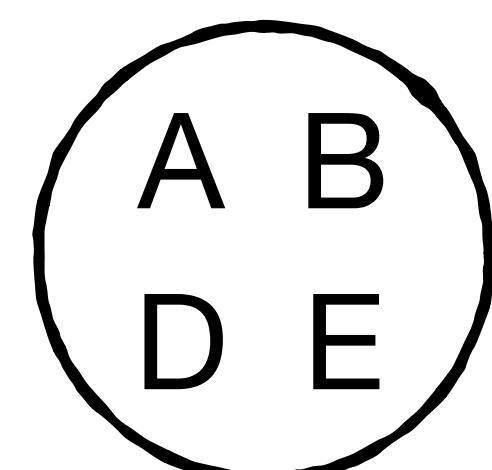
Tried 2 entries. 2 succeeded

Total bitmap size?

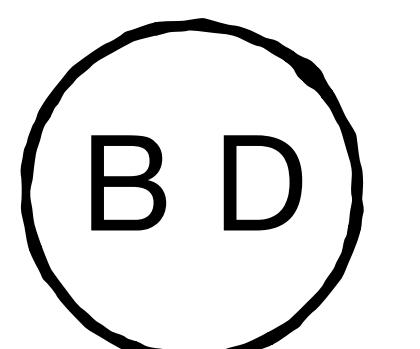
trials across all entries until success



Tried 6 entries. 2 succeeded



Tried 4 entries. 2 succeeded

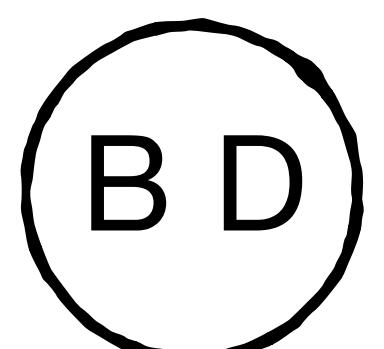
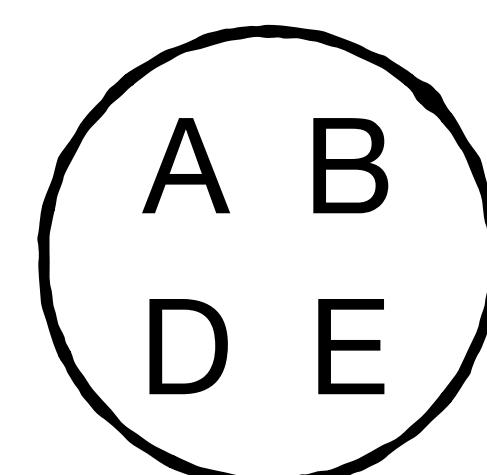
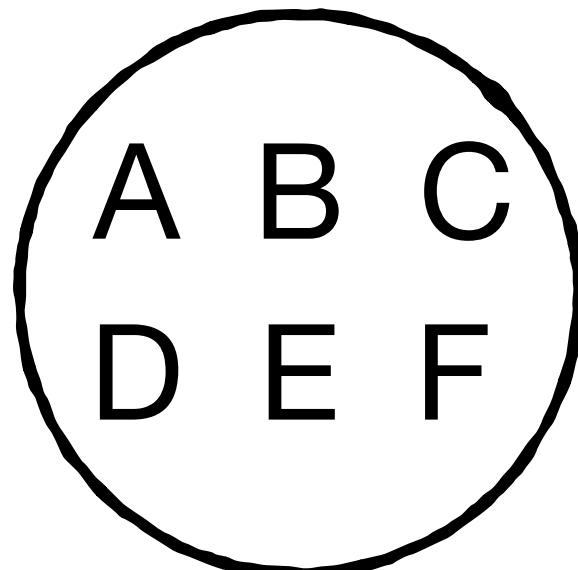
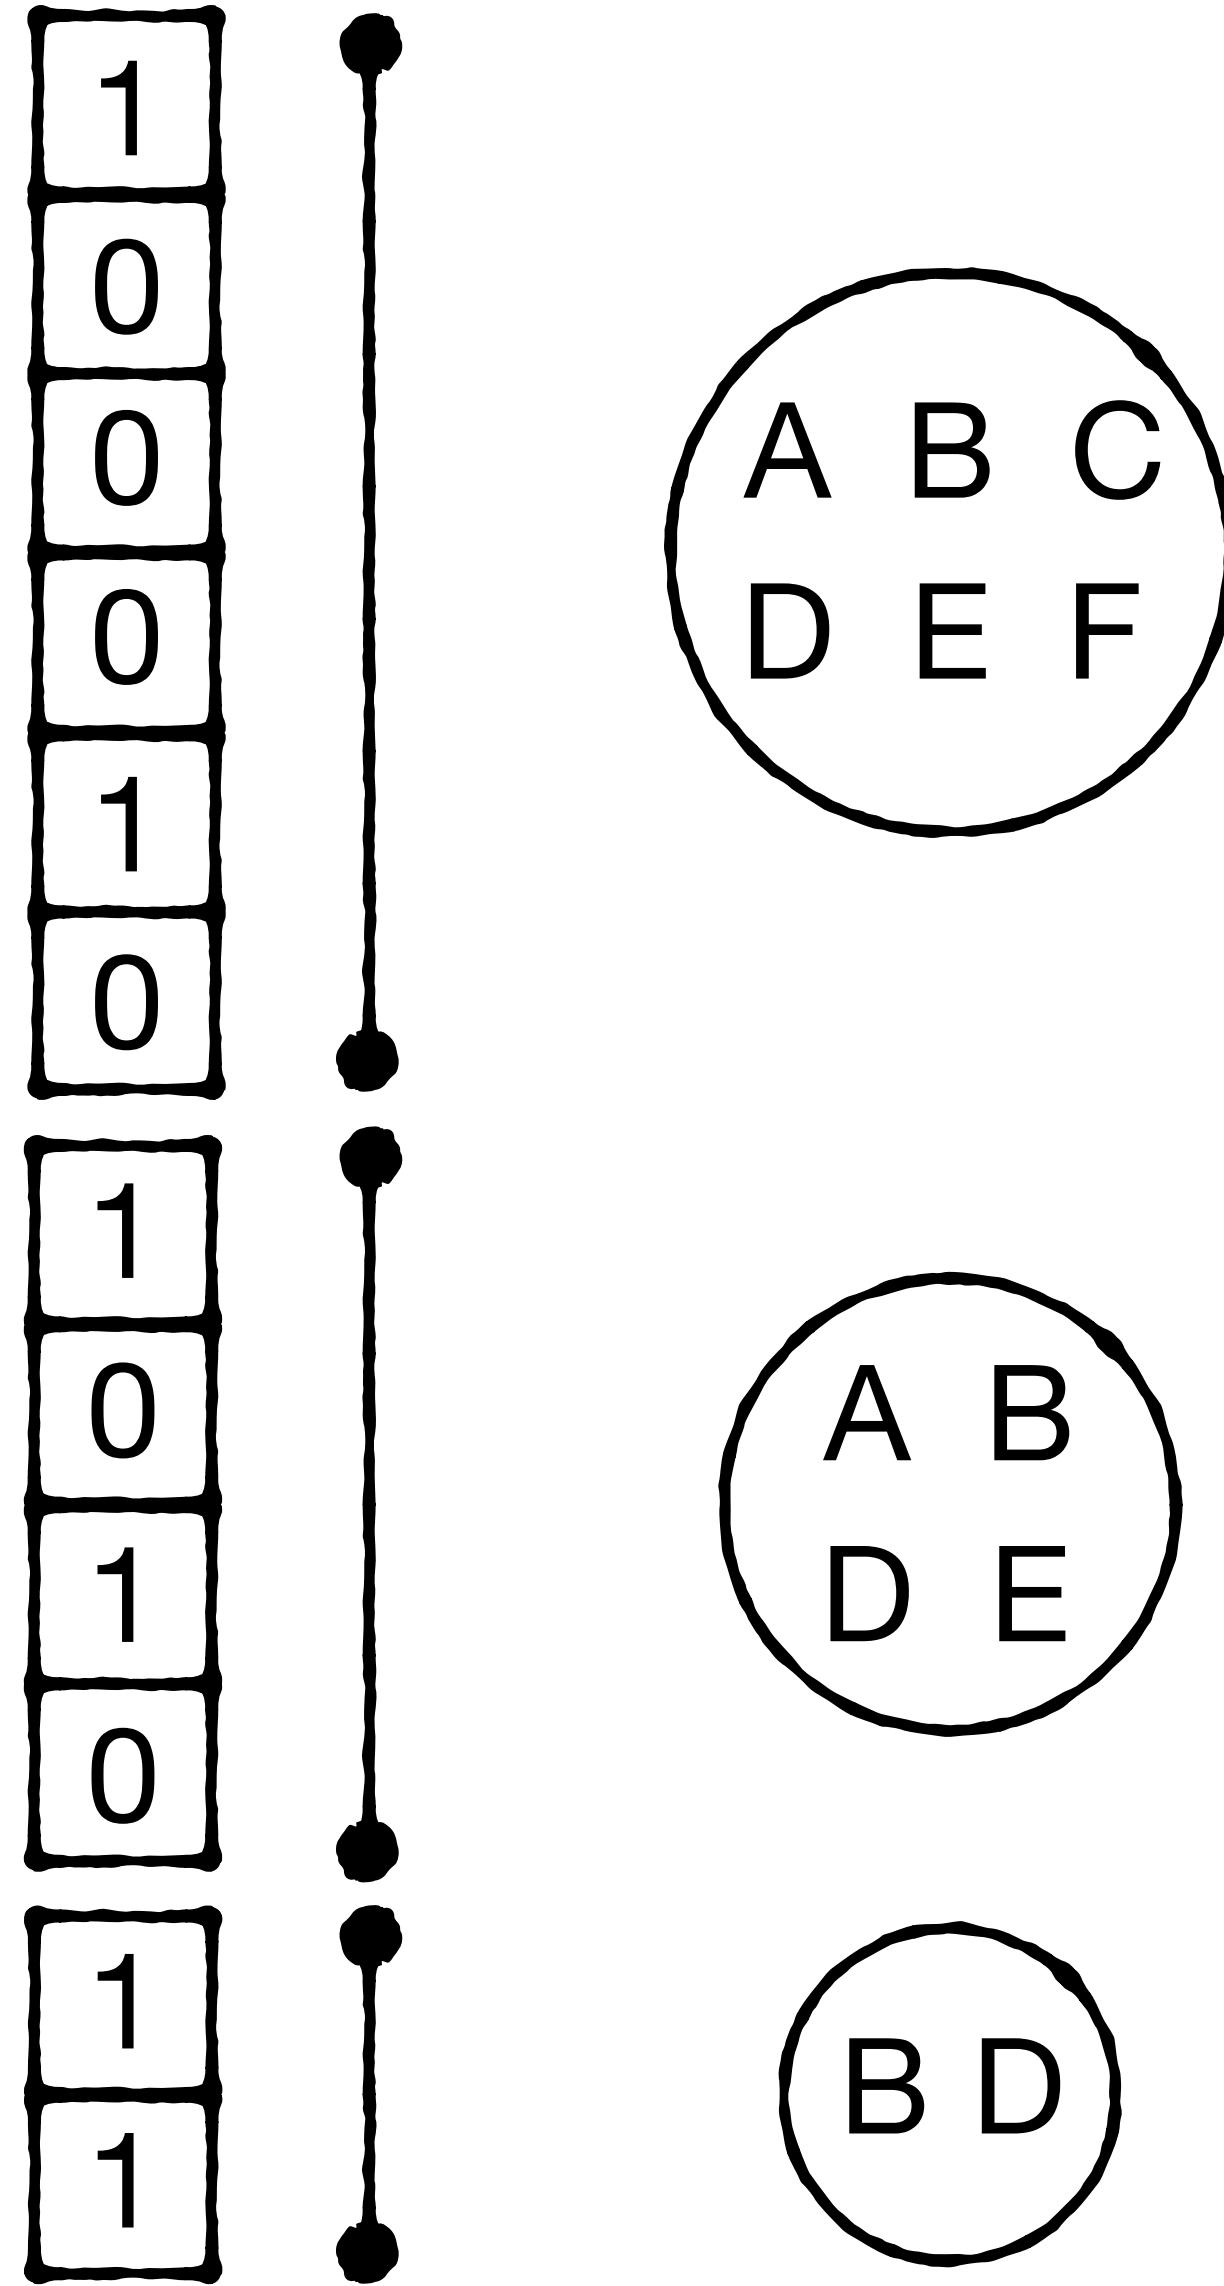


Tried 2 entries. 2 succeeded

Total bitmap size?

trials across all entries until success

$P[\text{entry succeeds in given iteration}]?$



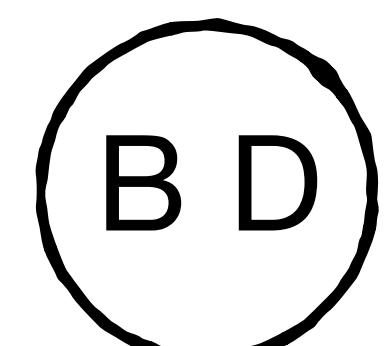
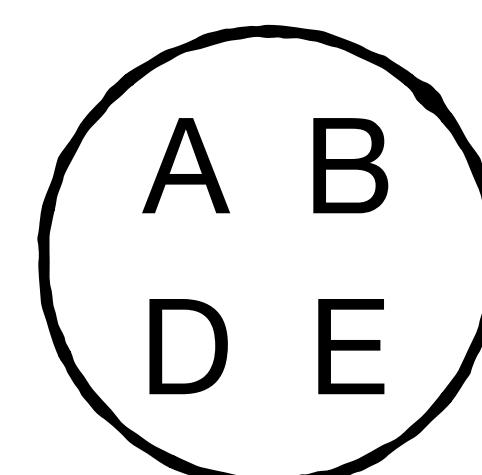
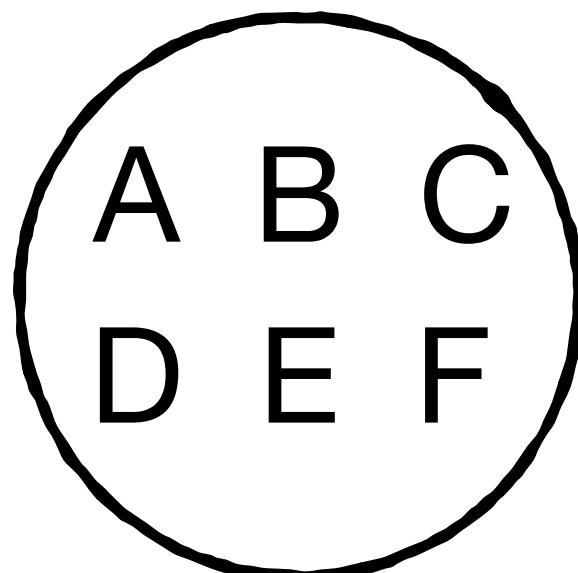
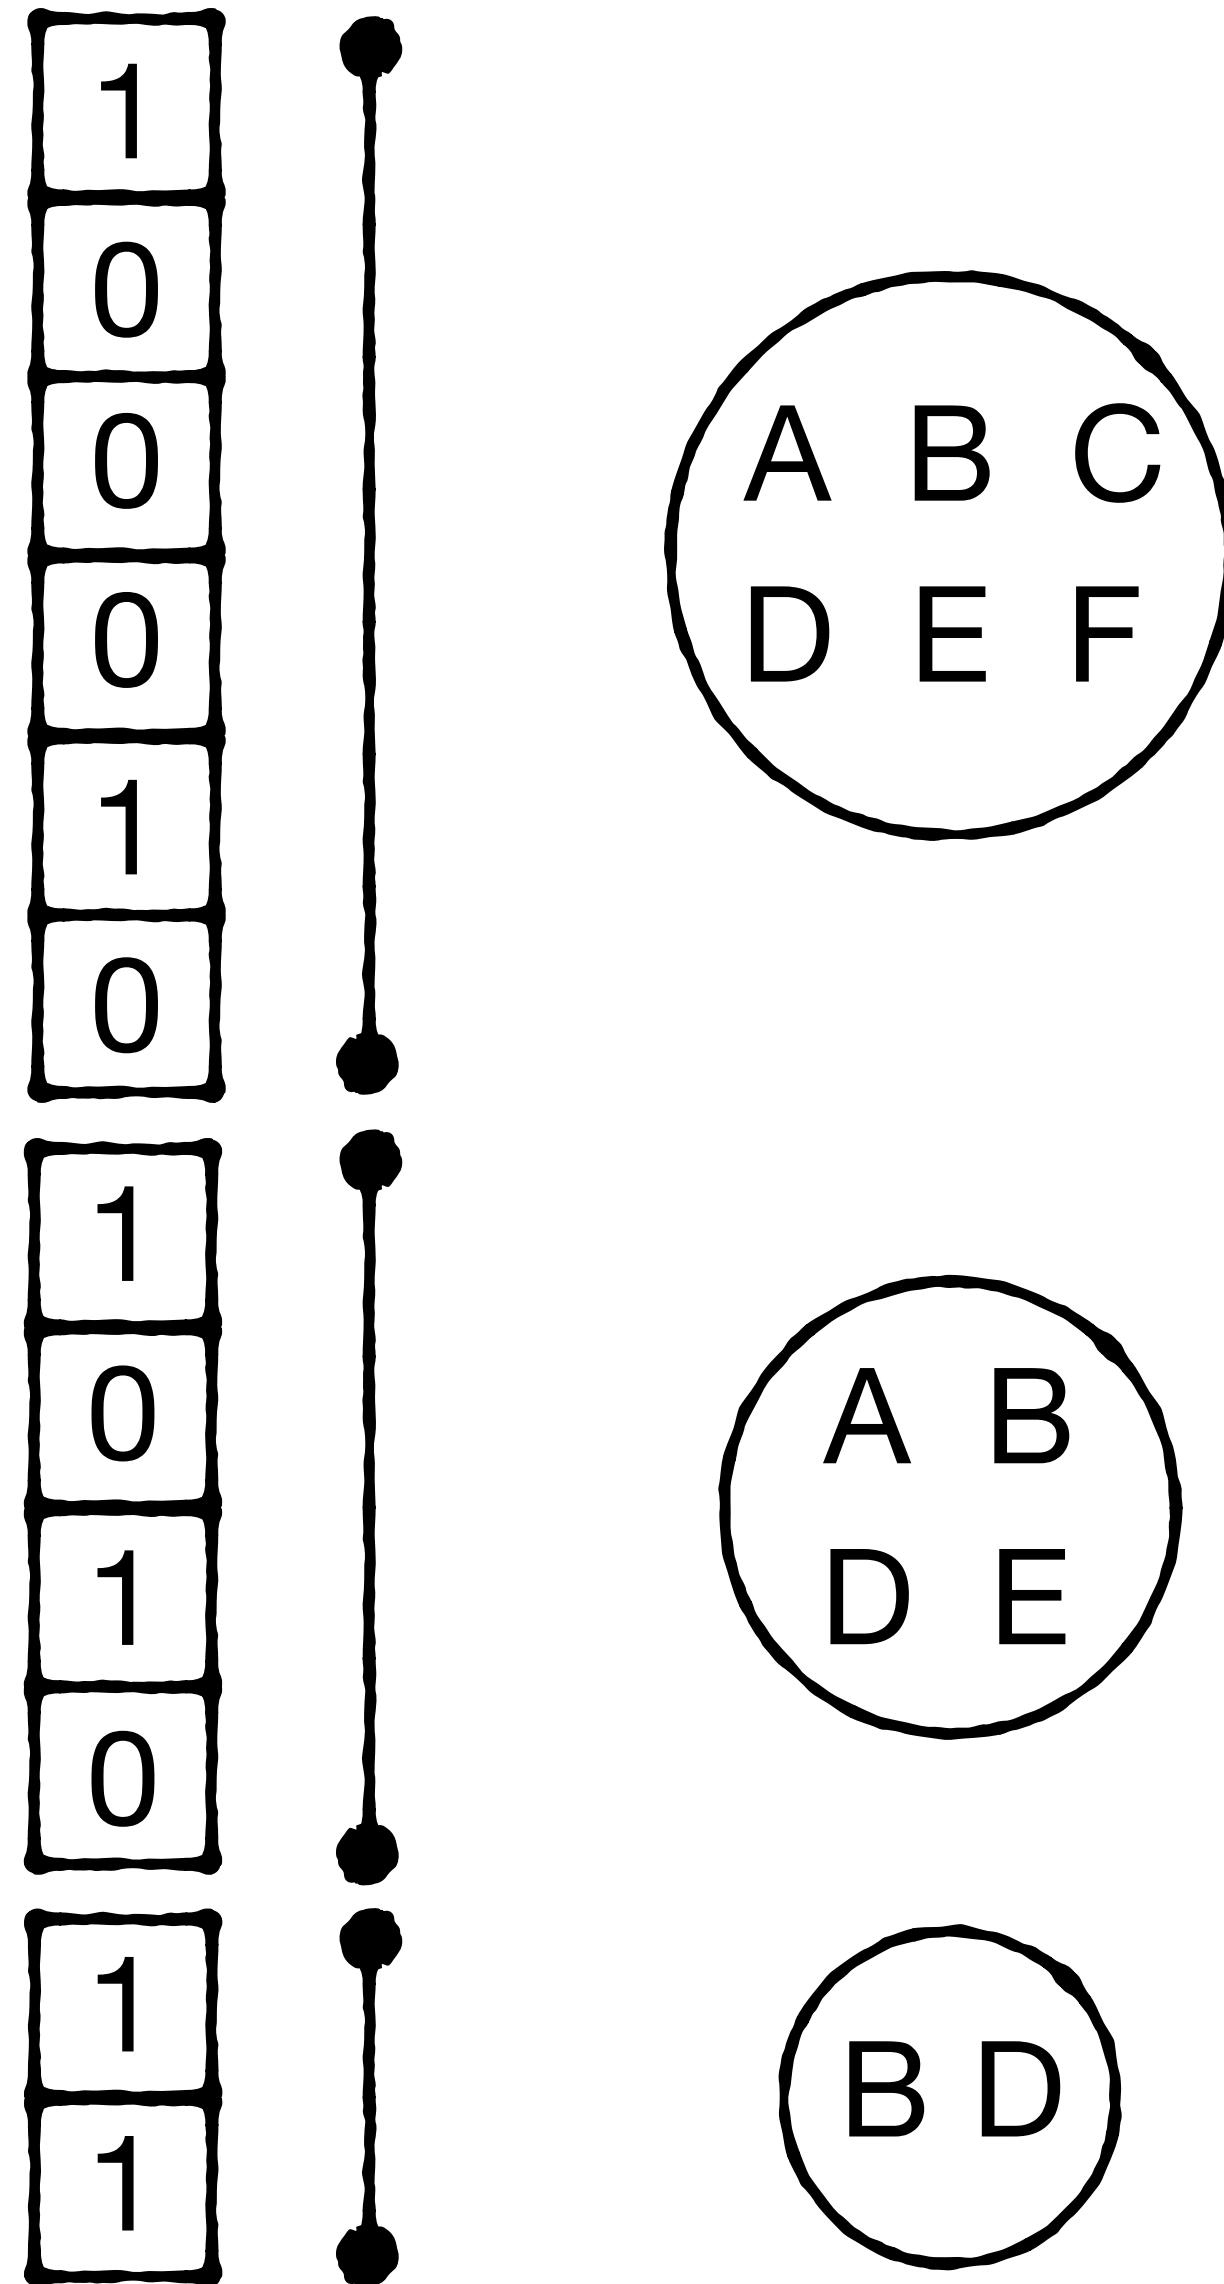
Total bitmap size?

trials across all entries until success

$P[\text{entry succeeds in given iteration}]$

Poisson(λ , 1)

$\lambda = 1$



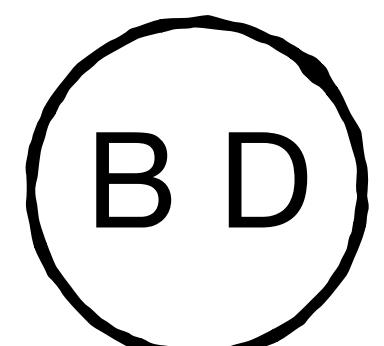
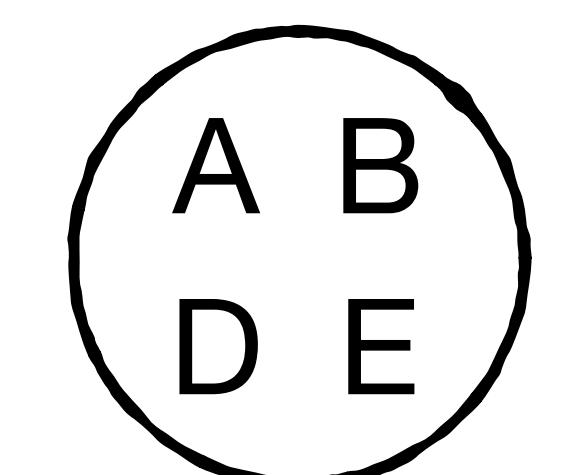
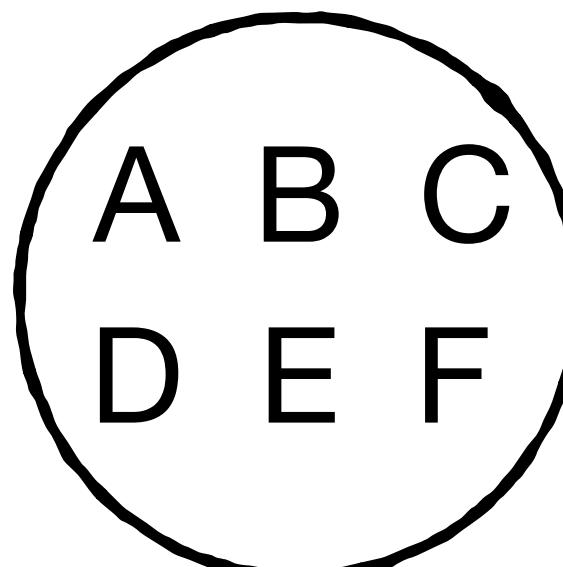
Total bitmap size?

trials across all entries until success

$P[\text{entry succeeds in given iteration}]$

Poisson(1,1) = e^{-1}

1
0
0
0
1
0
1
0
1
0
1
0
1
1



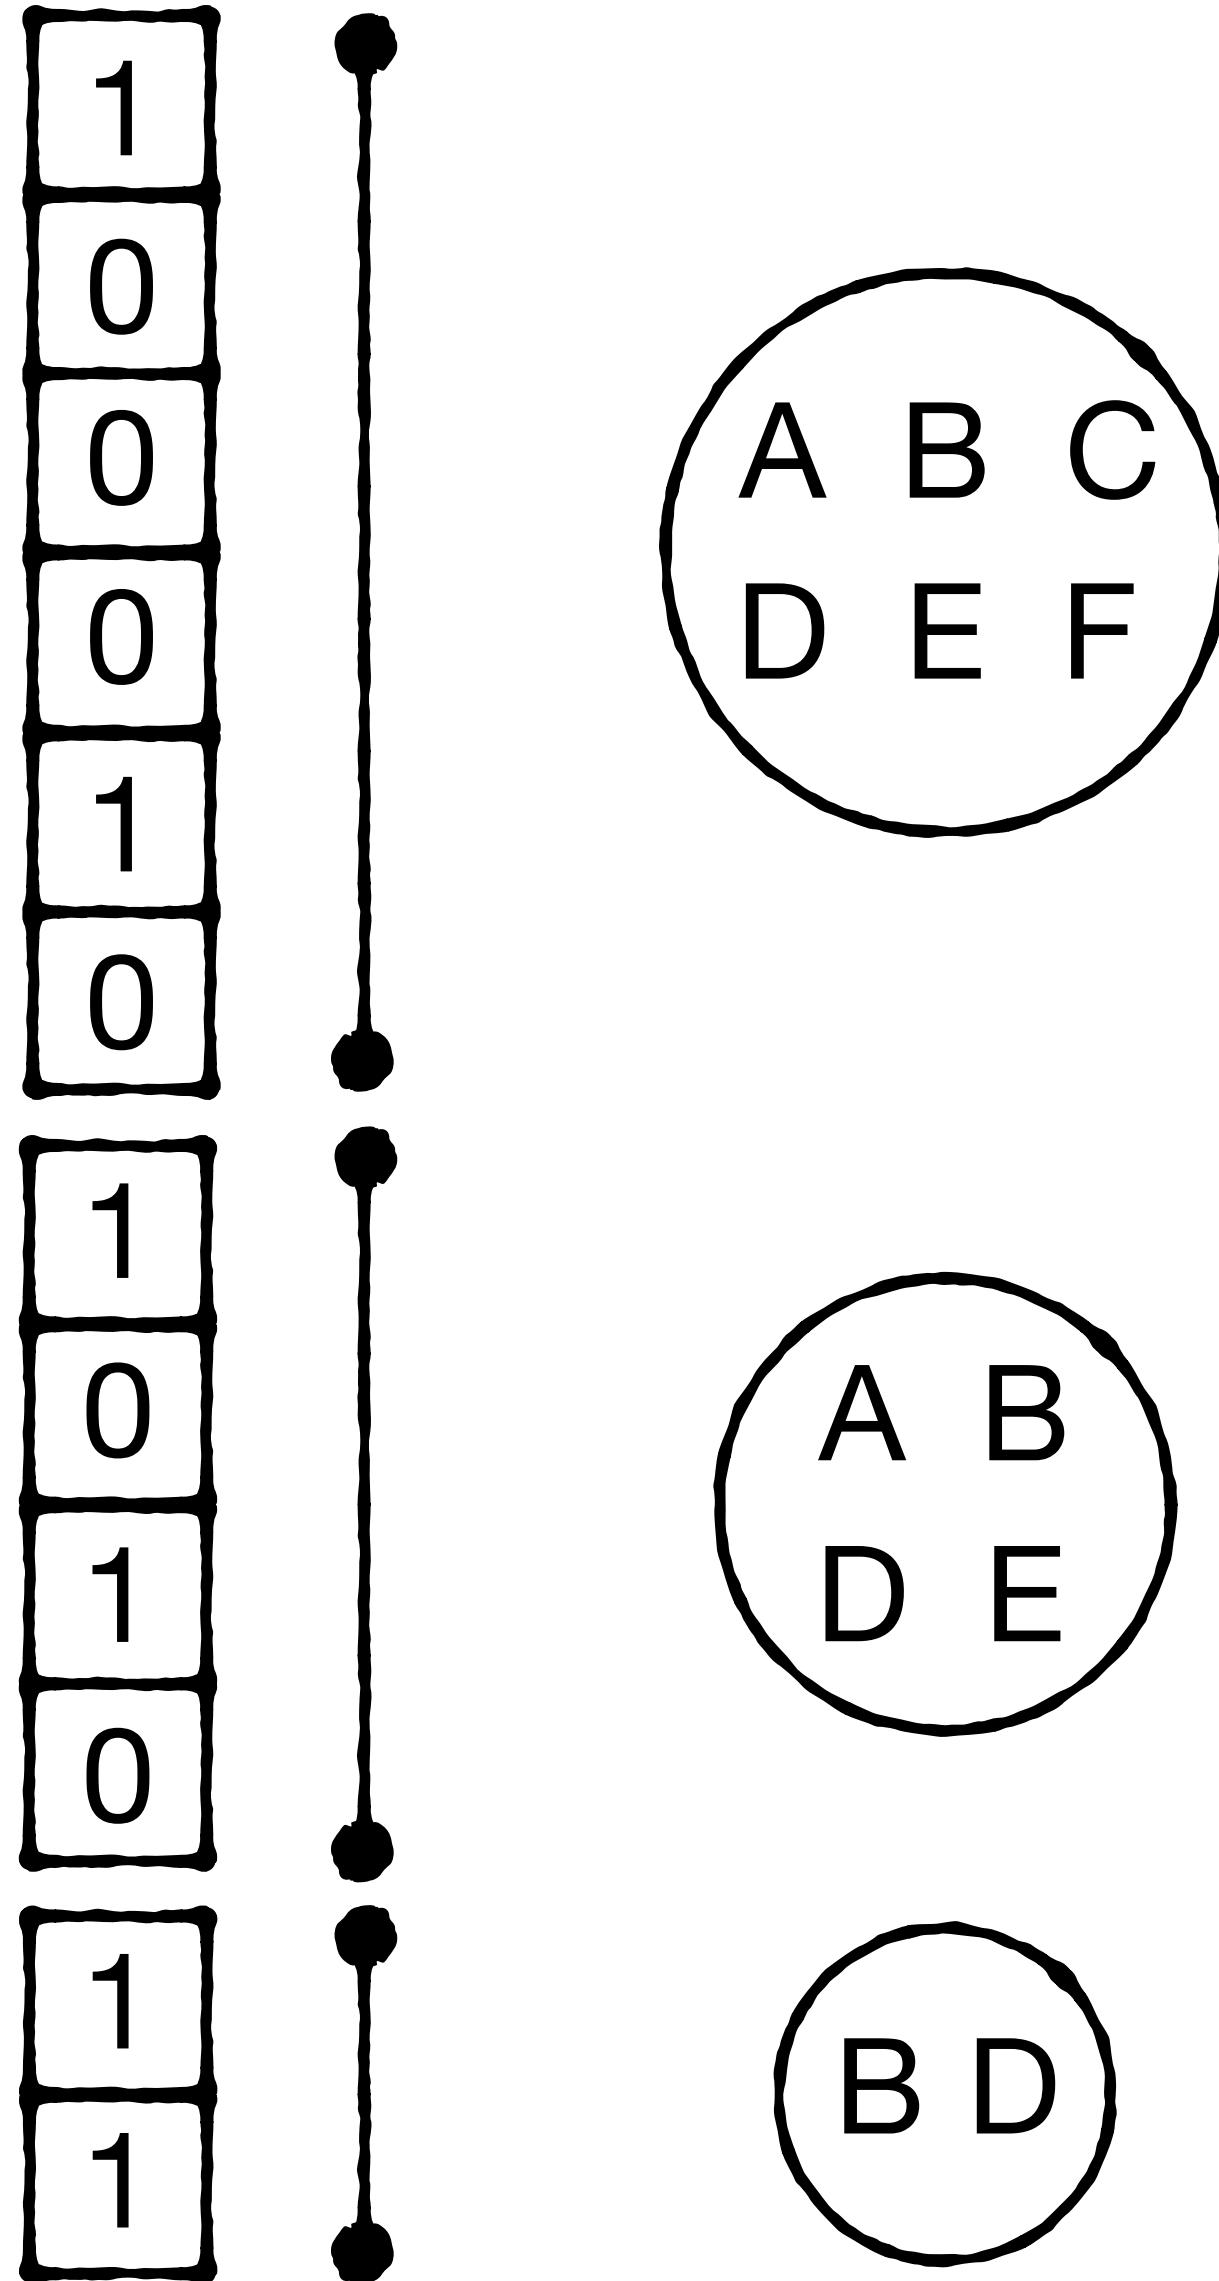
Total bitmap size?

trials across all entries until success

$P[\text{entry succeeds in given iteration}]$

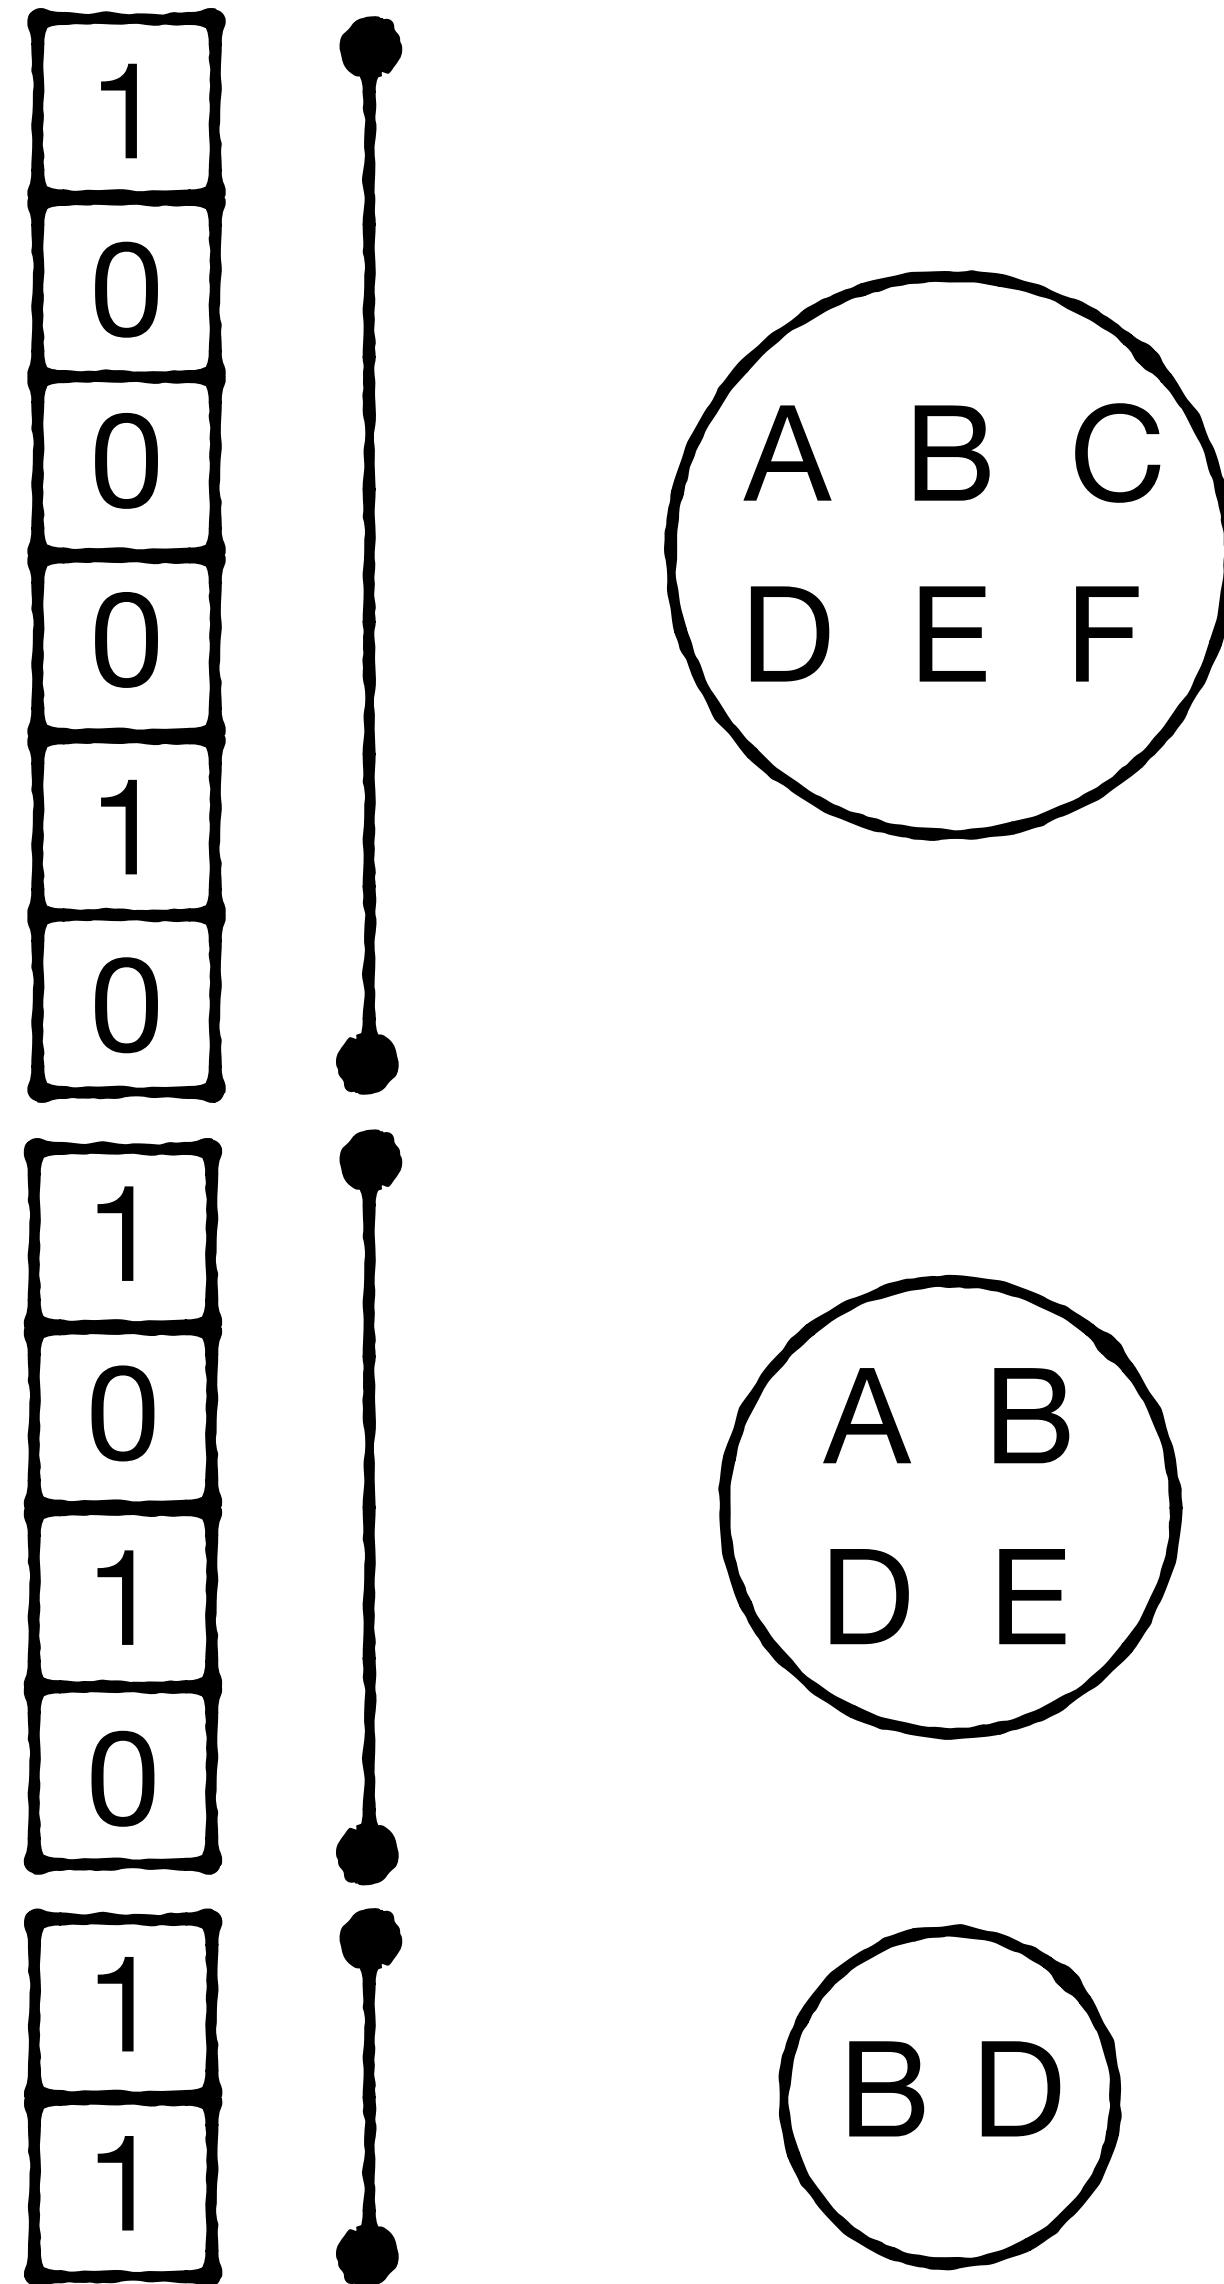
$$\text{Poisson}(1,1) = e^{-1}$$

$E[\#\text{ iterations until succeeding}]?$



Total bitmap size?

trials across all entries until success

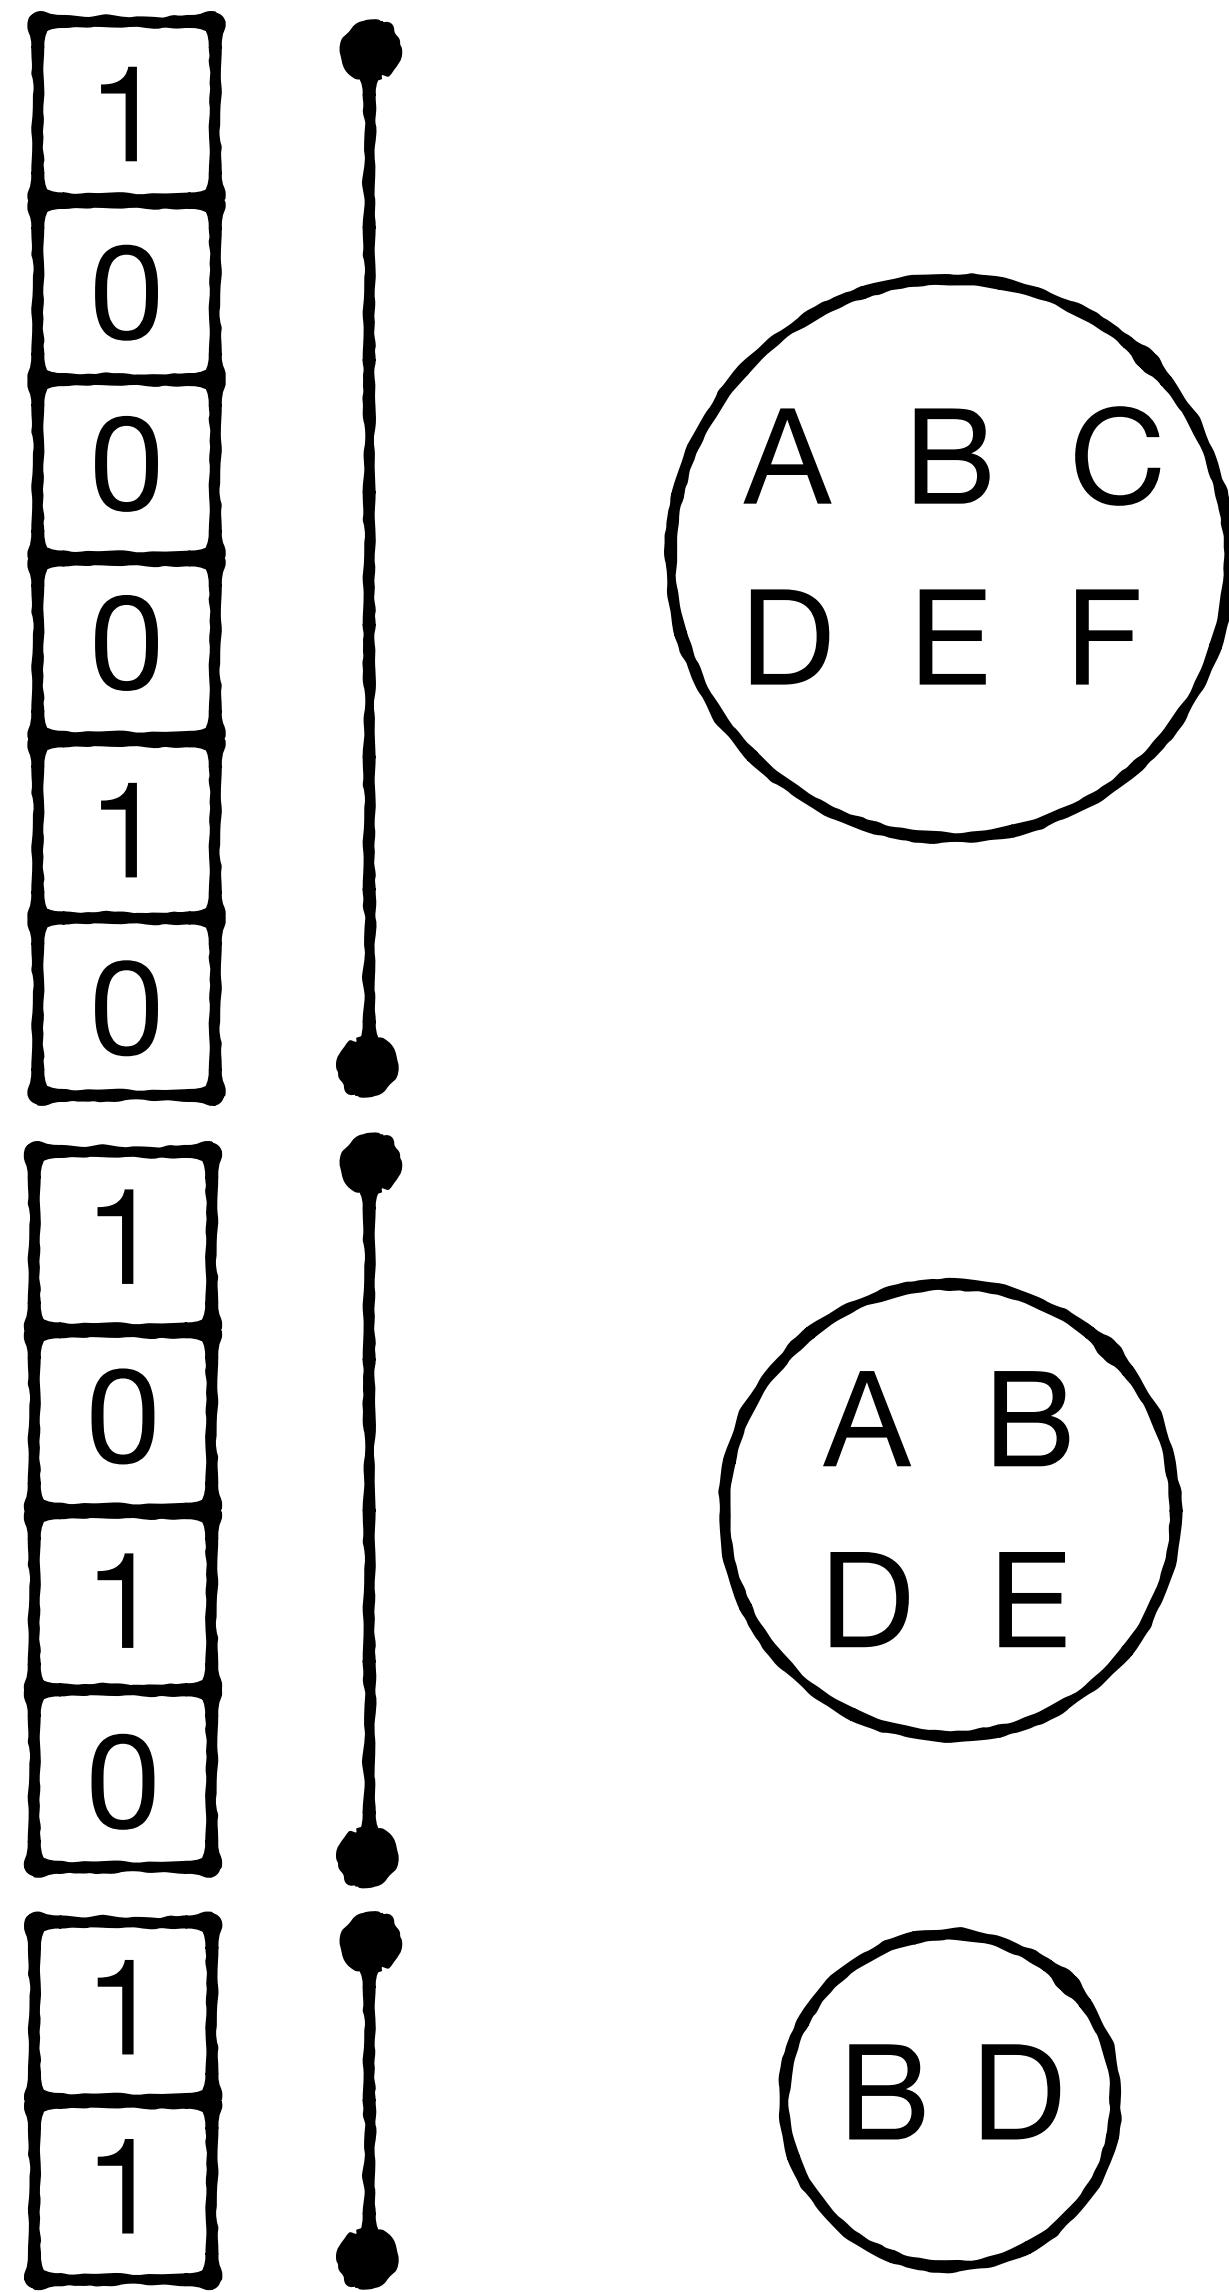


$E[\# \text{ iterations until succeeding}]?$

$$\frac{1}{e^{-1}}$$

Total bitmap size?

trials across all entries until success



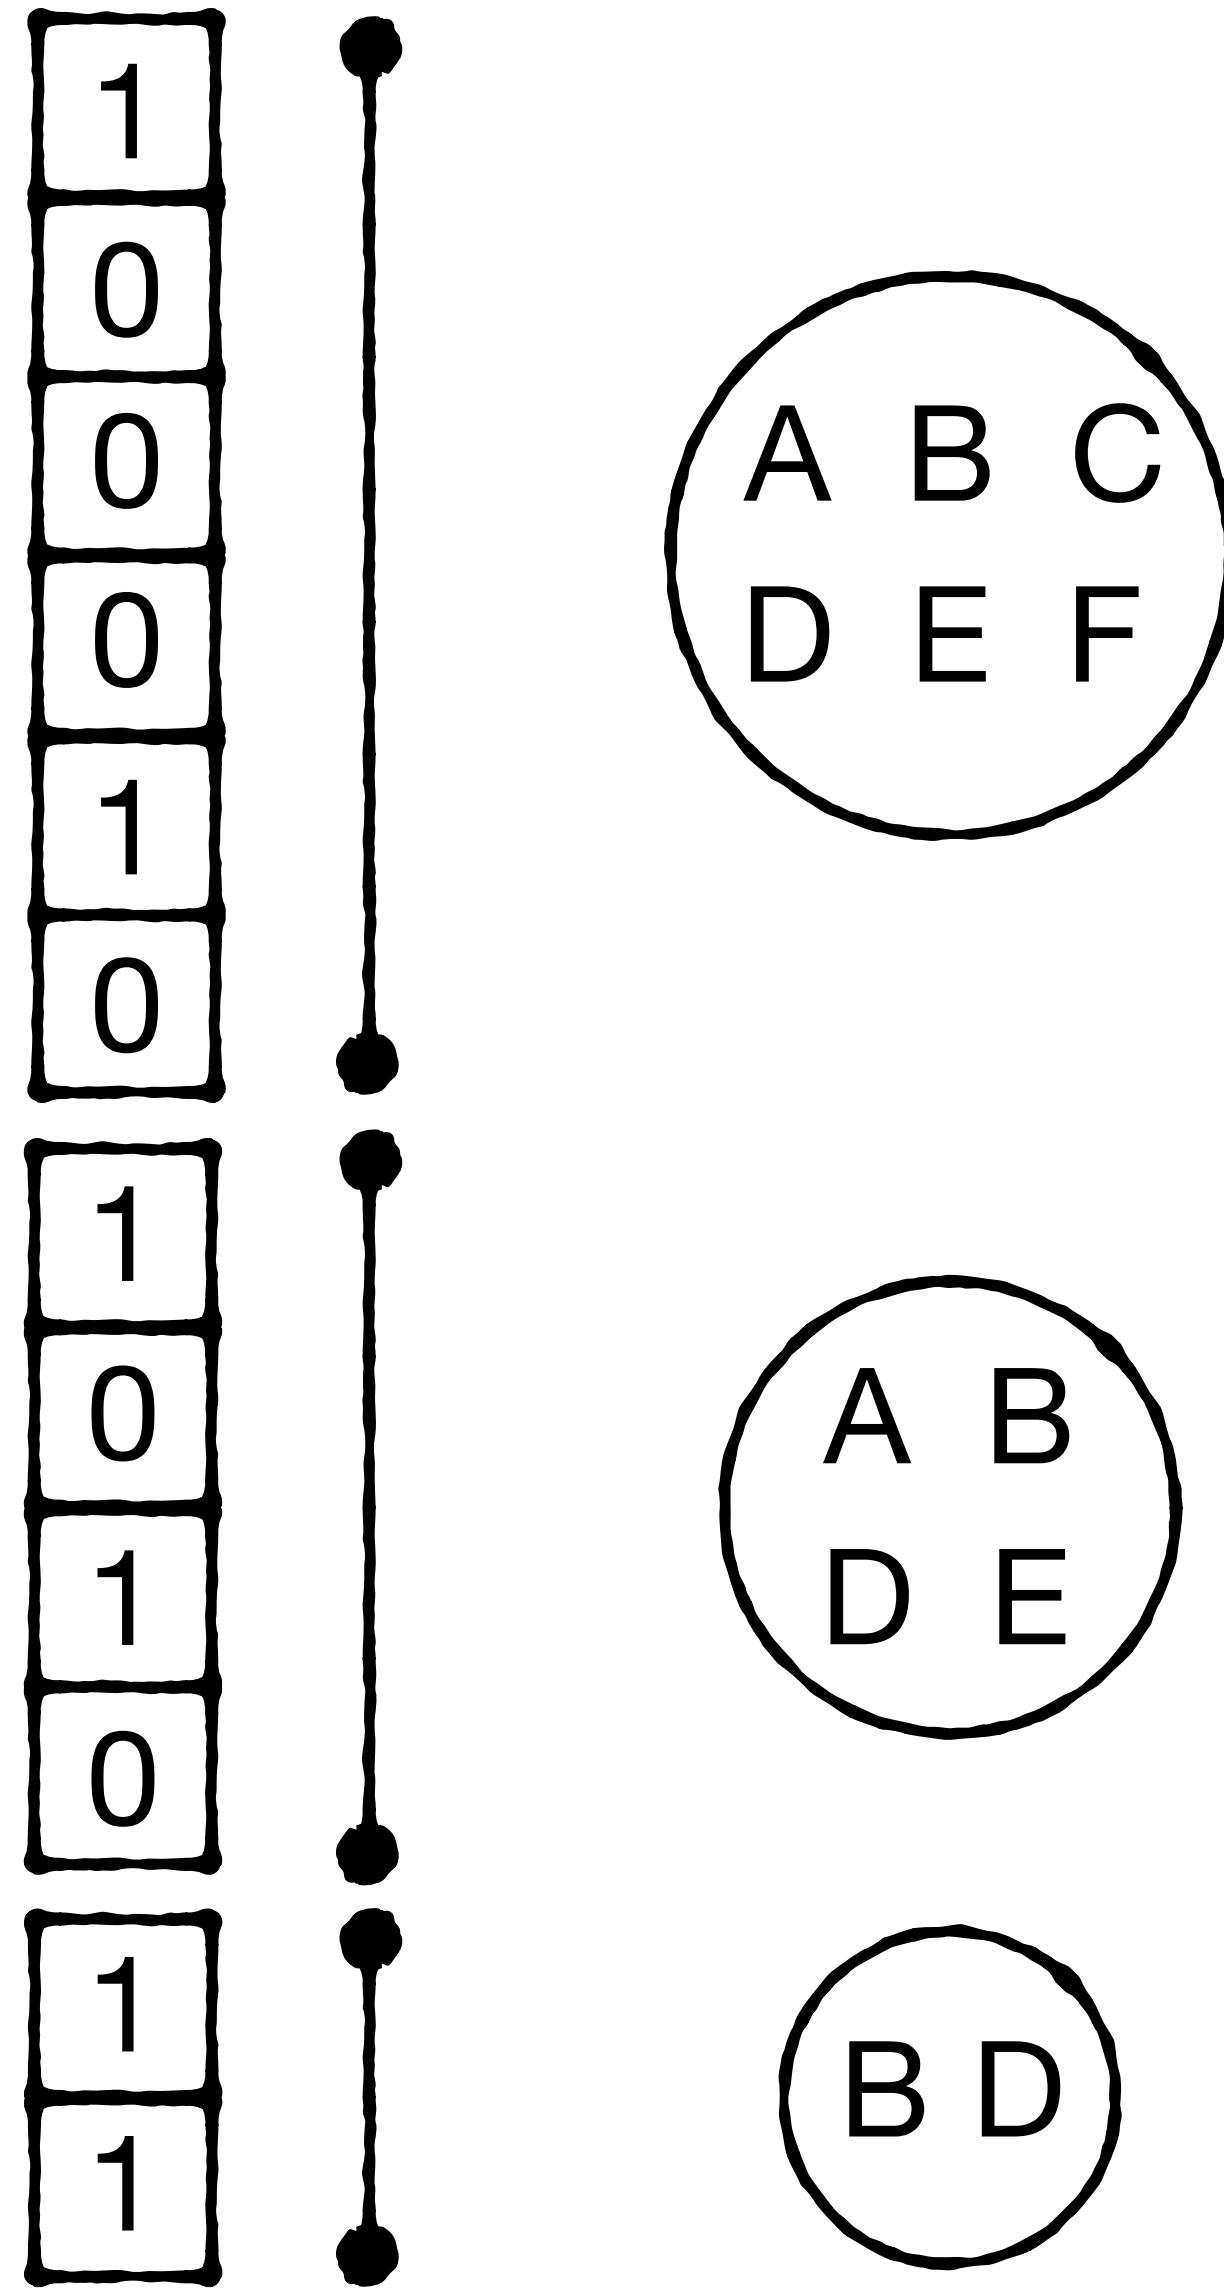
$E[\# \text{ iterations until succeeding}]?$

e

Total bitmap size?

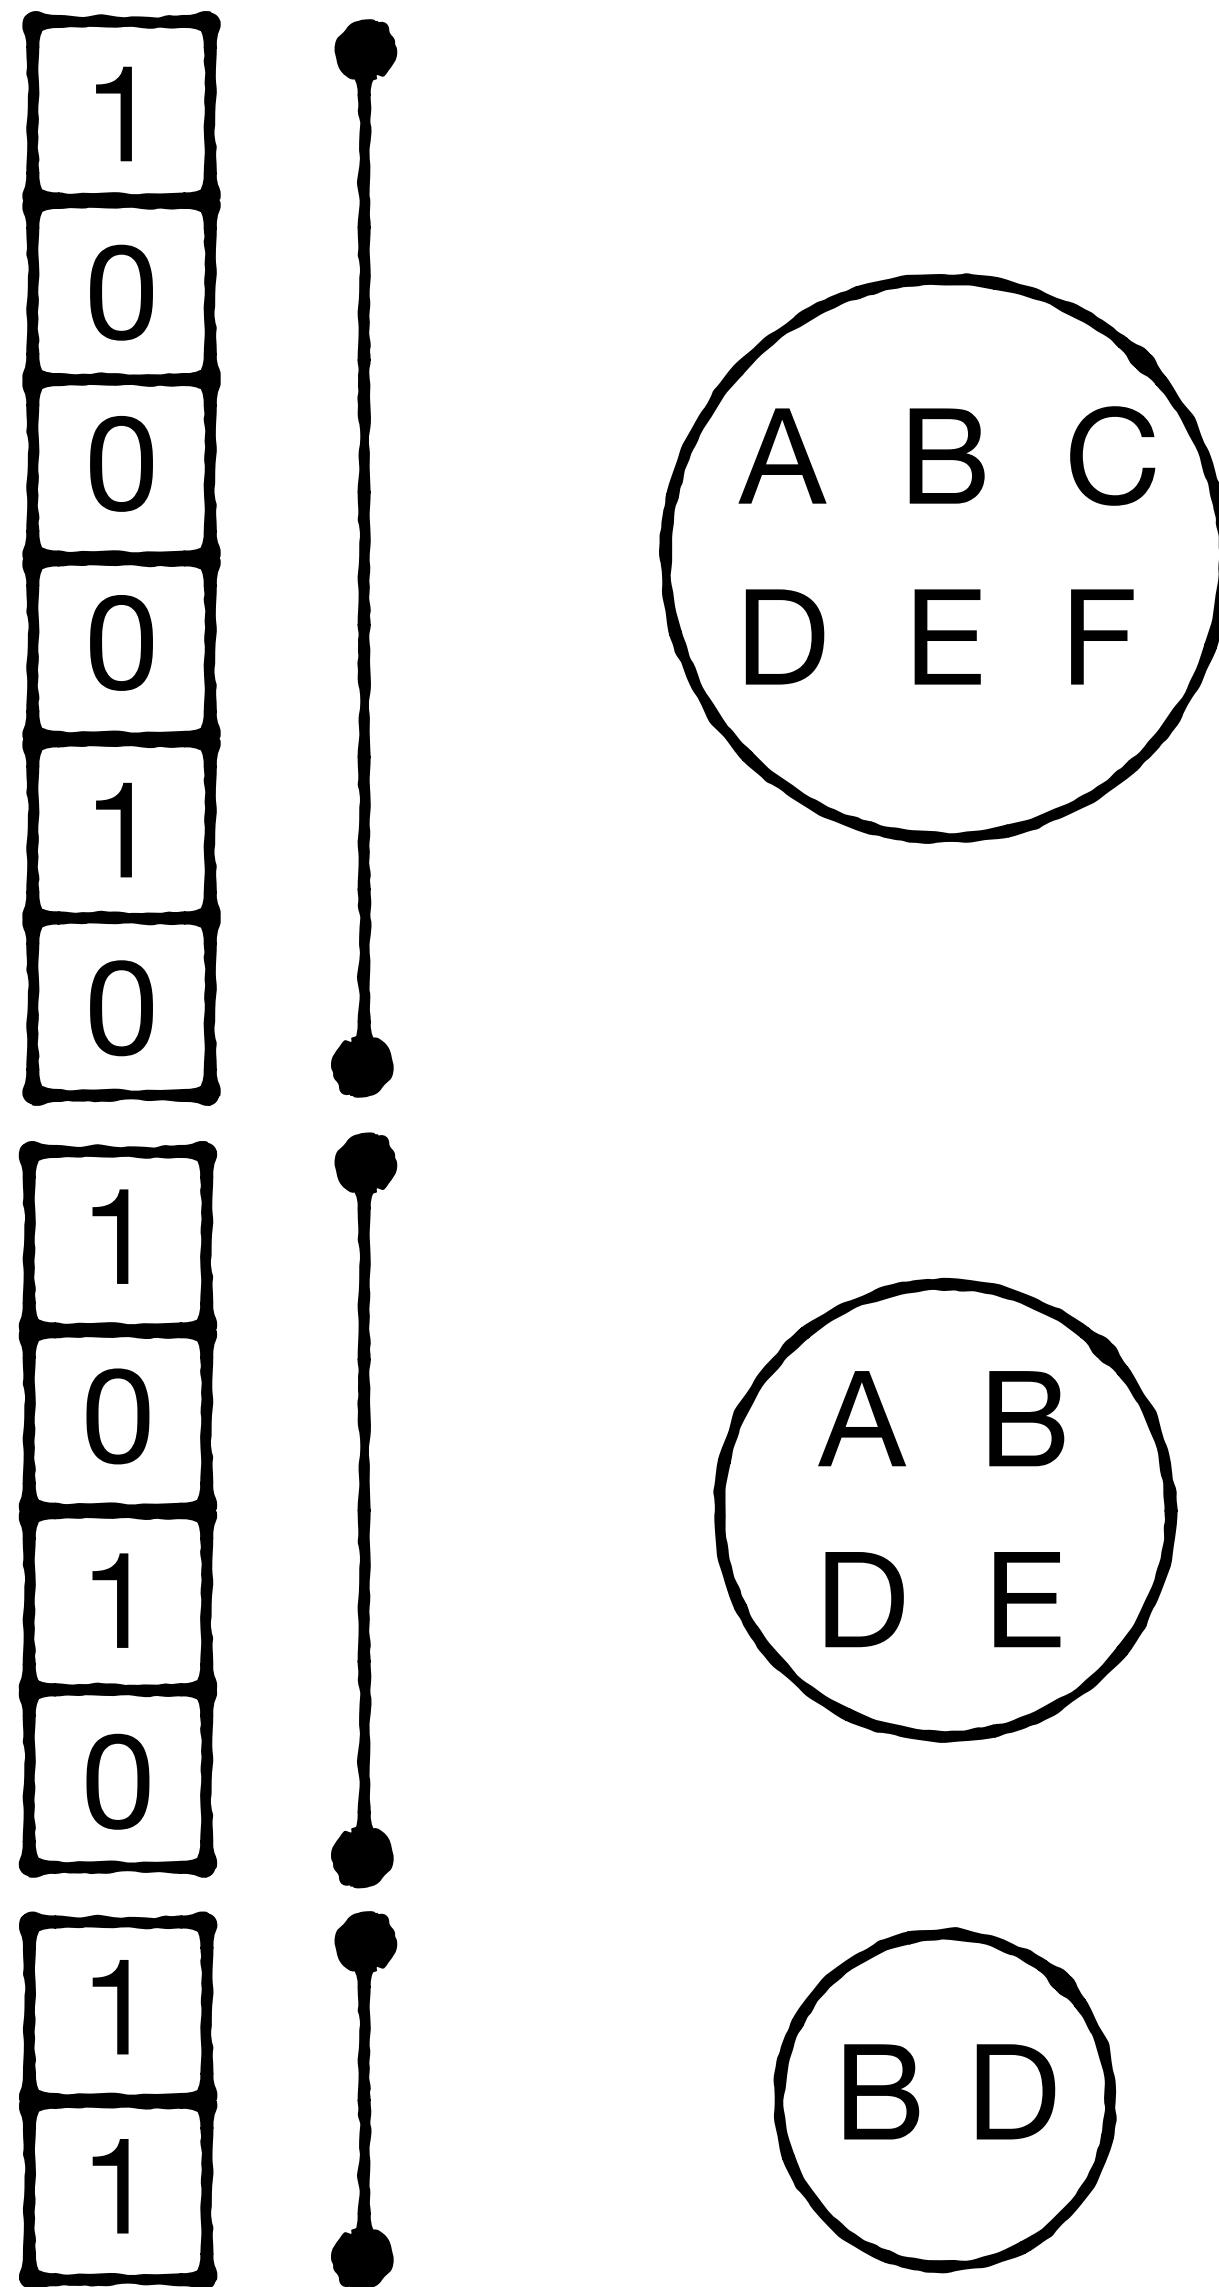
trials across all entries until success

$N \cdot e$ bits



Total bitmap size?

trials across all entries until success



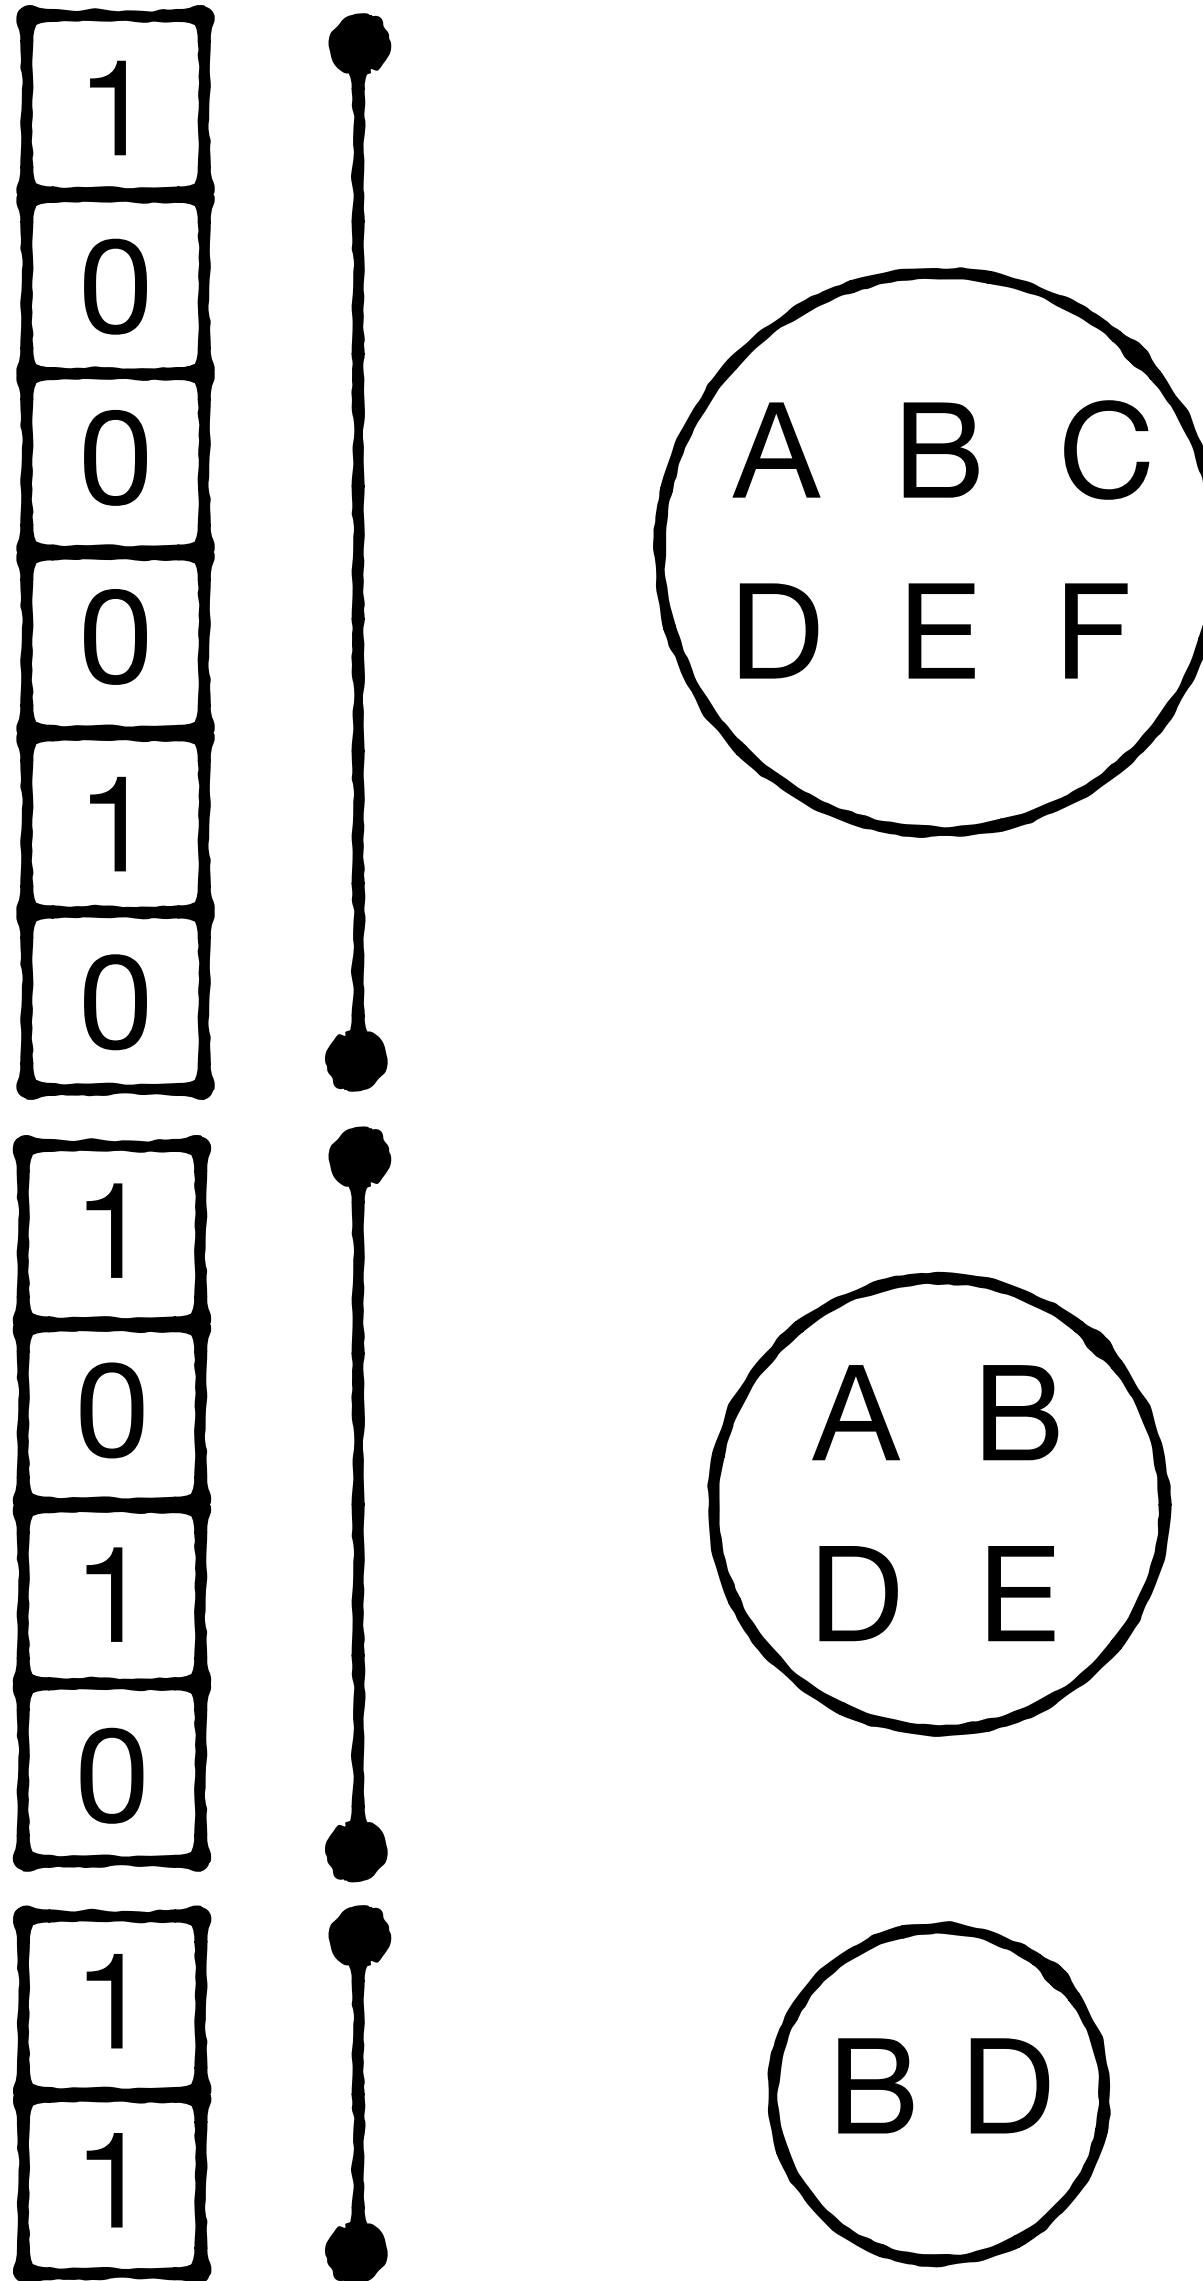
$N \cdot e$



Also our construction time :)

Total bitmap size?

trials across all entries until success



$N \cdot e$



Also our construction time :)

times each entry is hashed

Worst Case Query cost?

1
0
0
0
1
0
1
0
1
1
0
1
1

Worst Case Query cost?

1
0
0
0
1
0
1
0
1
0
1
0
1
1

= # bitmaps to traverse

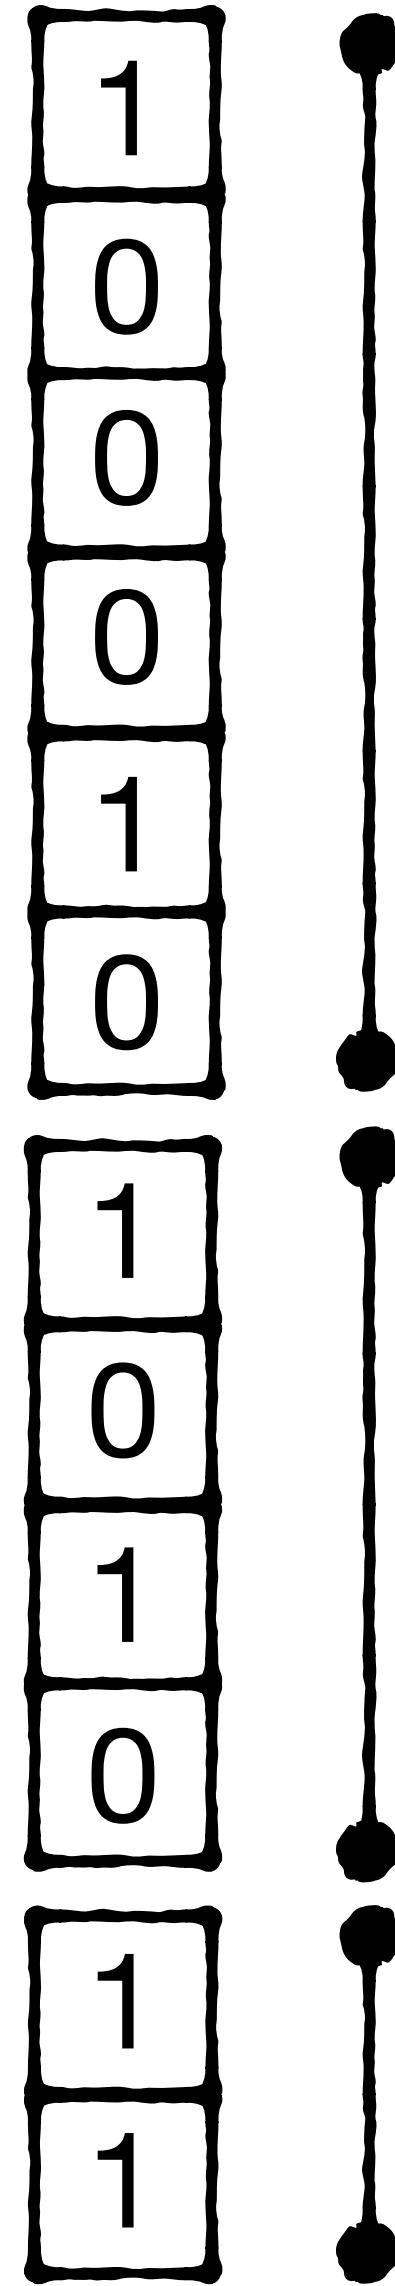
Worst Case Query cost?

1
0
0
0
1
0
1
0
1
0
1
0
1
1

= # bitmaps to traverse

= $\log_{\text{base}} (N)$

Worst Case Query cost?



= # bitmaps to traverse

= $\log_{\text{base}} (N)$

↑
?

Worst Case Query cost?

1
0
0
0
1
0
1
0
1
0
1
0
1
1

= # bitmaps to traverse

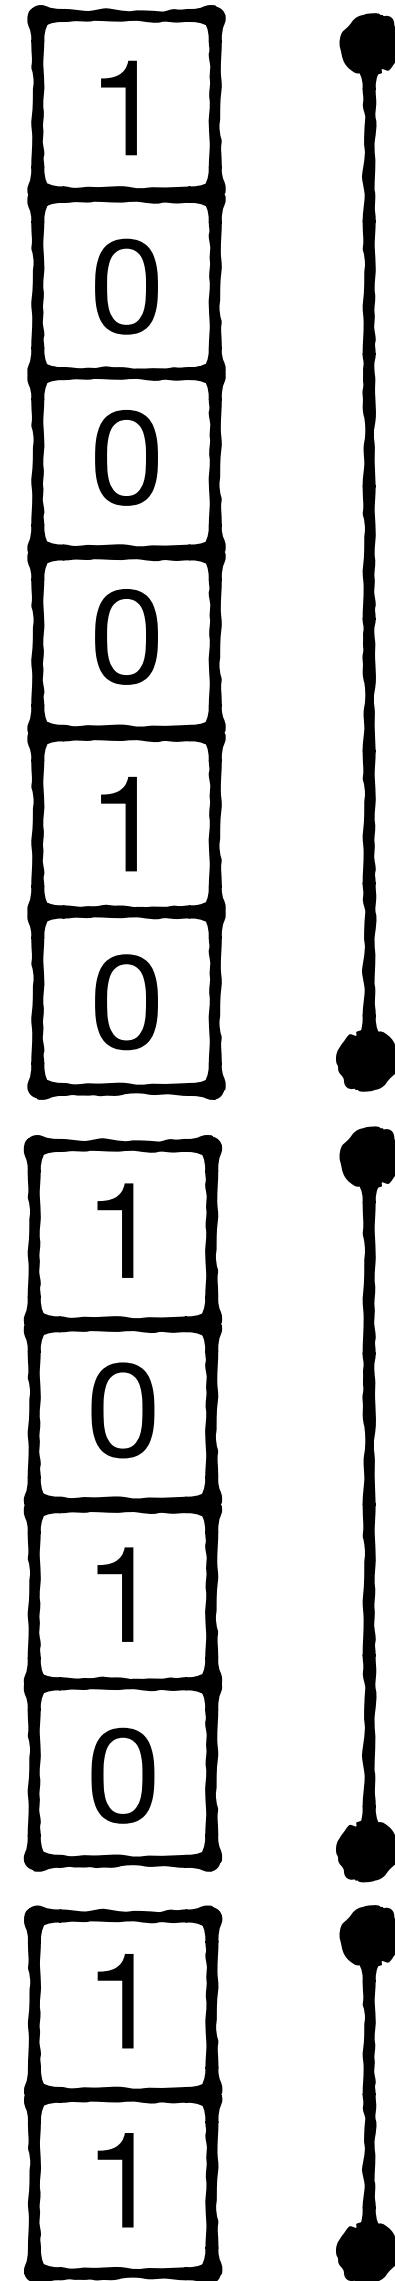
= $\log_{\text{base}}(N)$

Fraction of entries failing each iteration

$$\rightarrow \frac{1}{1 - e^{-1}}$$

↑

Worst Case Query cost?



= # bitmaps to traverse

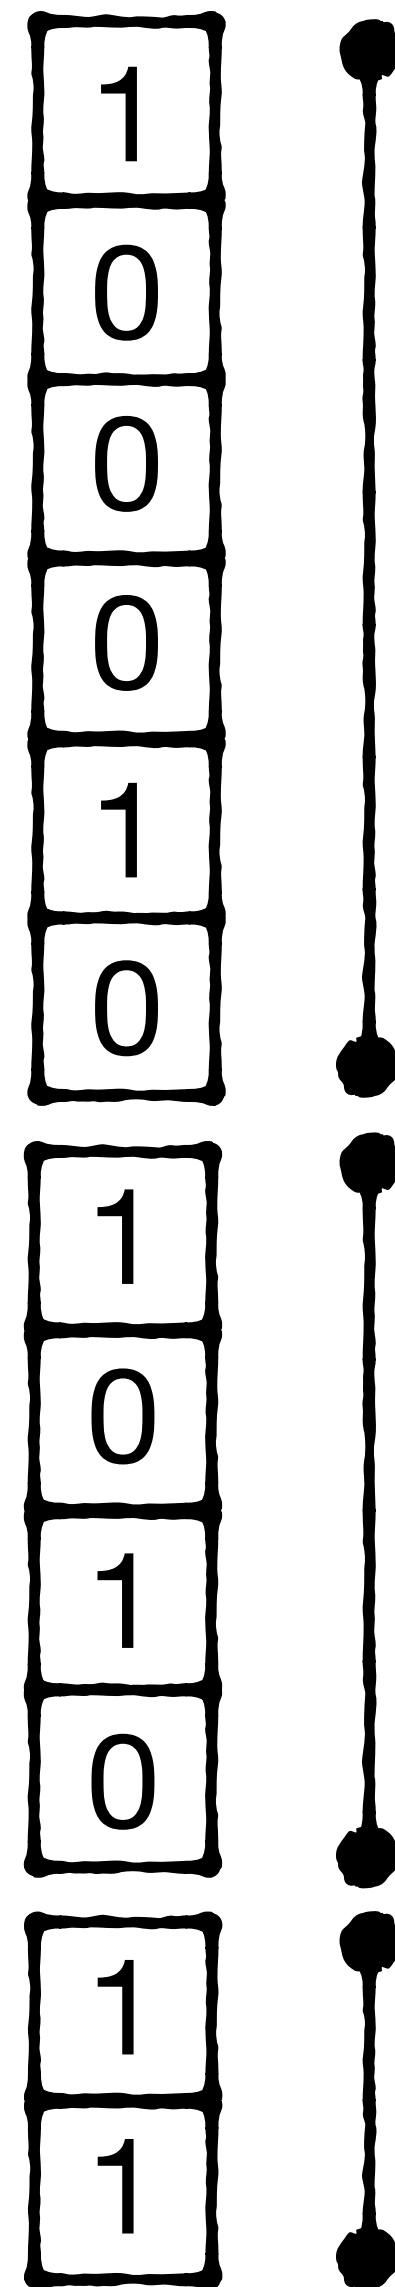
= $\log_{\text{base}} (N)$



1.58

Worst Case Query cost?

$$= \log 1.58 (N)$$



Worst Case Query cost?

$$= \log 1.58 (N)$$

Is it really so bad?

1
0
0
0
1
0
1
0
1
0
1
1
0
1
1

Is it really so bad?

1
0
0
0
1
0
1
0
1
0
1
0
1
1

← Entries that hit 1 terminate immediately.

← Only for very few entries we must go to end

expected worst-case query cost?

1
0
0
0
1
0

1
0
1
0

1
1

expected worst-case query cost?

1
0
0
0
1
0
1
0
1
0
1
1
0
1
1

1 access to first bit map

expected worst-case query cost?

1
0
0
0
1
0
1
0
1
1
0
1
0
1
1

1 access to first bit map

$(1 - e^{-1})$ chance to access second

expected worst-case query cost?

1
0
0
0
1
0

1 access to first bit map

1
0
1
0

$(1 - e^{-1})$ chance to access second

1
1

$(1 - e^{-1})^2$ chance to access third

expected worst-case query cost?

1
0
0
0
1
0

1
0
1
0
1
0

1
1

$$= 1 + (1 - e^{-1}) + (1 - e^{-1})^2 + (1 - e^{-1})^3 + \dots$$

expected worst-case query cost?

1
0
0
0
1
0
1
0
1
0
1
1

$$= 1 + (1 - e^{-1}) + (1 - e^{-1})^2 + (1 - e^{-1})^3 + \dots$$

$$= \frac{1}{e^{-1}} = 2.72 = O(1)$$

1
0
0
0
1
0

Memory (bits)

$N \cdot e$

worst-case query cost

$O(\log N)$

1
0
1
0

expected worst-case query cost

$O(1)$

1
1

Construction

$O(N)$

Is this the best we can do?

1
0
0
0
1
0

Memory (bits)

1
0
1
0

worst-case query cost

1
1

Construction



$N \cdot e$

$O(\log N)$

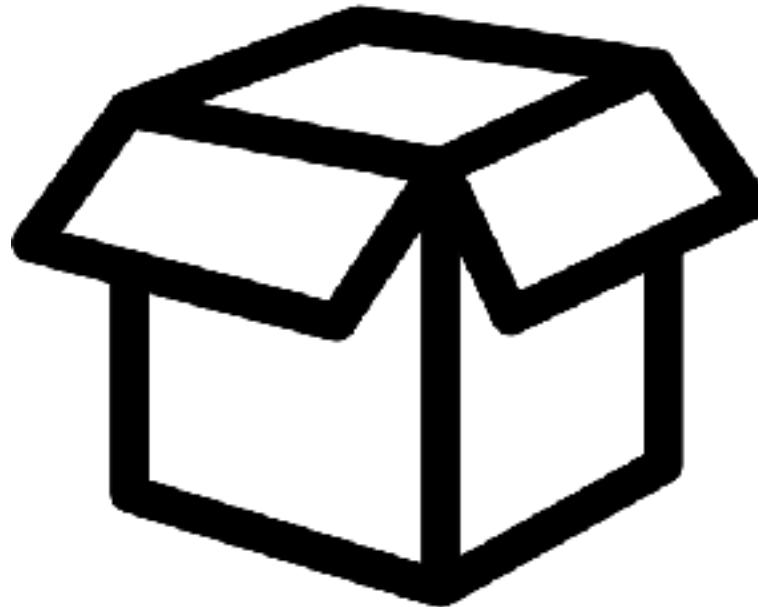
$O(1)$

$O(N)$

Lower Bound for Minimal Perfect Hashing (MPH)

Lower Bound for Minimal Perfect Hashing (MPH)

**Assume nothing
about implementation**

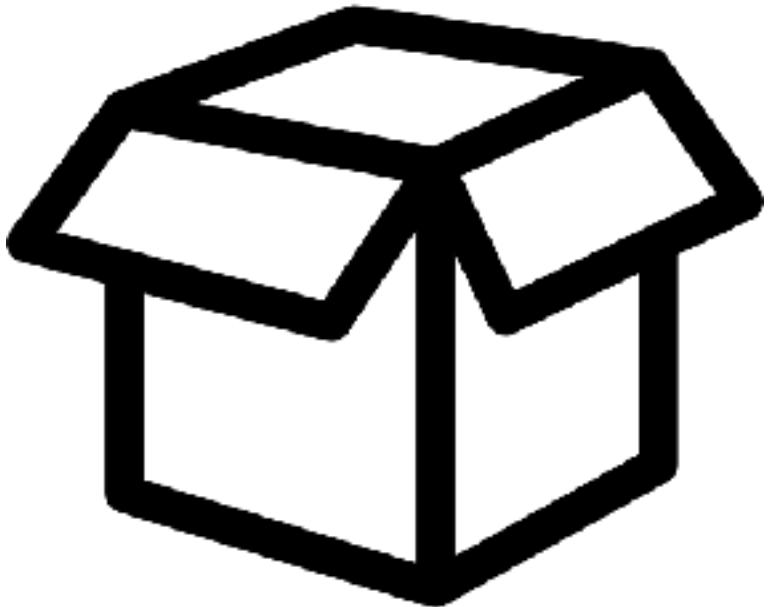


**Analyze with respect
to specification**



Lower Bound for Minimal Perfect Hashing (MPH)

Assume nothing
about implementation



Analyze with respect
to specification



$N - \#$ entries
Bijective (one-to-one)

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{IMPHI} + ???| \geq |\text{Permutation}|$$

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + \text{???} \geq |\text{Permutation}|$$



What data must we add to transform MPH into permutation? :)

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + \text{???} \geq |\text{Permutation}|$$



What data must we add to transform MPH into permutation? :)

The data itself! $N \cdot \log_2(N)$

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + N \cdot \log_2(N) \geq |\text{Permutation}|$$

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + N \cdot \log_2(N) \geq |\text{Permutation}|$$



How big is this?

Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + N \cdot \log_2(N) \geq \log_2(N!)$$



How big is this?

Lower Bound for Minimal Perfect Hashing (MPH)

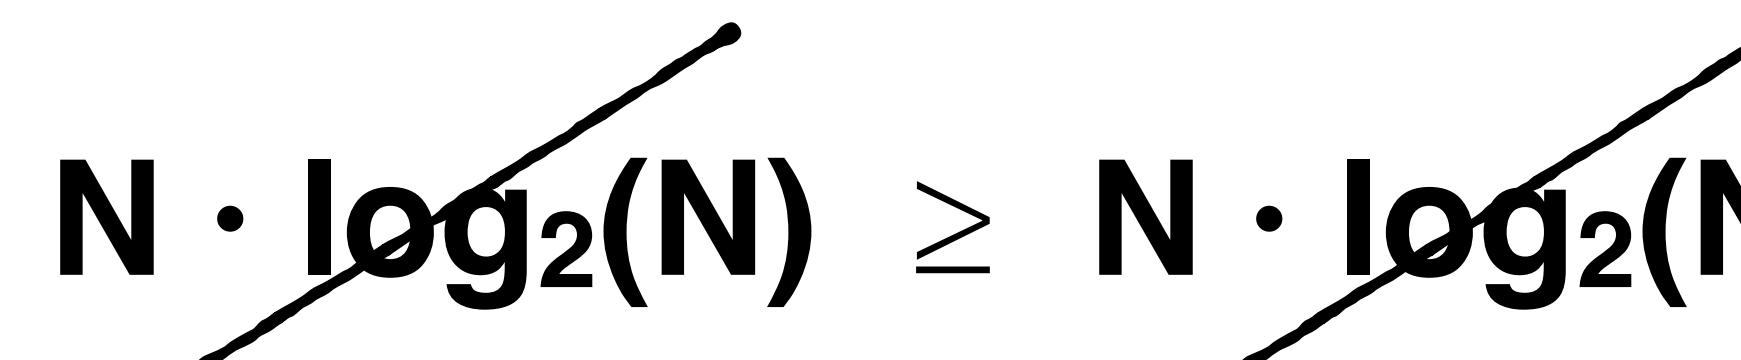
$$|\text{MPH}| + N \cdot \log_2(N) \geq N \cdot \log_2(N) + N \cdot \log_2(e)$$



By Stirling's approximation



Lower Bound for Minimal Perfect Hashing (MPH)

$$|\text{MPH}| + N \cdot \log_2(N) \geq N \cdot \log_2(N) + N \cdot \log_2(e)$$


Lower Bound for Minimal Perfect Hashing (MPH)

$$|MPHI| \geq N \log_2(e)$$

We're done :)

Lower Bound

$$N \cdot \log_2(e)$$

Fingerprinting

$$N \cdot e$$

Lower Bound

$$N \cdot \log_2(e)$$

1.44 bits / key

Fingerprinting

$$N \cdot e$$

2.71 bits / key

Lower Bound

$$N \cdot \log_2(e)$$

Fingerprinting

$$N \cdot e$$

Not far off, but also not there...

Lower Bound

$$N \cdot \log_2(e)$$

Fingerprinting

$$N \cdot e$$

Not far off, but also not there...

Other methods push memory footprint lower :)

Perfect Hashing



Minimal

space-efficient
static data



Dynamic

more space
supports updates

Dynamic Perfect Hashing

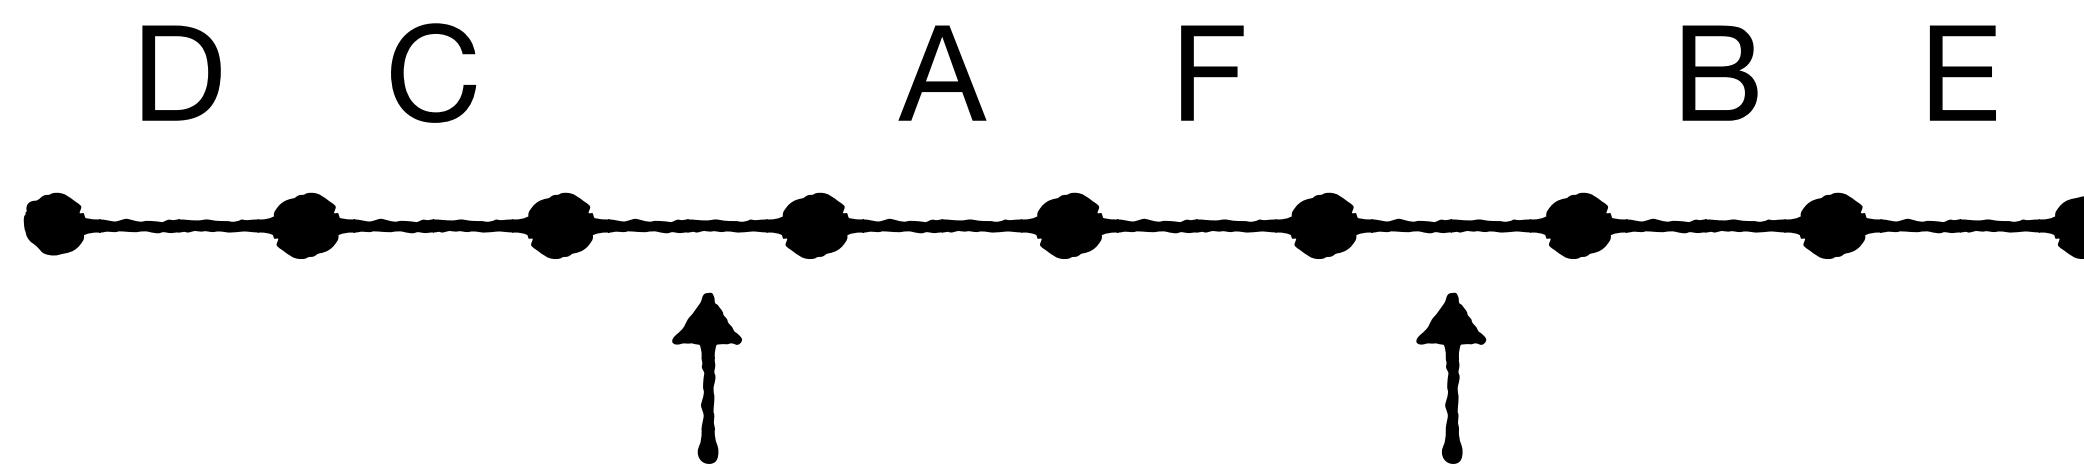
N keys

A B C D E F



more than N slots

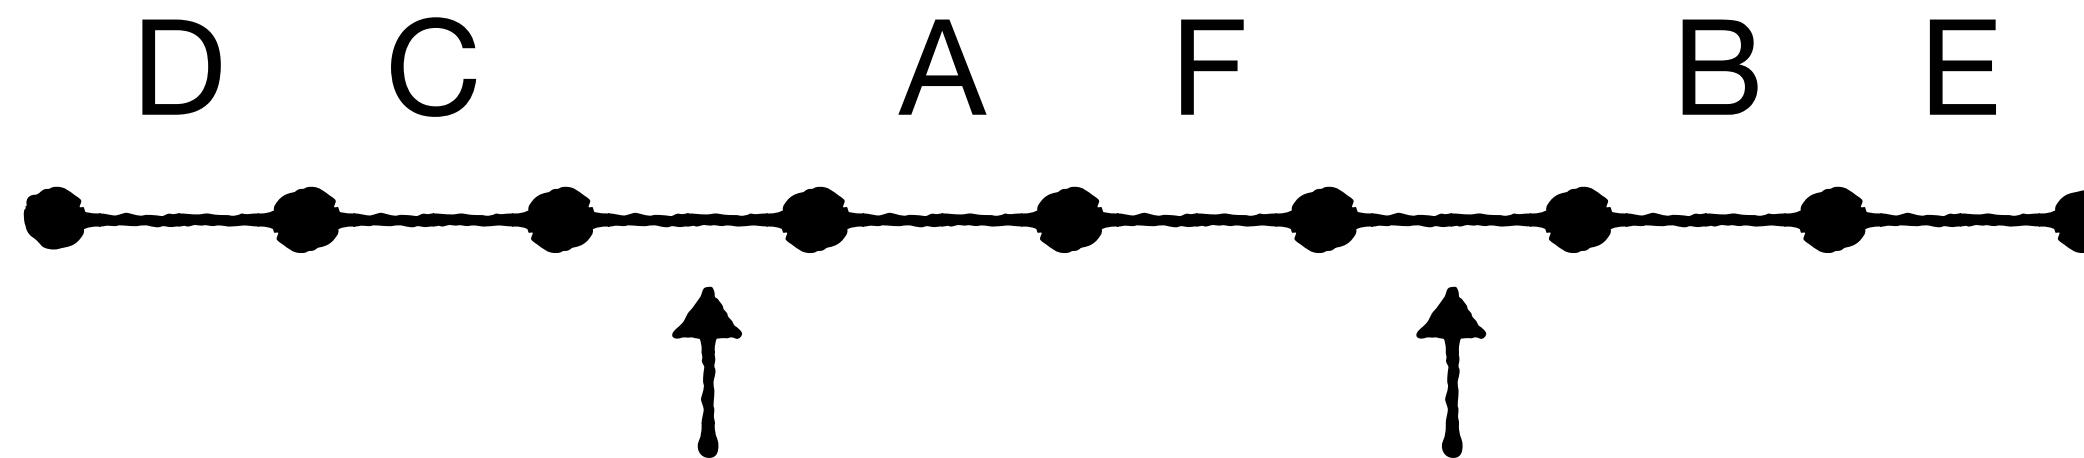
Dynamic Perfect Hashing



Some slots can stay free

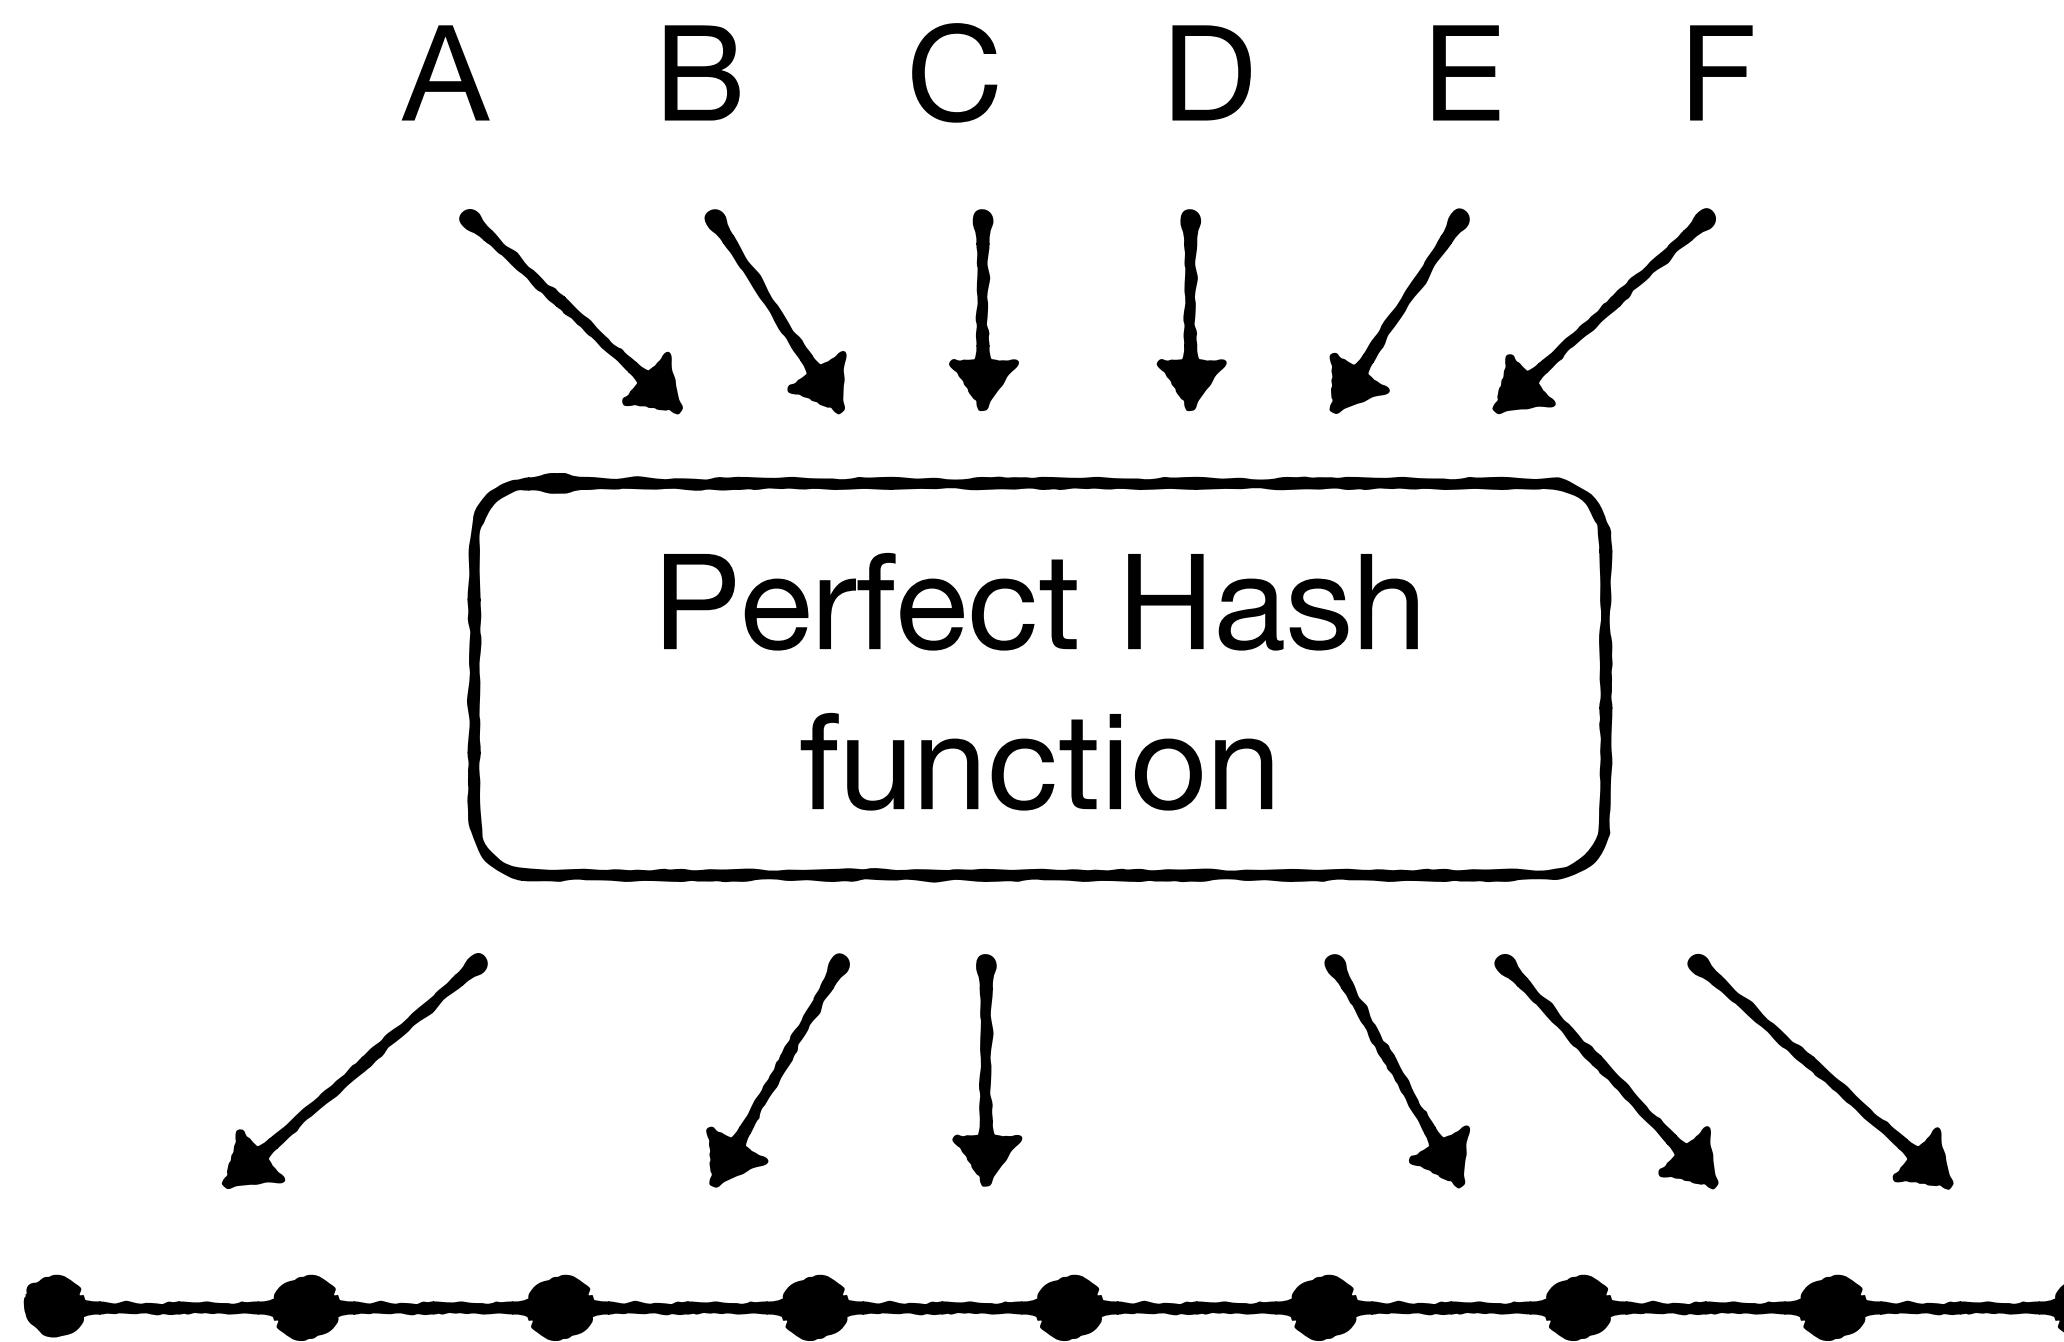
Dynamic Perfect Hashing

More flexibility :)

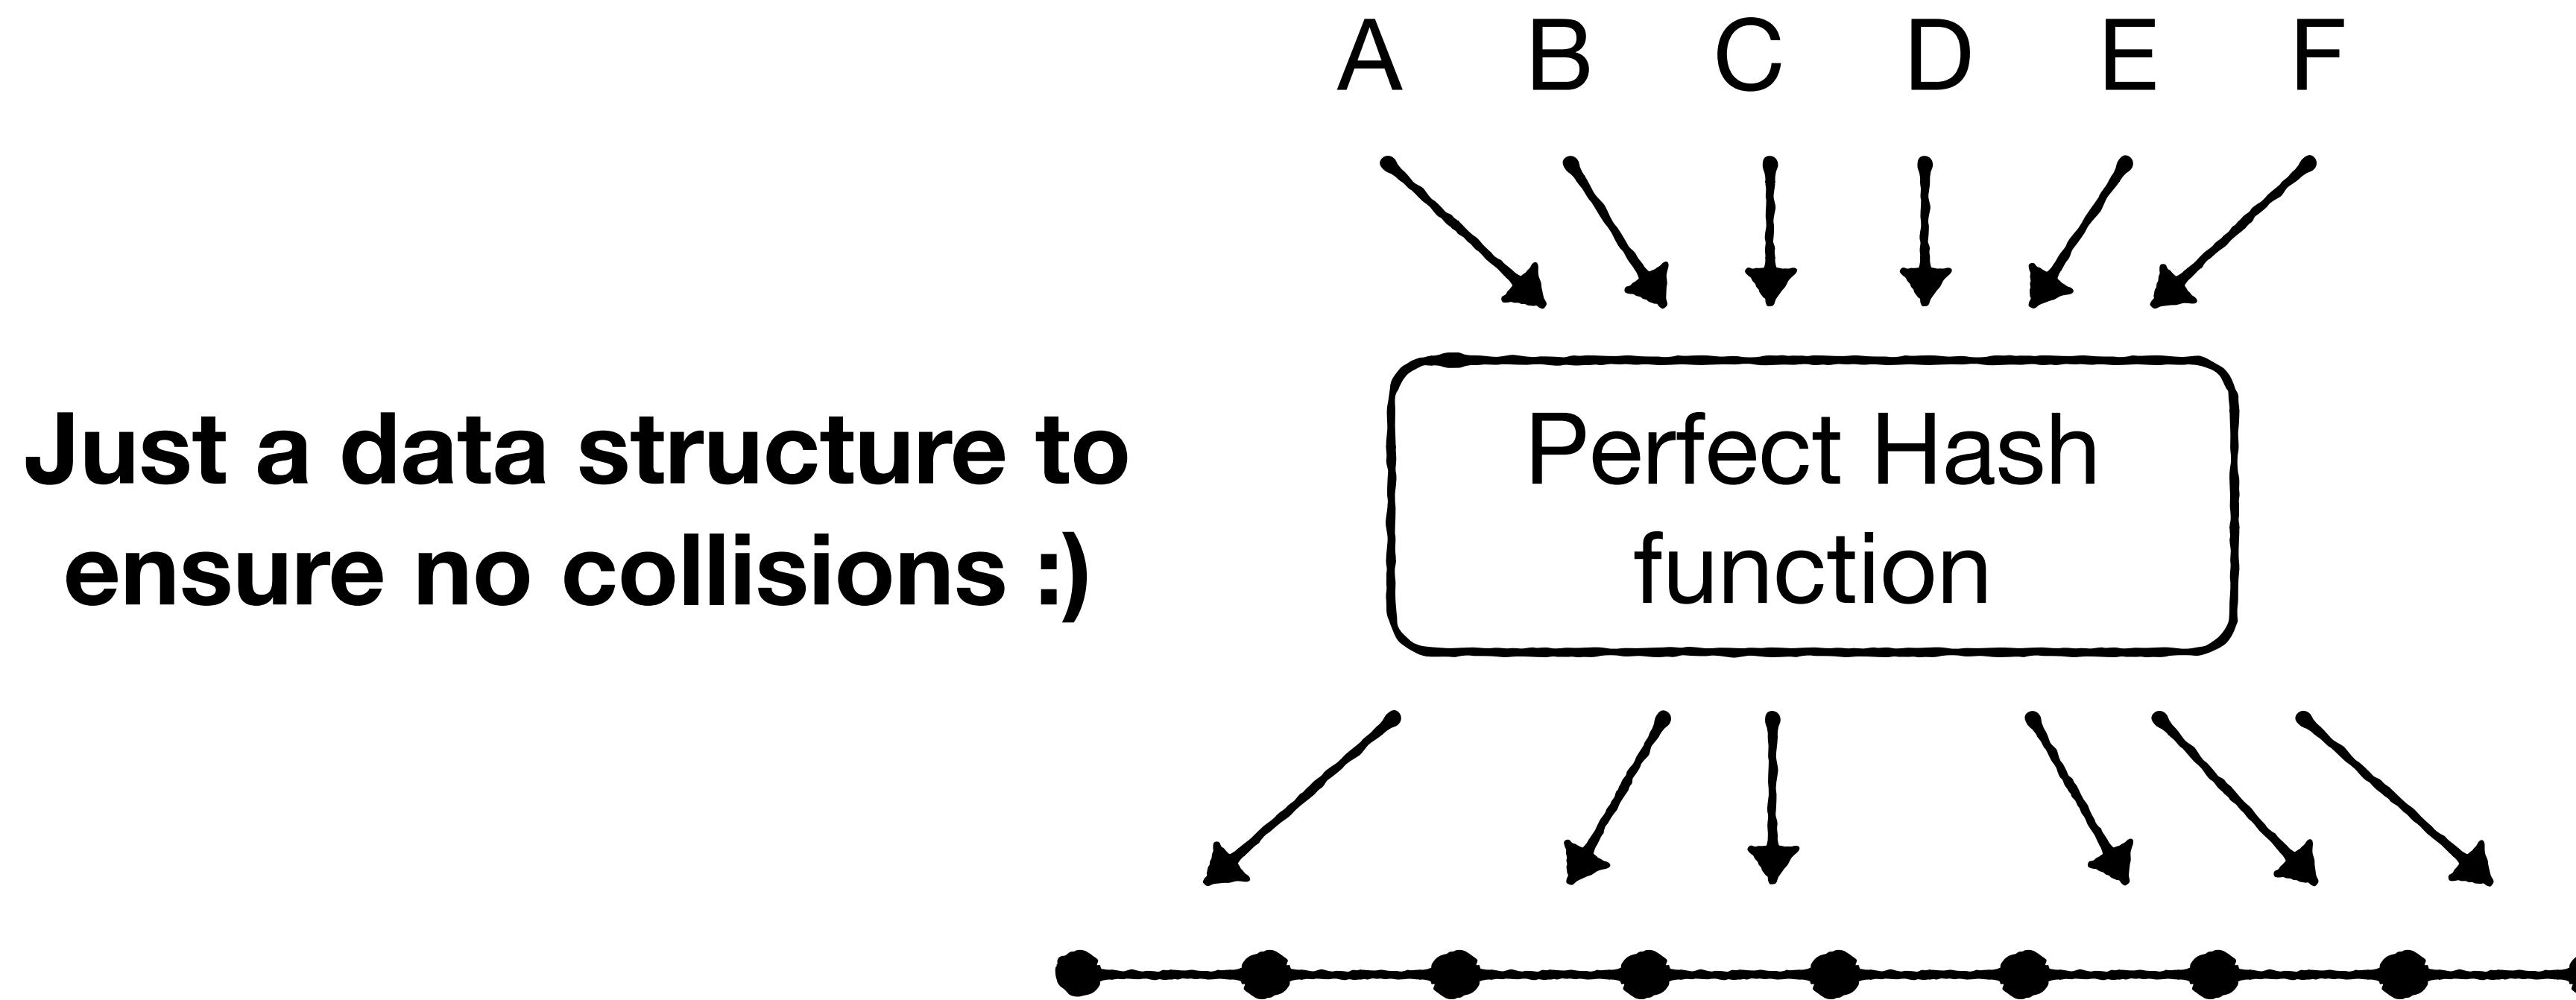


Some slots can stay free

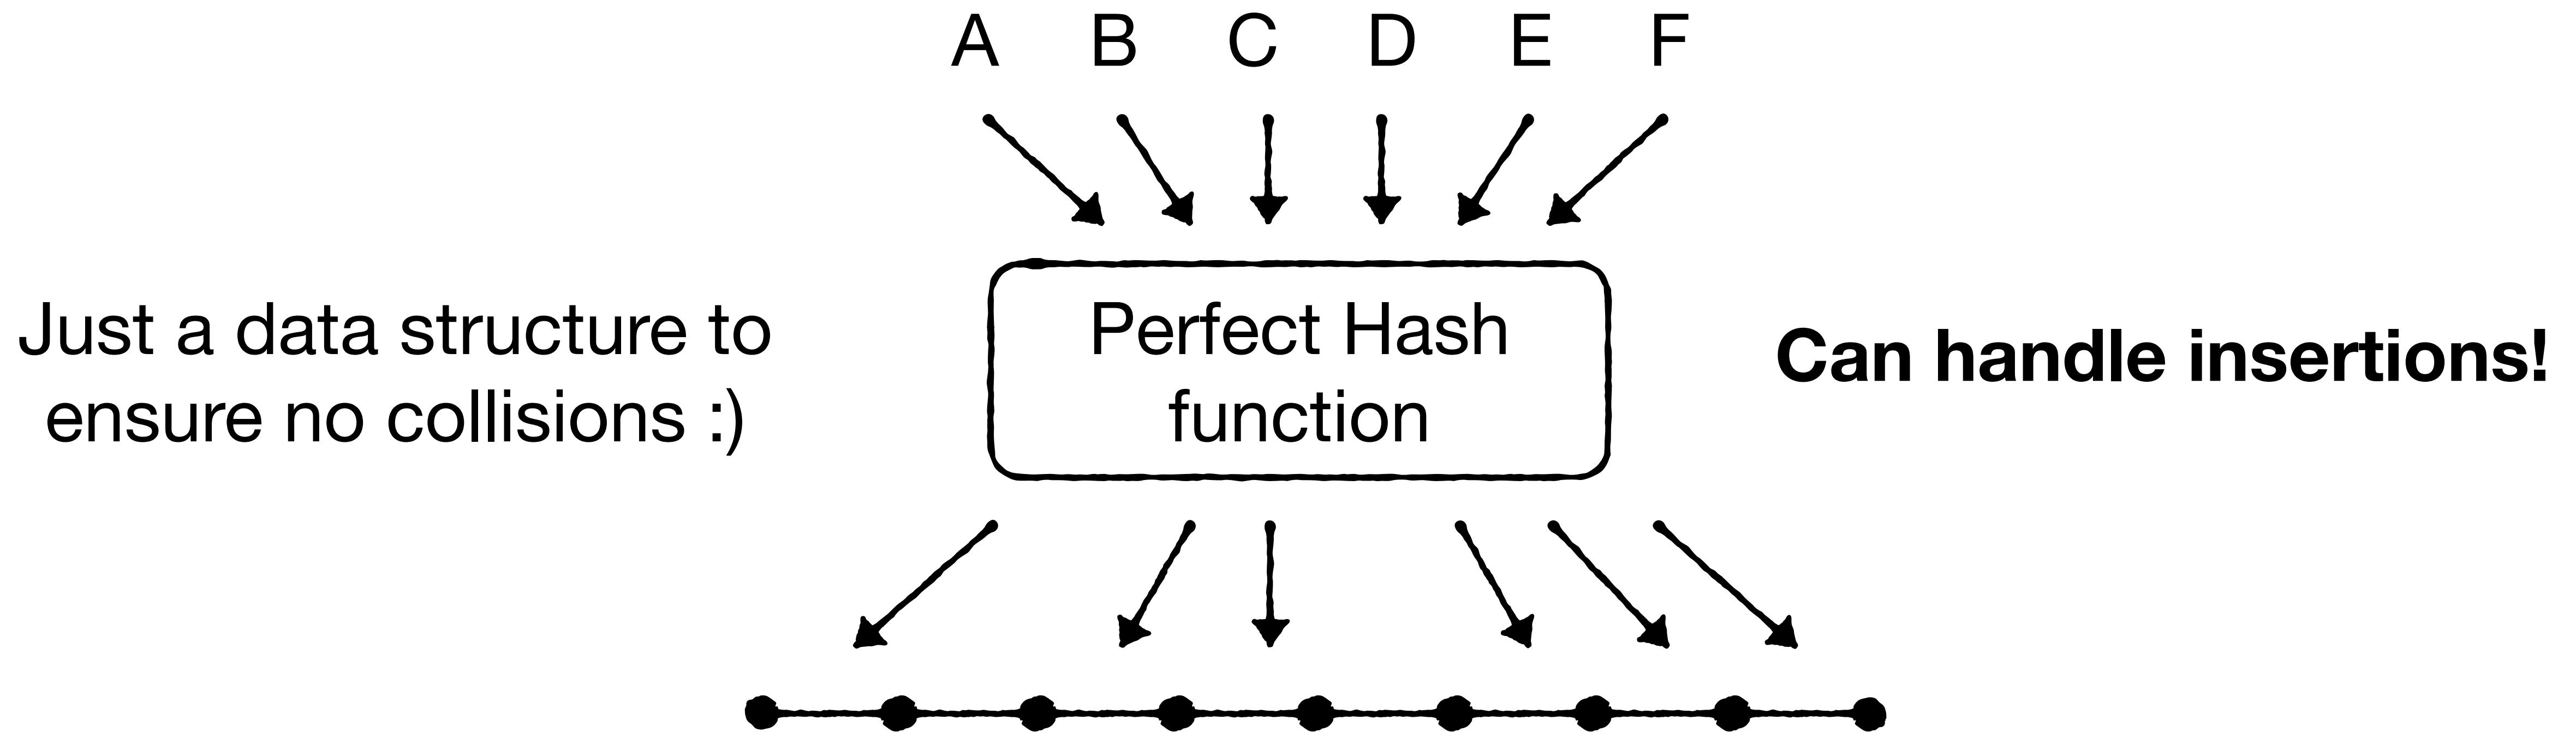
Dynamic Perfect Hashing



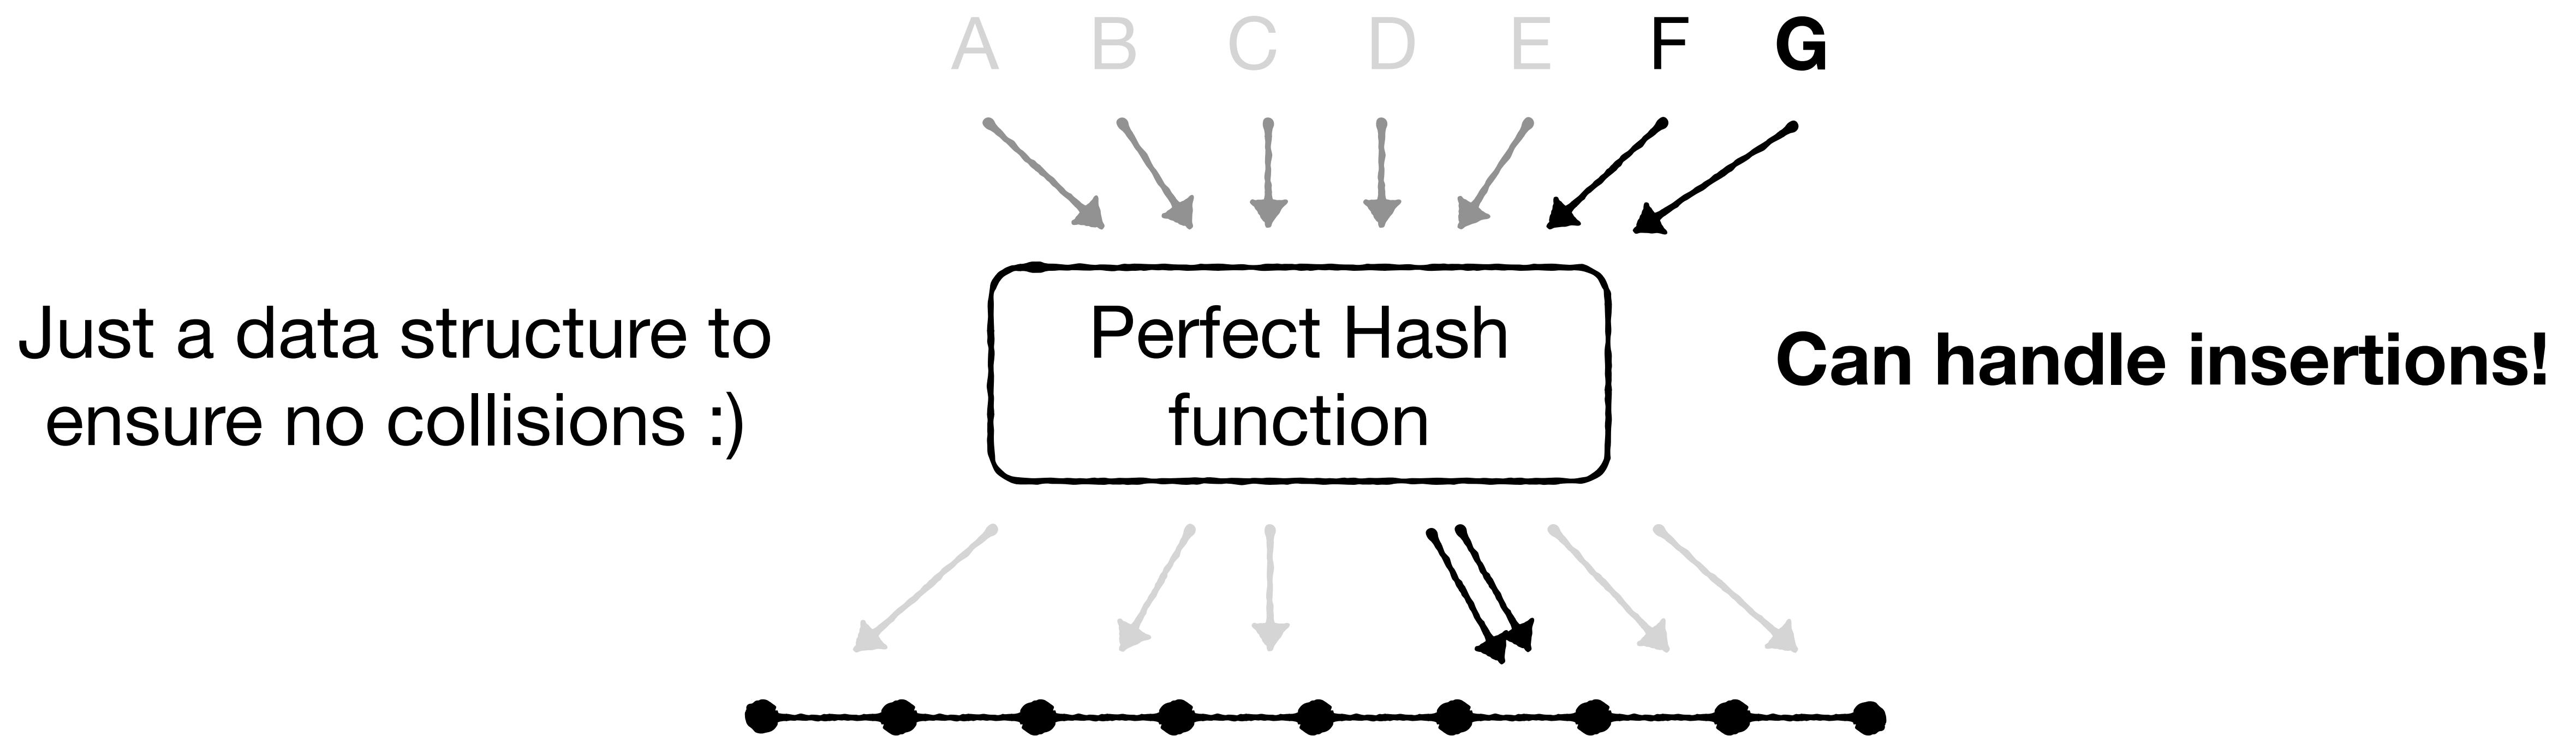
Dynamic Perfect Hashing



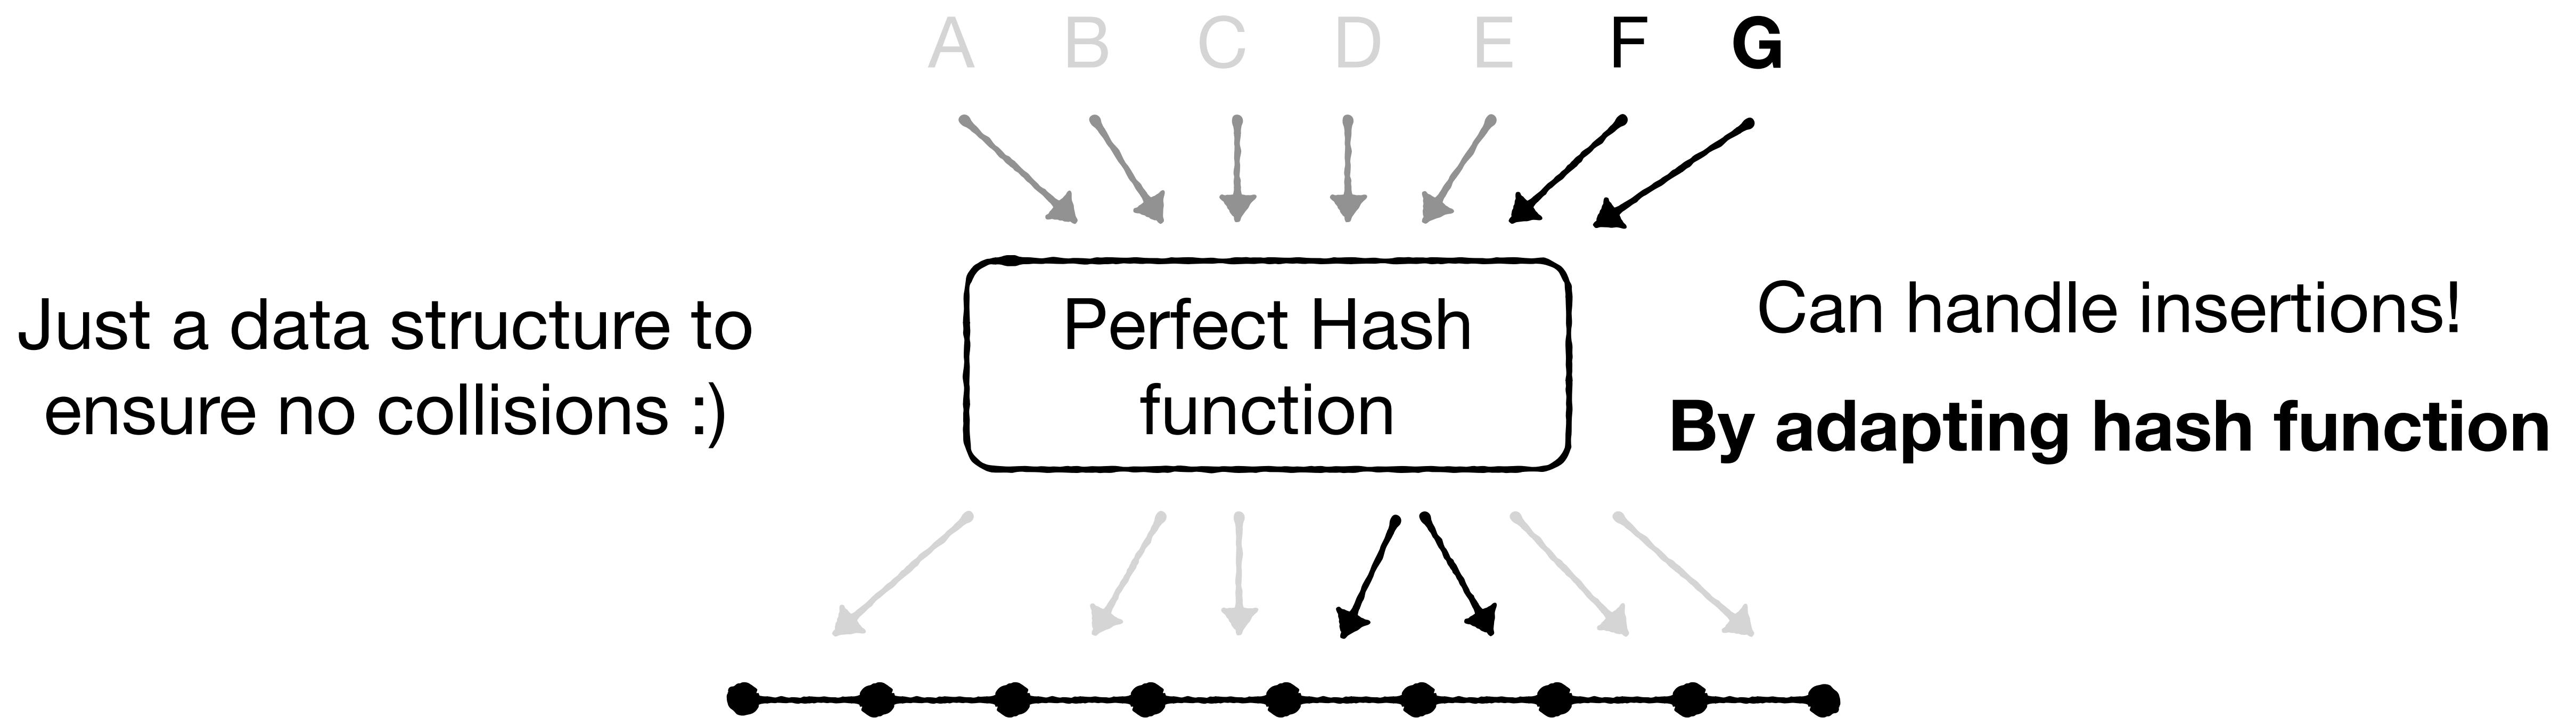
Dynamic Perfect Hashing



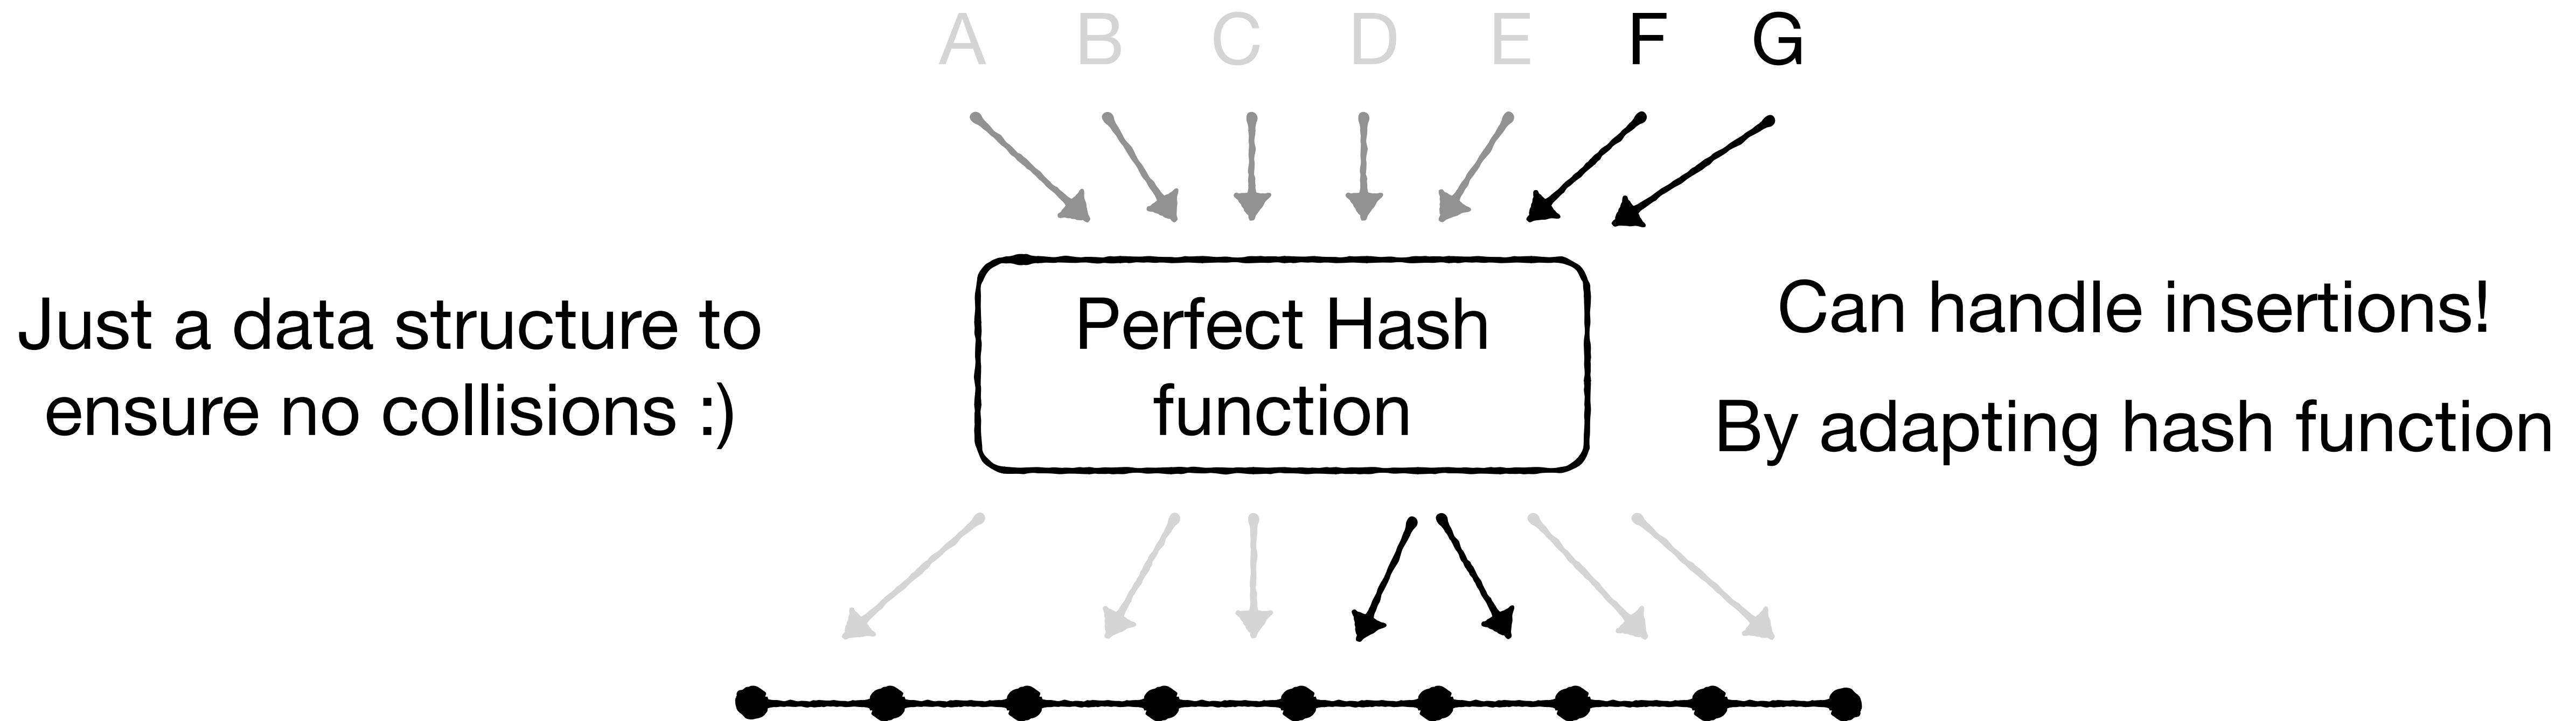
Dynamic Perfect Hashing



Dynamic Perfect Hashing



Dynamic Perfect Hashing



Catch: must be able to fetch original keys to resolve new collisions

Dynamic Perfect Hashing

Storing a Sparse Table with $O(1)$ Worst Case Access Time. JACM 1984.

ML Fredman, J Komlós, E Szemerédi

The End of Moore's Law and the Rise of the Data Processor. VLDB 2021.

Niv Dayan, Moshe Twitto, Yuval Rochman, Uri Bejtler, Itai Ben Zion, Edward Bortnikov, Shmuel Dashevsky, Ofer Frishman, Evgeni Ginzburg, Igal Maly, Avraham Meir, Mark Mokry, Iddo Naiss, Noam Rabinovich

Many more...

Dynamic Perfect Hashing

The End of Moore's Law and the Rise of the Data Processor. VLDB 2021.

Niv Dayan, Moshe Twitto, Yuval Rochman, Uri Beiter, Itai Ben Zion, Edward Bortnikov, Shmuel Dashevsky, Ofer Frishman, Evgeni Ginzburg, Igal Maly, Avraham Meir, Mark Mokry, Iddo Naiss, Noam Rabinovich

space-efficient, used in practice

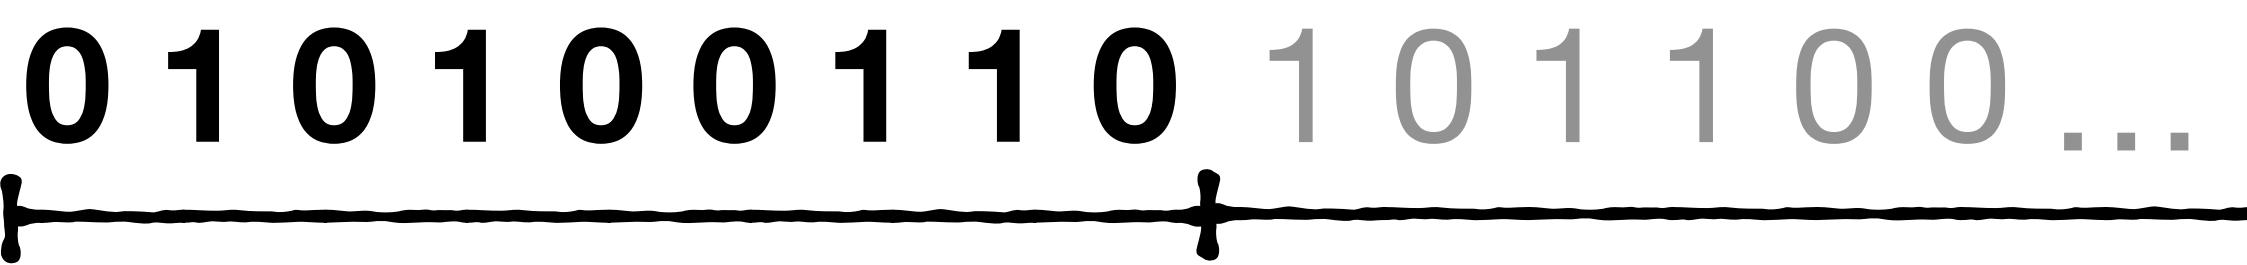
Delta Hash Table

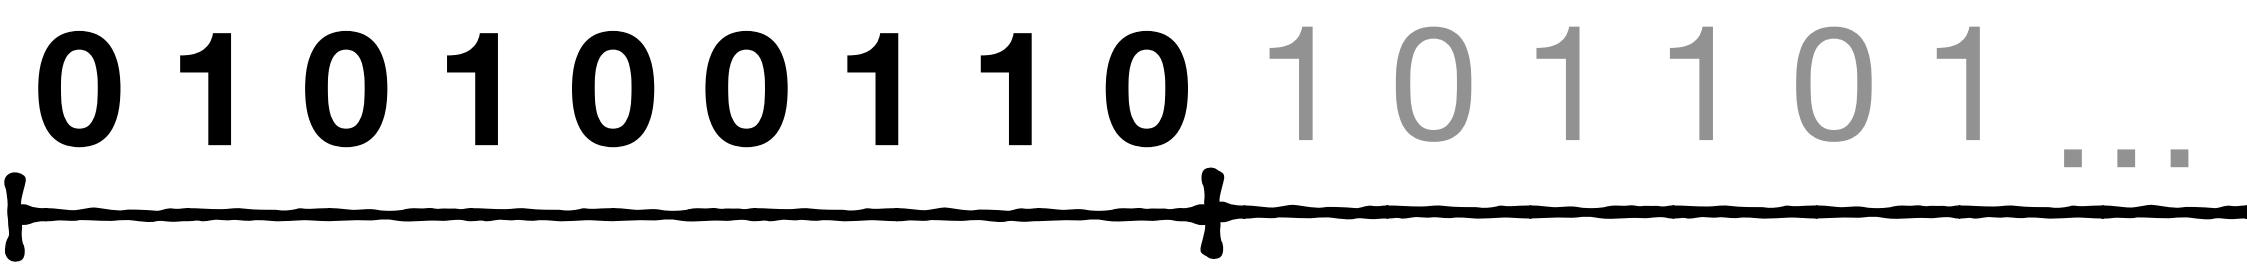
hash(X) = 0 1 0 1 0 0 1 1 0 1 0 1 1 0 0 ...

hash(Y) = 0 1 0 1 0 0 1 1 0 1 0 1 1 0 1 ...



Delta Hash Table

hash(X) =  0 1 0 1 0 0 1 1 0 1 0 1 1 0 0 ...

hash(Y) =  0 1 0 1 0 0 1 1 0 1 0 1 1 0 1 ...



Same slot

Delta Hash Table

hash(X) = 1 0 1 1 0 0 ...

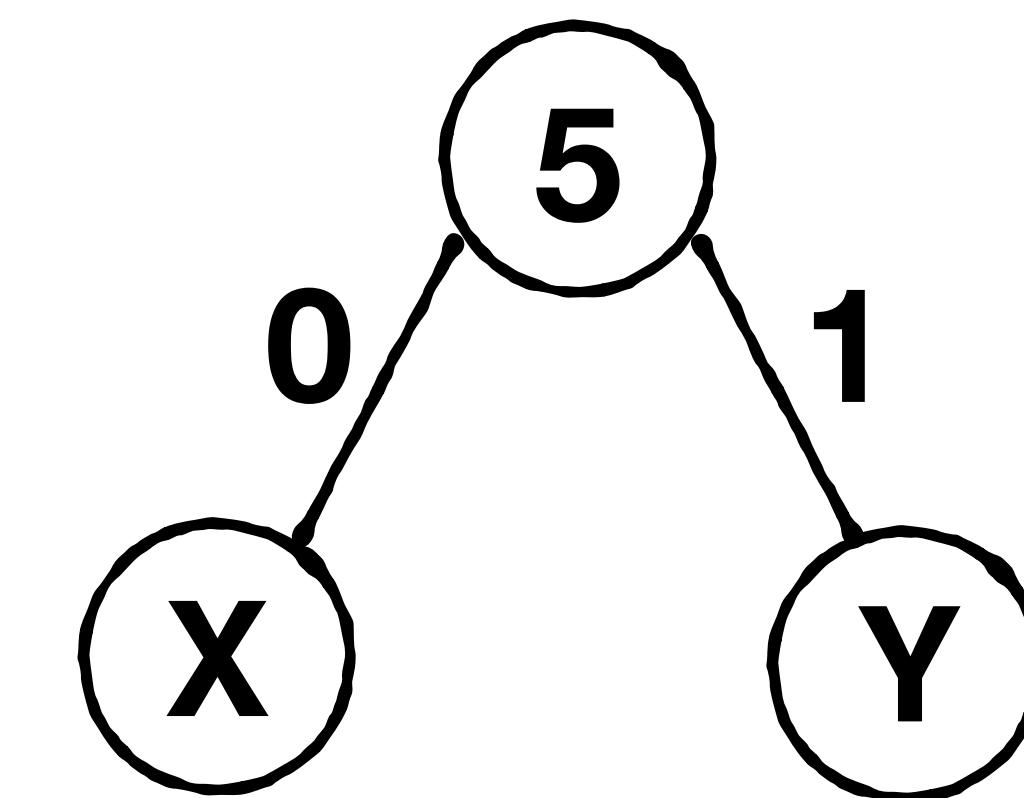
hash(Y) = 1 0 1 1 0 1 ...



Build Trie capturing index of first different bit

hash(X) = 1 0 1 1 0 0 ...

hash(Y) = 1 0 1 1 0 1 ...

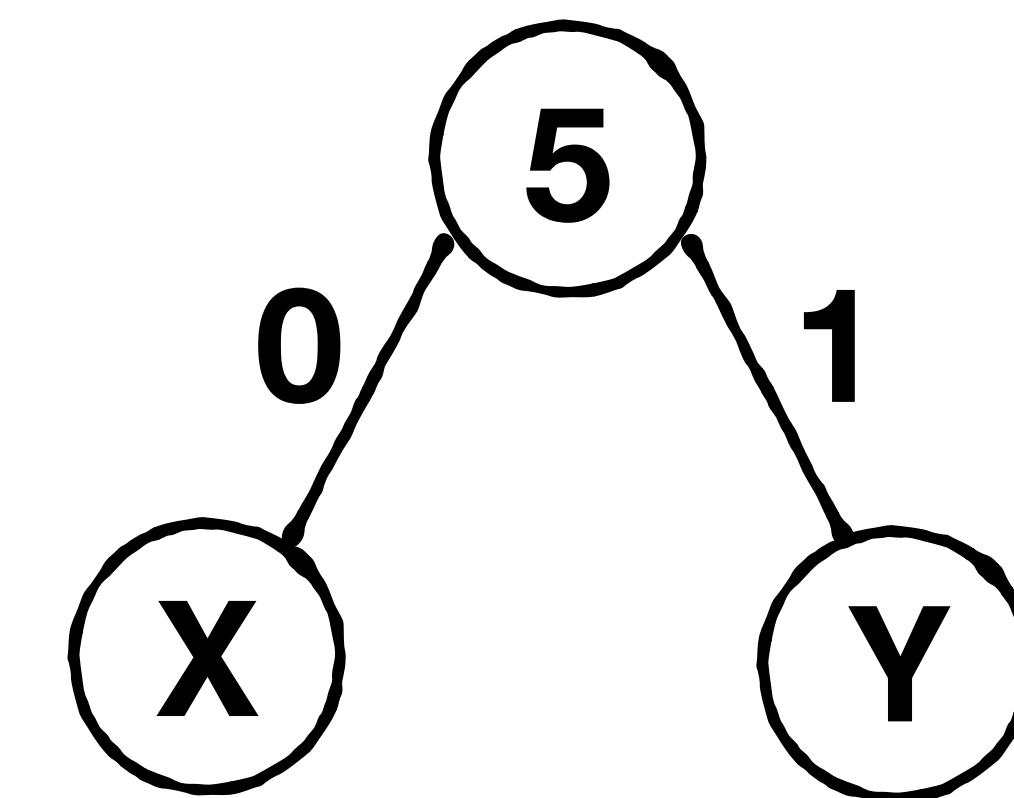


Build Trie capturing index of first different bit

hash(X) =  1 0 1 1 0 0 ...

hash(Y) =  1 0 1 1 0 1 ...

hash(Z) =  0 1 0 0 1 1 ...

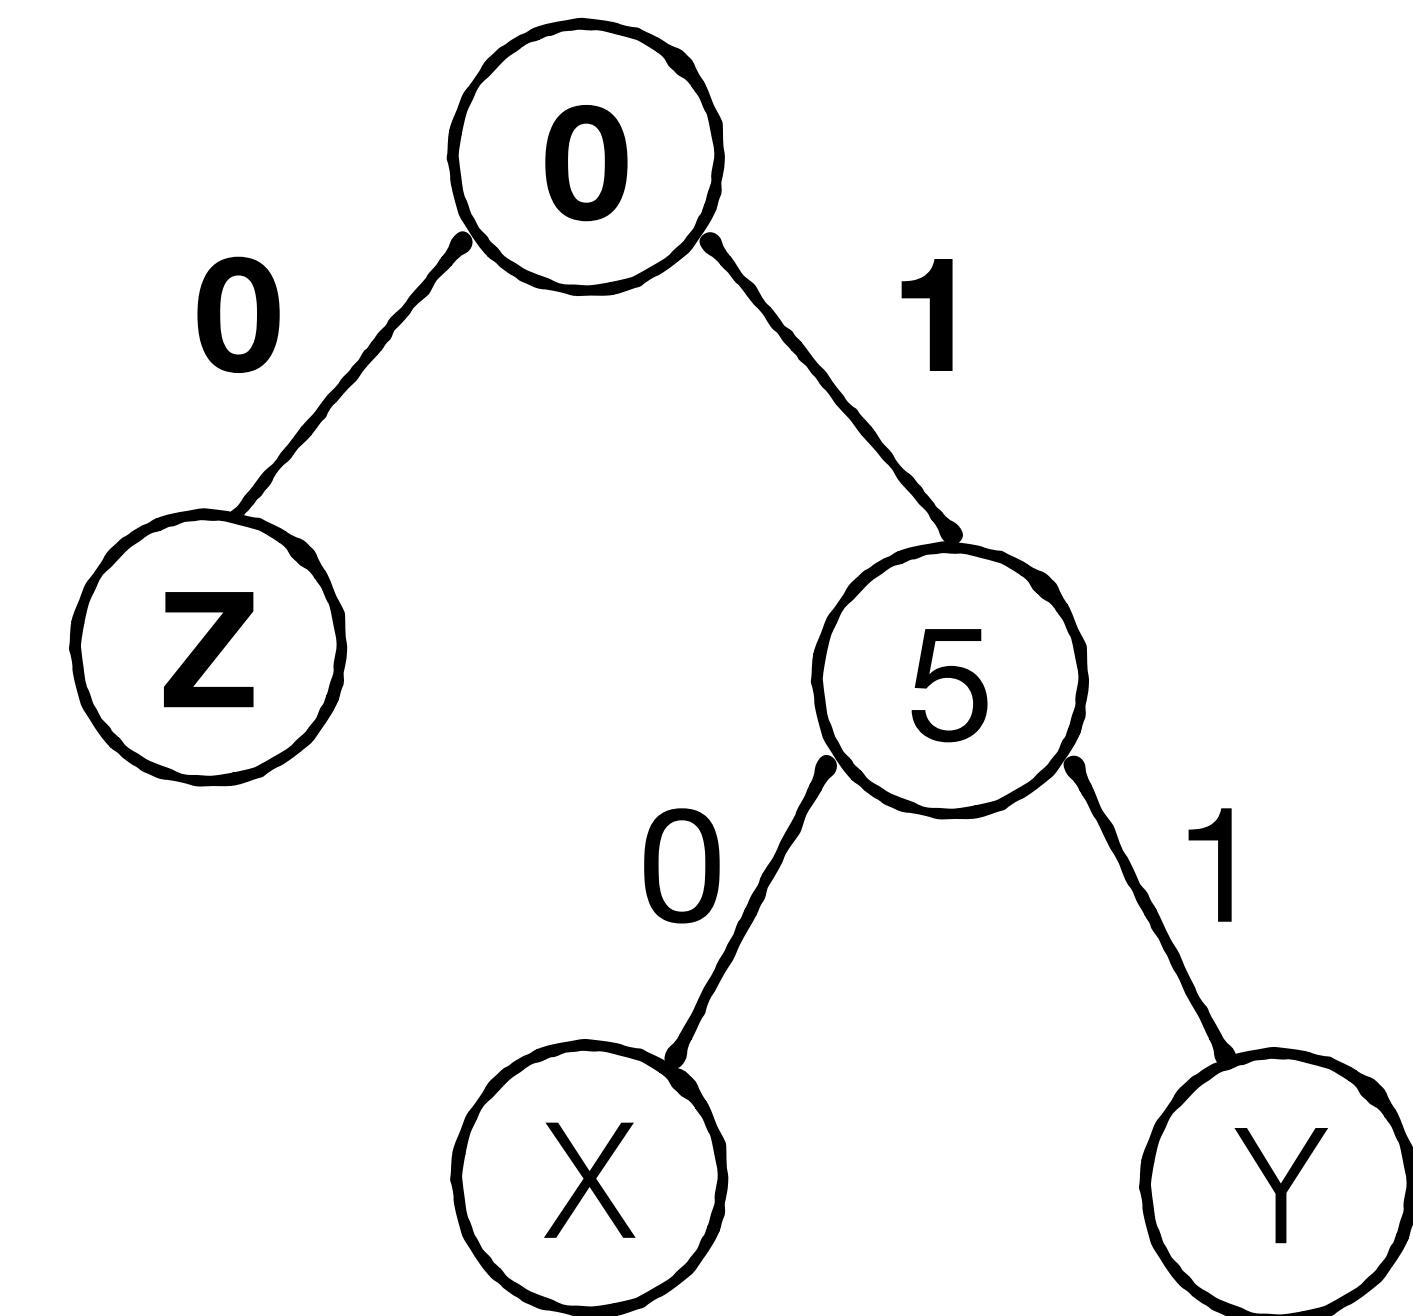


Build Trie capturing index of first different bit

hash(X) =  1 0 1 1 0 0 ...

hash(Y) =  1 0 1 1 0 1 ...

hash(Z) =  0 1 0 0 1 1 ...



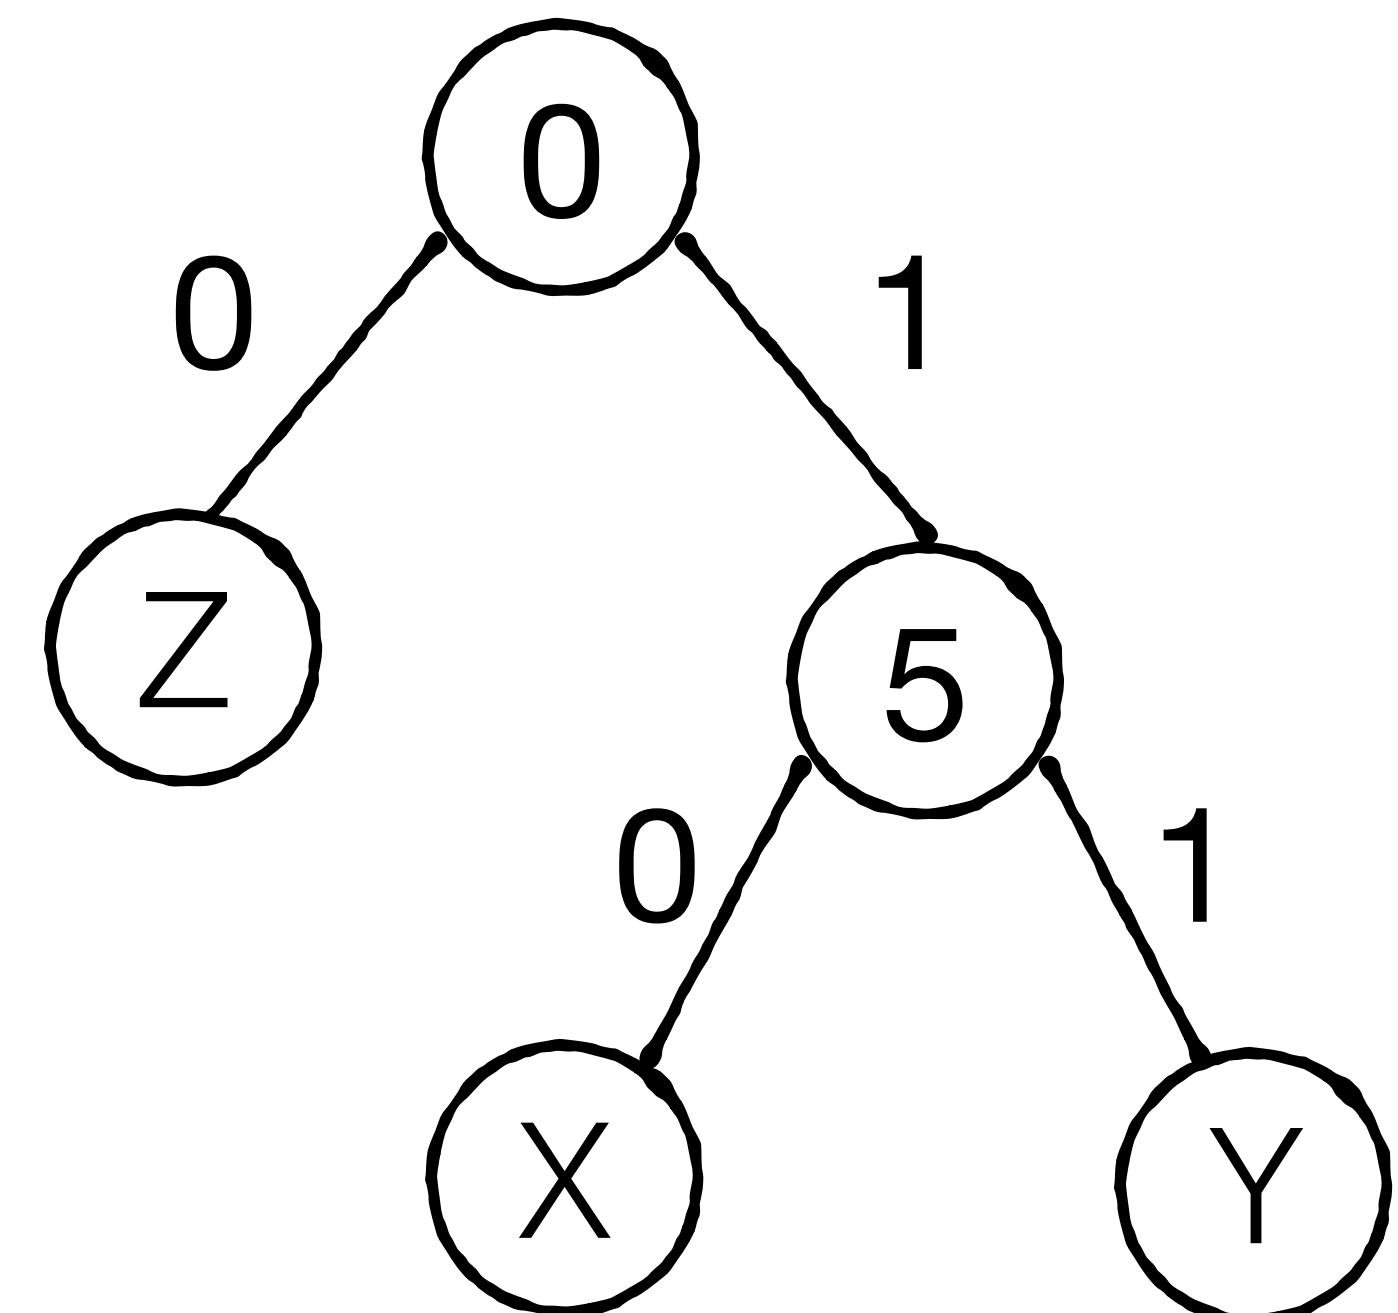
Build Trie capturing index of first different bit

hash(X) = 1 0 1 1 0 0 ...

hash(Y) = 1 0 1 1 0 1 ...

hash(Z) = 0 1 0 0 1 1 ...

hash(Q) = 0 1 1 0 0 1 ...



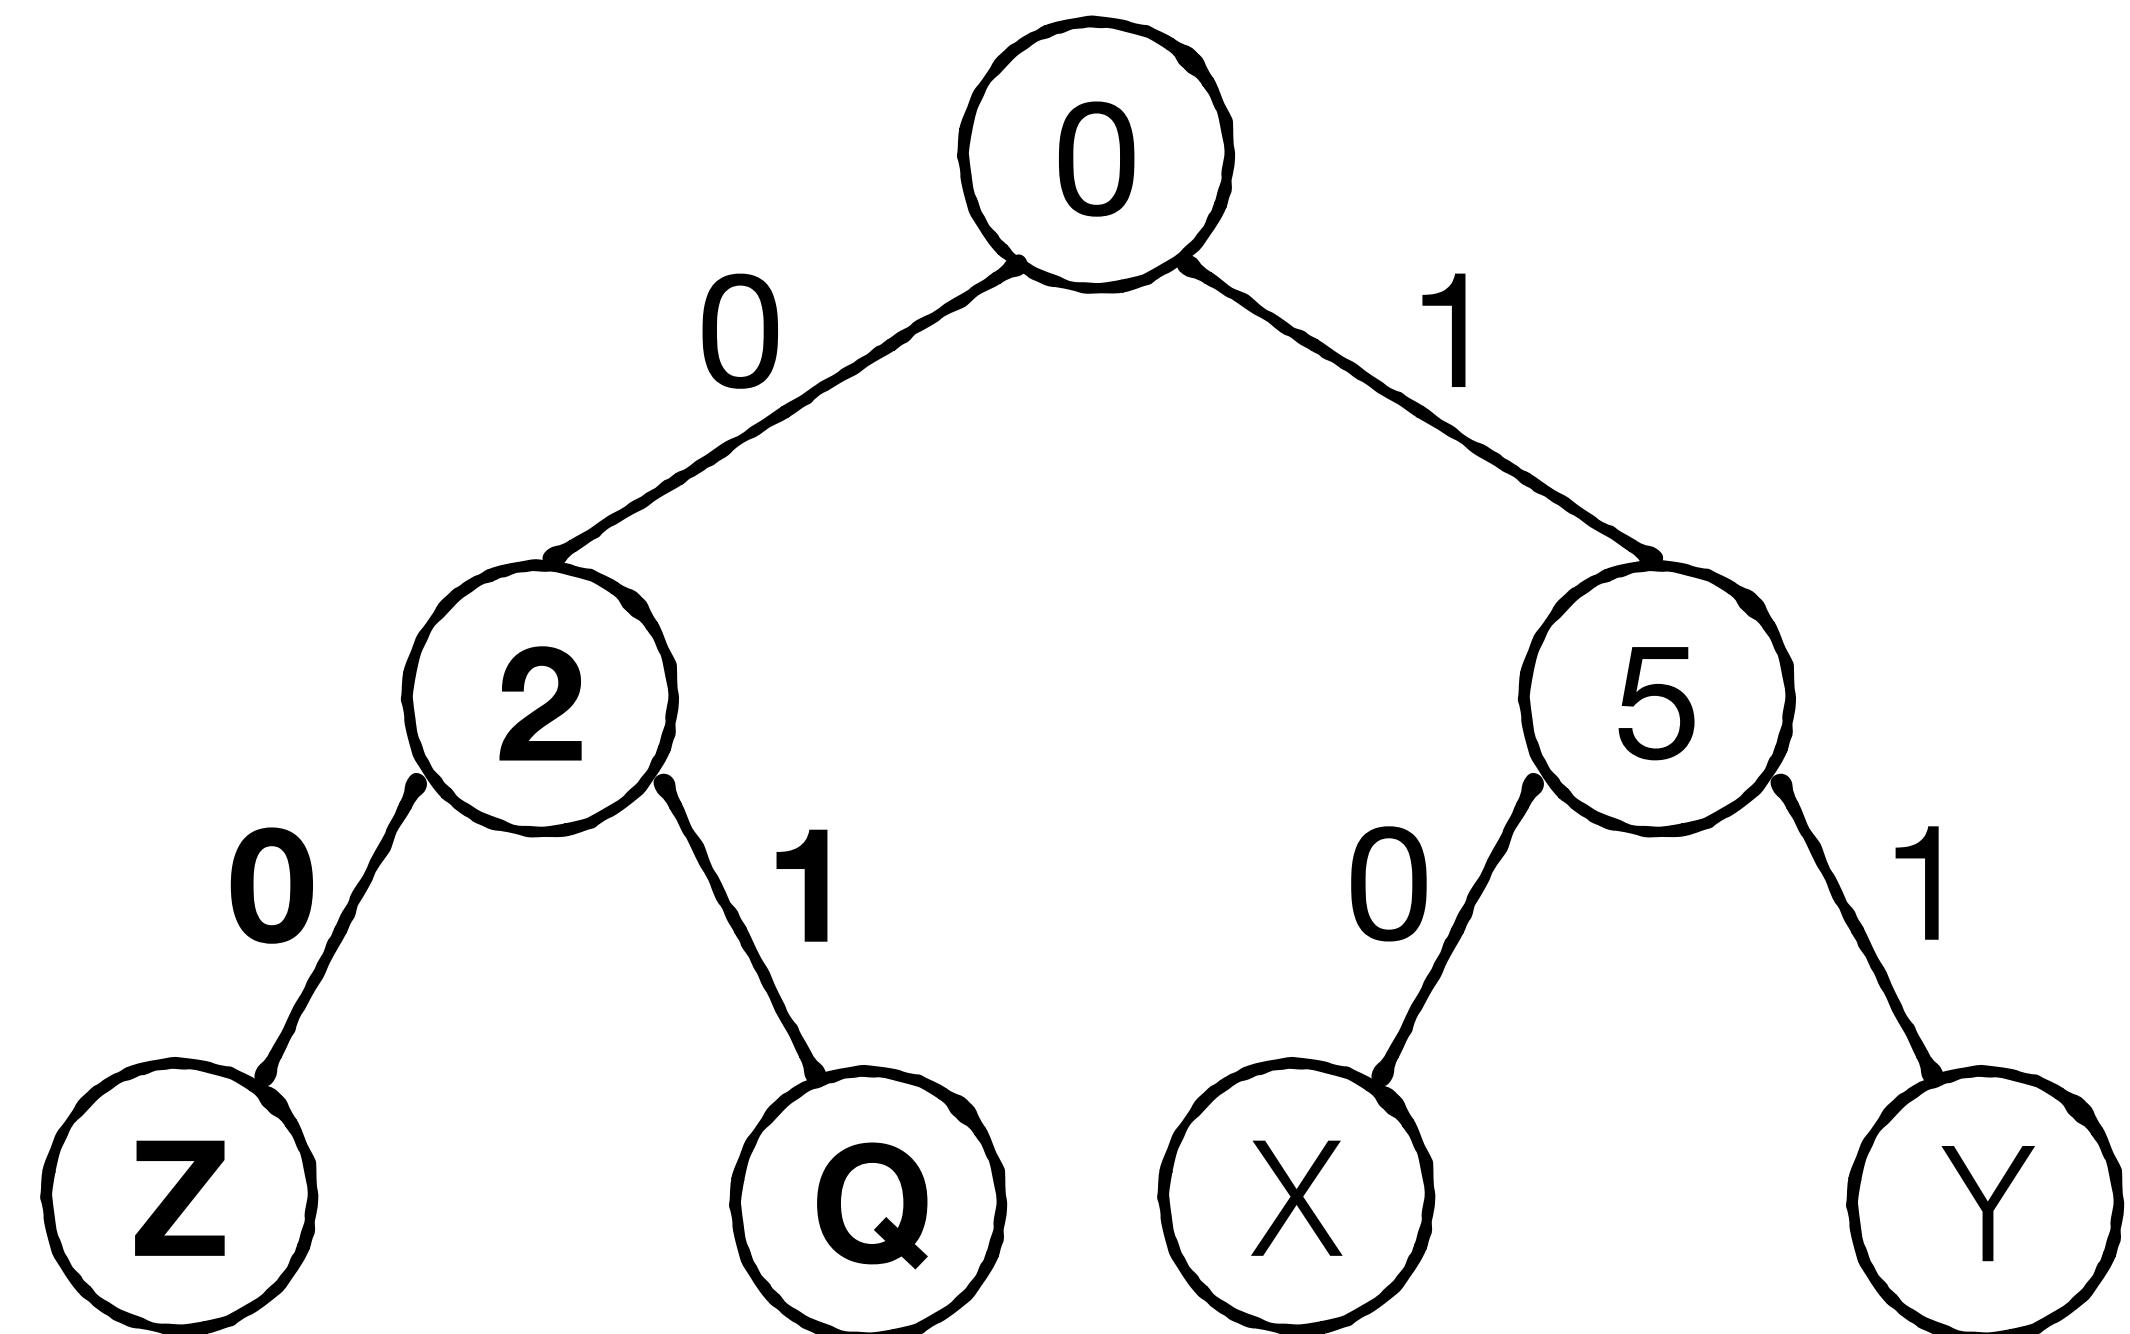
Build Trie capturing index of first different bit

hash(X) = 1 0 1 1 0 0 ...

hash(Y) = 1 0 1 1 0 1 ...

hash(Z) = 0 1 0 0 1 1 ...

hash(Q) = 0 1 1 0 0 1 ...



Build Trie capturing index of first different bit

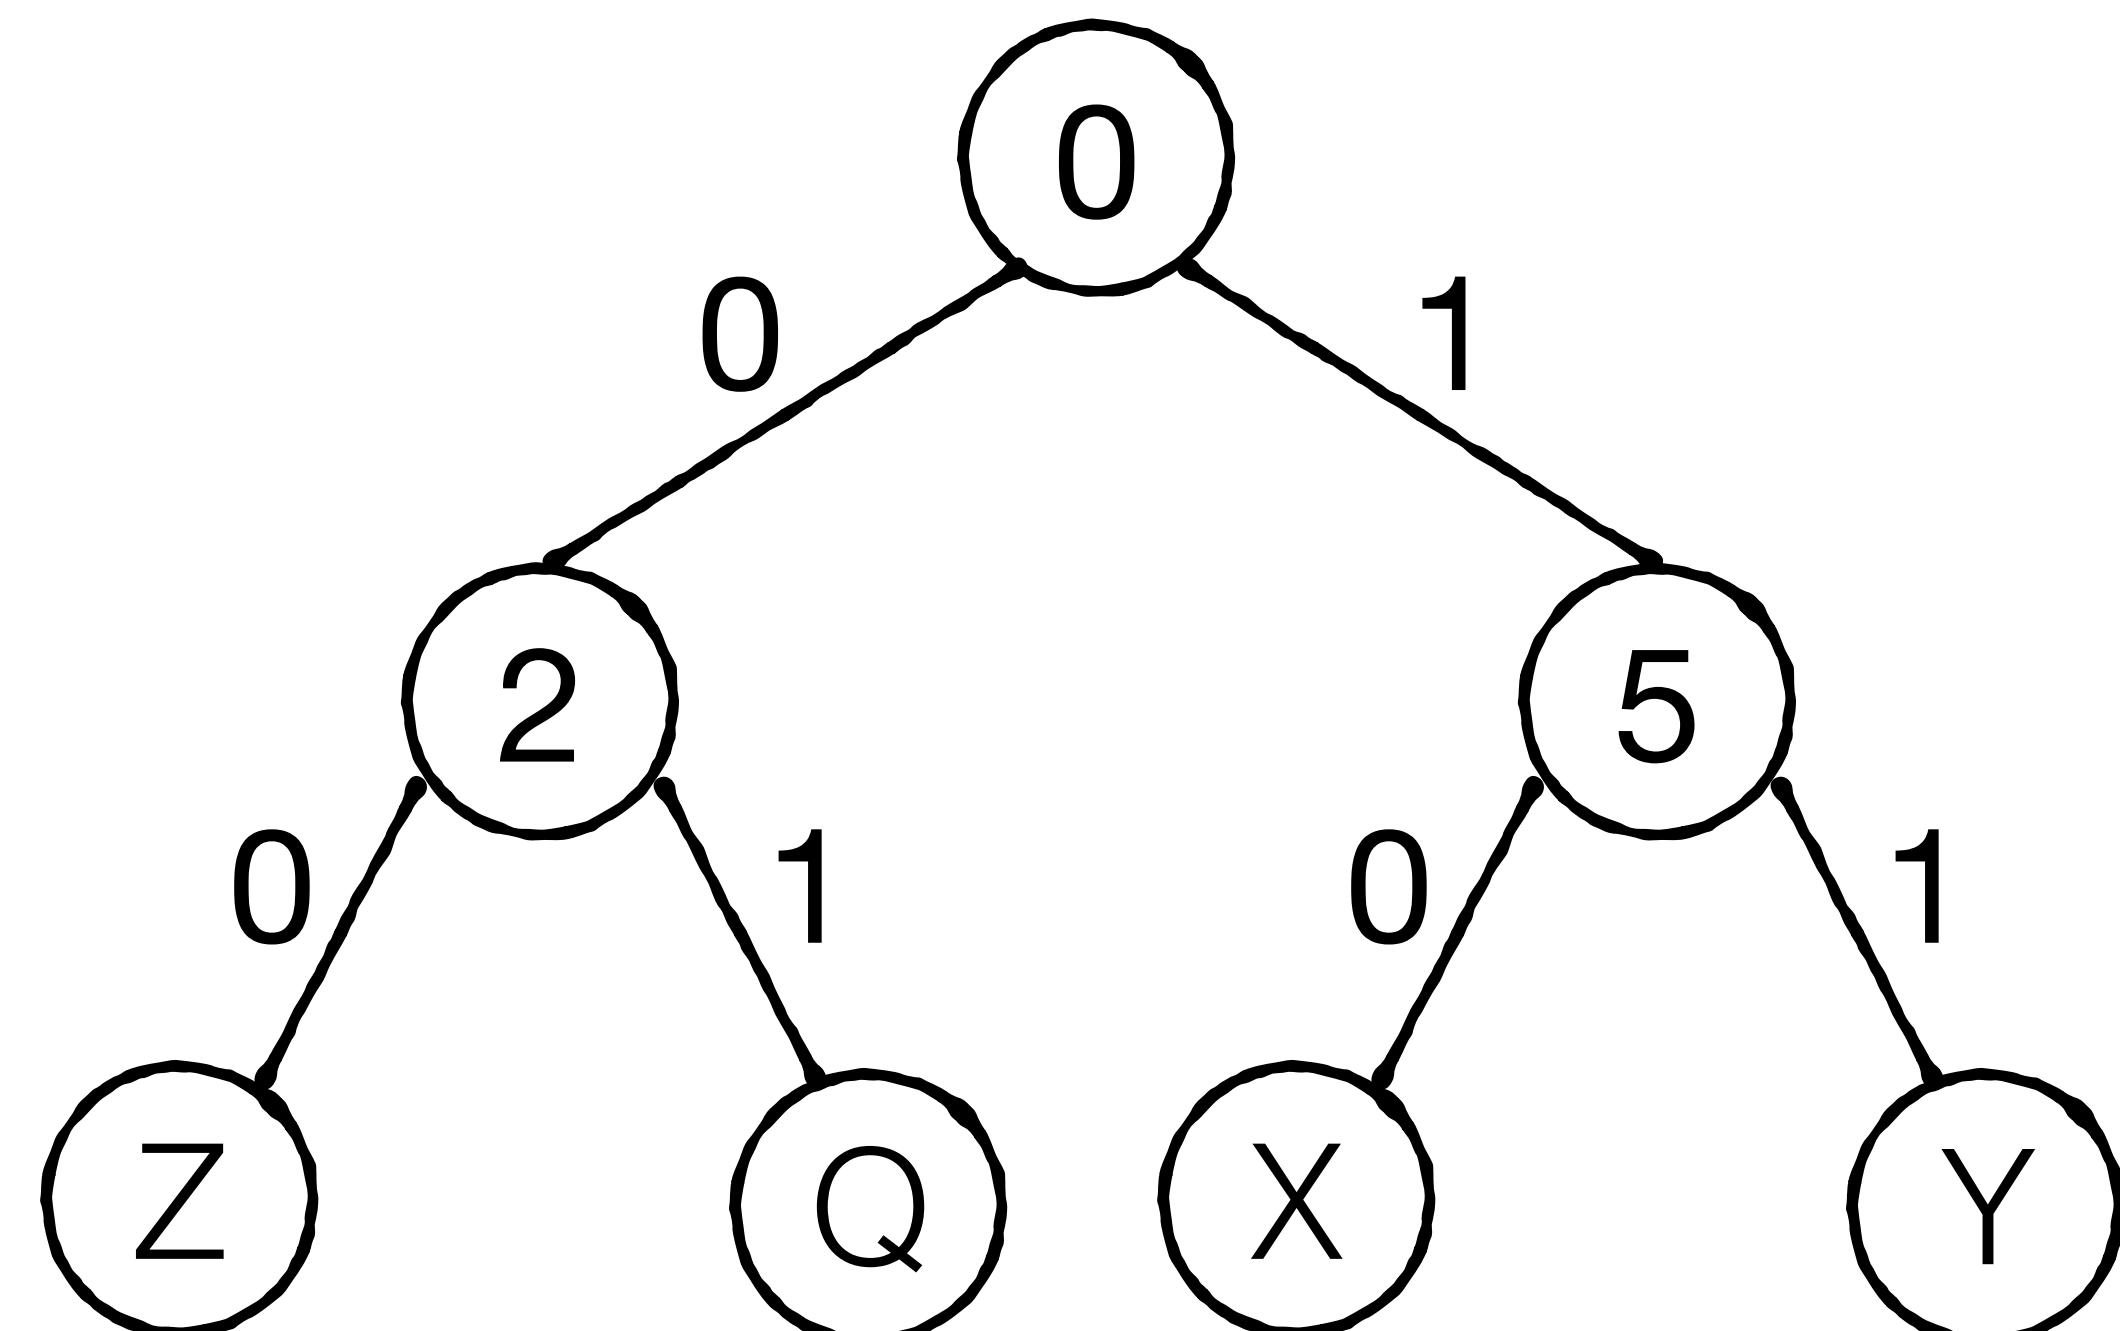
hash(X) = 1 0 1 1 0 0 ...

hash(Y) = 1 0 1 1 0 1 ...

hash(Z) = 0 1 0 0 1 1 ...

hash(Q) = 0 1 1 0 0 1 ...

hash(W) = 0 0 1 0 0 1 ...



Build Trie capturing index of first different bit

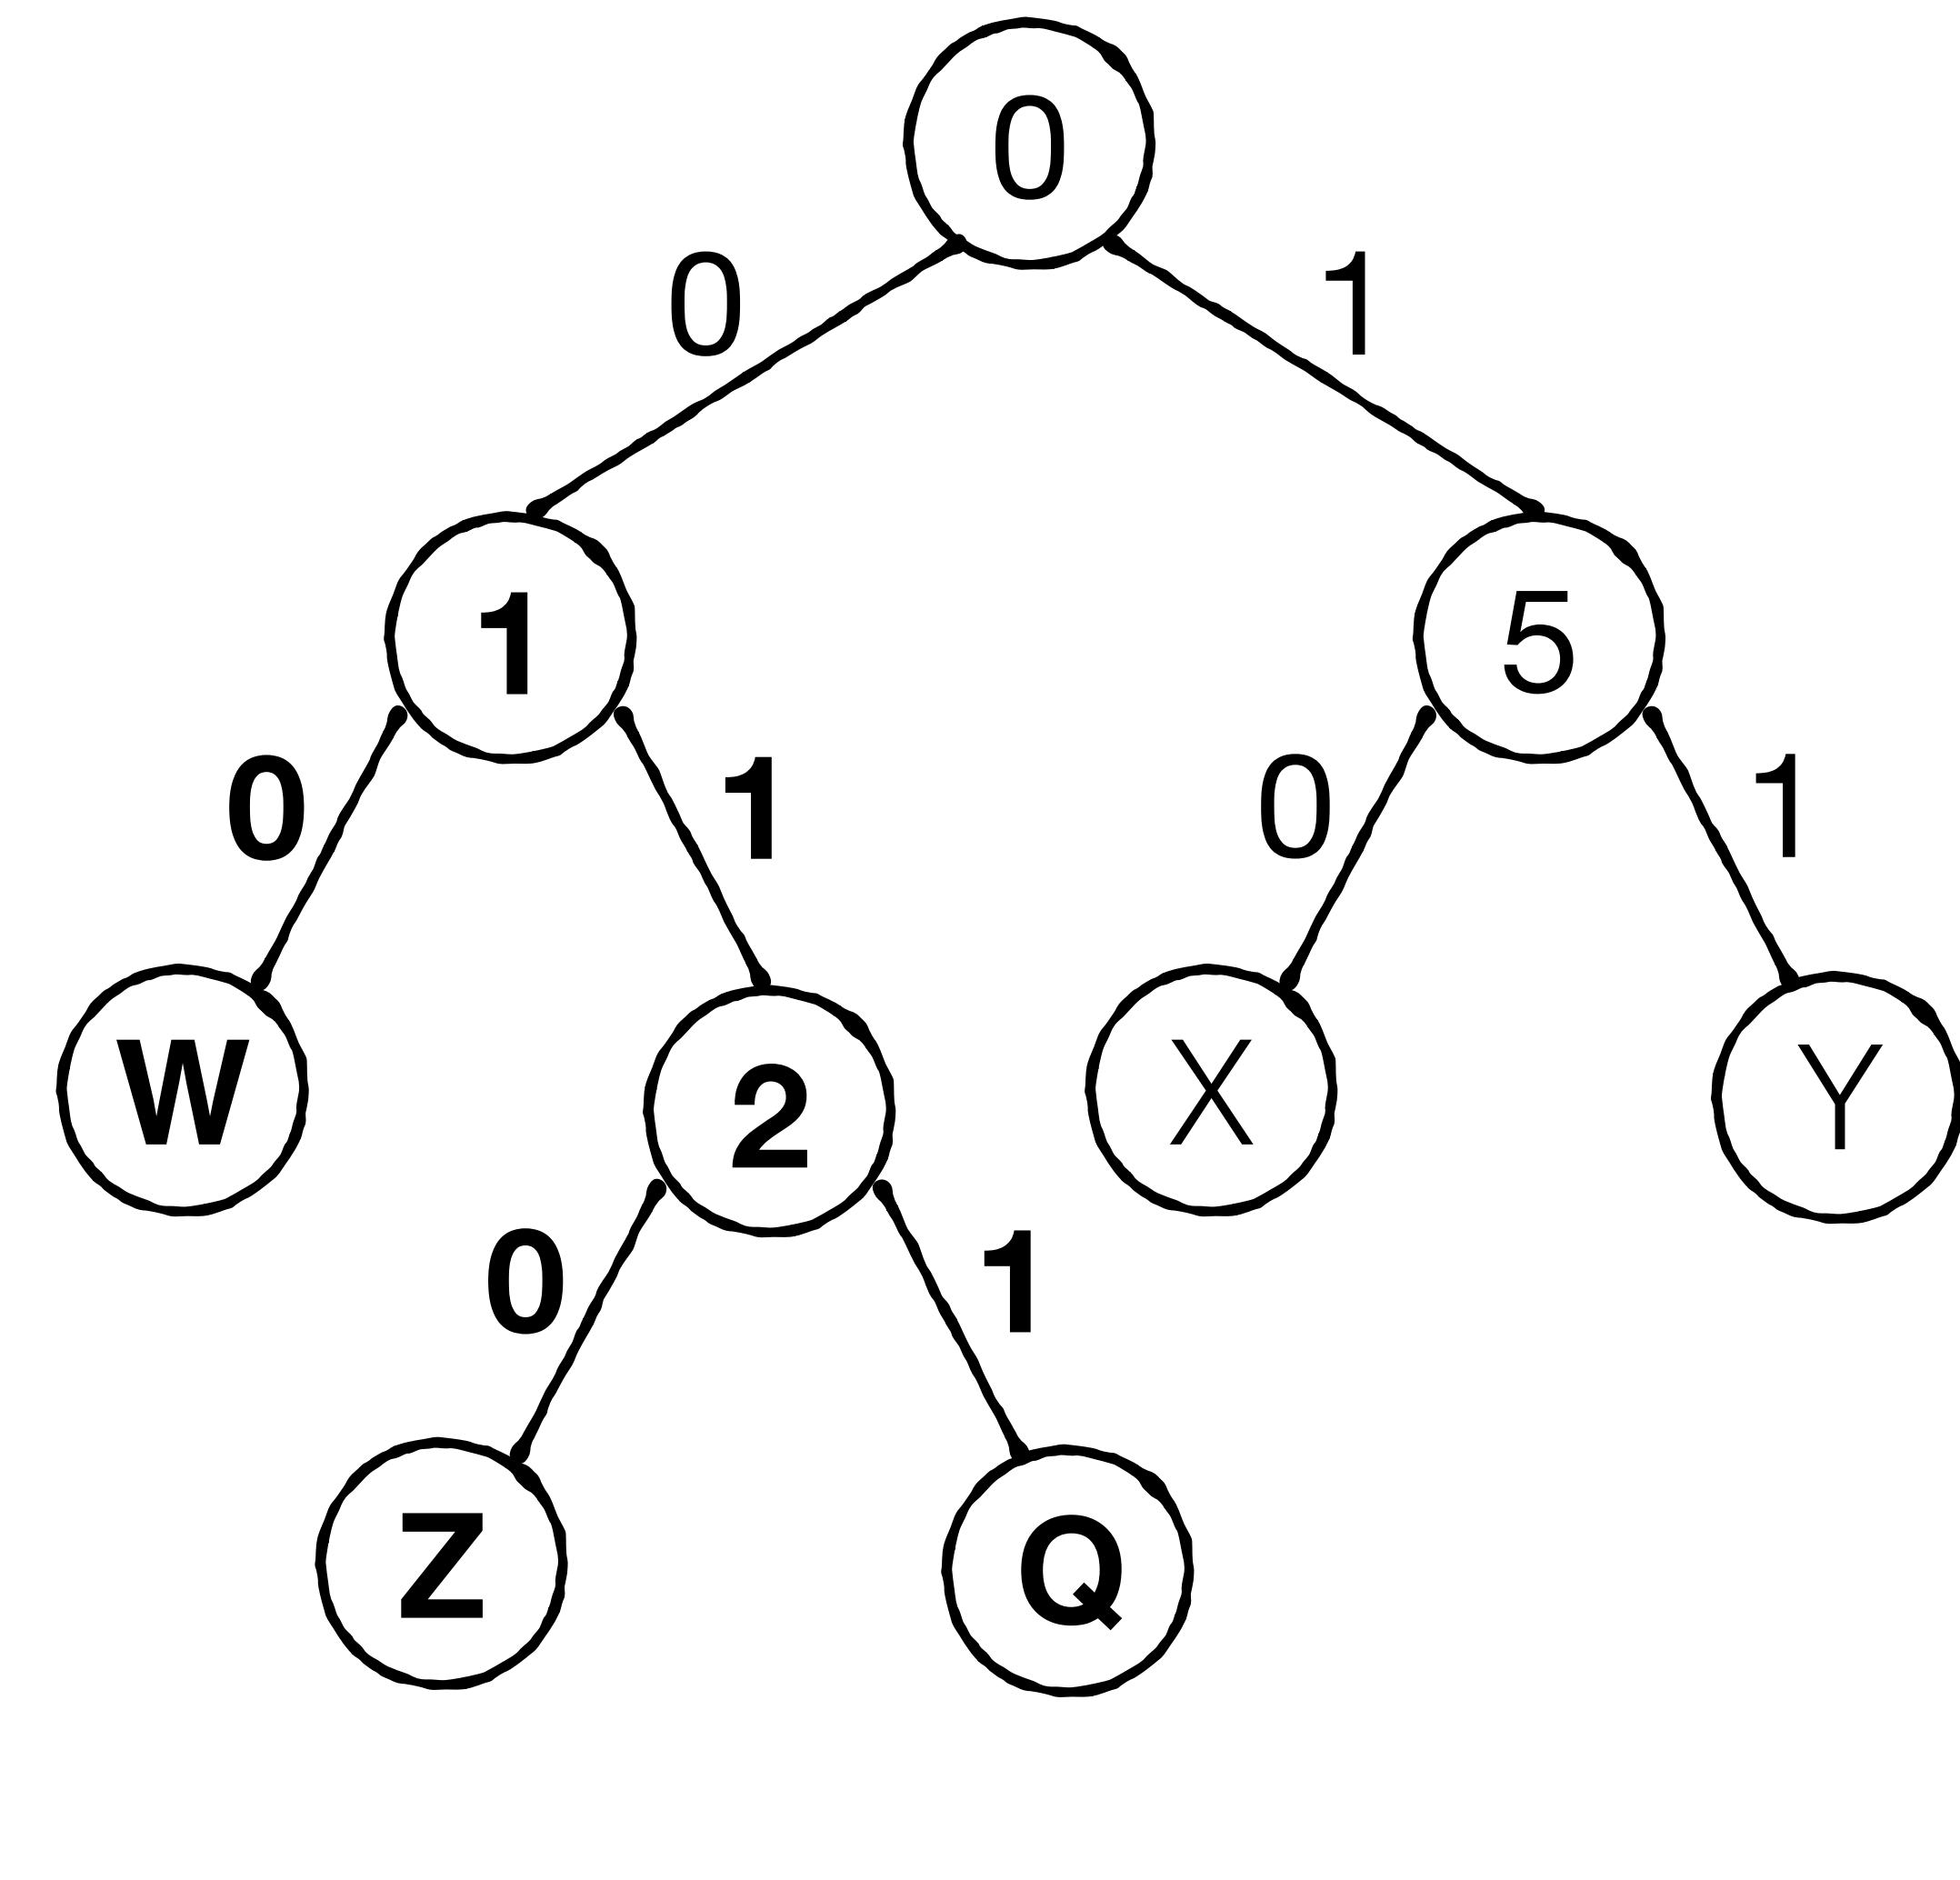
hash(X) = 1 0 1 1 0 0 ...

hash(Y) = 1 0 1 1 0 1 ...

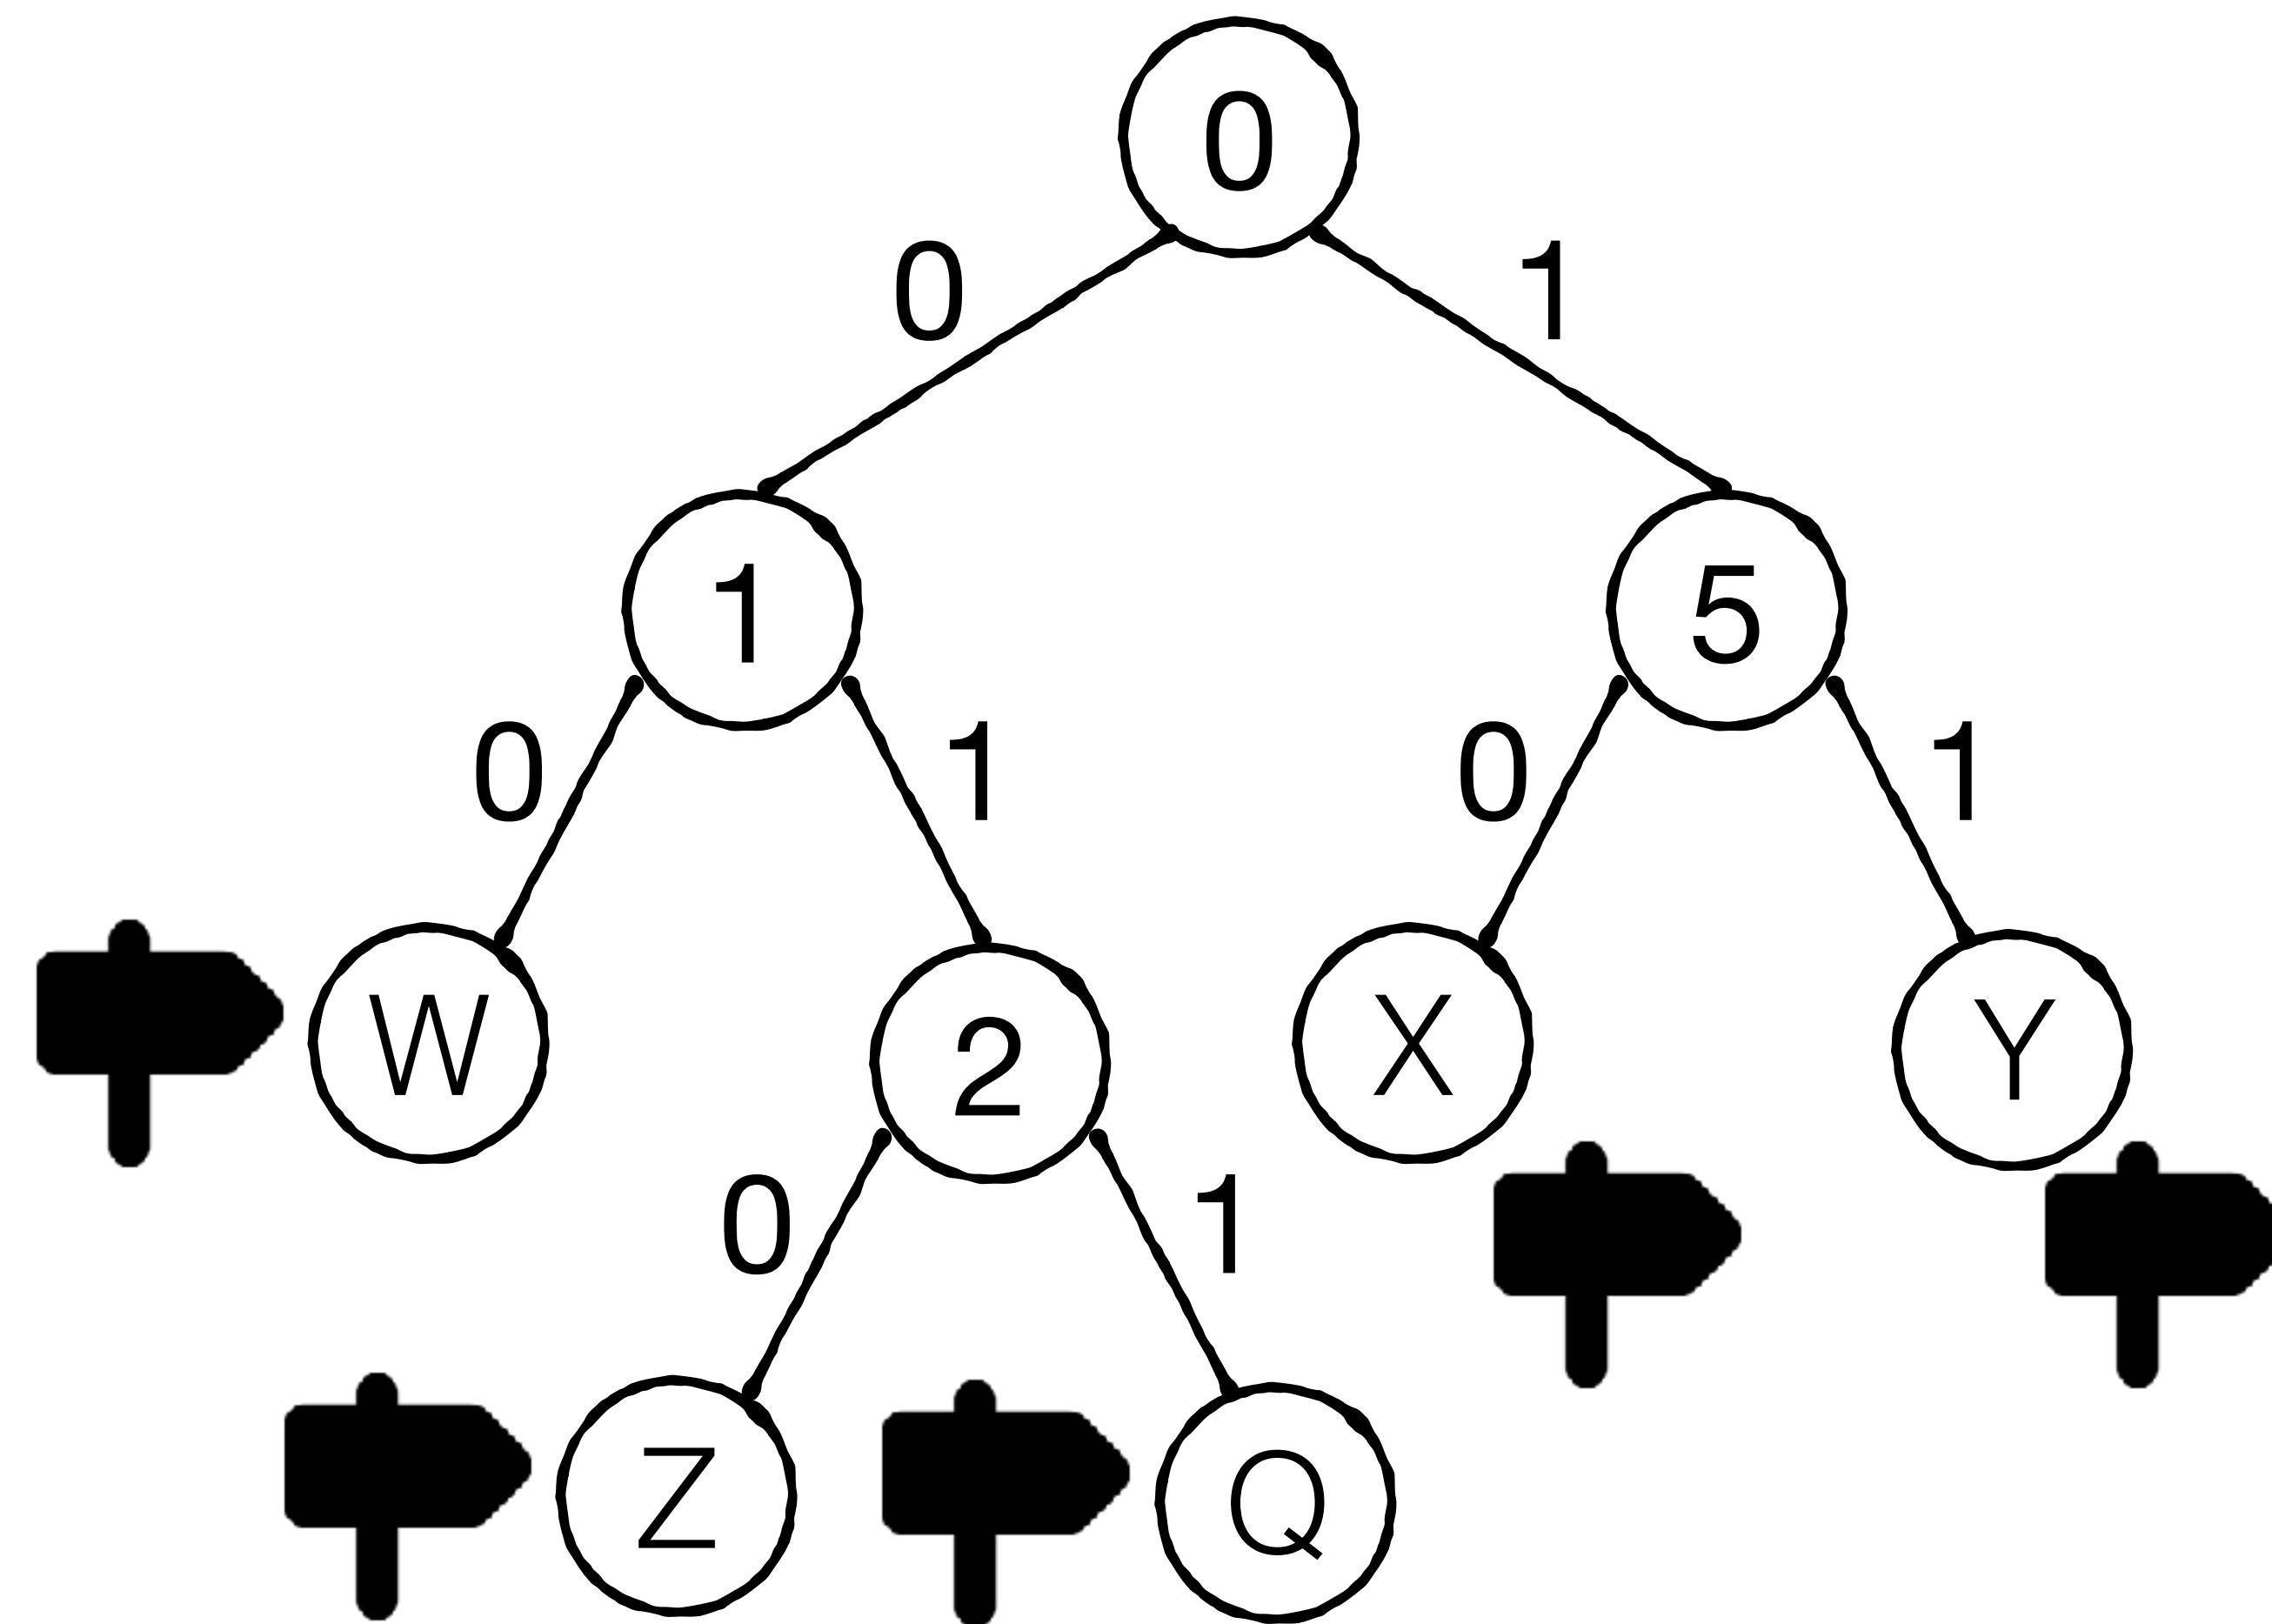
hash(Z) = 0 1 0 0 1 1 ...

hash(Q) = 0 1 1 0 0 1 ...

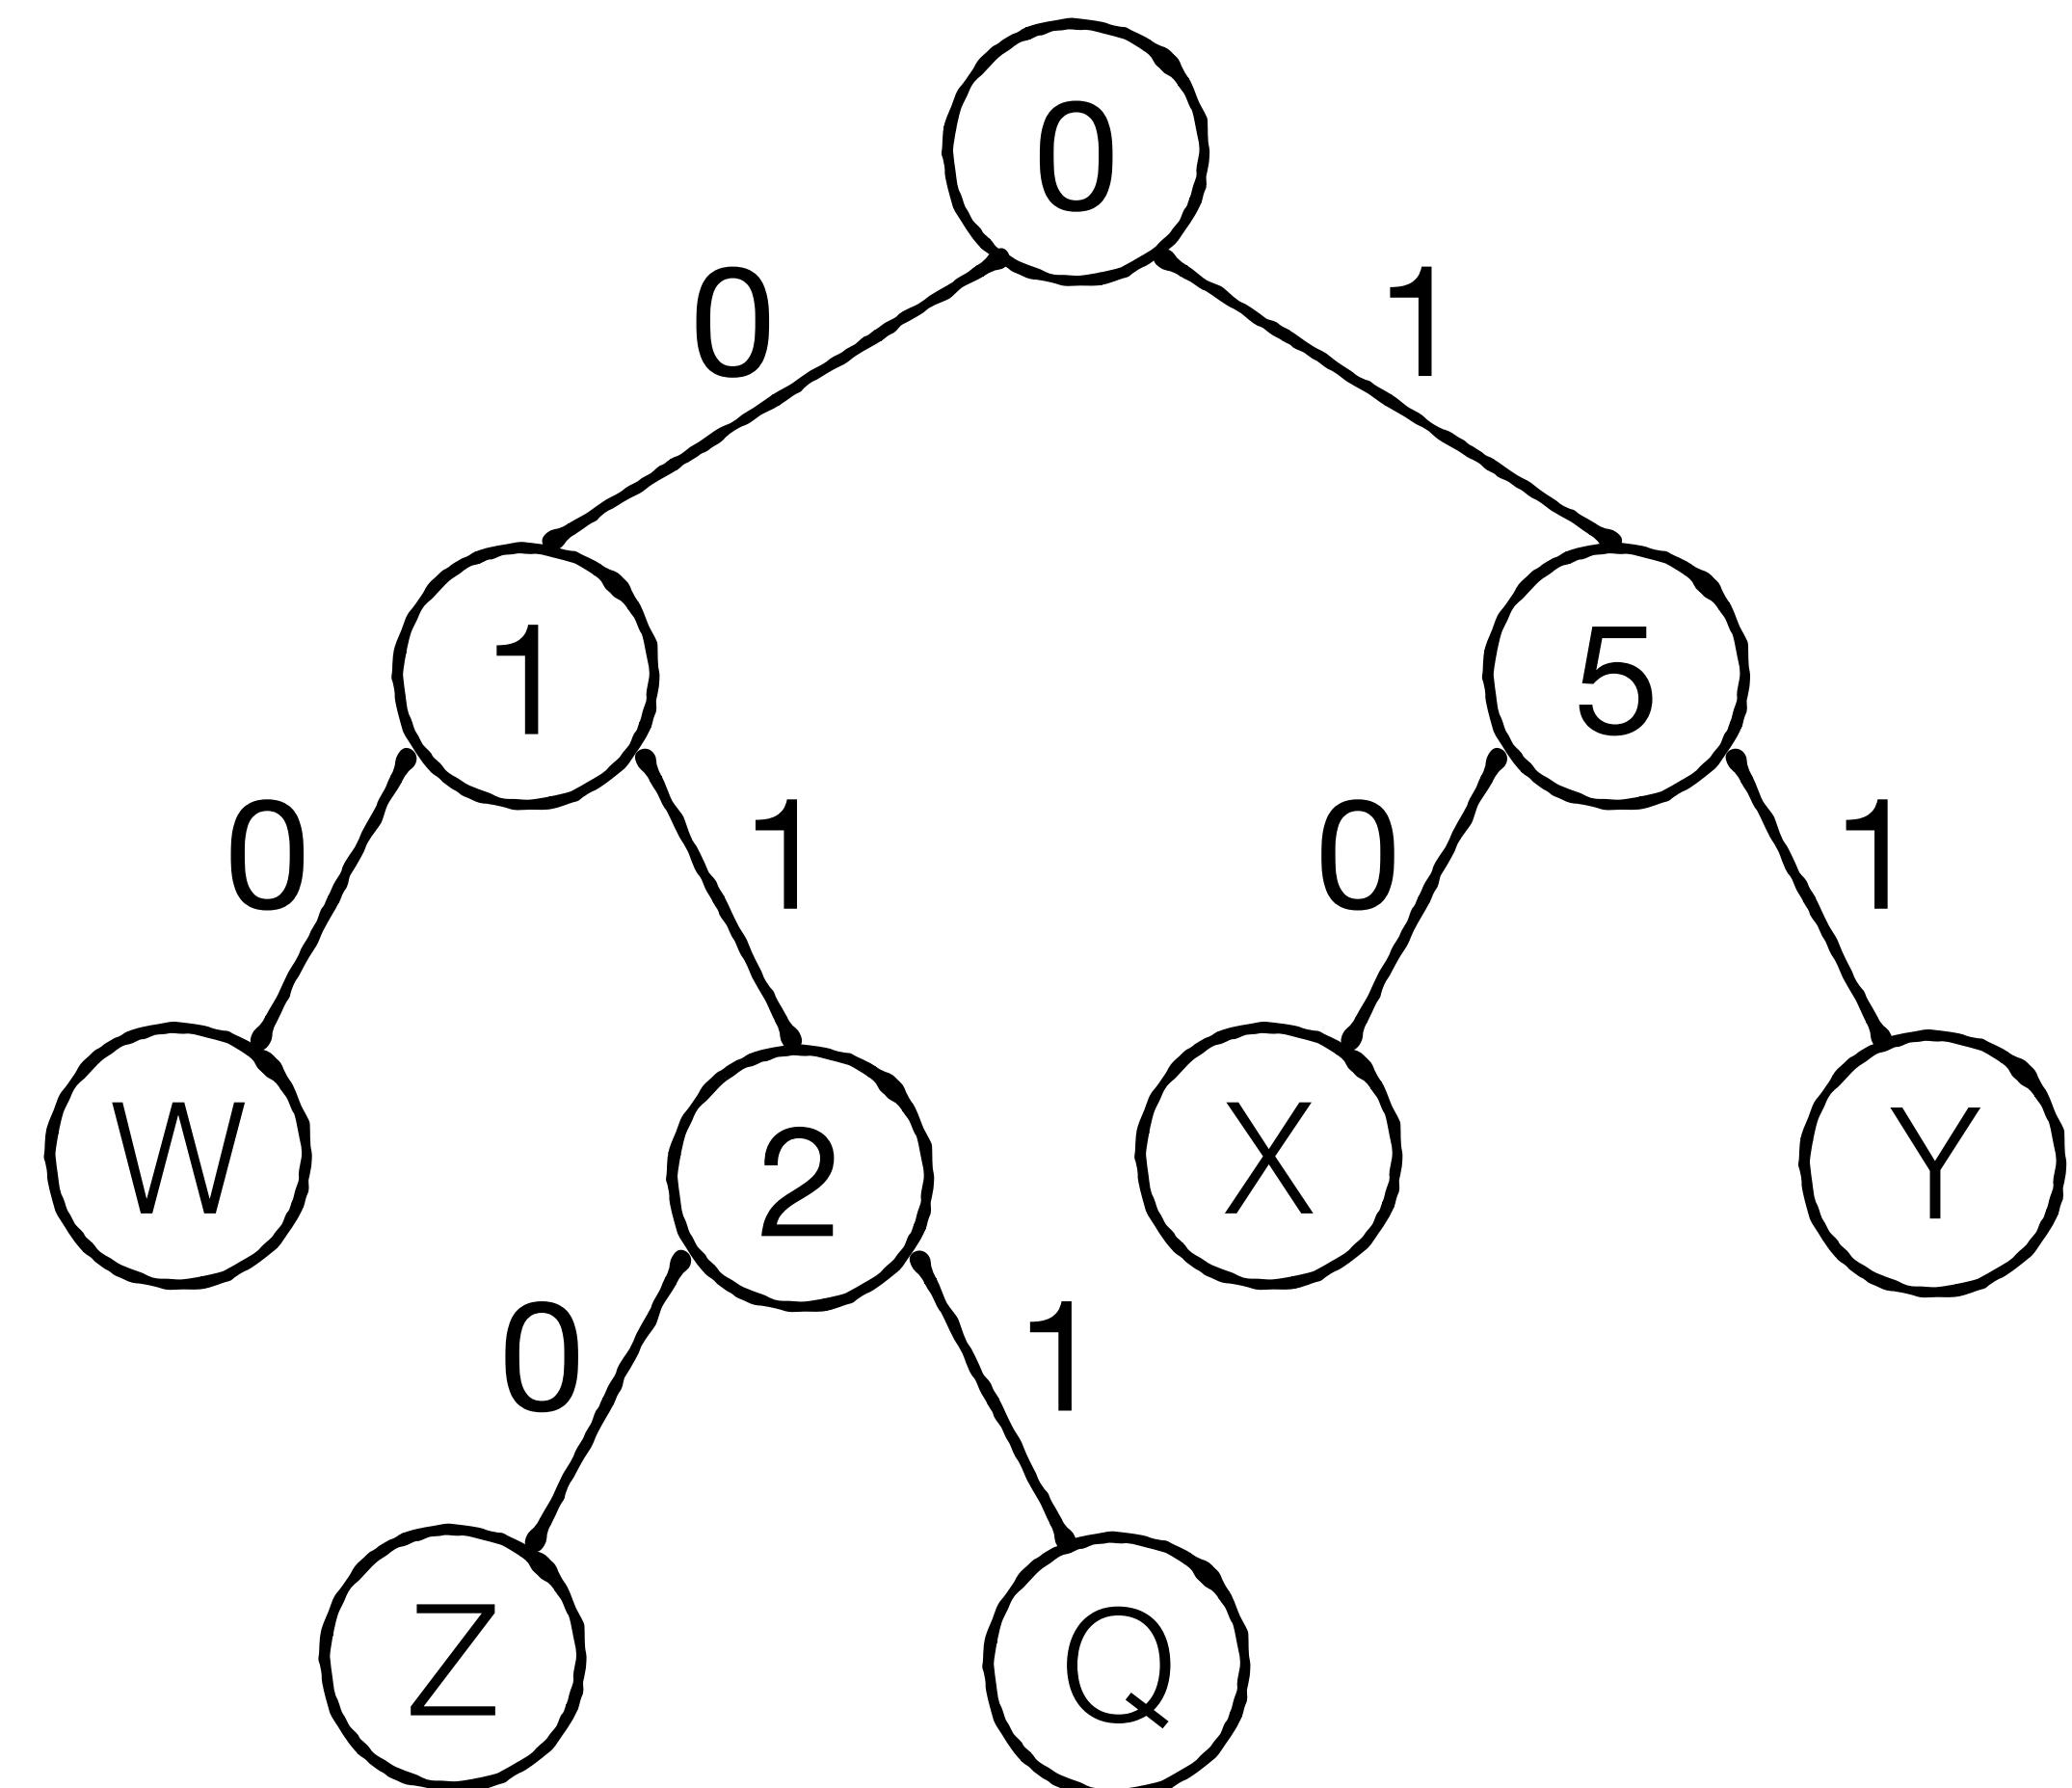
hash(W) = 0 0 1 0 0 1 ...



Place pointer/payload in leafs

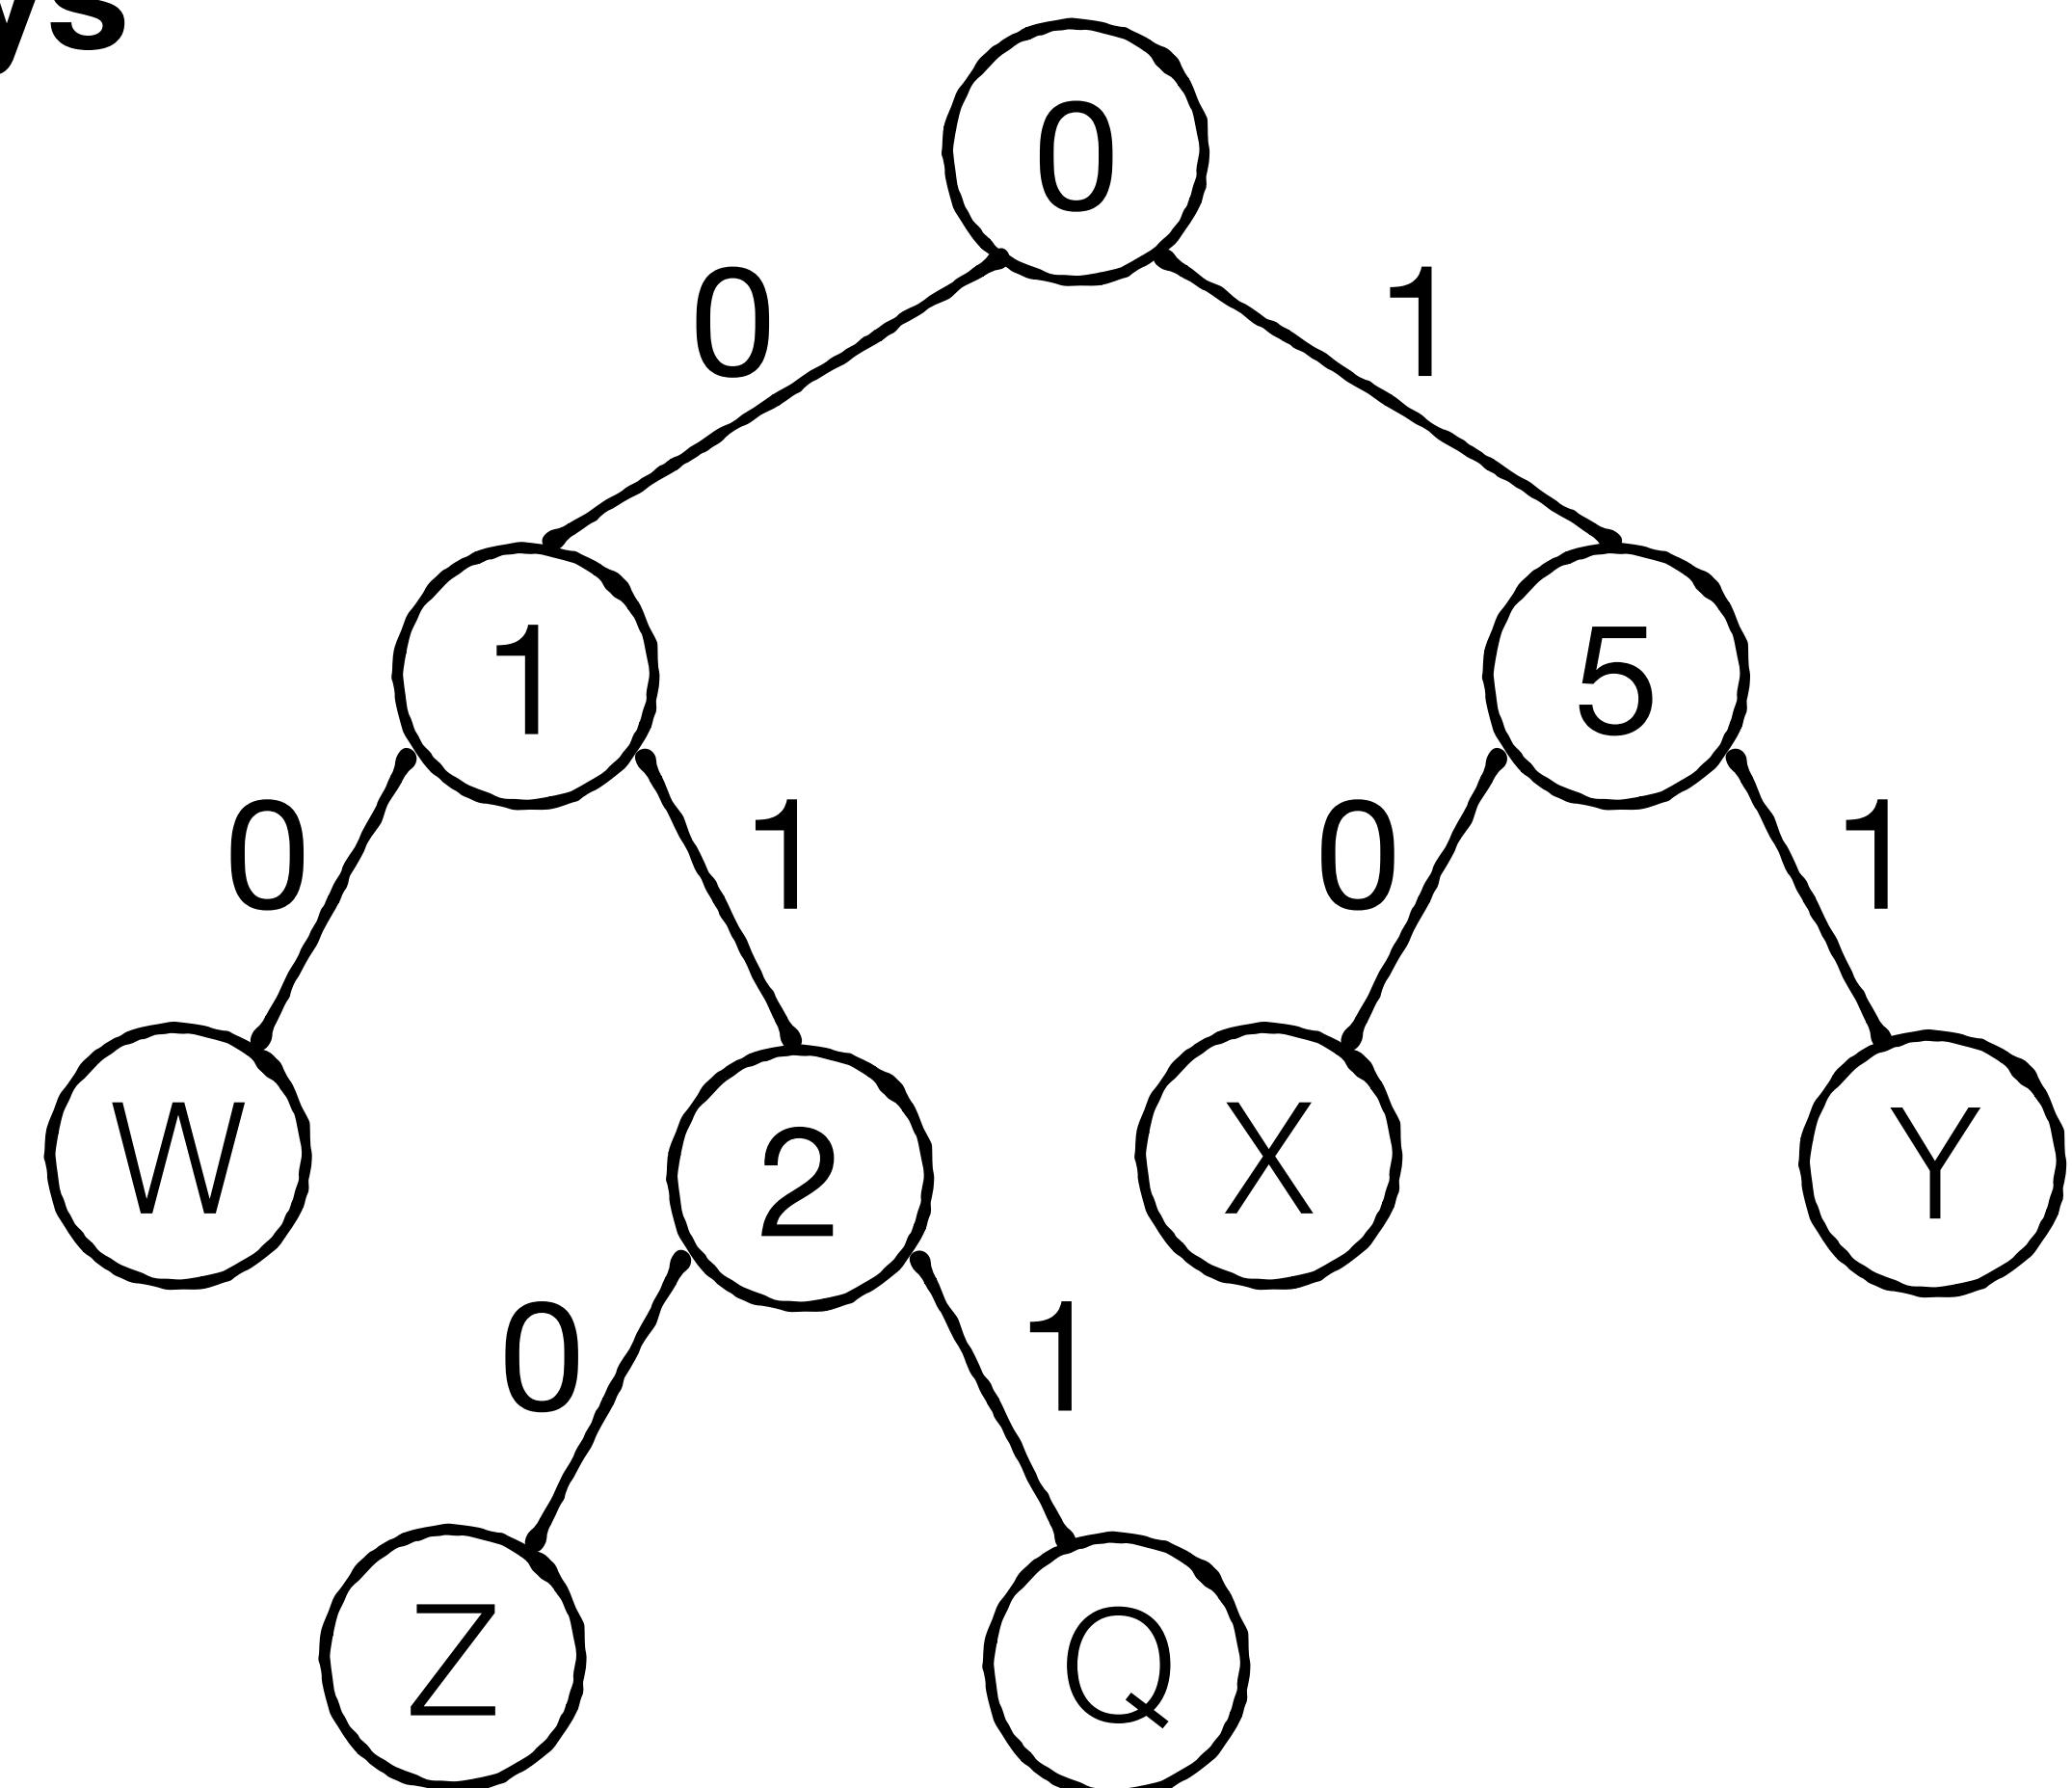


Existing keys are fully differentiated

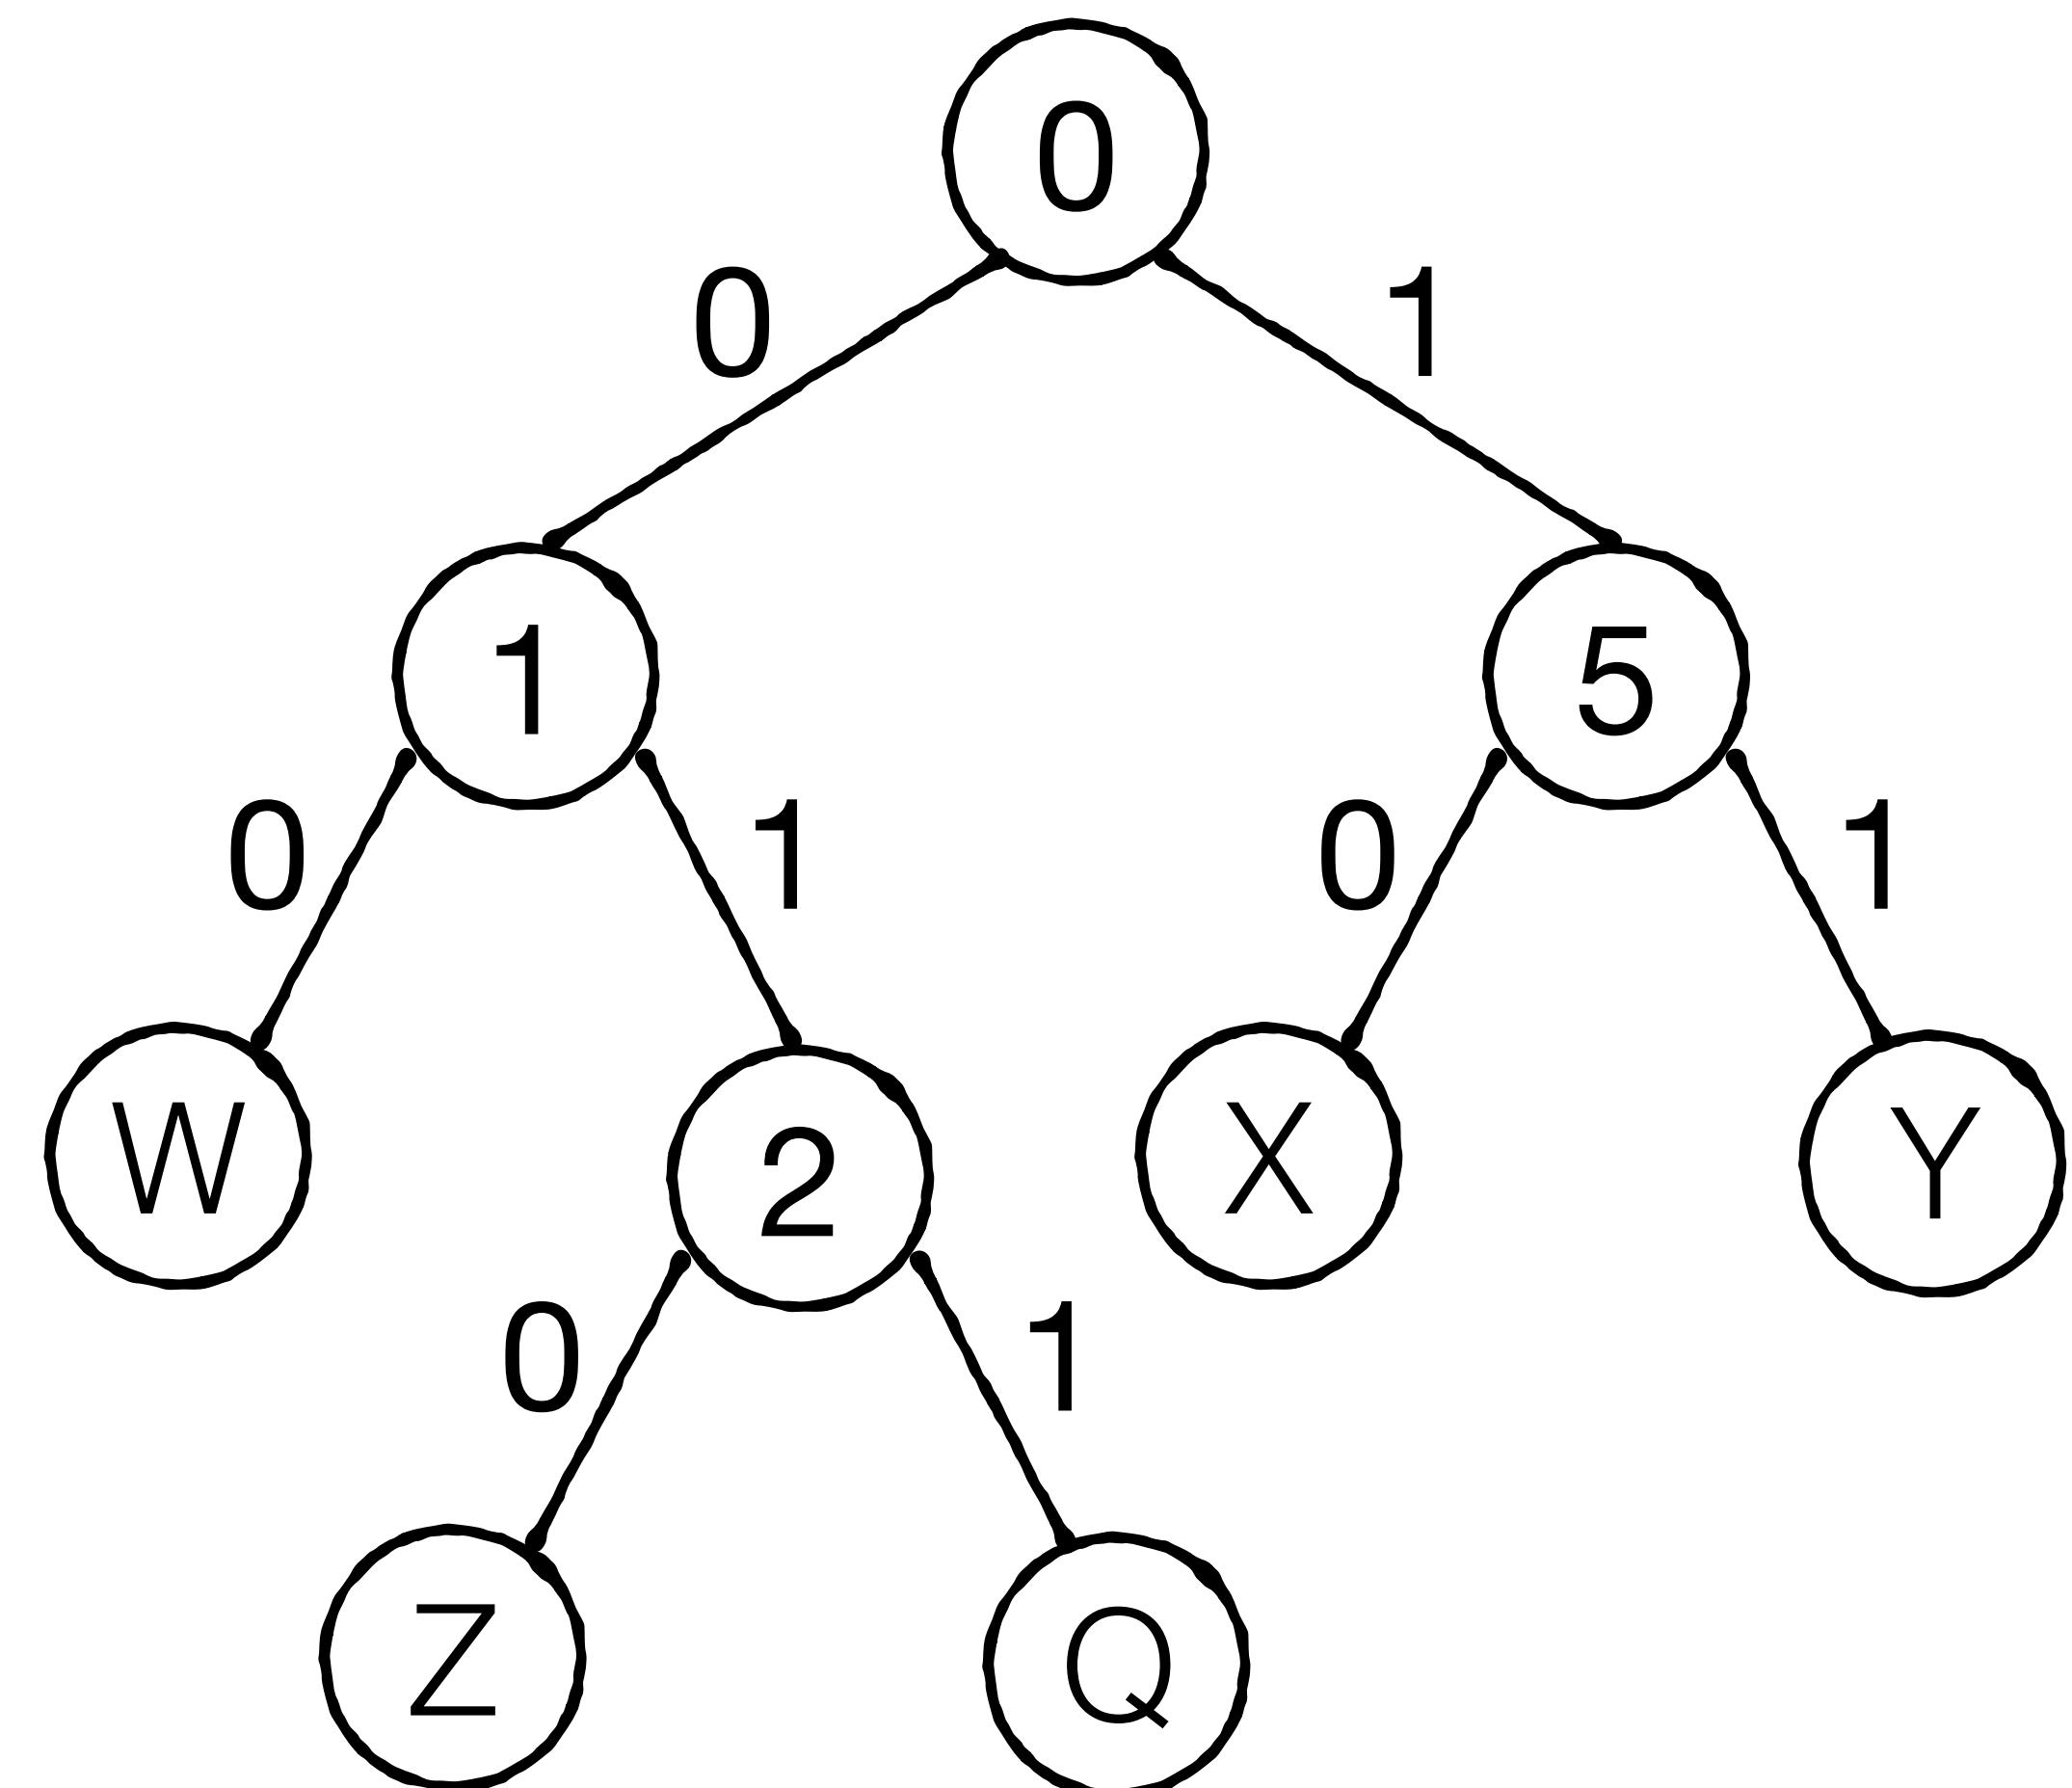


Existing keys are fully differentiated

**A query to an existing key always
finds correct payload**

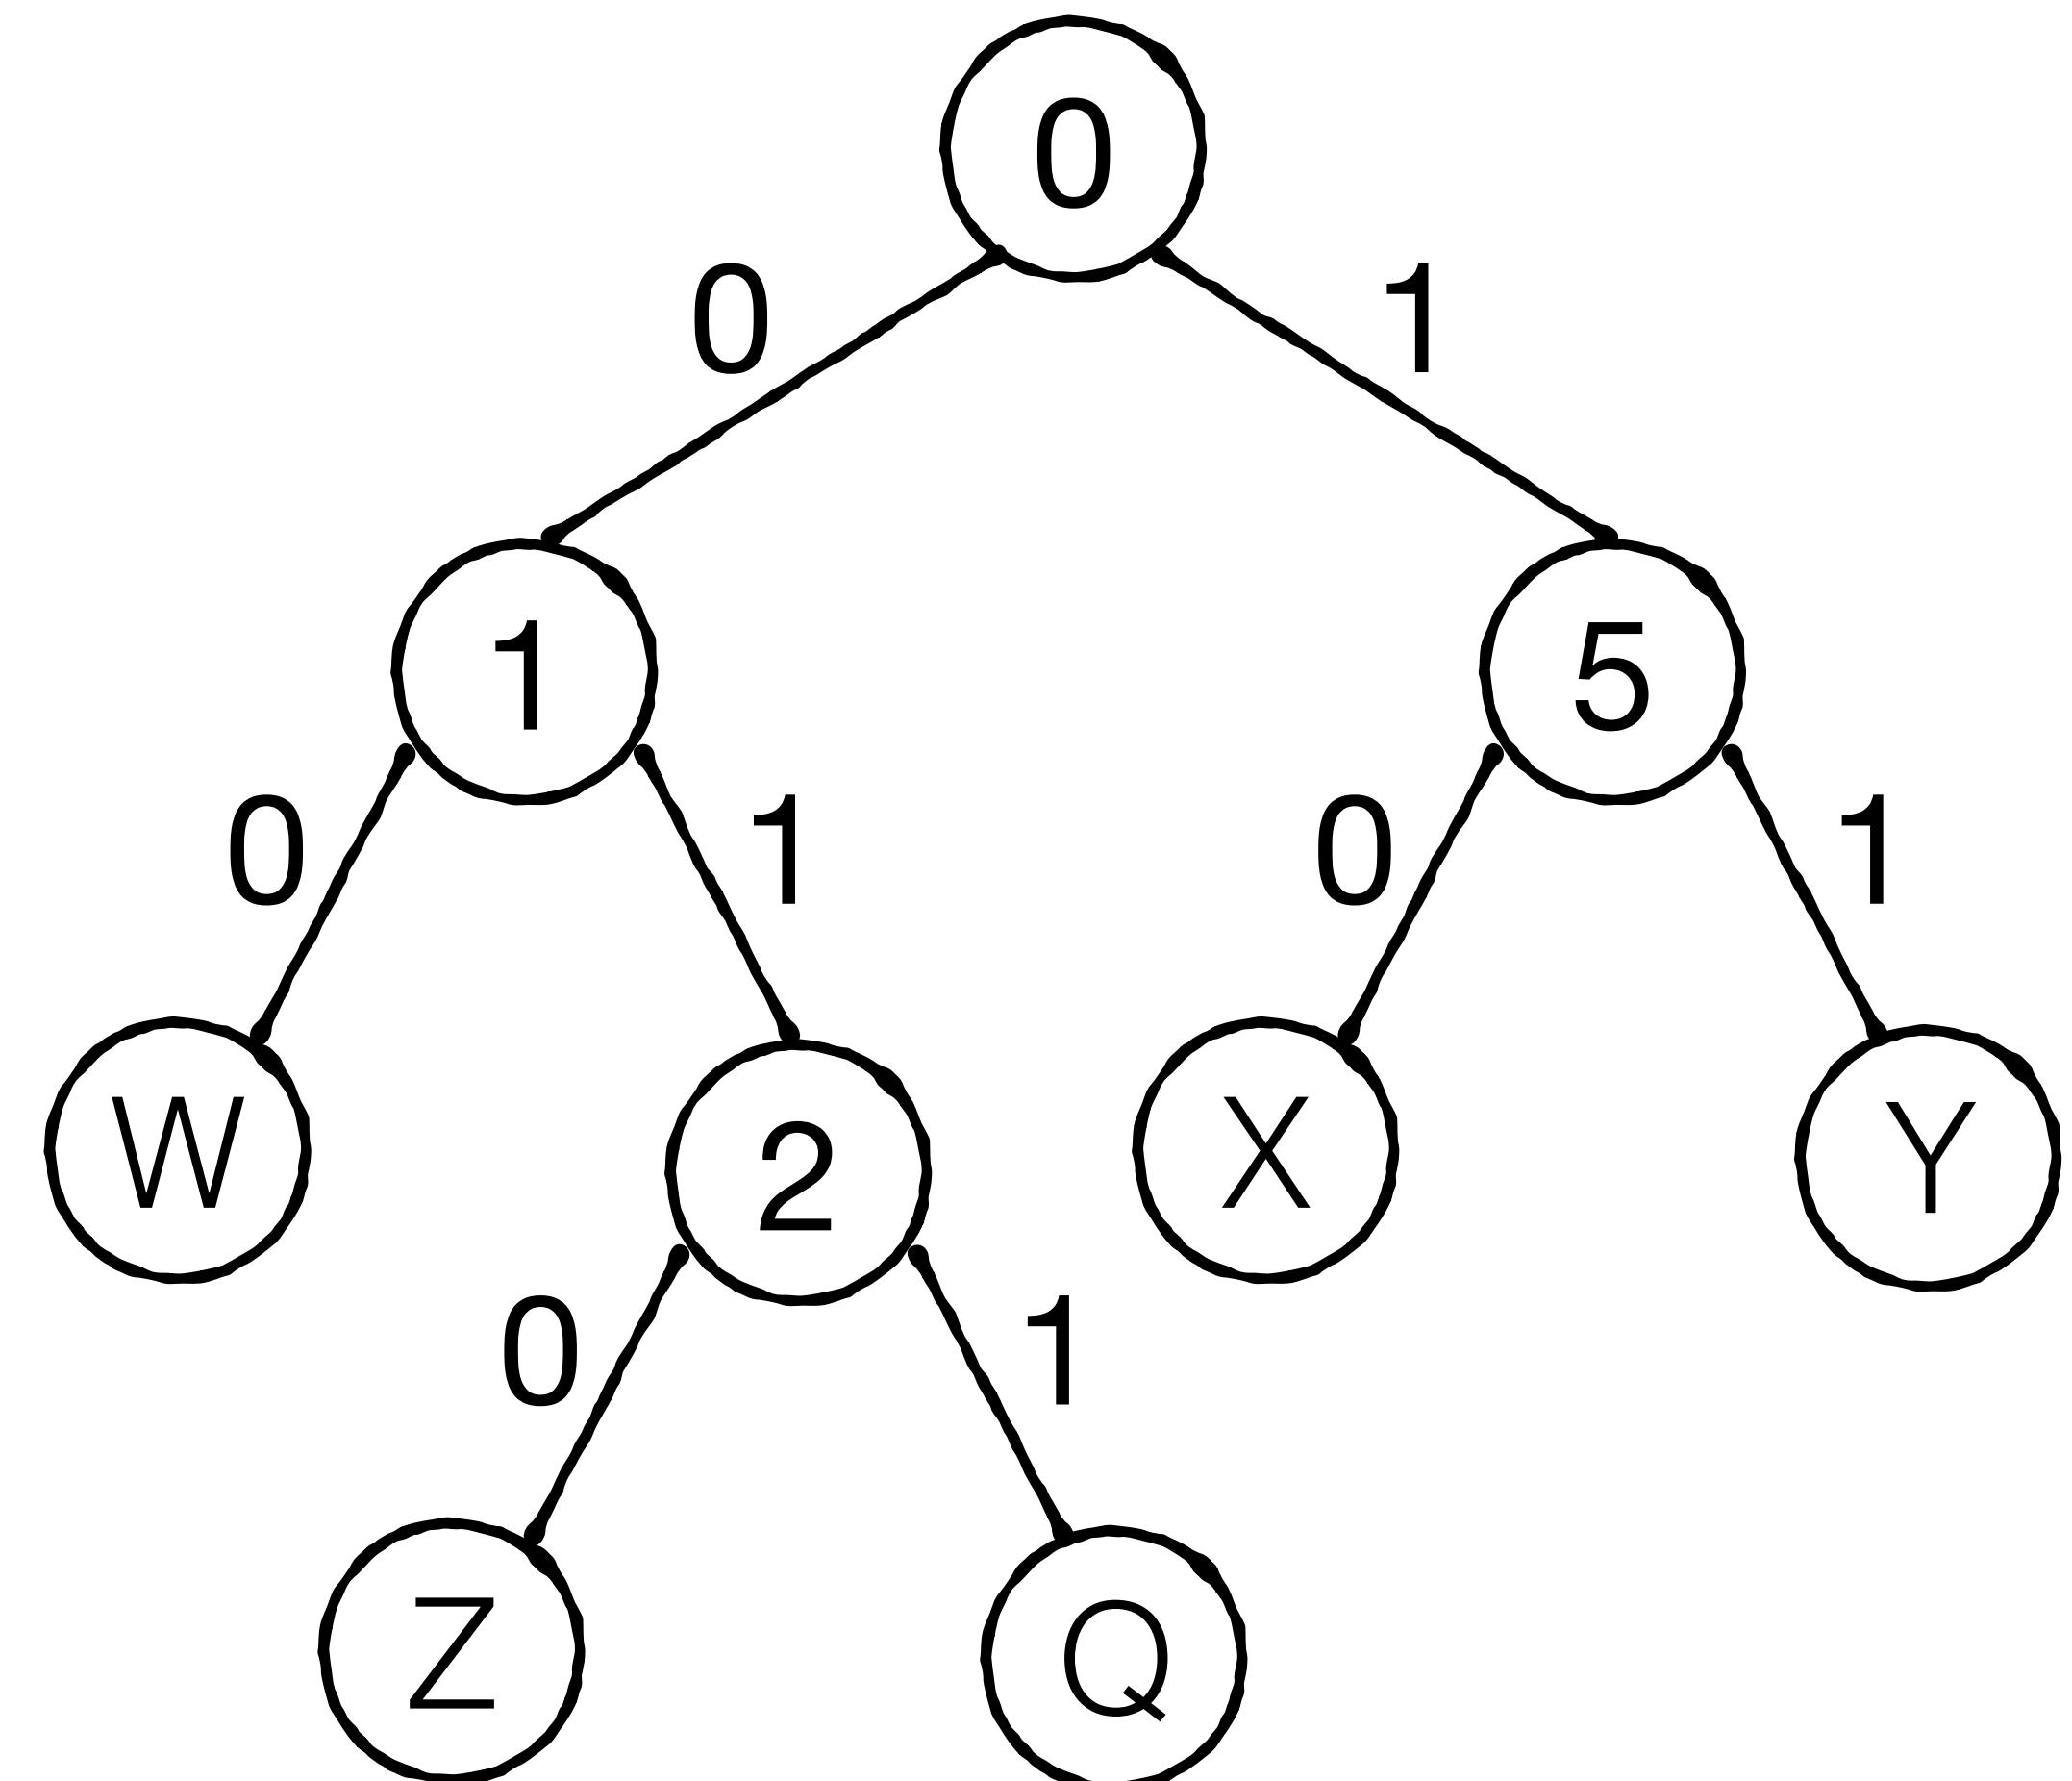


How about a query to non-existing key?

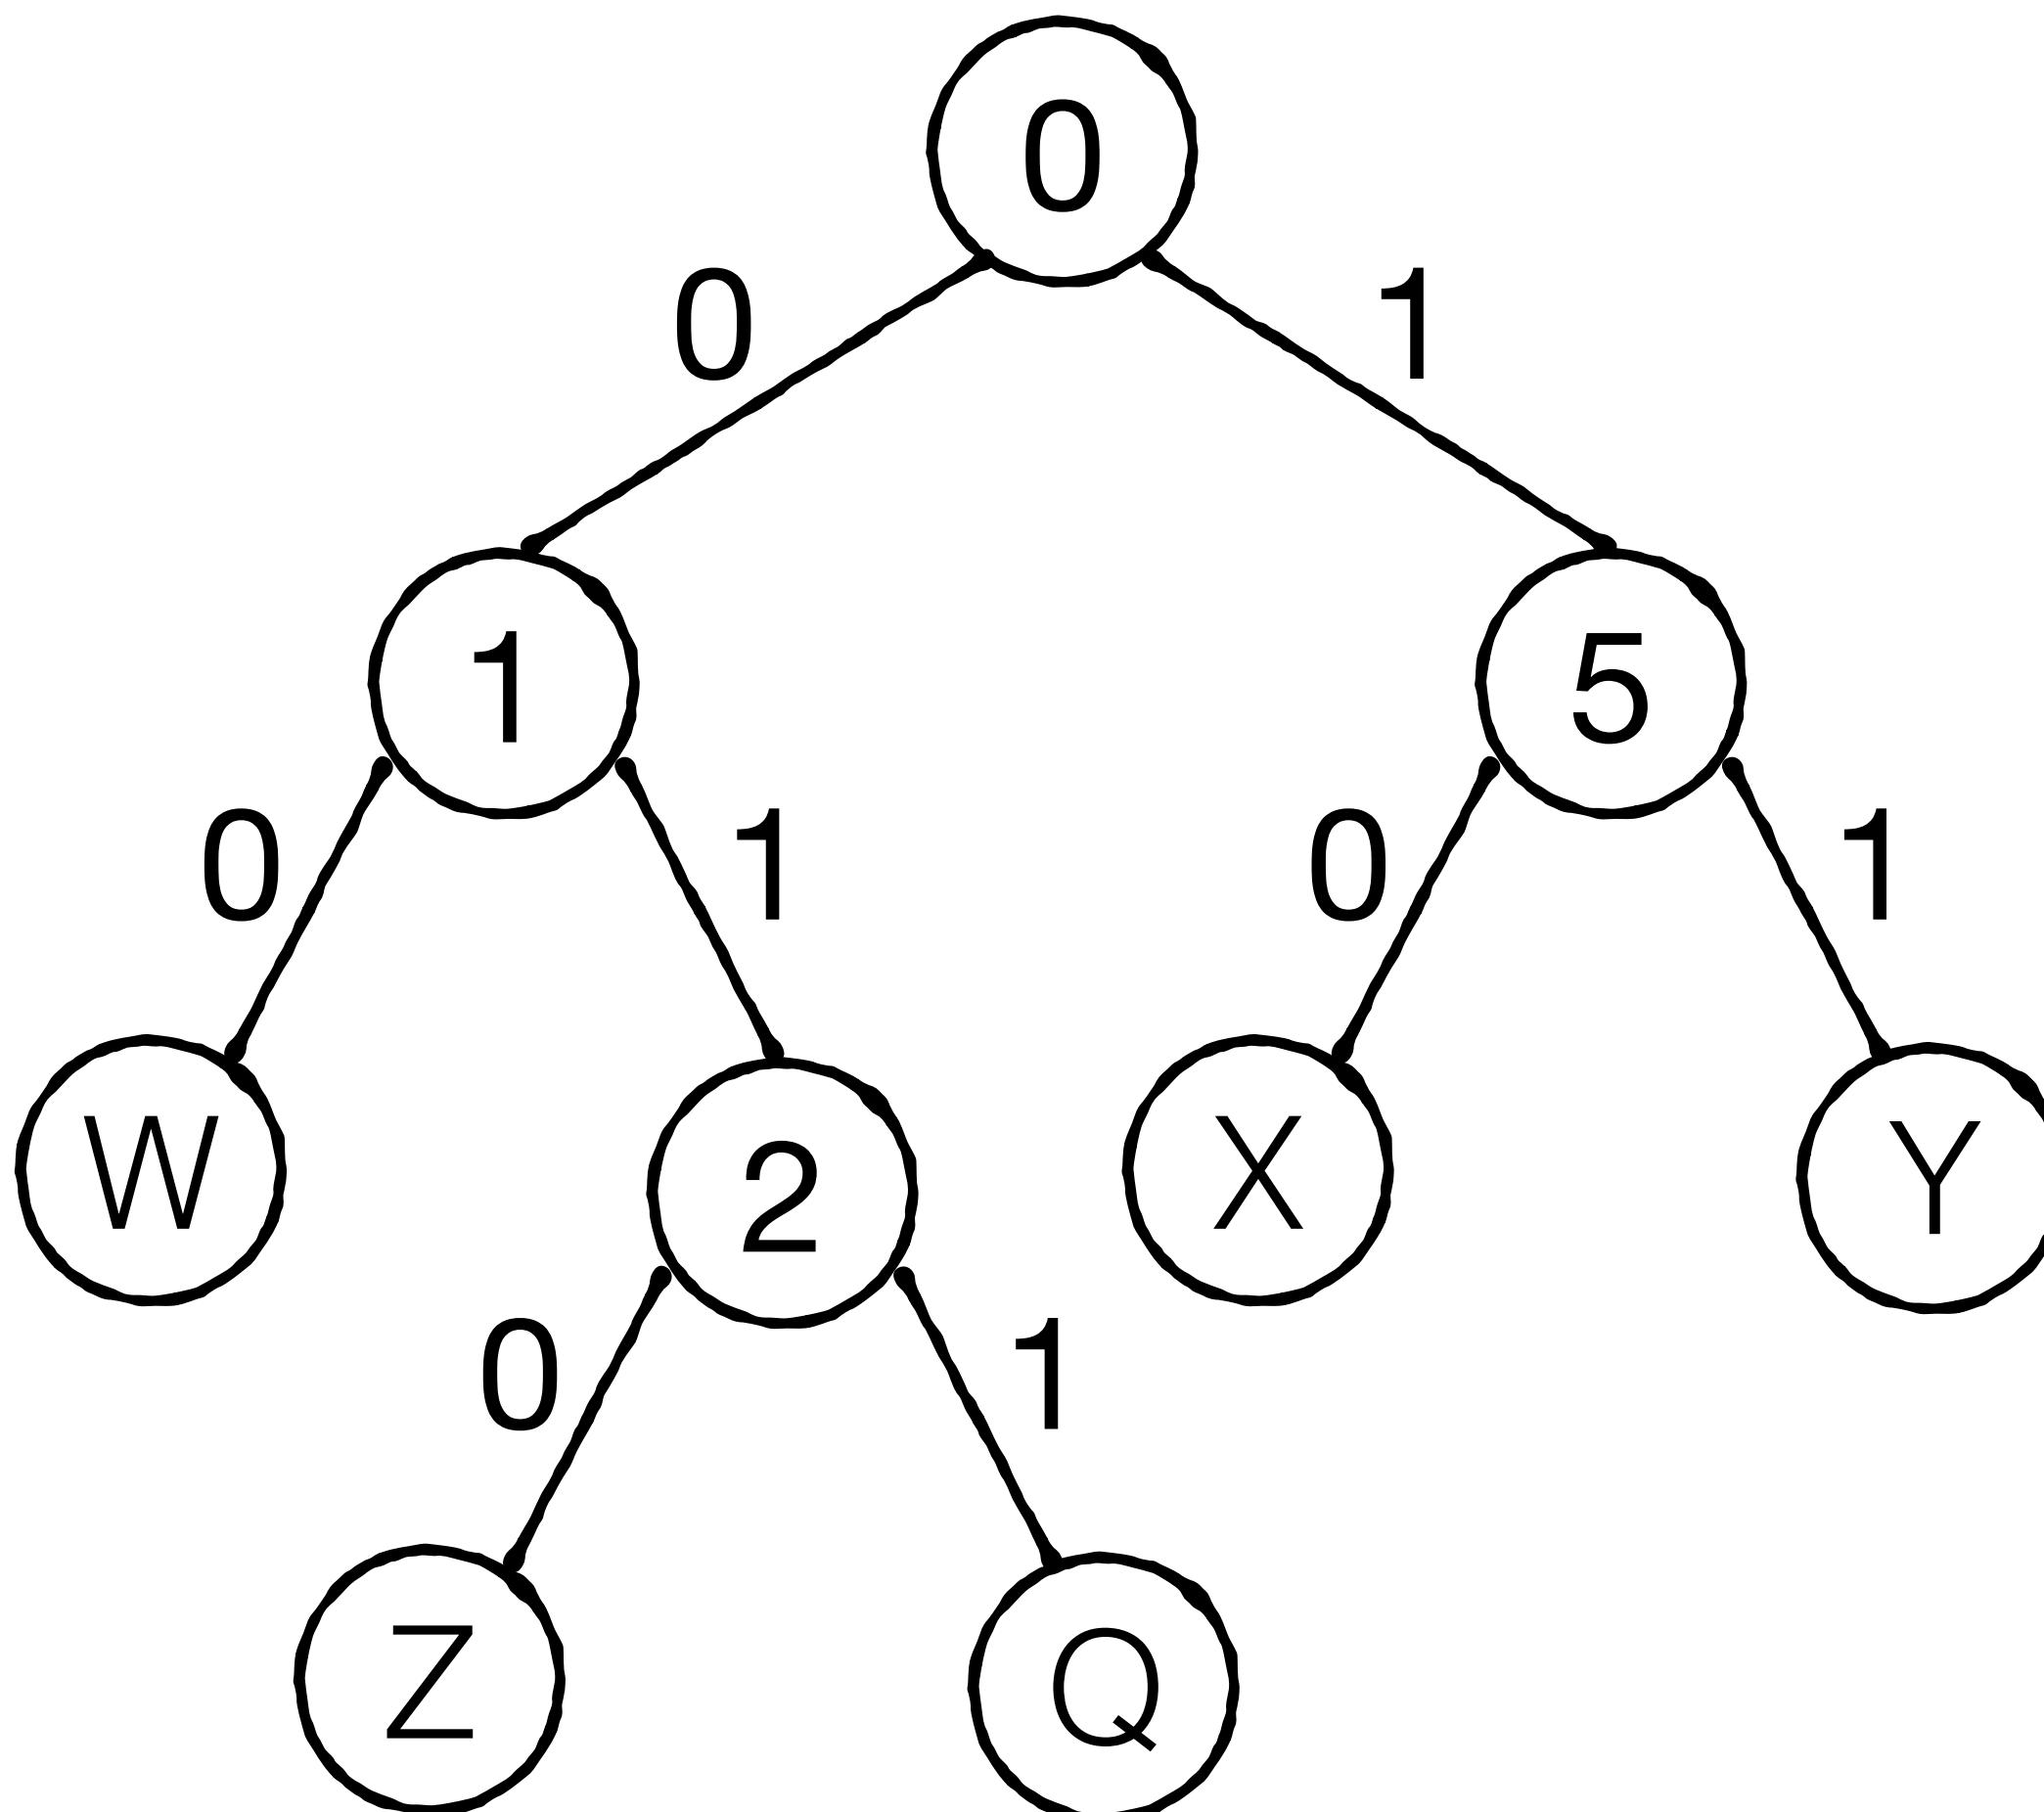


How about a query to non-existing key?

**Finds empty slot or returns
payload associated with some
other key**

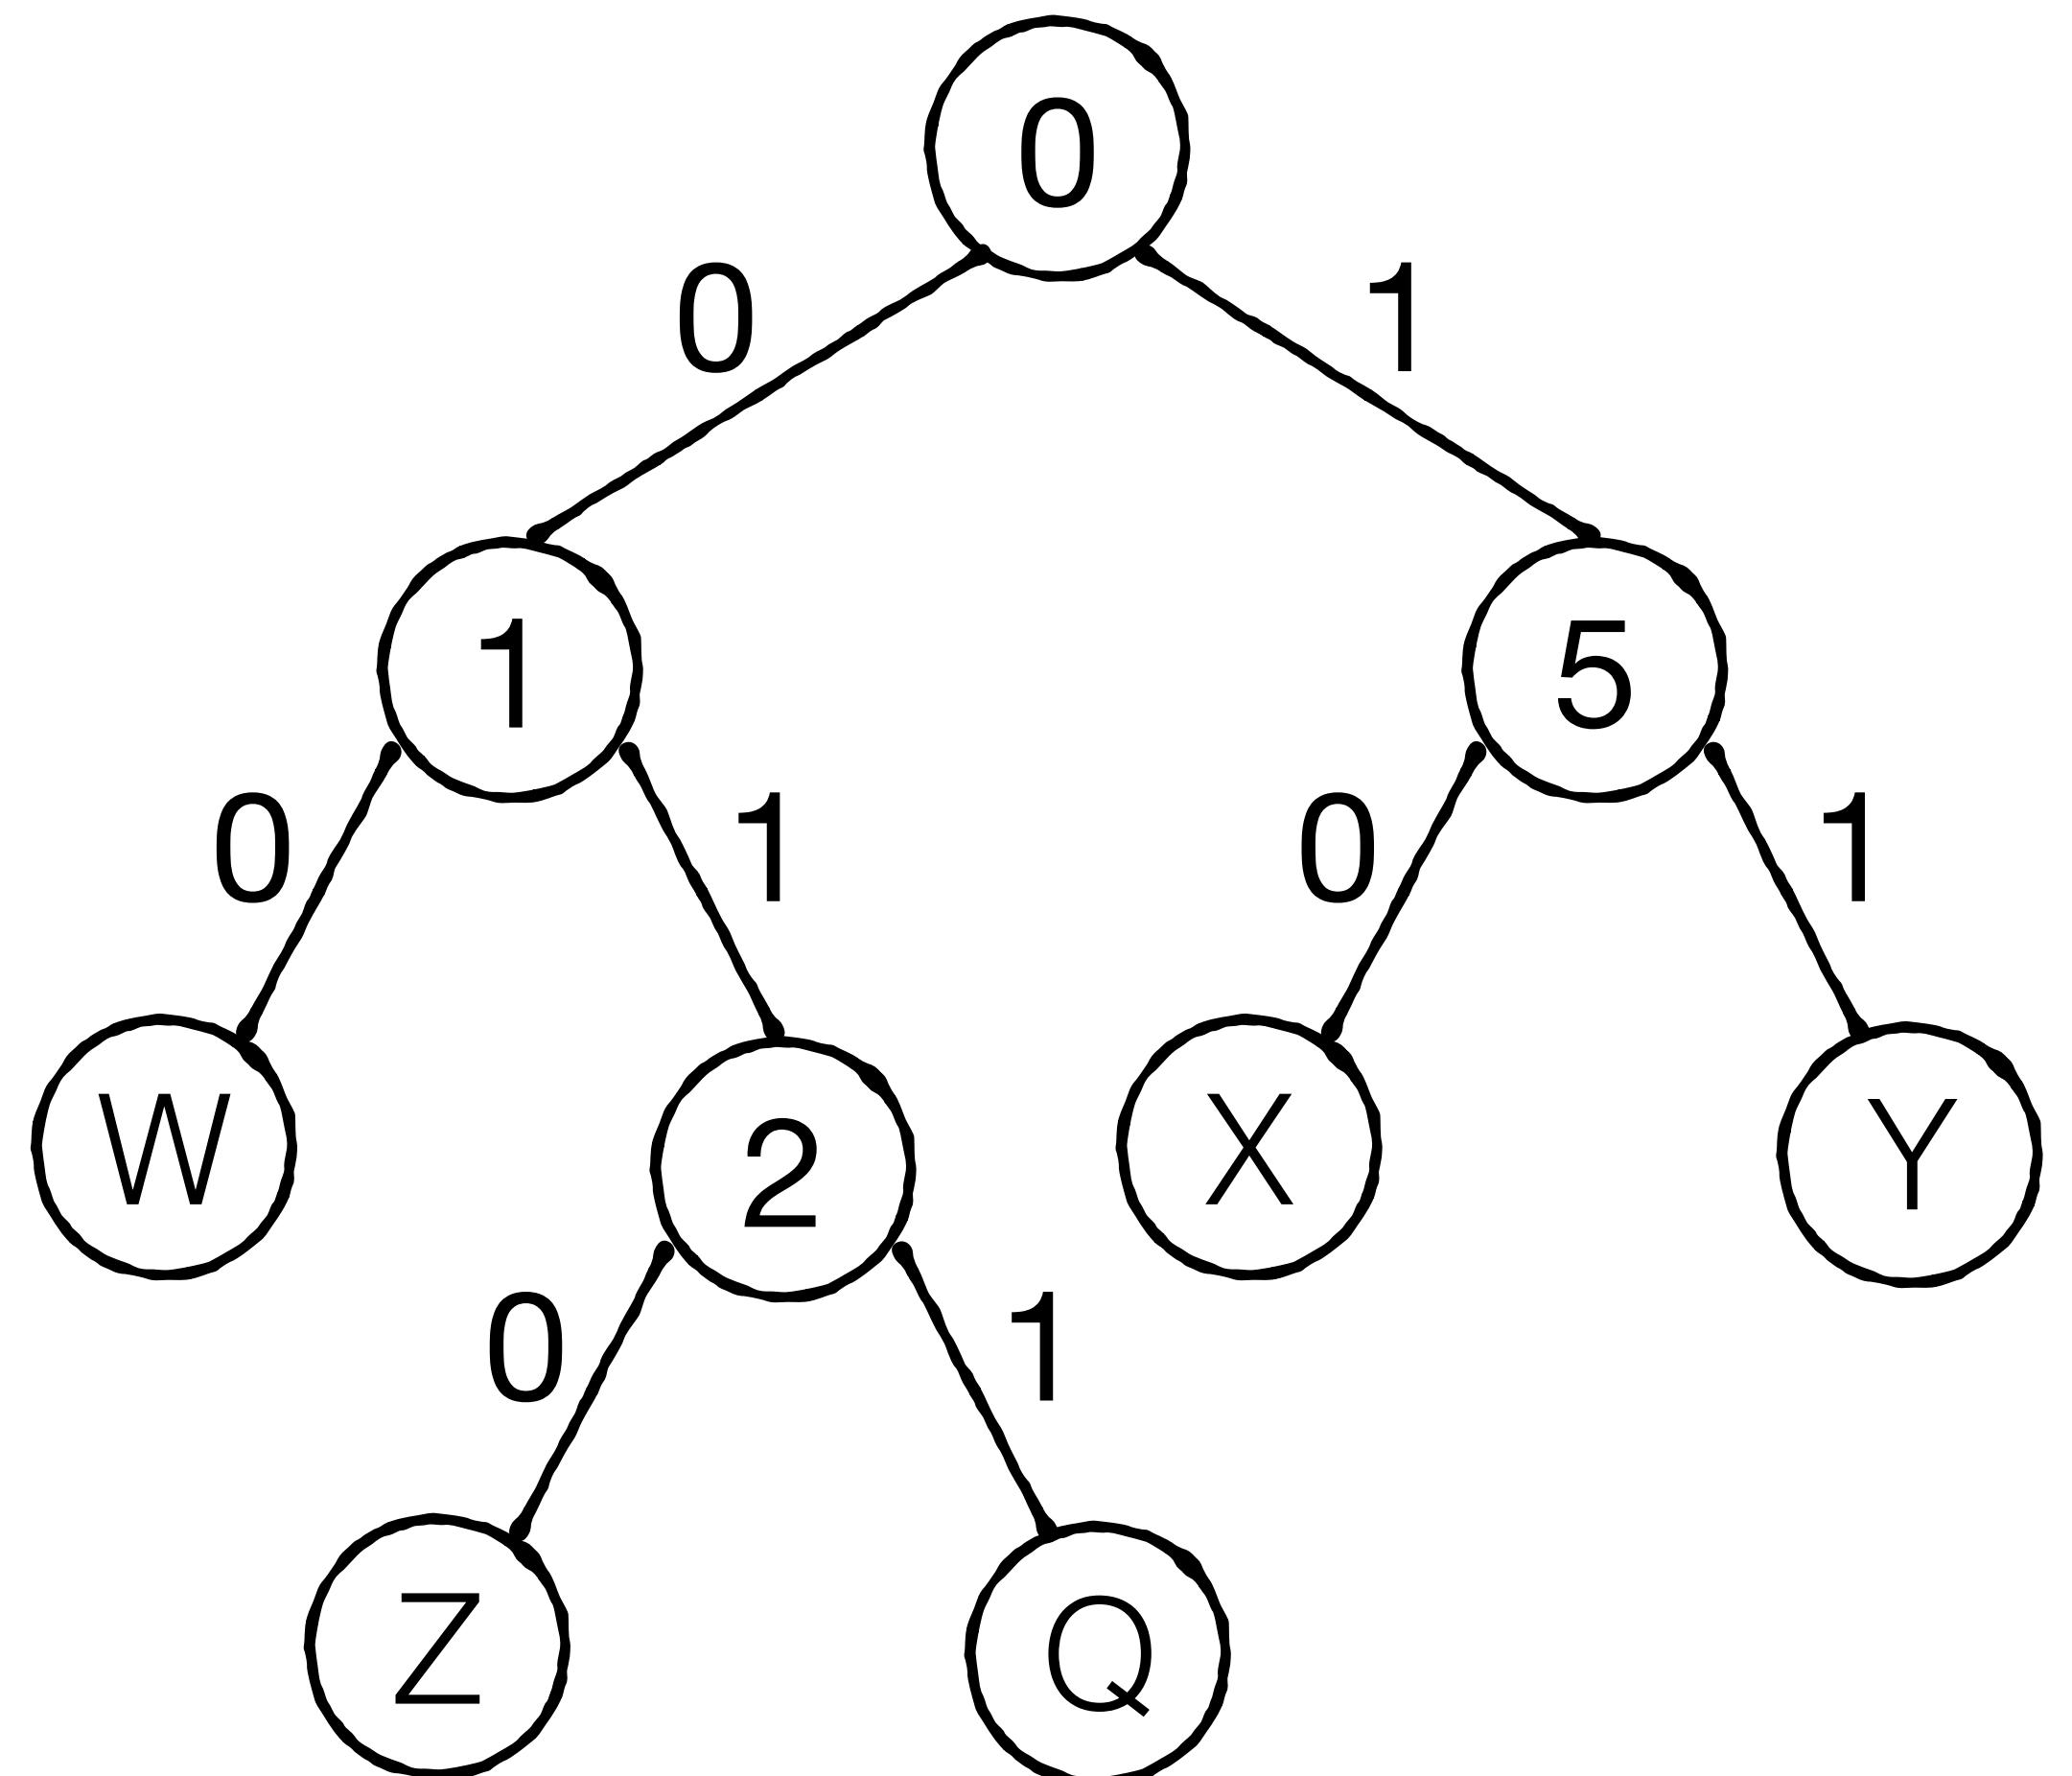


How to encode this trie efficiently within a hash table bucket?

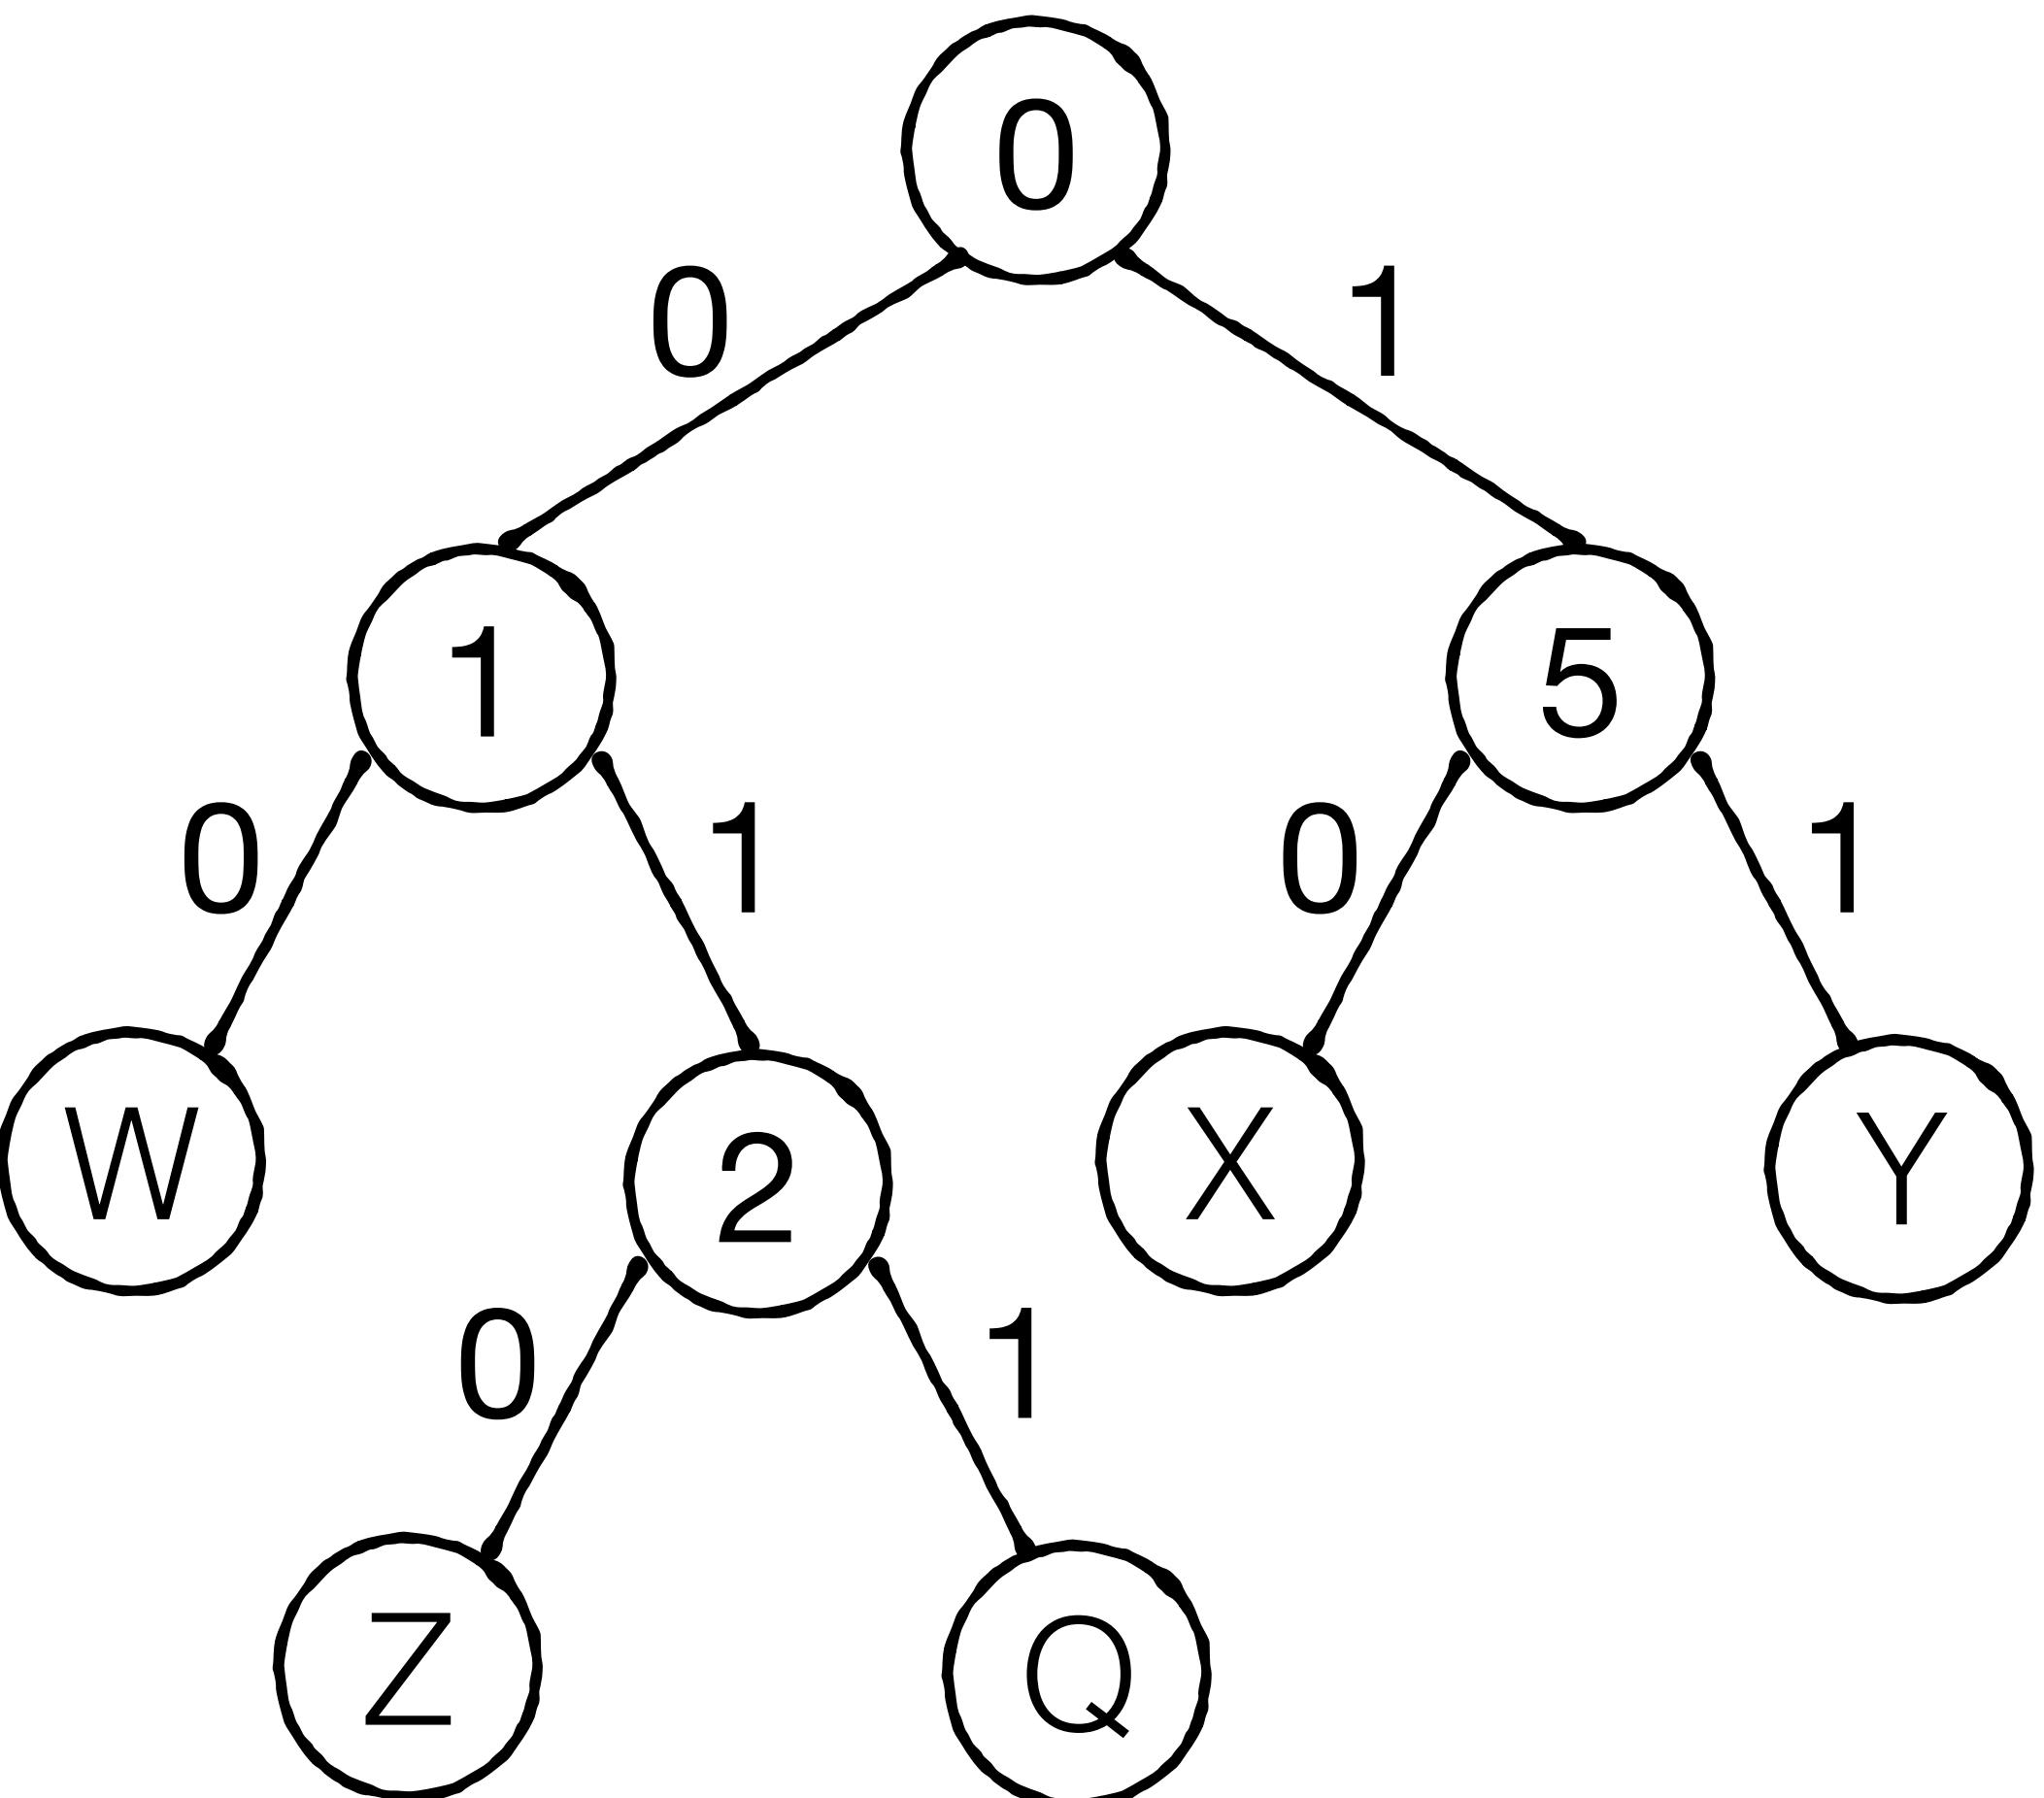


How to encode this trie efficiently within a hash table bucket?

Encode # entries in unary: 111110



How to encode this trie efficiently within a hash table bucket?

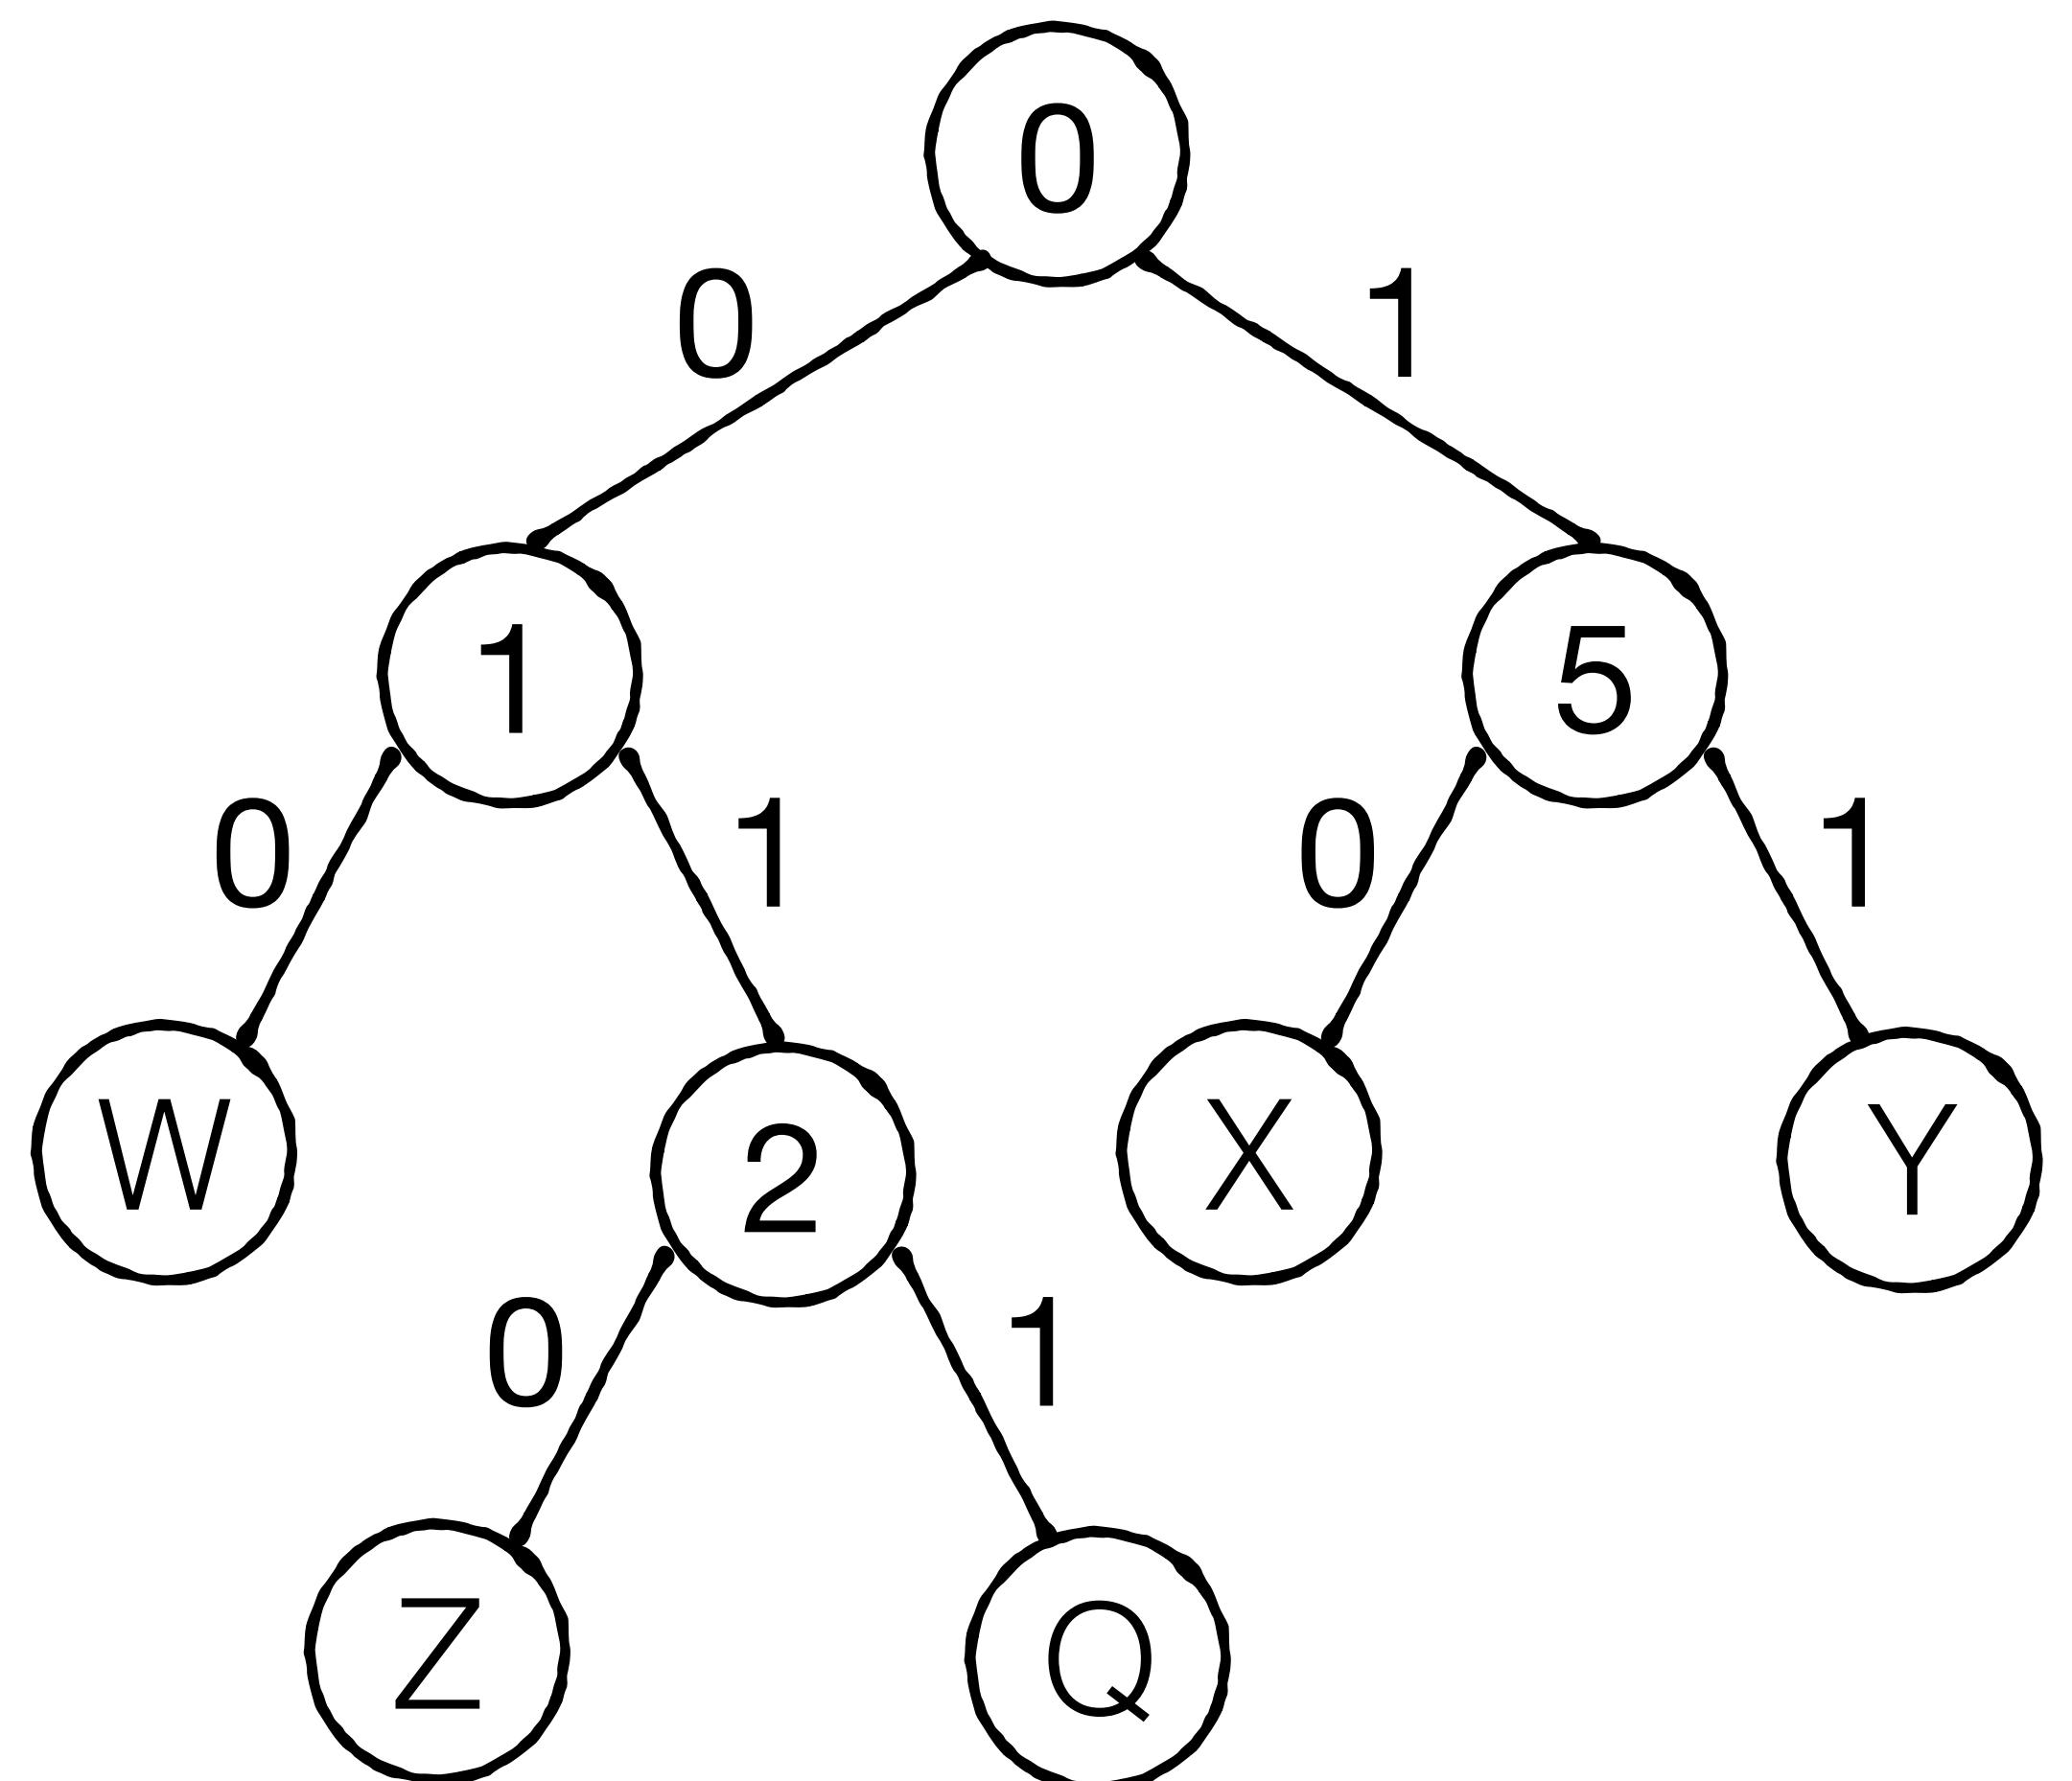


Encode # entries in unary: 111110

Unary is ideal since trie is usually small (0-2 entries)

entries

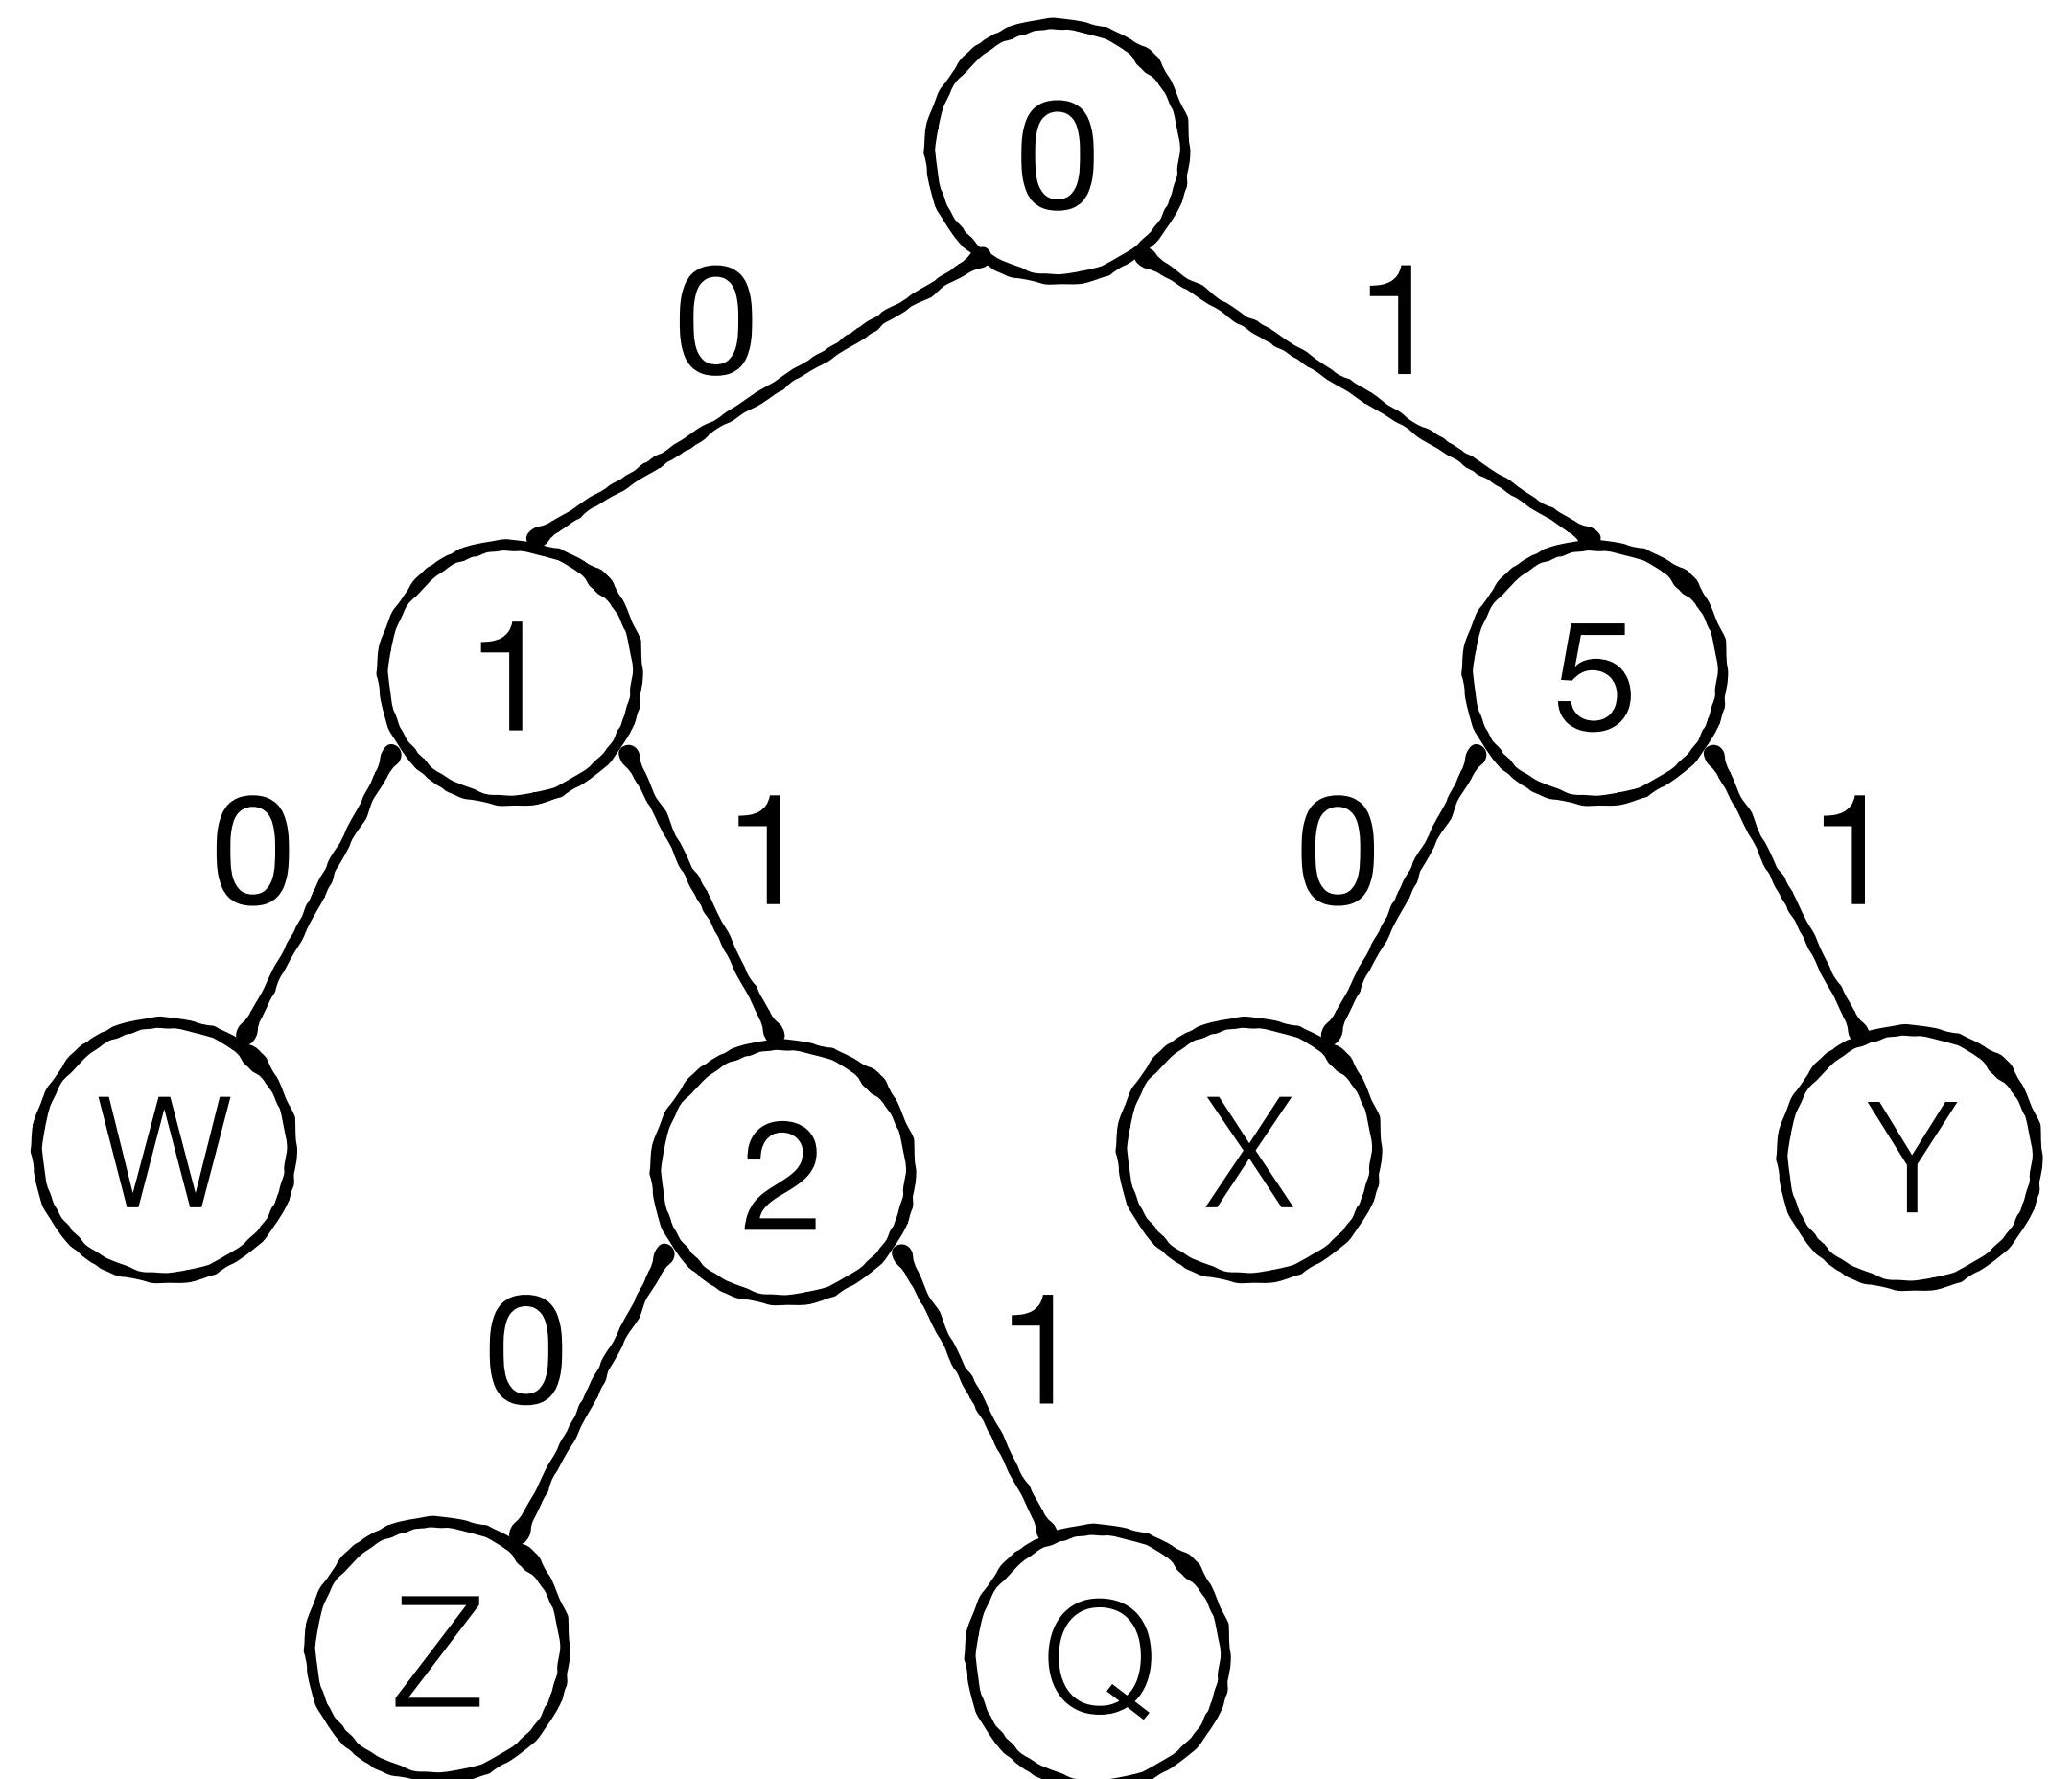
111110



entries

111110

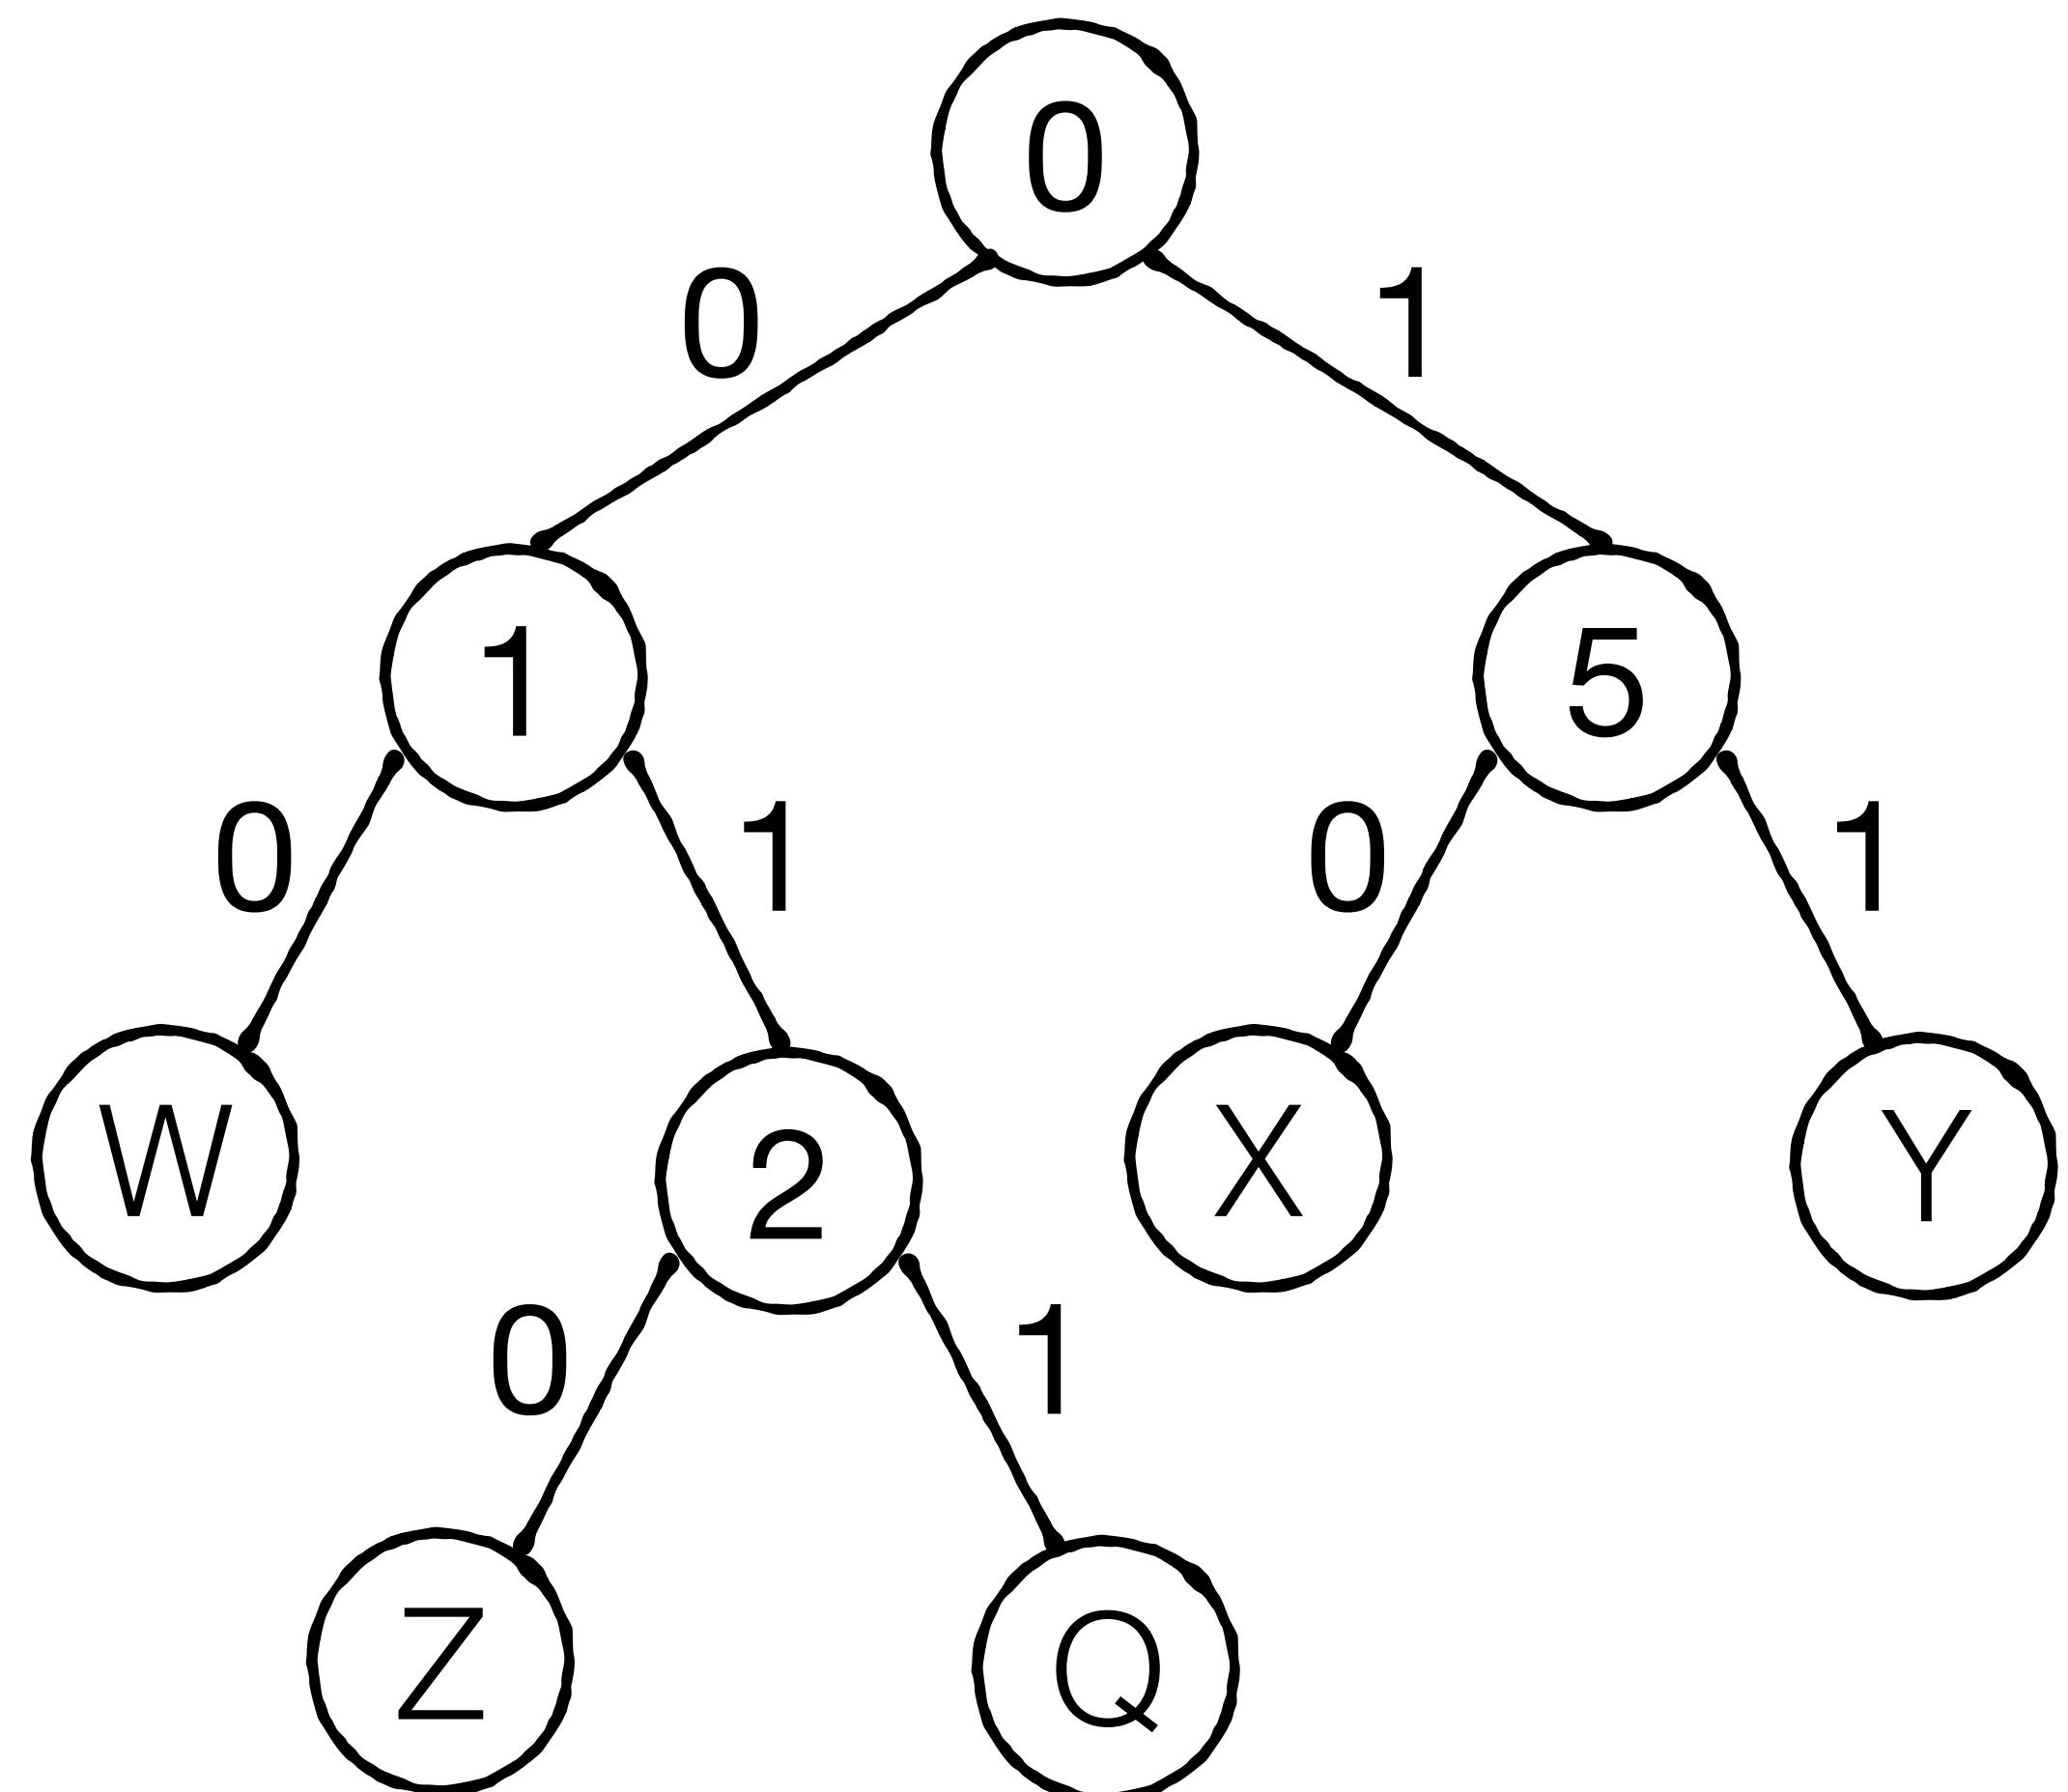
Trie topology?



entries

111110

Trie topology?



00 - no children

01 - one right child node

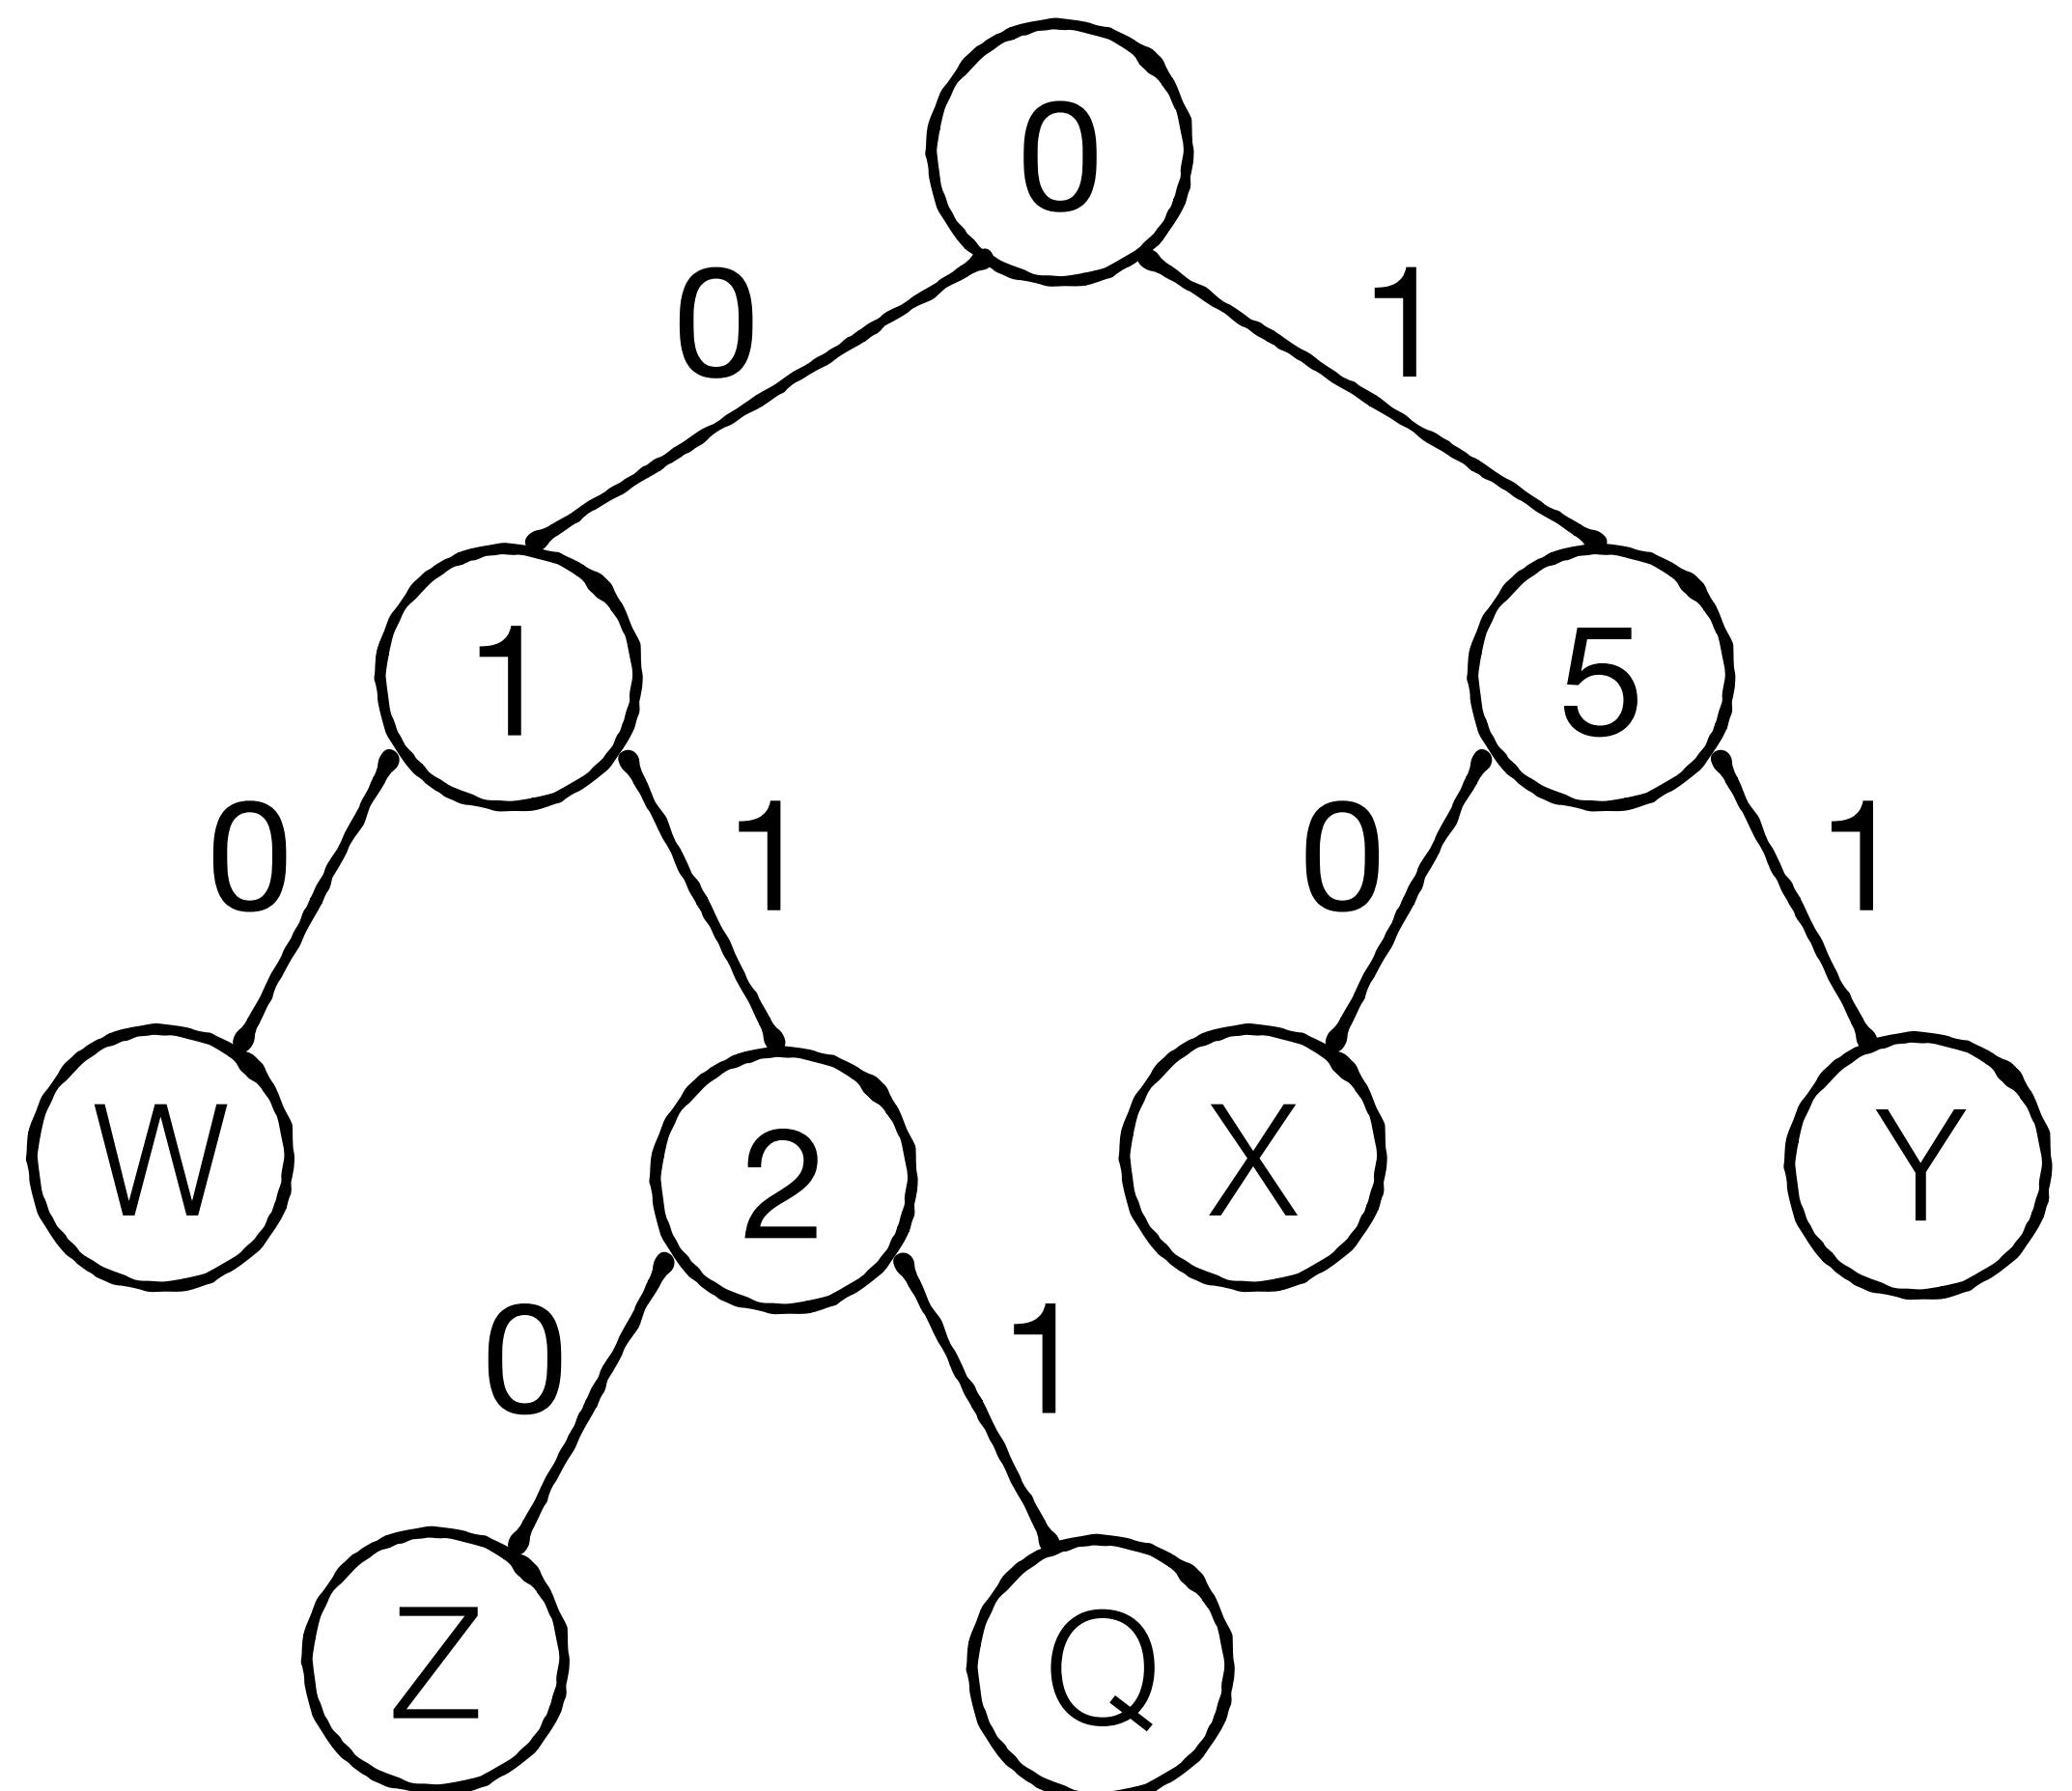
10 - one left child node

11 - two children

entries

111110

Trie topology?



00 - no children

01 - one right child node

10 - one left child node

11 - two children

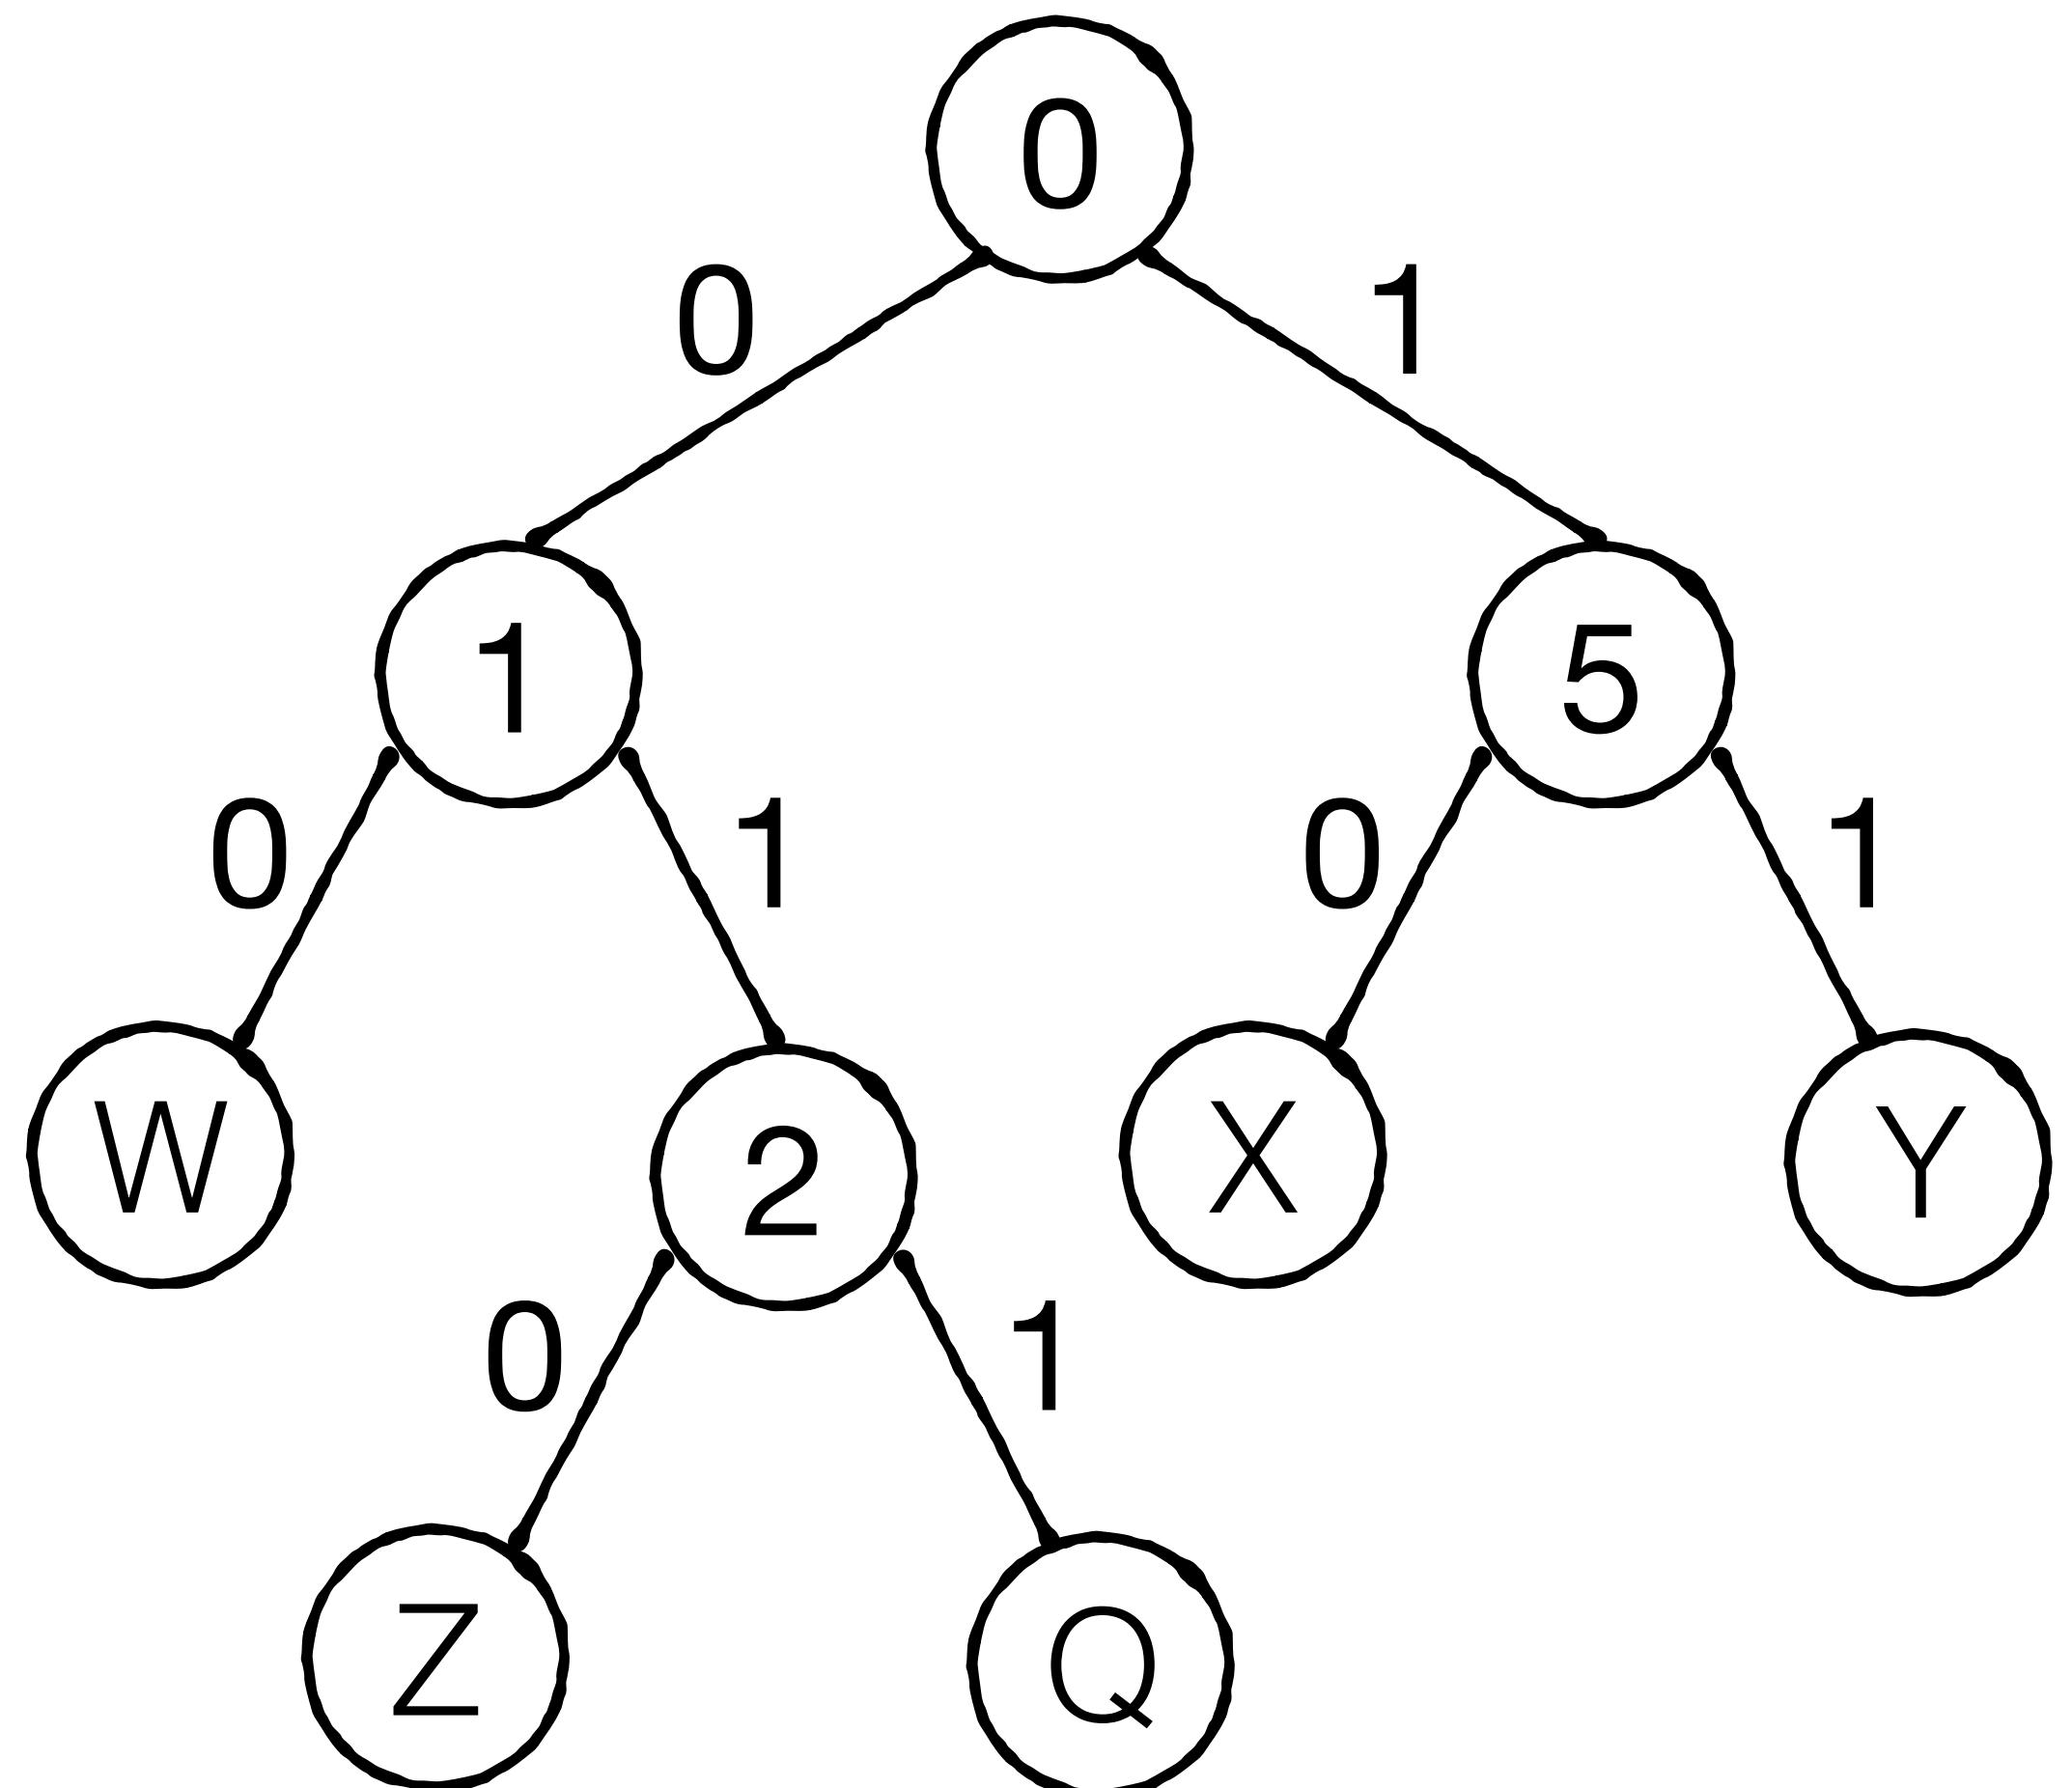
Encode structure depth-first

11 01 00 00

entries

111110

Trie topology?



00 - no children

01 - one right child node

10 - one left child node

11 - two children

Encode structure depth-first

11 01 00 00

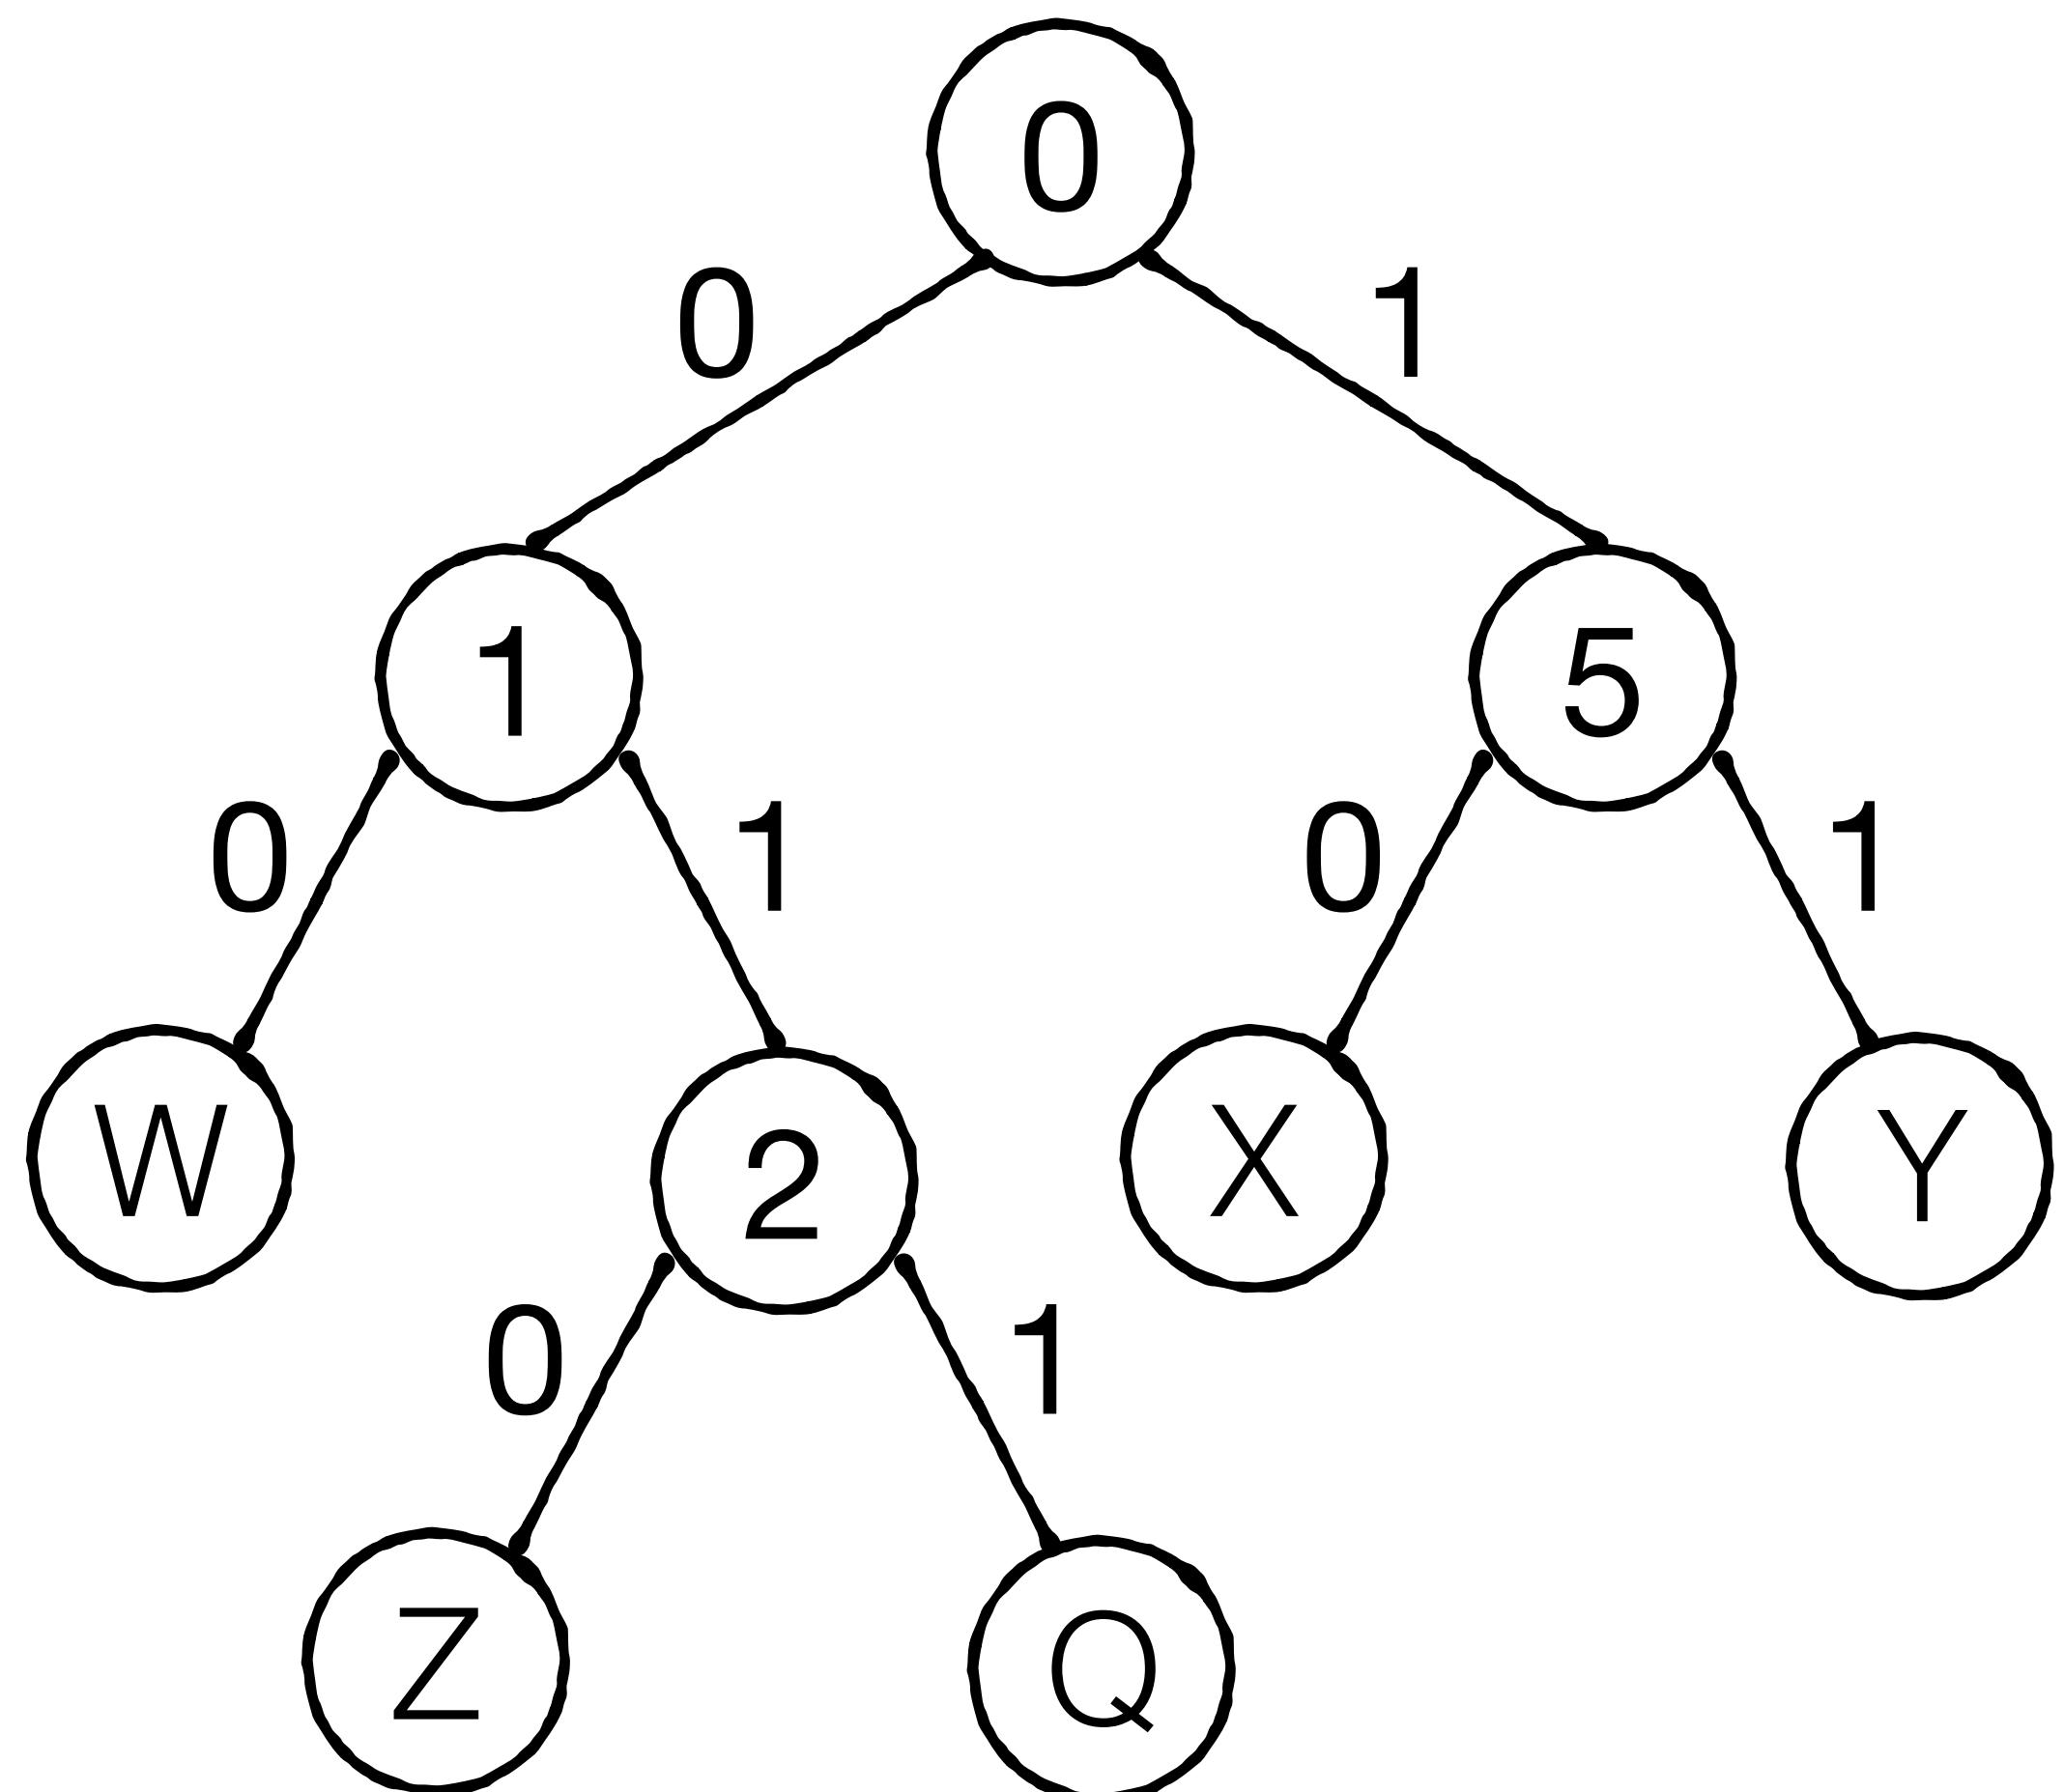


Last bits always zero

entries

111110

Trie topology?



00 - no children

01 - one right child node

10 - one left child node

11 - two children

Encode structure depth-first

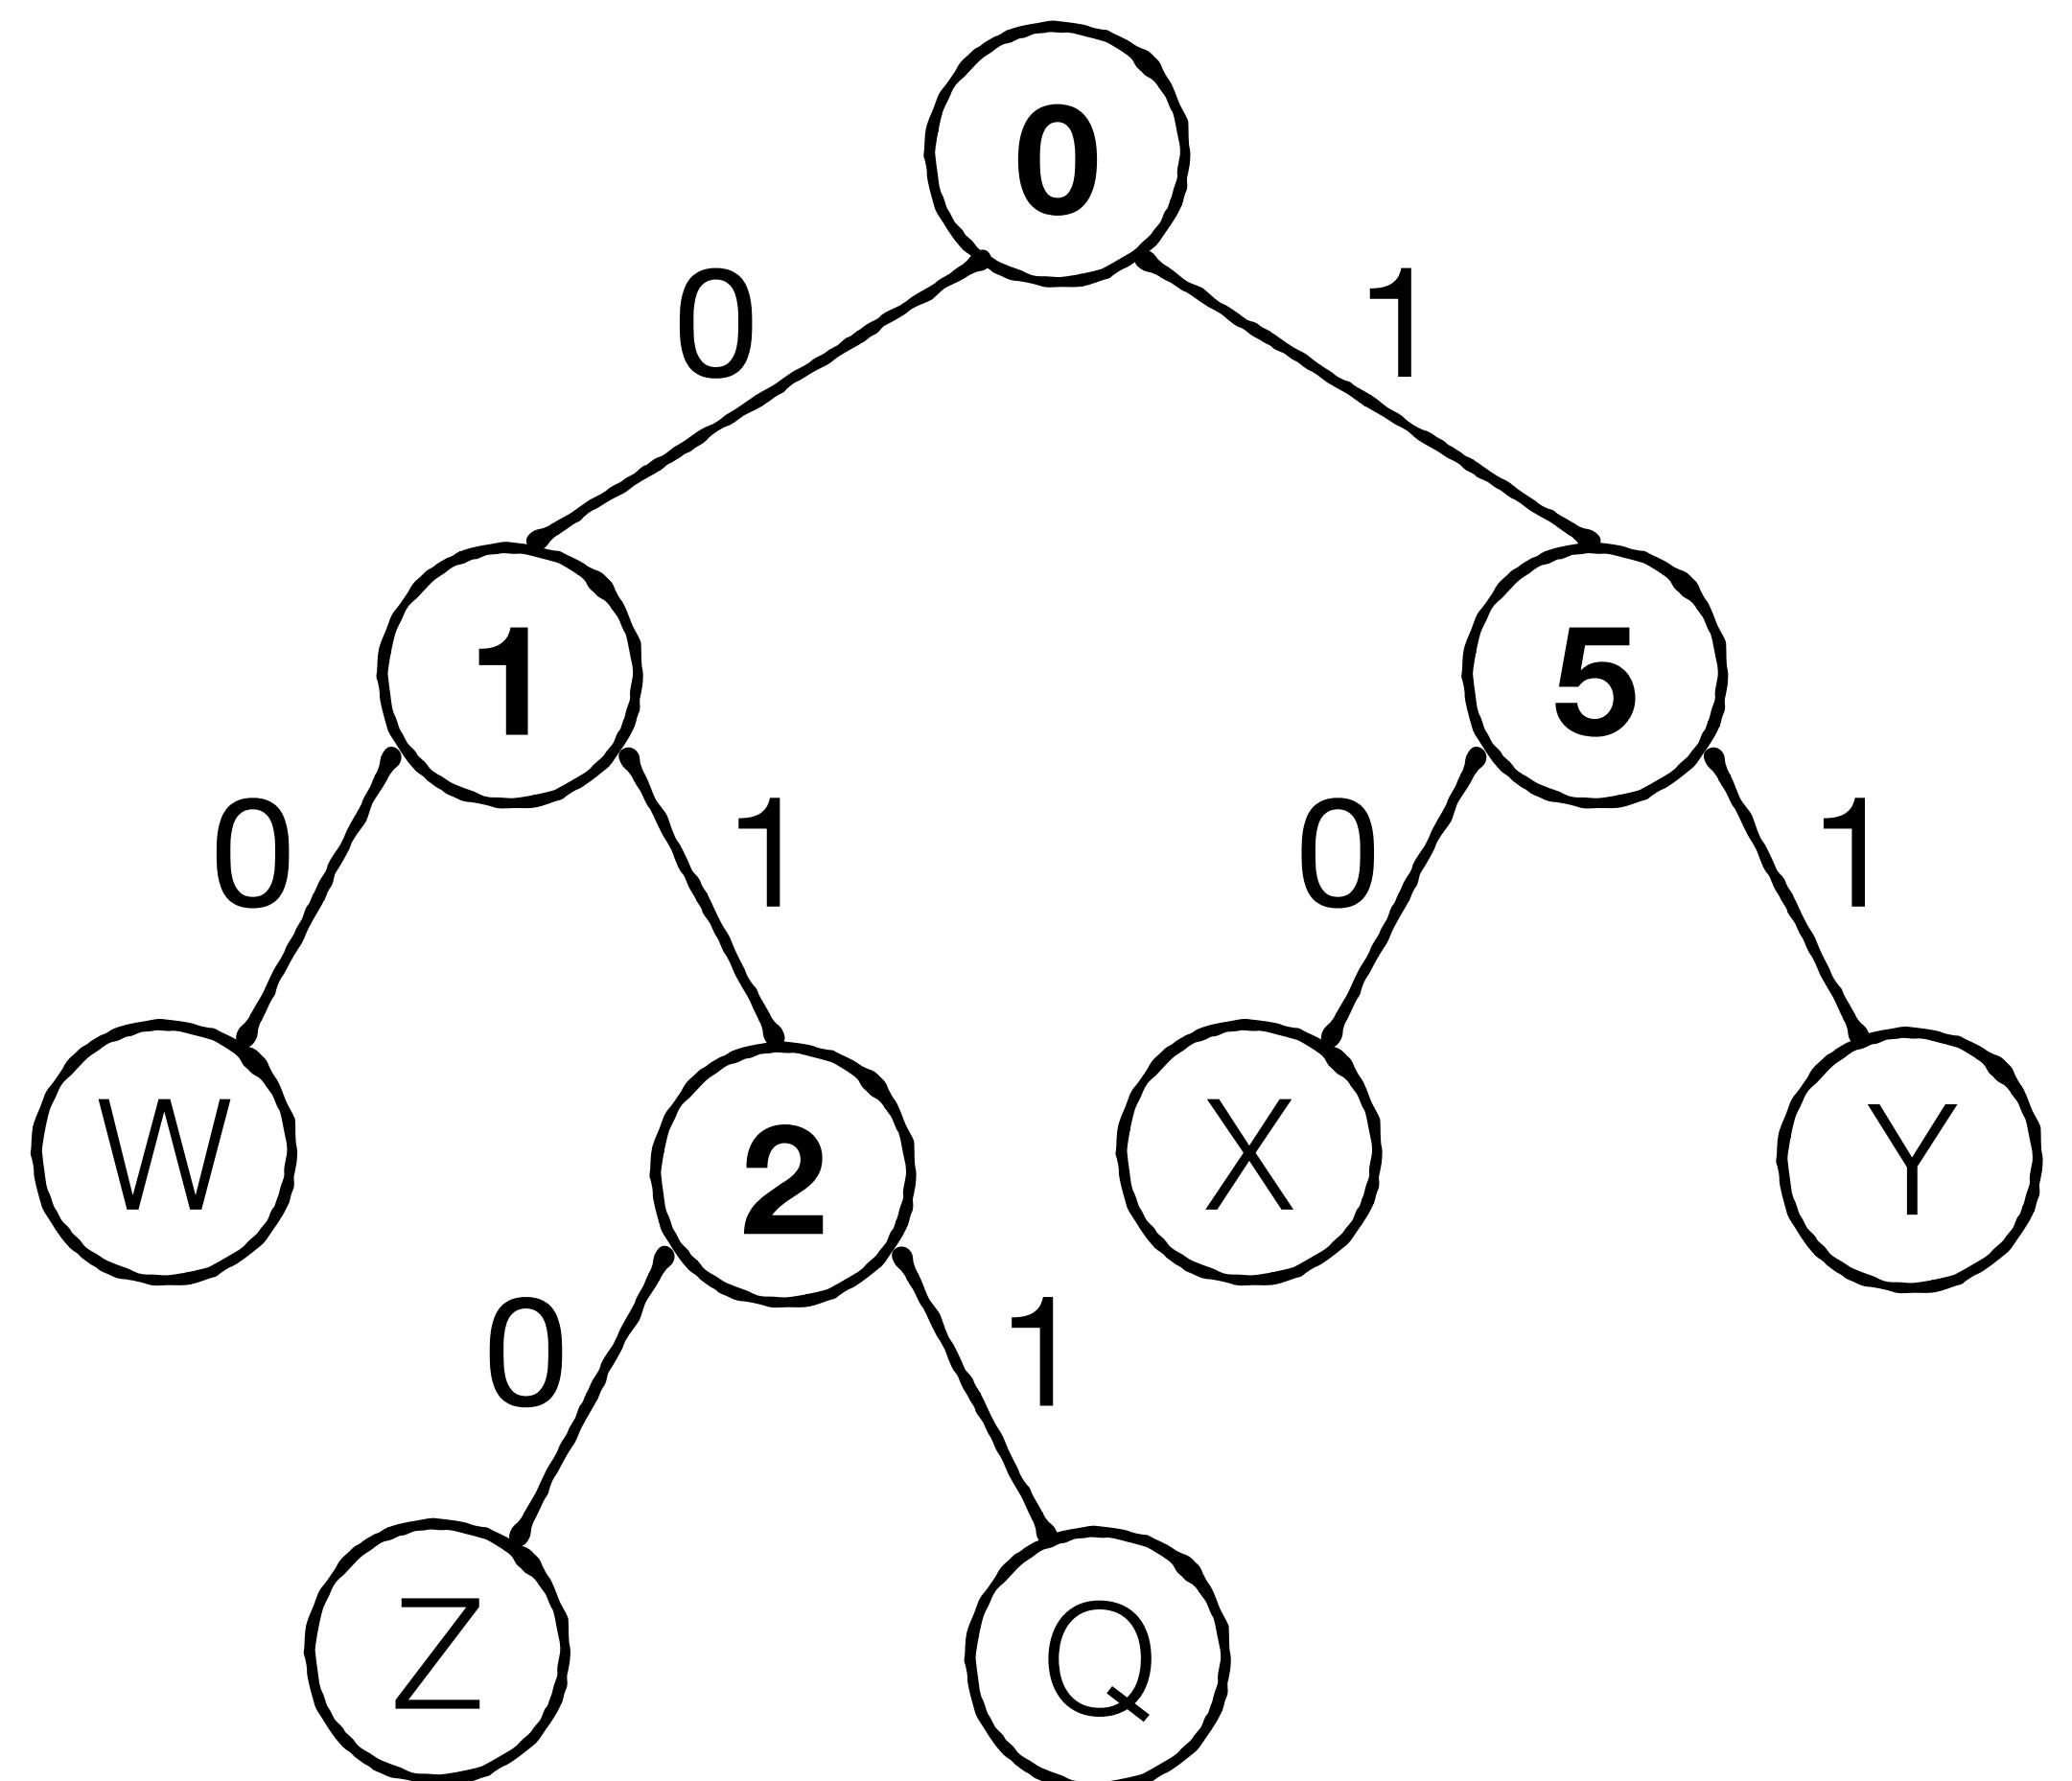
11 01 00

entries

111110

Topology

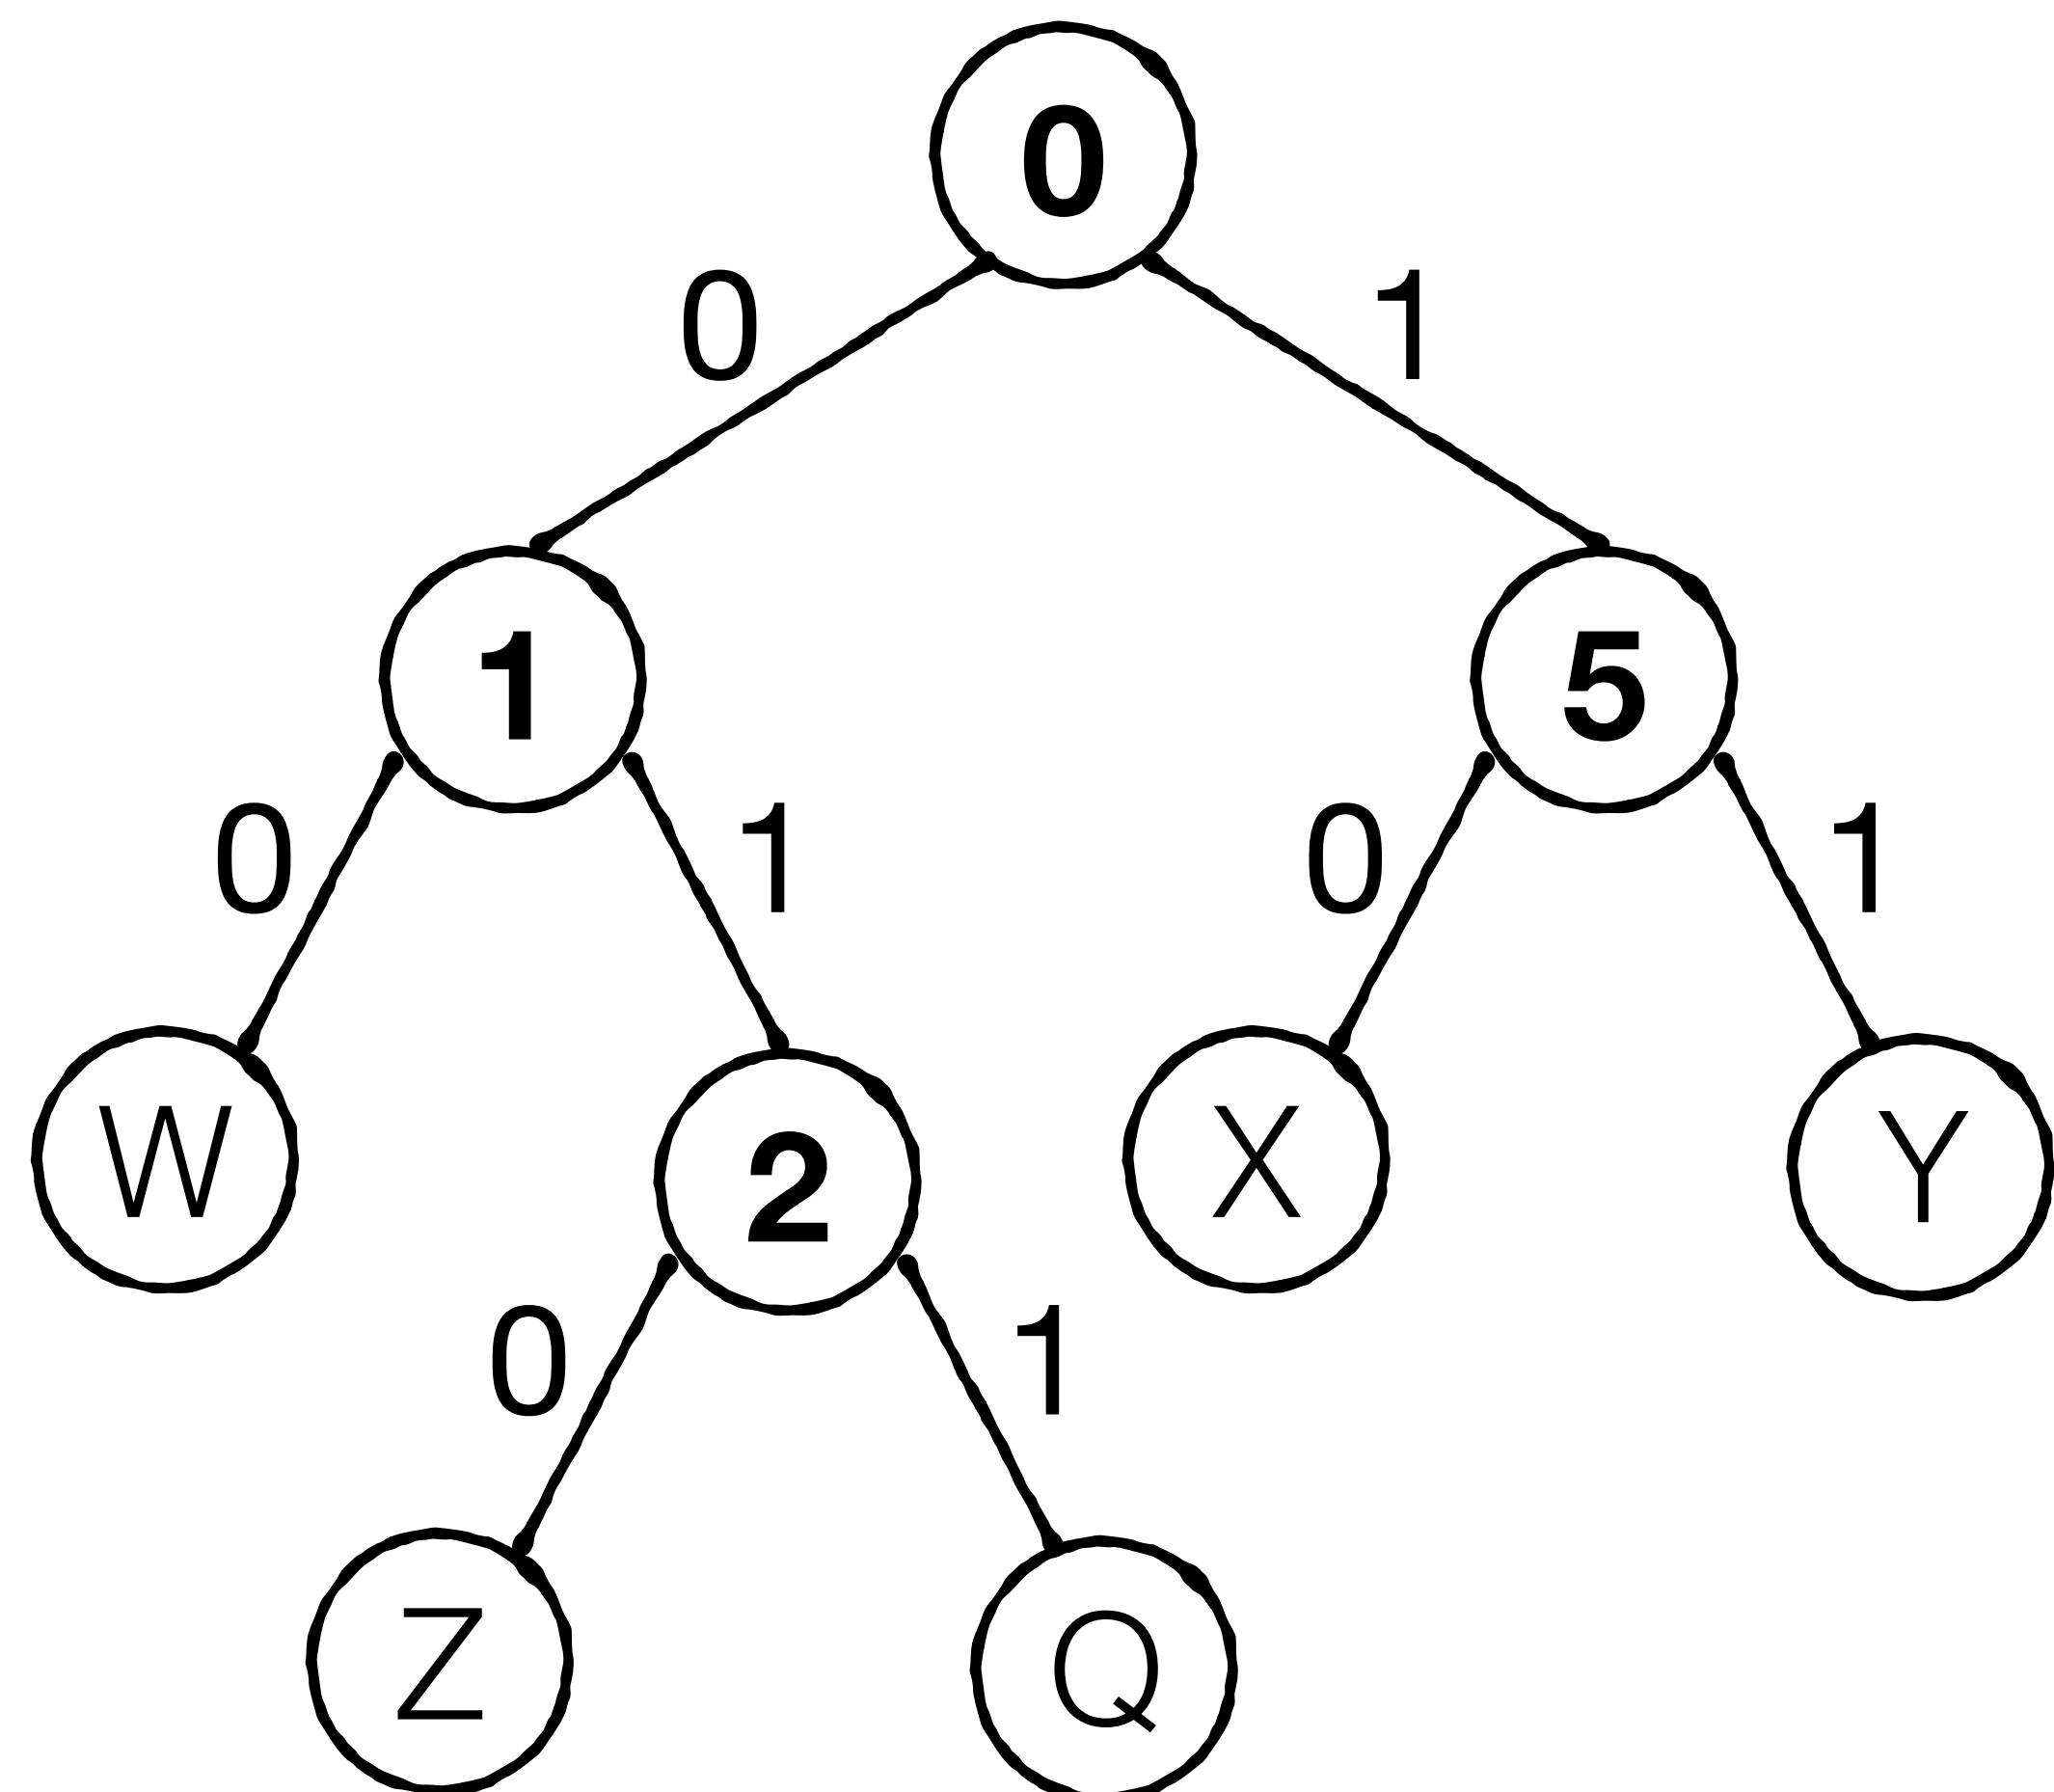
11 01 00



entries
111110

Topology
11 01 00

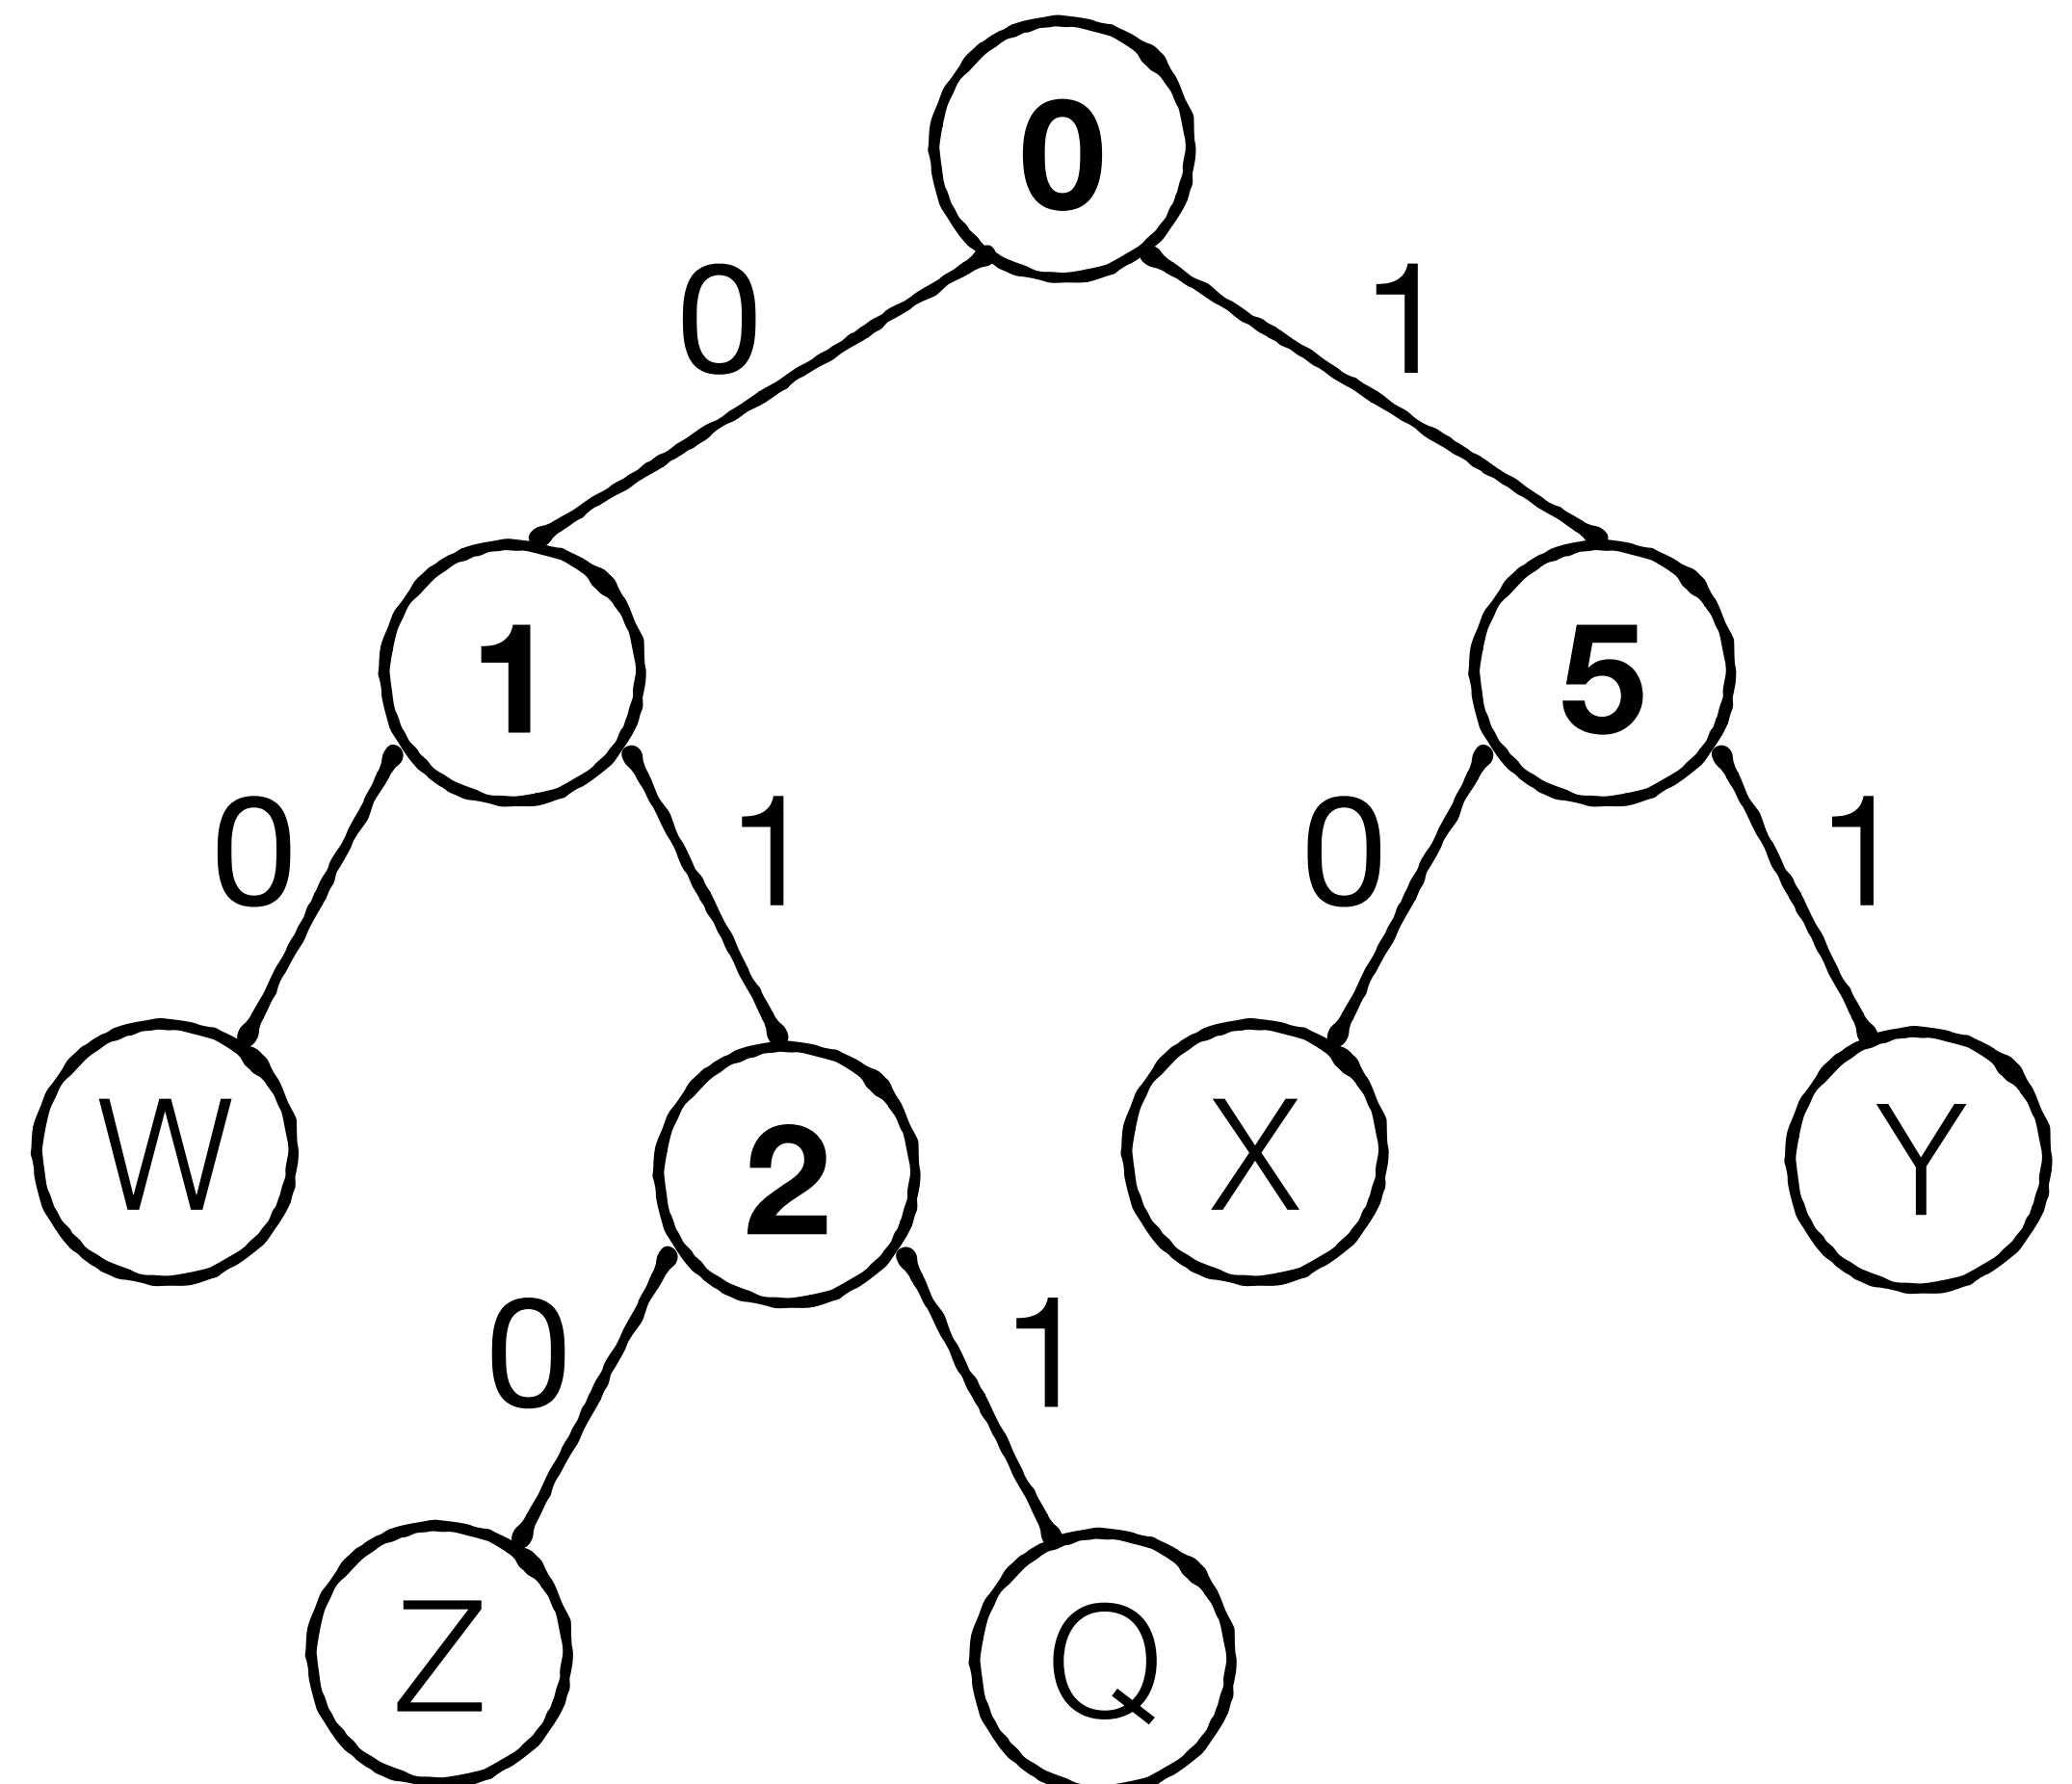
Differentiating bit indexes?



entries
111110

Topology
11 01 00

Differentiating bit indexes?
Depth-first, unary
0 10 110 111110



entries

111110

Topology

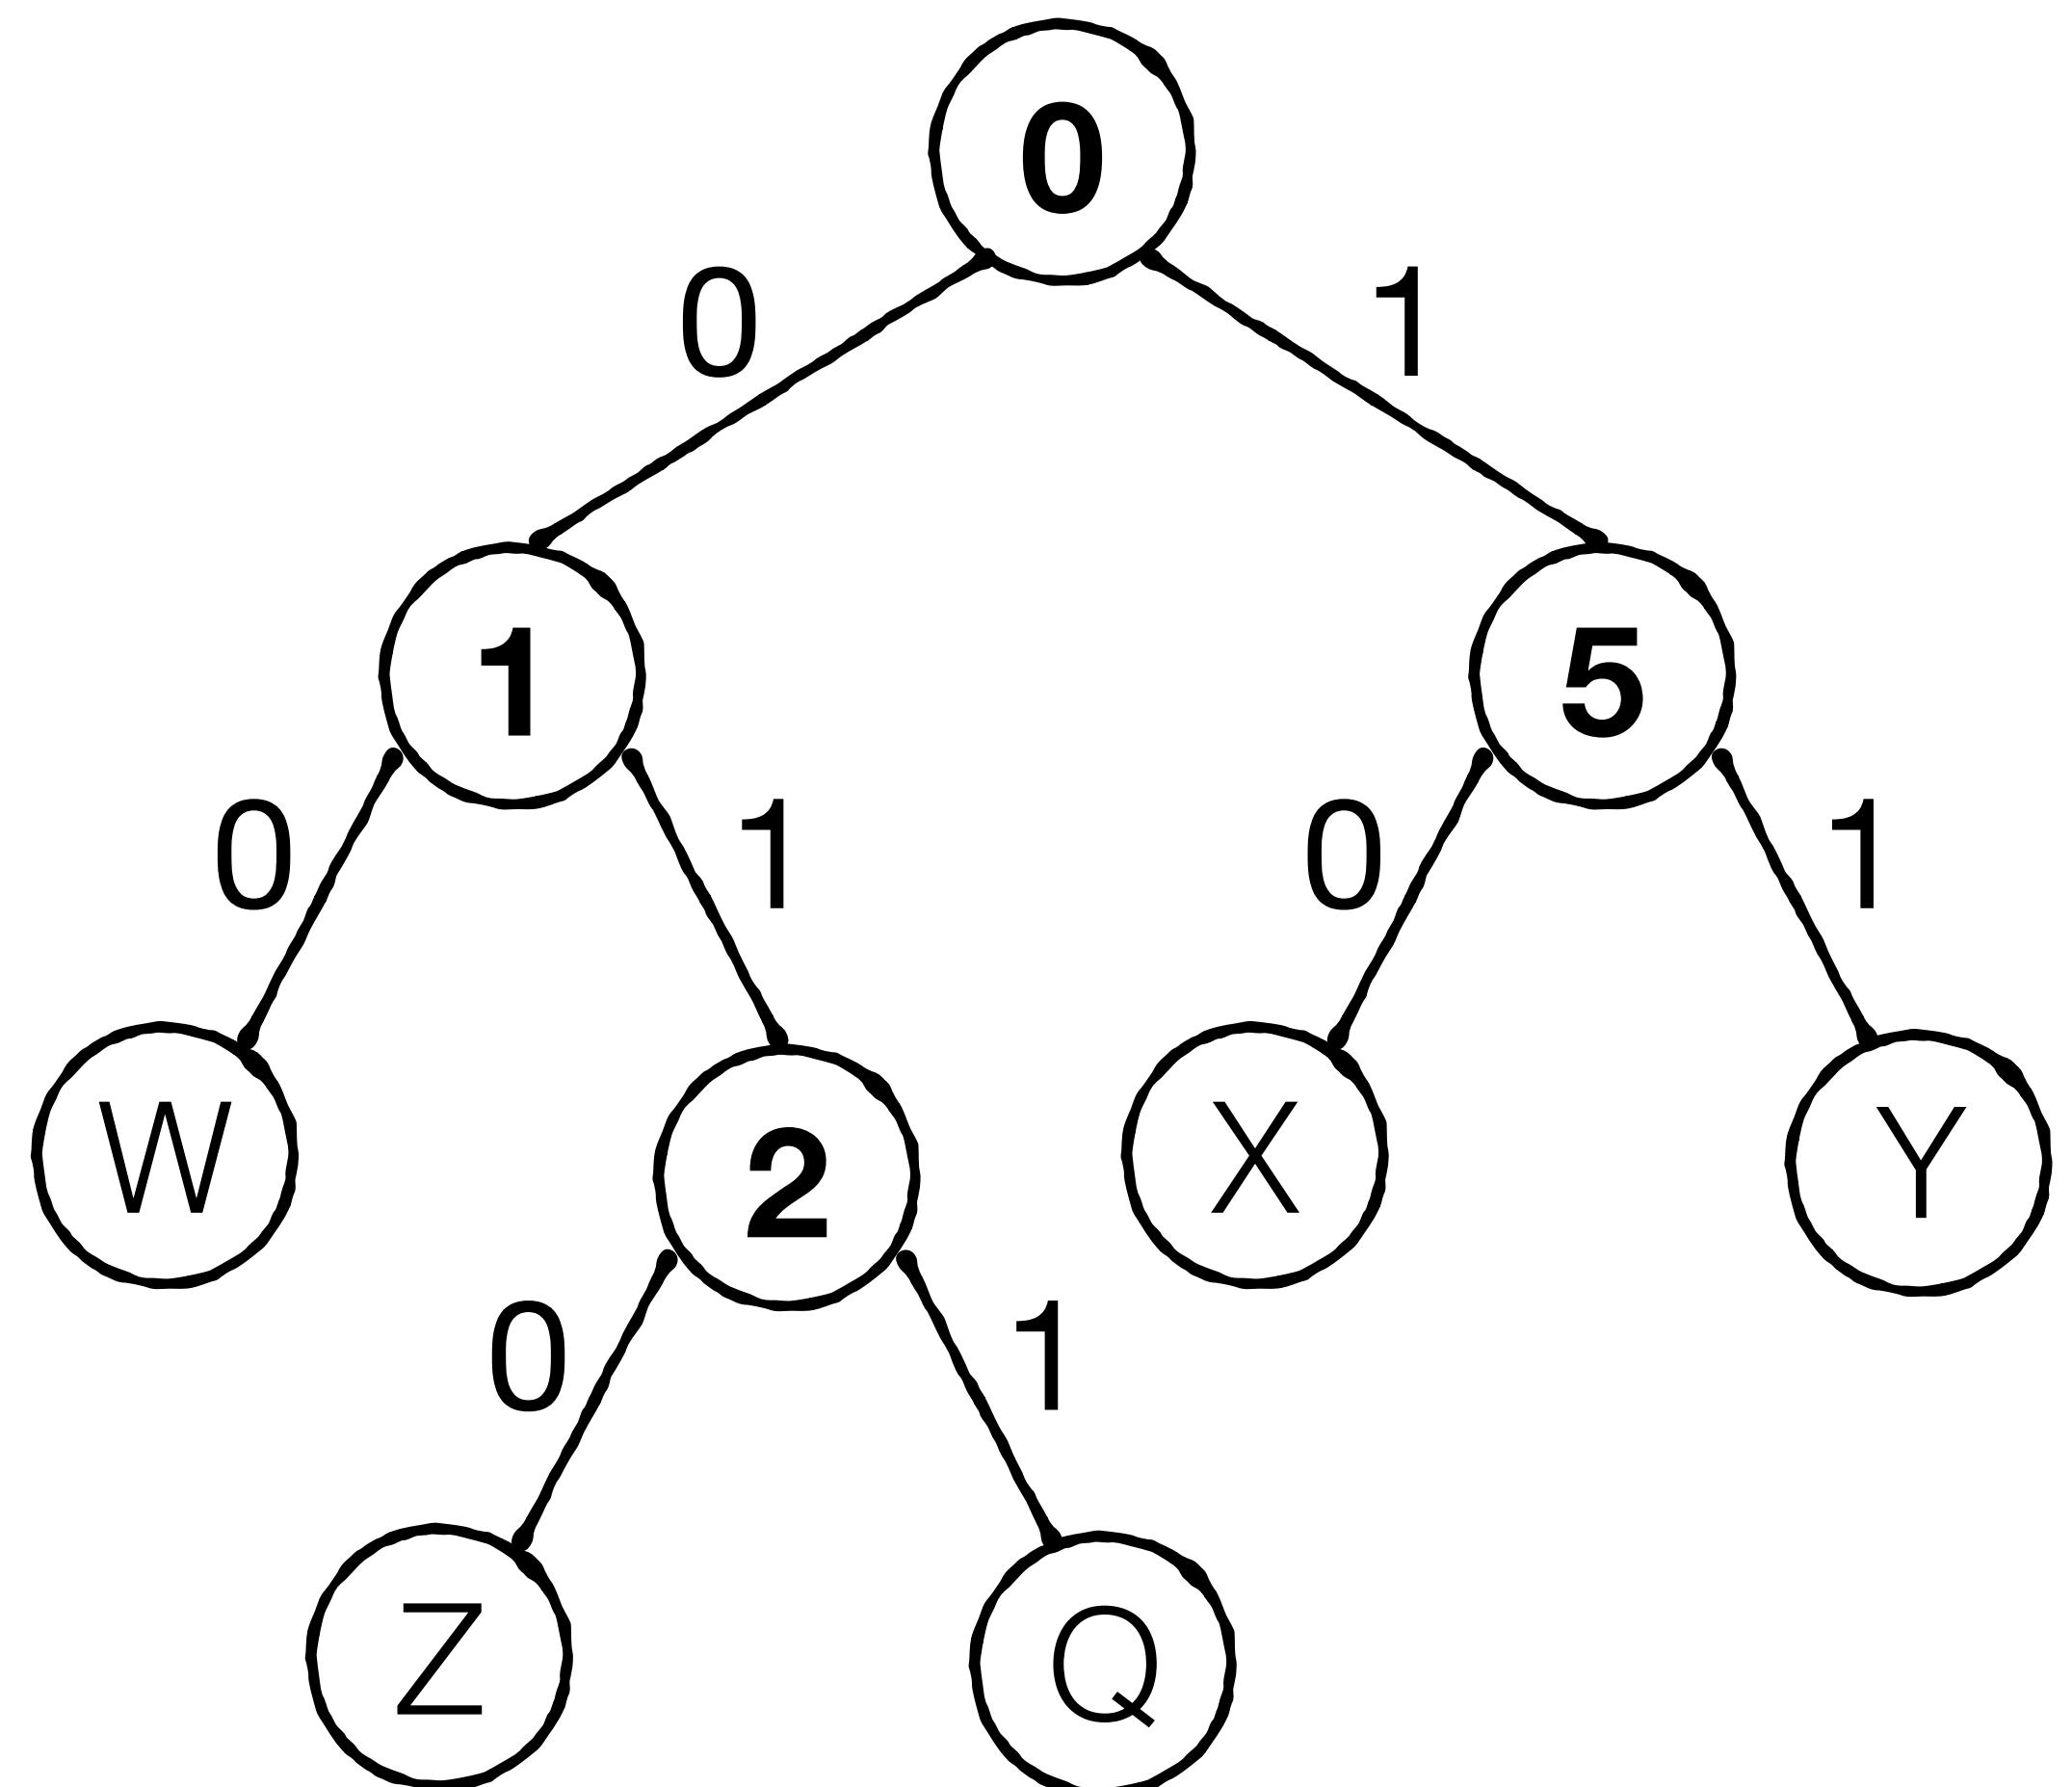
11 01 00

Differentiating bit indexes?

Depth-first, unary

0 10 110 111110

↑ ↑ ↑ ↑
0 1 2 5



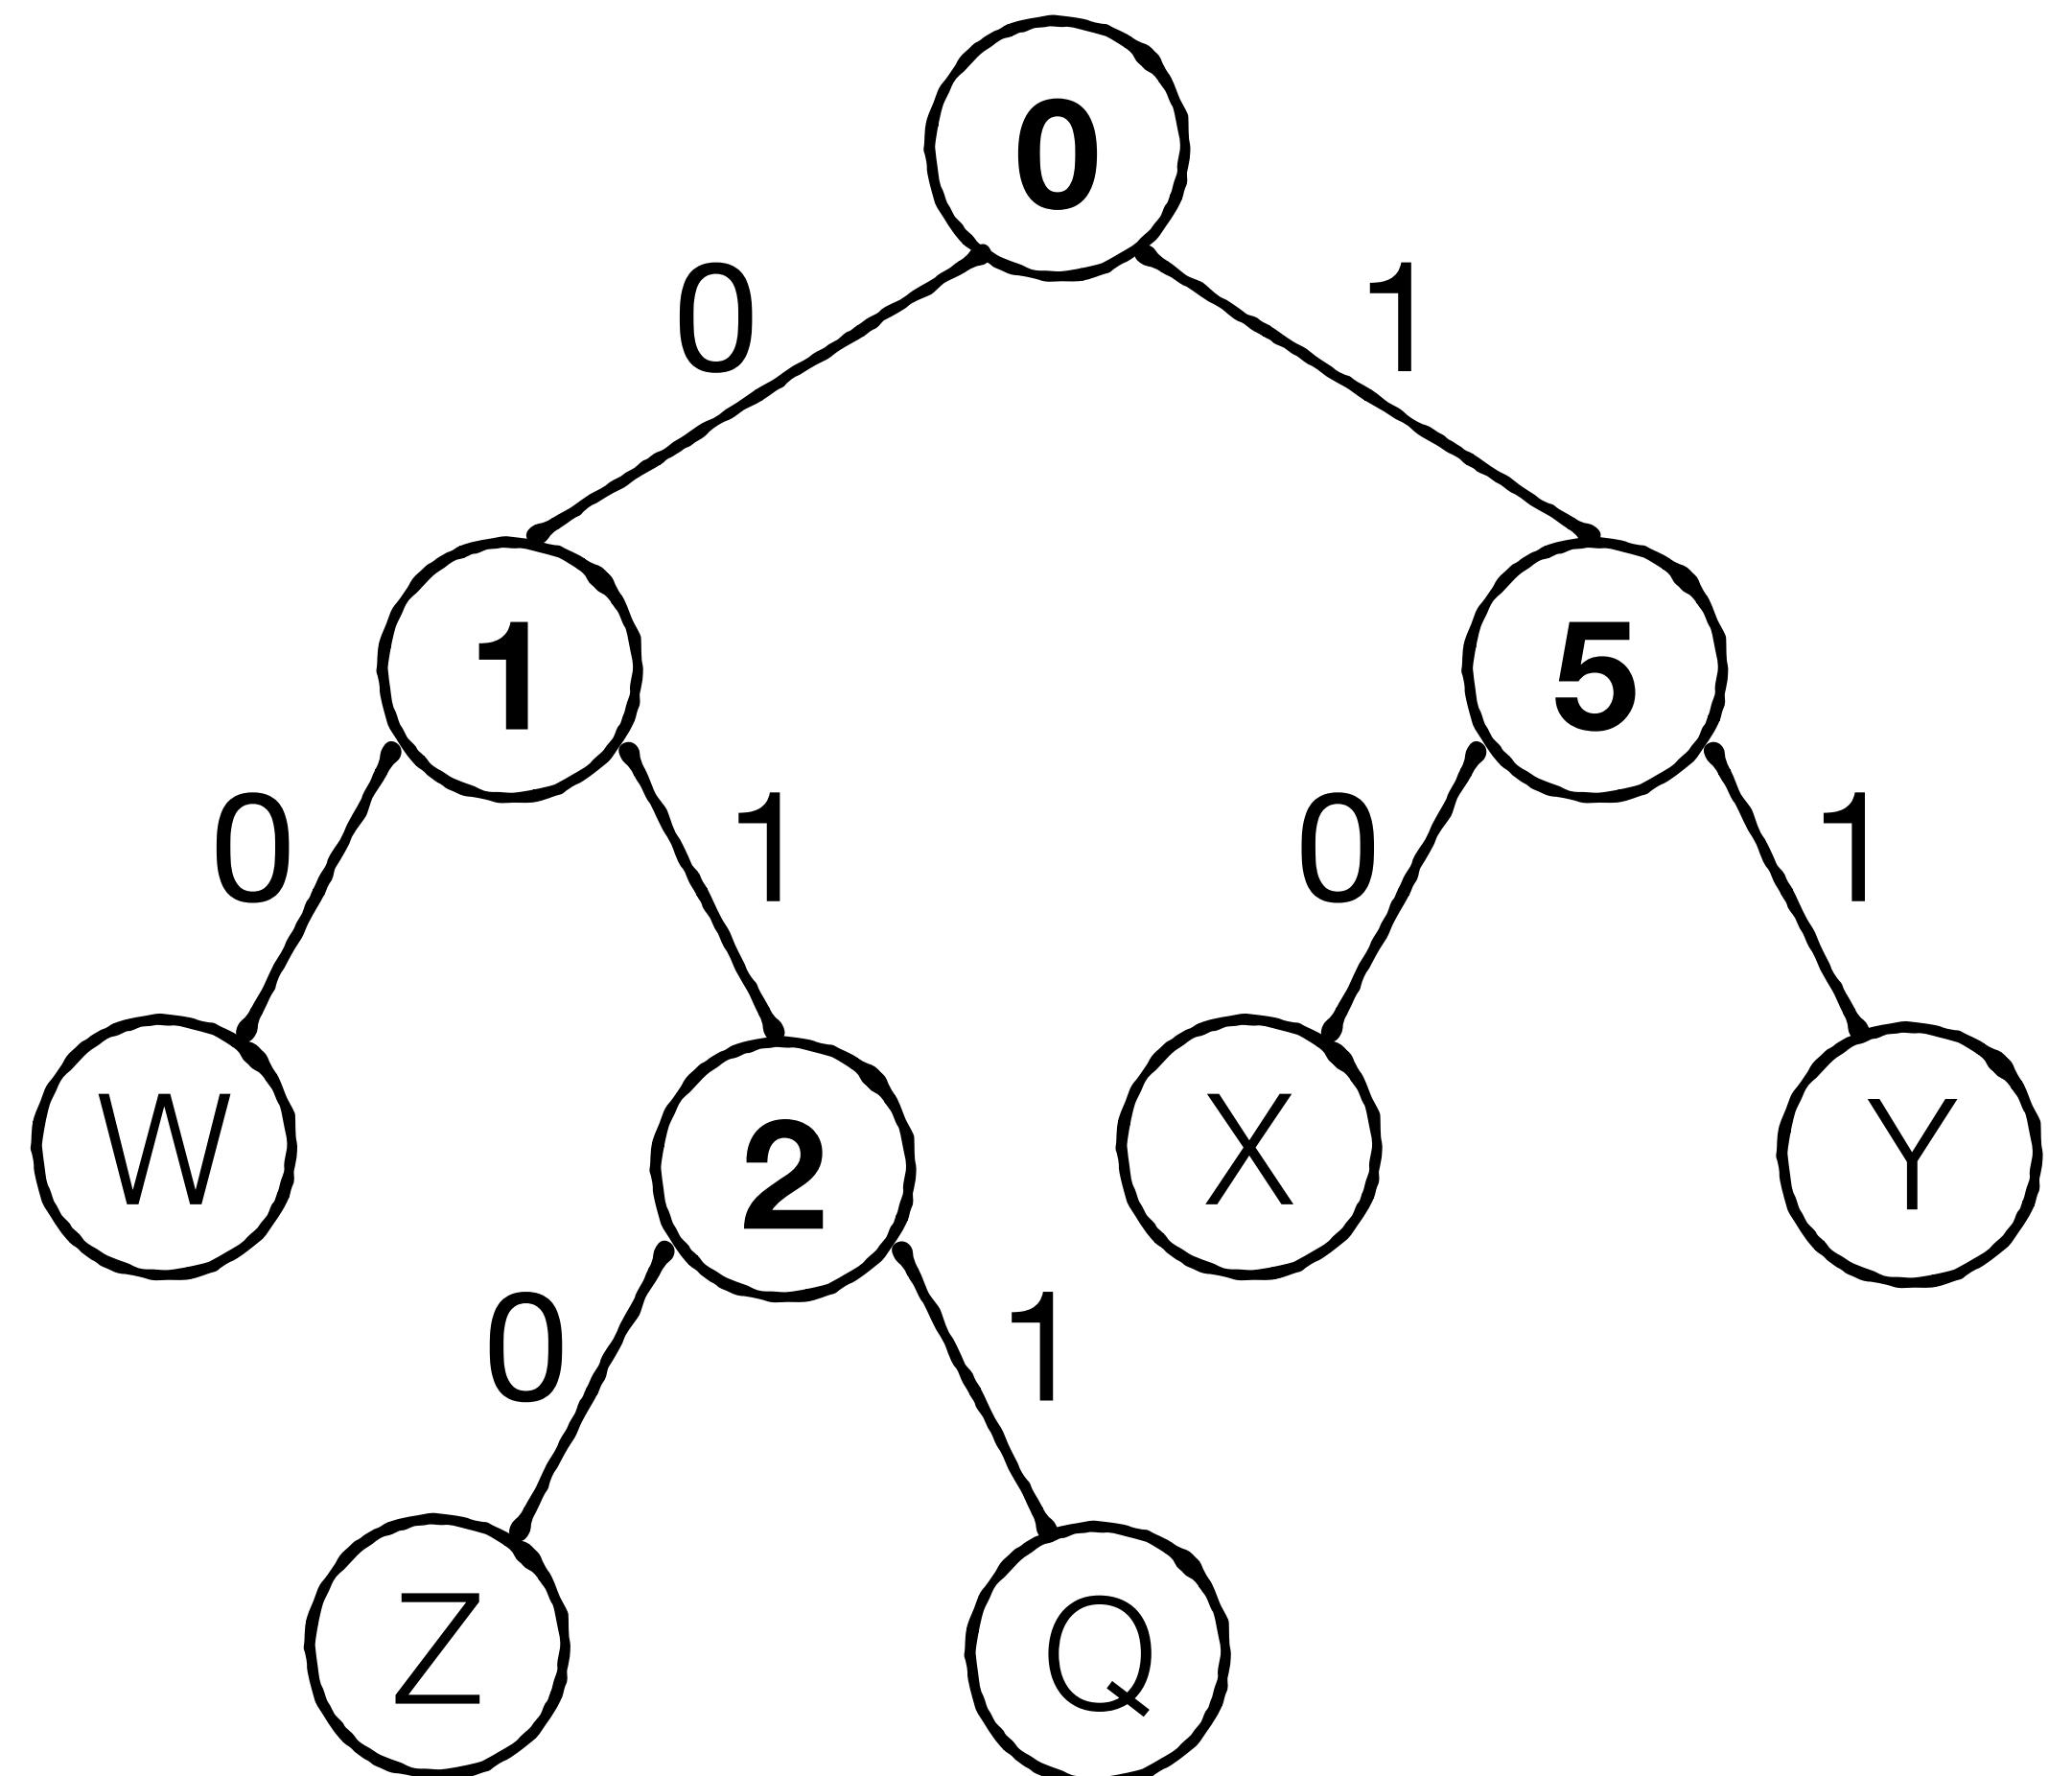
entries
111110

Topology

Differentiating bit indexes?

Depth-first, unary

0 10 110 111110



child's index is larger than parent's by at least 1

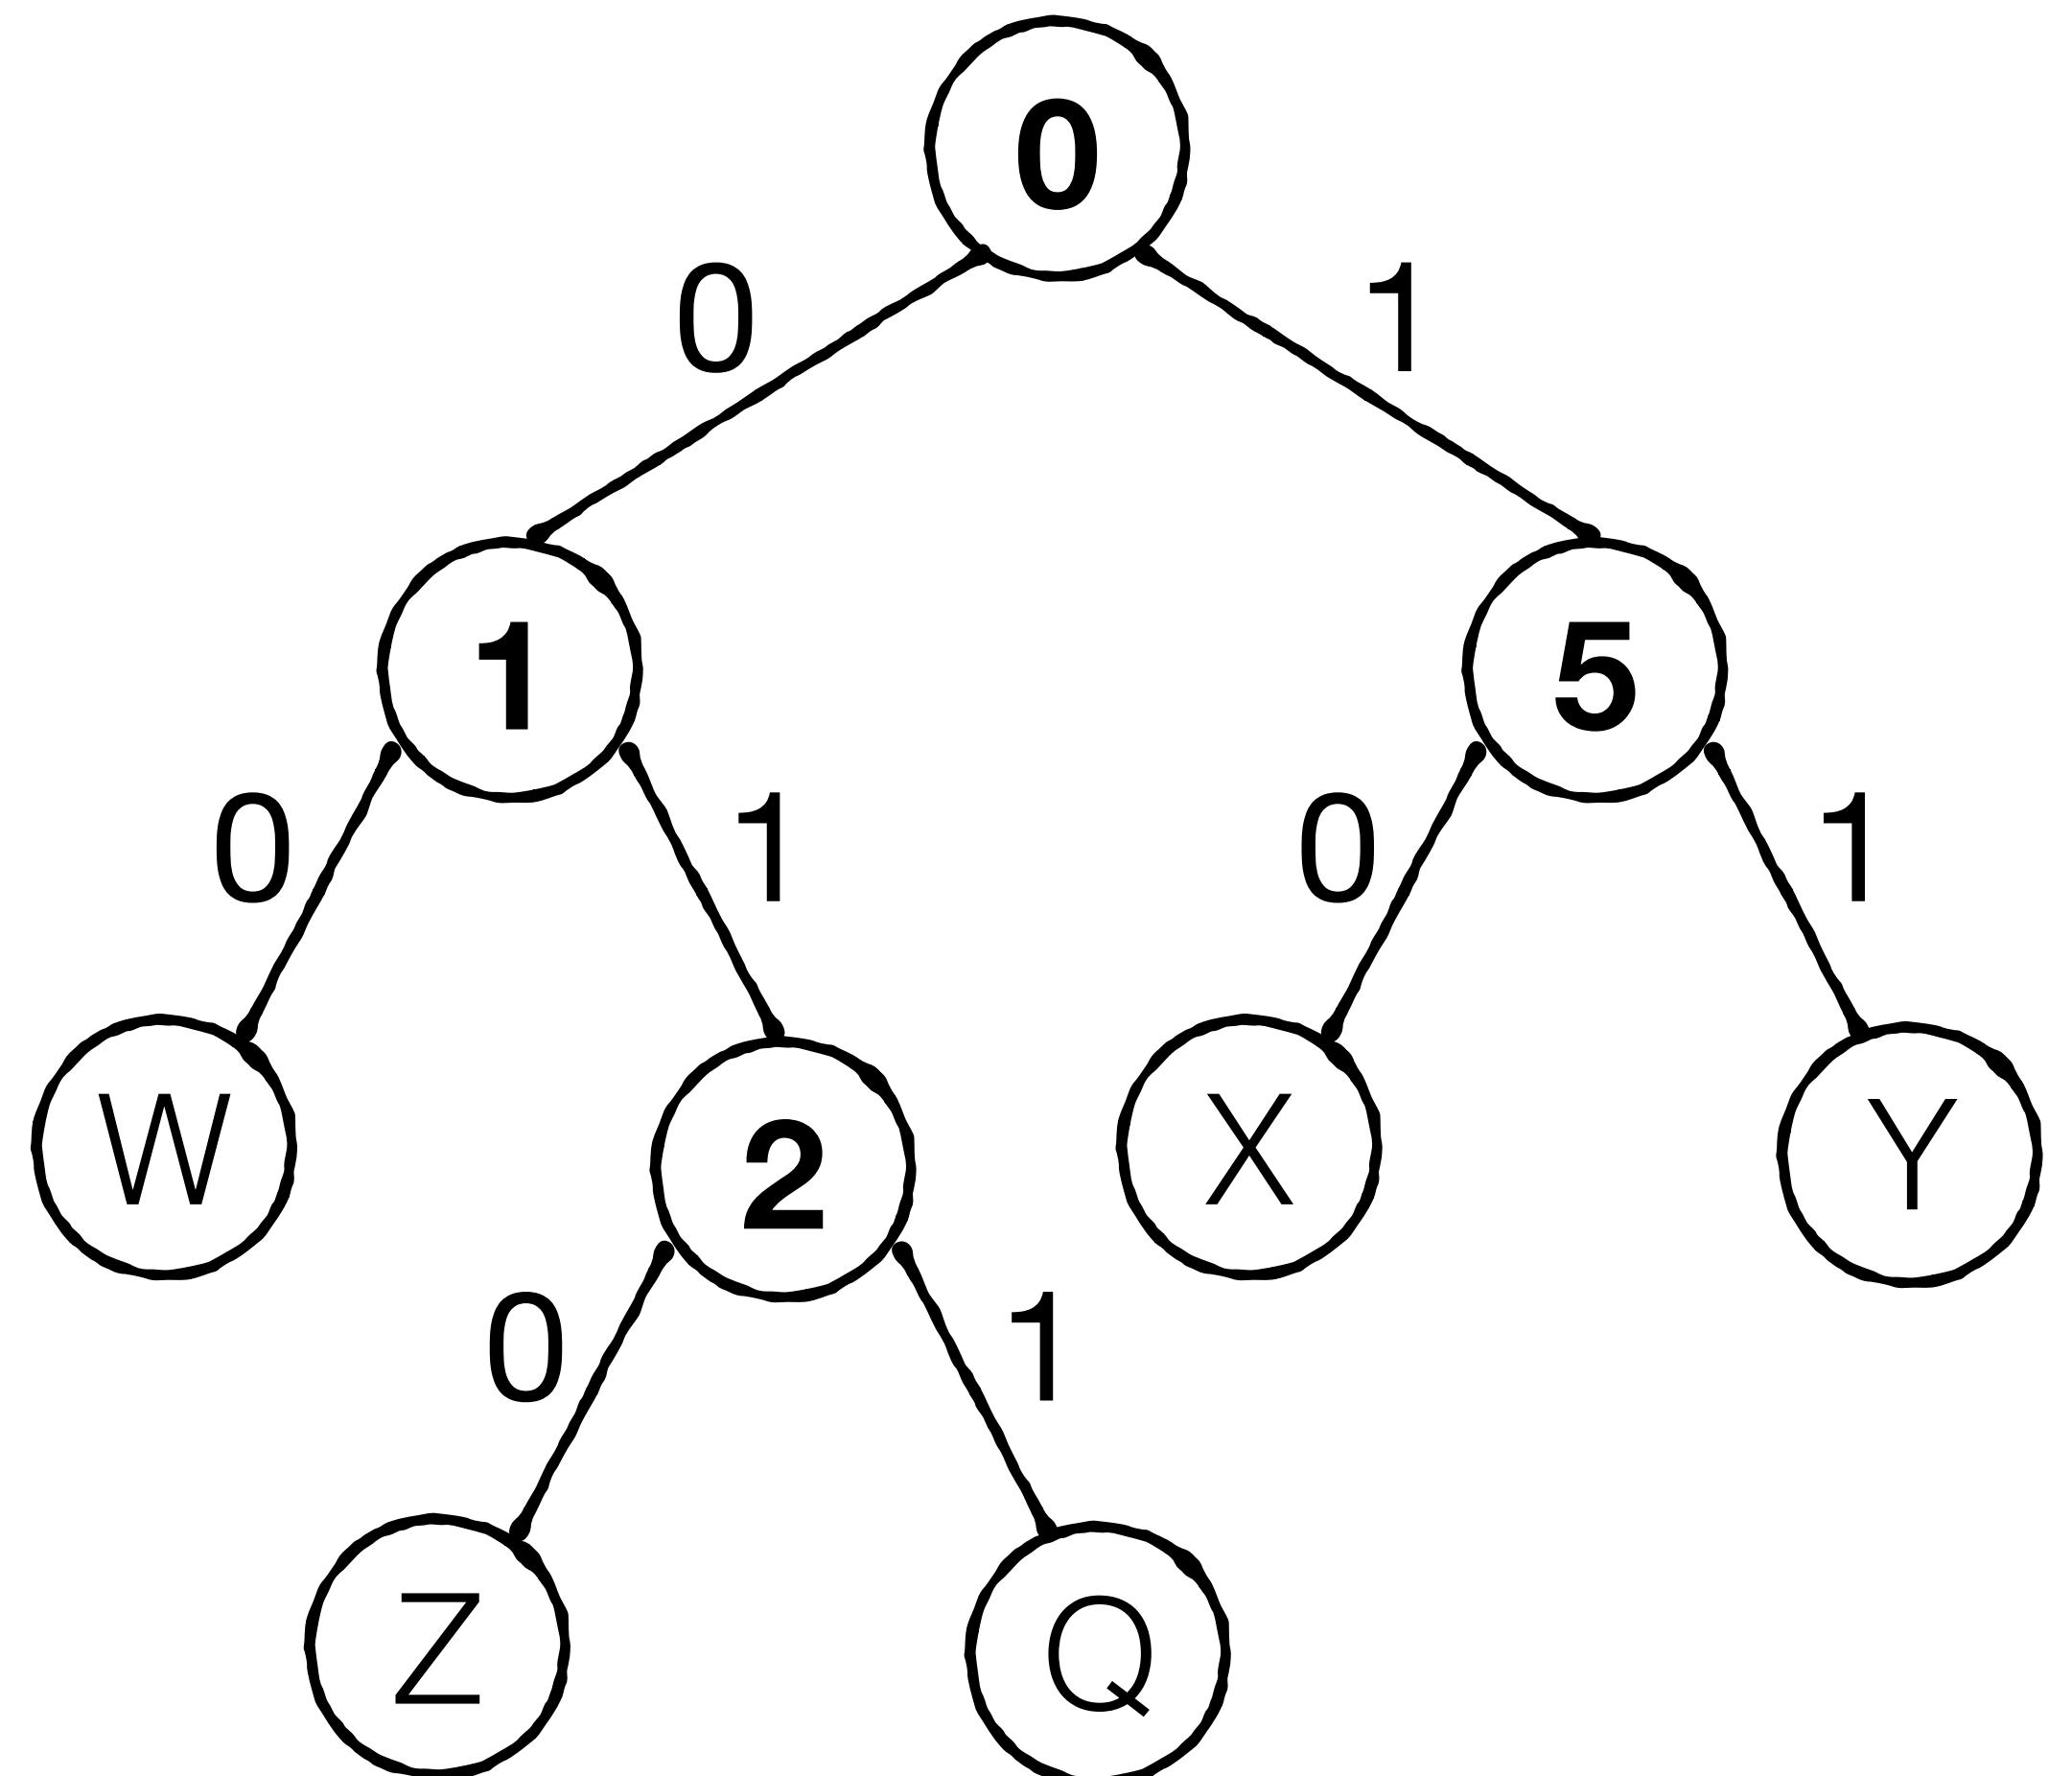
entries
111110

Topology
11 01 00

Differentiating bit indexes?

Depth-first, unary

0 10 110 111110



child's index is larger than parent's by at least 1

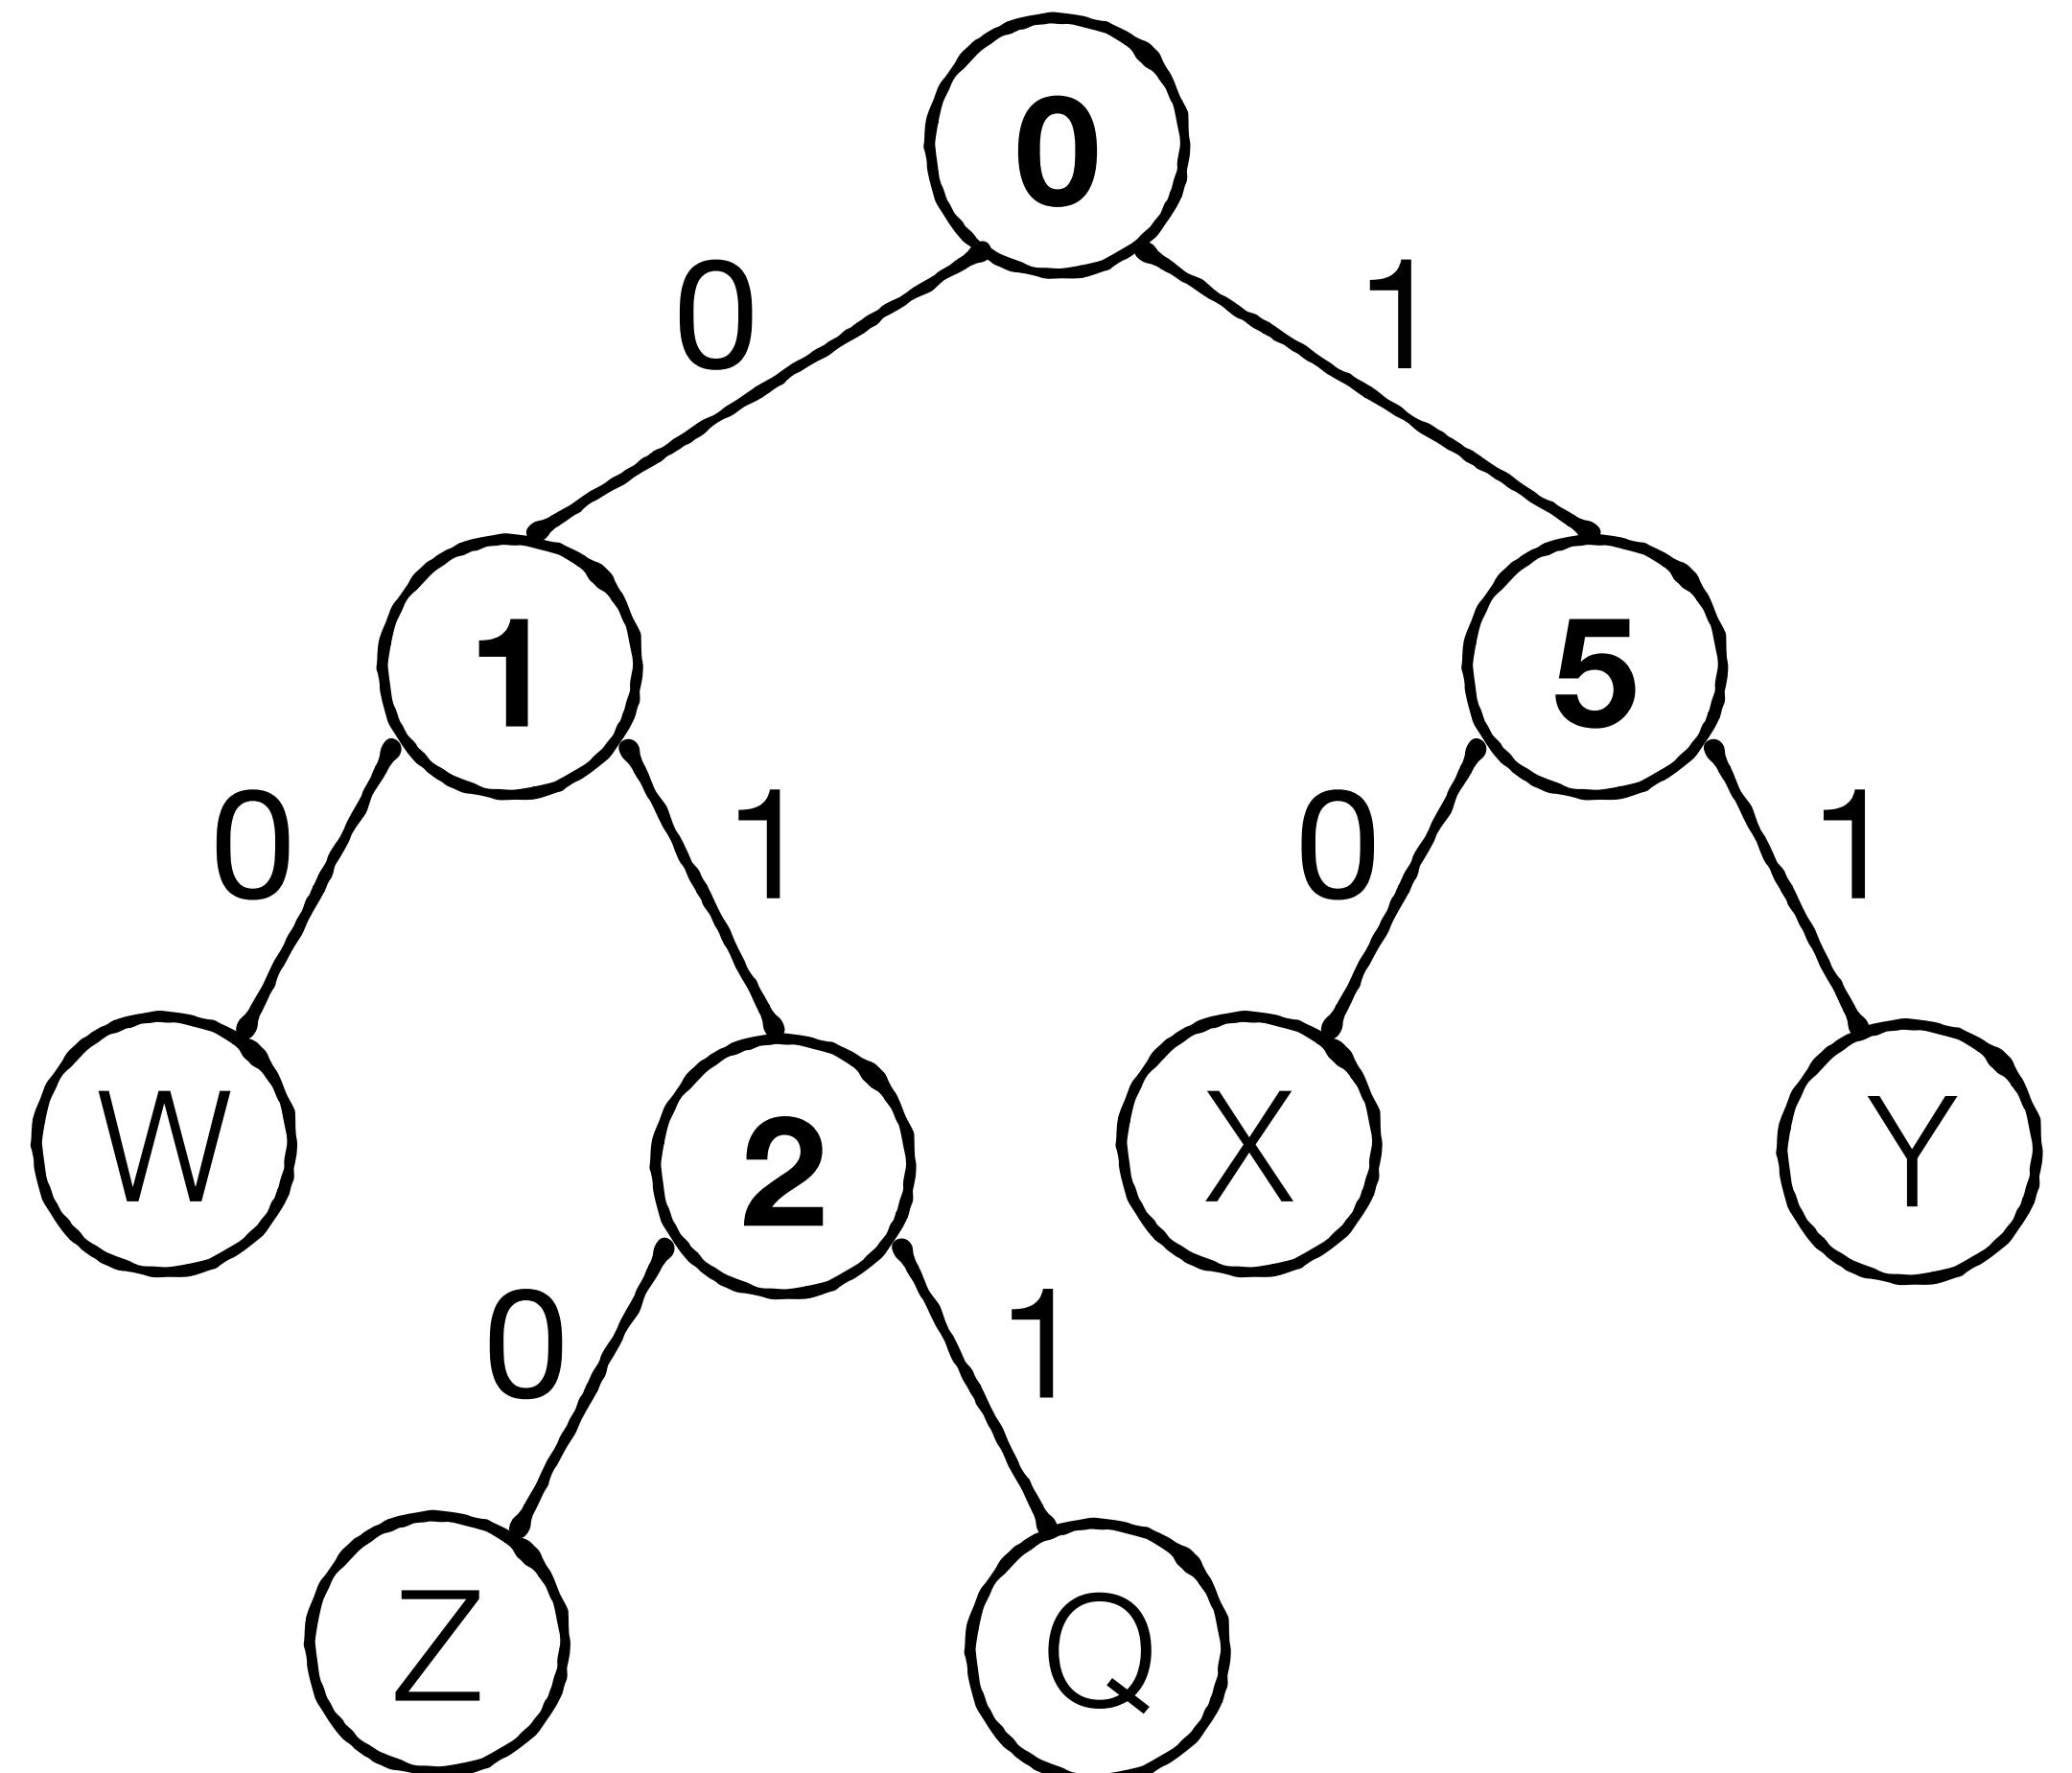
So let's encode index as $\Delta - 1$

entries
111110

Topology
11 01 00

Differentiating bit indexes?
Depth-first, unary

0 ~~4~~0 ~~4~~10 ~~4~~11110



child's index is larger than parent's by at least 1

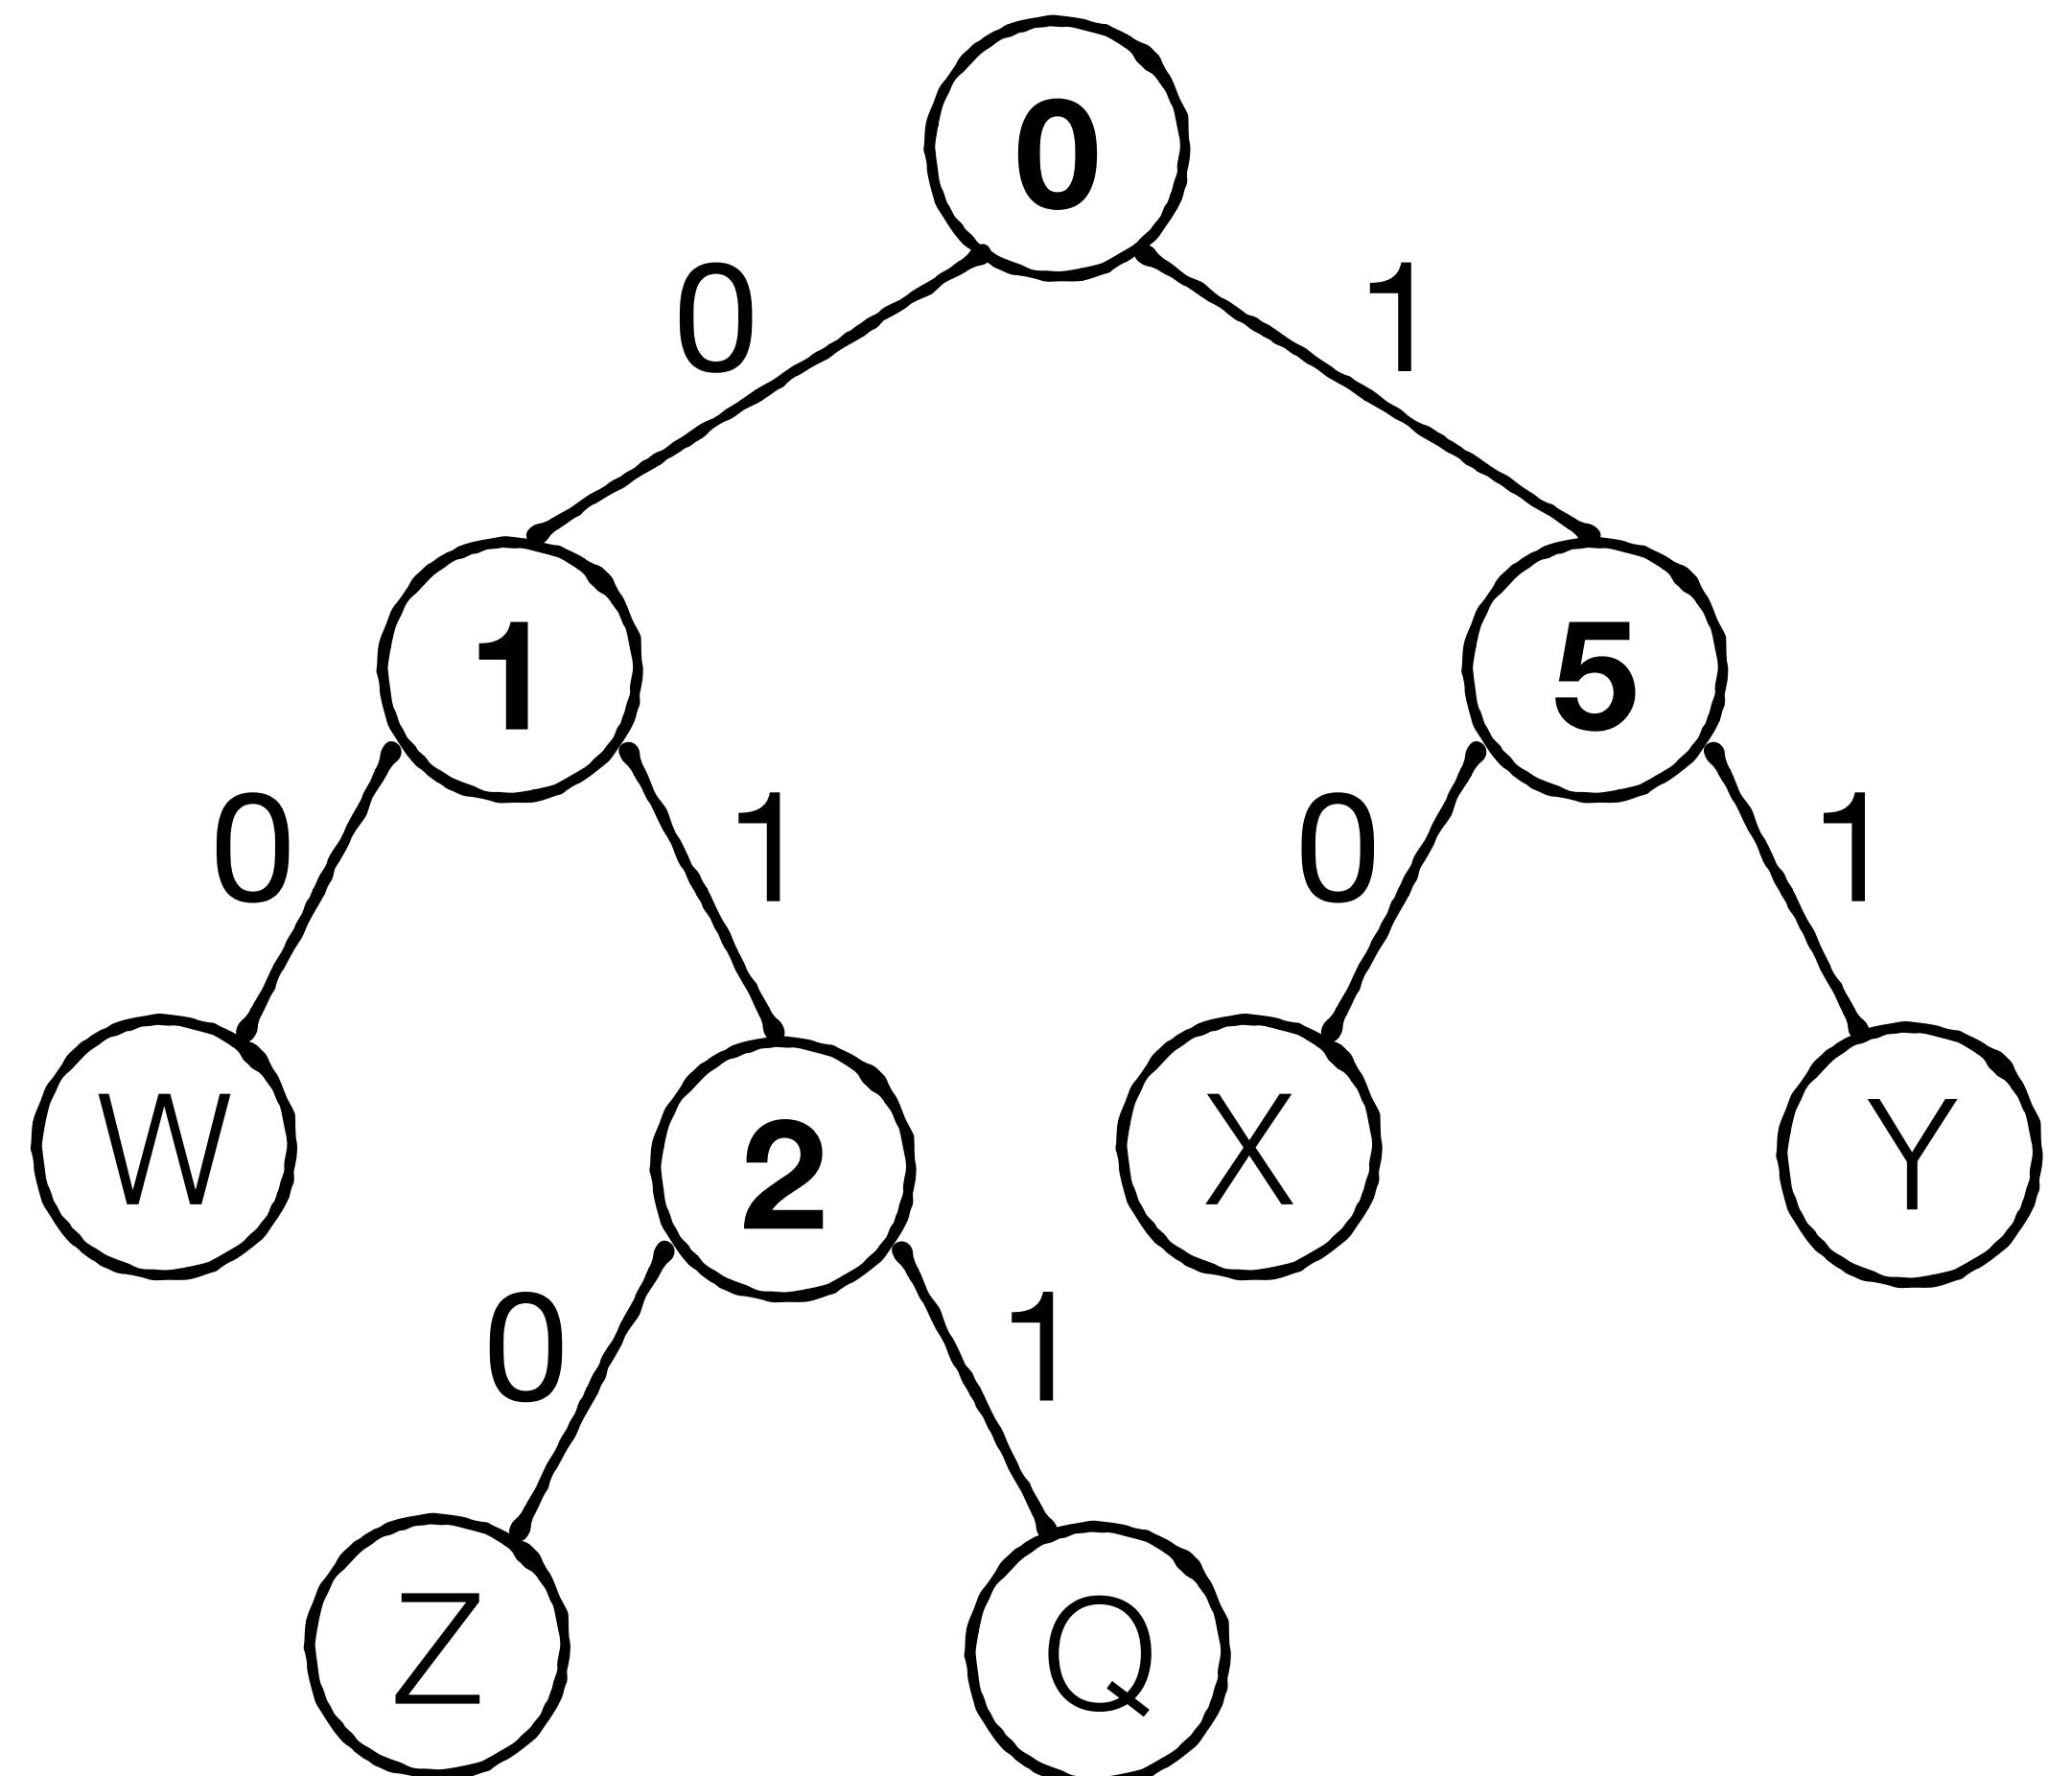
So let's encode index as $\Delta - 1$

entries
111110

Topology
11 01 00

Differentiating bit indexes?
Depth-first, unary

0 0 0 11110



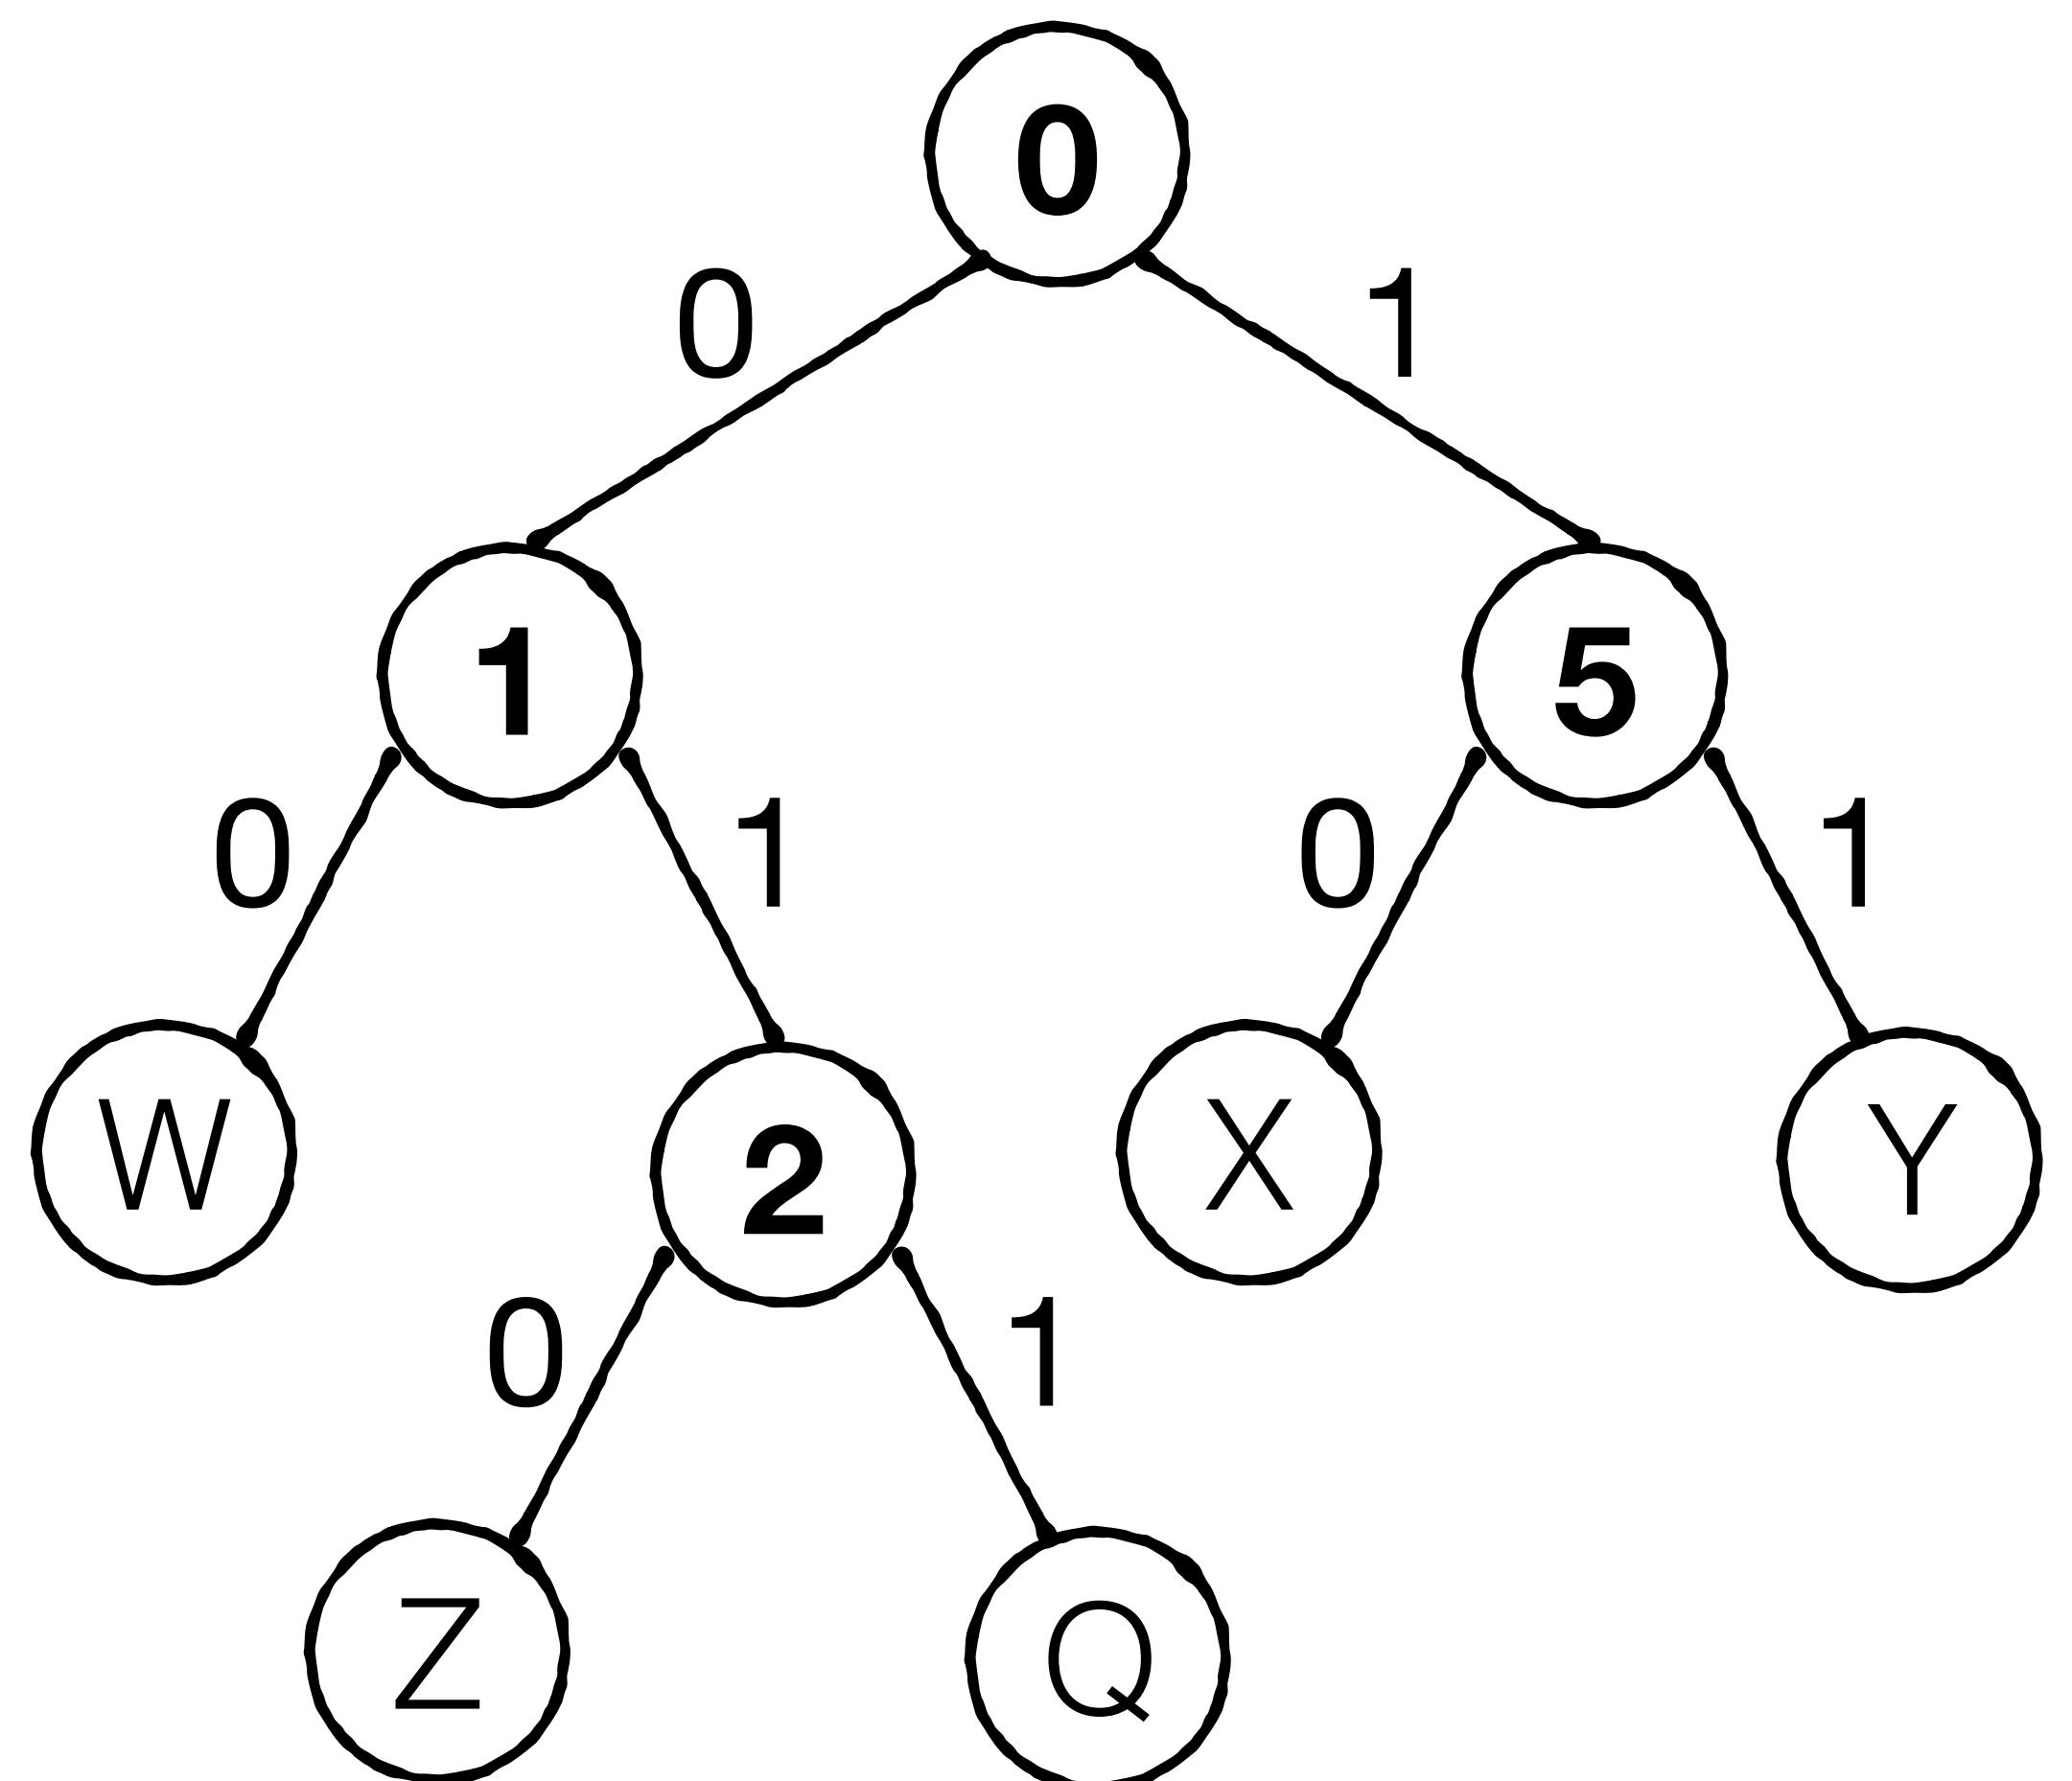
child's index is larger than
parent's by at least 1

So let's encode index as $\Delta - 1$

entries
111110

Topology
11 01 00

Indices
0 0 0 11110

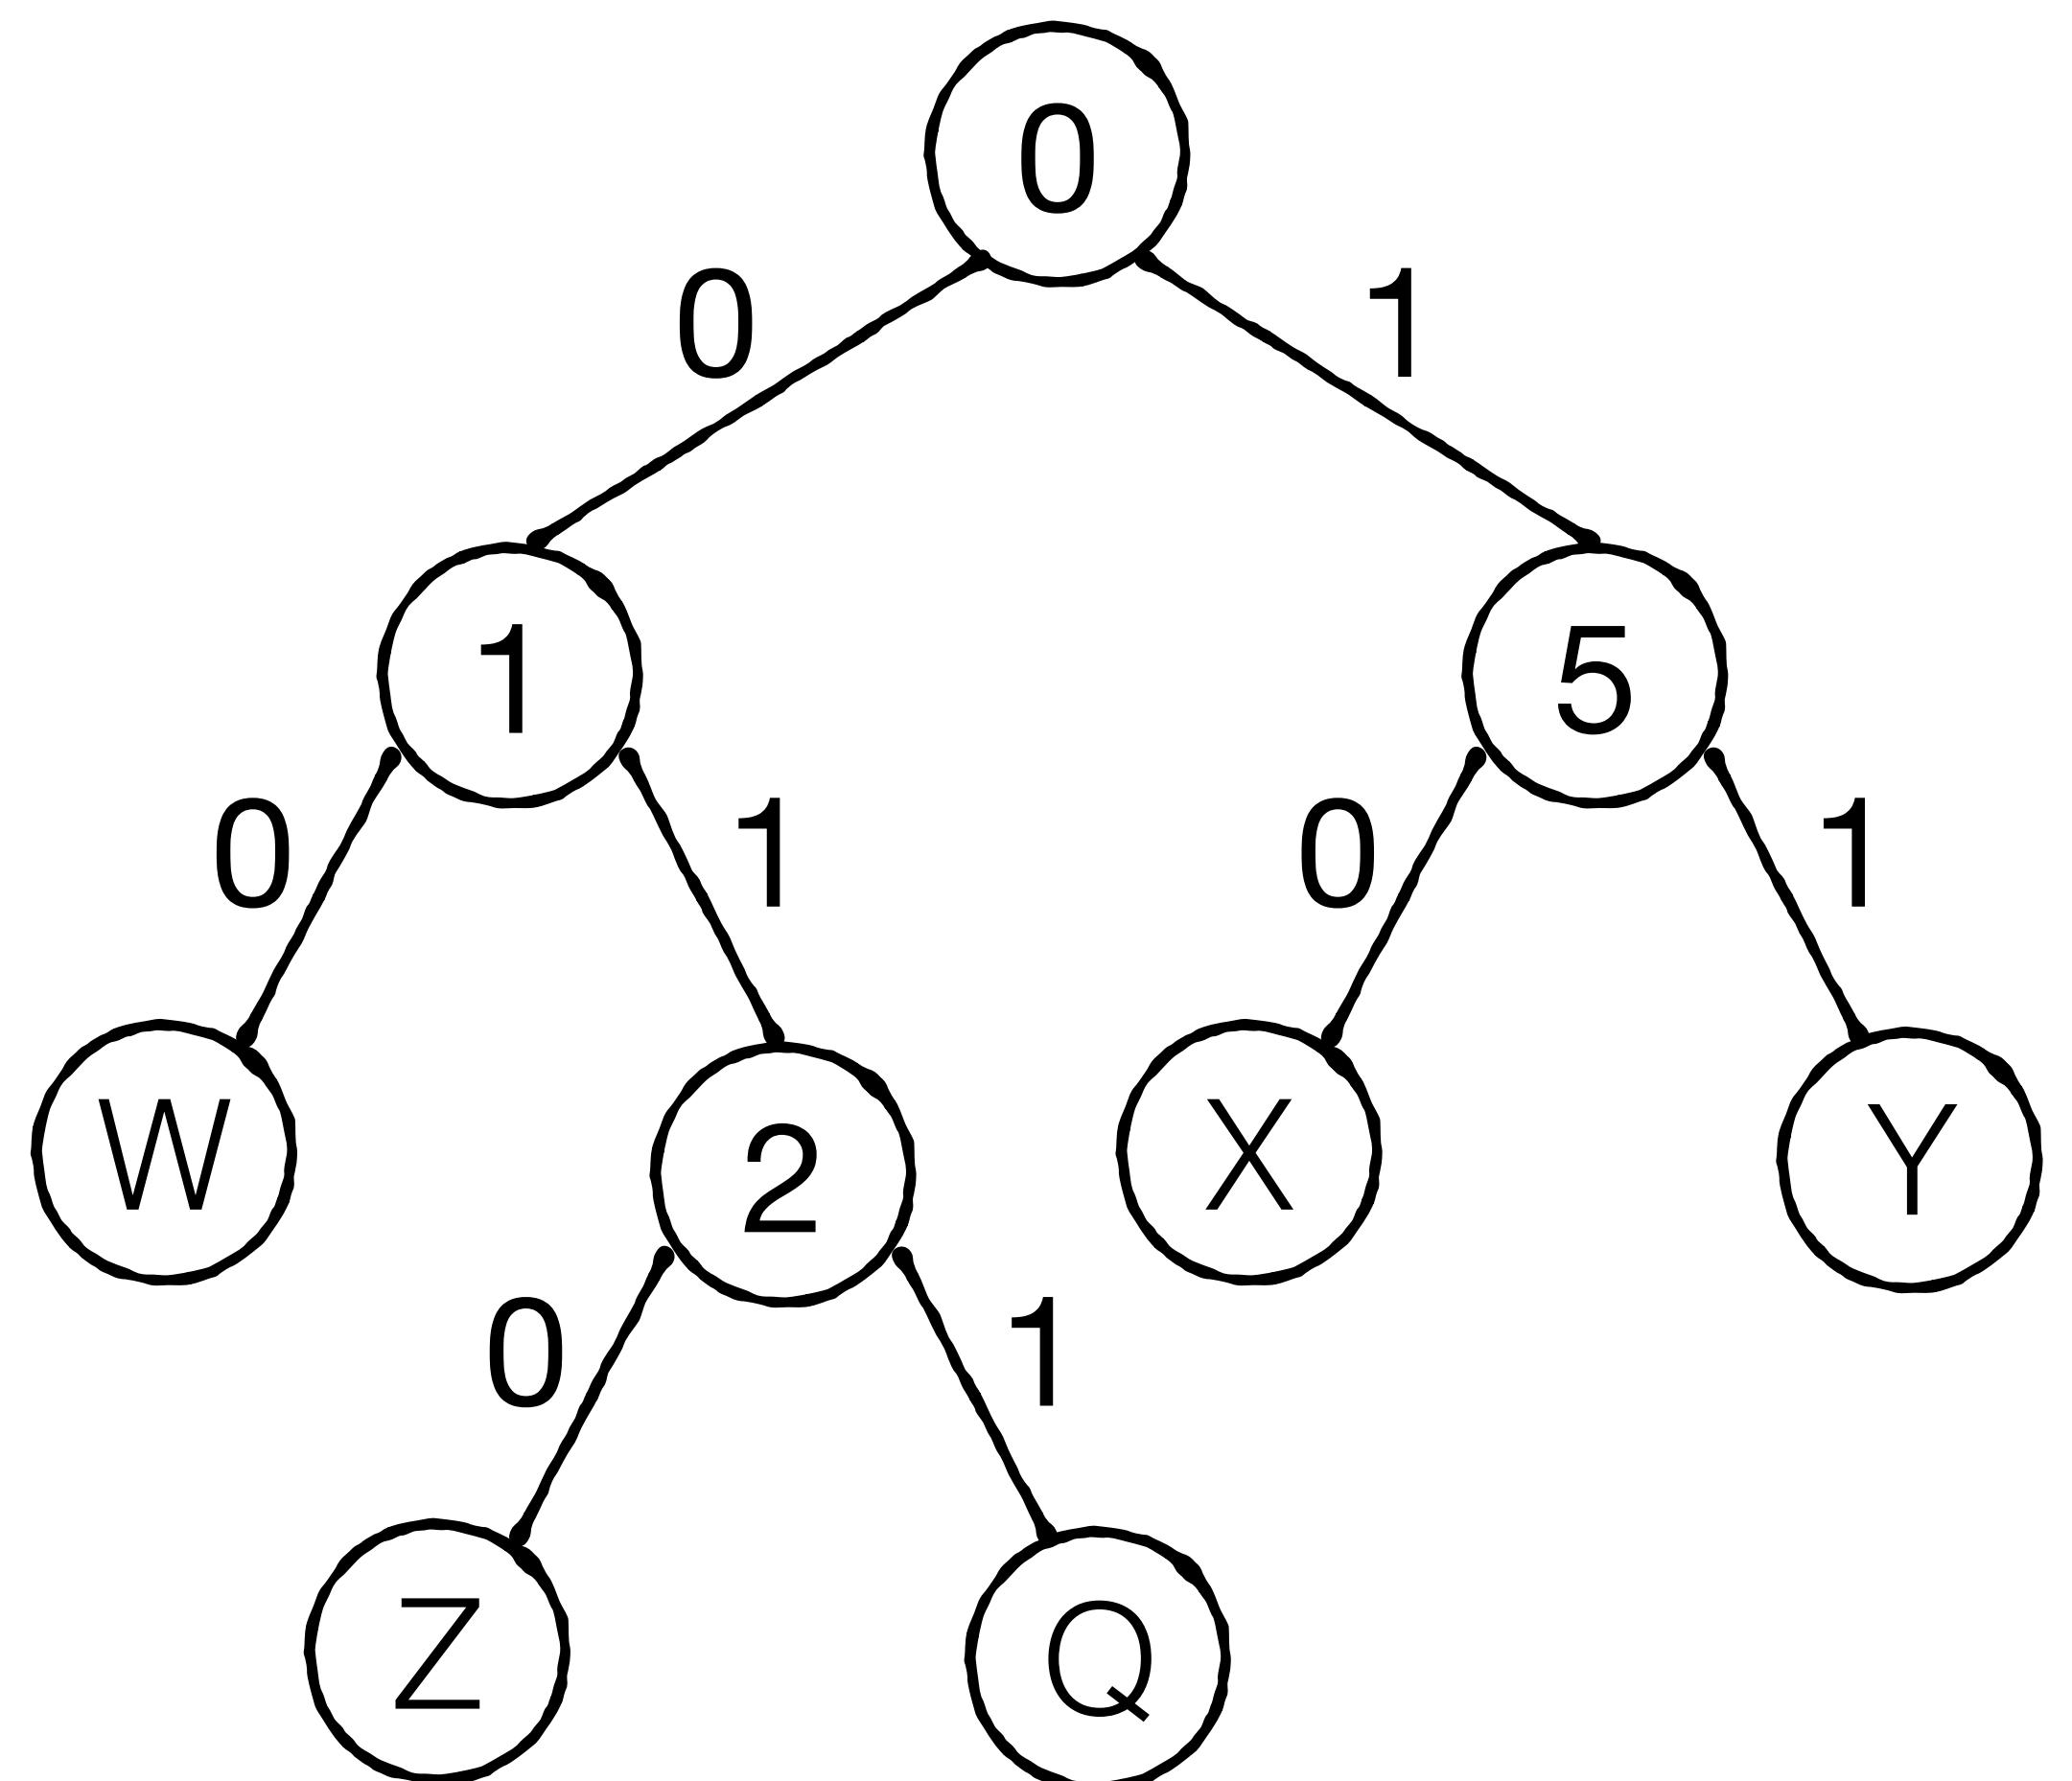


entries
111110

Topology
11 01 00

Indices
0 0 0 11110

Pointers
W Z Q X Y

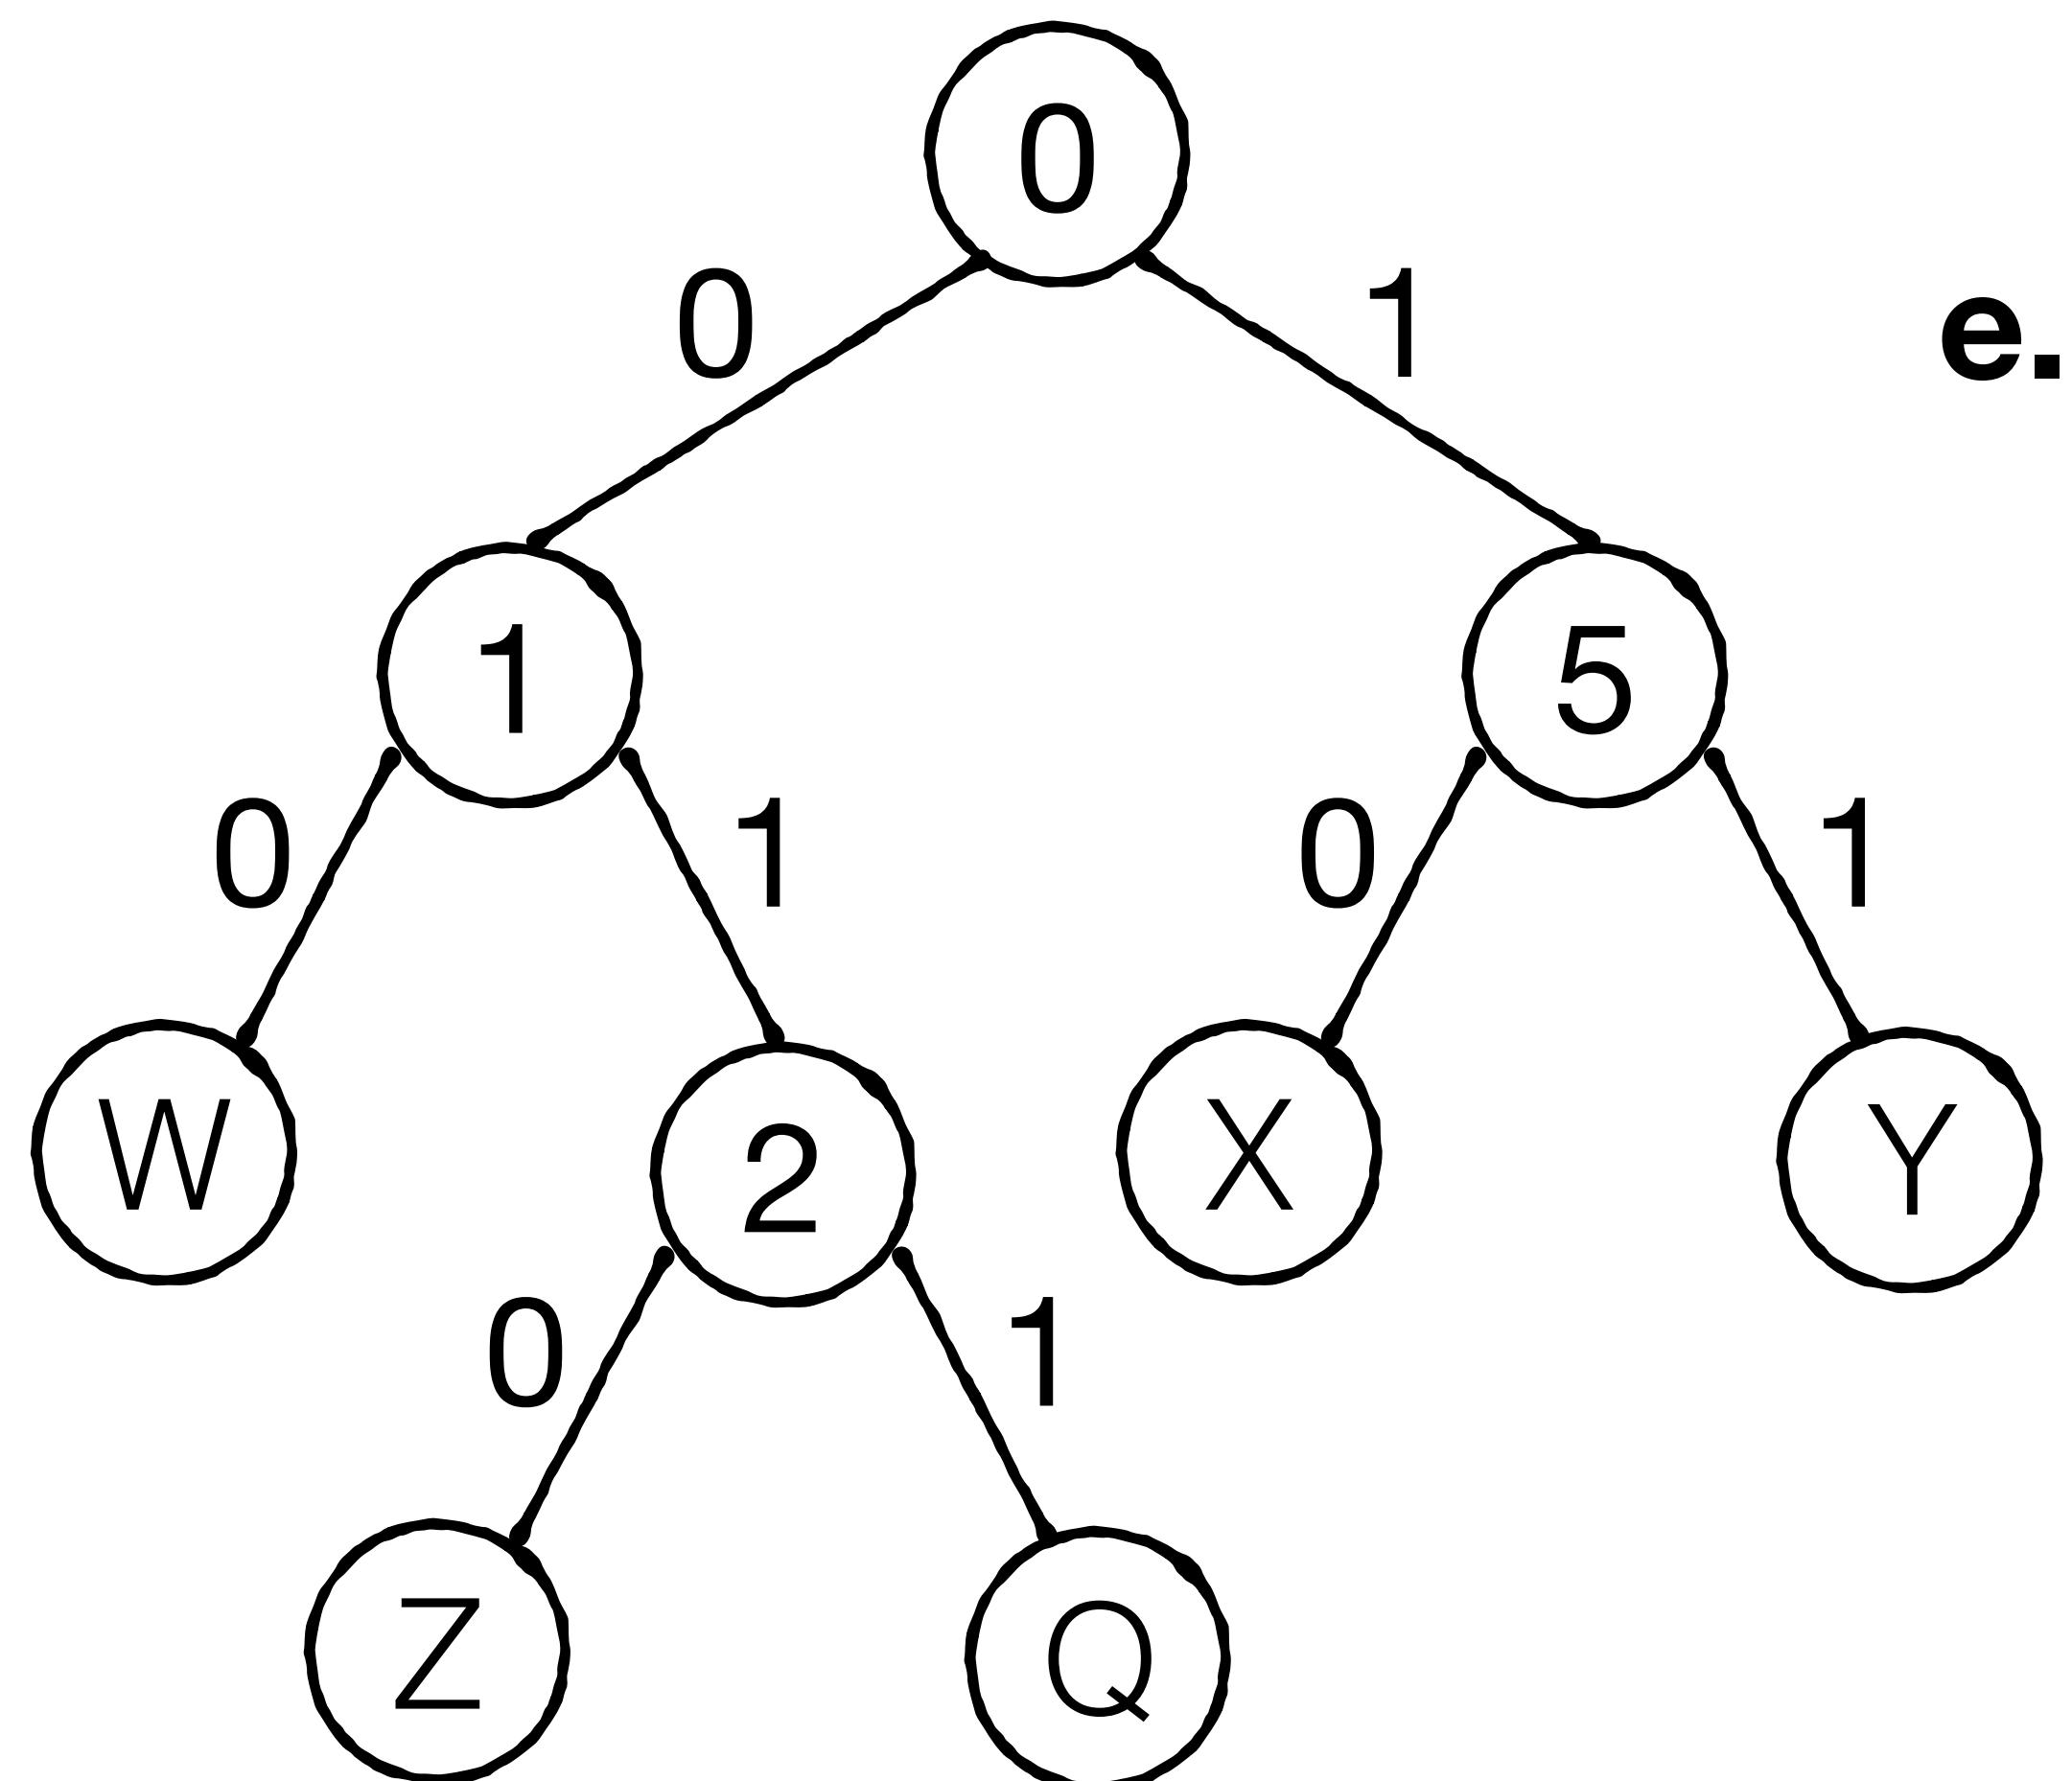


entries
111110

Topology
11 01 00

Indices
0 0 0 11110

Pointers
W Z Q X Y



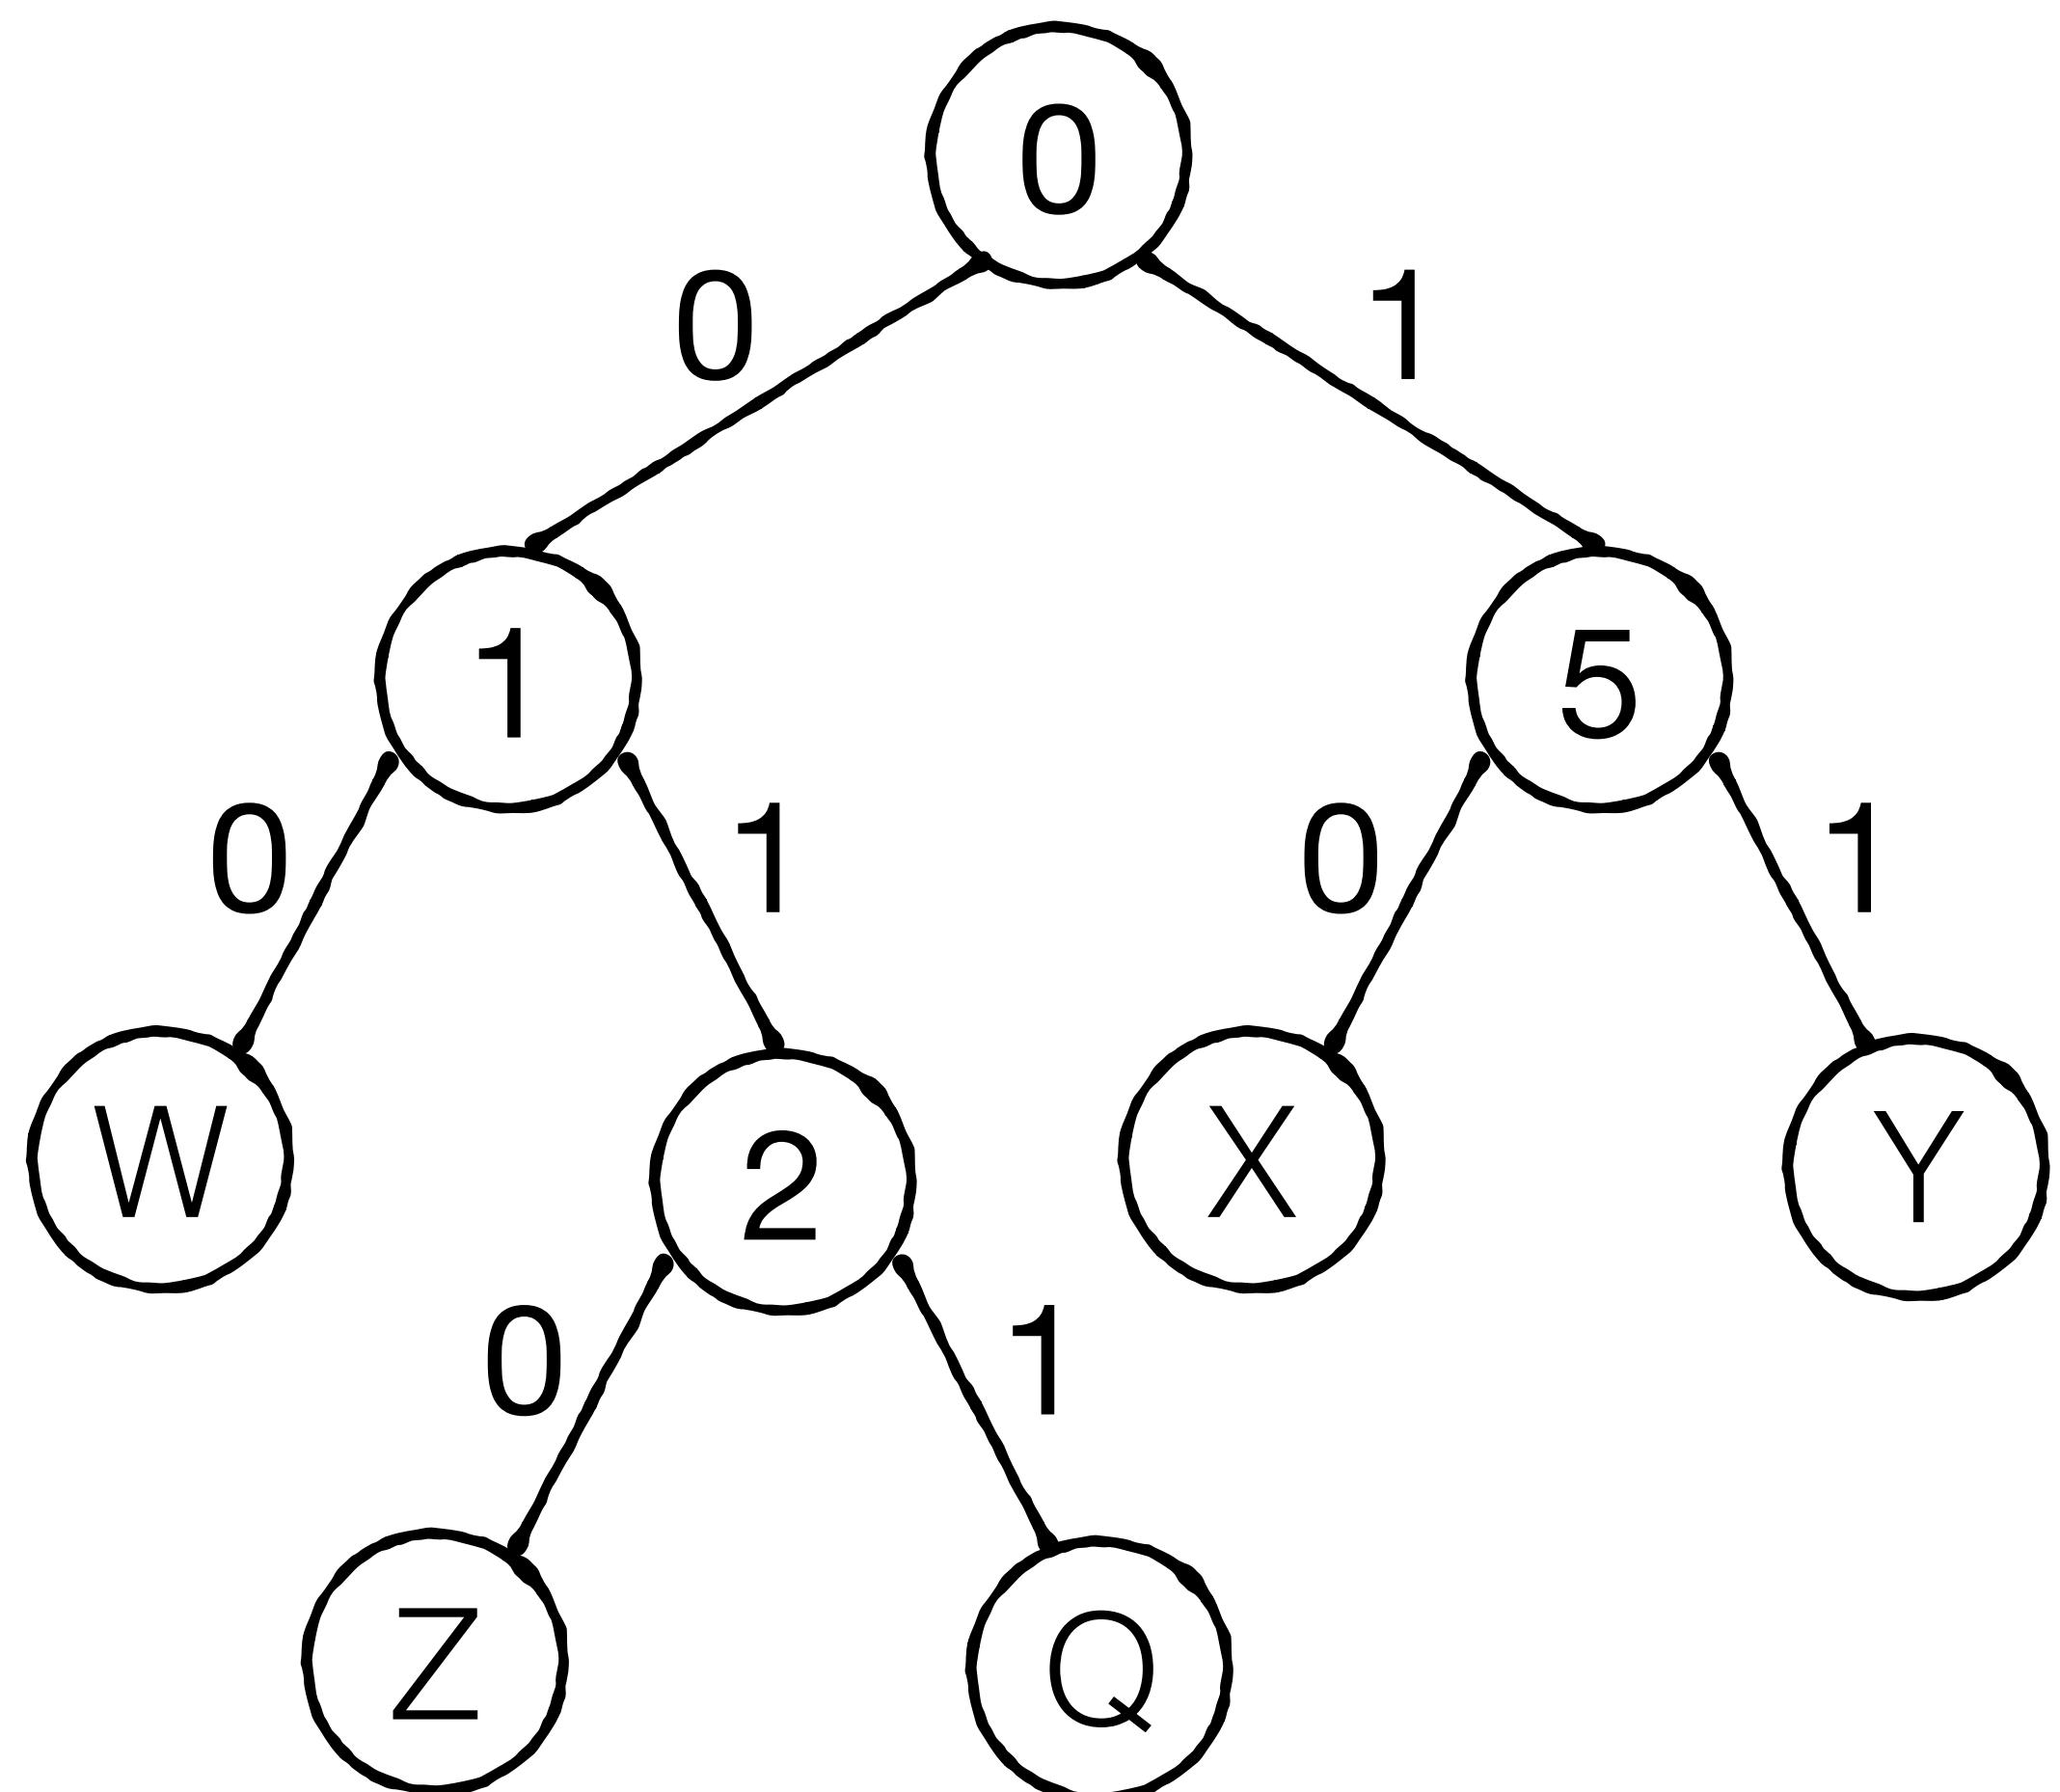
e.g., **get(Q)** where **FP(Q) = 0 1 1 0 0 1**

entries
111110

Topology
11 01 00

Indices
0 0 0 11110

Pointers
W Z Q X Y



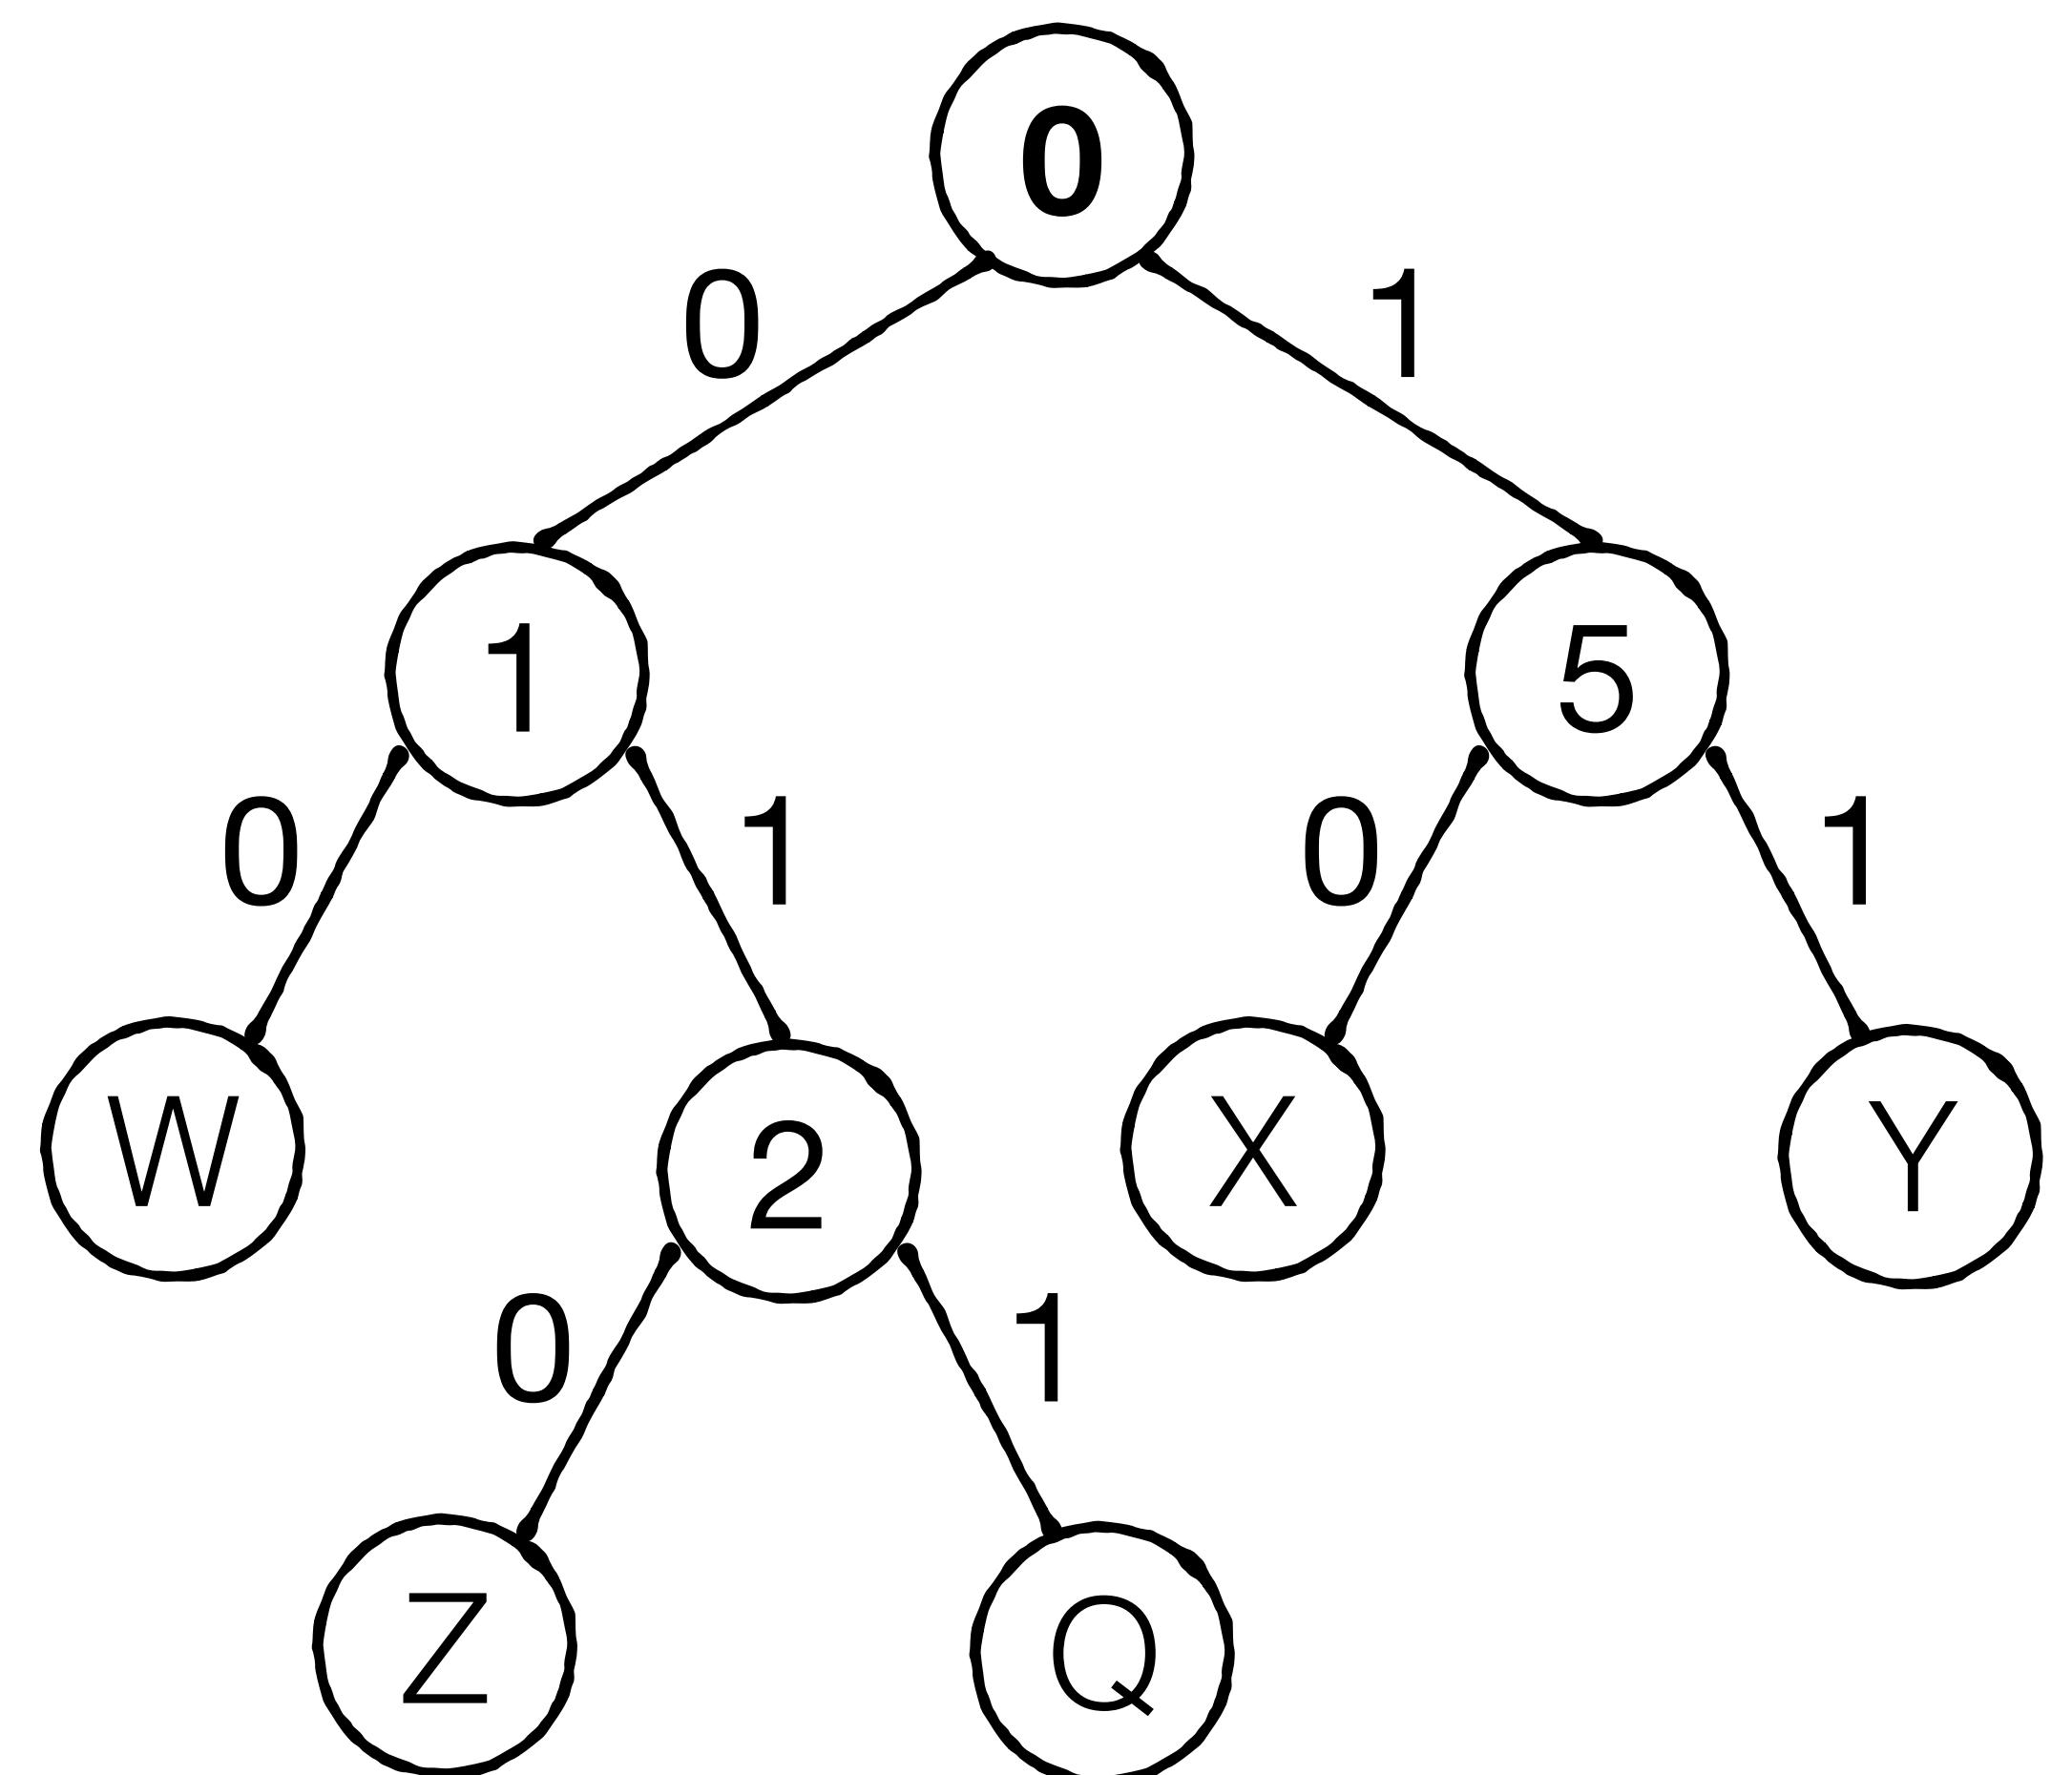
$FP(Q) = 011001$

entries

111110

Topology

11 01 00



Indices

0 0 0 11110

Pointers

W Z Q X Y

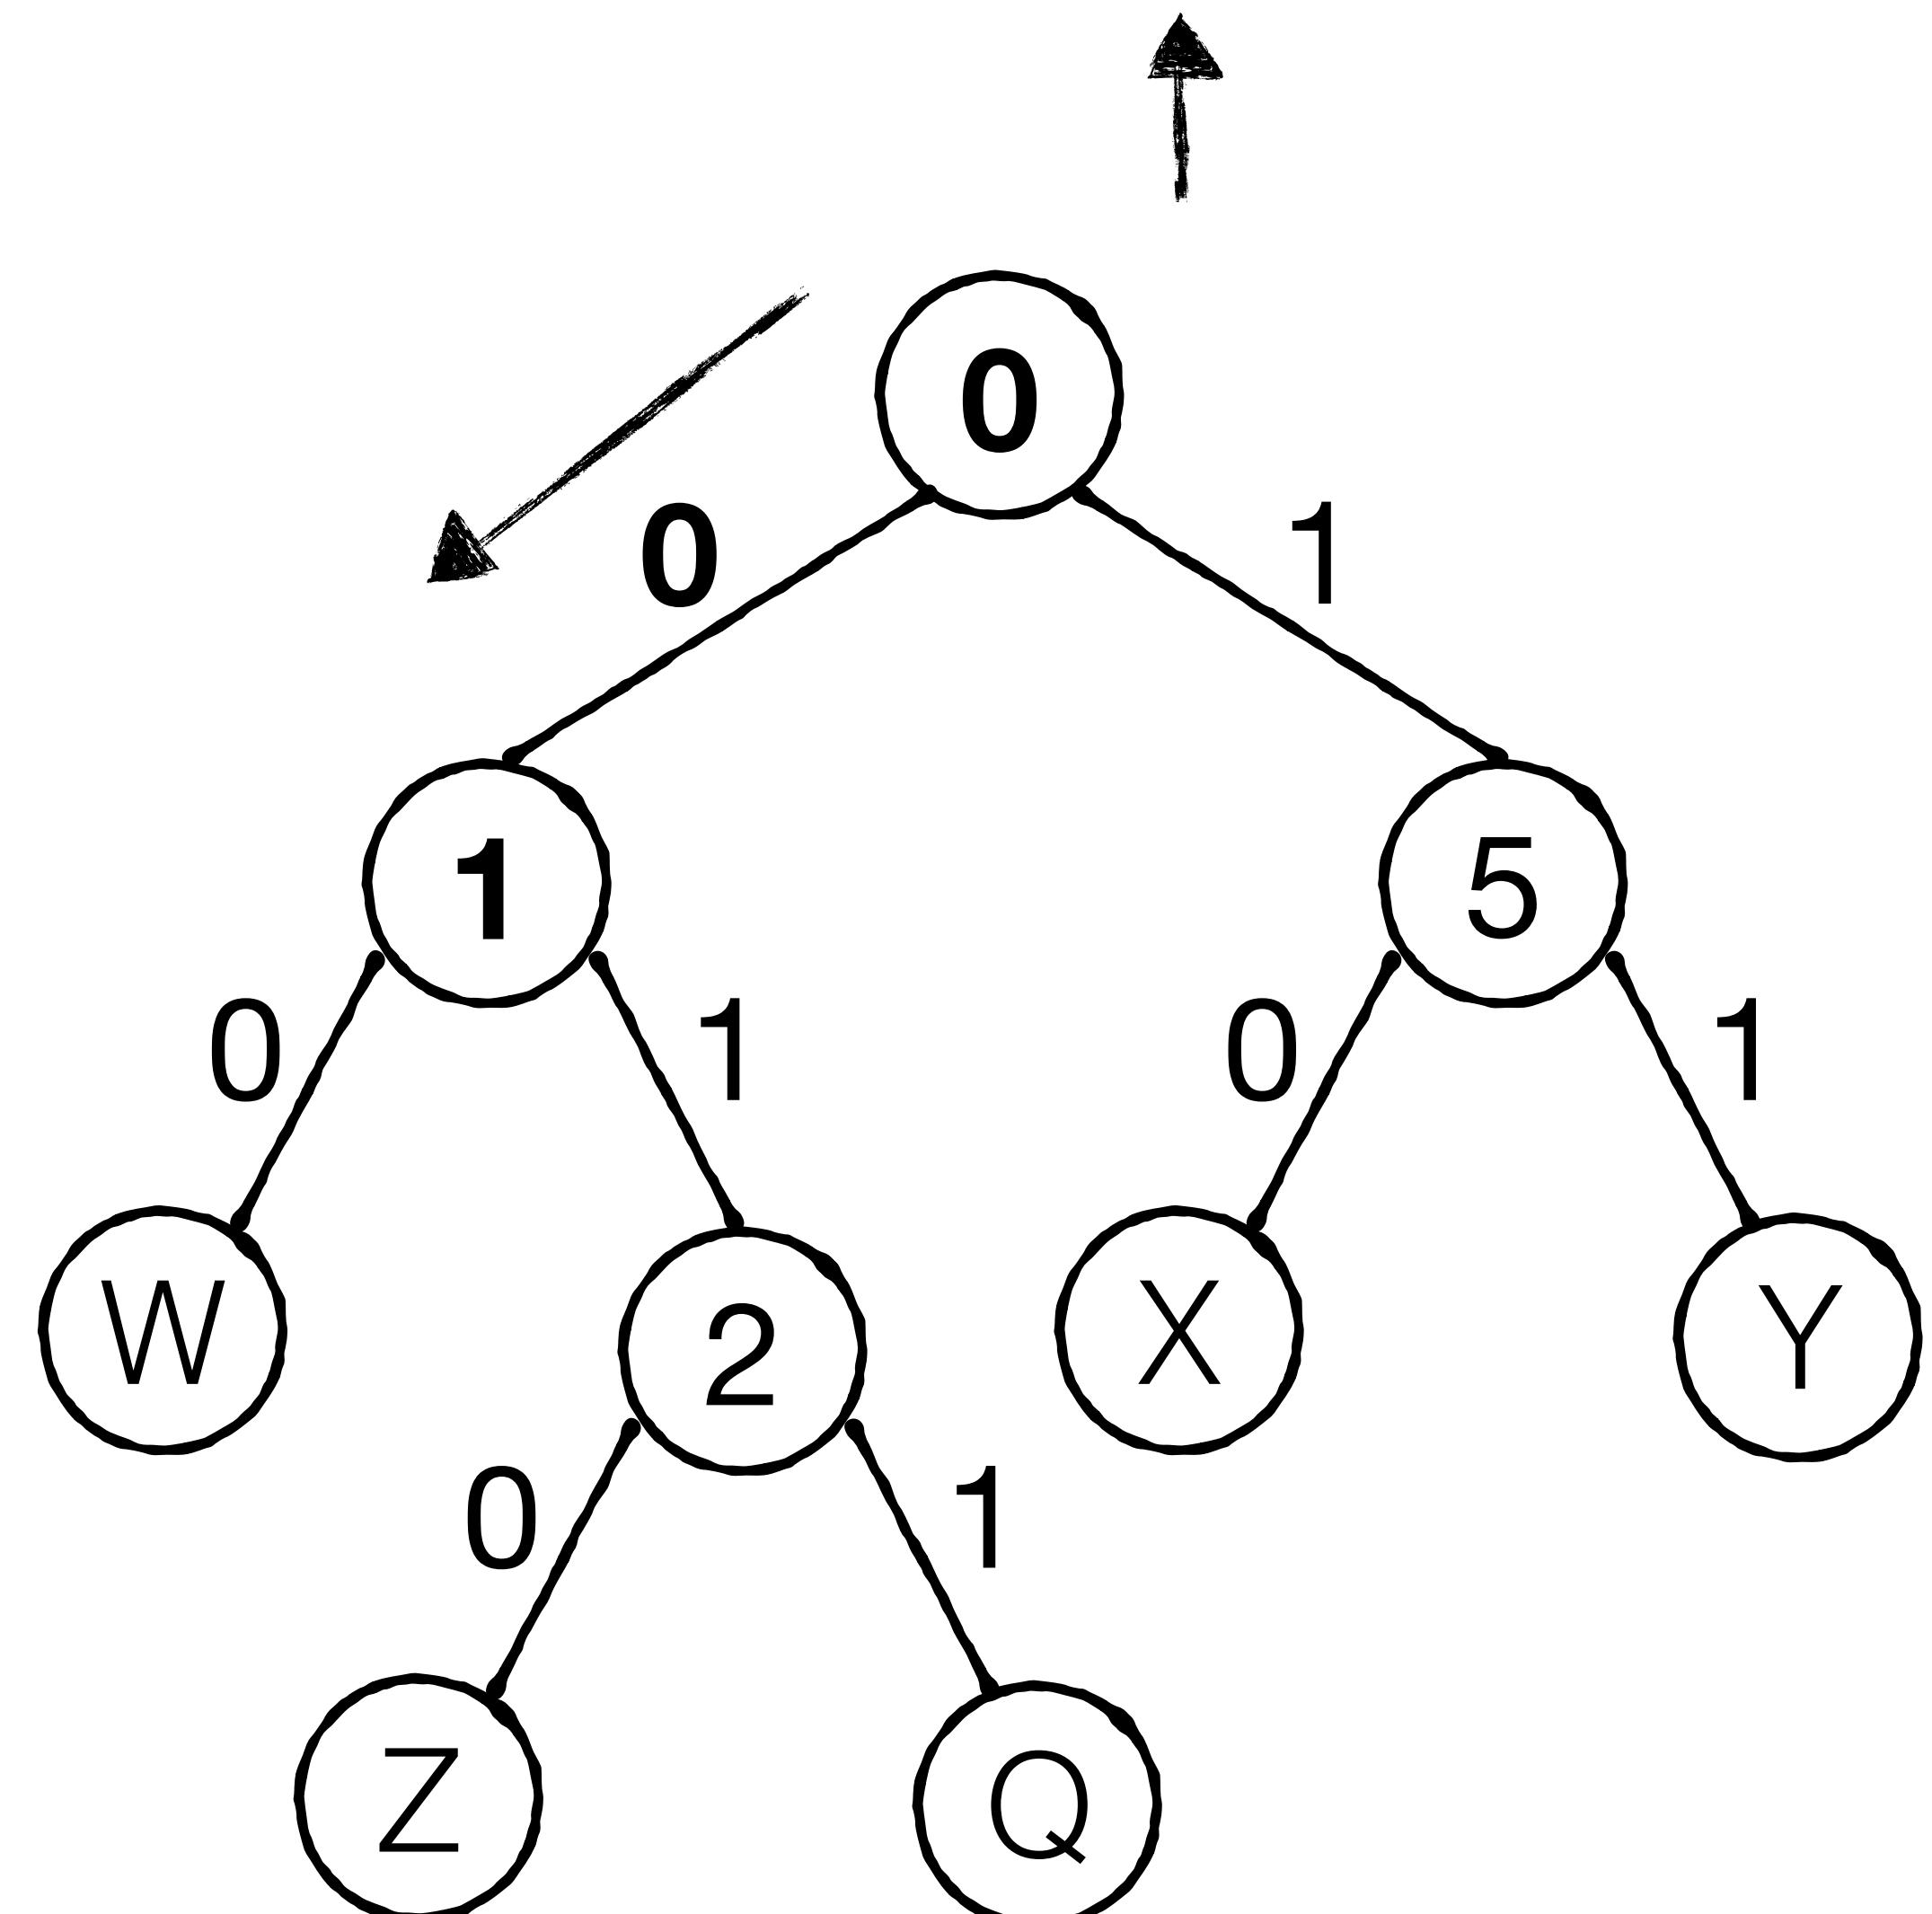
$FP(Q) = 011001$

entries

111110

Topology

11 01 00



Indices

0 0 0 11110

Pointers

W Z Q X Y

$FP(Q) = 011001$

entries

Topology

Indices

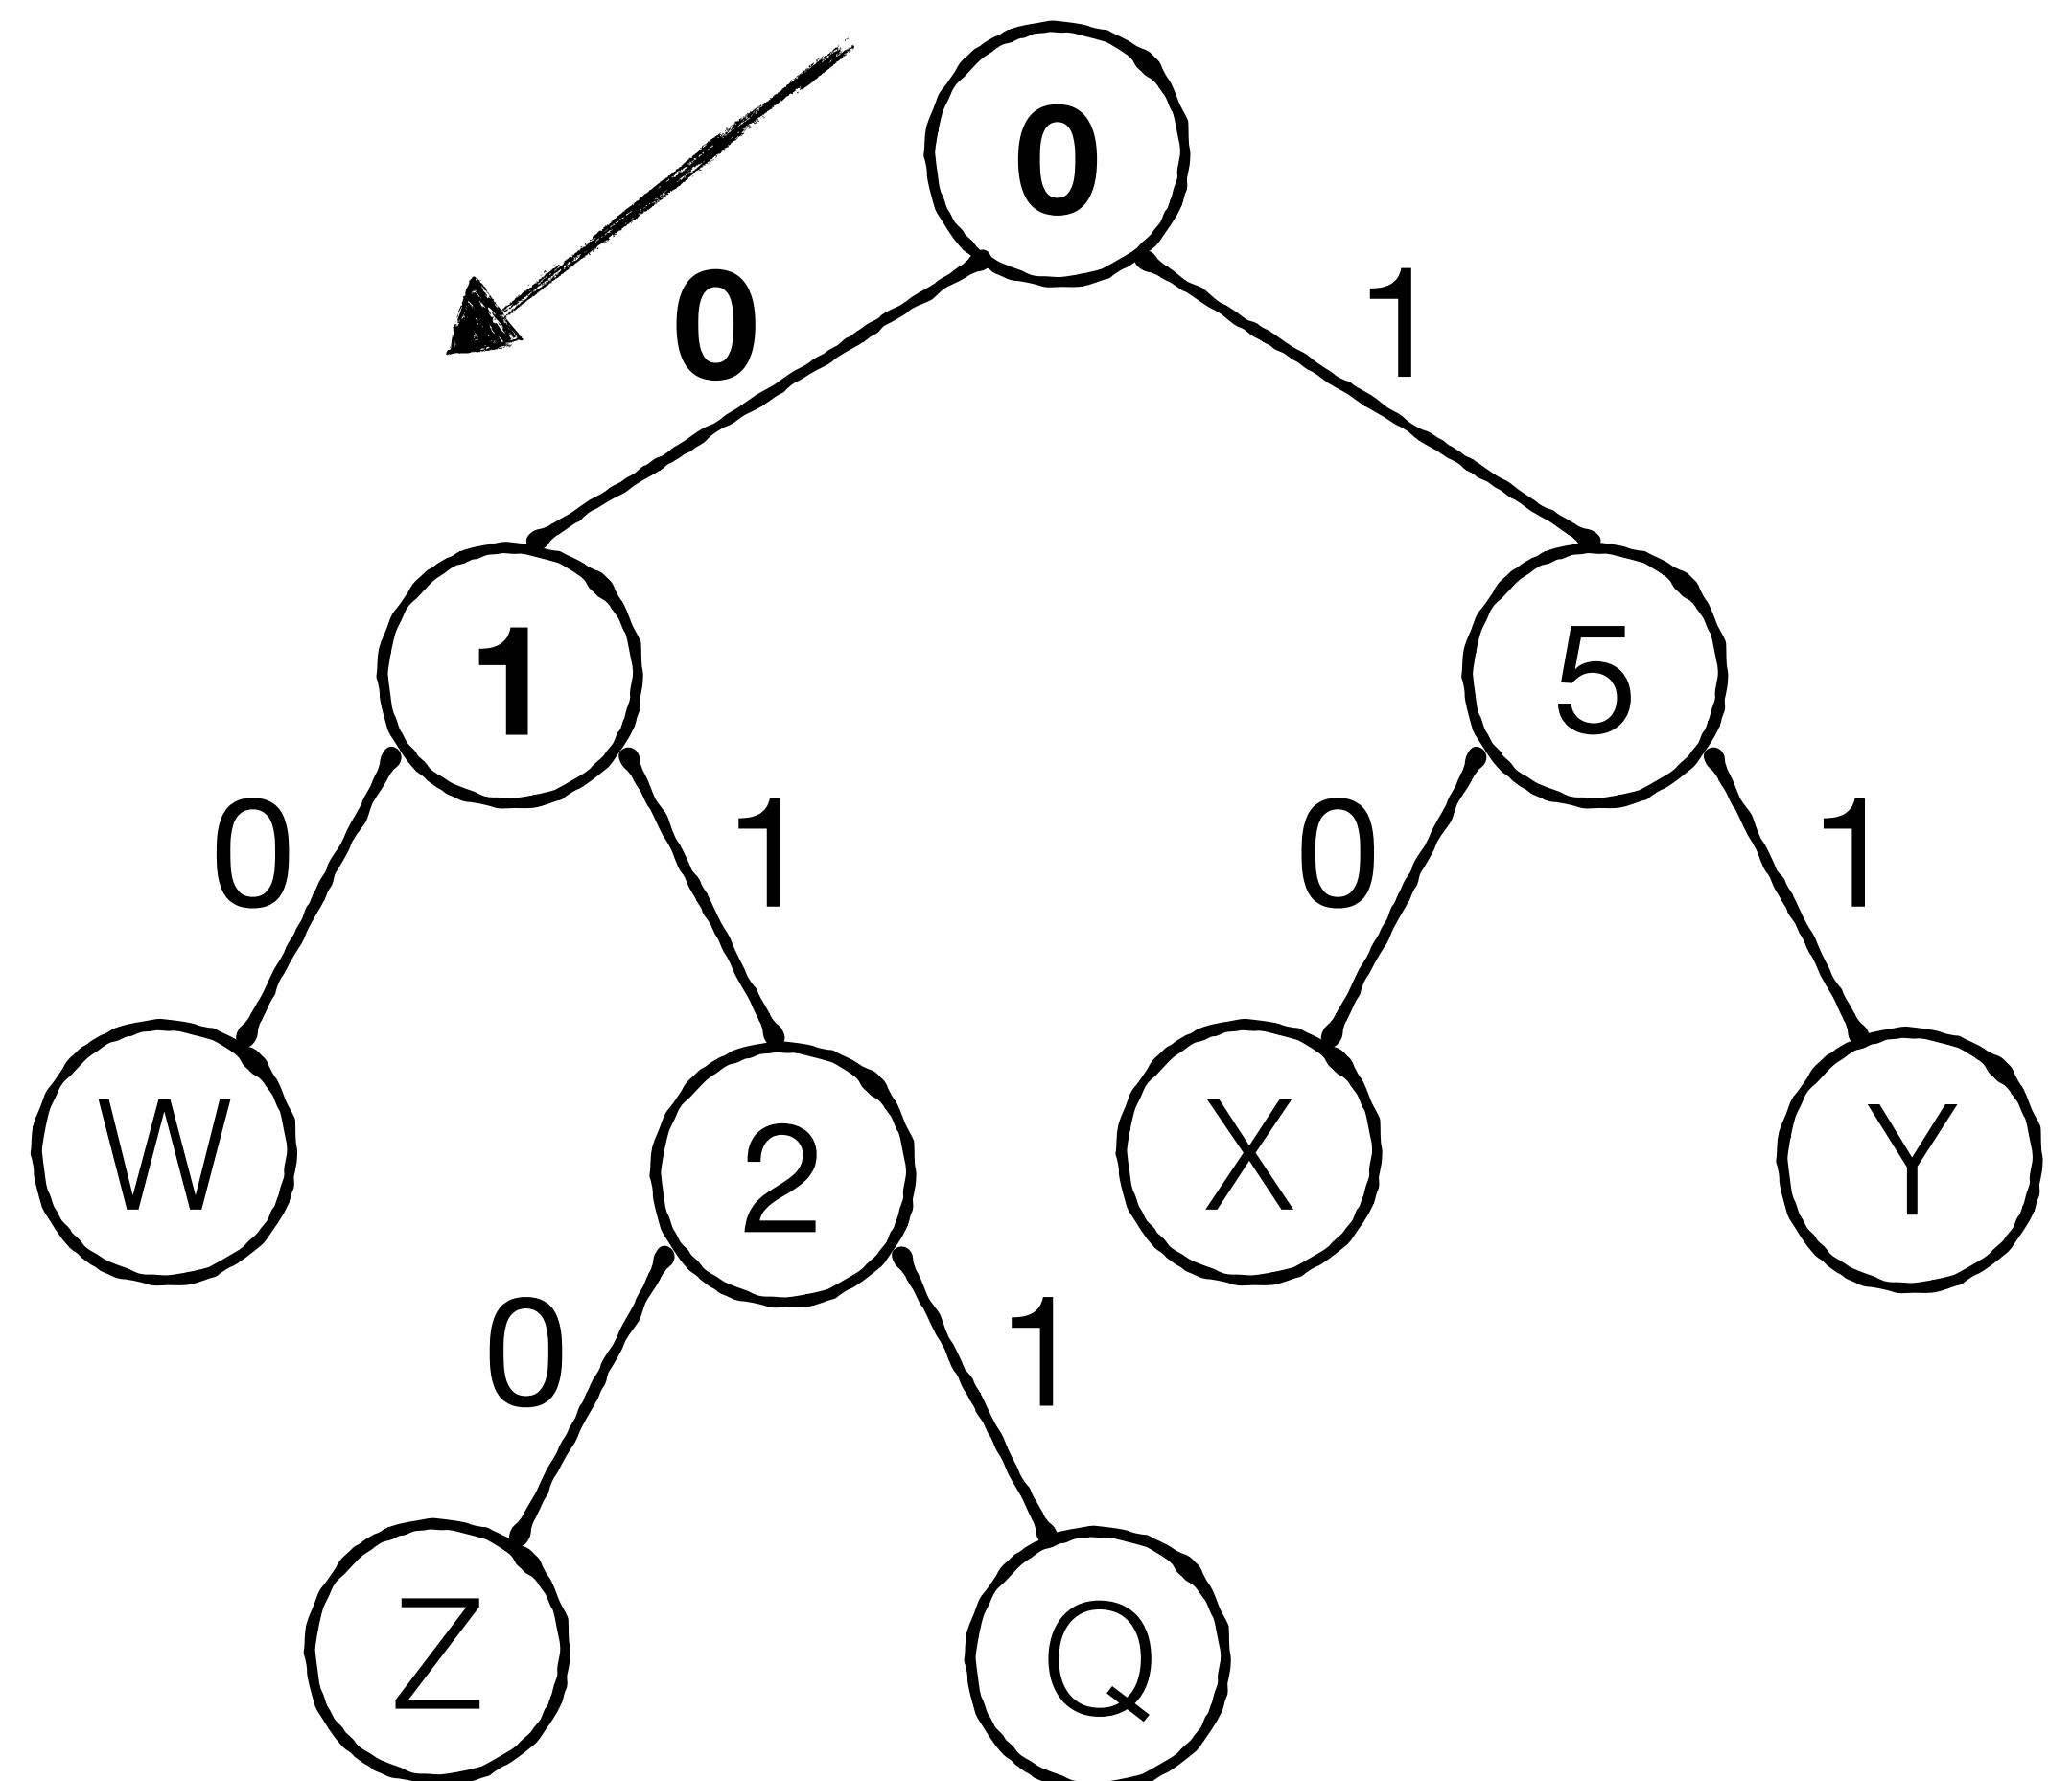
Pointers

111110

110100

00011110

W Z Q X Y



$$FP(Q) = 0\ 1\ 1\ 0\ 0\ 1$$

entries

111110

Topology

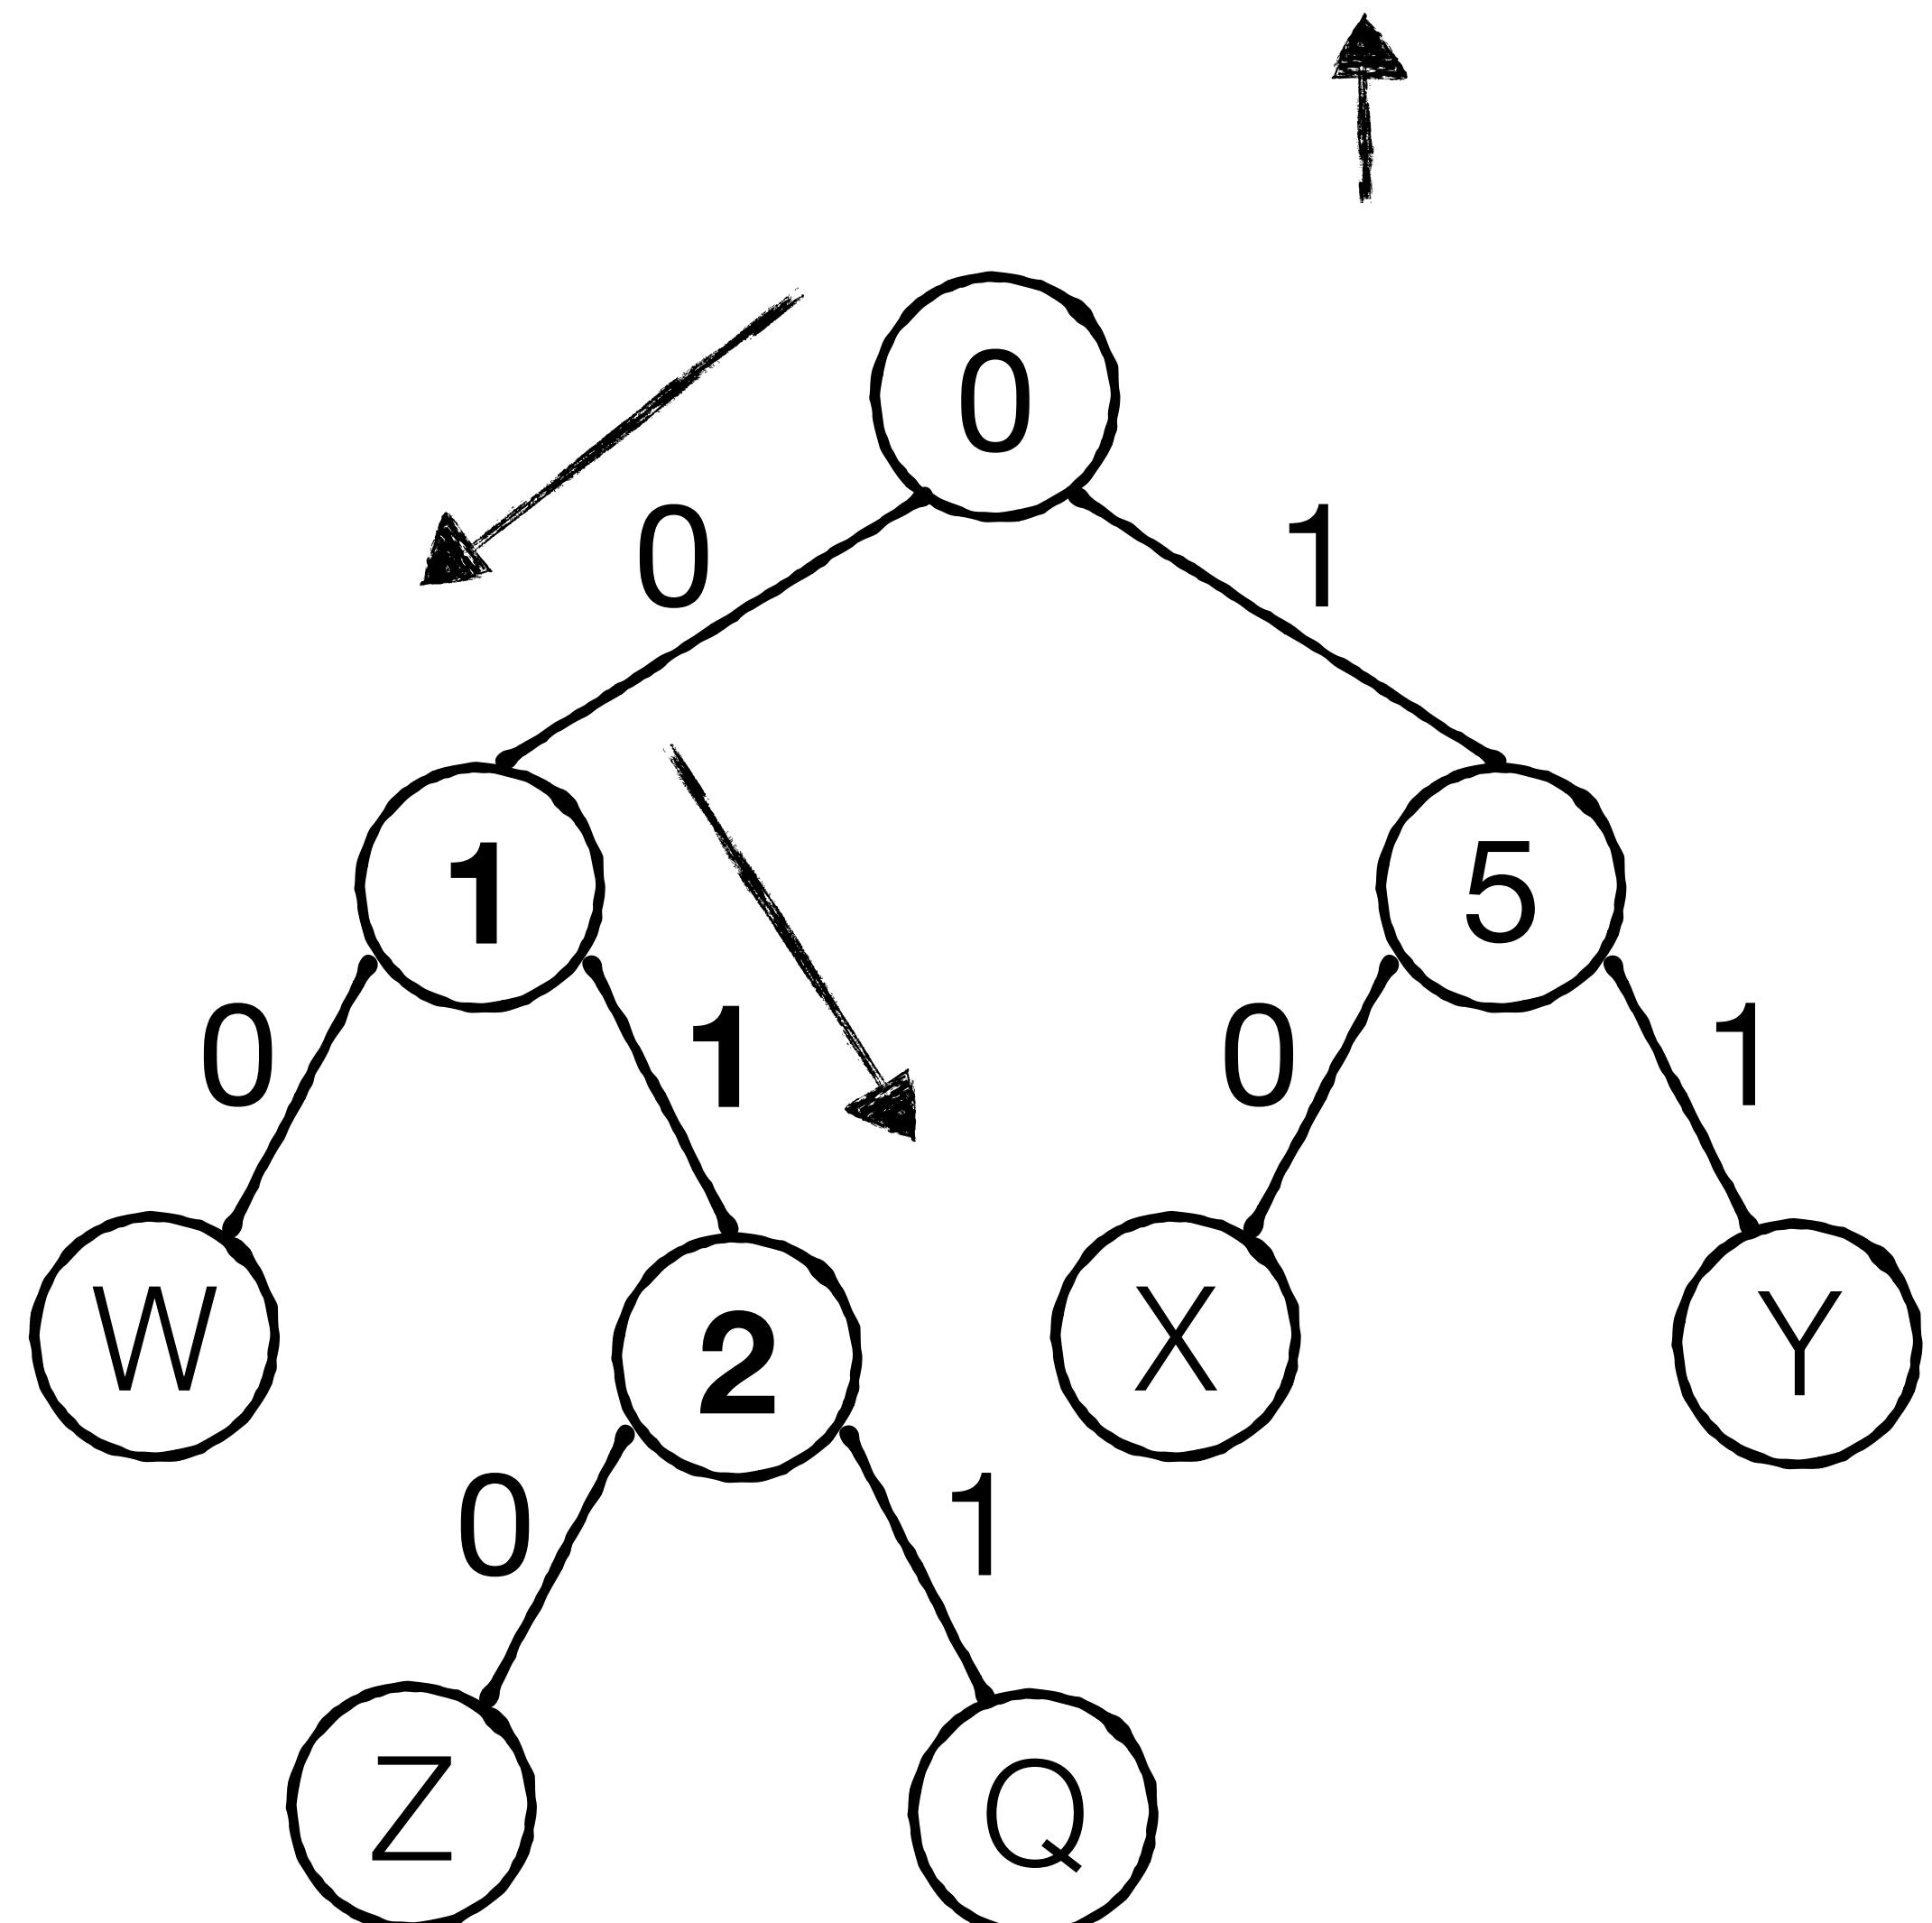
11 01 00

Indices

0 0 0 11110

Pointers

W Z Q X Y



entries

Topology

Indices

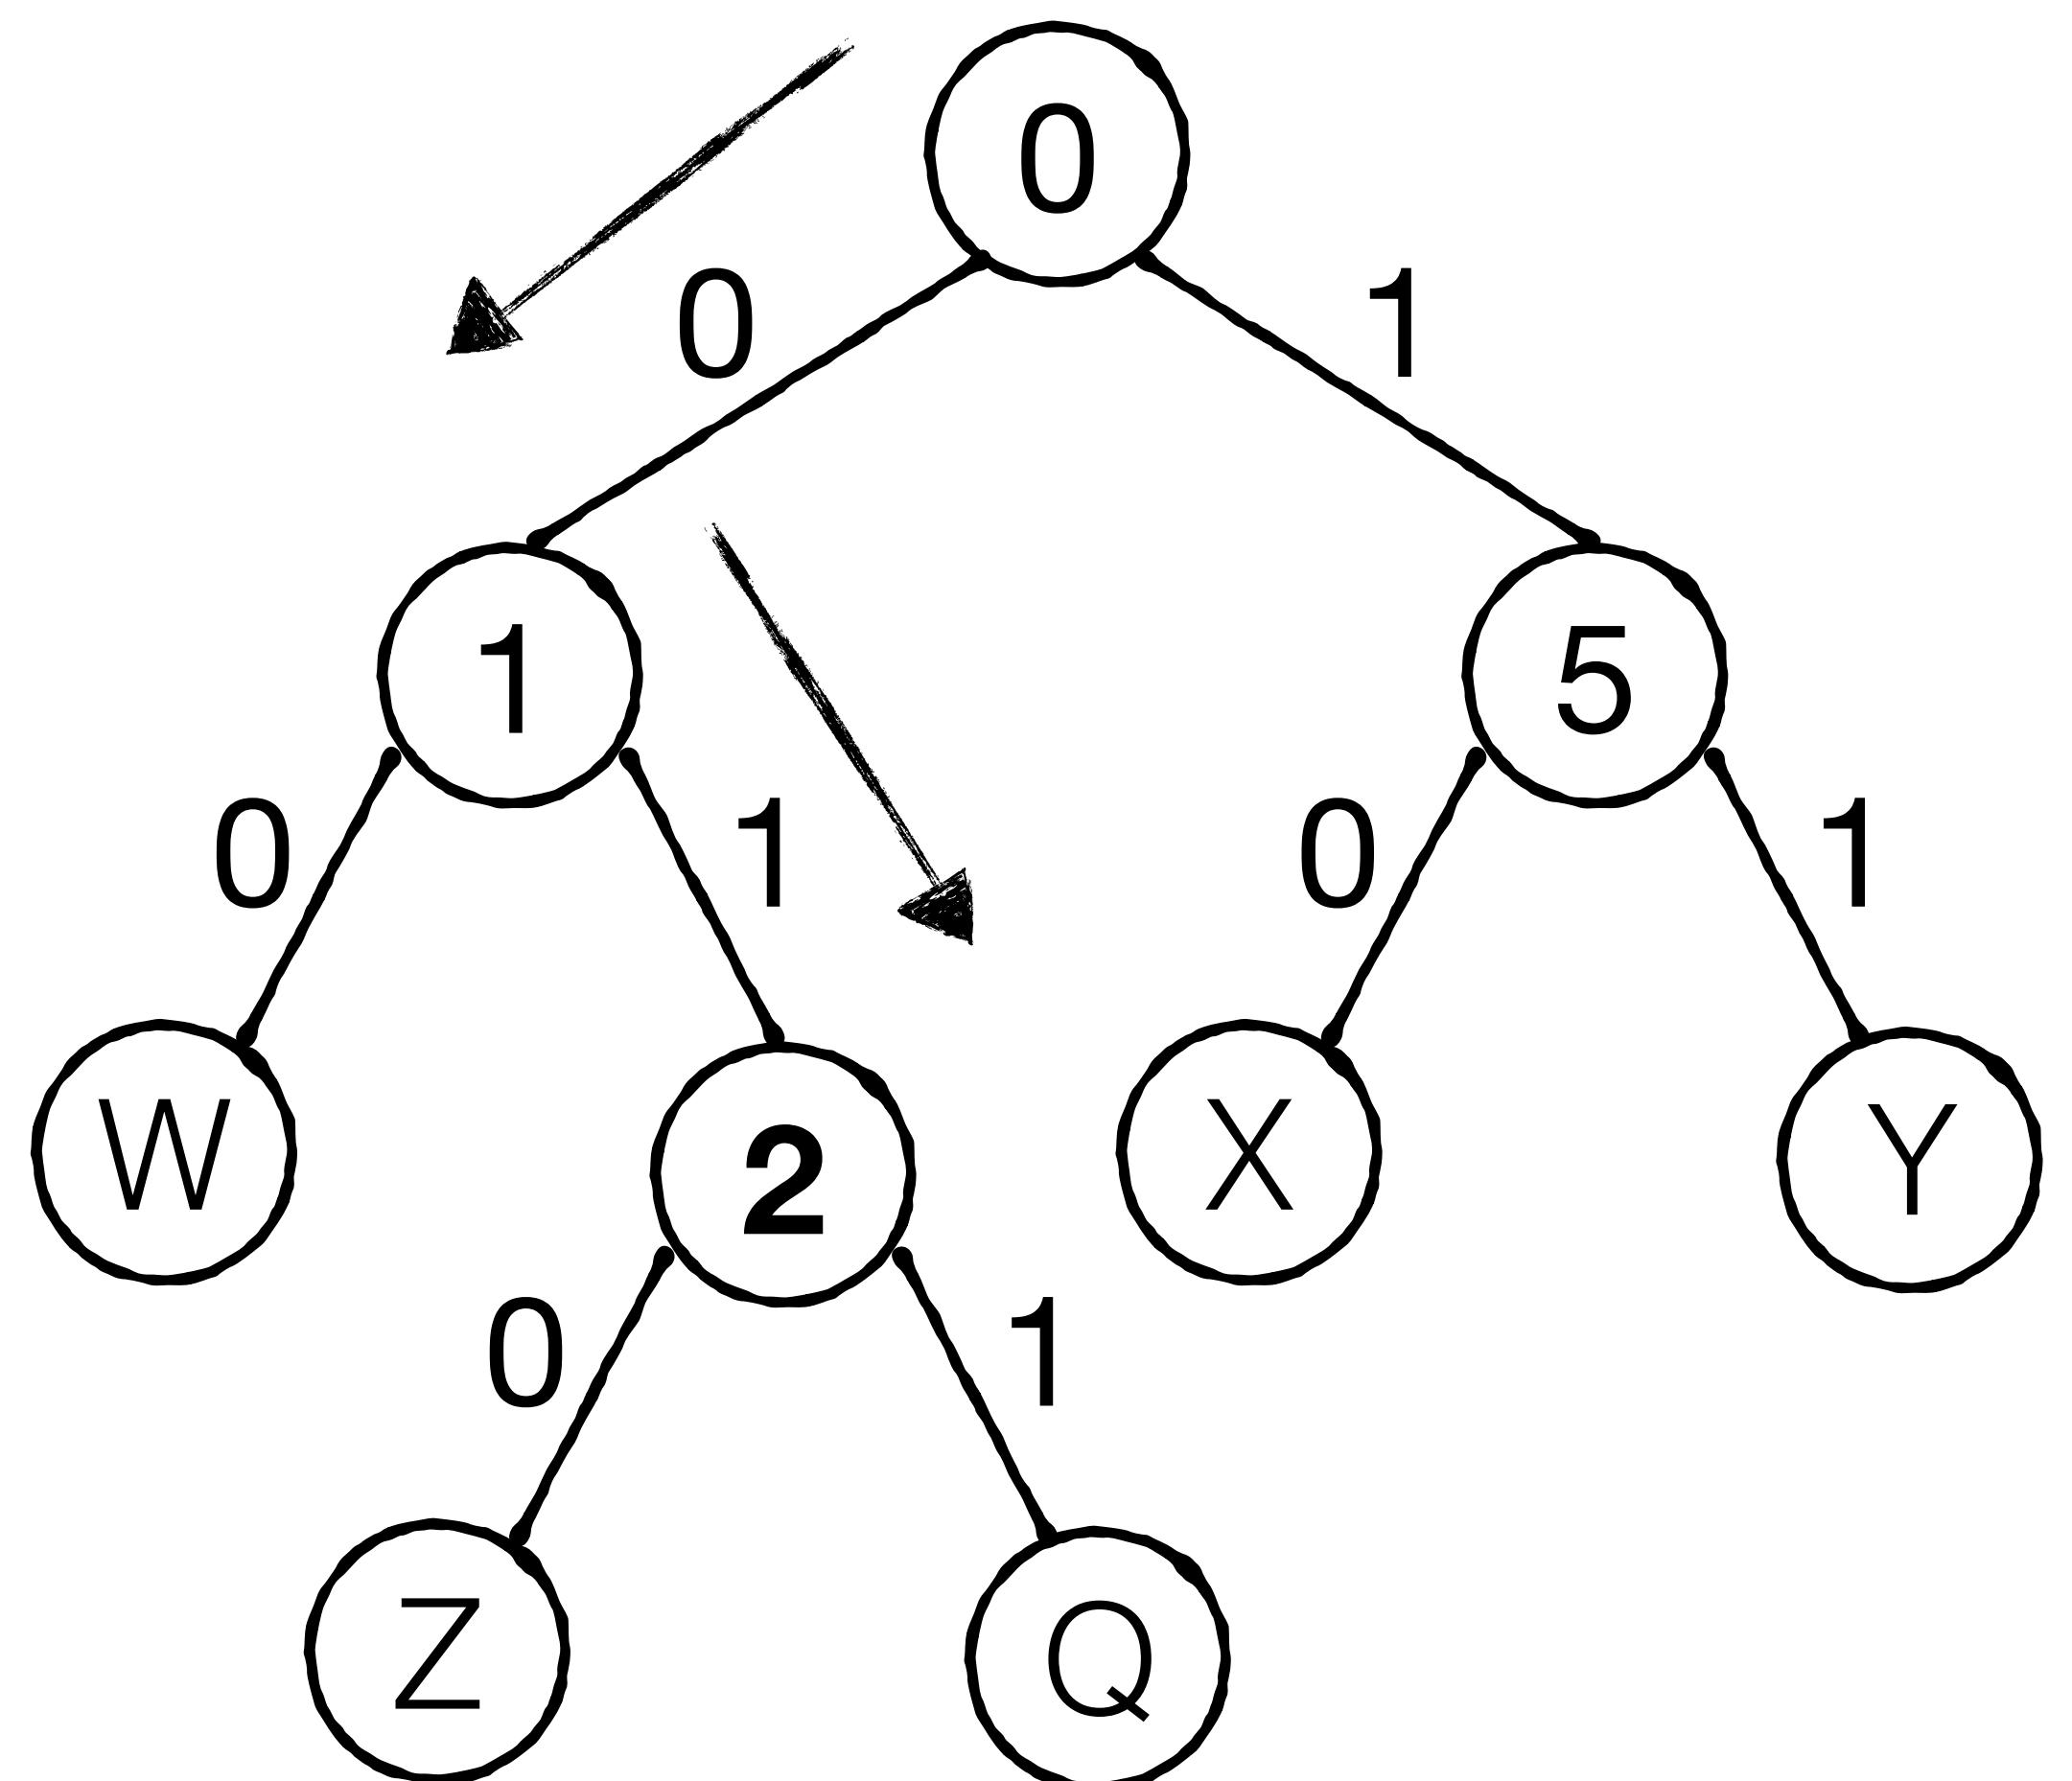
Pointers

111110

11 01 00

0 0 0 1 1 1 1 0

W Z Q X Y



$$FP(Q) = 011001$$

entries

111110

Topology

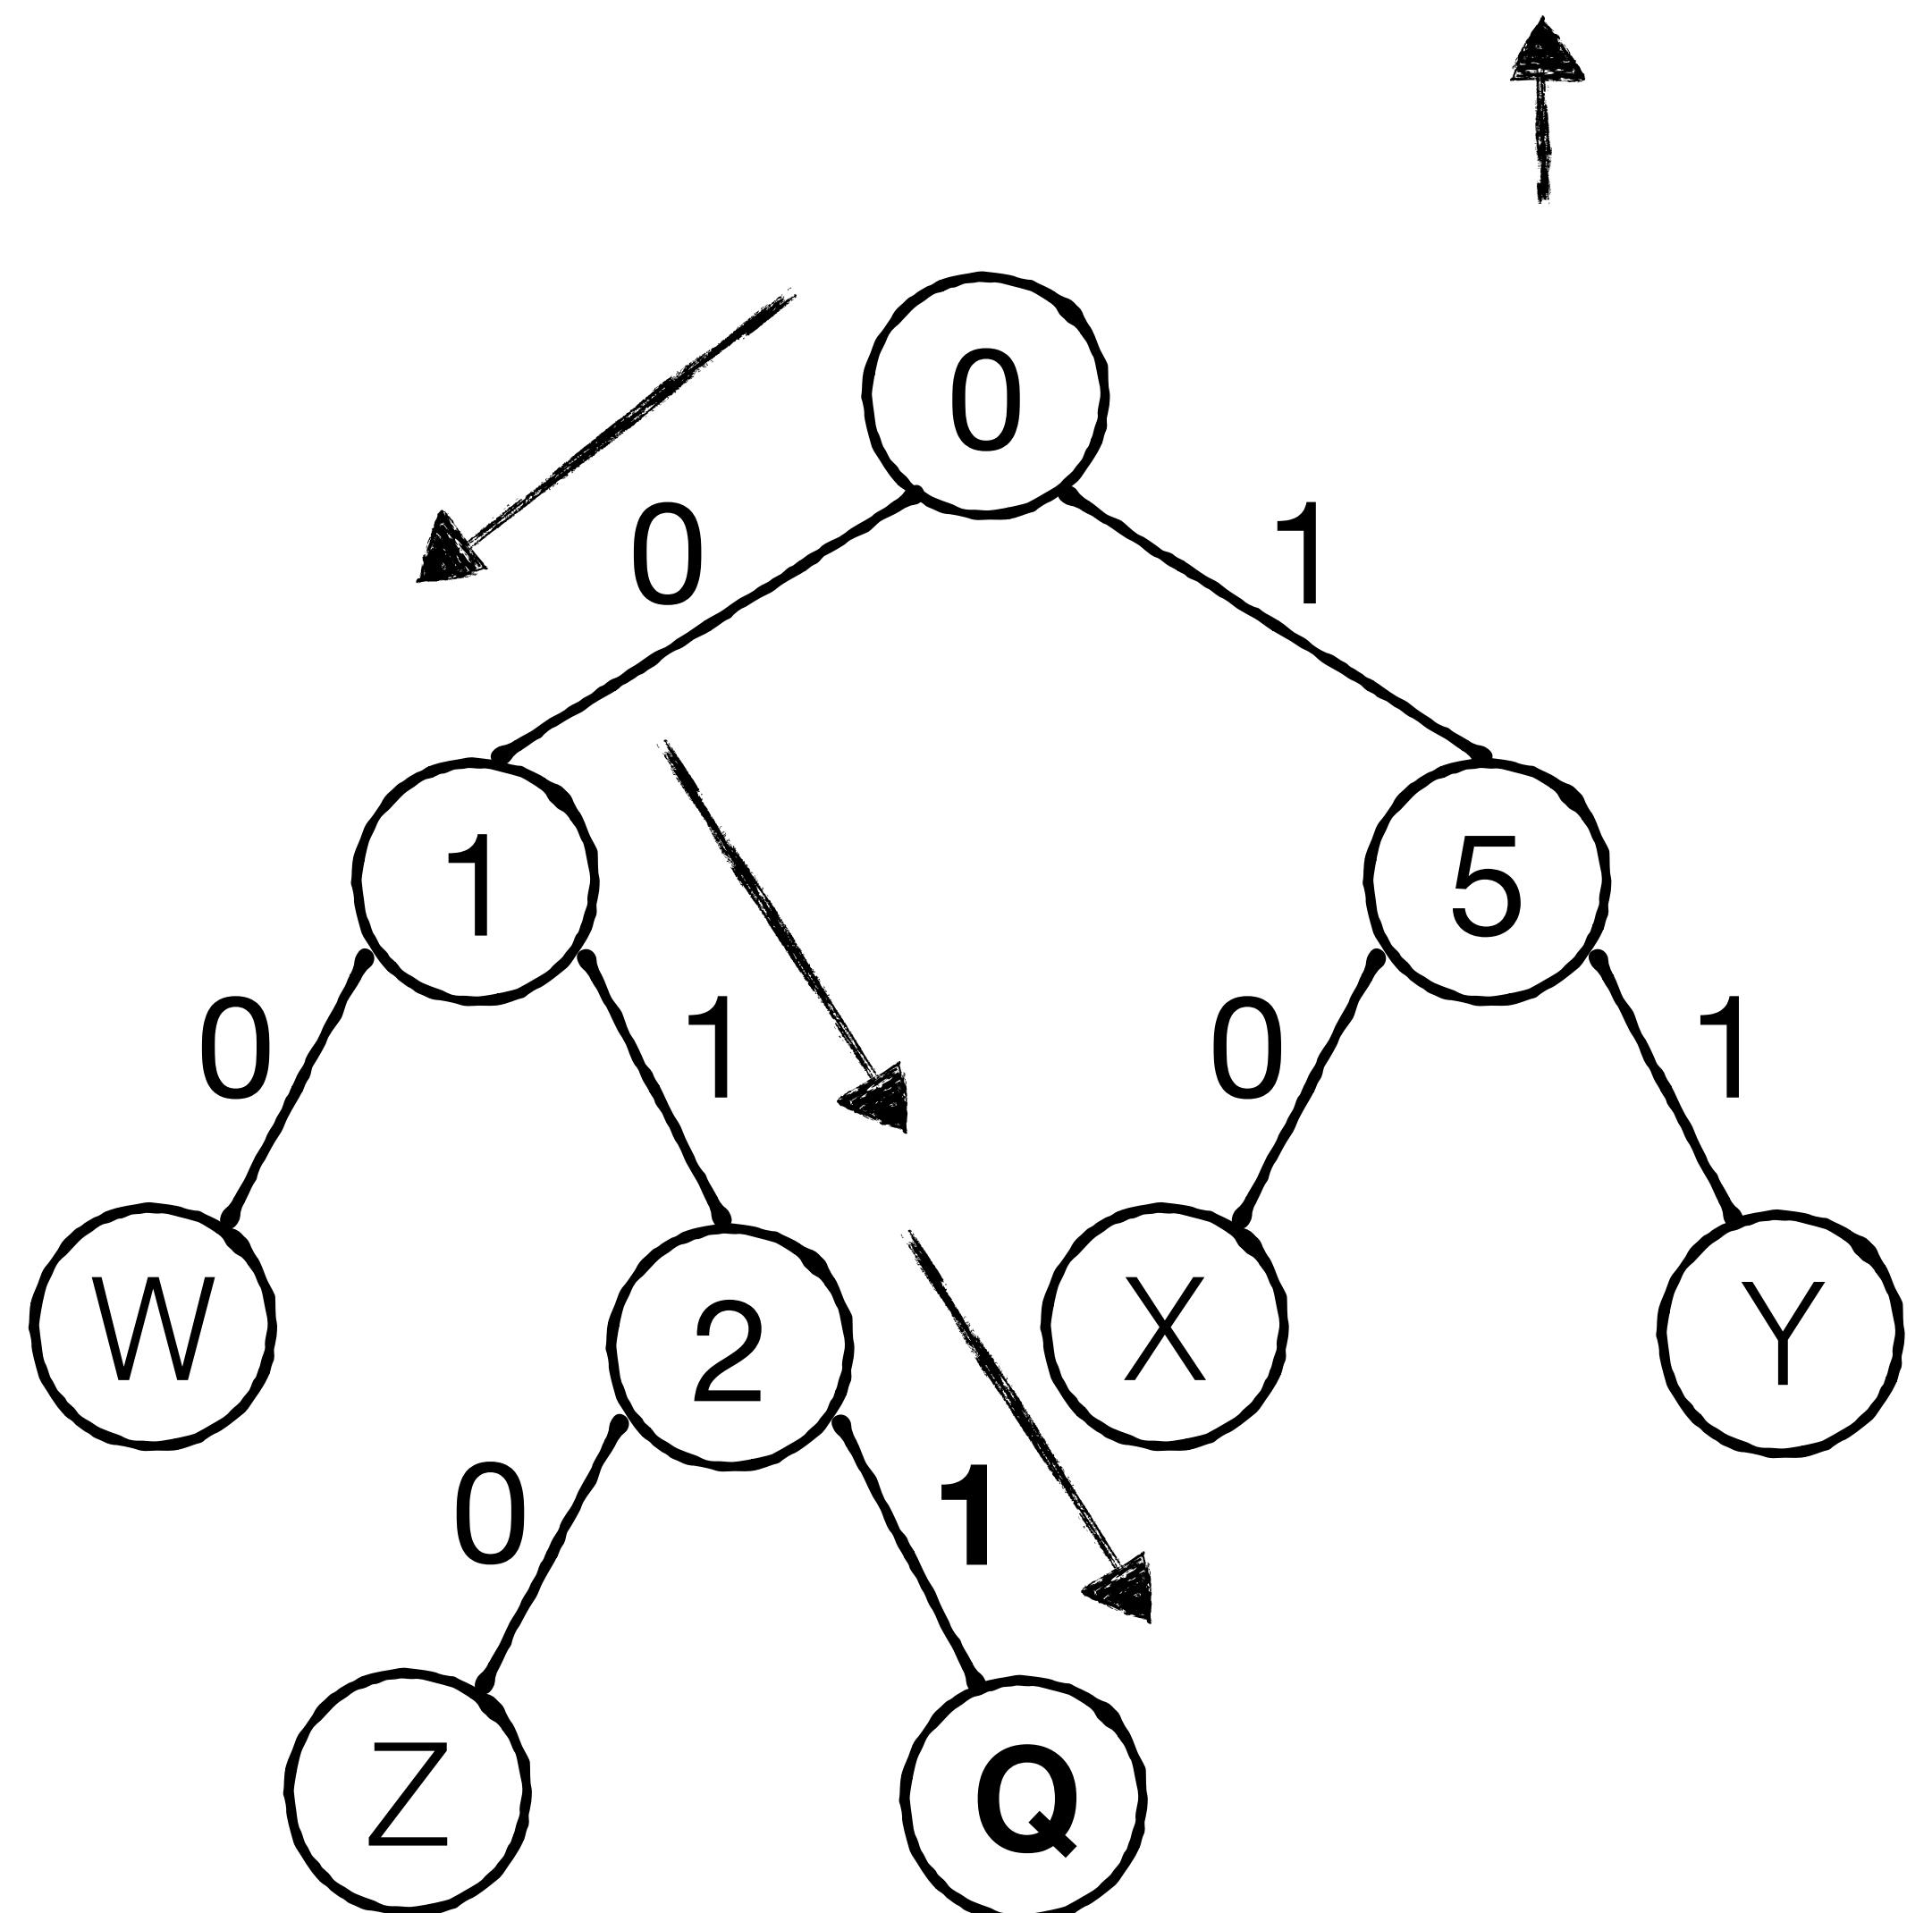
11 01 00

Indices

0 0 0 11110

Pointers

W Z Q X Y



$FP(Q) = 011001$

entries

Topology

Indices

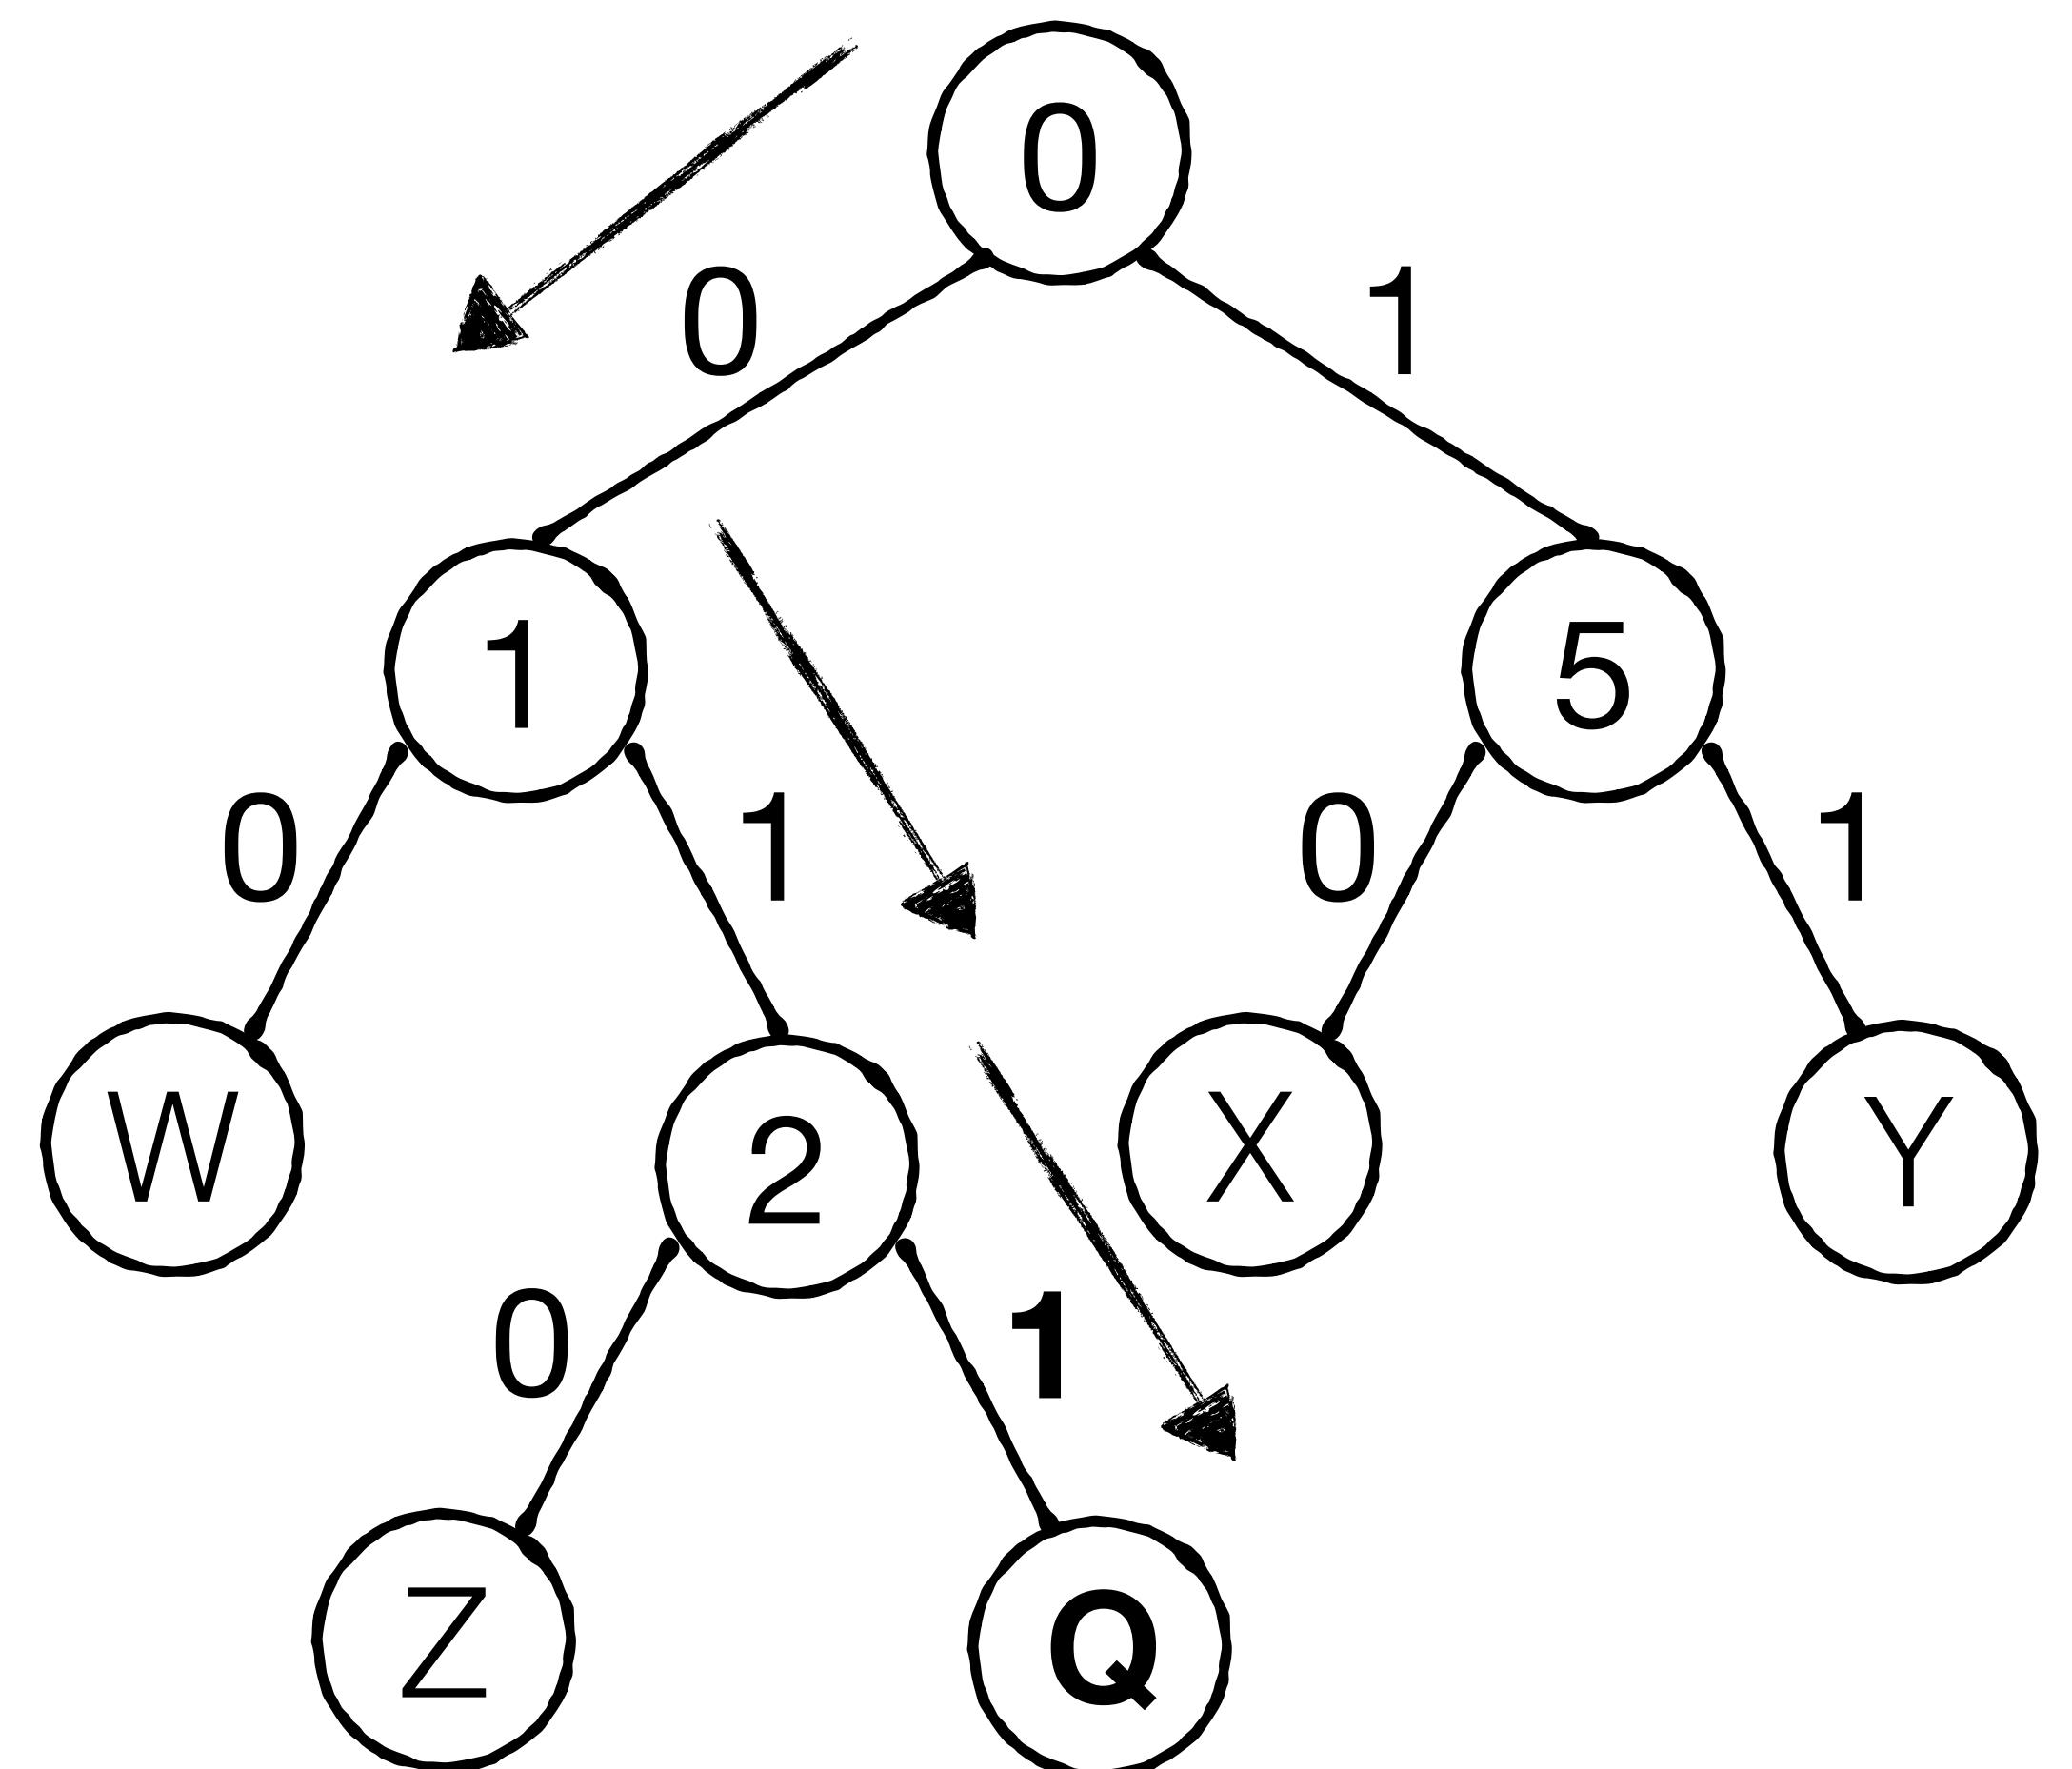
Pointers

111110

11 01 00

0 0 0 1 1 1 1 0

W Z Q X Y

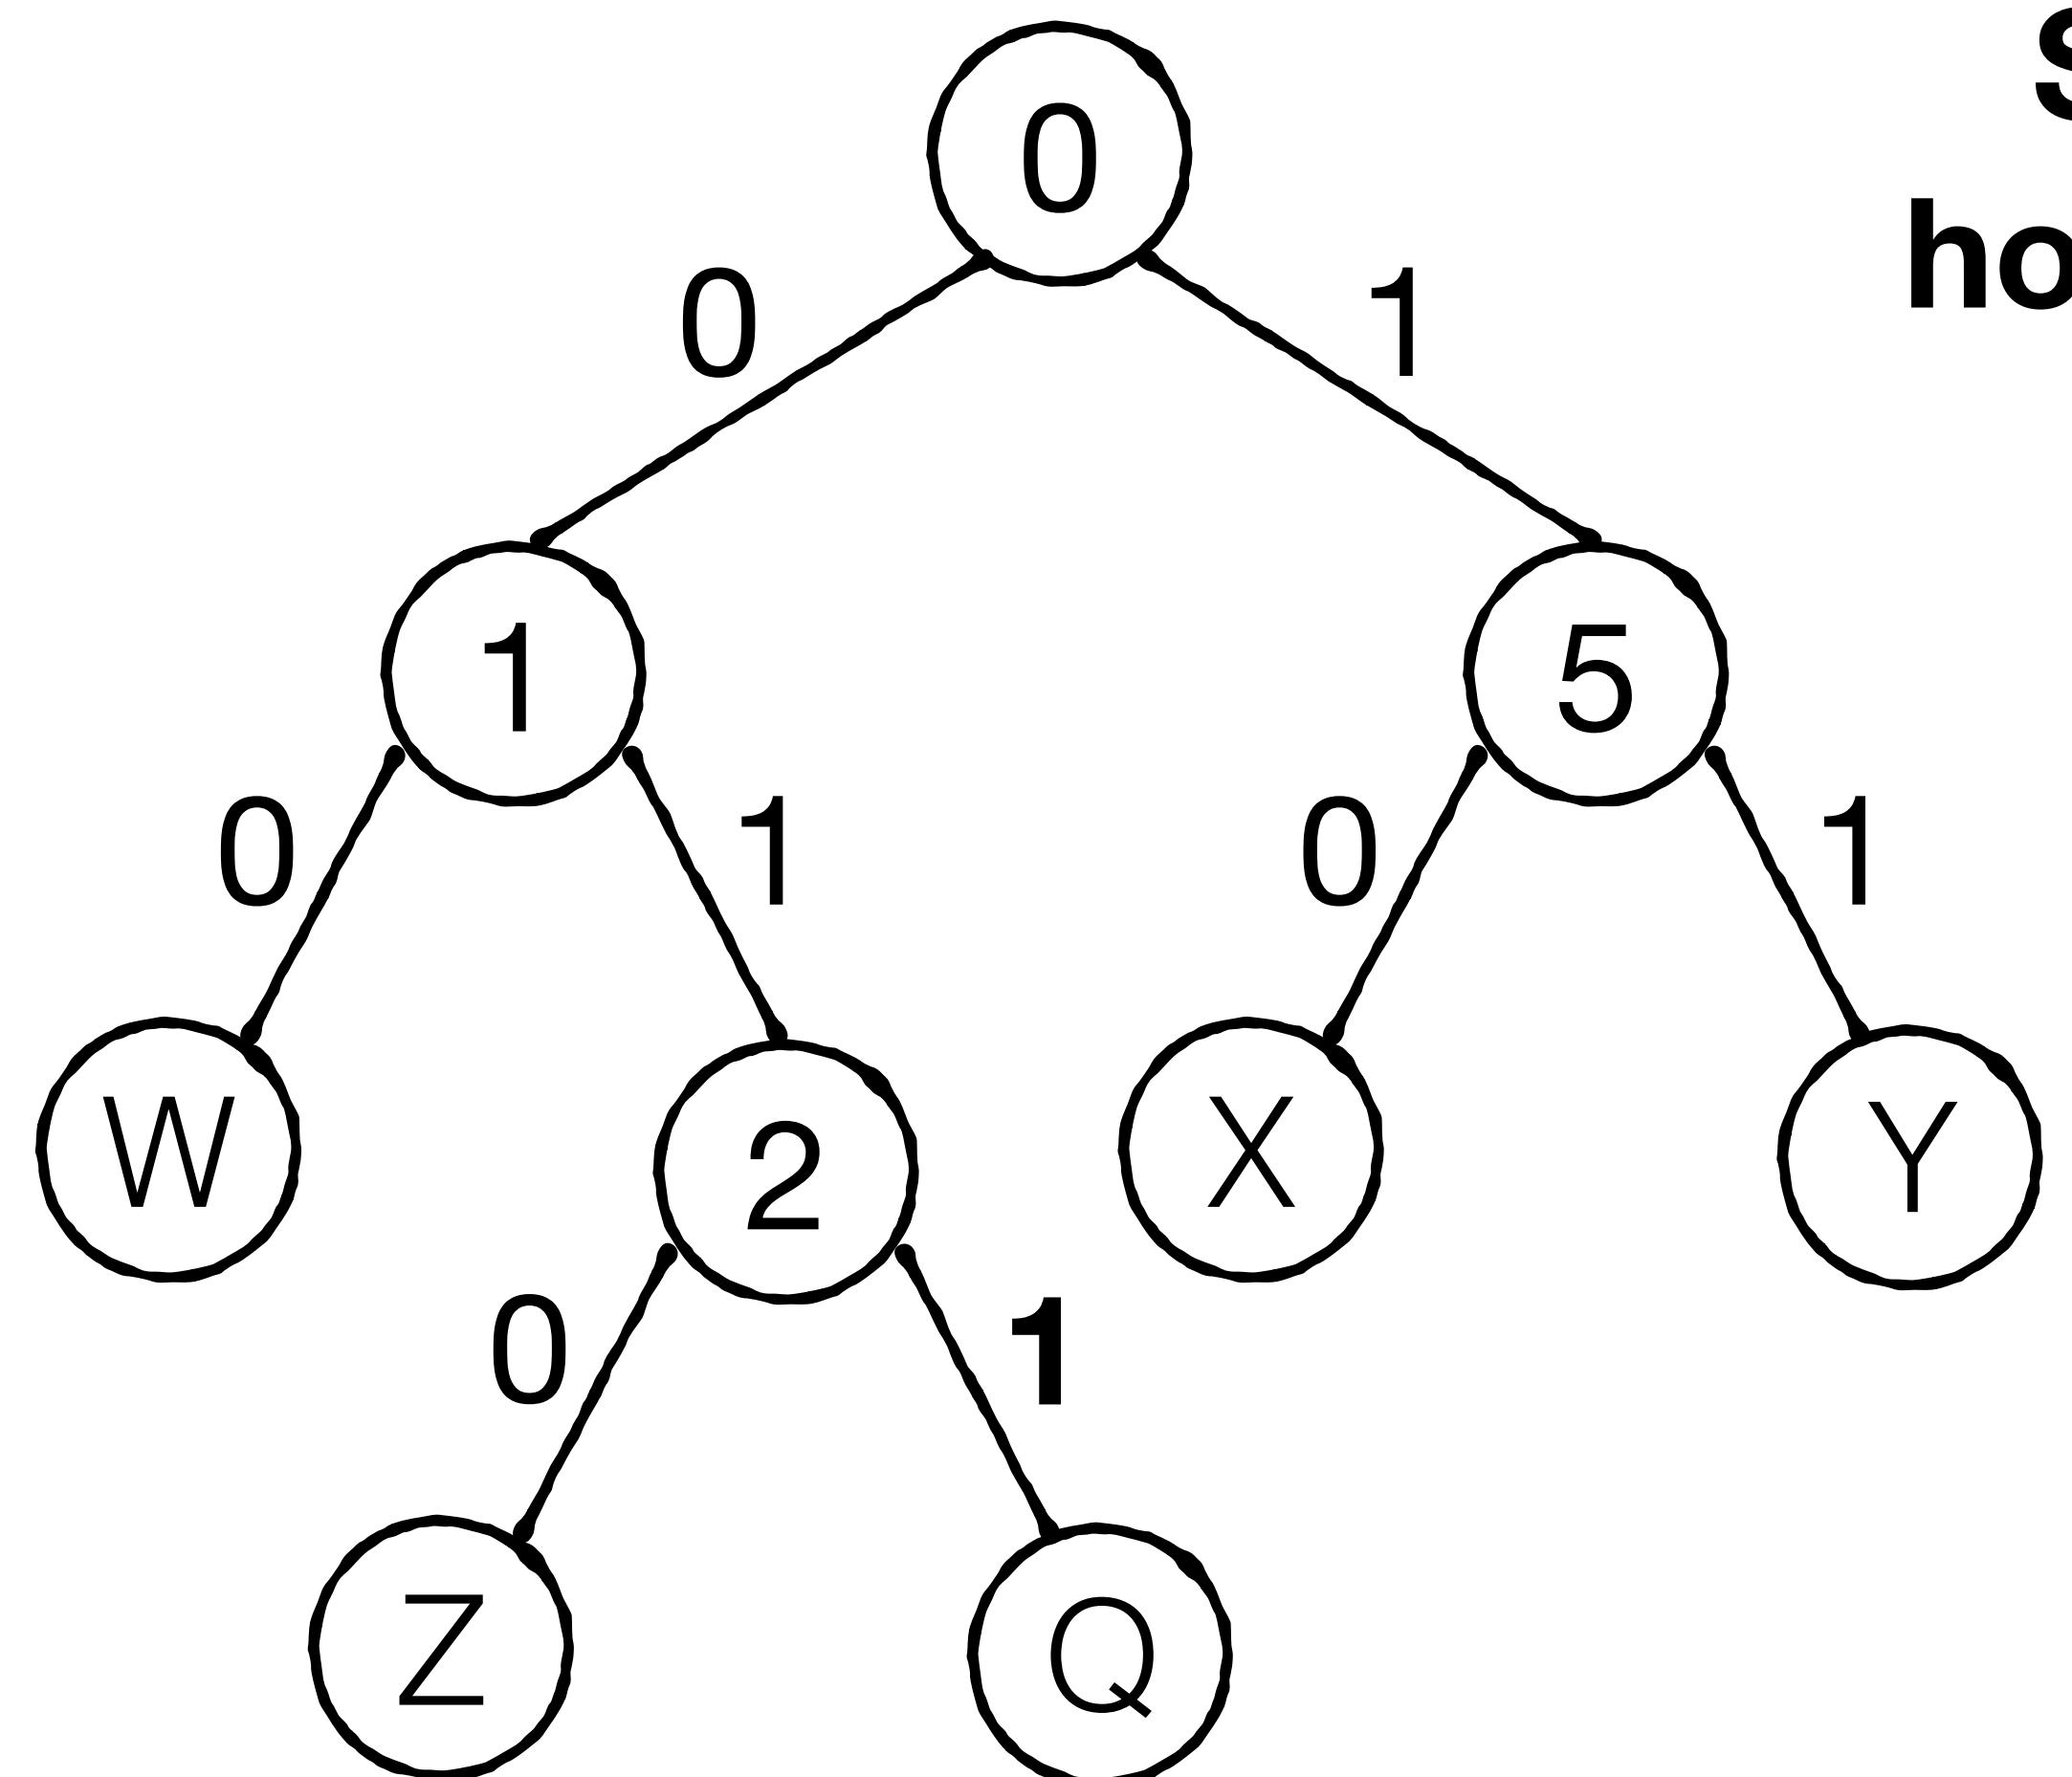


entries
111110

Topology
11 01 00

Indices
0 0 0 11110

Pointers
W Z Q X Y



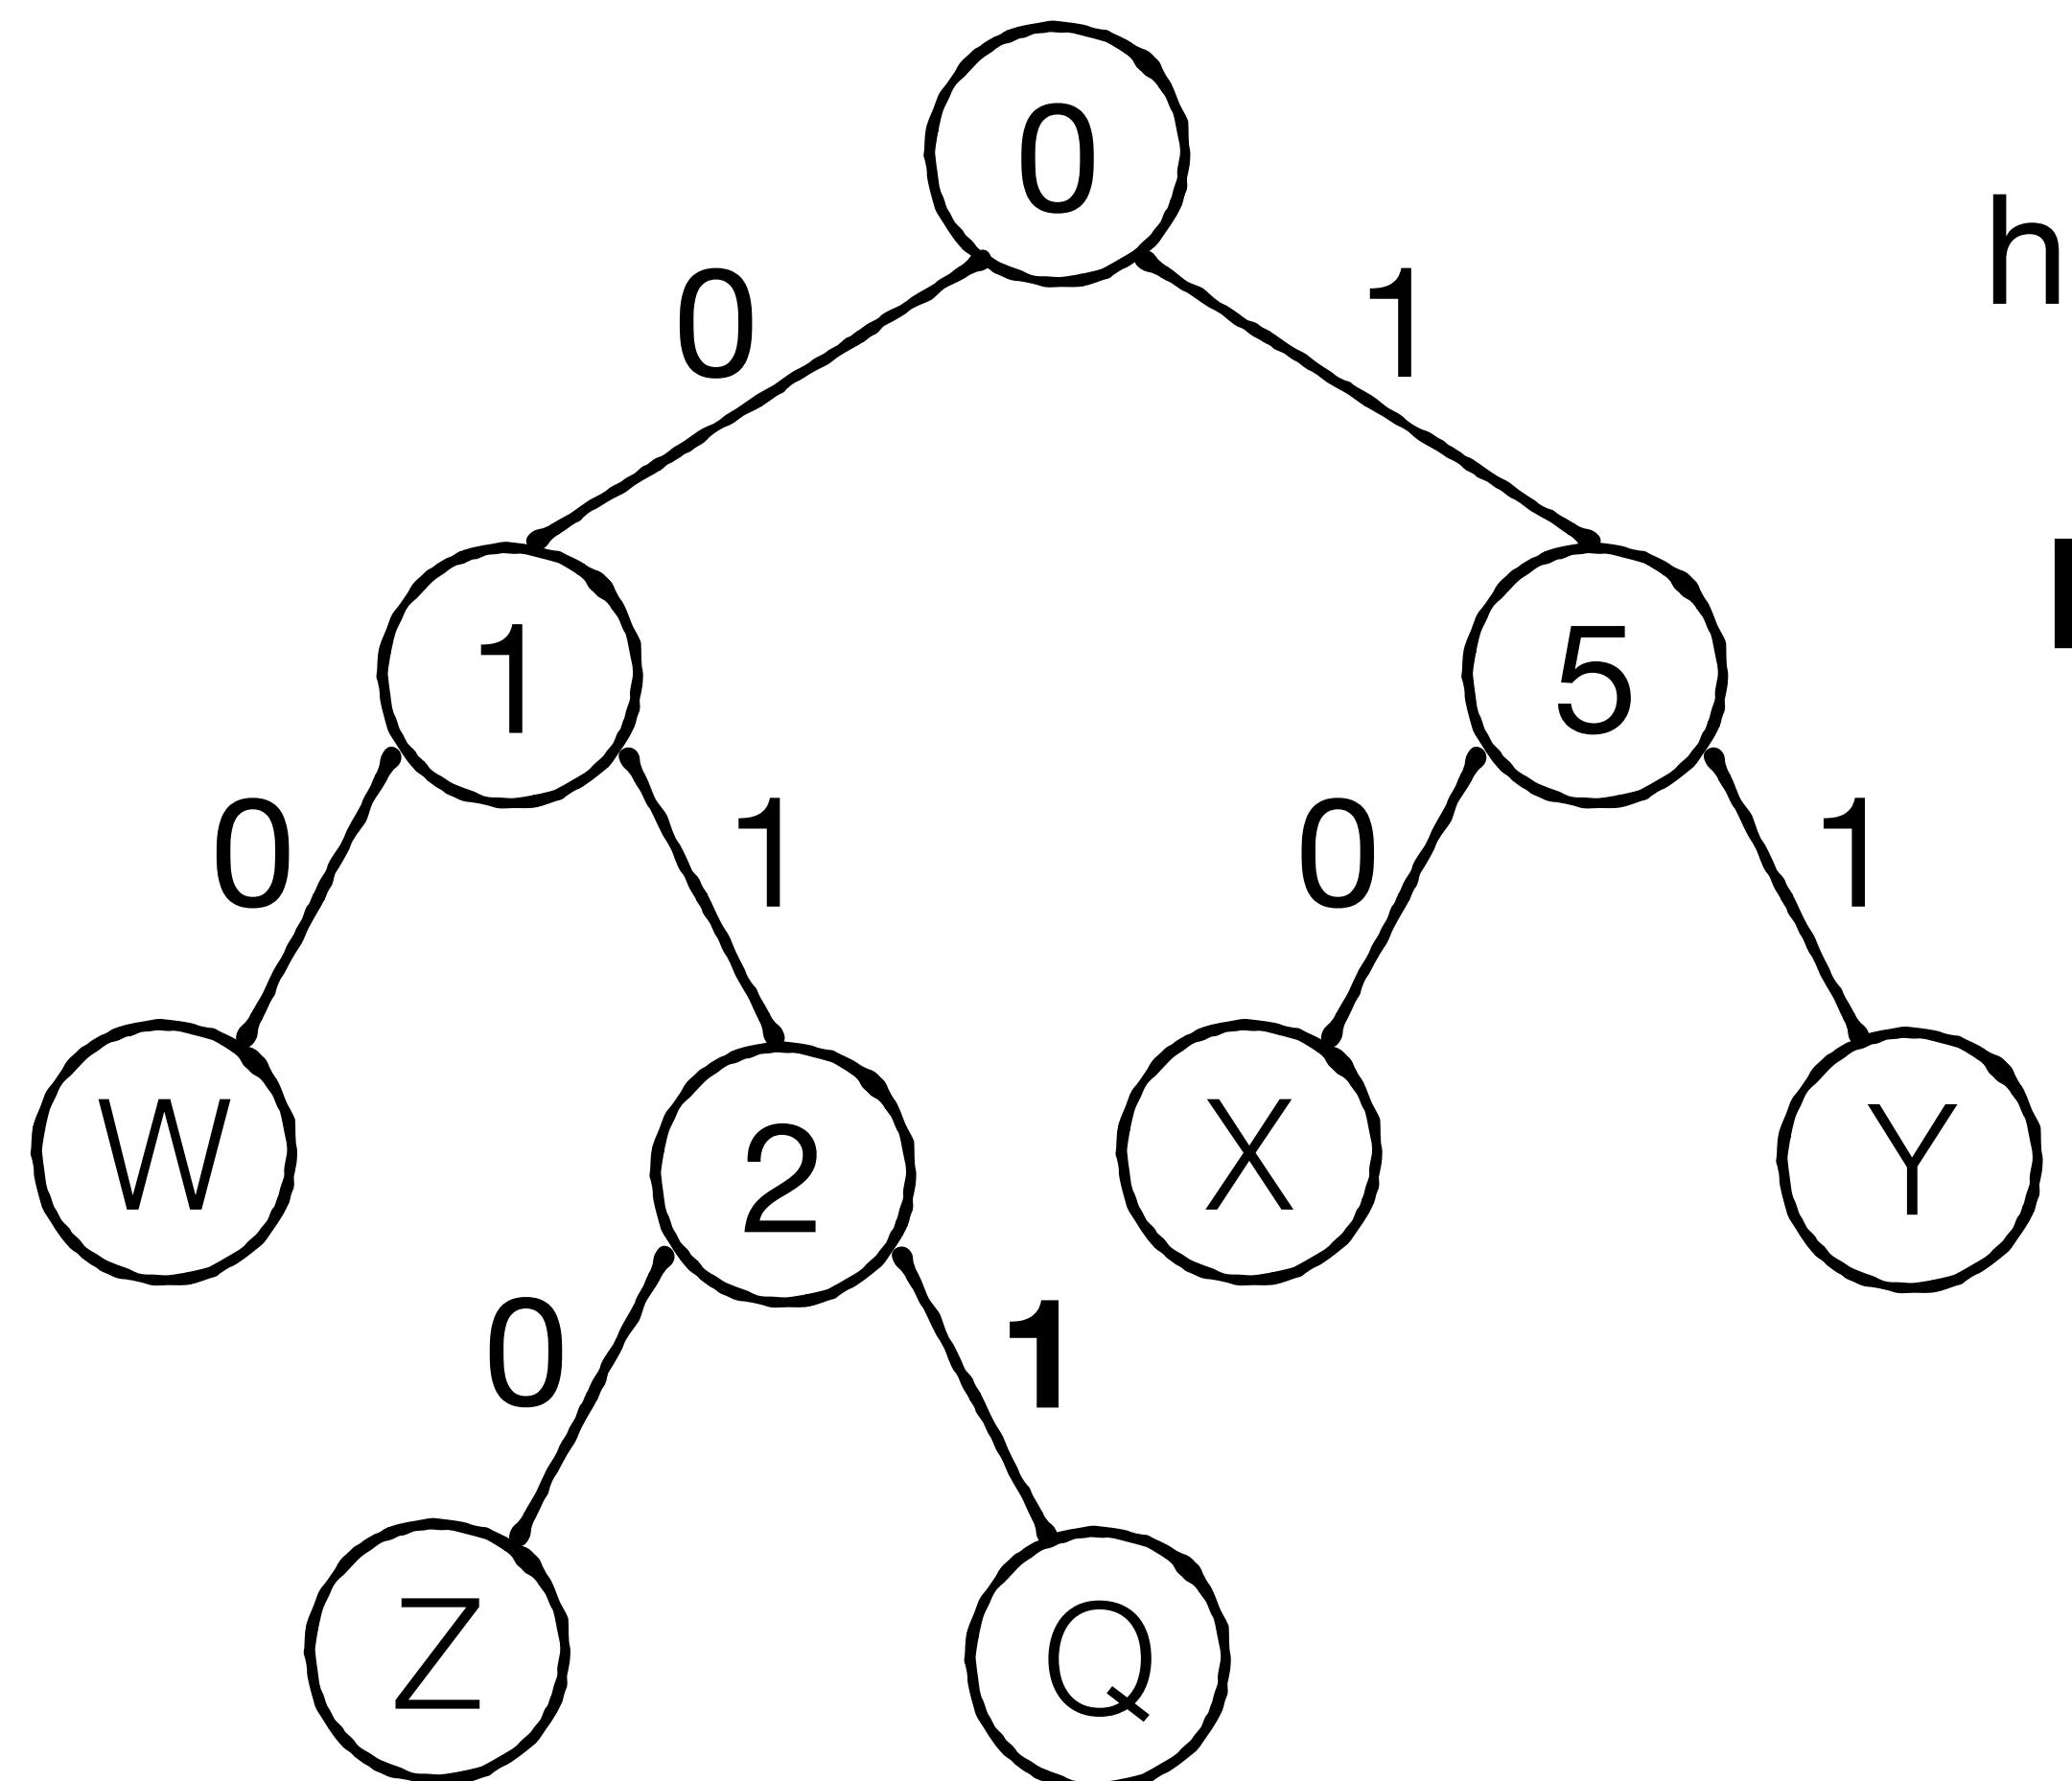
**Since this encoding is var-length,
how do we store it in the hash table?**

entries
111110

Topology
11 01 00

Indices
0 0 0 11110

Pointers
W Z Q X Y



Since this encoding is var-length,
how do we store it in the hash table?

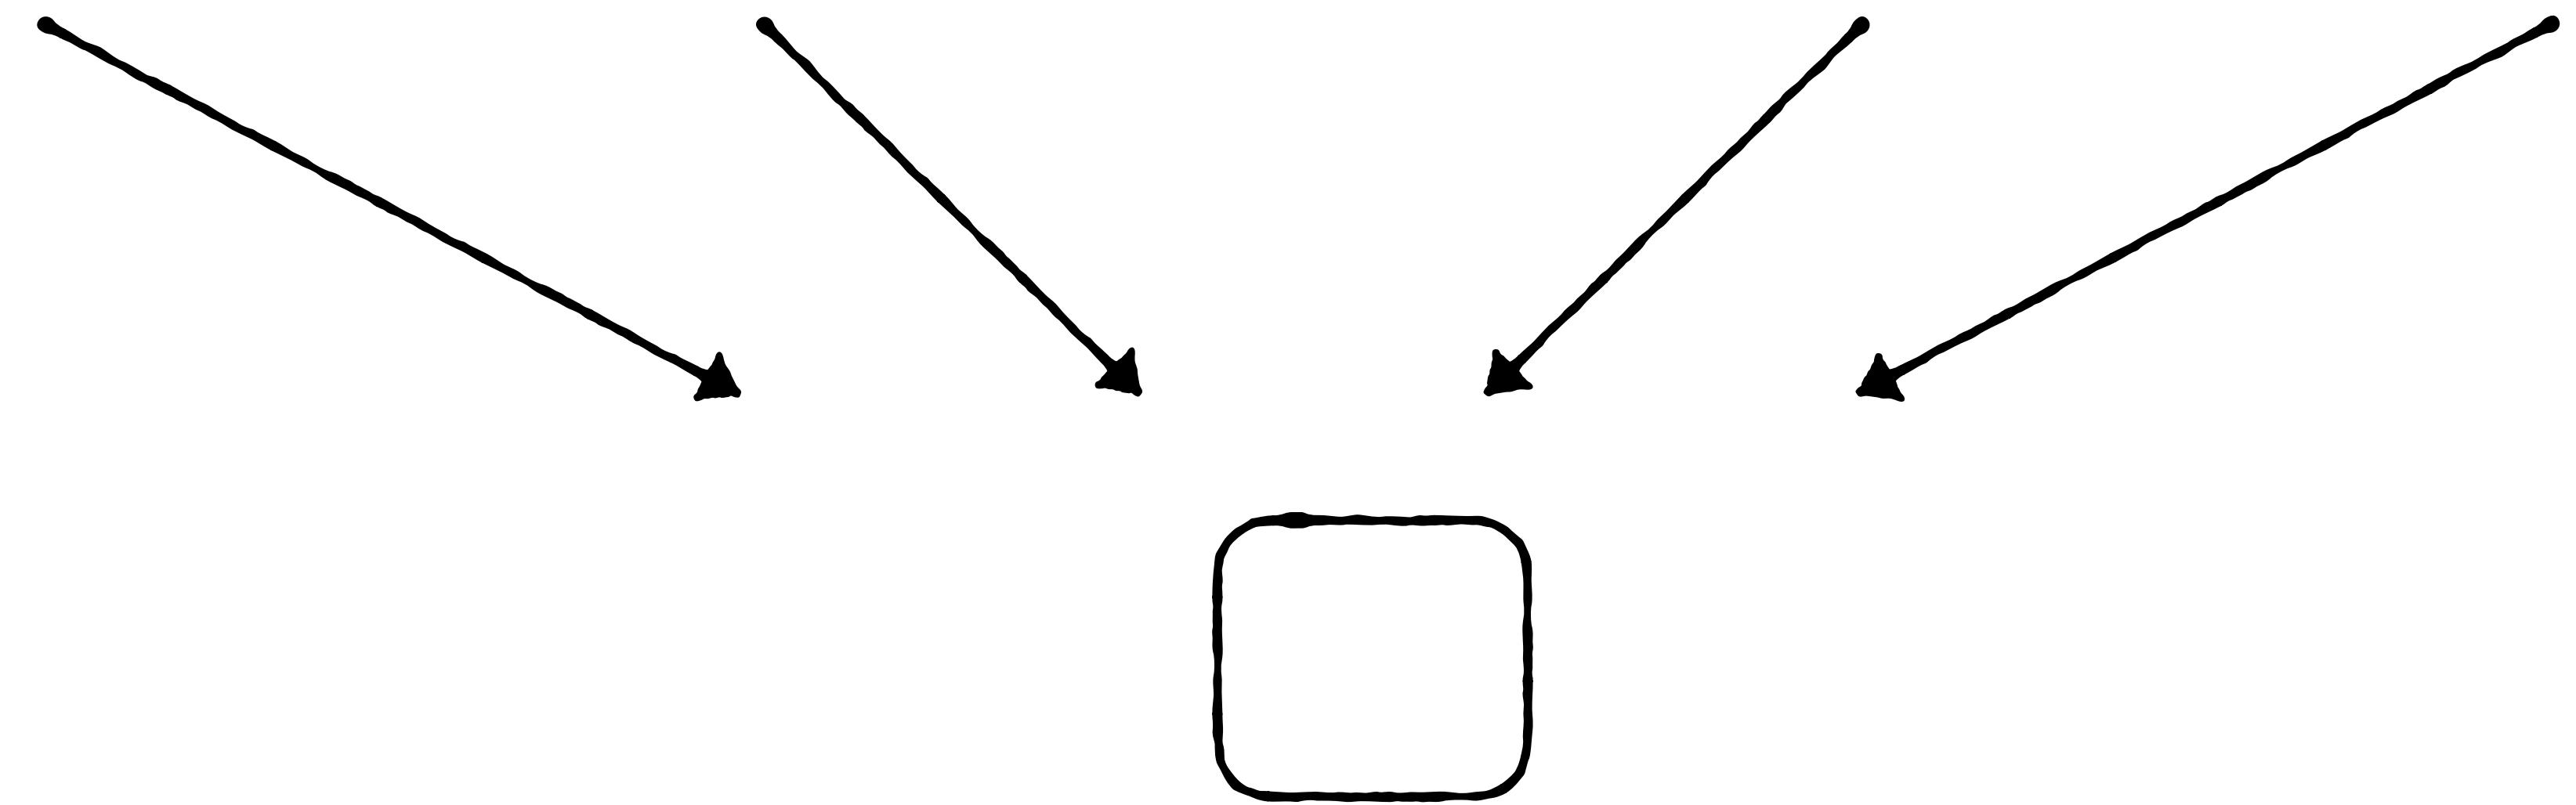
Note: all fields are self-delimiting

entries
111110

Topology
11 01 00

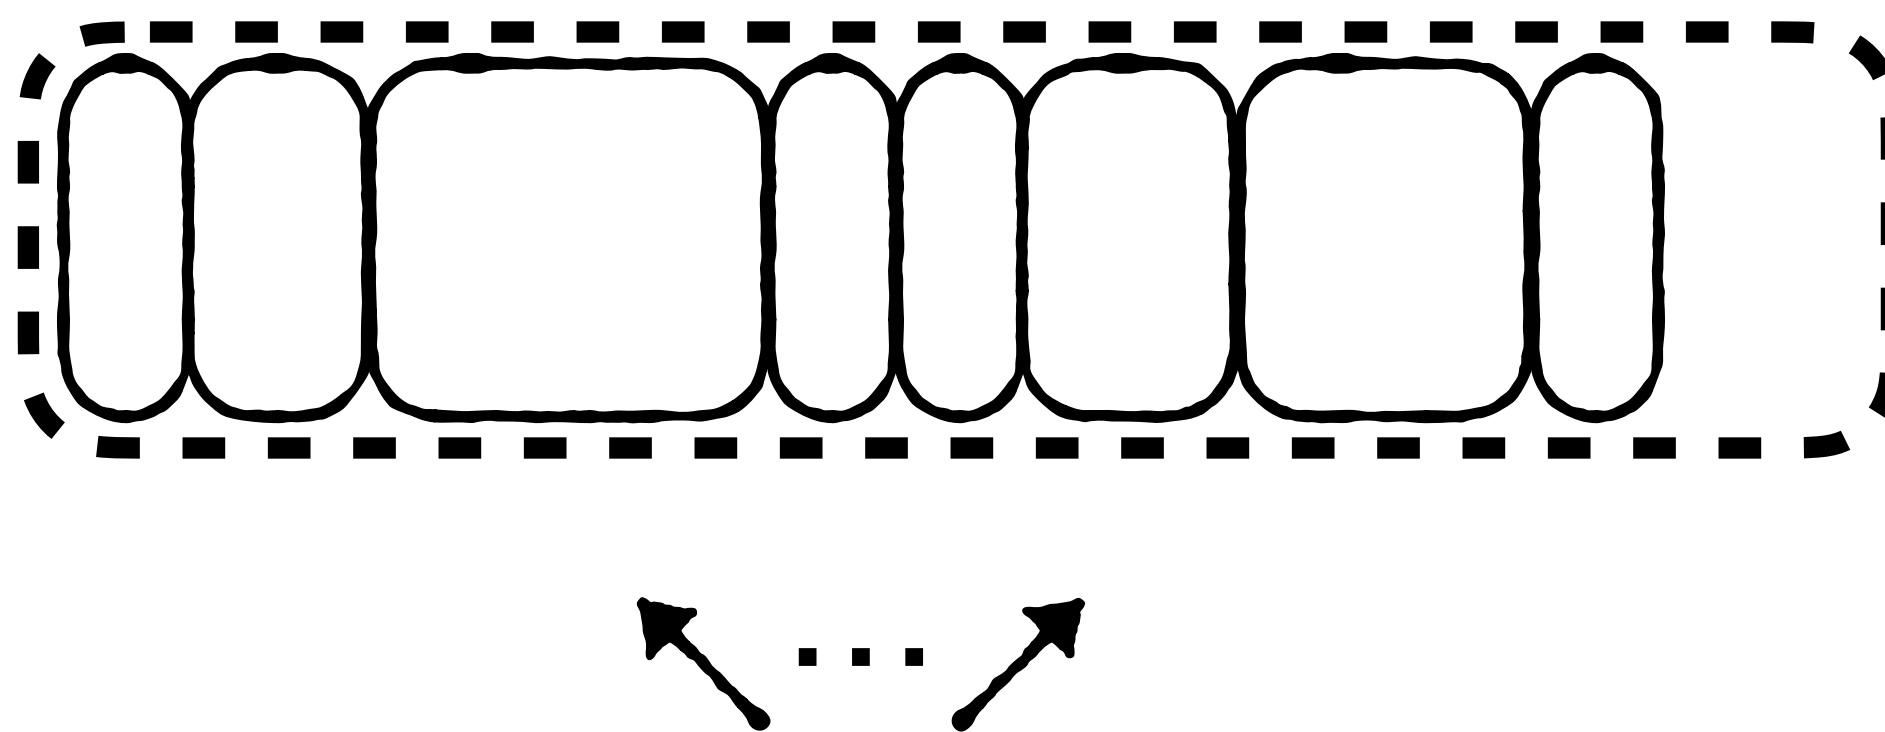
Indices
0 0 0 11110

Pointers
W Z Q X Y



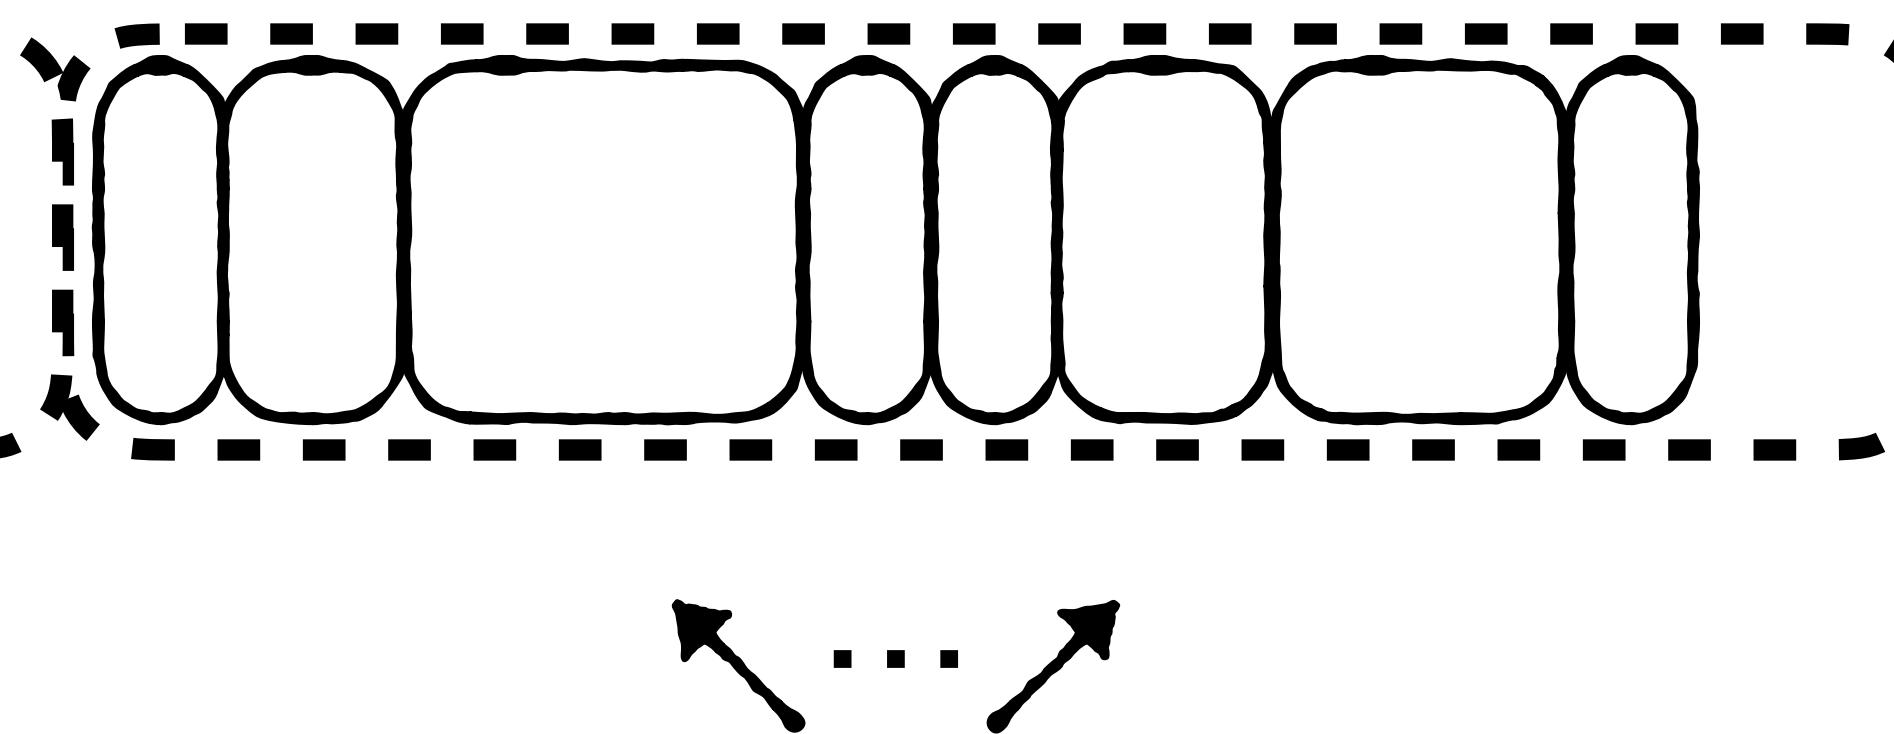
Variable-length slot

Fixed-sized block



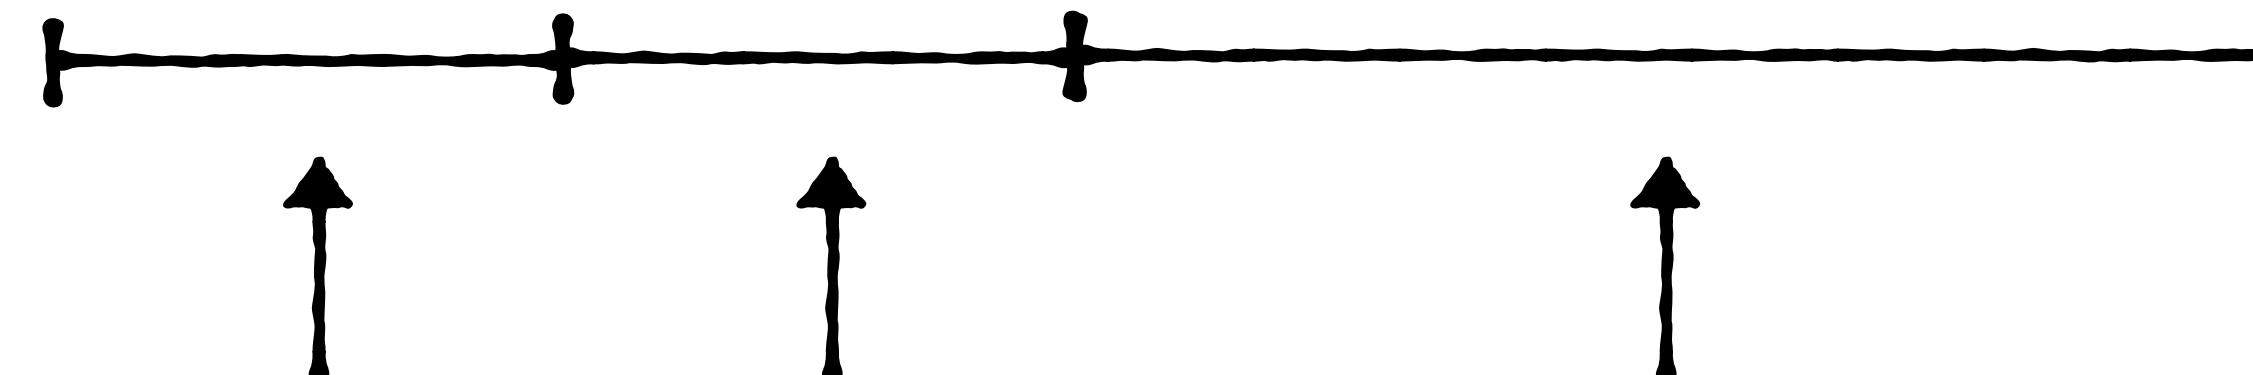
Multiple Variable-length slots

blocks

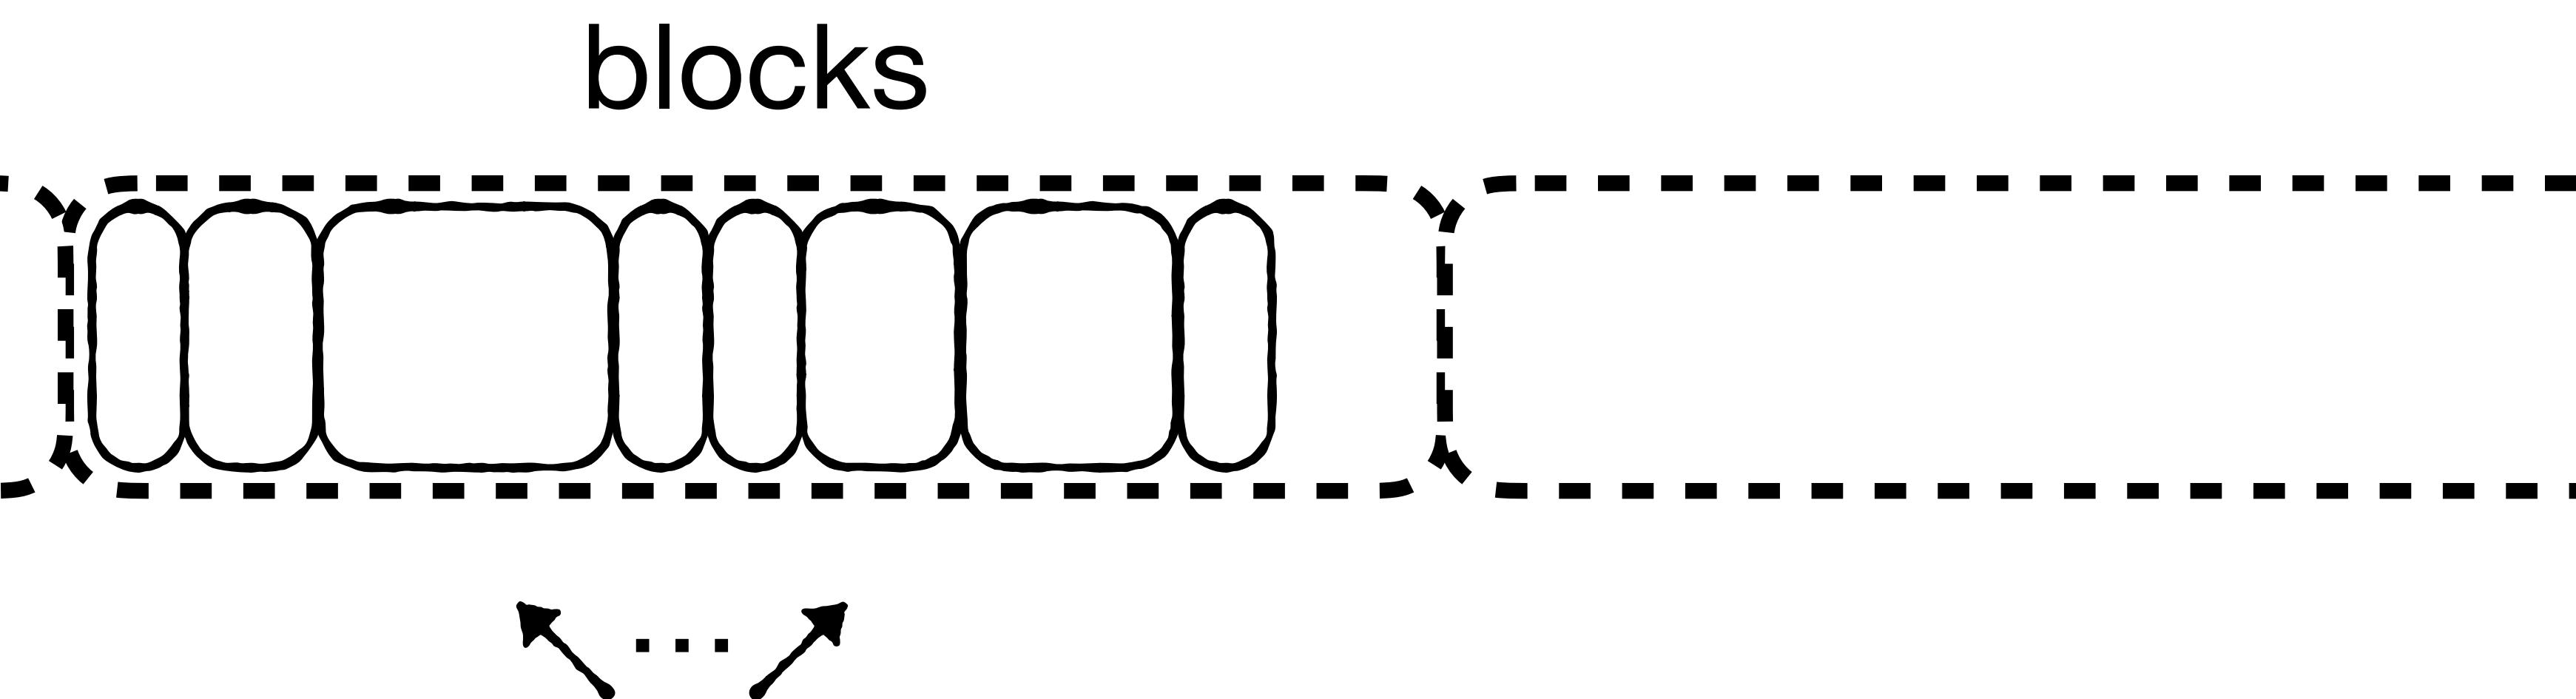


Multiple Variable-length slots

hash(X) = 0 1 0 1 0 0 1 1 0 1 0 1 1 0 0 ...



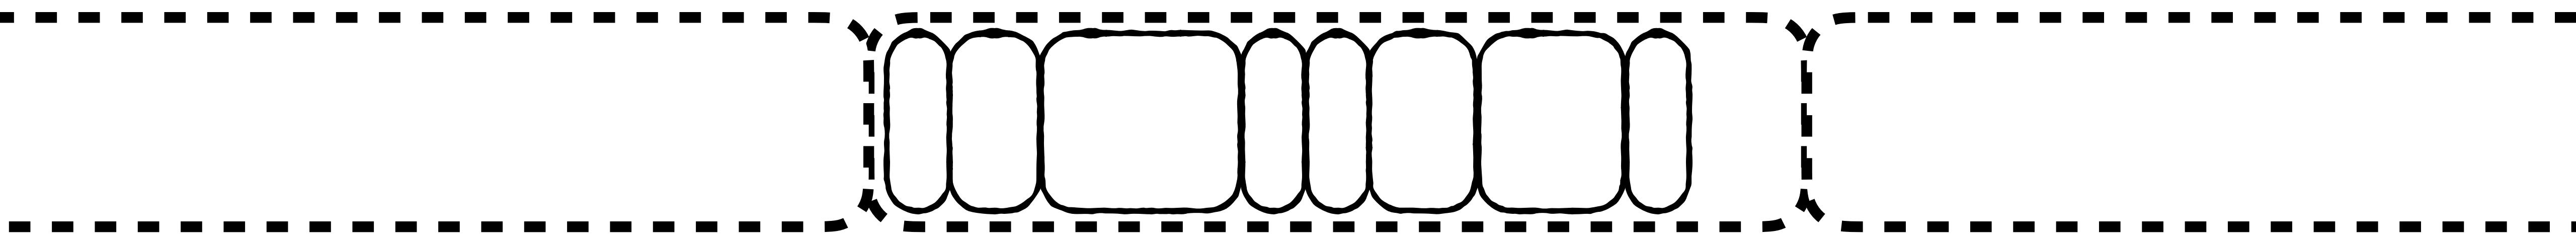
Block ID Slot ID Fingerprint



Multiple Variable-length slots

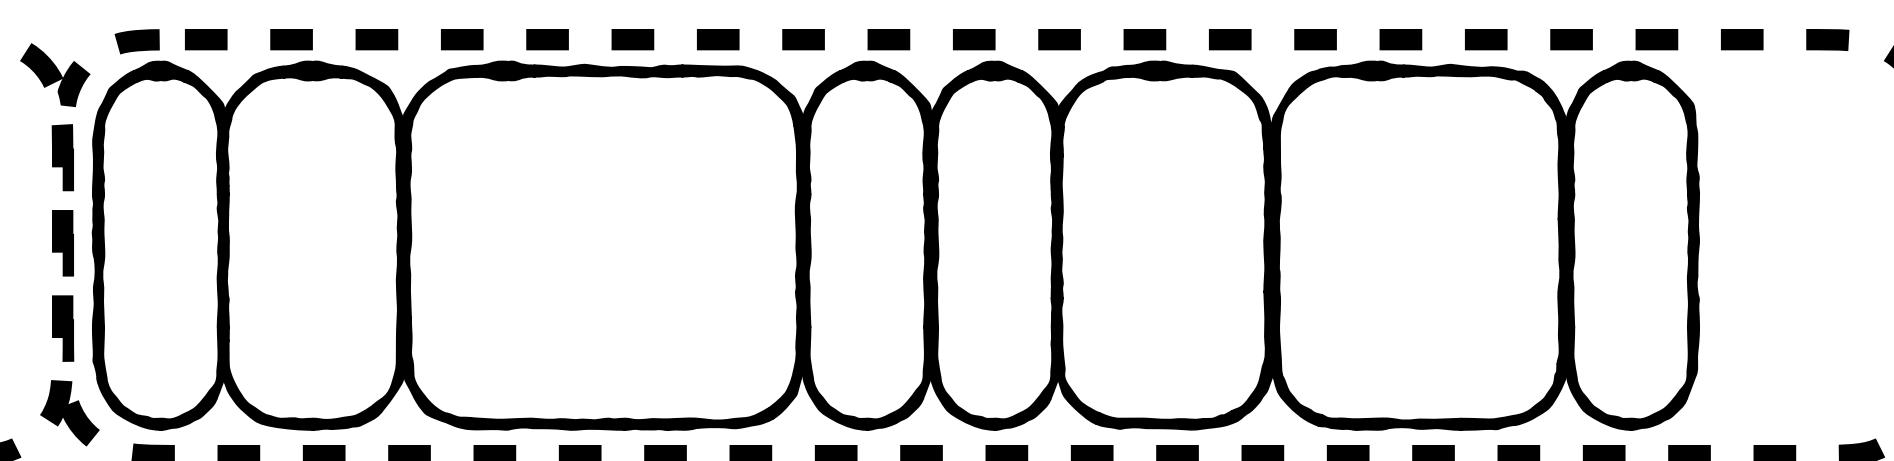
X

**Hash entry to some slot in
some block**



X

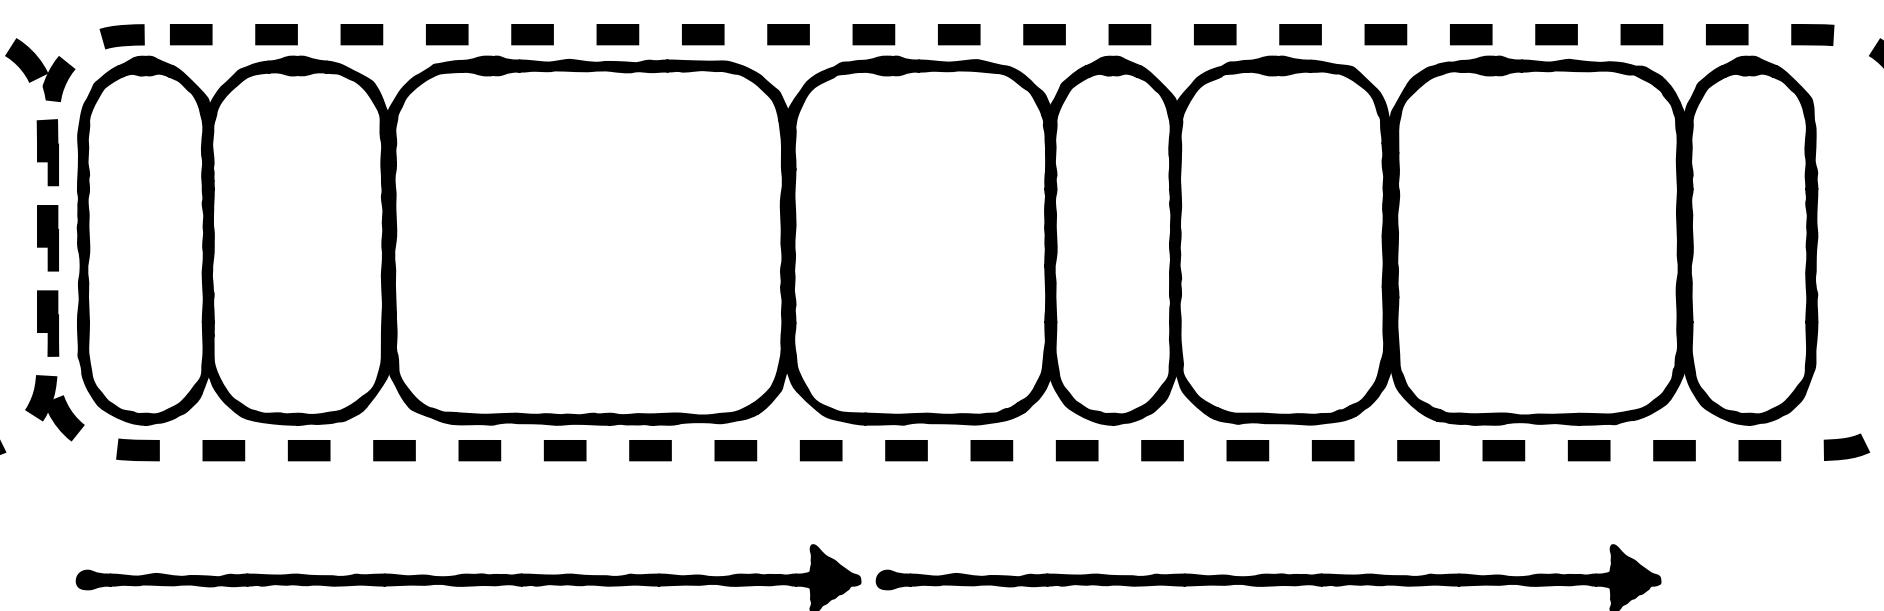
**Hash entry to some slot in
some block**



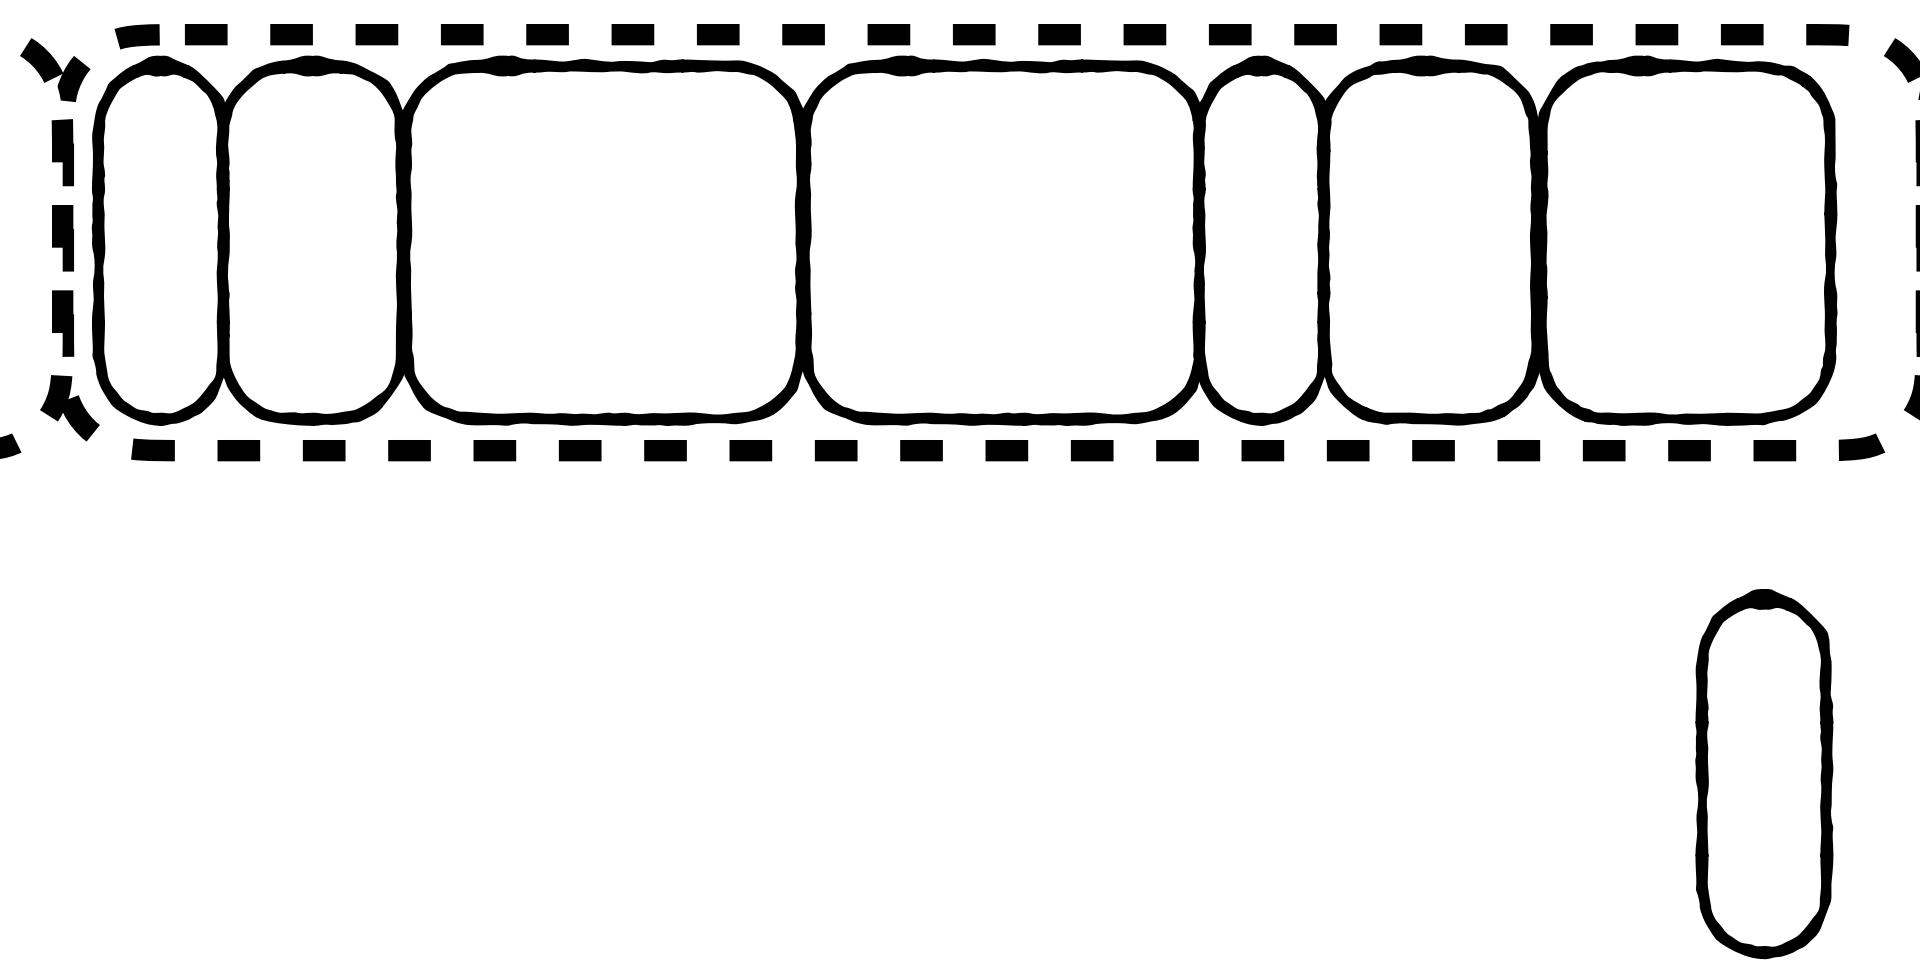
**Scan from
start of block**

X

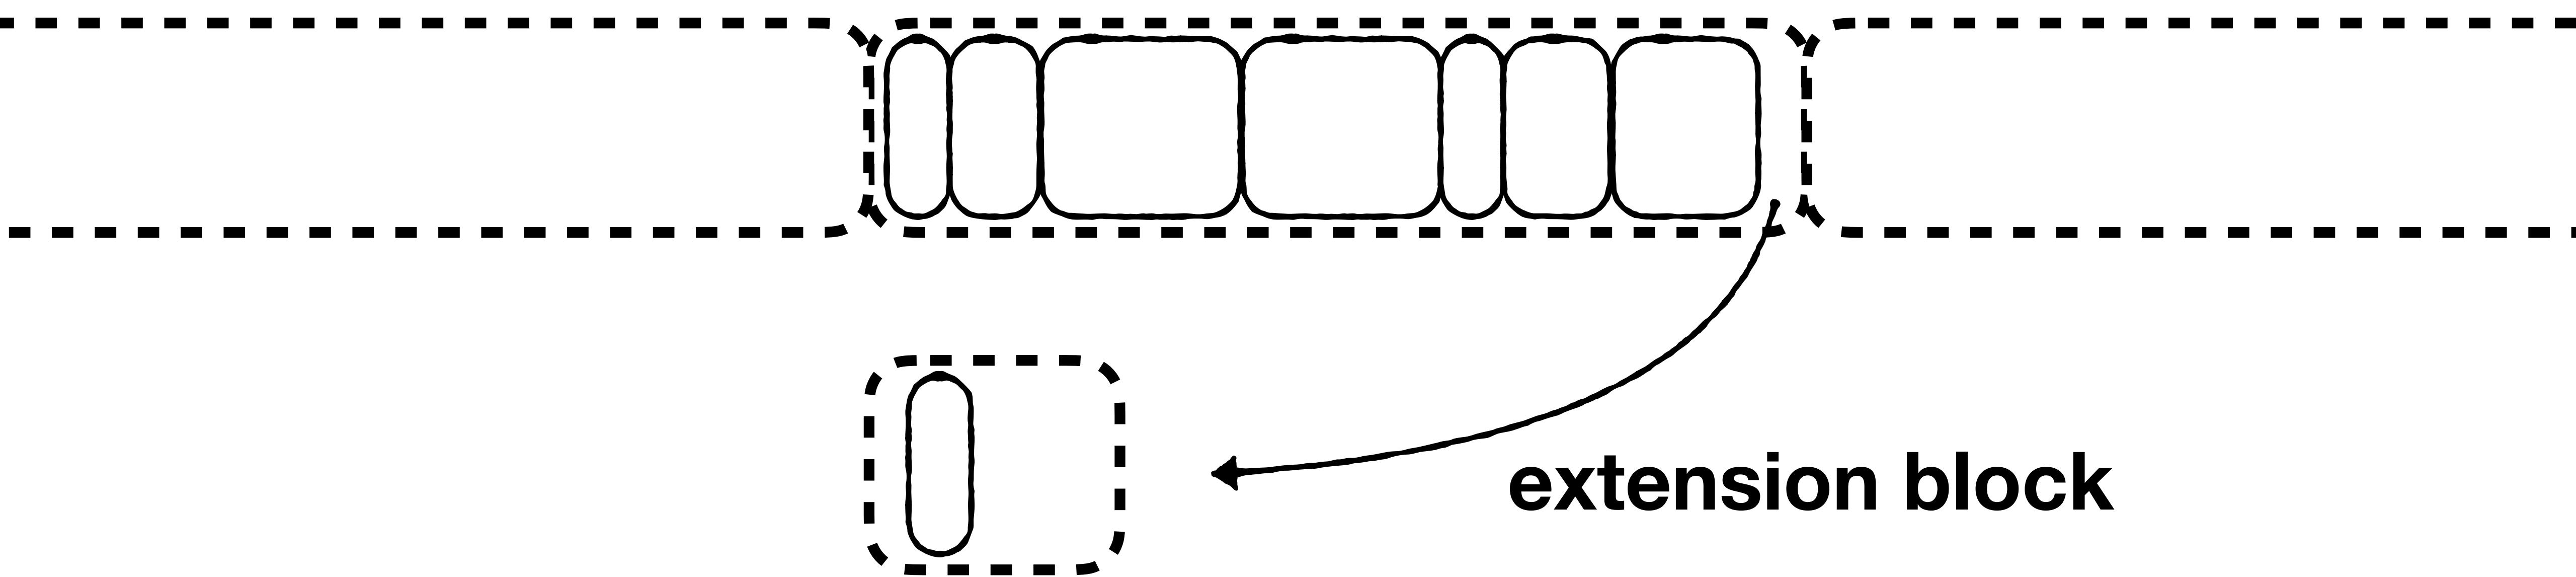
**Hash entry to some slot in
some block**



**insertion pushes all
other slots to right**



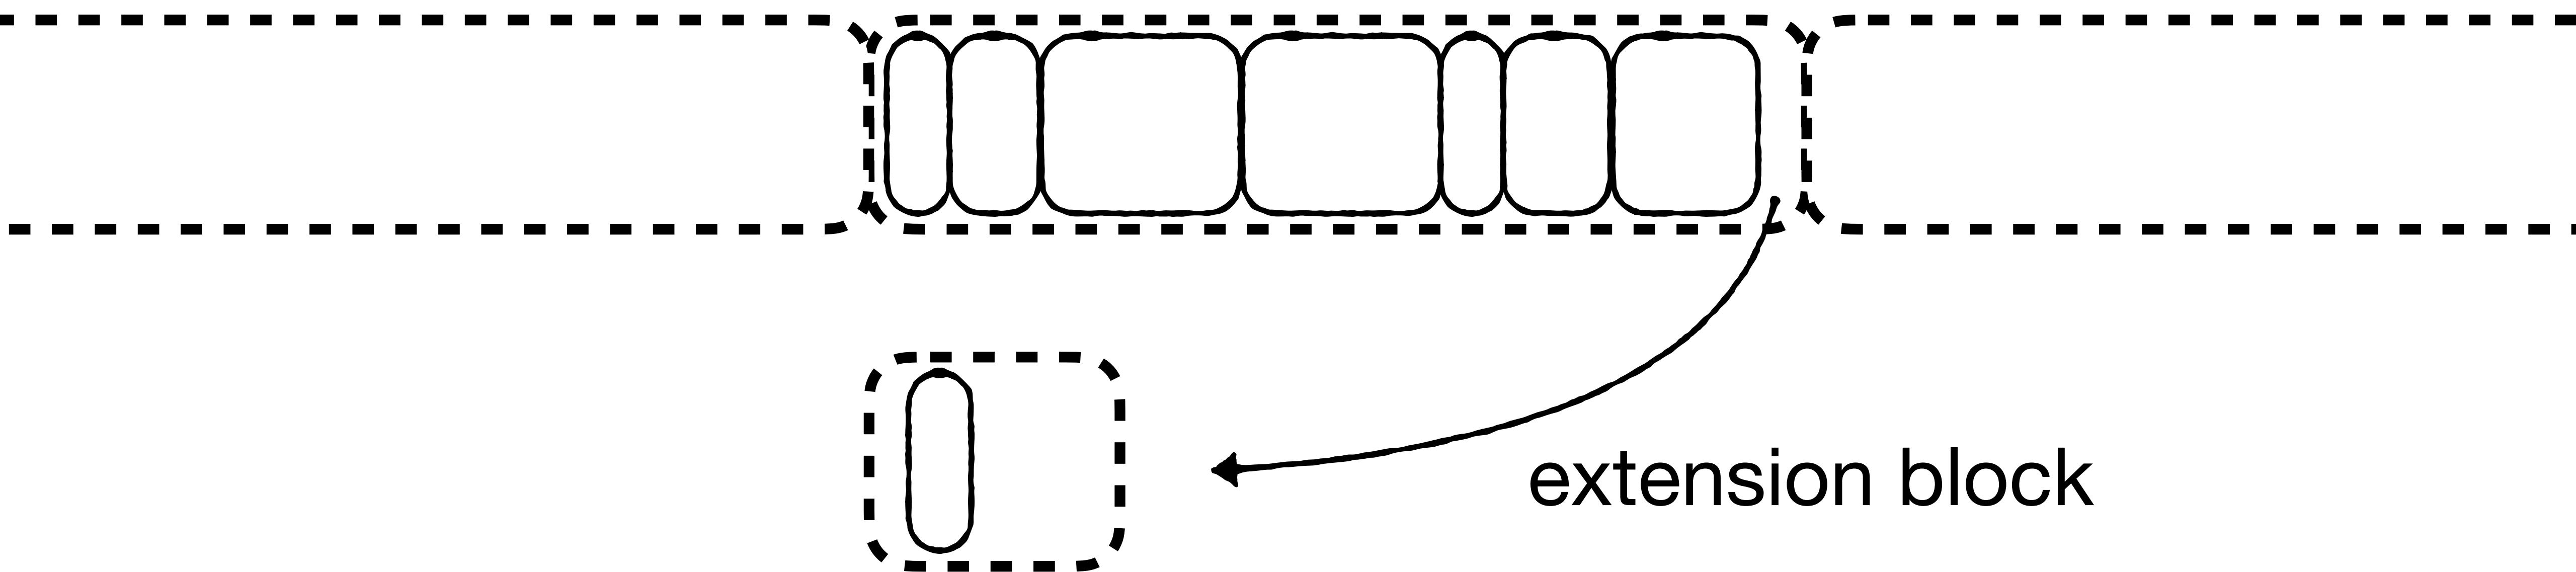
Overflow



With more slots per block...

(1)

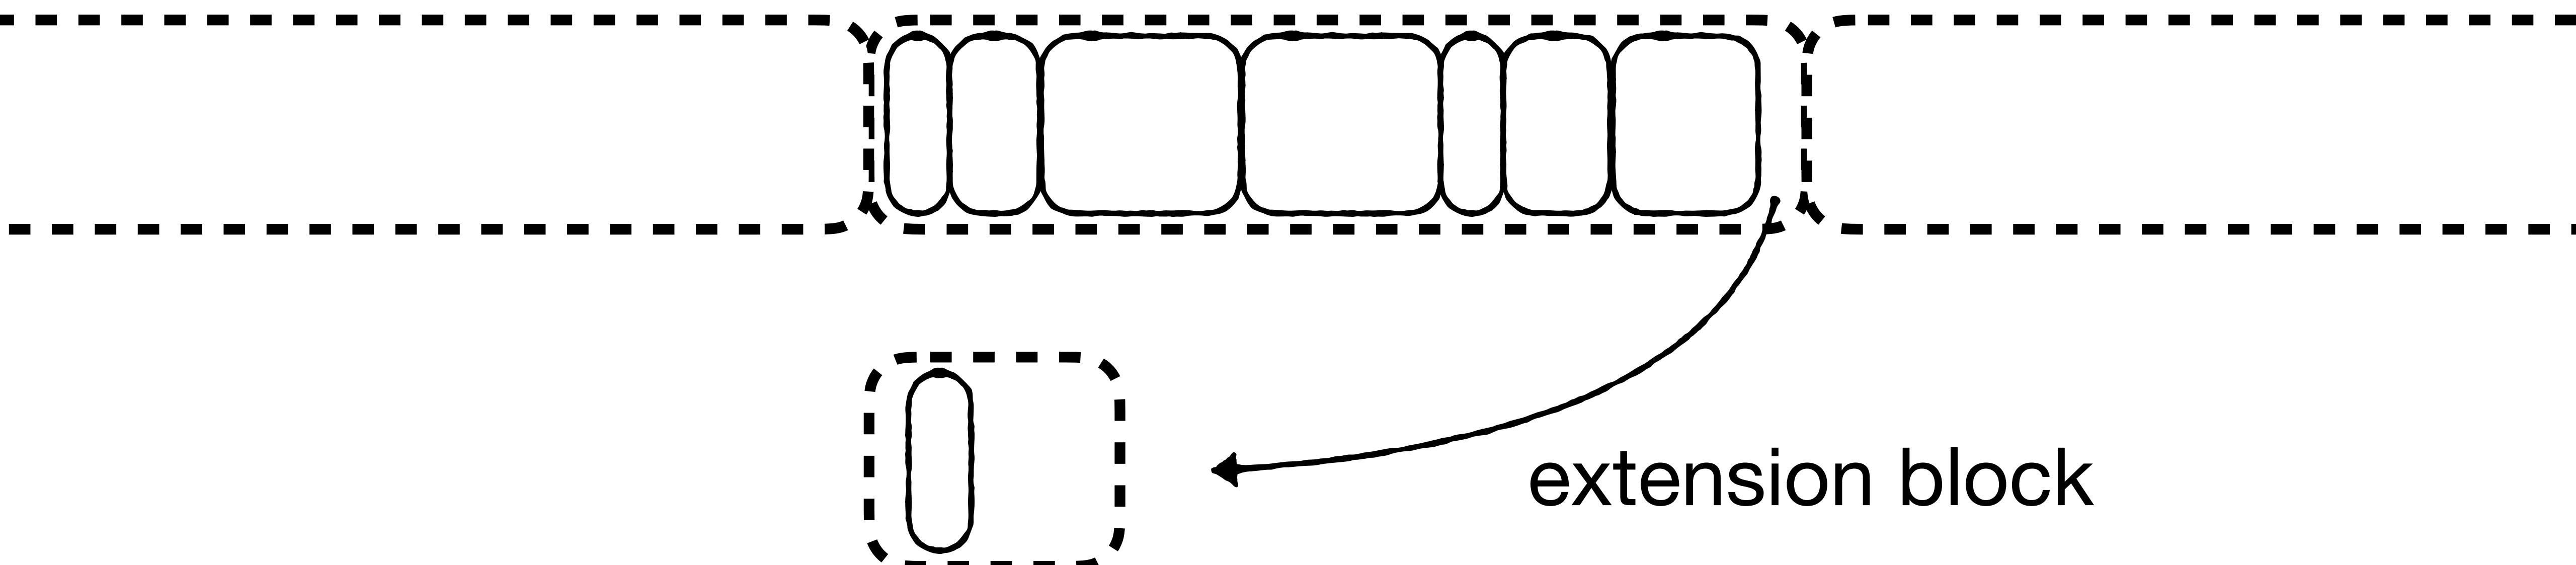
(2)



With more slots per block...

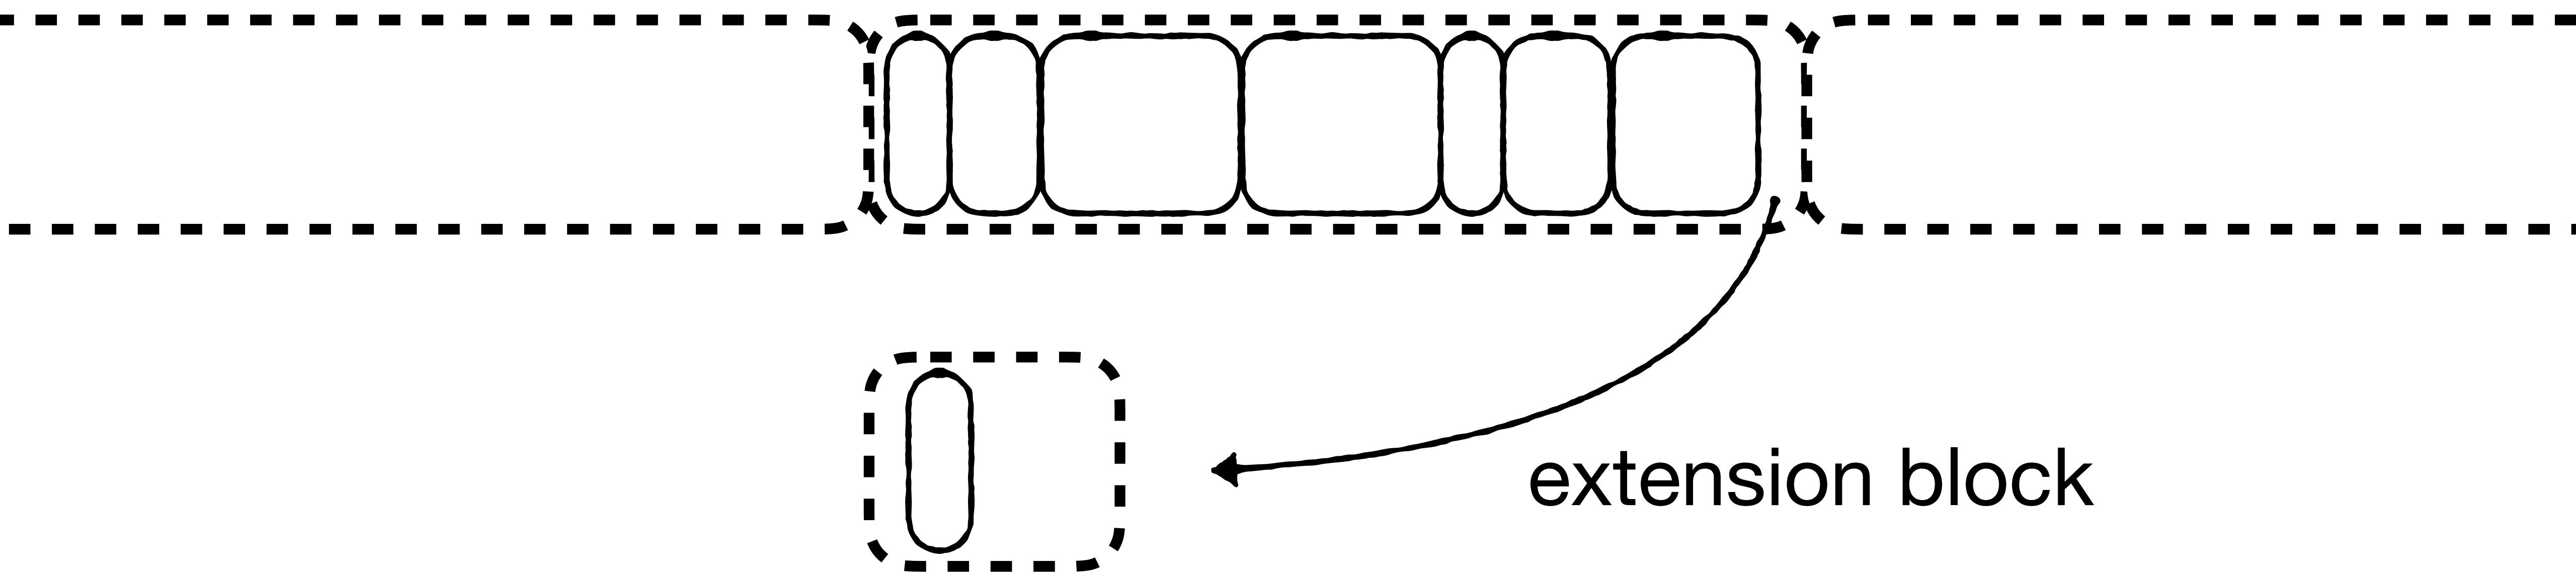
(1) Less size variability
-> fewer overflows

(2) More to traverse
for queries/inserts

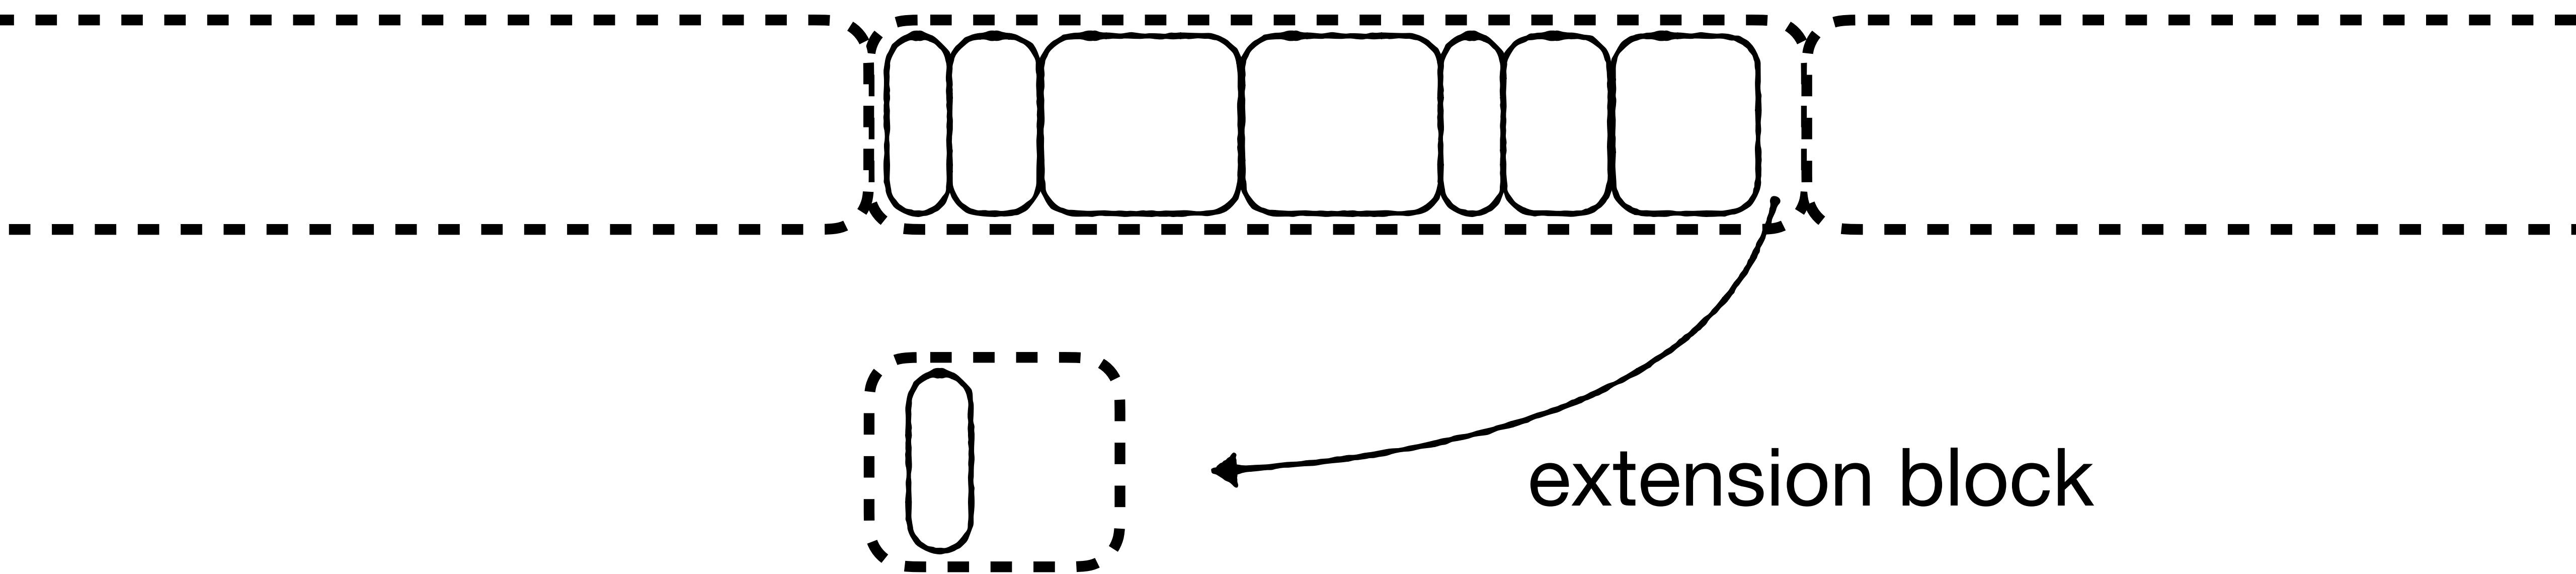


**The structure is not expandable and tuned
to support on avg. 1 entry per slot**

e.g., 64 slots per block



How large is each slot?



How large is each slot?

#entries

Topology

Indices

Pointers

#entries

Slot size X

Encoding

#Bits

#entries

Slot size X

Encoding

#Bits

0

0

1

Slot size X	Encoding	#Bits
0	0	1
1	10	2

#entries

Slot size X

Encoding

#Bits

0

0

1

1

10

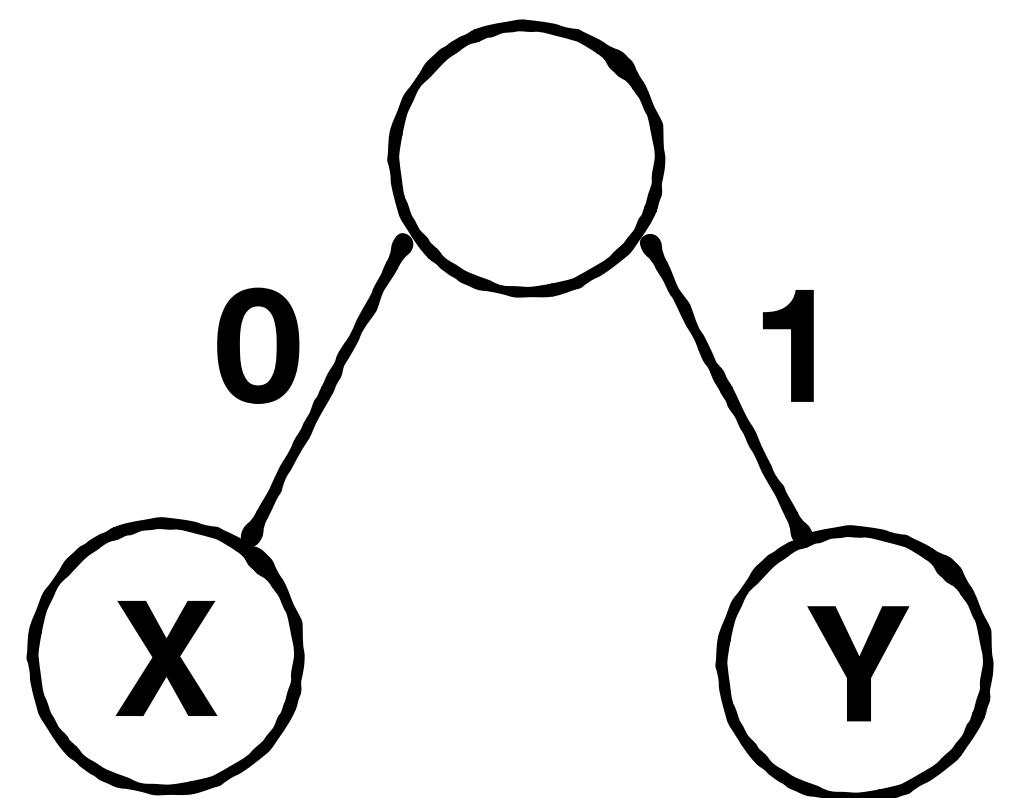
2

2

110

3

#entries



Slot size X

Encoding

#Bits

0

0

1

1

10

2

2

110

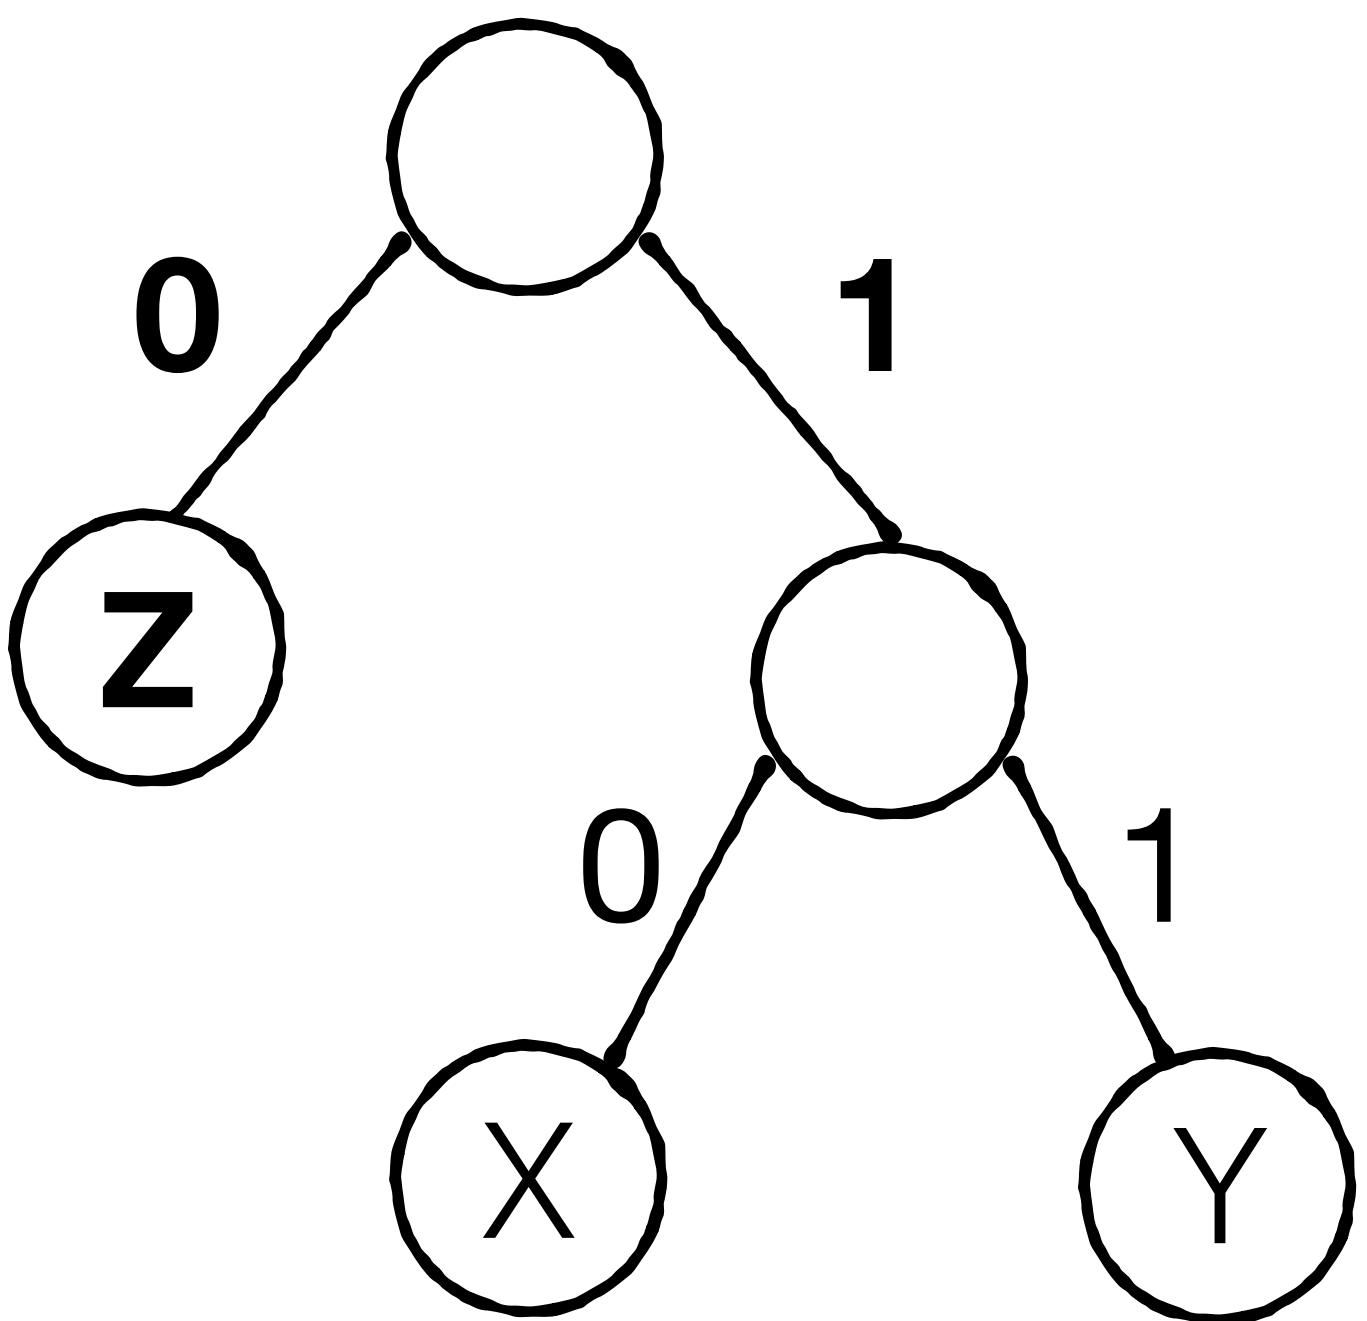
3

3

1110

4

#entries



Slot size X

Encoding

#Bits

0

0

1

1

10

2

2

110

3

3

1110

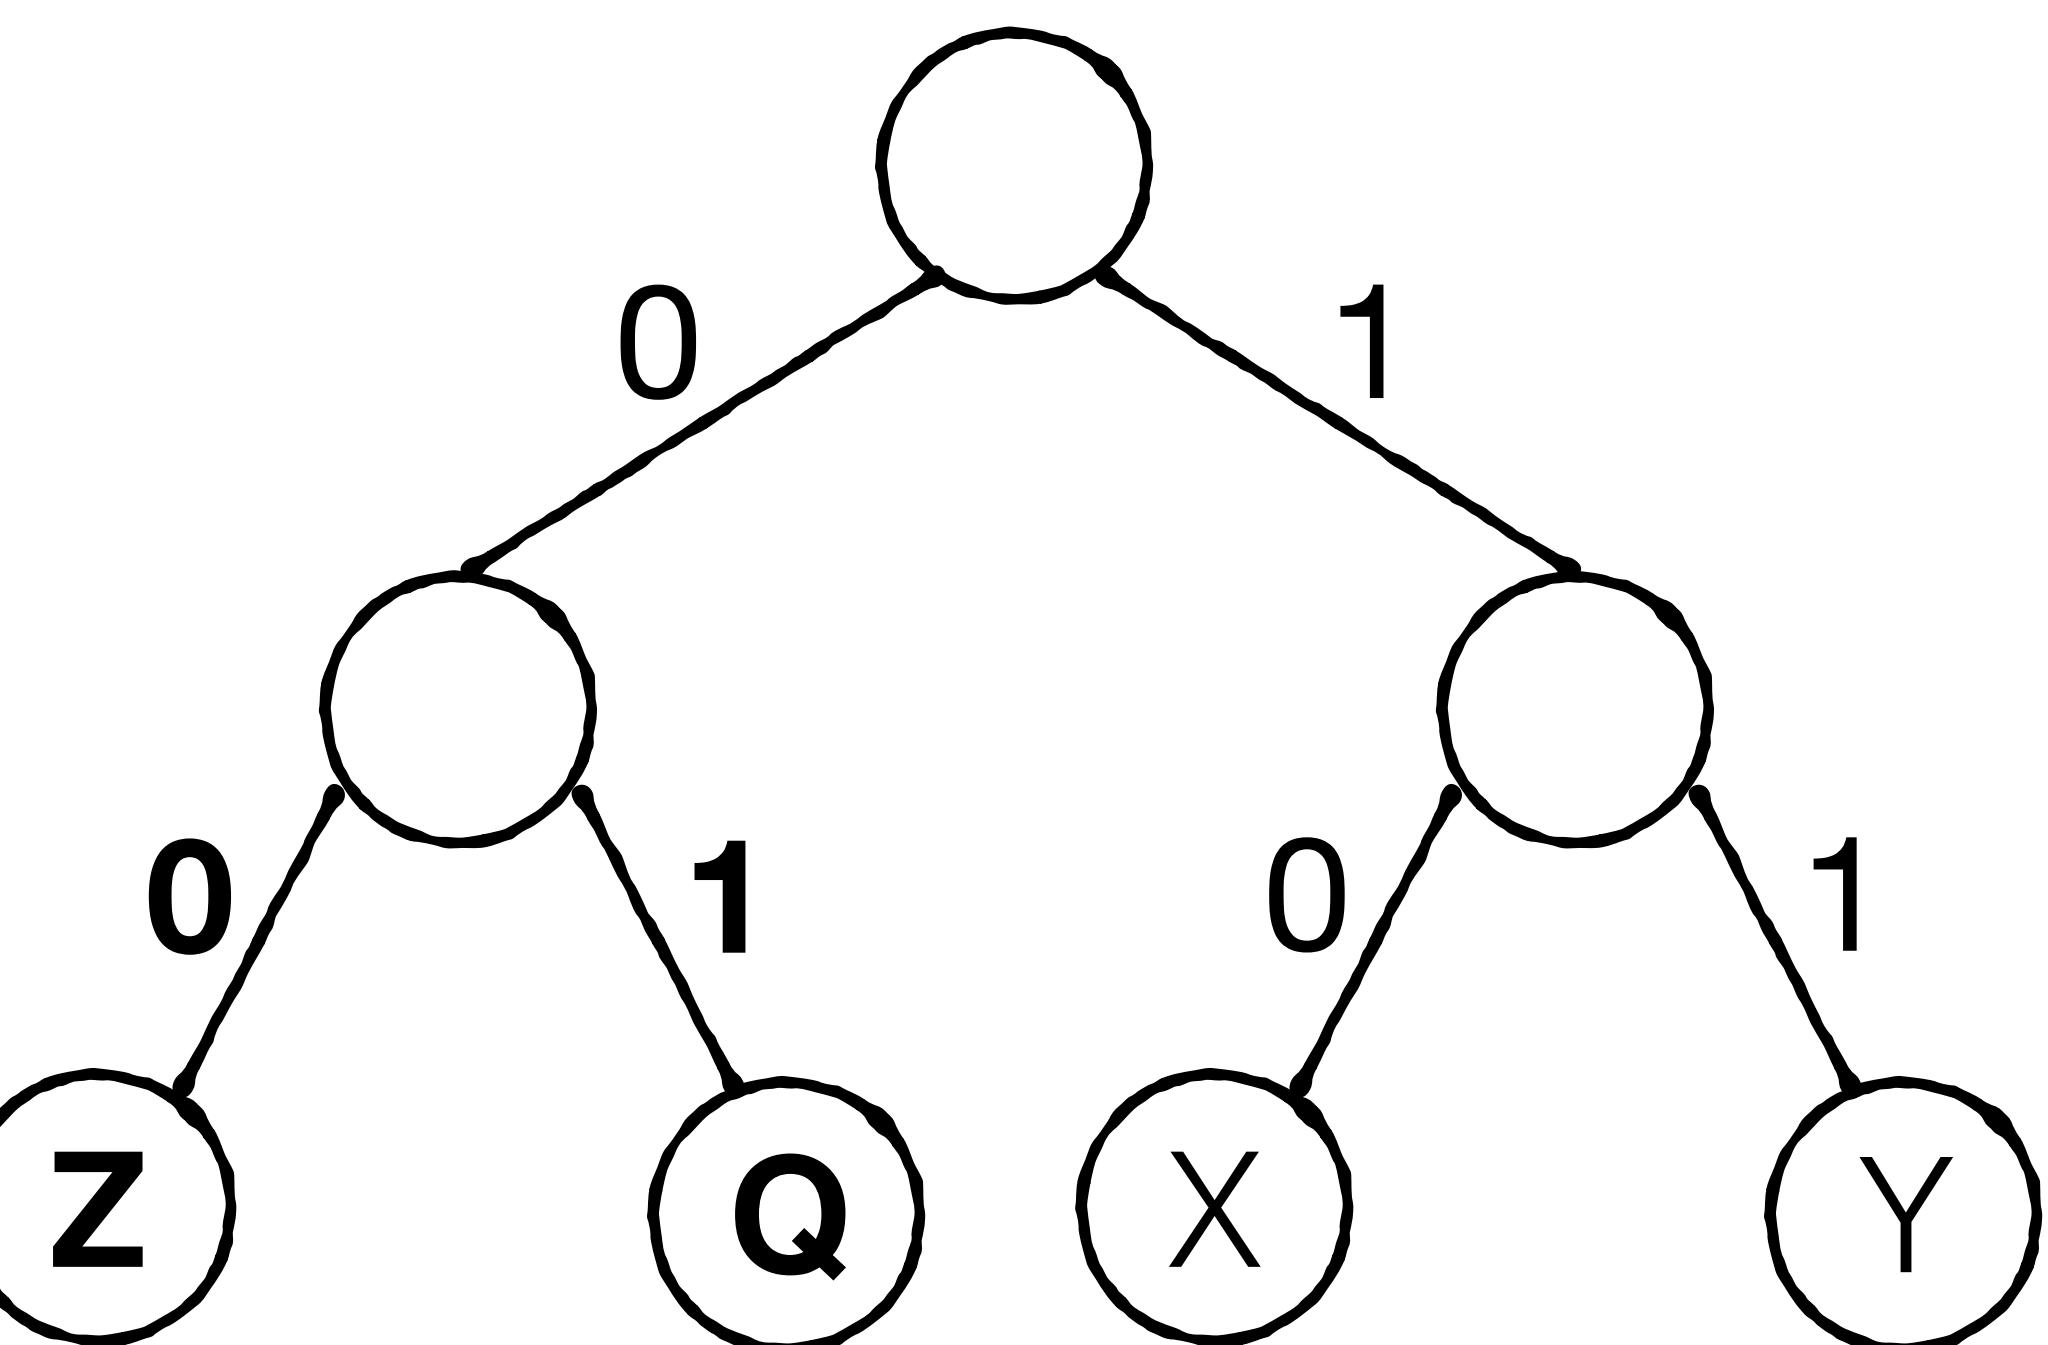
4

4

11110

5

#entries



Slot size X

Encoding

#Bits

0

0

1

10

2

10

3

110

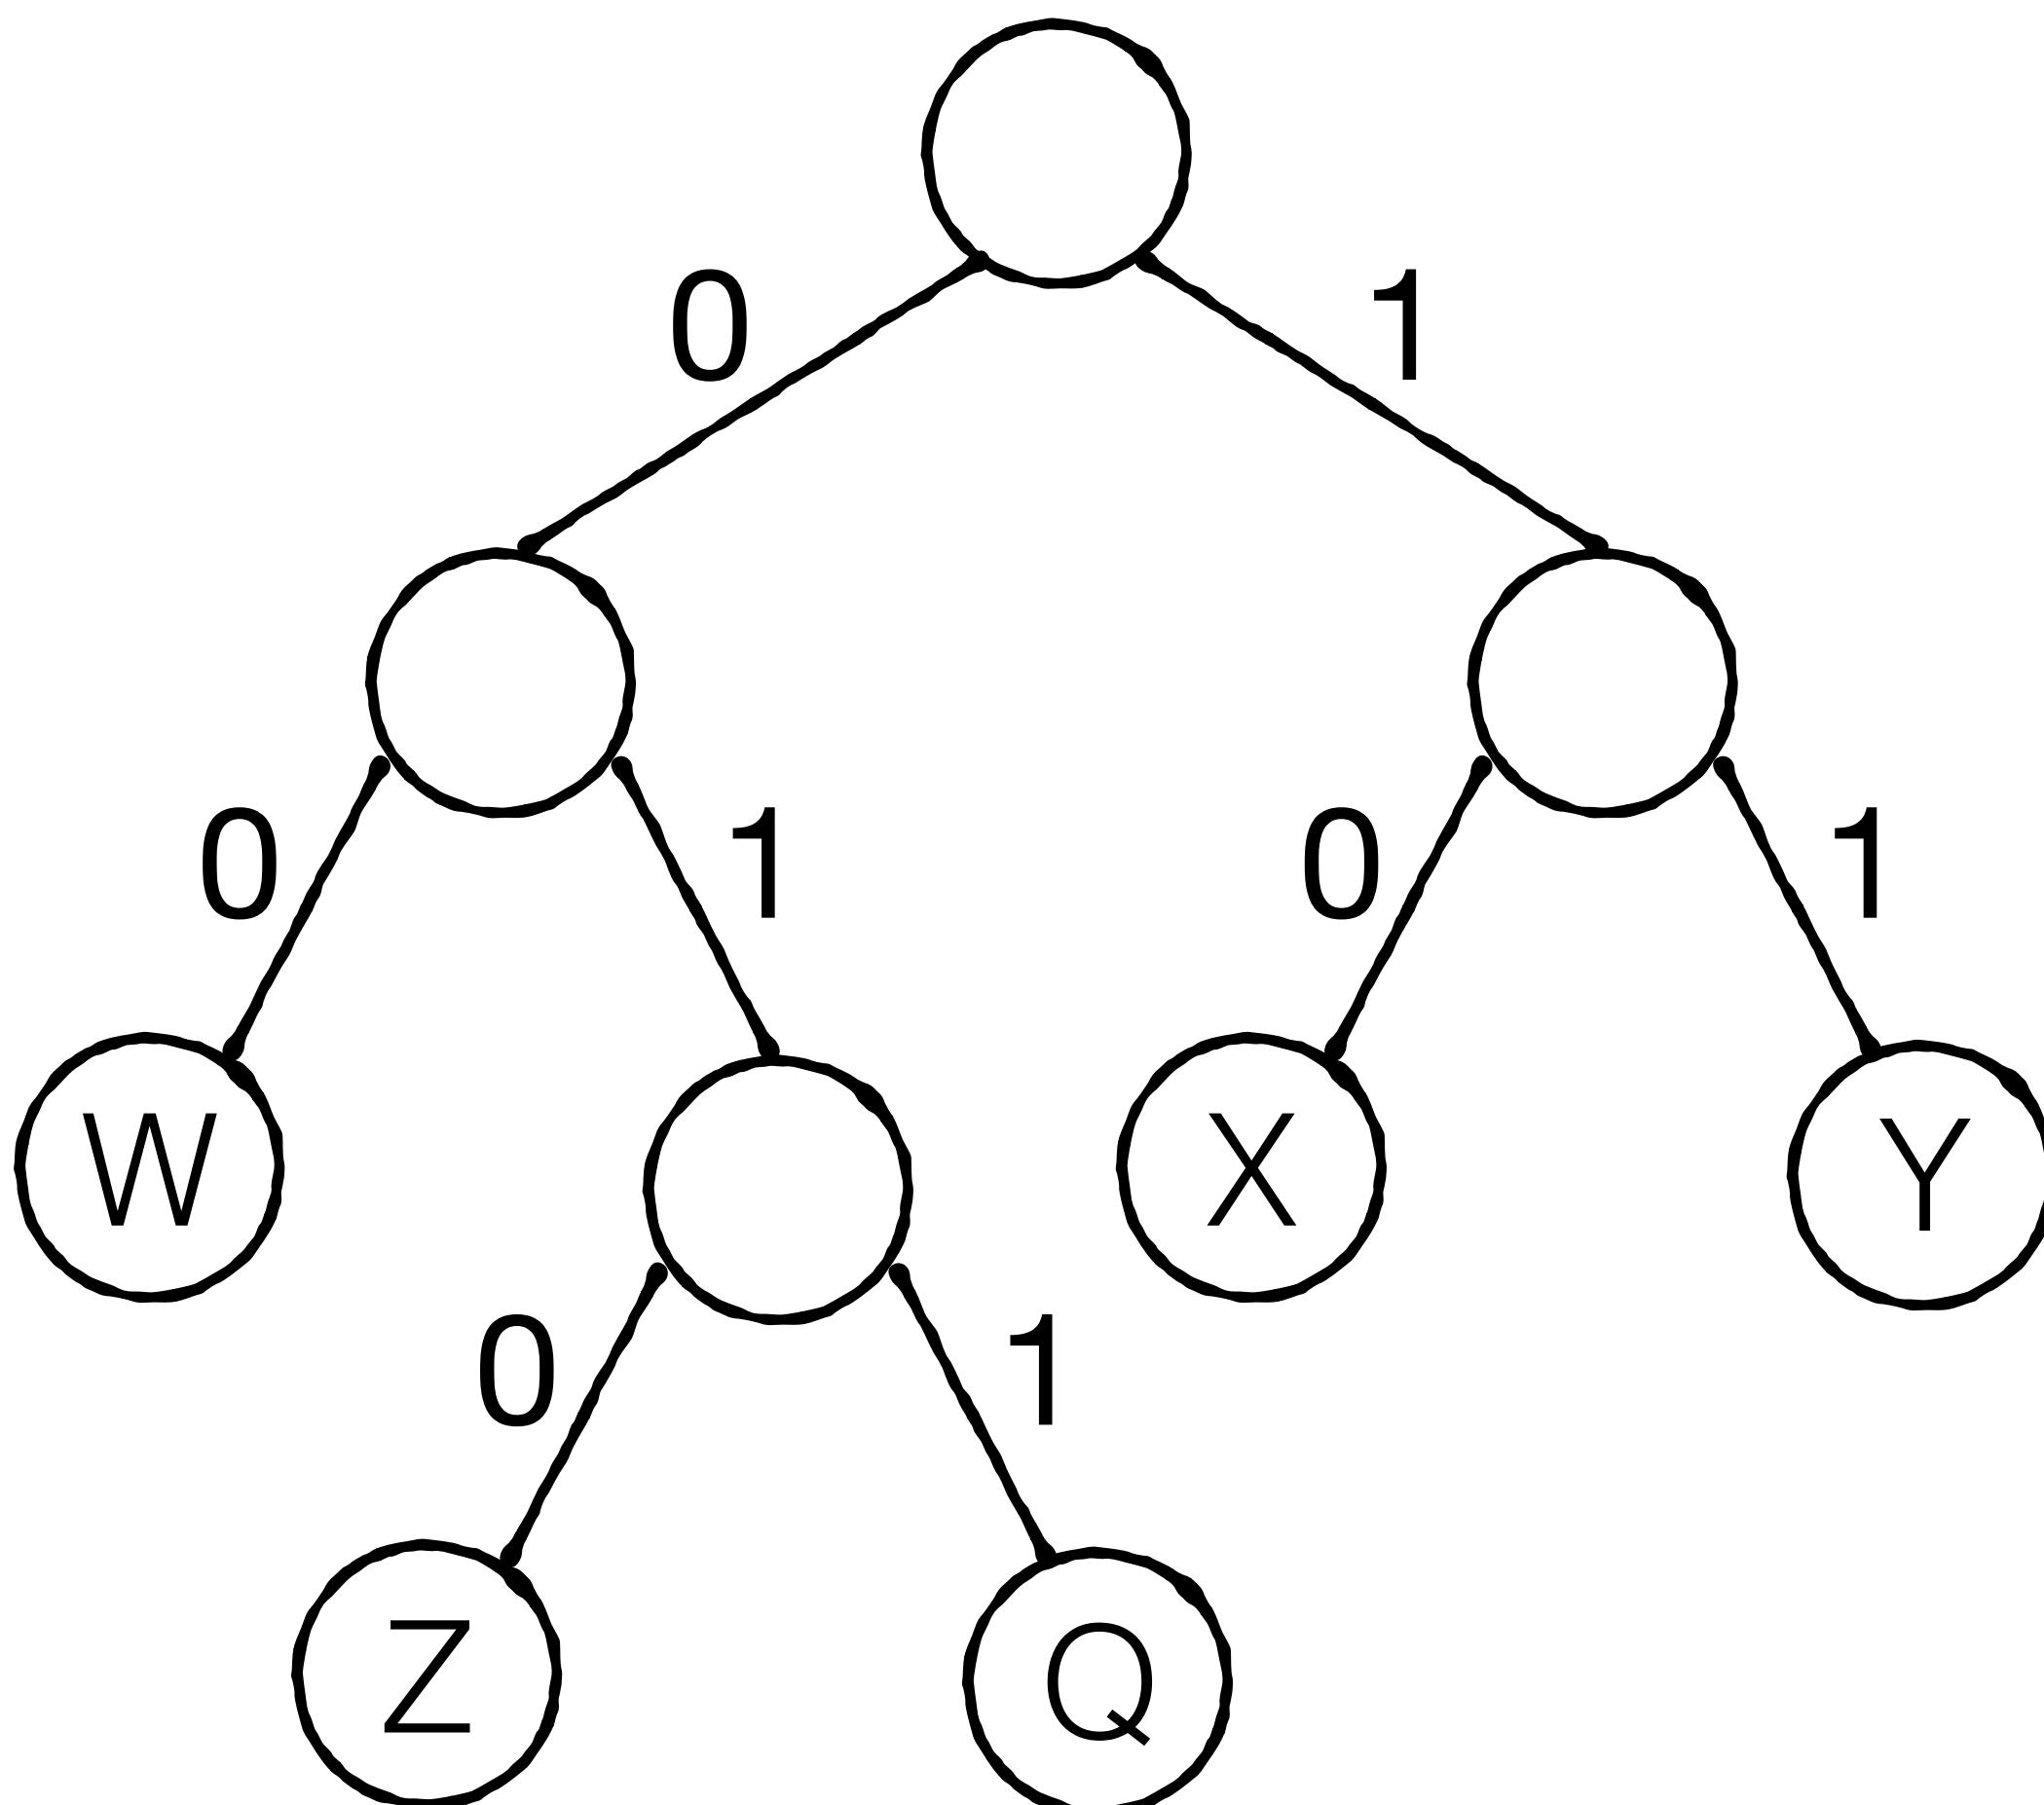
4

1110

5

1110

#entries



How large is each slot?

#entries

Topology

Indices

Pointers

Topology

Slot size X **Encoding** **#Bits**

0

-

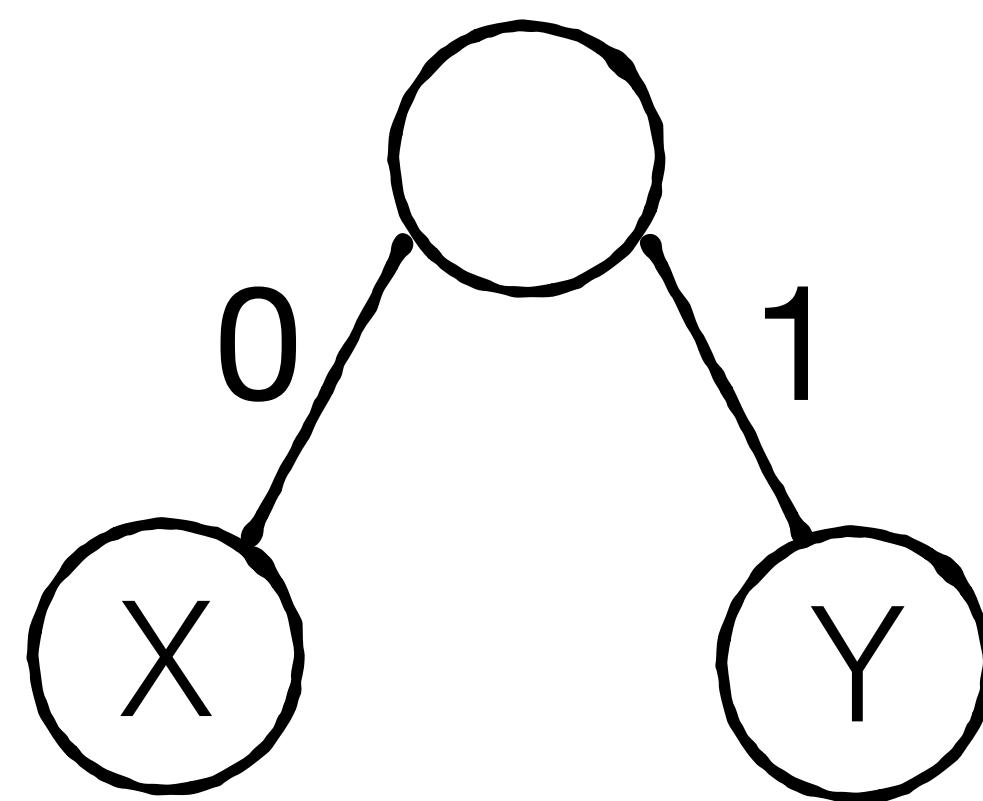
0

Topology

Slot size X	Encoding	#Bits
0	-	0
1	-	0

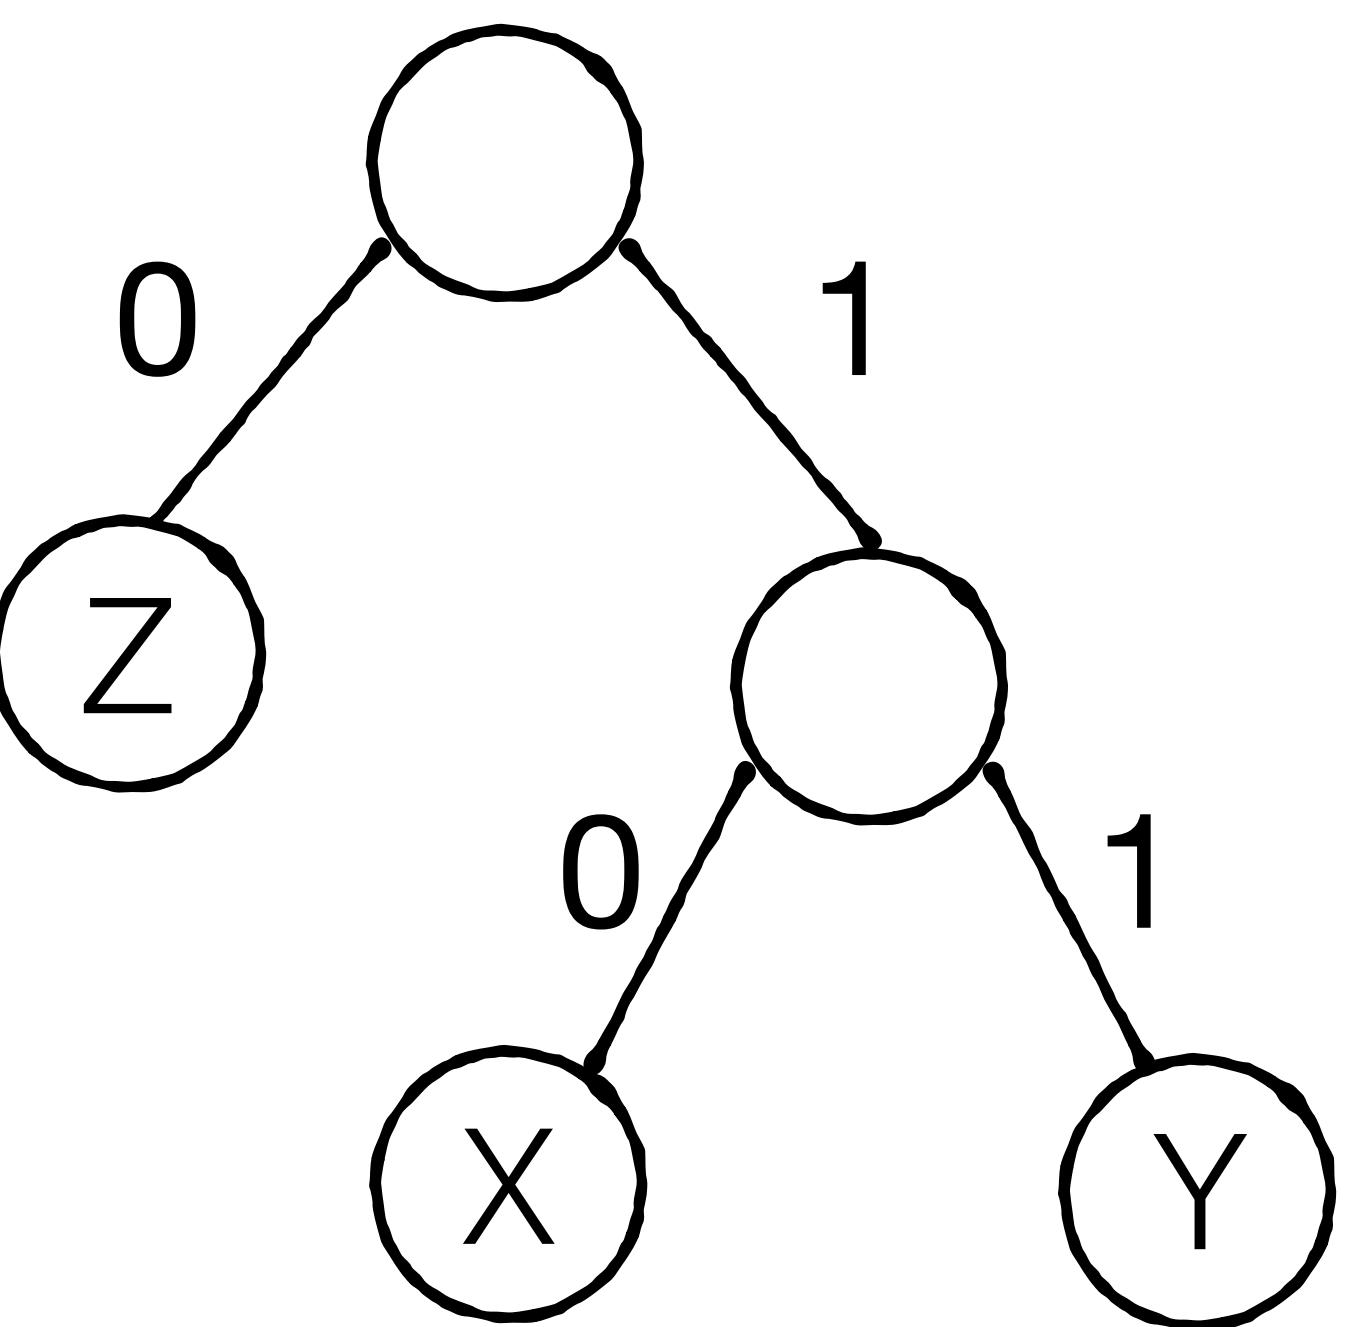
Topology

Slot size X	Encoding	#Bits
0	-	0
1	-	0
2	00	0



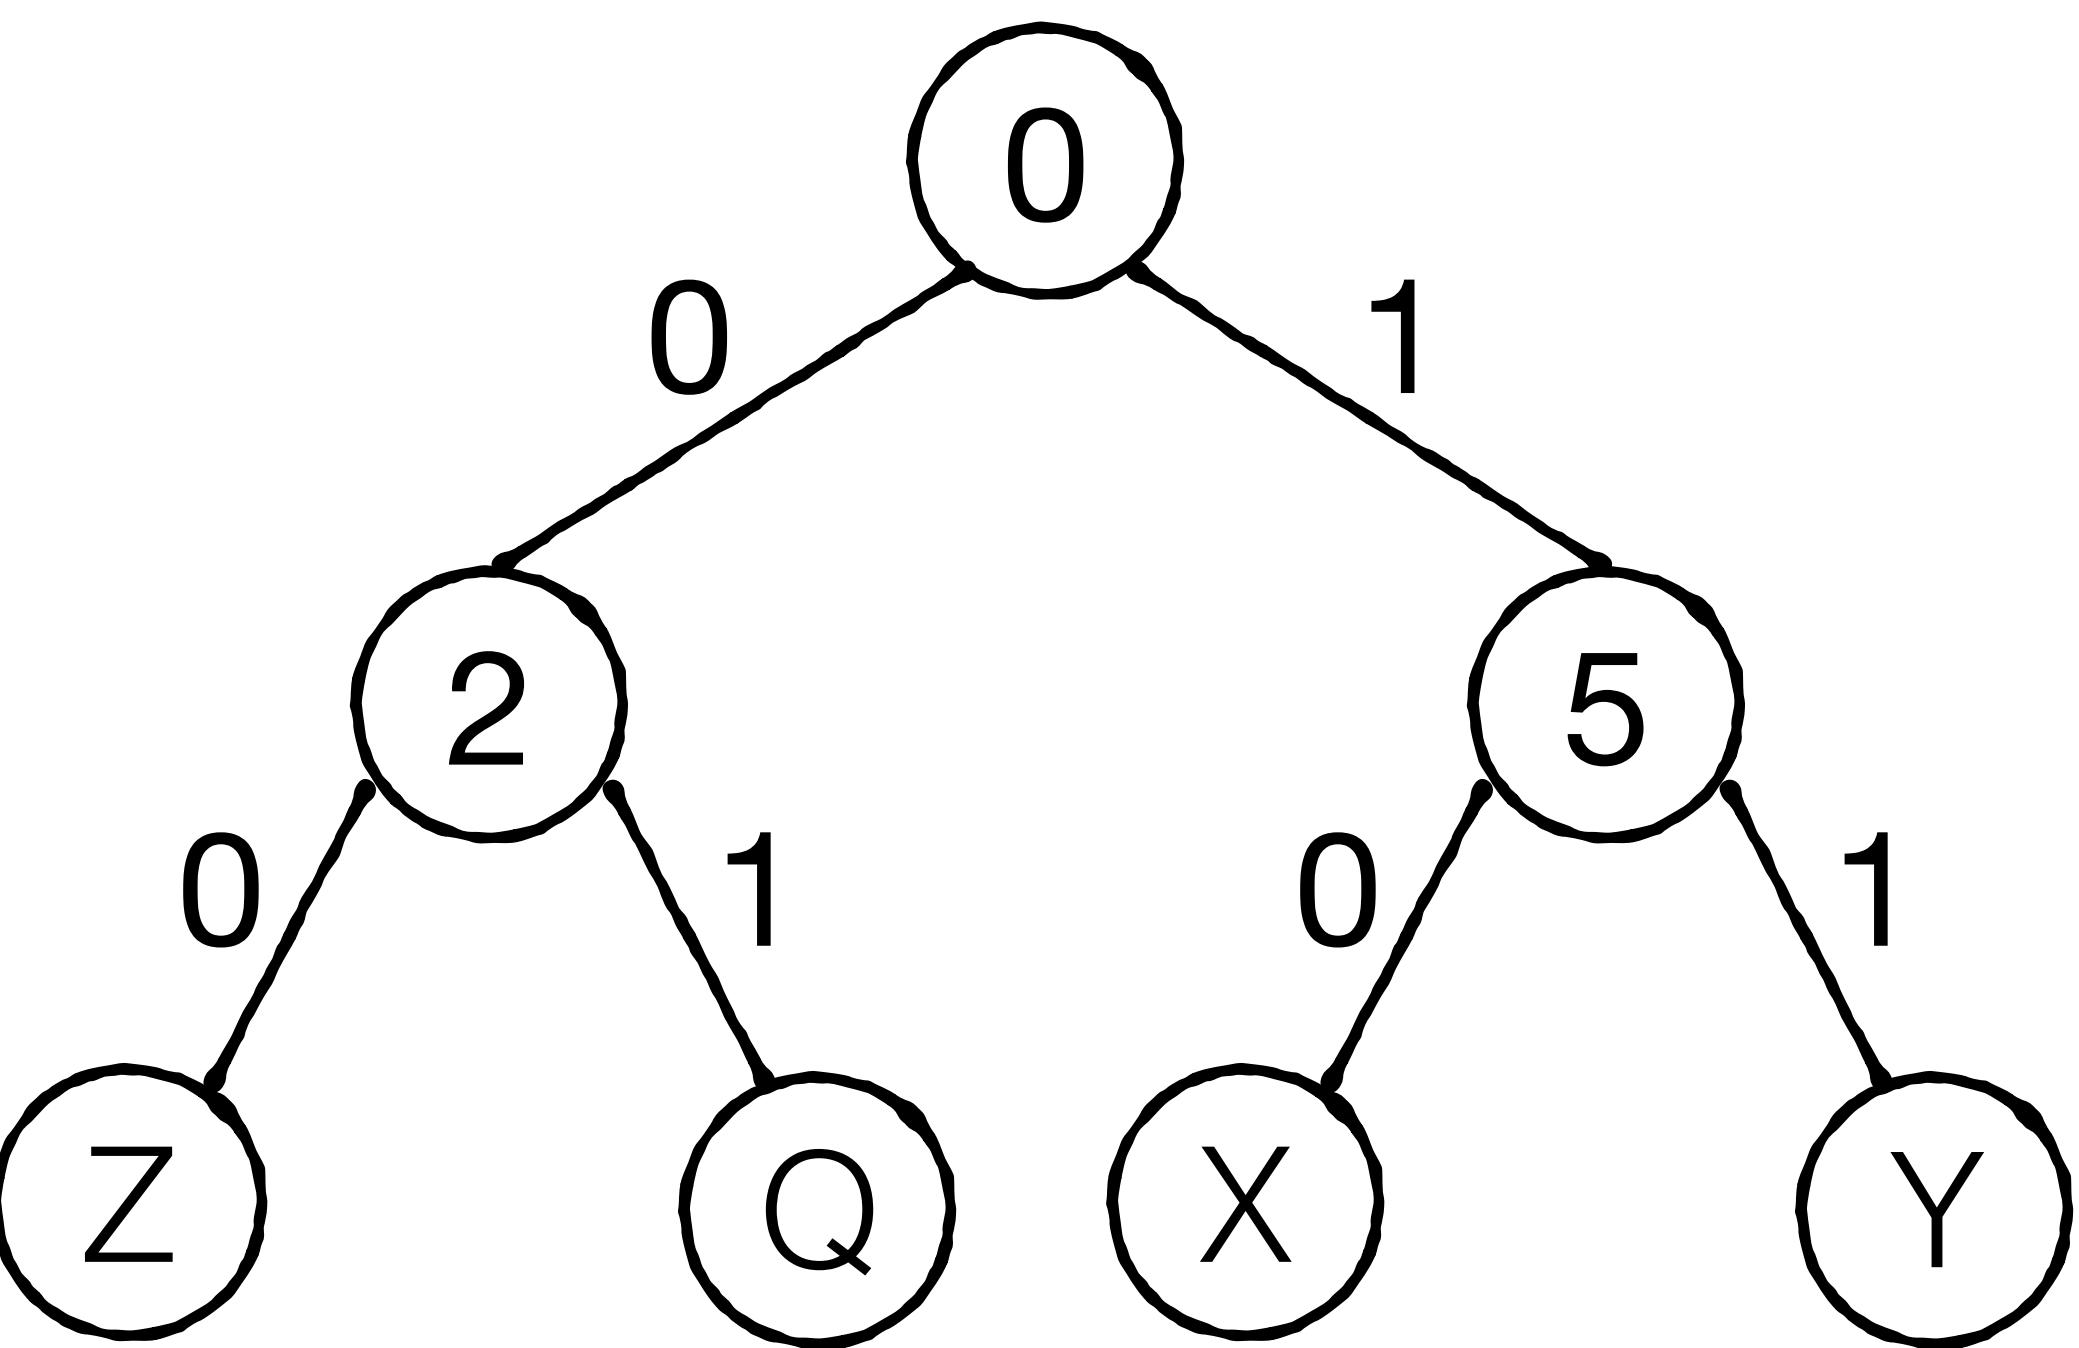
Topology

Slot size X	Encoding	#Bits
0	-	0
1	-	0
2	00	0
3	01- 00	2



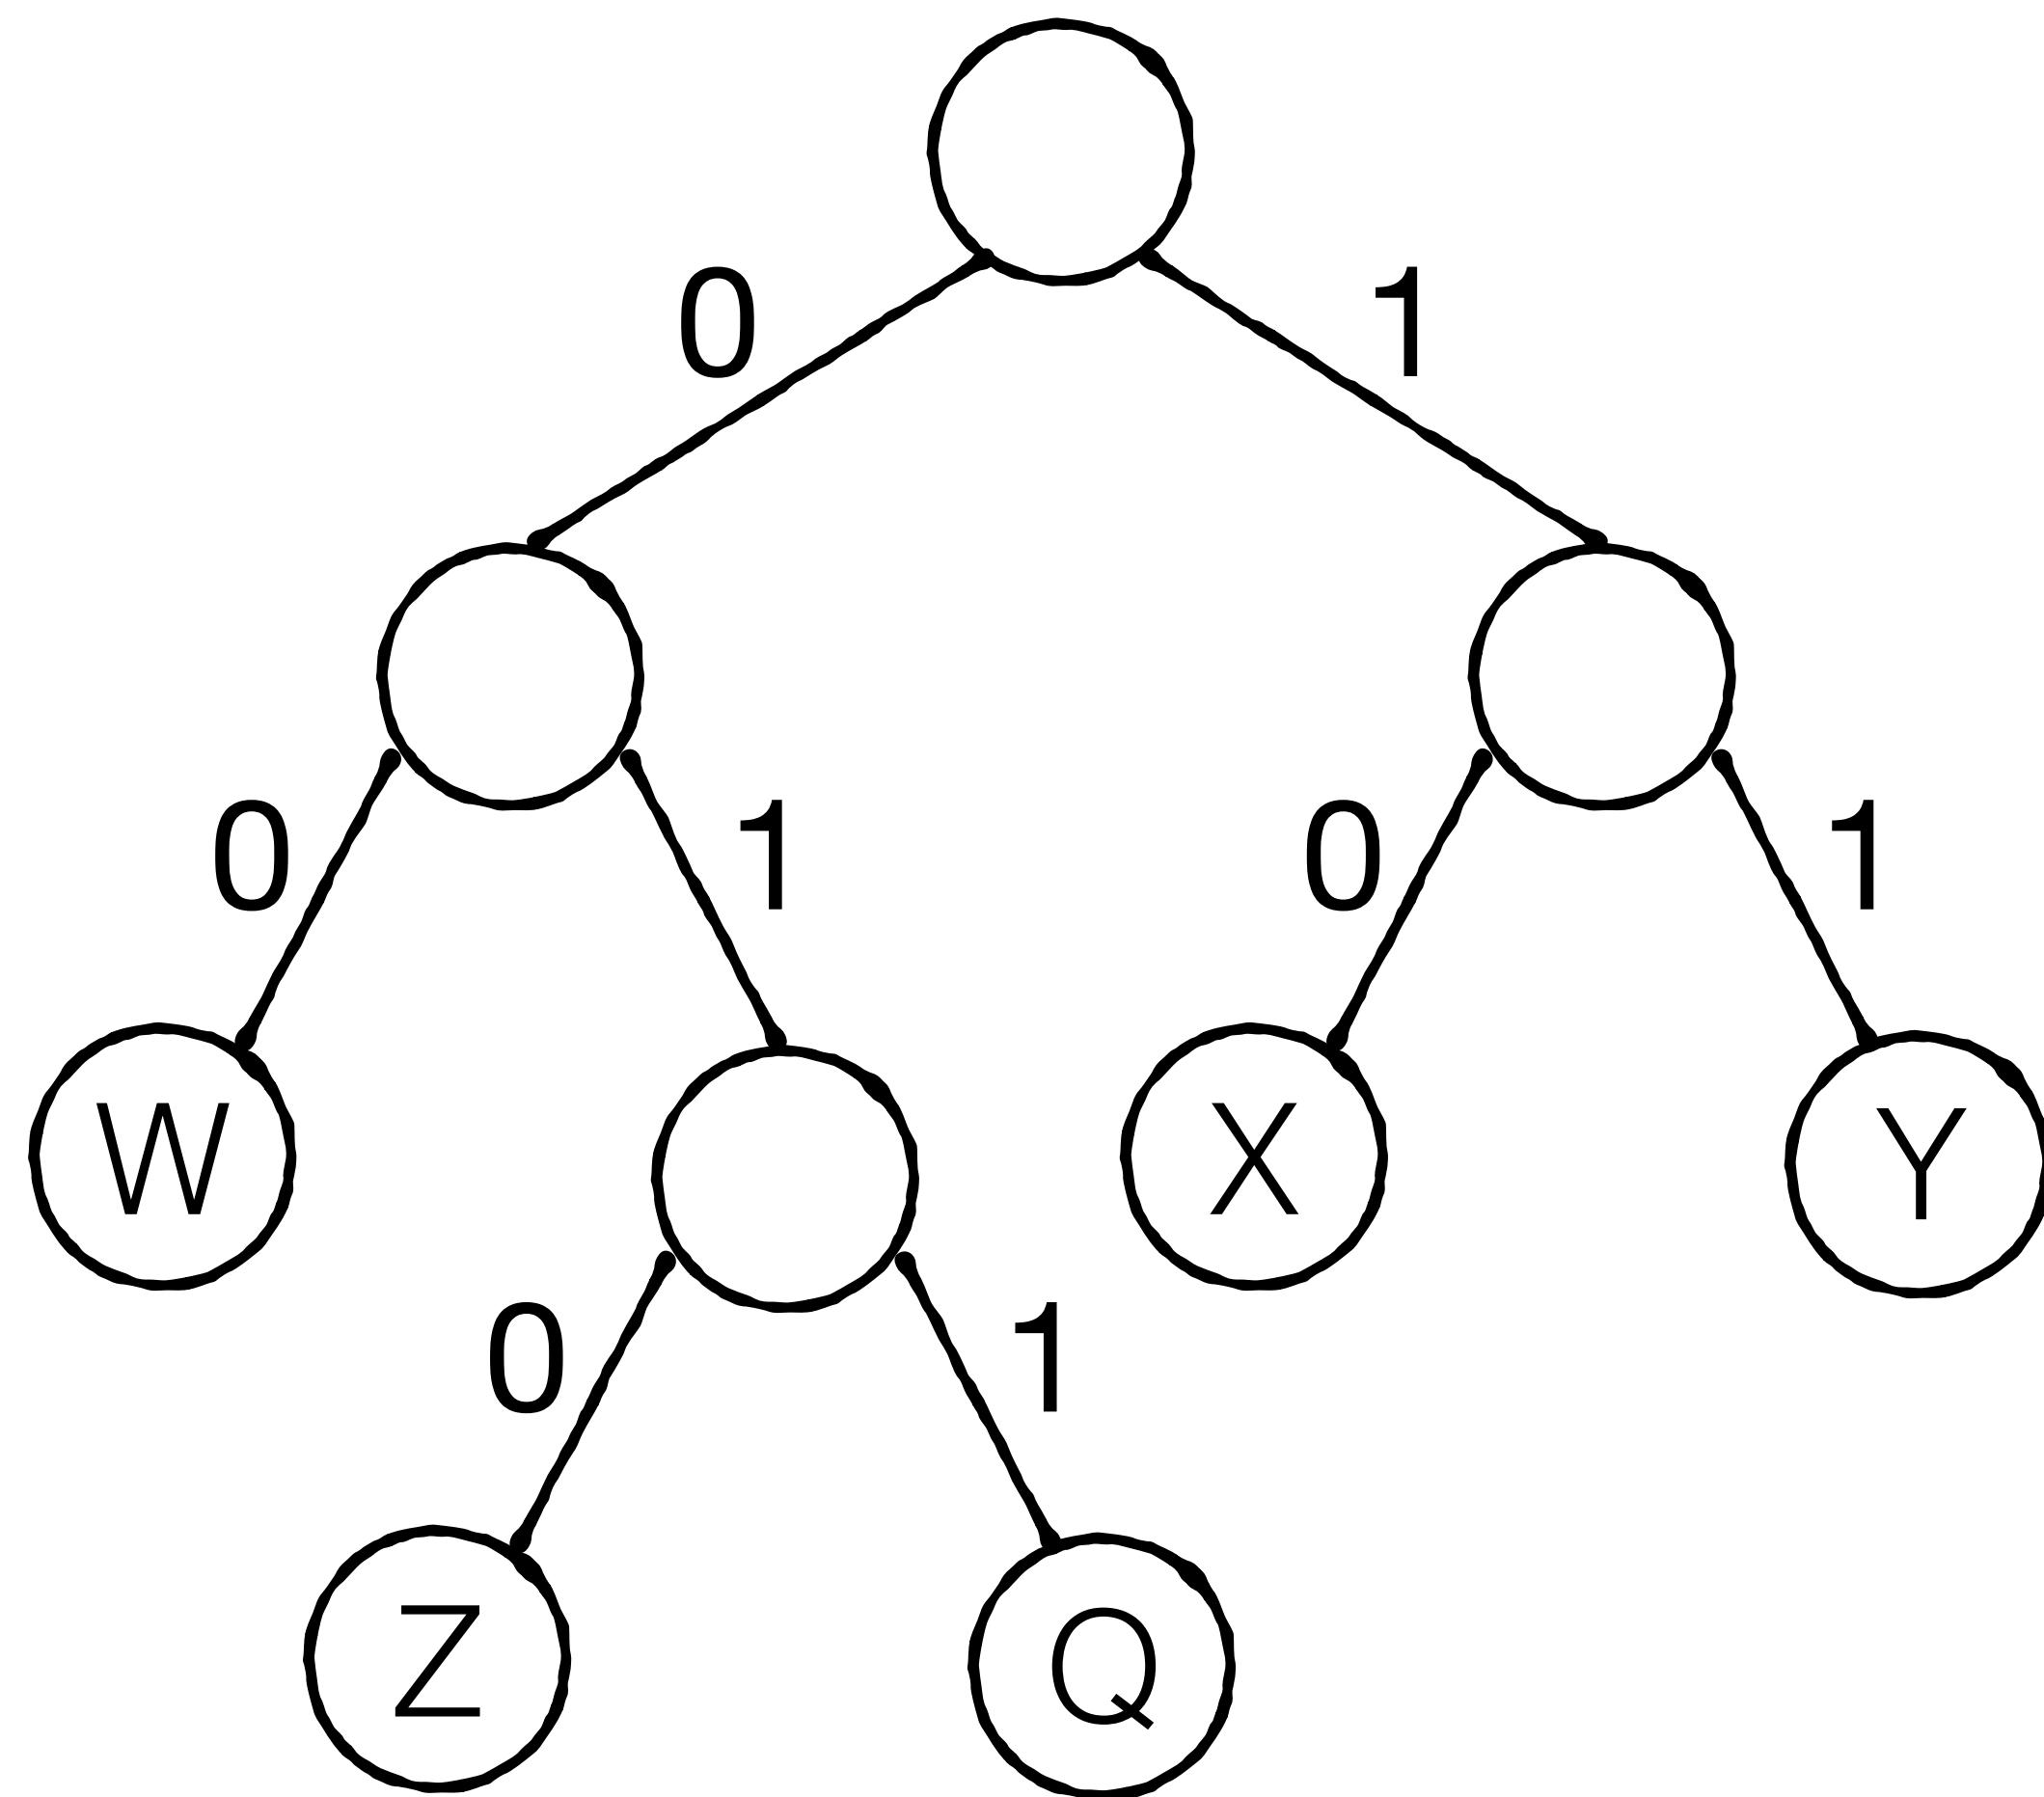
Topology

Slot size X	Encoding	#Bits
0	-	0
1	-	0
2	00	0
3	01-00	2
4	11 00-00	4



Topology

Slot size X	Encoding	#Bits
0	-	0
1	-	0
2	00	0
3	01-00	2
4	11 00-00	4
5	11 01 00-00	6



How large is each slot?

#entries

Topology

Indices

Pointers

Indices

Slot size X

#Bits

Indices

Slot size X

0

#Bits

0

Indices

Slot size X

#Bits

0

0

1

0

Indices

Slot size X

#Bits

0

0

hash(X) = 0 1 0 1 0 ...

1

0



2

?

diff with prob 0.5

Indices

Slot size X

#Bits

0

0

hash(X) = 0 1 0 1 0 ...

1

0

2

?

diff with prob 0.5



Indices

Slot size X

#Bits

0

0

hash(X) = 0 1 0 1 0 ...

1

0



2

?

diff with prob 0.5

Indices

Slot size X

#Bits

0	0
1	0
2	?

hash(X) = 0 1 0 1 0 ...



diff with prob 0.5

**First diff bit occurs after
2 bits in expectation**

Indices

Slot size X

#Bits

0

0

hash(X) = 0 1 0 1 0 ...

1

0



2

?

diff with prob 0.5

First diff bit occurs after
2 bits in expectation

Avg. of geometric dist. with prob 0.5

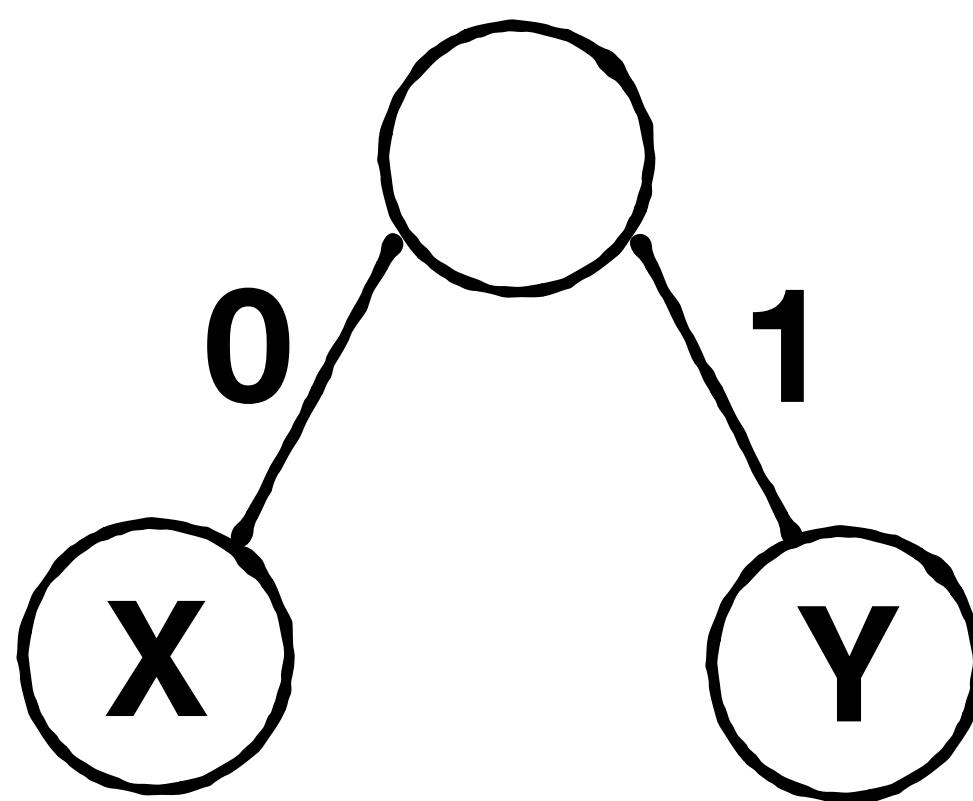
Indices

Slot size X

Avg. #Bits

0
1
2

0
0
2



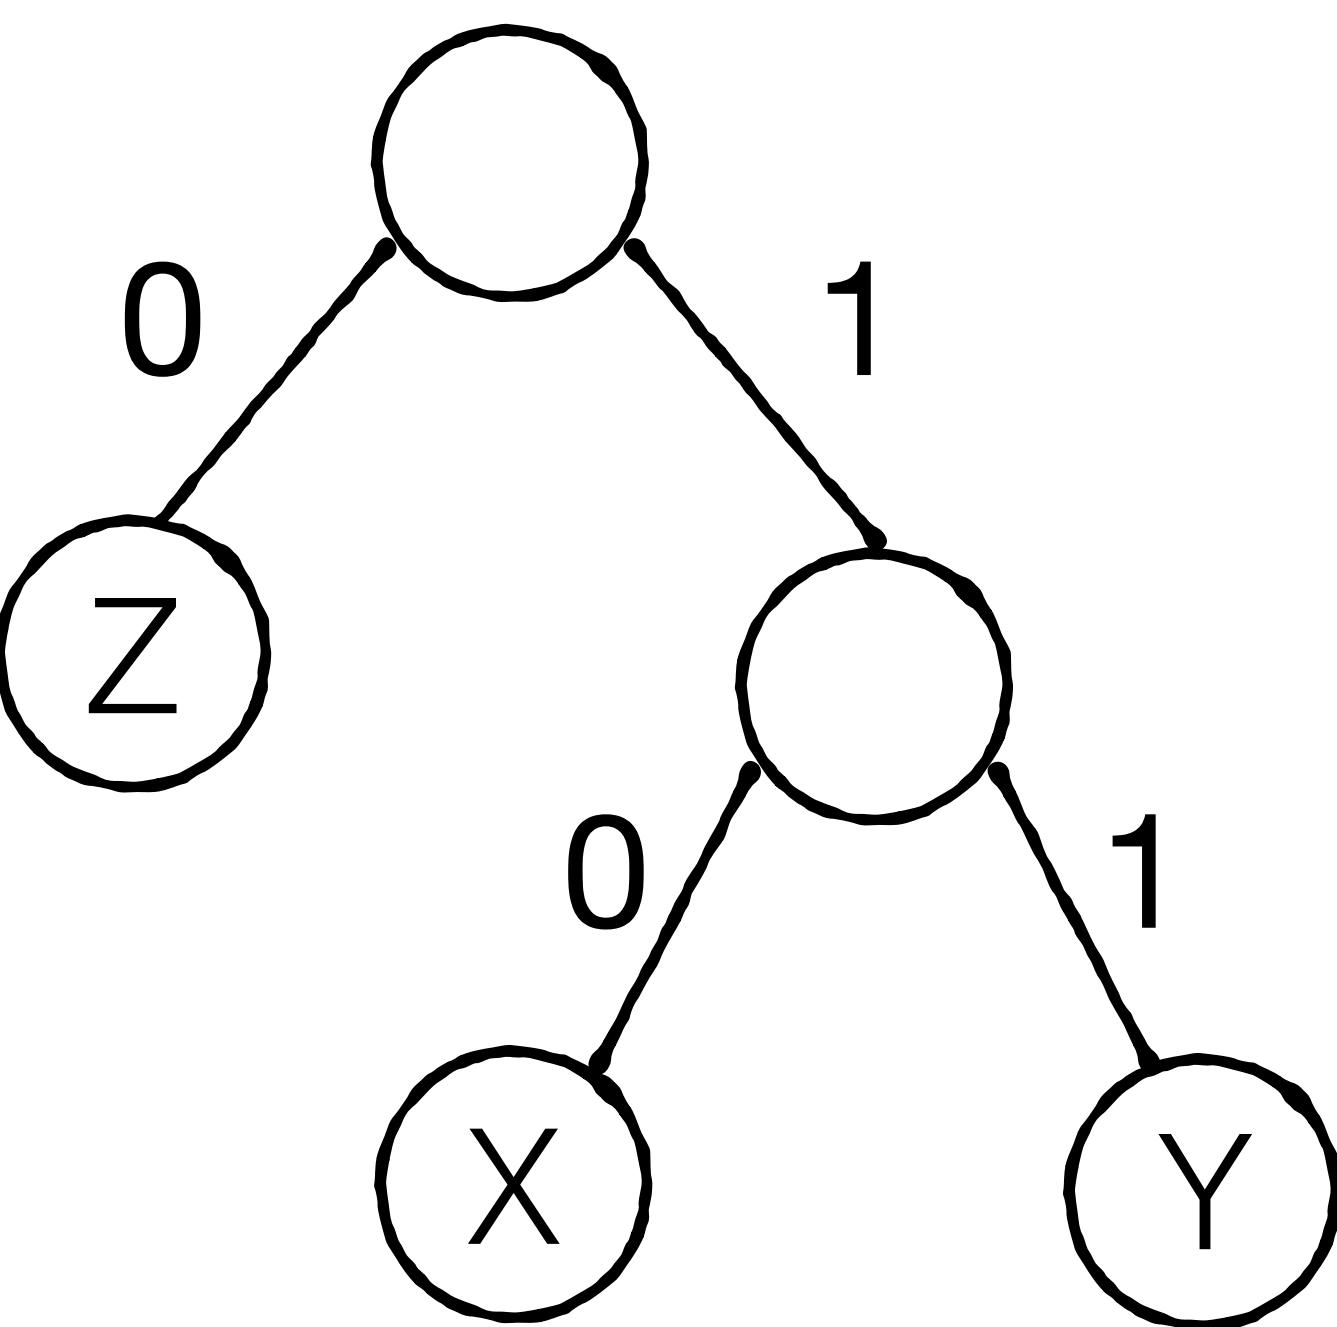
Indices

Slot size X

0
1
2
3

Avg. #Bits

0
0
2
4



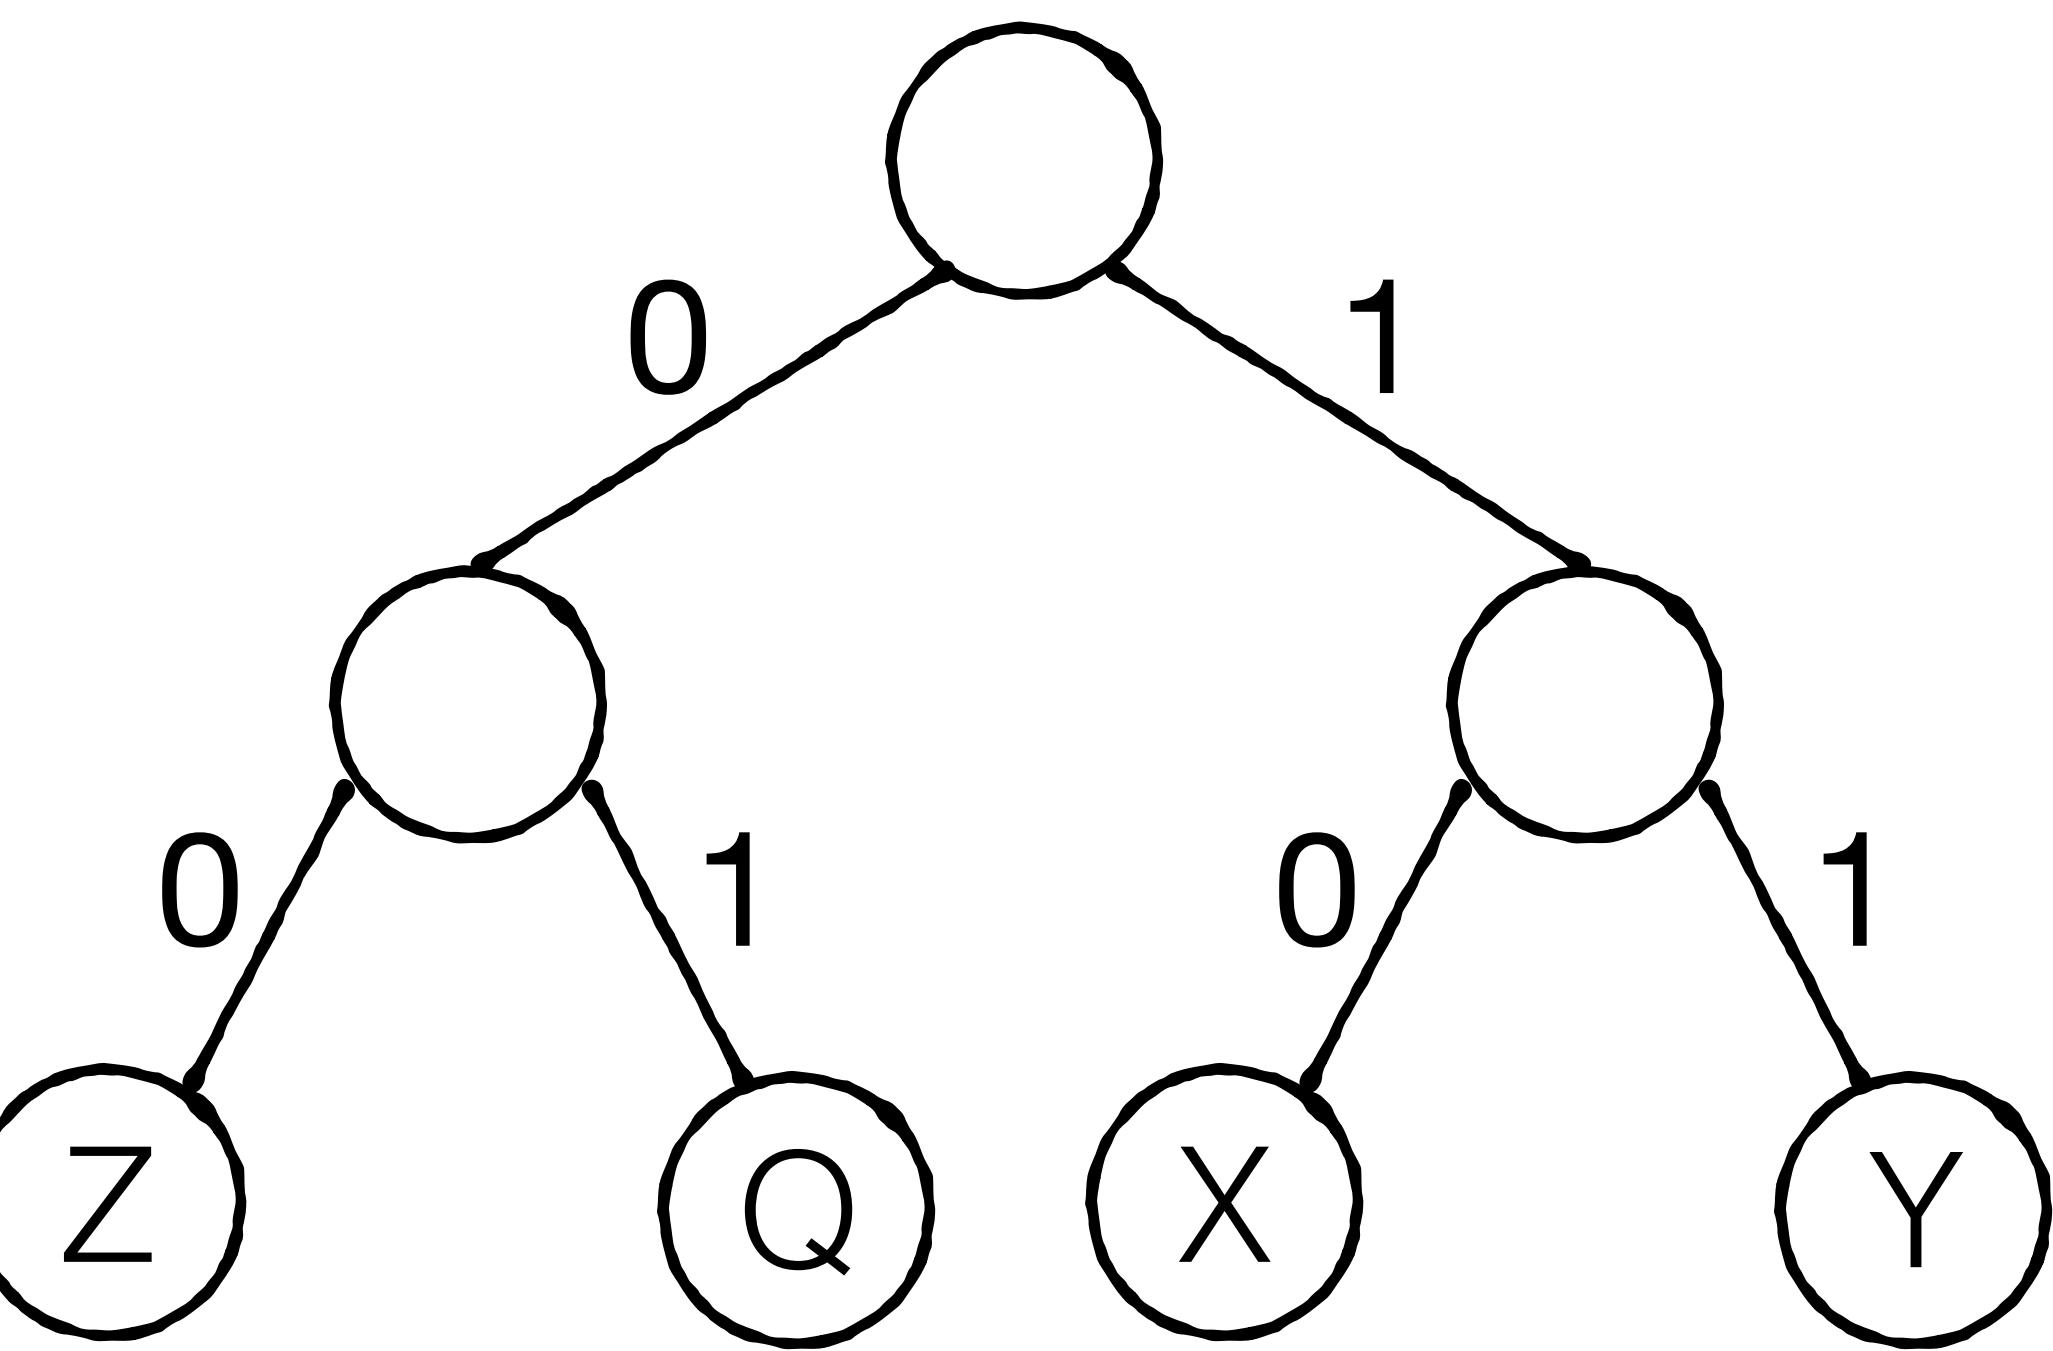
Slot size X

0
1
2
3
4

Avg. #Bits

0
1
2
4
6

Indices



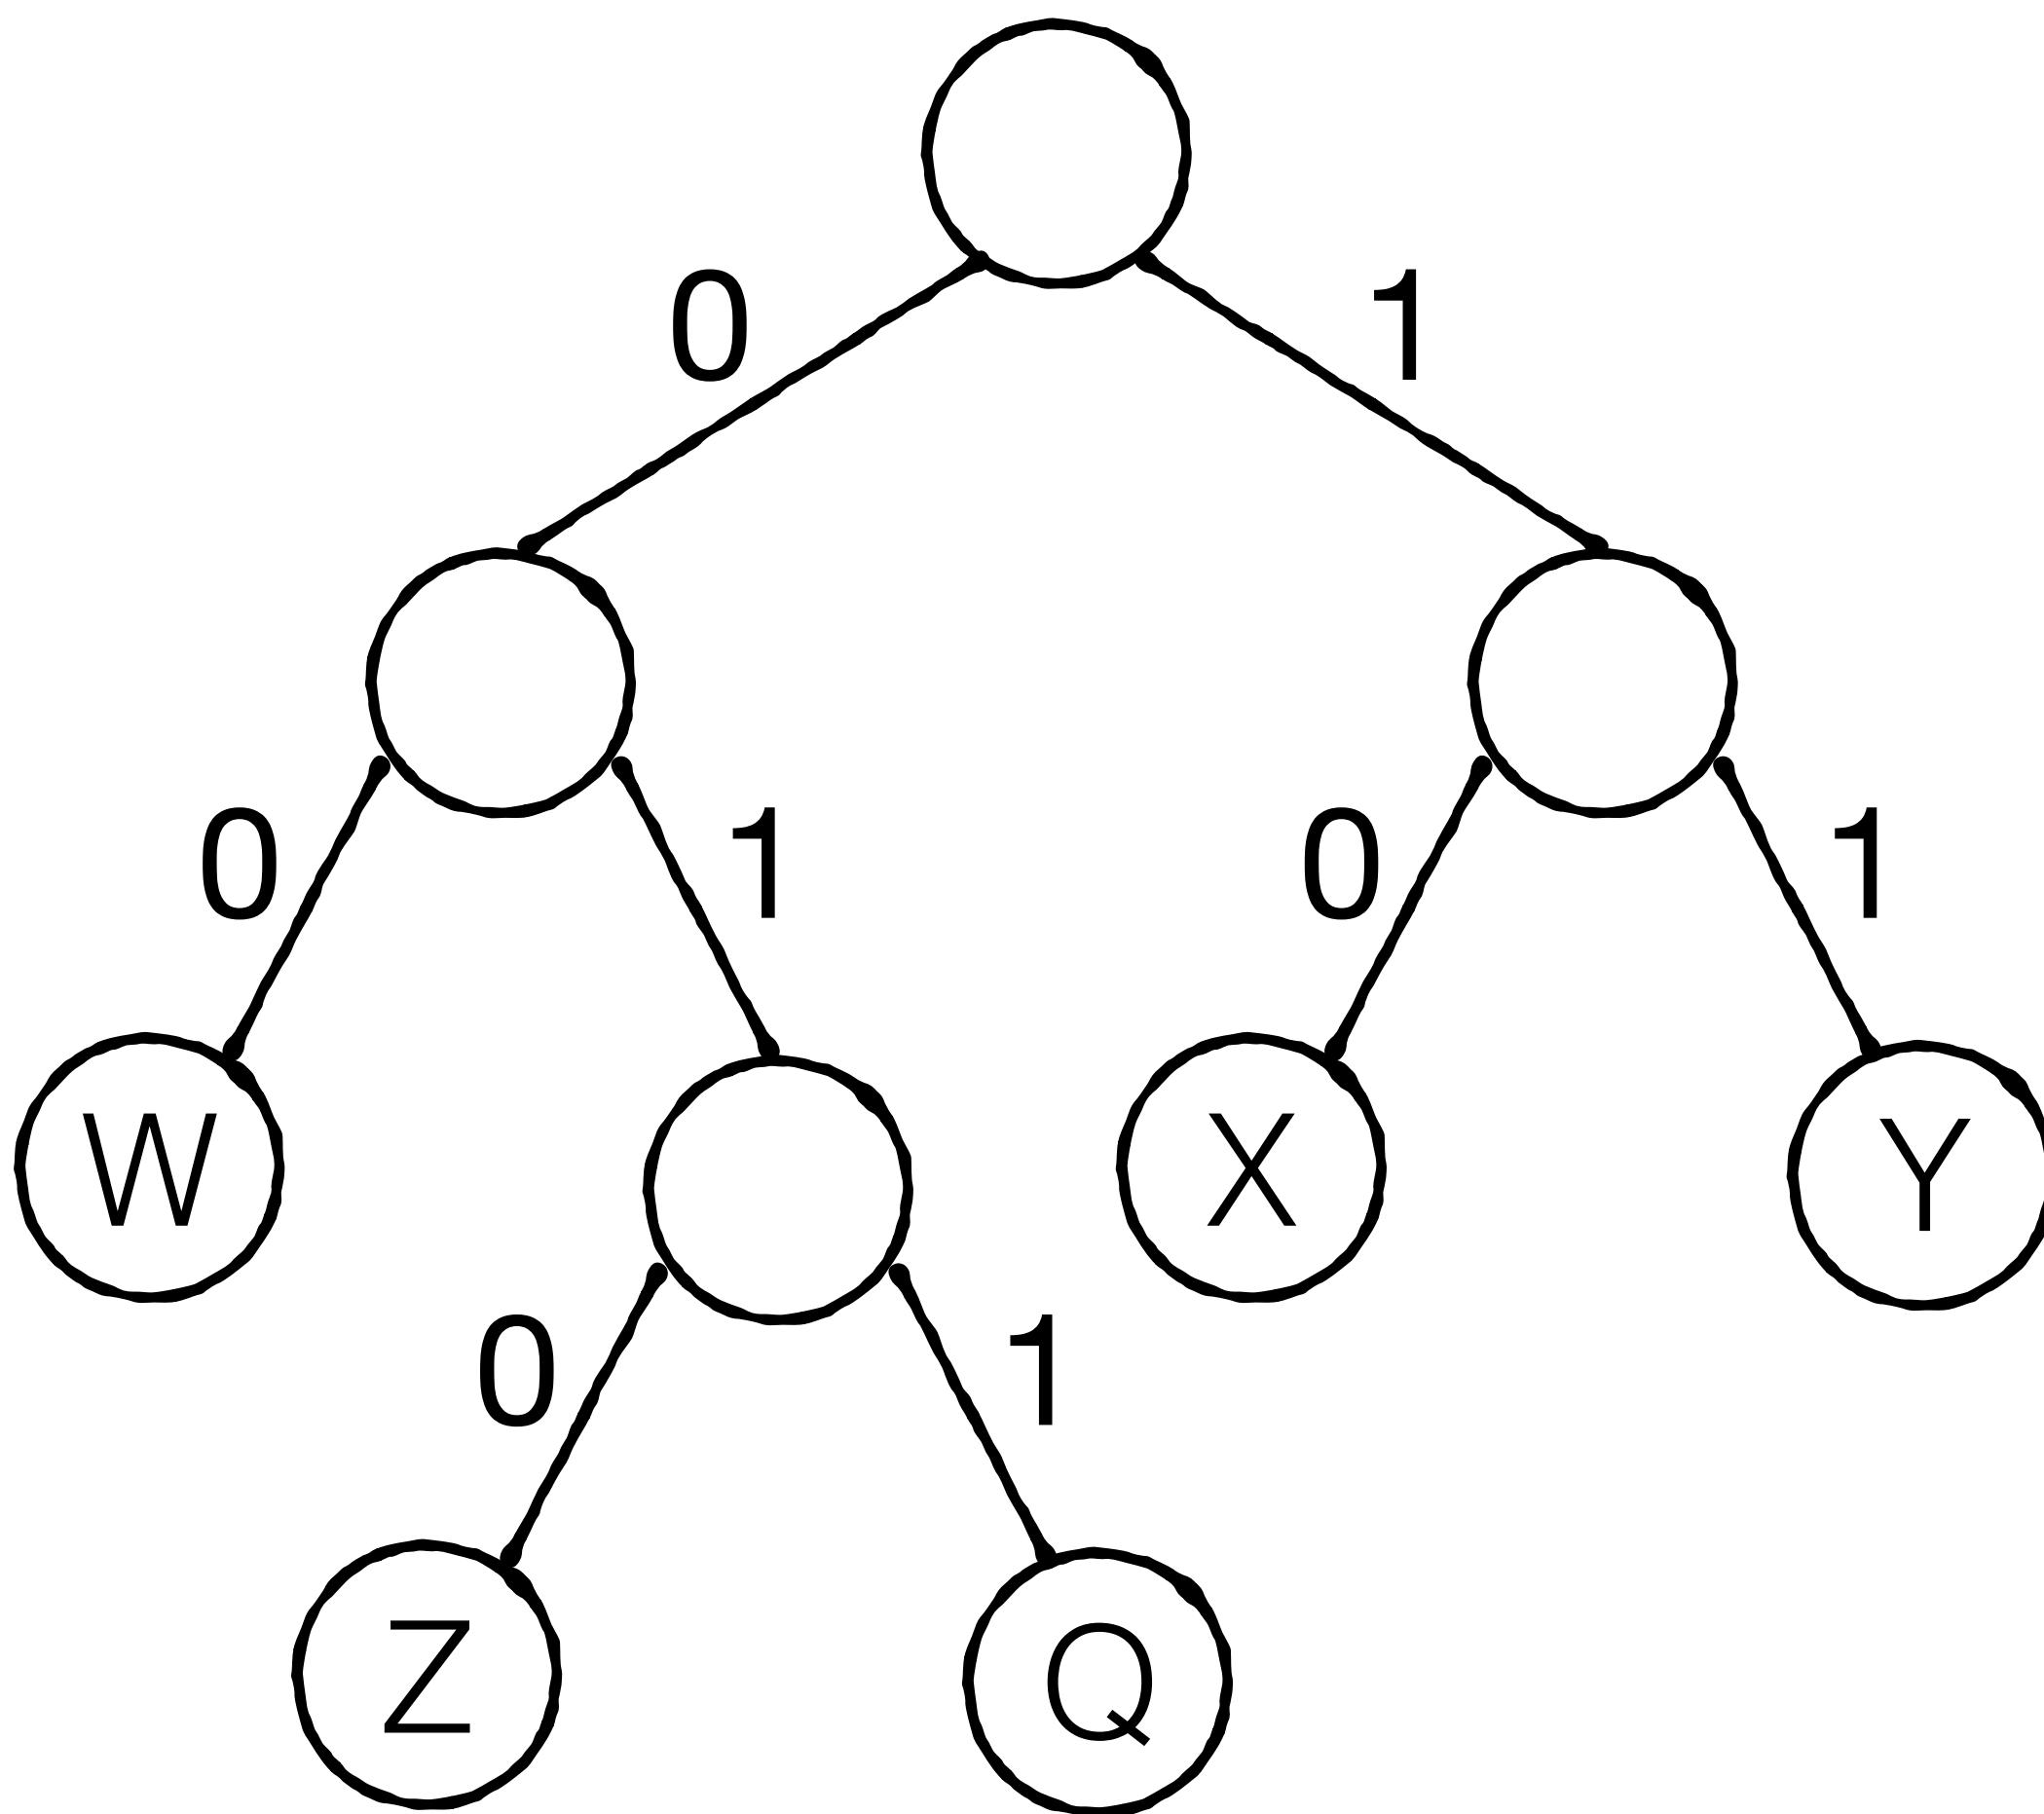
Indices

Slot size X

Avg. #Bits

0
1
2
3
4
5

0
1
2
4
6
8



How large is each slot?

#entries

Topology

Indices

Average

$$(\text{#entries}(i) + \text{Topology}(i) + \text{Indices}(i))$$

Average

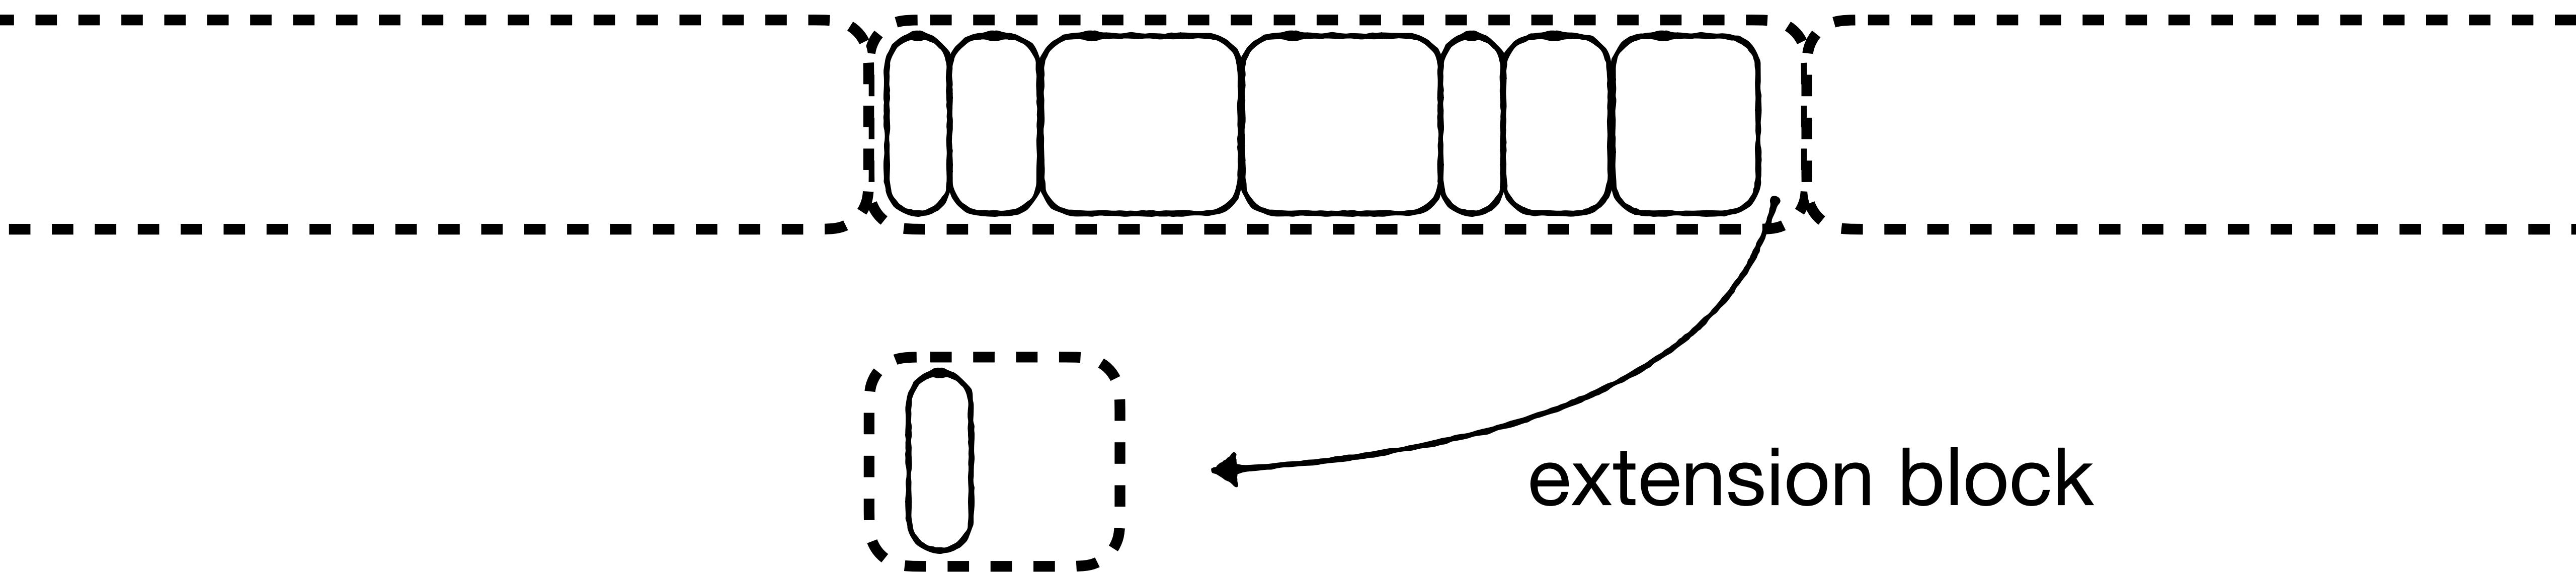
$$\sum_{i=0}^{\infty} \text{Poisson}(1, i) \cdot (\# \text{entries}(i) + \text{Topology}(i) + \text{Indices}(i))$$

Average

$$\sum_{i=0}^{\infty} \text{Poisson}(1, i) \cdot (\text{\#entries}(i) + \text{Topology}(i) + \text{Indices}(i)) \approx 3 \text{ bits}$$

bits in block containing X slots should be at least

$$X \cdot (3 + \text{pointer_size})$$

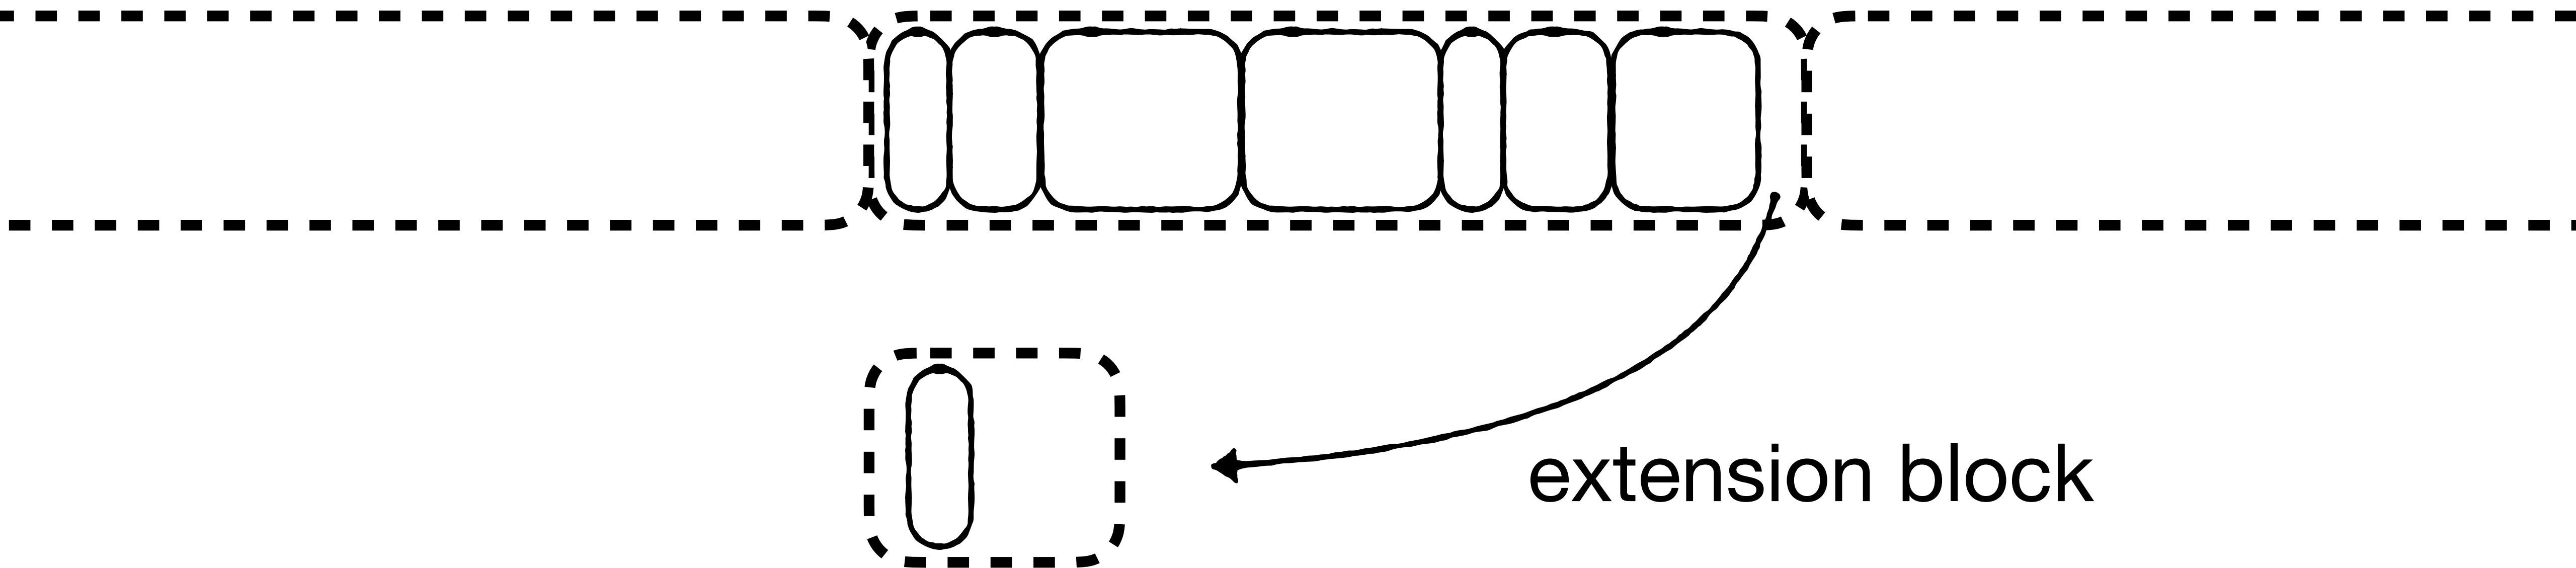


bits in block containing X slots should be at least

$$X \cdot (3 + \text{pointer_size} + 1)$$



Reduce overflows



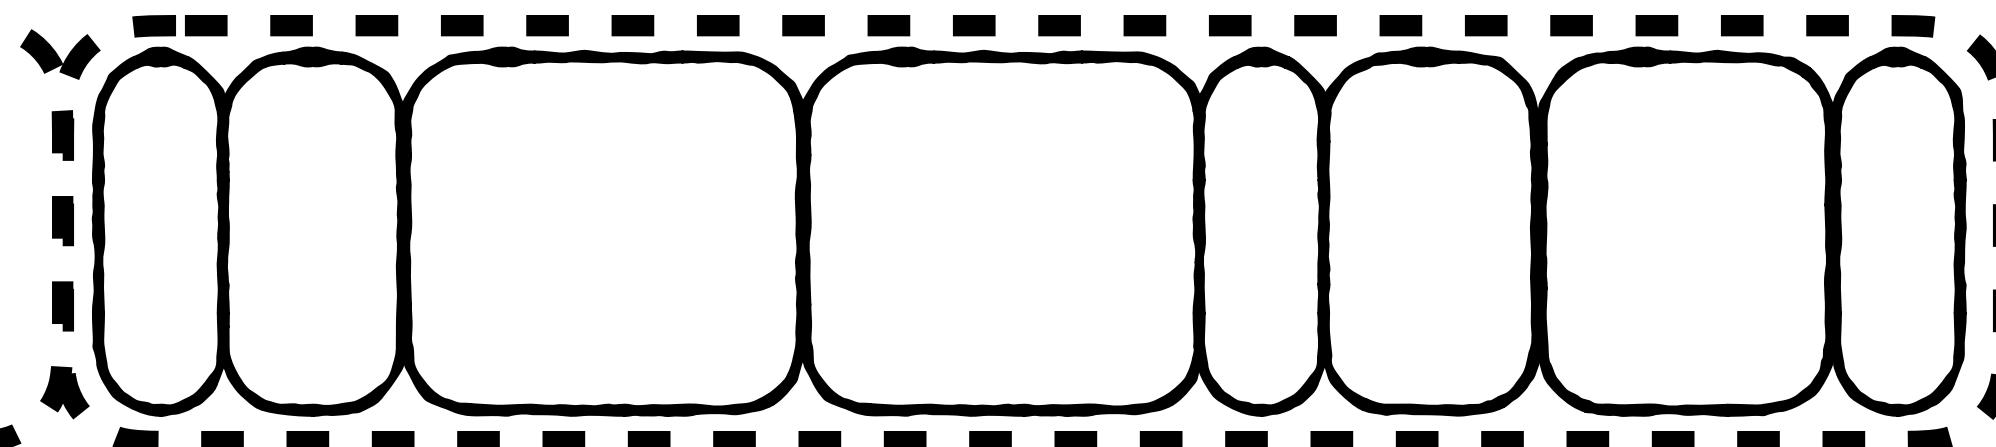
extension block

bits in block containing X slots should be at least

$$X \cdot (3 + \text{pointer_size} + 1)$$



Reduce overflows



Summary

Fingerprinting

Memory

$\approx e \cdot N$

Load factor

100%

Avg. query

$O(1)$

Insertions

N/A

Delta Hash Table

$\approx 4 \cdot N$

$\approx 90\%$

$O(1)$

$O(1)$

Fingerprinting

Memory

$\approx e \cdot N$

$\approx 4 \cdot N$

Load factor

100%

$\approx 90\%$

Avg. query

O(1)

O(1)

Insertions

N/A

O(1)

Construction

O(N)

O(N)

Delta Hash Table

**And now, a student
presentation**