

# Range Filters

(SuRF, Memento Filter, Diva)

Niv Dayan - CSC2525: Research Topics in Database Management

# Last lecture



This was fun...


# Logistics



Projects due  
on April 15




Exam  
April 30



Course  
Evaluation

# Course Evaluation



**Helps us improve**

# Course Evaluation

● Live

FAS Winter 2026 Grad

**RSH TOPICS IN DATA MANAGEMENT**  
**CSC2525H-S-LEC0101**

20 Invited  
0 Started  
0 Responded  
0 Opted out

Evaluation ends on:  
**2026-04-17**



Changes allowed until 2026-04-17



# Course Evaluation

● Live

FAS Winter 2026 Grad

**RSH TOPICS IN DATA MANAGEMENT**  
**CSC2525H-S-LEC0101**

20 Invited  
0 Started  
0 Responded  
0 Opted out

Evaluation ends on:  
**2026-04-17**




Changes allowed until 2026-04-17




**No excuses :)**

# What is a Filter?

**Does X exist?**




If key X does not exist



If key X does not exist




Does key  
X exist



**True  
negative  
with prob  $1-\epsilon$**


Memory



**False  
positive  
with prob  $\epsilon$**

Data





**Saves storage accesses & network hops**

**only support point queries (to one key)**

**Does X exist?**




**Set**



only support point queries (to one key)

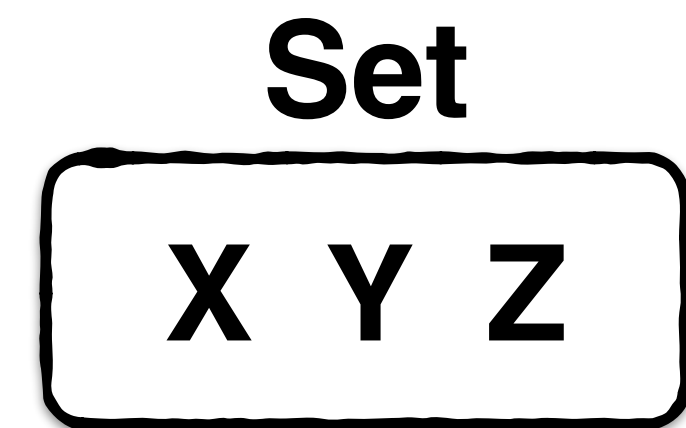

**How about a range?**

**Does X exist?**




How about a range?

**Does anything in  
[A, B] exist?**



How about a range?


**Does anything in  
[3, 5] exist?**



How about a range?

Does anything in  
**[3, 5]** exist?

**True positive**




Does anything in  
**[5, 8]** exist?




**Set**  
**1 4 9**


Does anything in  
**[5, 8]** exist?



**True negative  
with prob  $1-\epsilon$**



Does anything in  
**[5, 8]** exist?




**True negative  
with prob  $1-\epsilon$**

**False positive  
with prob  $\epsilon$**

# Applications?

Does anything in  
[5, 8] exist?



True negative  
with prob  $1-\epsilon$

False positive  
with prob  $\epsilon$




Does anything in  
[5, 8] exist?

### Filters




Memory

### LSM-tree



Storage


**Does anything in  
[A, B] exist?**



**Filters**




LSM-tree




# Ideally skip runs that don't contain relevant key range


Does anything in  
[A, B] exist?




**Skip**



**Skip**




**Skip**




# How to implement a range filter?


Does anything in  
[A, B] exist?




Skip



Skip




Skip



How to implement a range filter?

**Can we use a point filter (Bloom, Quotient) to answer range queries?**

Does anything in  
[A, B] exist?



Does anything in  
**[5, 8]** exist?



**Bloom**



Set



# Query every key in range


Does anything in **[5, 8]** exist?

- Check 5** →
- Check 6** →
- Check 7** →
- Check 8** →

Bloom

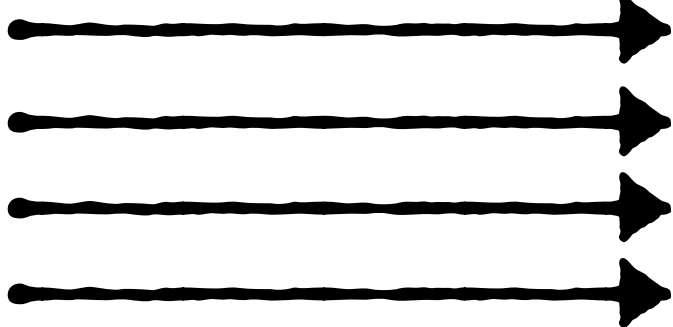



Set



Query every key in range

Does anything in  
**[A, B]** exist?

**R queries**  
  
 **$R = B - A + 1$**

Bloom  



Set  


**If all queries return negative, we know key does not exist**

Query every key in range

Does anything in  
**[A, B]** exist?

**R queries**



$$R = B - A$$

Bloom



Set




If all queries return negative, we know key does not exist

**If at least one returns a positive, the overall outcome is a positive**


Query every key in range

## Problems?



Does anything in  
[A, B] exist?

R queries



Bloom



Set




If all queries return negative, we know key does not exist

If at least one returns a positive, the overall outcome is a positive

# Problem 1: Query Cost

Does anything in  
[A, B] exist?

$O(R)$



Bloom




Set



## Problem 2: Does not work for infinite universe (e.g., strings)

Does anything in  
[A, B] exist?



Bloom

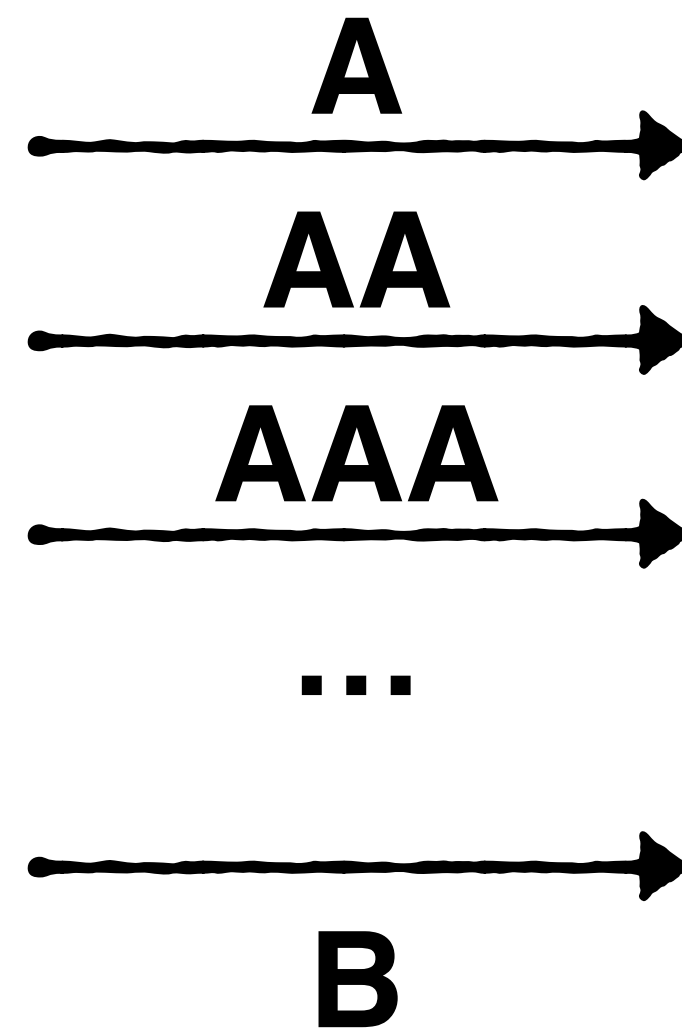


Set



## Problem 2: Does not work for infinite universe (e.g., strings)

Does anything in  
[A, B] exist?



Bloom

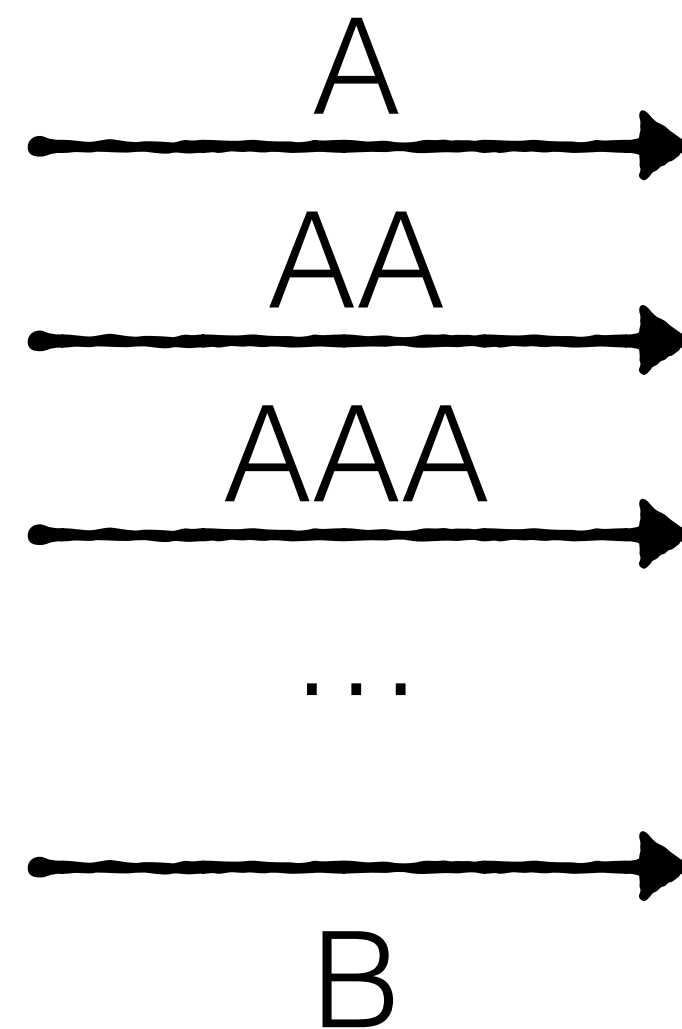


Set



Problem 2: Does not work for infinite universe (e.g., strings)

Does anything in  
[A, B] exist?



Bloom




Set



**There are infinitely many possible keys in a range :)**

# Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?



Bloom




Set



# Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?

**R queries**



**FPR  $\epsilon$**




Set



# Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?

**R queries**



**FPR  $\epsilon$**




Set




**P[false positive]?**

### Problem 3: Higher False Positive Rate




$$\mathbf{P[\text{false positive}] = P[\text{at least one query returns positive}]}$$

### Problem 3: Higher False Positive Rate




$$\begin{aligned} P[\text{false positive}] &= P[\text{at least one query returns positive}] \\ &= \mathbf{1 - P[\text{all queries return negative}]} \end{aligned}$$

### Problem 3: Higher False Positive Rate




$$\begin{aligned} P[\text{false positive}] &= P[\text{at least one query returns positive}] \\ &= 1 - P[\text{all queries return negative}] \\ &= \mathbf{1 - P[\text{one query return negative}]^R} \end{aligned}$$

### Problem 3: Higher False Positive Rate




$$\begin{aligned} P[\text{false positive}] &= P[\text{at least one query returns positive}] \\ &= 1 - P[\text{all queries return negative}] \\ &= 1 - P[\text{one query return negative}]^R \\ &= \mathbf{1 - (1 - \epsilon)^R} \end{aligned}$$

### Problem 3: Higher False Positive Rate



$$\begin{aligned} P[\text{false positive}] &= P[\text{at least one query returns positive}] \\ &= 1 - P[\text{all queries return negative}] \\ &= 1 - P[\text{one query return negative}]^R \\ &= 1 - (1-\epsilon)^R \\ &\approx \mathbf{1 - e^{\epsilon \cdot R}} \end{aligned}$$

### Problem 3: Higher False Positive Rate




$$\begin{aligned} P[\text{false positive}] &= P[\text{at least one query returns positive}] \\ &= 1 - P[\text{all queries return negative}] \\ &= 1 - P[\text{one query return negative}]^R \\ &= 1 - (1 - \epsilon)^R \\ &\approx 1 - e^{-\epsilon \cdot R} \\ &\prec \epsilon \cdot R \quad \text{By union bound} \end{aligned}$$

# Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?

R queries



FPR  $\epsilon$



Set




$$\mathbf{P[\text{false positive}] < \epsilon \cdot R}$$

# Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?

R queries



FPR  $\epsilon$



Set




$$P[\text{false positive}] < \epsilon \cdot R$$

**Useless for large R**

### Problem 3: Higher False Positive Rate

Does anything in  
[A, B] exist?

R queries



FPR  $\epsilon$



Set



$$P[\text{false positive}] < \epsilon \cdot R$$

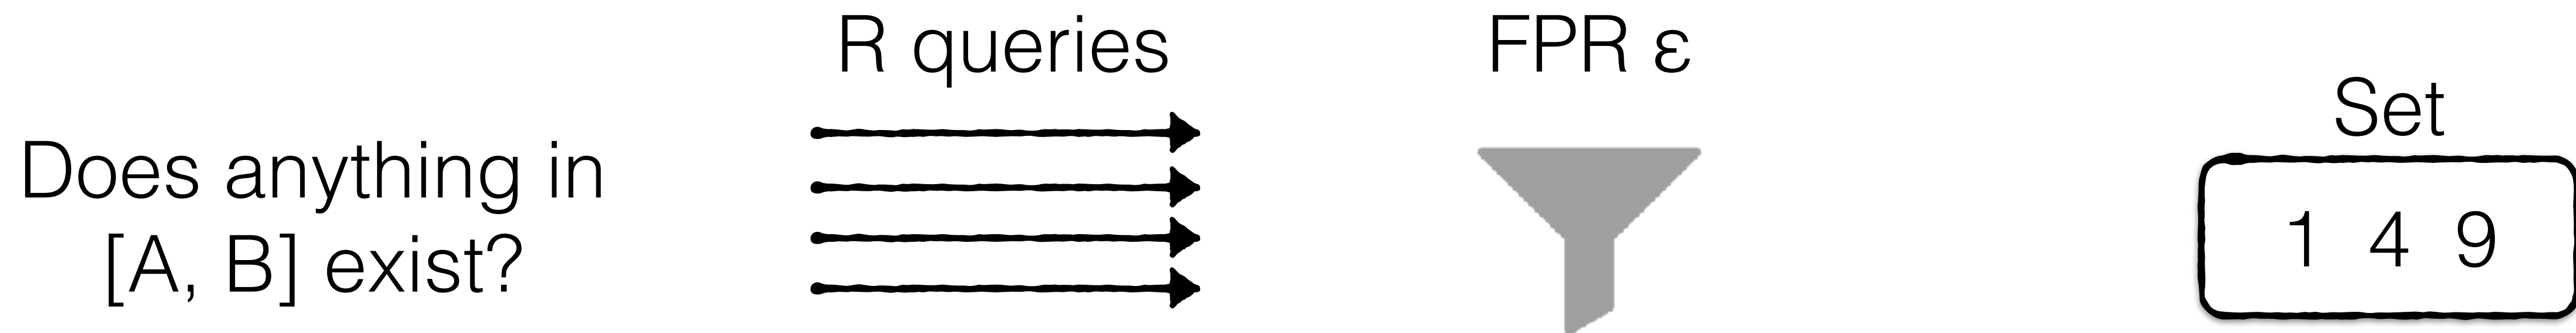


**Recall that:**

$$\epsilon = 2^{-F}$$

**where F = bits / entry**

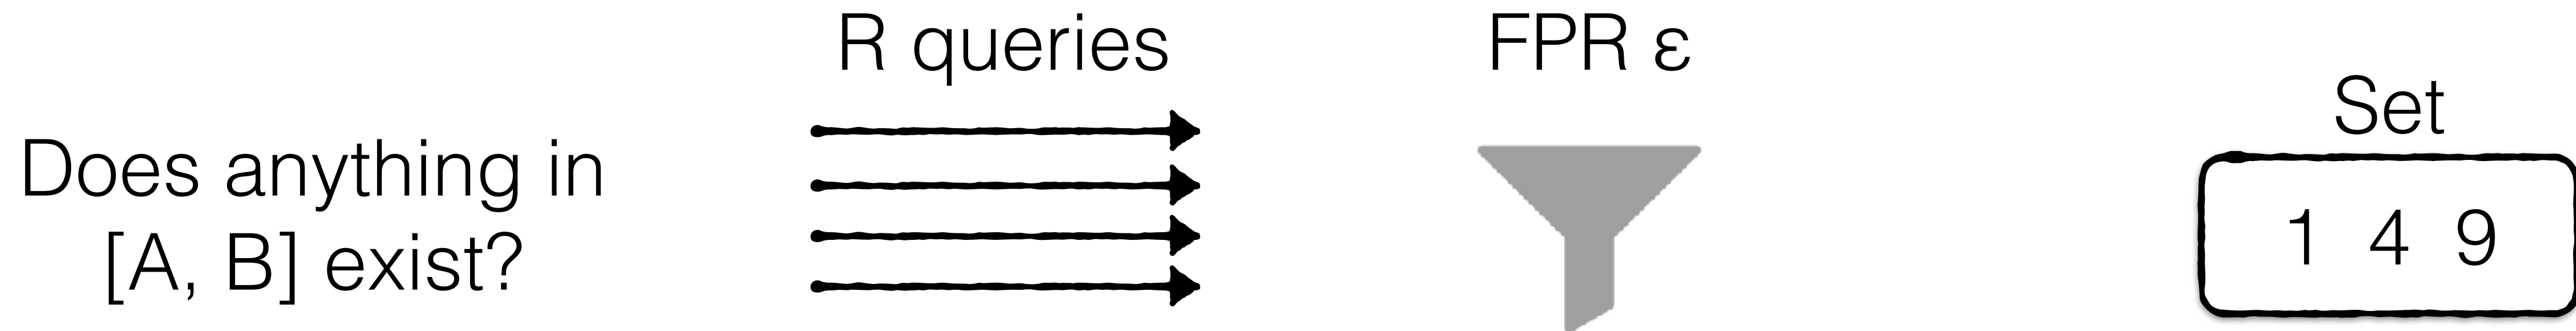
### Problem 3: Higher False Positive Rate



$$P[\text{false positive}] < \mathbf{2^{-F}} \cdot R$$

**How many extra bits per entry do we need to make up for R?**

### Problem 3: Higher False Positive Rate

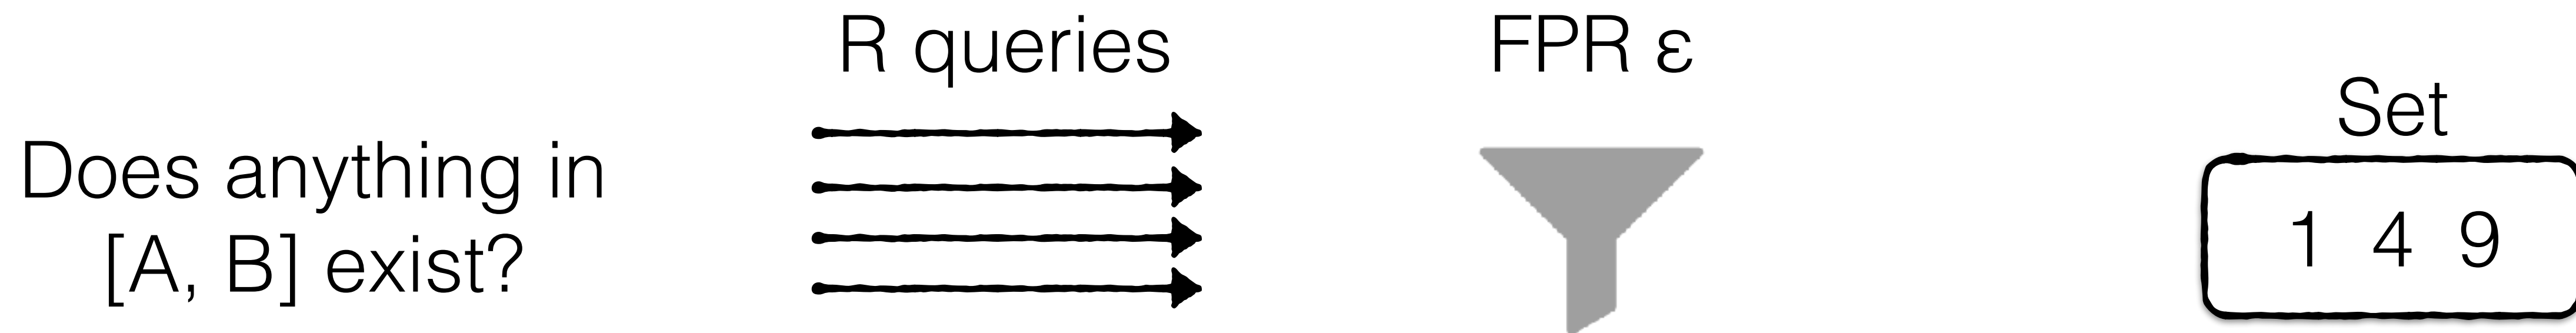


$$P[\text{false positive}] < \mathbf{2^{-F} \cdot R}$$

**How many extra bits per entry do we need to make up for R?**

$$\mathbf{\log_2(R)}$$


### Problem 3: Higher False Positive Rate



$$P[\text{false positive}] < \mathbf{2^{-F - \log_2(R)} \cdot R}$$

**How many extra bits per entry do we need to make up for R?**


### Problem 3: Higher False Positive Rate



$$P[\text{false positive}] < 2^{-F - \log_2(R)} \cdot R$$

**How many extra bits per entry do we need to make up for R?**


### Problem 3: Higher False Positive Rate



$$P[\text{false positive}] < 2^{-F - \log_2(R)} \cdot R$$

**FPR is higher by factor of R, or we need  $\log_2(R)$  extra bits / entry to keep it stable**

### Problem 3: Higher False Positive Rate



$$P[\text{false positive}] < 2^{-F - \log_2(R)} \cdot R$$

FPR is higher by factor of  $R$ , or we need  $\log_2(R)$  extra bits / entry to keep it stable

**But this also requires us to know max range query length in advance**


# Problems




**query cost**  
 **$O(R)$**



**FPR**  
 **$O(\epsilon \cdot R)$**



**Fixed-length**  
**keys**



**Bounded**  
**queries**

# Problems


**How much can we improve these?**




query cost  
 $O(R)$



FPR  
 $O(\epsilon \cdot R)$



Fixed-length  
keys



Bounded  
queries

# How much can we improve these?




**query cost**  
 **$O(1)$**




**?**



FPR  
 $O(\epsilon \cdot R)$



Fixed-length  
keys



Bounded  
queries

# How much can we improve these?




query cost  
 $O(1)$




**FPR**  
 $O(\varepsilon \cdot R)$

↑  
?



Fixed-length  
keys



Bounded  
queries


# How much can we improve these?




query cost  
 $O(1)$



FPR  
 $O(\epsilon)$



**Var-length  
keys**



Bounded  
queries



?


# How much can we improve these?




query cost  
 $O(1)$



FPR  
 $O(\epsilon)$



Var-length  
keys



**Var-length  
queries**







query cost  
 $O(1)$



FPR  
 $O(\epsilon)$



Var-length  
keys



Var-length  
queries



**+ Dynamic operations**  
**(Inserts, deletes, expansions, contractions)**

## Hot research topic in past decade

**SuRF**  
SIGMOD18

**Rosetta**  
SIGMOD2020

**SNARF**  
SIGMOD22

**Proteus**  
SIGMOD22

**REncoder**  
ICDE23

**BloomRF**  
EDBT23

**Grafite**  
SIGMOD24

**Oasis**  
VLDB24

Hot research topic in past decade

**SuRF**  
SIGMOD18

**Rosetta**  
SIGMOD2020

**SNARF**  
SIGMOD22



**Proteus**  
SIGMOD22

**REncoder**  
ICDE23

**BloomRF**  
EDBT23

**Grafite**  
SIGMOD24

**Oasis**  
VLDB24



SuRF  
SIGMOD18

Rosetta  
SIGMOD2020

SNARF  
SIGMOD22



Proteus  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

**Grafite**  
SIGMOD24

Oasis  
VLDB24



SuRF  
SIGMOD18

**Rosetta**  
SIGMOD2020

**SNARF**  
SIGMOD22




**Proteus**  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

**Grafite**  
SIGMOD24

**Oasis**  
VLDB24



**SuRF**  
SIGMOD18

Rosetta  
SIGMOD2020

**SNARF**  
SIGMOD22



Proteus  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

Grafite  
SIGMOD24

**Oasis**  
VLDB24



**SuRF**  
SIGMOD18

Rosetta  
SIGMOD2020

SNARF  
SIGMOD22



Proteus  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

Grafite  
SIGMOD24

Oasis  
VLDB24



SuRF  
SIGMOD18

Rosetta  
SIGMOD2020

SNARF  
SIGMOD22




Proteus  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

Grafite  
SIGMOD24

Oasis  
VLDB24



SuRF  
SIGMOD18

Rosetta  
SIGMOD2020

SNARF  
SIGMOD22

Proteus  
SIGMOD22

REncoder  
ICDE23

BloomRF  
EDBT23

Grafite  
SIGMOD24

Oasis  
VLDB24

**None achieve all goals simultaneously**

**SuRF**  
**SIGMOD18**



**Memento**  
**SIGMOD24**



**Diva**  
**VLDB25**



**SuRF**  
**SIGMOD18**



**Me**

Memento  
SIGMOD24



Diva  
VLDB25




**SuRF**  
**SIGMOD18**



Memento  
SIGMOD24



Diva  
VLDB25



**Navid**

**SuRF**



**Memento**



**Diva**



**Var-length keys  
& queries**

**FPR guarantee**

**Query speed**

**Dynamic**

**SuRF**



**Memento**



**Diva**



**Var-length keys  
& queries**

**Yes**

**FPR guarantee**

**None**

**Query speed**

**$O(L)$   
( $L = \text{key length}$ )**

**Dynamic**

**No**

**SuRF**



**Memento**



**Diva**



**Var-length keys  
& queries**

Yes

**No**

**FPR guarantee**

None

**Robust**

**Query speed**

$O(L)$

**$O(1)$**

( $L$  = key length)

**Dynamic**

No

**Yes**

## SuRF



## Memento



## Diva



**Var-length keys  
& queries**

Yes

No

**Yes**

**FPR guarantee**

None

Robust

**Semi-Robust**

**Query speed**

$O(L)$

$O(1)$

**$O(\log L)$**

( $L$  = key length)

**Dynamic**

No


Yes

**Yes**

# SuRF: Practical Range Query Filtering with Fast Succinct Tries

Huanchen Zhang, Hyeontaek Lim, Viktor Leis, David G. Andersen,  
Michael Kaminsky, Kimberly Keeton, Andrew Pavlo

**SIGMOD18**




# Starting Point

## Keys

SIGAI

SIGMOD


SIGOPS




Starting Point

Keys  
SIGAI  
SIGMOD  
SIGOPS

**Fanout - 256**  
(1 byte)




# Problems?



# Problems?


1. Stores full keys




# Problems?

**1. Stores full keys**

**2. Pointers take space**

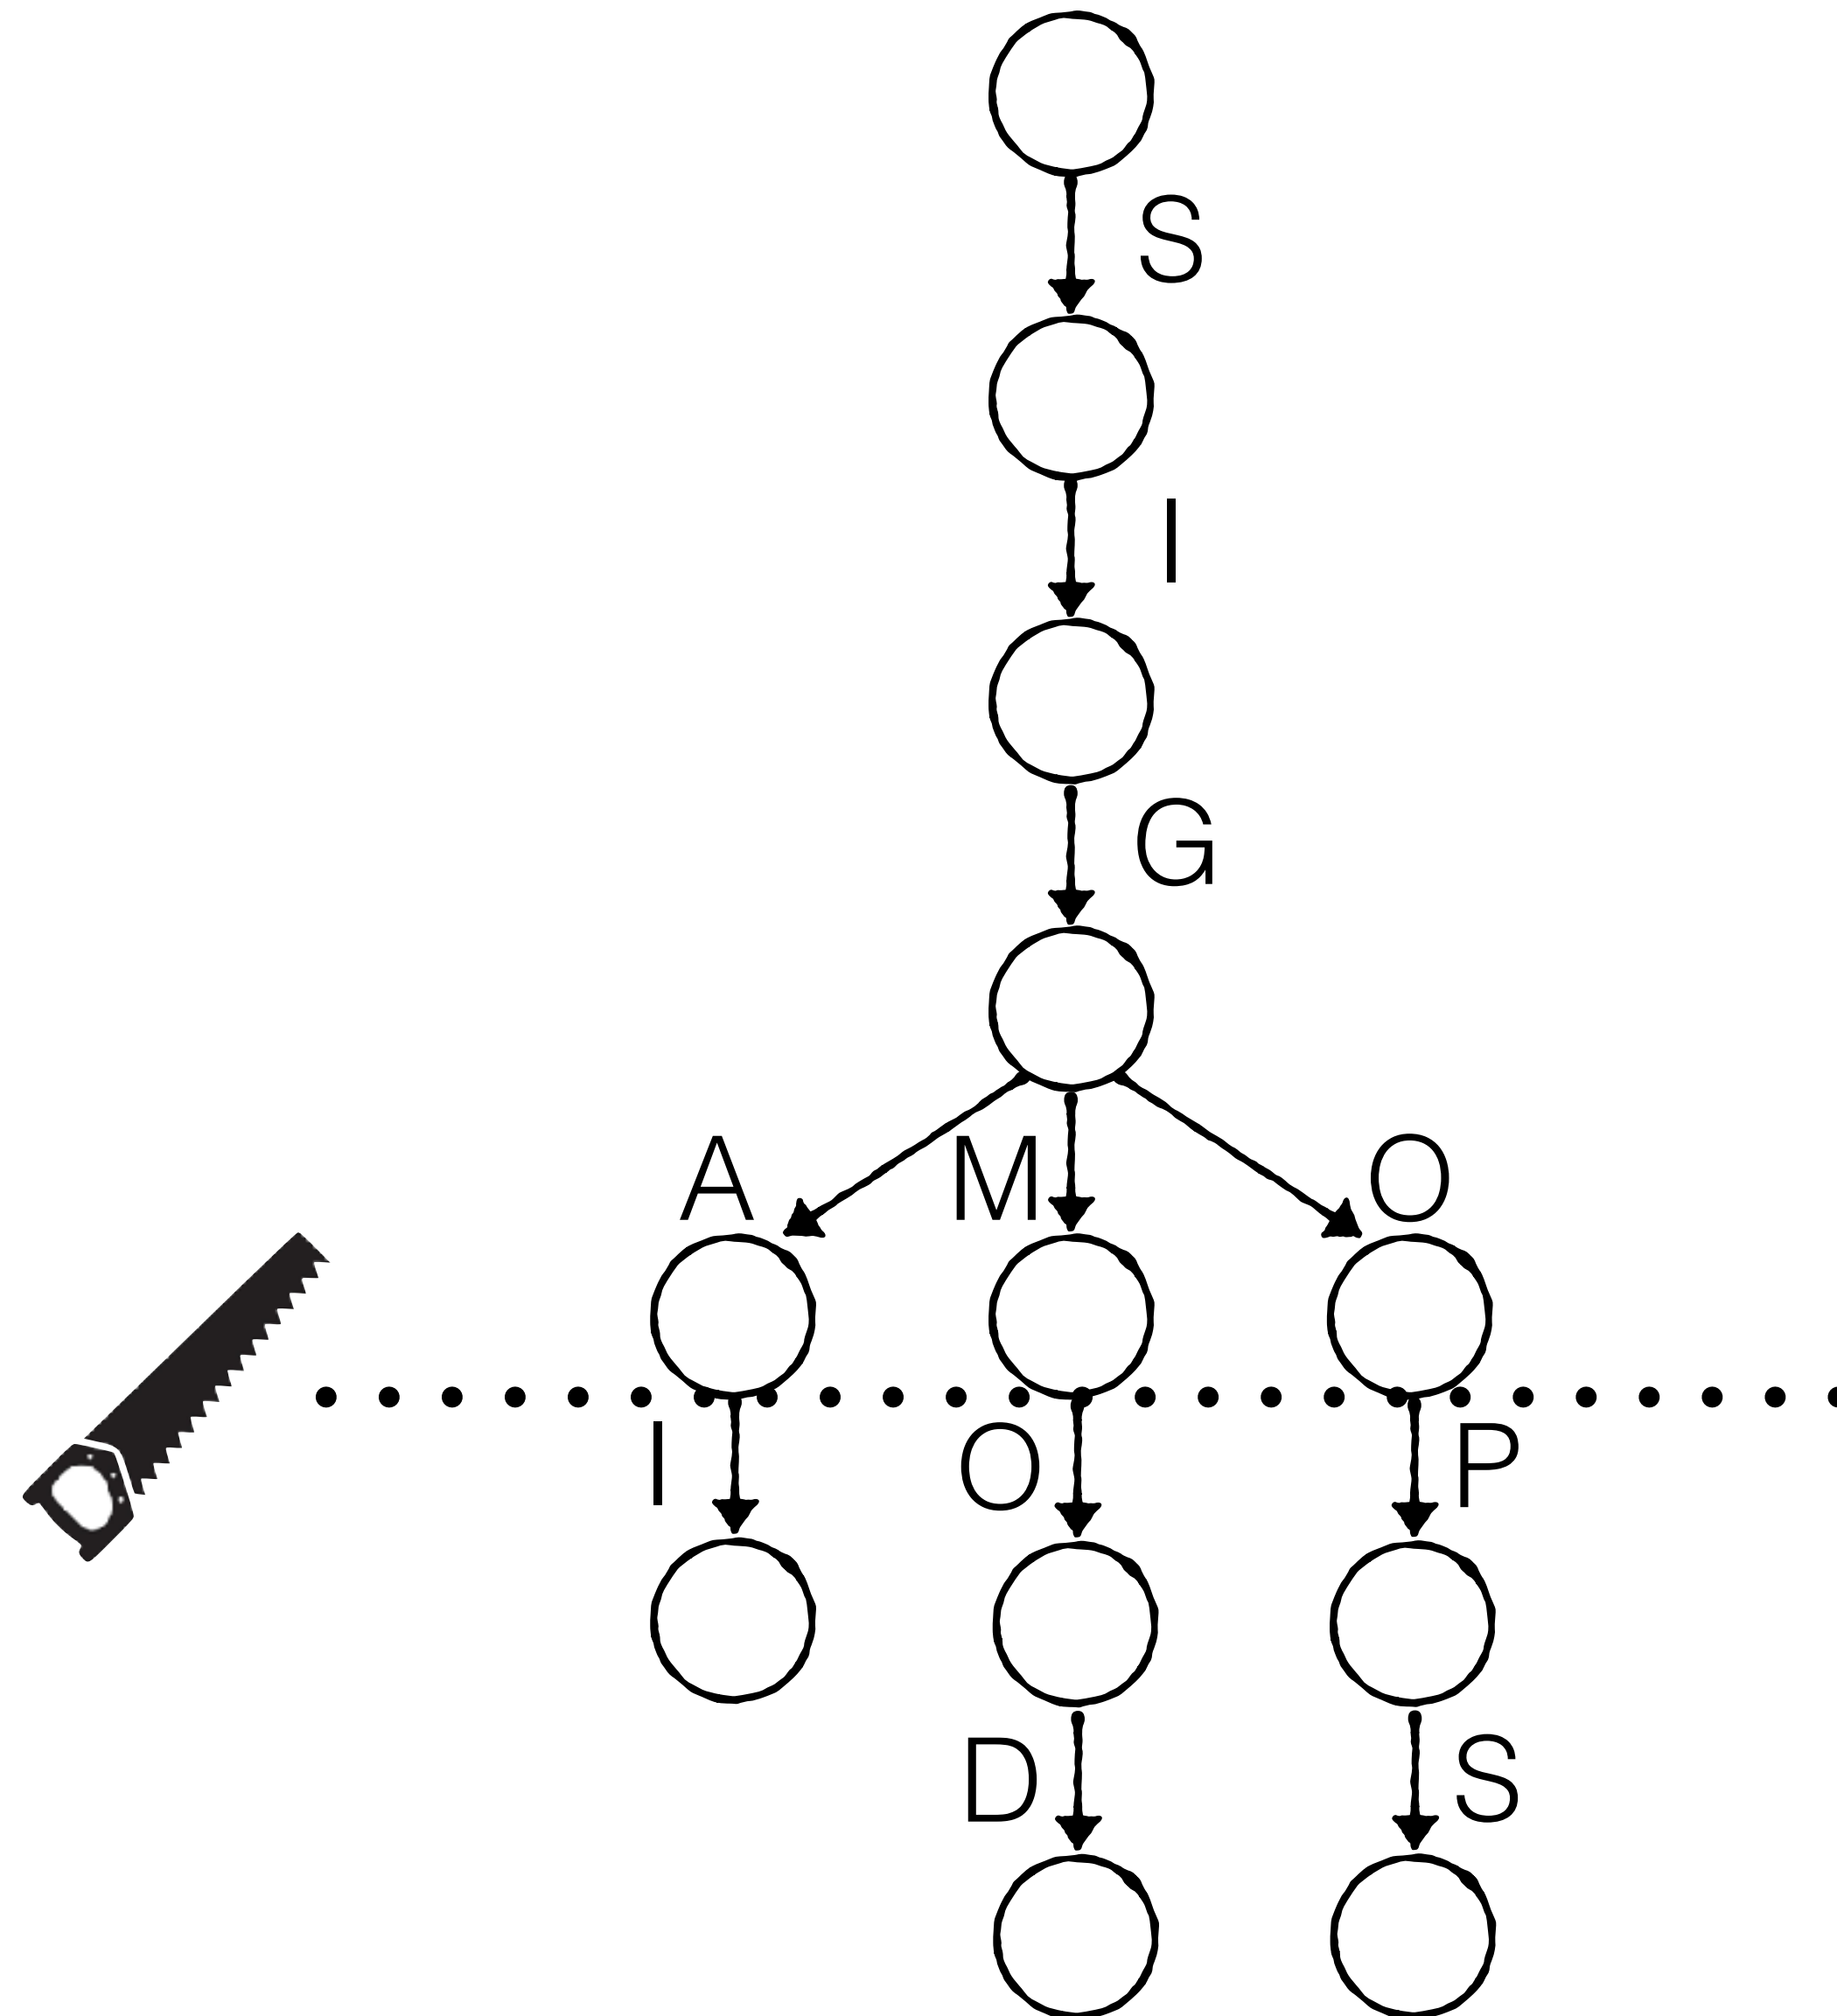


1. Stores full keys  
**truncation**

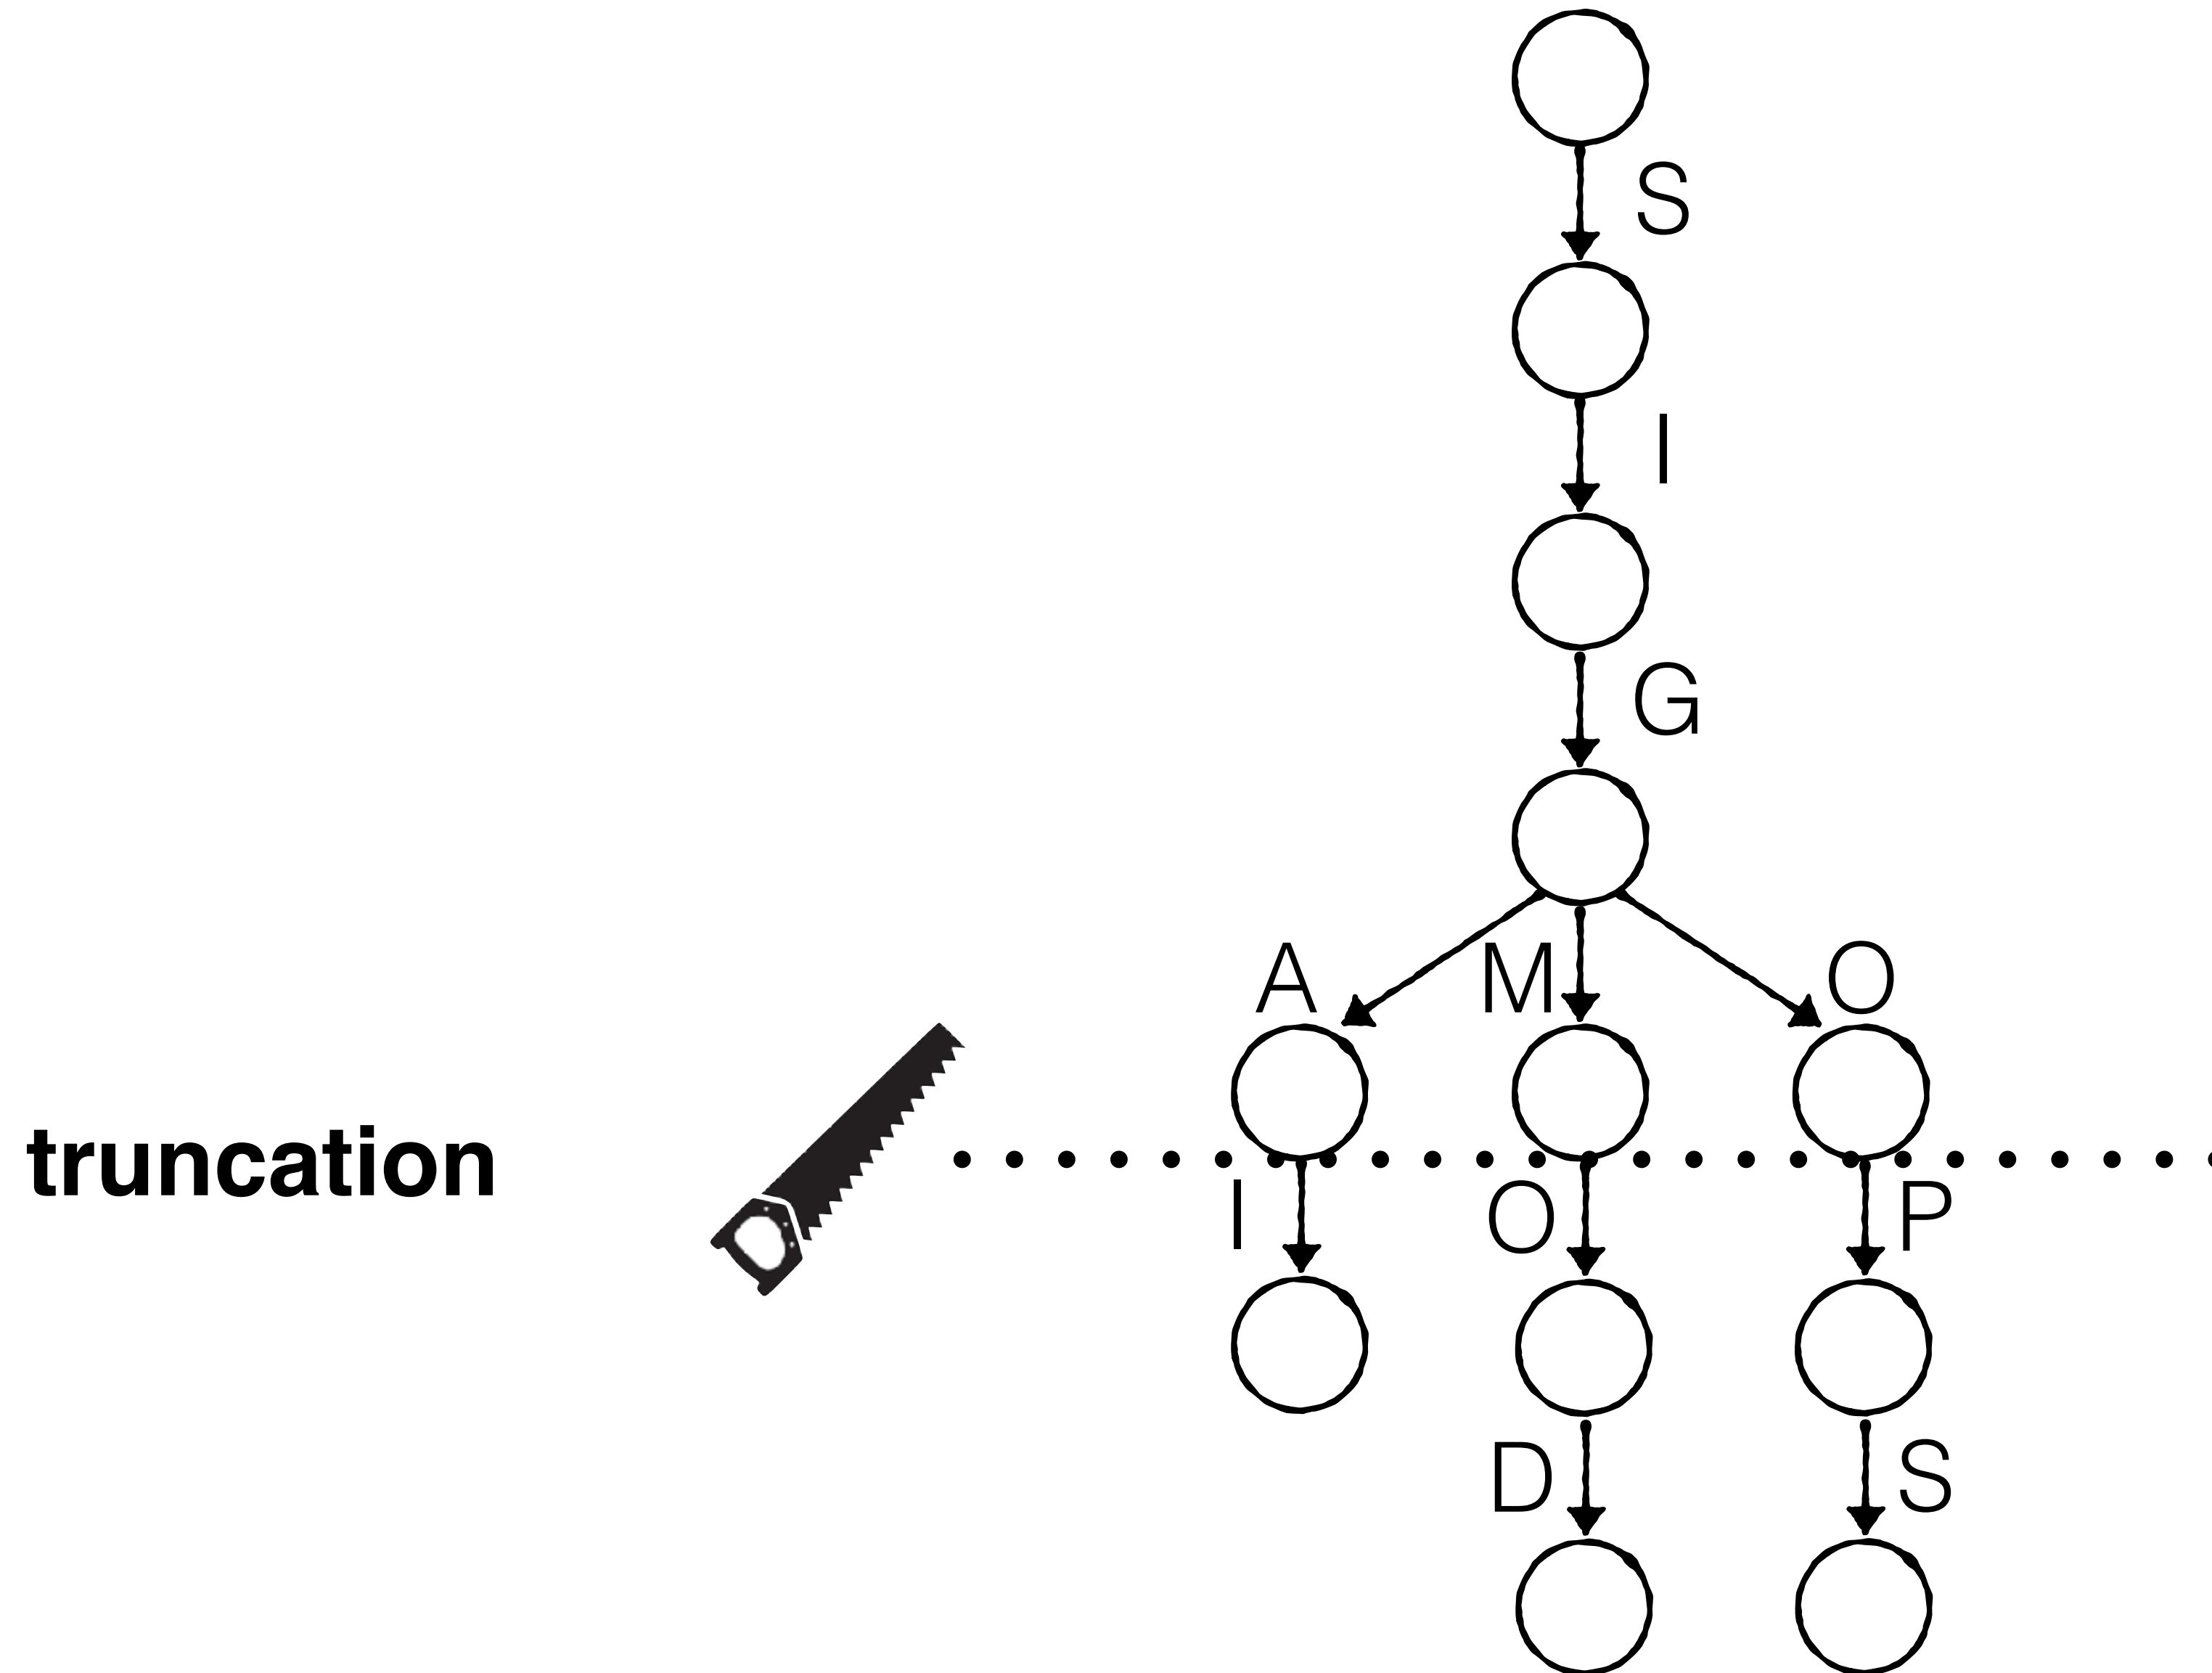


2. Pointers take space  
**Succinct encoding**

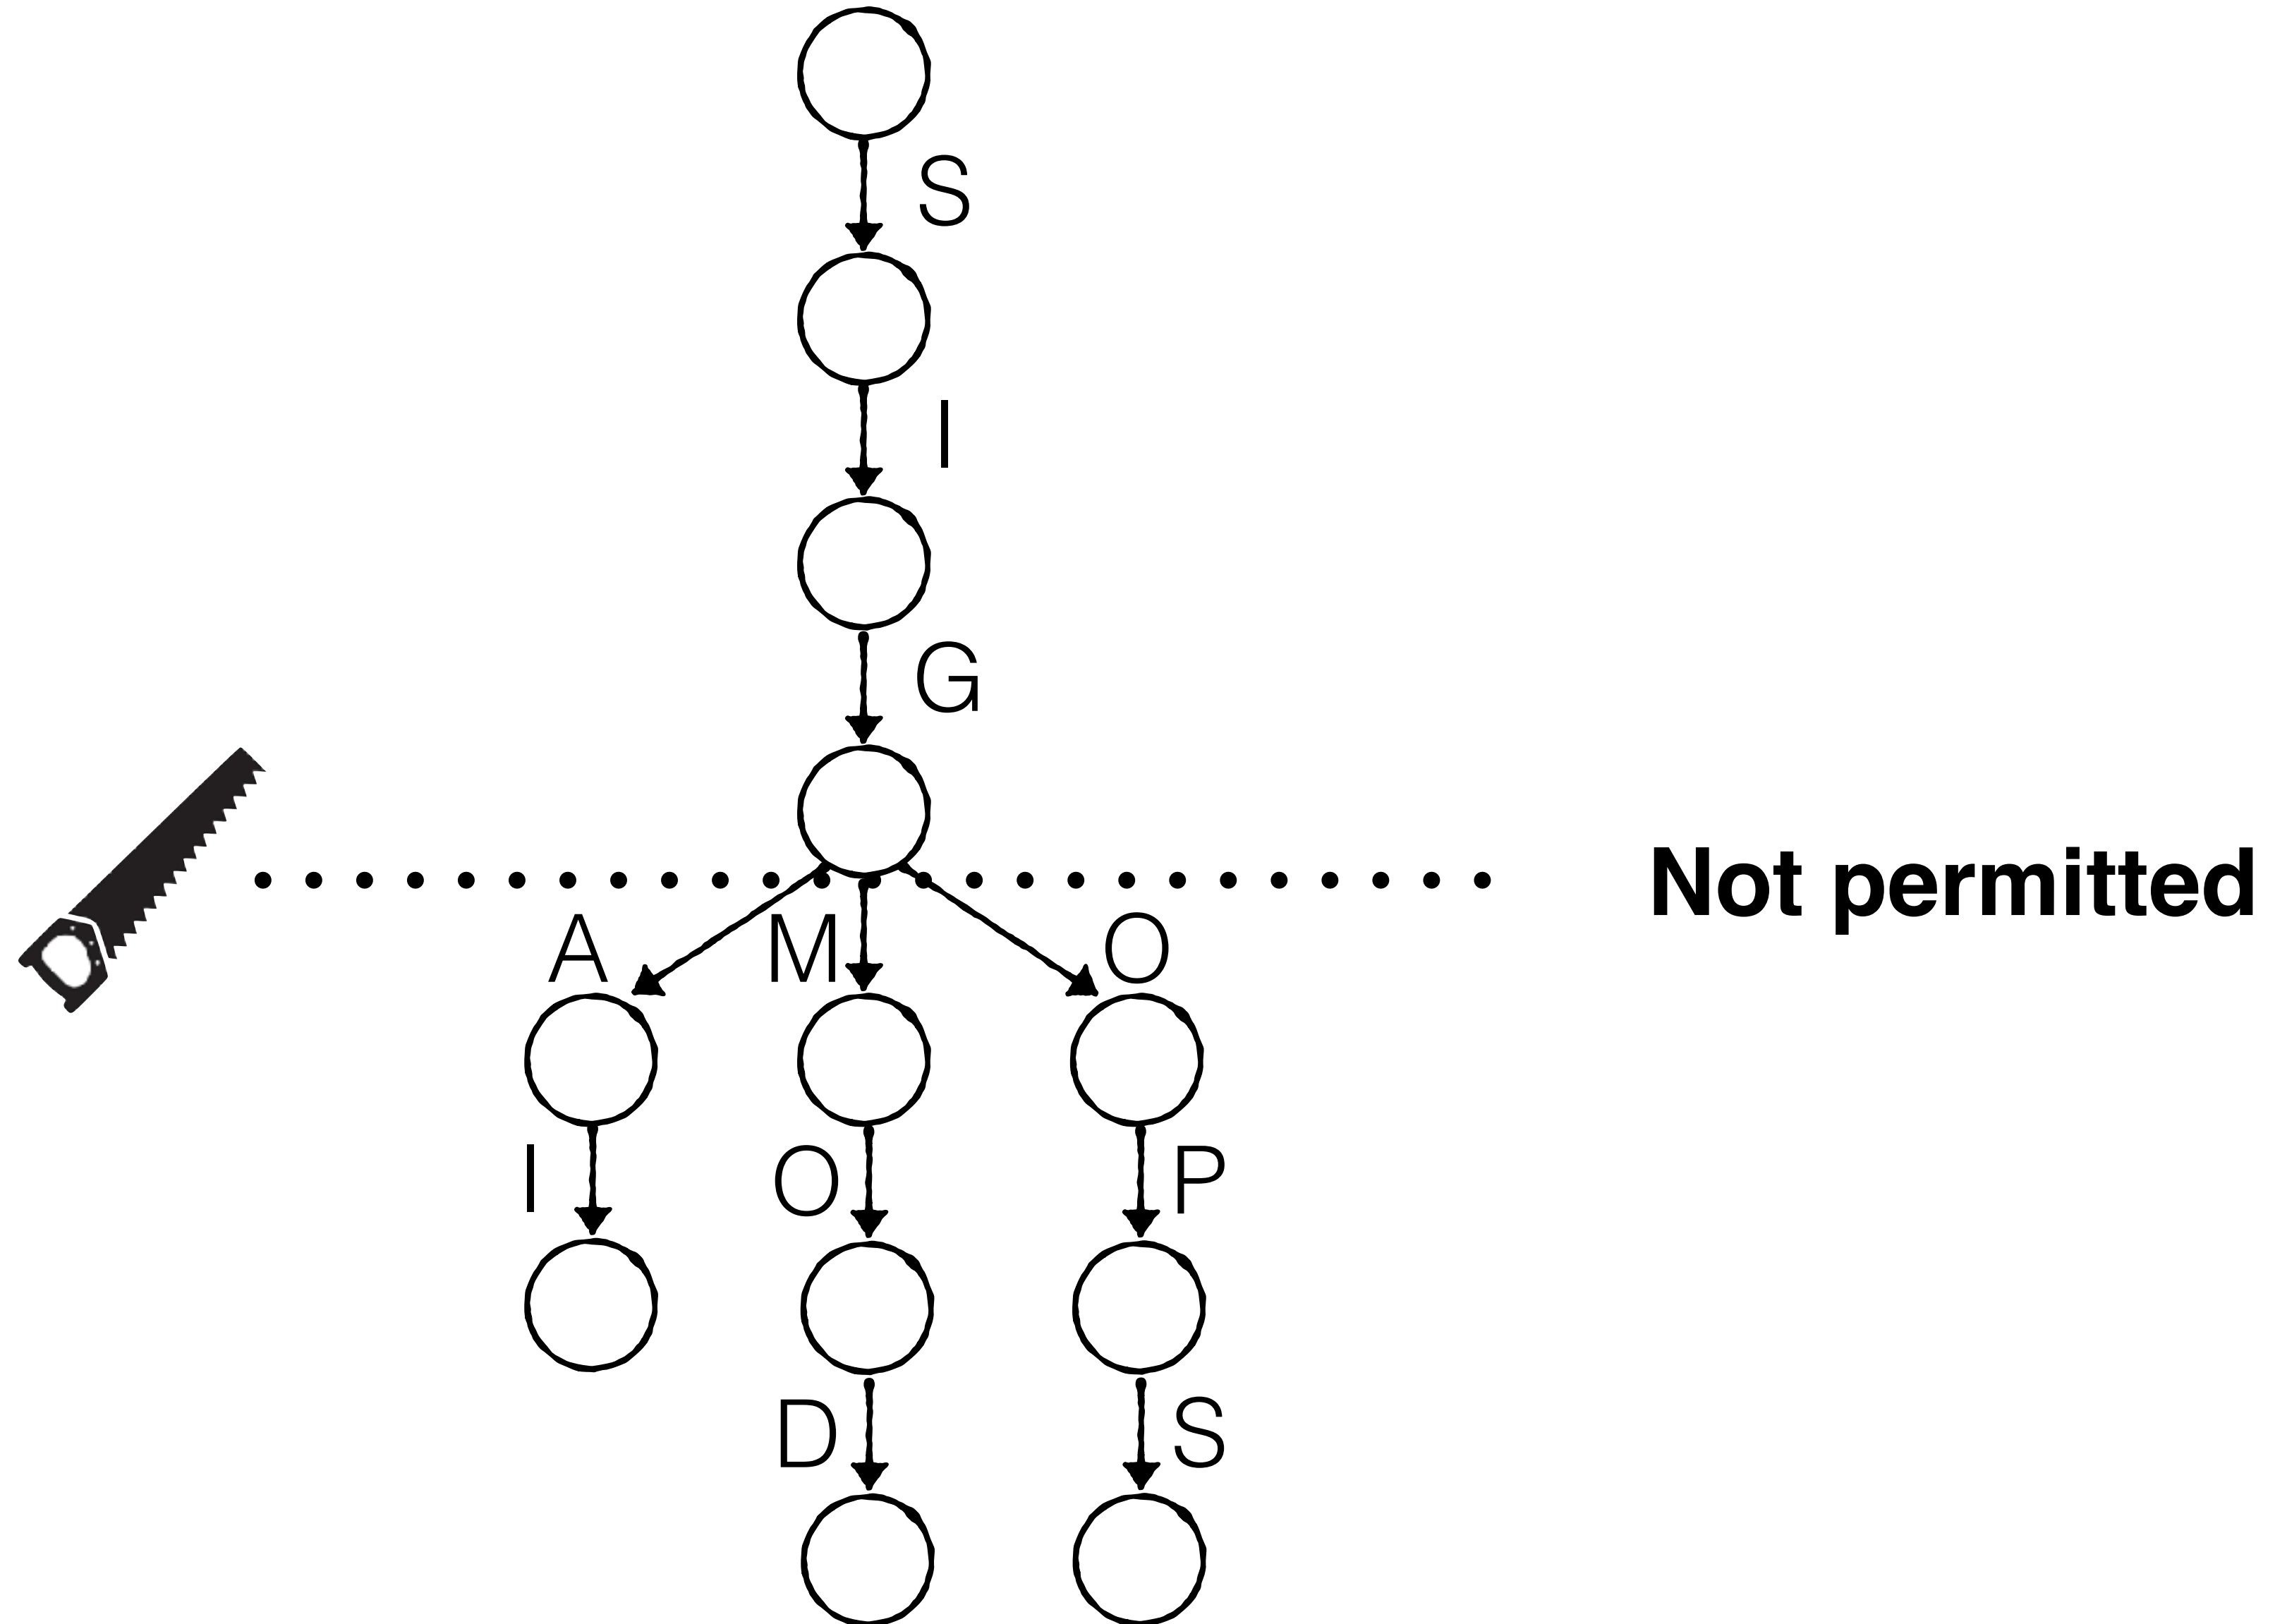
**truncation**




**Keep at least one unique byte for each key**



Keep at least one unique byte for each key

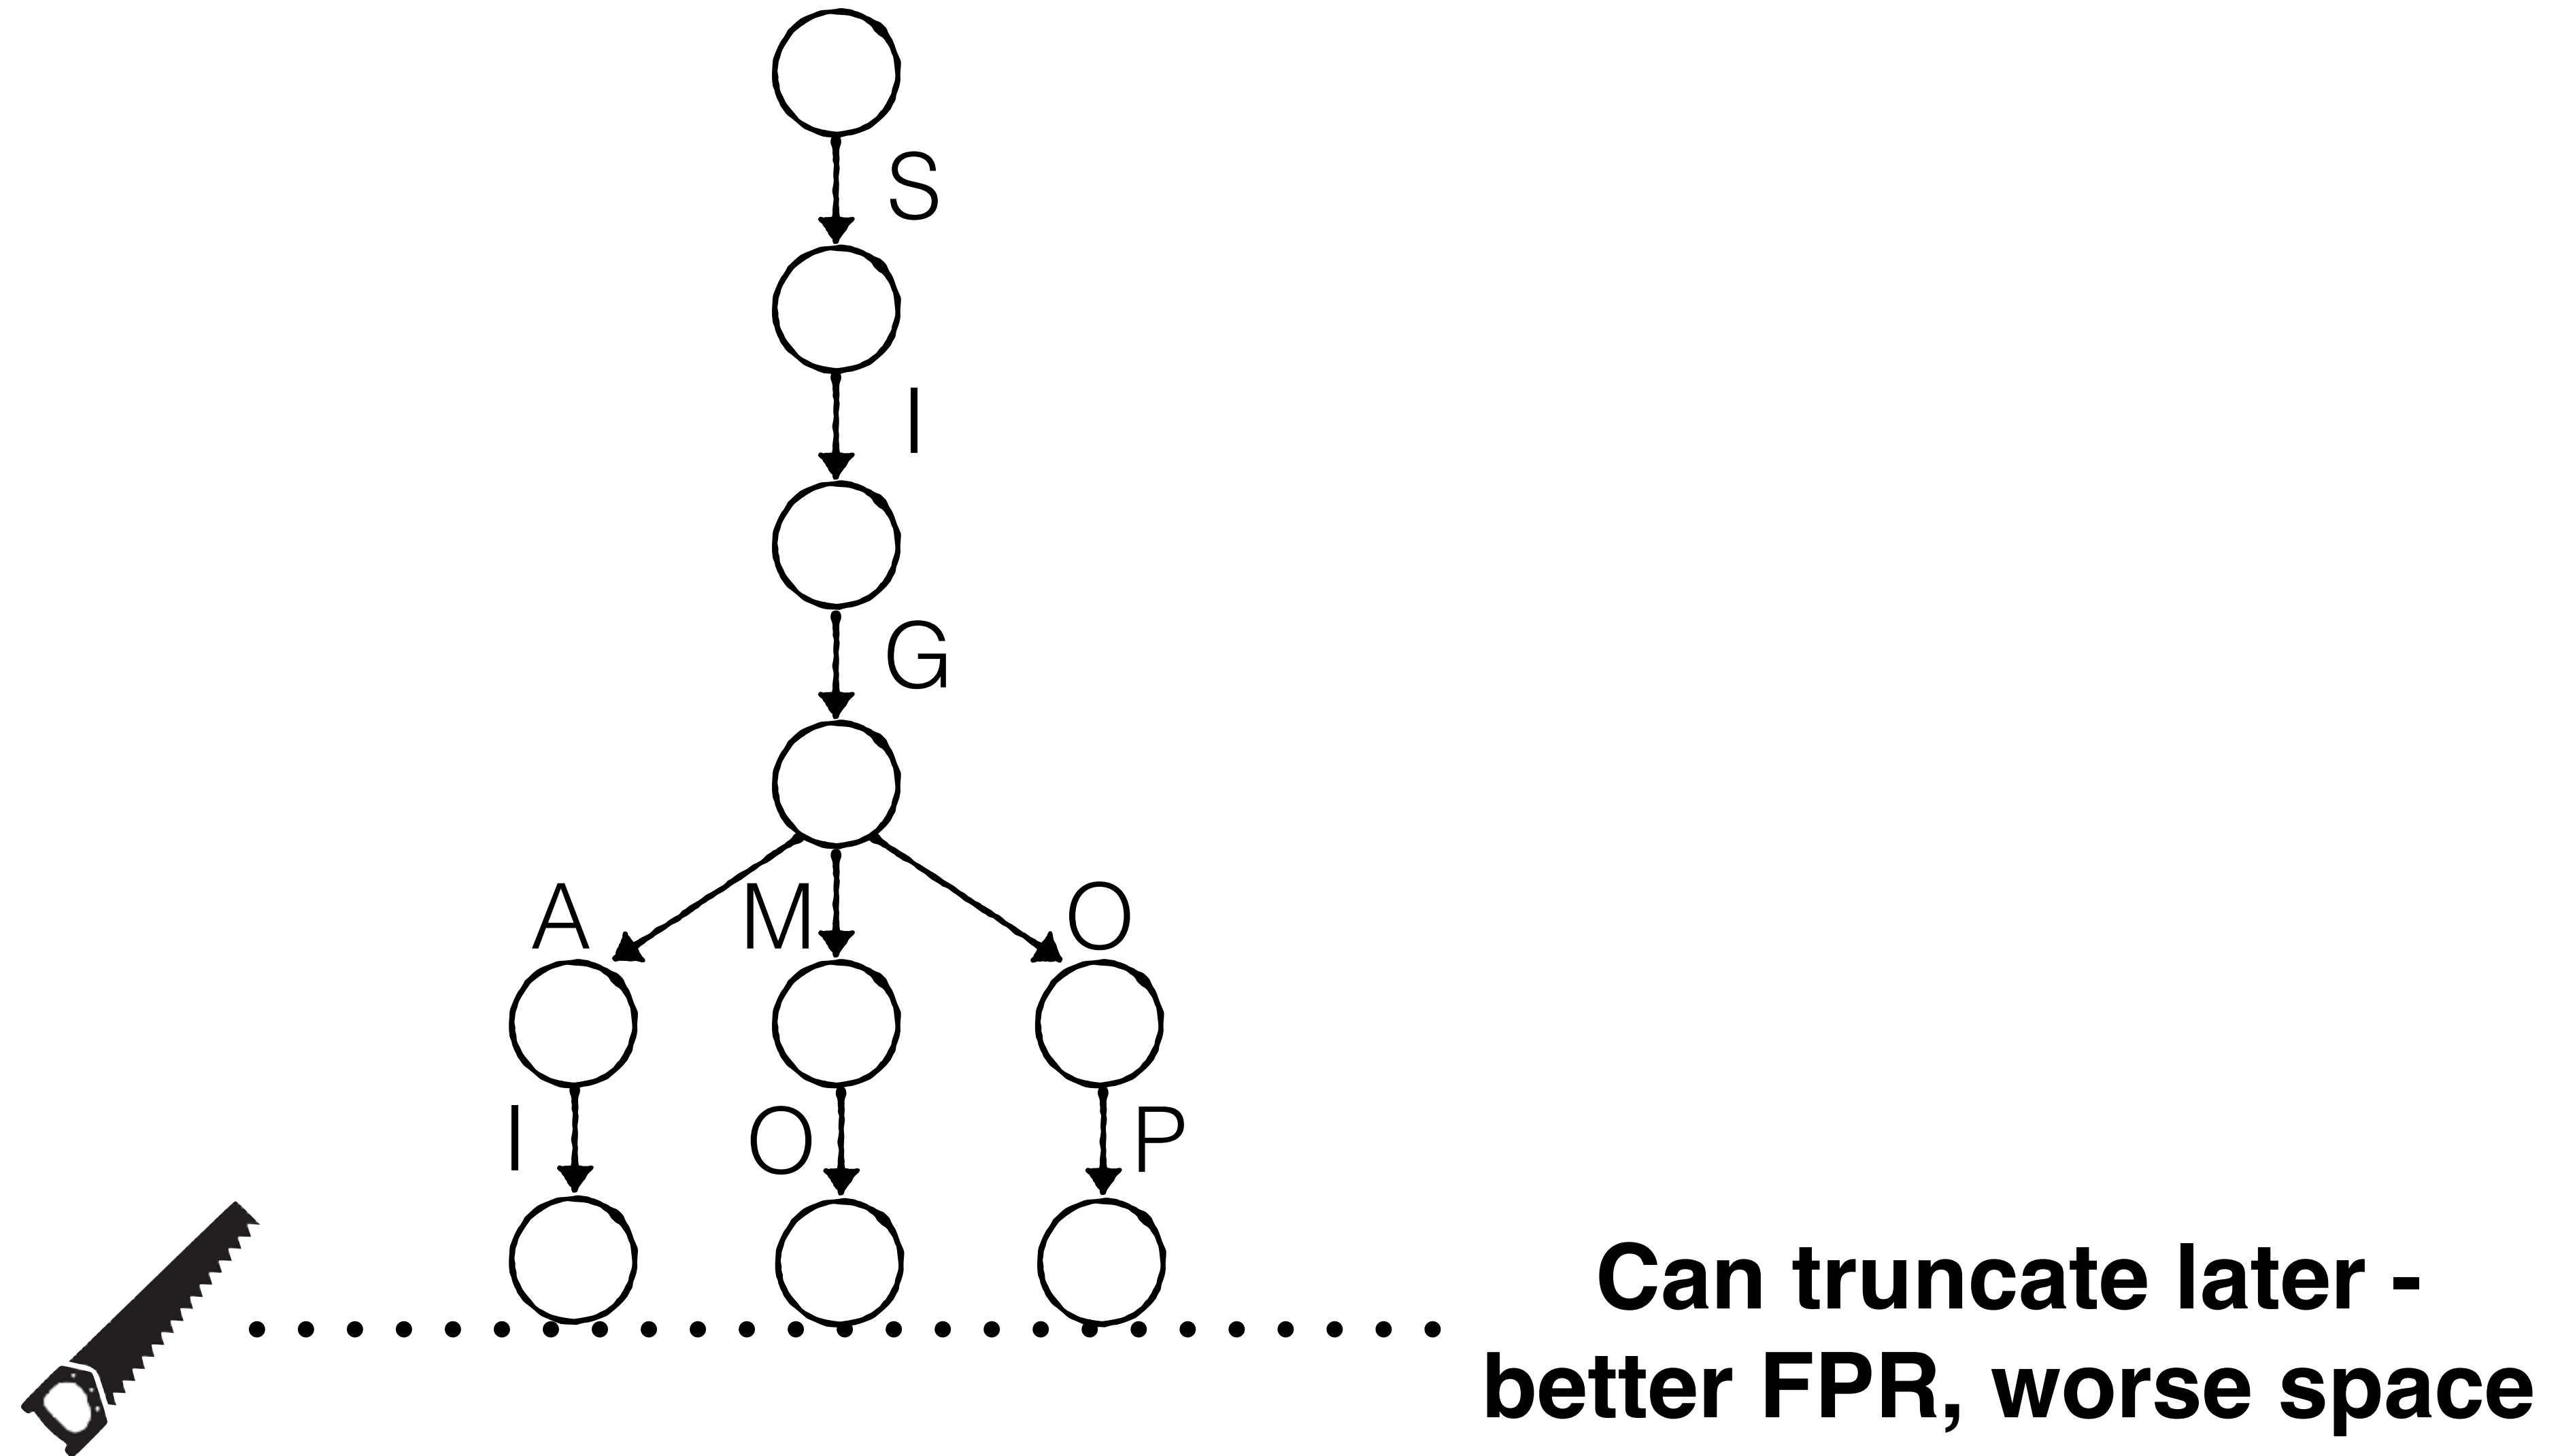


**We now get false positives (e.g., query SIGMETRICS)**




**Match**

We now get false positives (e.g., query SIGMETRICS)




We now get false positives (e.g., query **SIGMETRICS**)




**Negative :)**

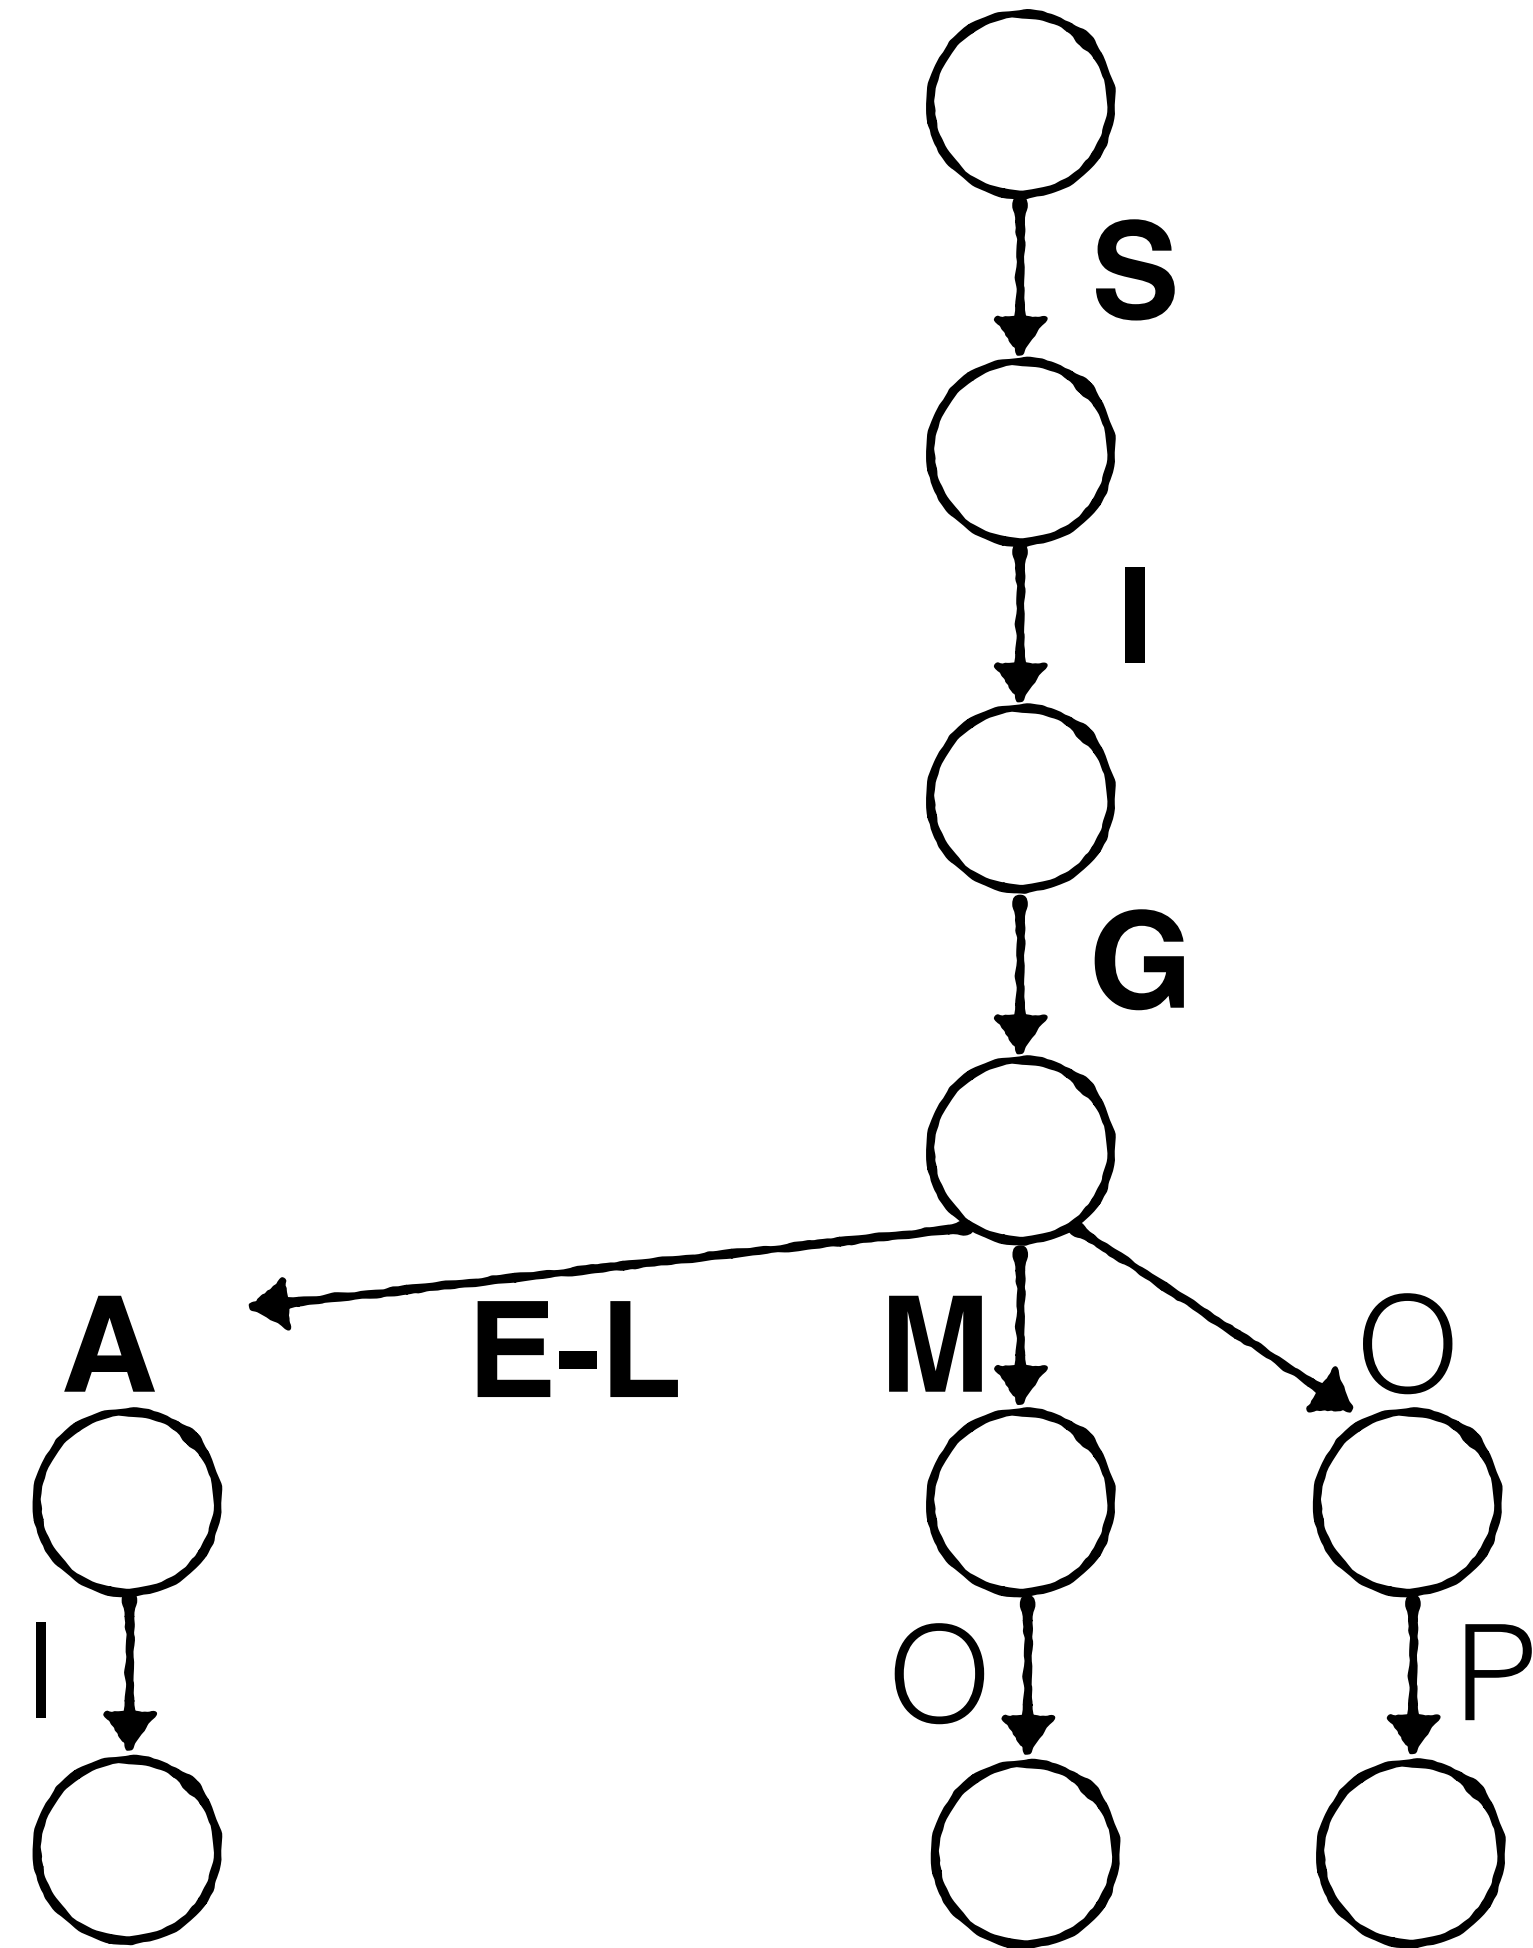
# Range query: SIGE - SIGL



Range query: SIG**E** - SIG**L**




Range query: SIG**E** - SIG**L**




**Negative :)**


Range query: **SIGMOR - SIGMOS**



Range query: **SIGMOR - SIGMOS**




**If many queries match existing prefixes, the FPR increases**



e.g., **SIGMA**  
**SIGMOID**


If many queries match existing prefixes, the FPR increases



e.g., SIGMA  
SIGMOID

**Correlated Workload**

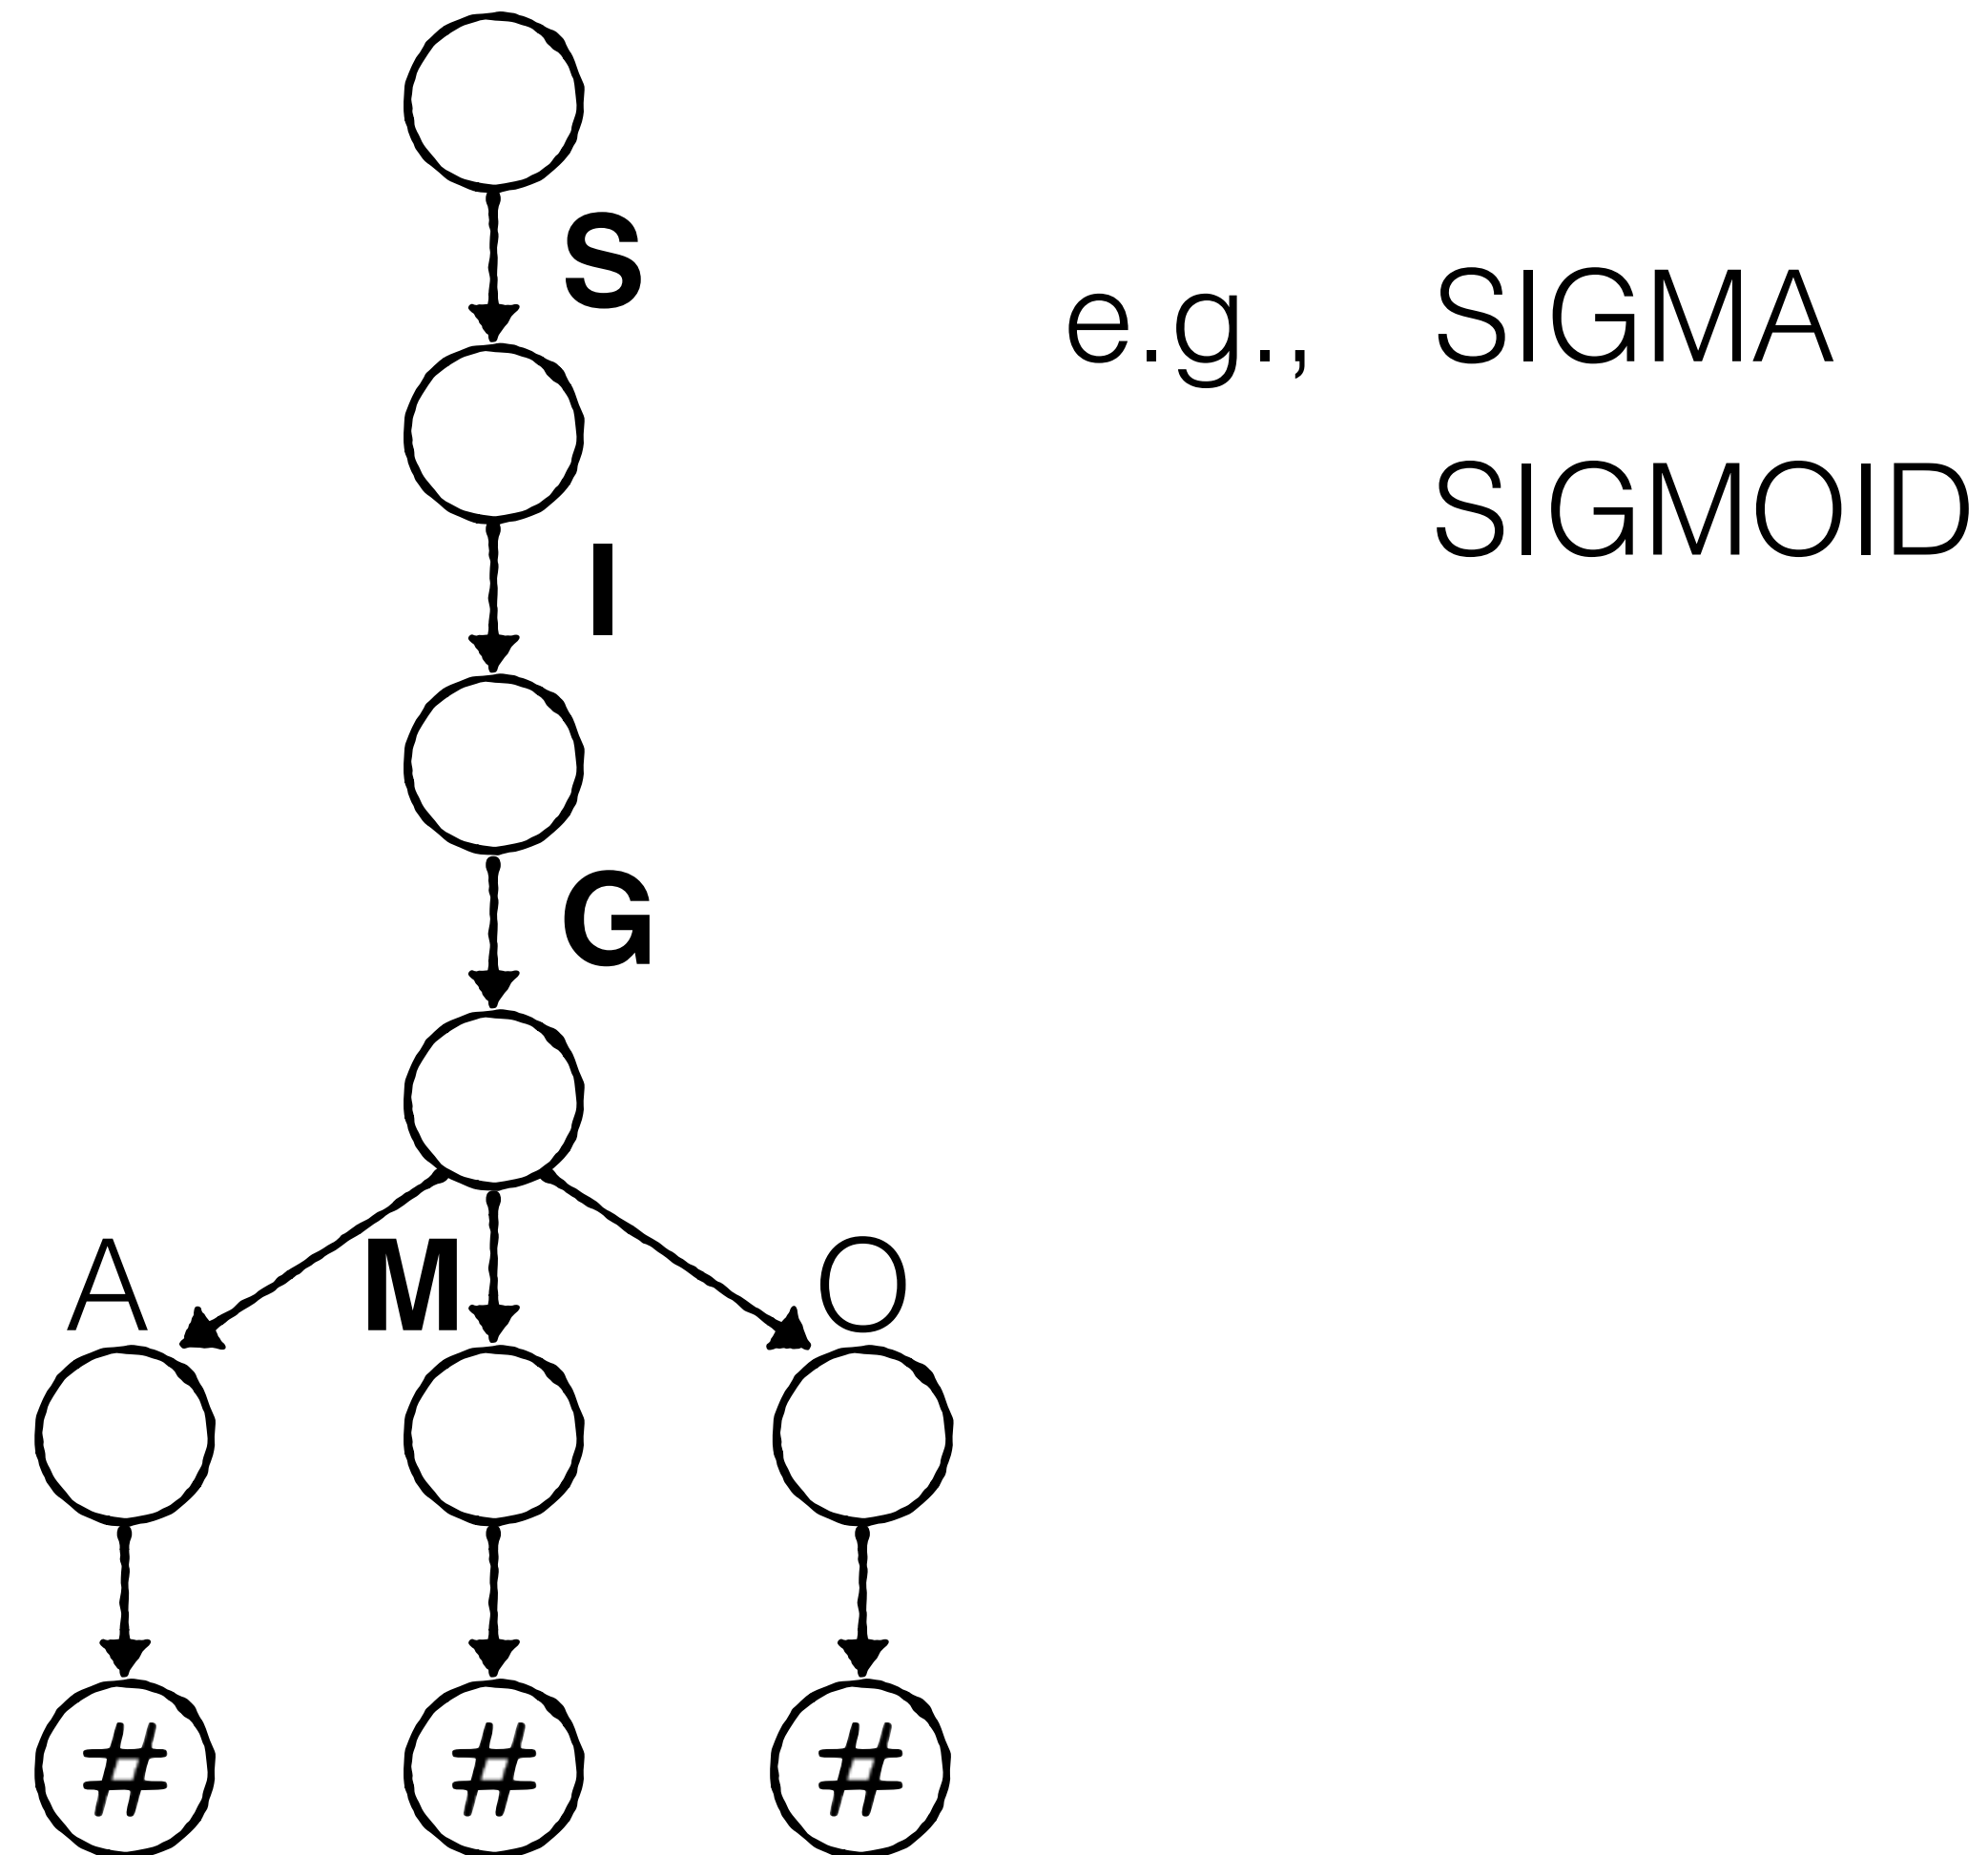
If many queries match existing prefixes, the FPR increases



Correlated Workload

**Can we alleviate this problem for point queries?**


If many queries match existing prefixes, the FPR increases




Can we alleviate this problem for point queries?

**Add 1 byte hash**


# How to encode trie logically?






**F, FAR, FAS, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY**

**Each key leads to leaf payload (e.g., pointer and/or hash)**




F, FAR, FAS, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY

**S leads to leaf: full key**  
**not prefix of any other key**




F, FAR, FAS, FAST, FAT, **S**, TOP, TOY, TRIE, TRIP, TRY

**T leads to internal node: not full key  
prefix of one or more other keys**




F, FAR, FAS, FAST, FAT, S, **TOP**, **TOY**, **TRIE**, **TRIP**, **TRY**

**F leads to internal node with \$ child: full key  
prefix of one or more other keys**




**F, FAR, FAS, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY**

F leads to internal node with \$ child: full key  
prefix of one or more other keys



**F, FAR, FAS, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY**

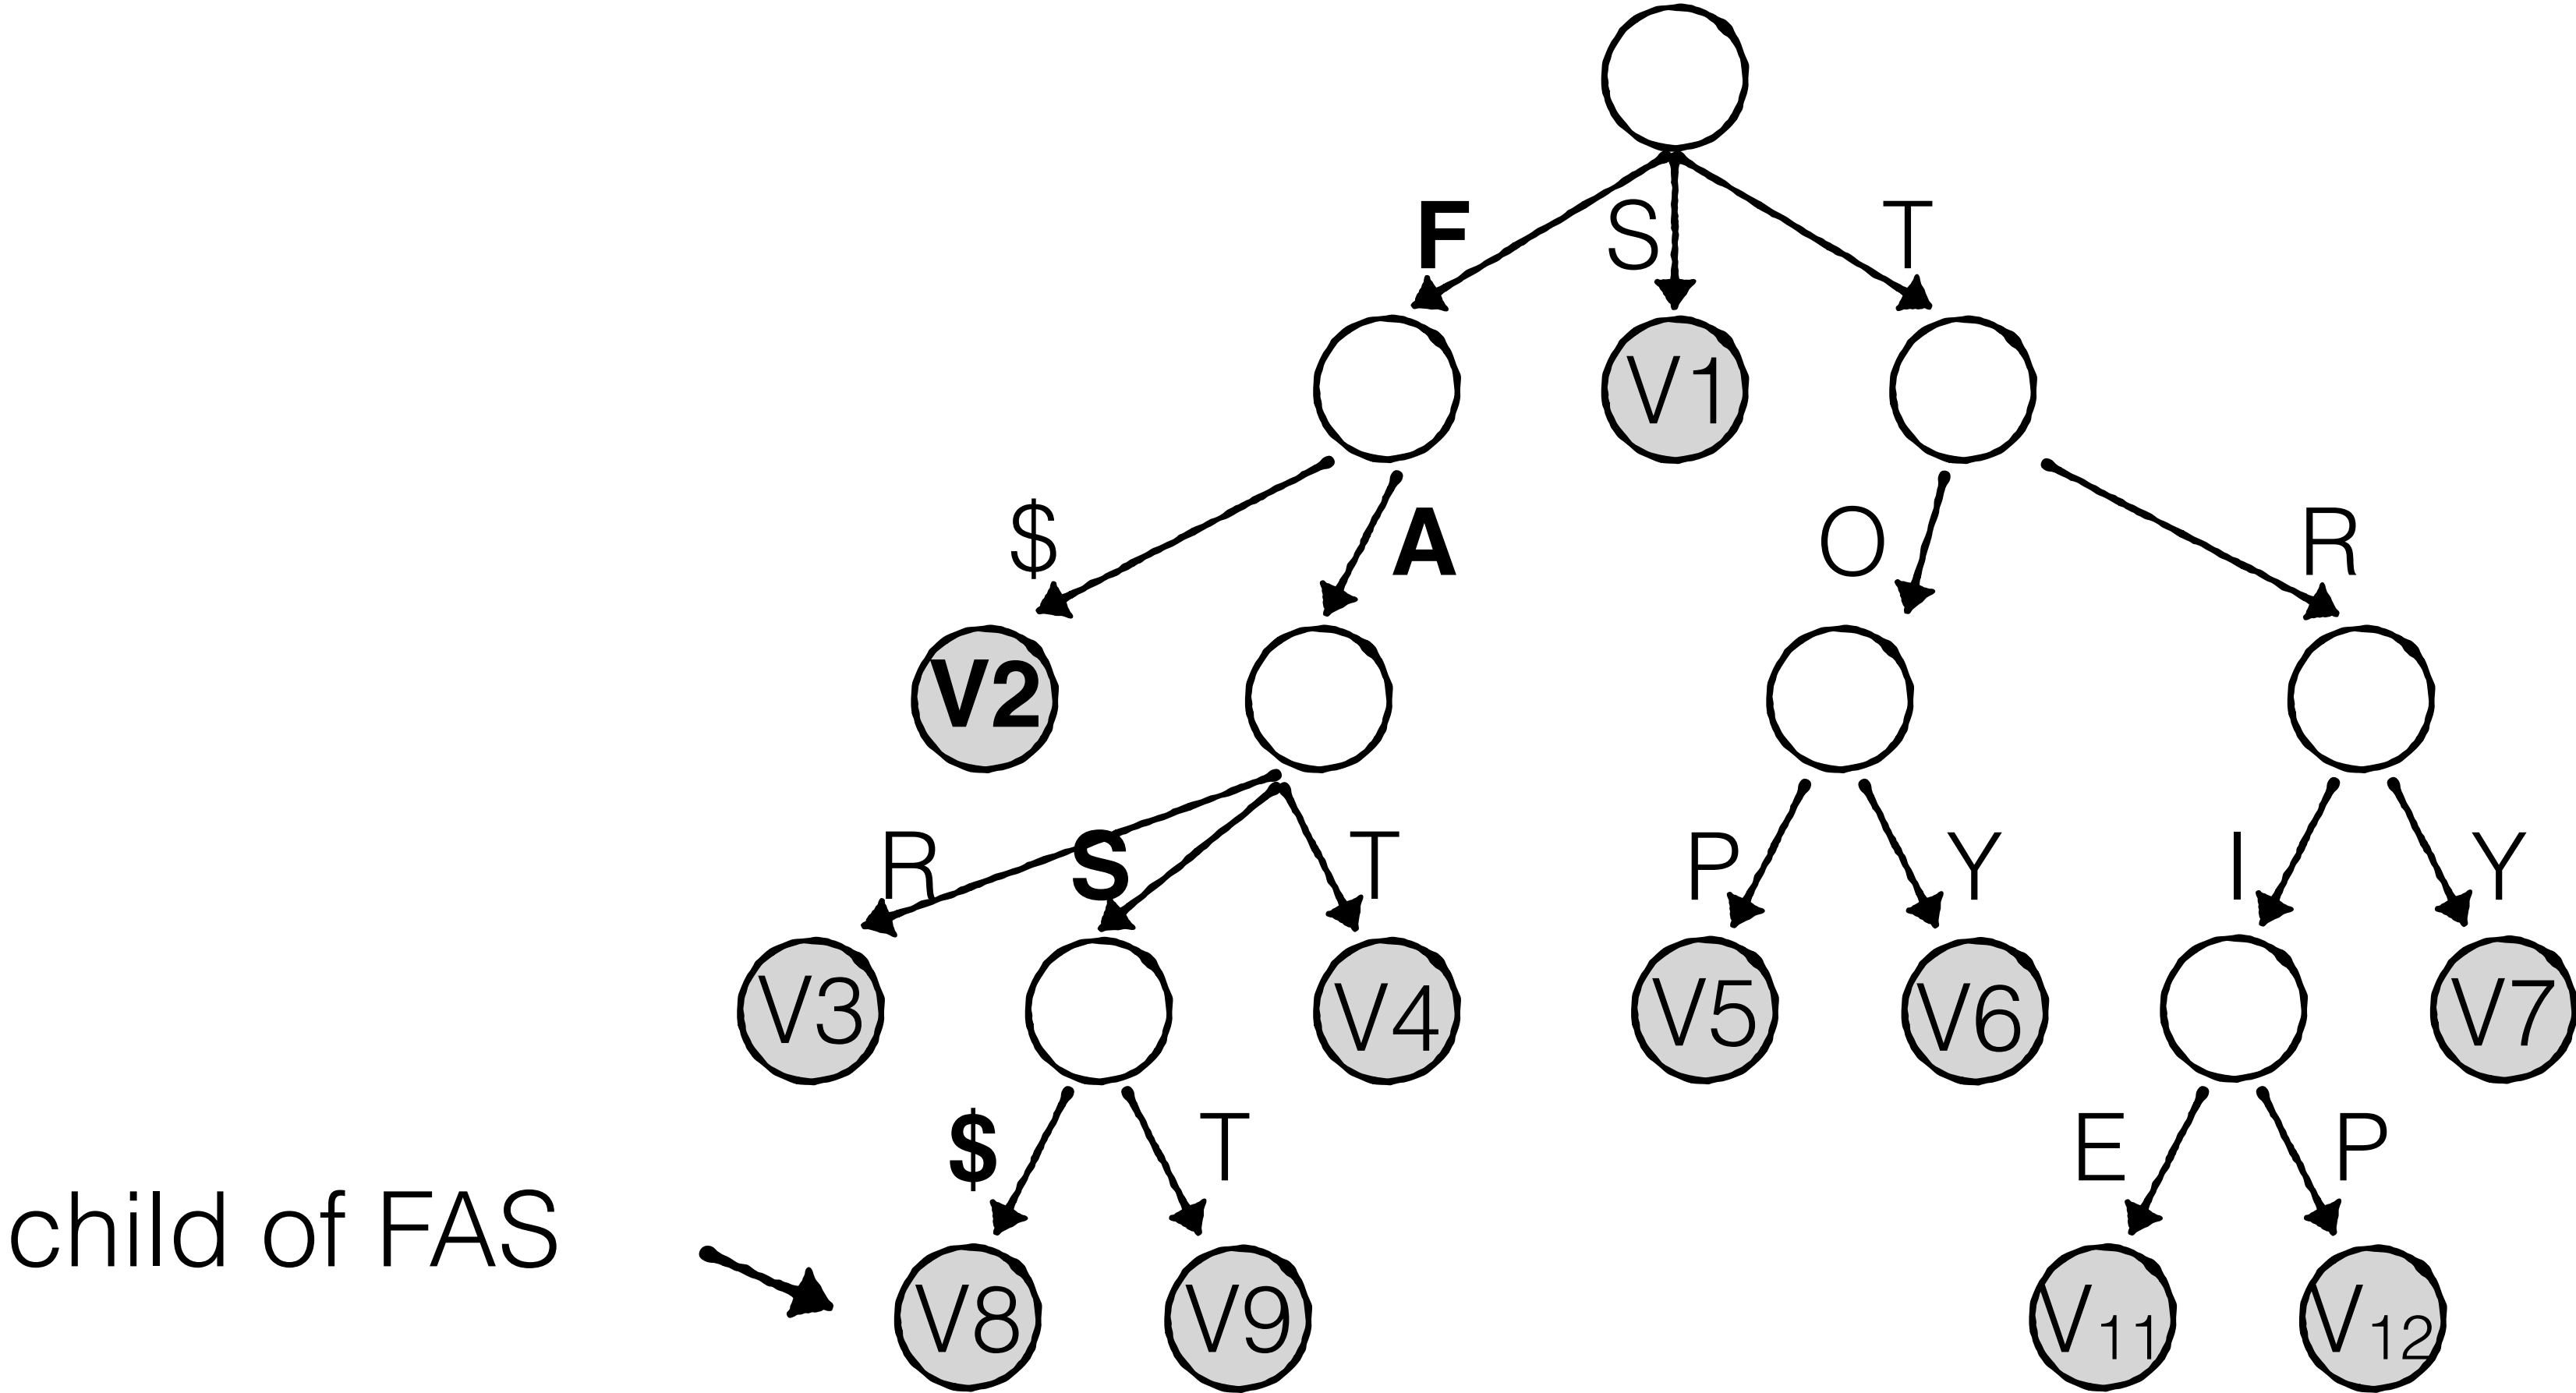
**FAS** leads to internal node with \$ child: full key  
prefix of one or more other keys



child of FAS


F, FAR, **FAS**, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY

# A prefix is always represented by an internal node




F, FAR, **FAS**, FAST, FAT, S, TOP, TOY, TRIE, TRIP, TRY

# How to encode physically and succinctly?




How to encode physically and succinctly?



**LOUDS: Level-Order Unary Degree Sequence**


How to encode physically and succinctly?




LOUDS: Level-Order Unary Degree Sequence


**Variation(!)**

# Label internal nodes in breadth-first order for clarity






**Internal nodes:      0      1      2      3      4      5      6      7**



<b>Internal nodes:</b>	<b>0</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>7</b>
<b>LOUDS:</b>	<b>100</b>	<b>10</b>	<b>100</b>	<b>100</b>	<b>10</b>	<b>10</b>	<b>10</b>	<b>10</b>




Internal nodes:

0    1    2    3    4    5    6    7

LOUDS:

100   10   100   100   10   10   10   10

**2 bits per node : )**




Internal nodes:	0	1	2	3	4	5	6	7
LOUDS:	100	10	100	100	10	10	10	10

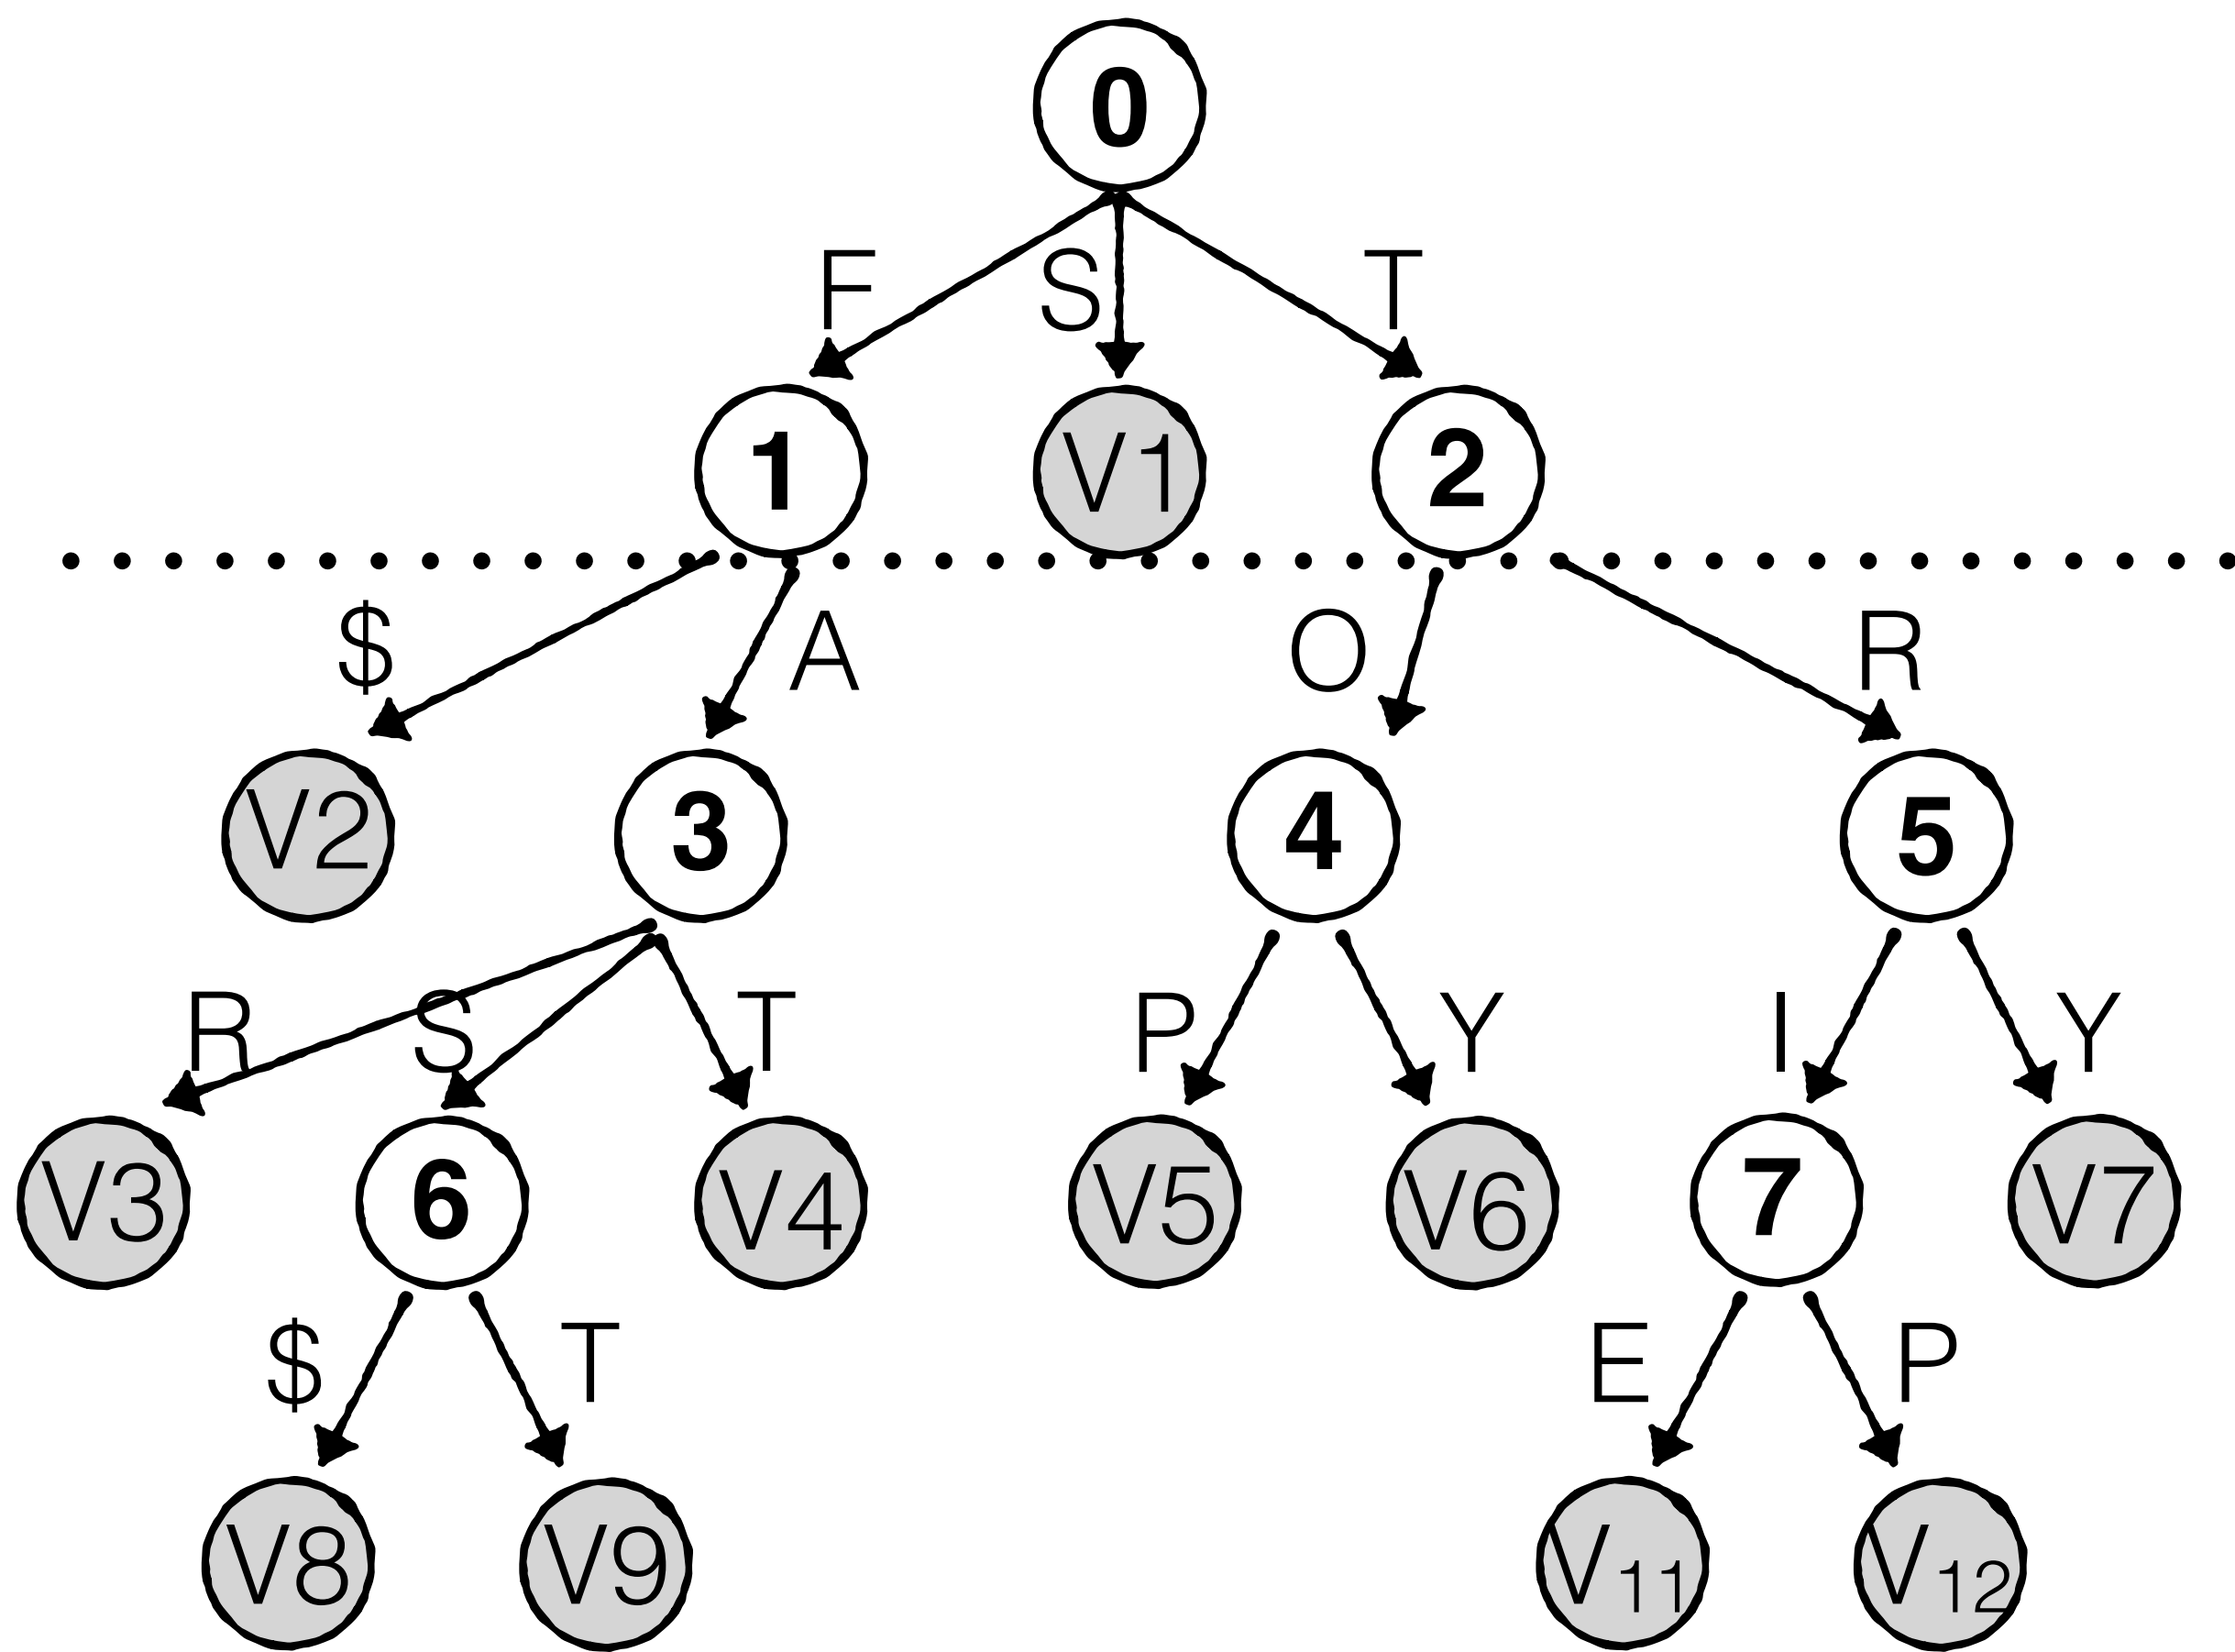
**2 bits per node : )**

**Also fast to traverse using rank/select (shortly)**

**But we also need to persist characters and values along paths...**



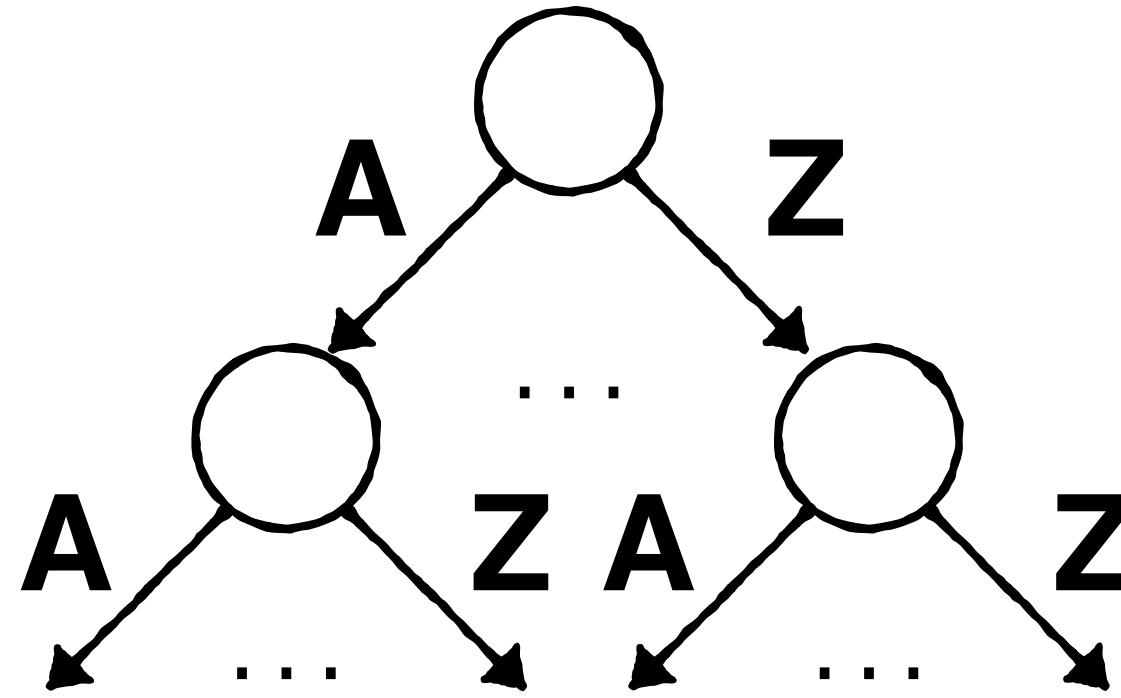
Internal nodes:	0	1	2	3	4	5	6	7
LOUDS:	100	10	100	100	10	10	10	10



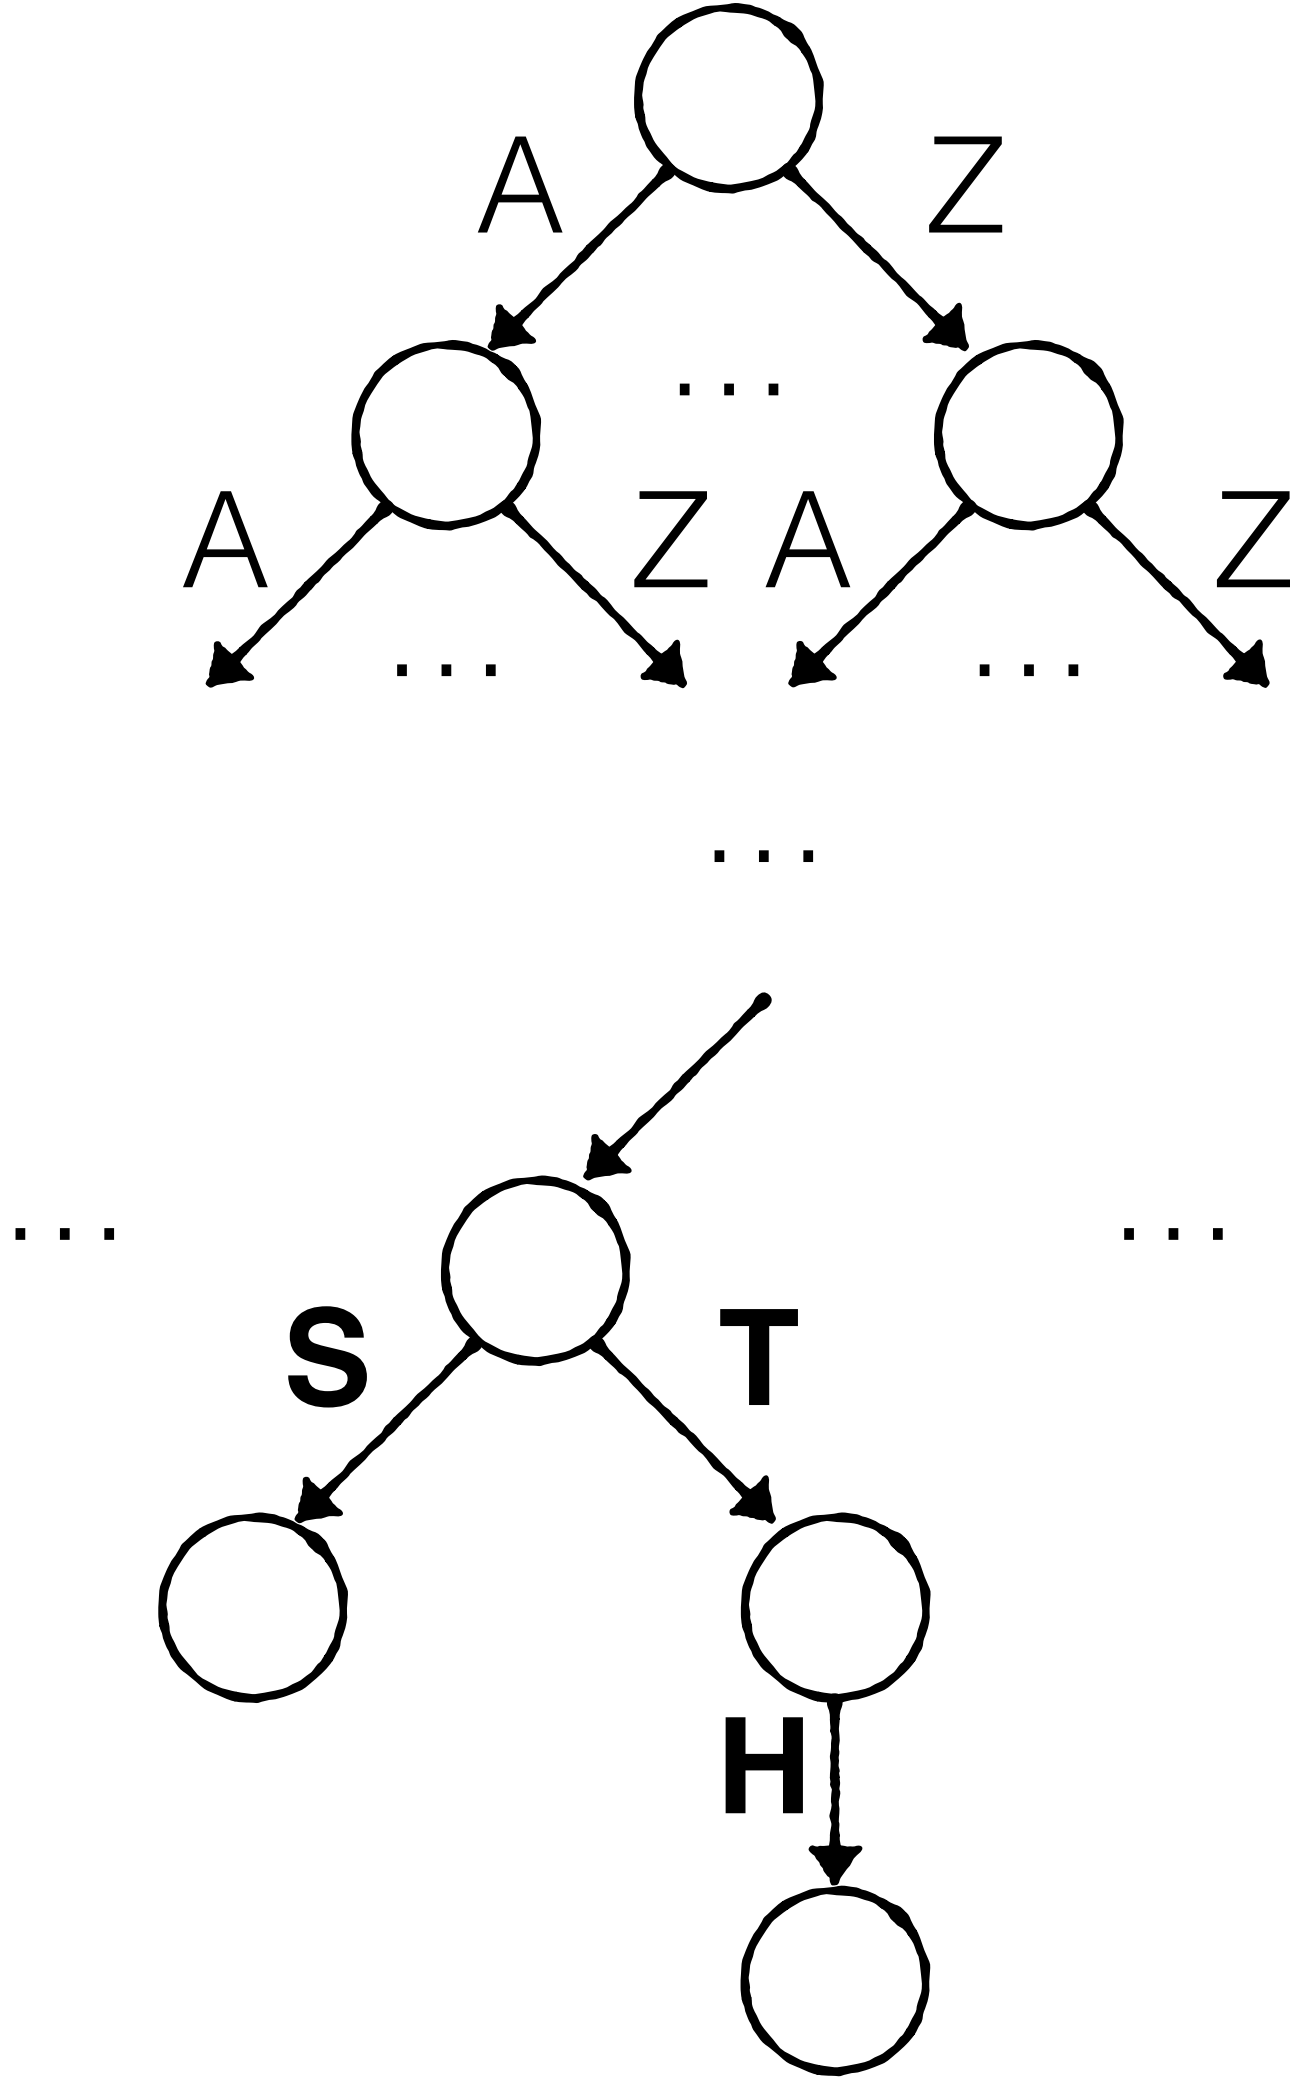
**Dense encoding**

**Sparse encoding**


**Typically, the upper levels of the trie are more full**




**Typically, the upper levels of the trie are more full**



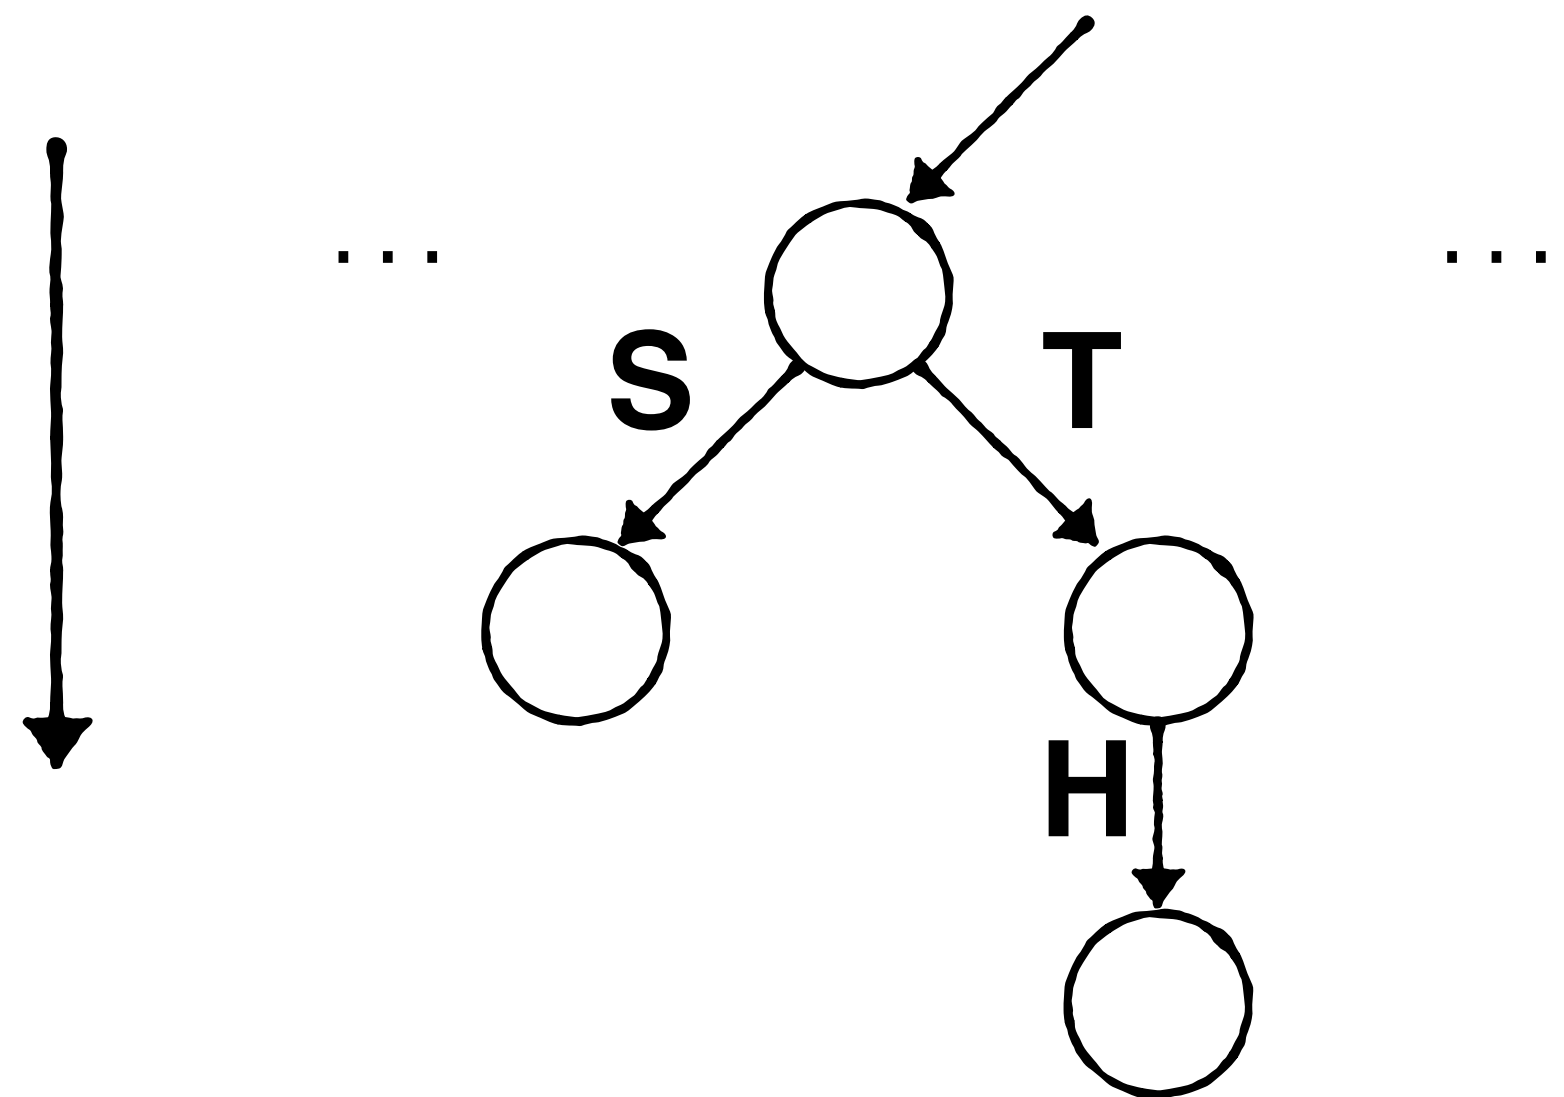
**More queries go  
through base layers**




More queries go through base layers




**A leaf is reached on avg. by fewer queries**




More queries go through base layers



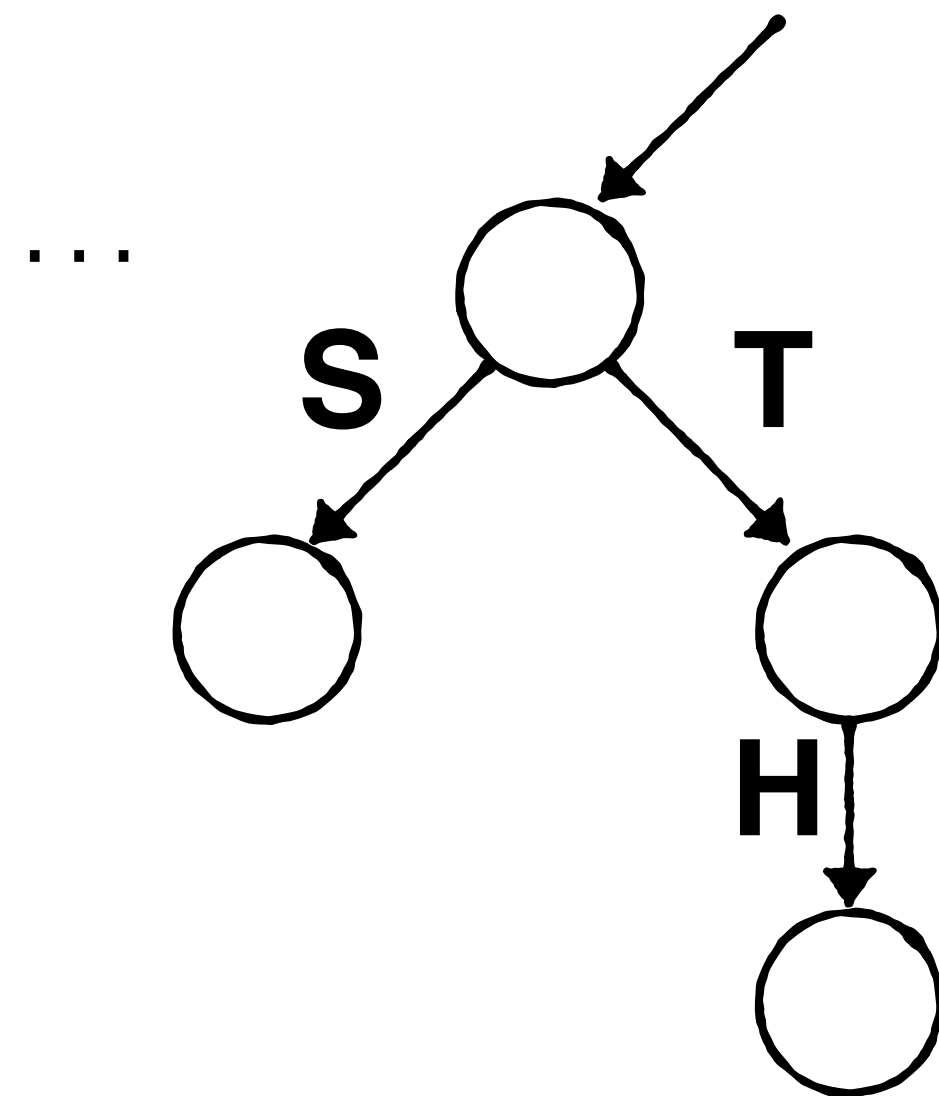
A leaf is reached on avg. by fewer queries




**Exponentially more nodes as we move down**




**Optimize for speed**




**Optimize for space**




**Edges**




**One bit set for every connected character**





Edges






Edges  
**Has-child**






Edges  
**Has-child**


	F	S	T
Edges			
Has-child			




Edges  
Has-child




**Continue in breadth-first order  
for each internal node**




Edges  
Has-child




Continue in breadth-first order  
for each internal node



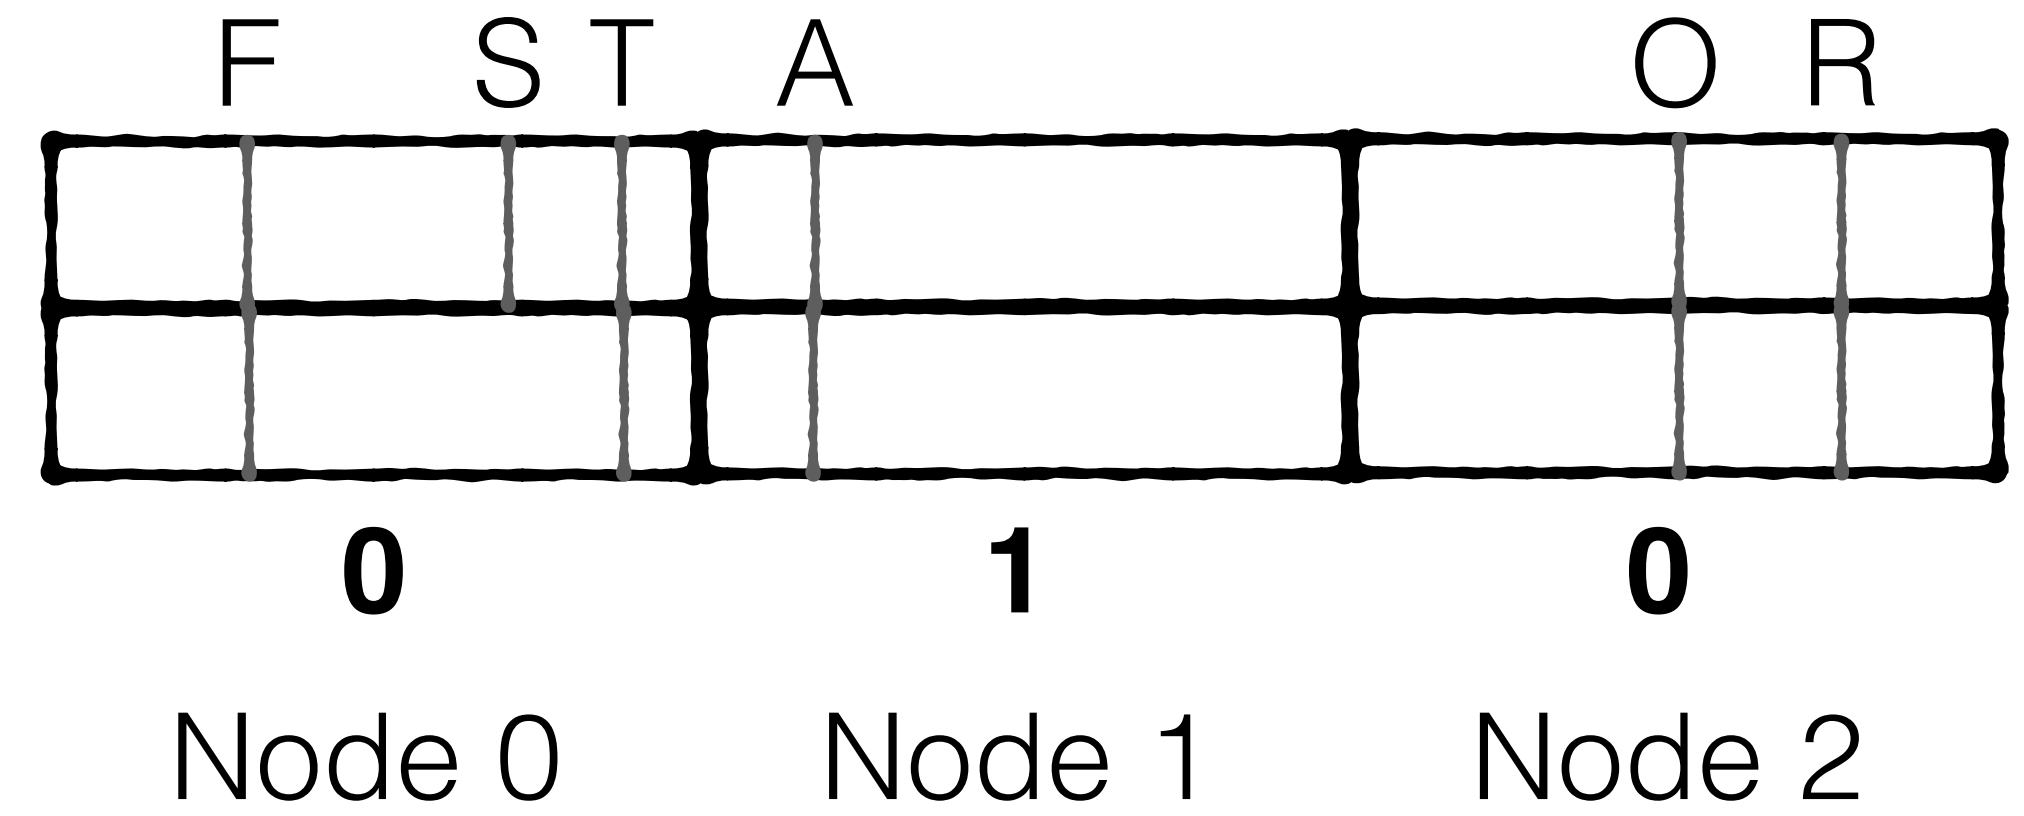
Edges  
Has-child




Continue in breadth-first order  
for each internal node

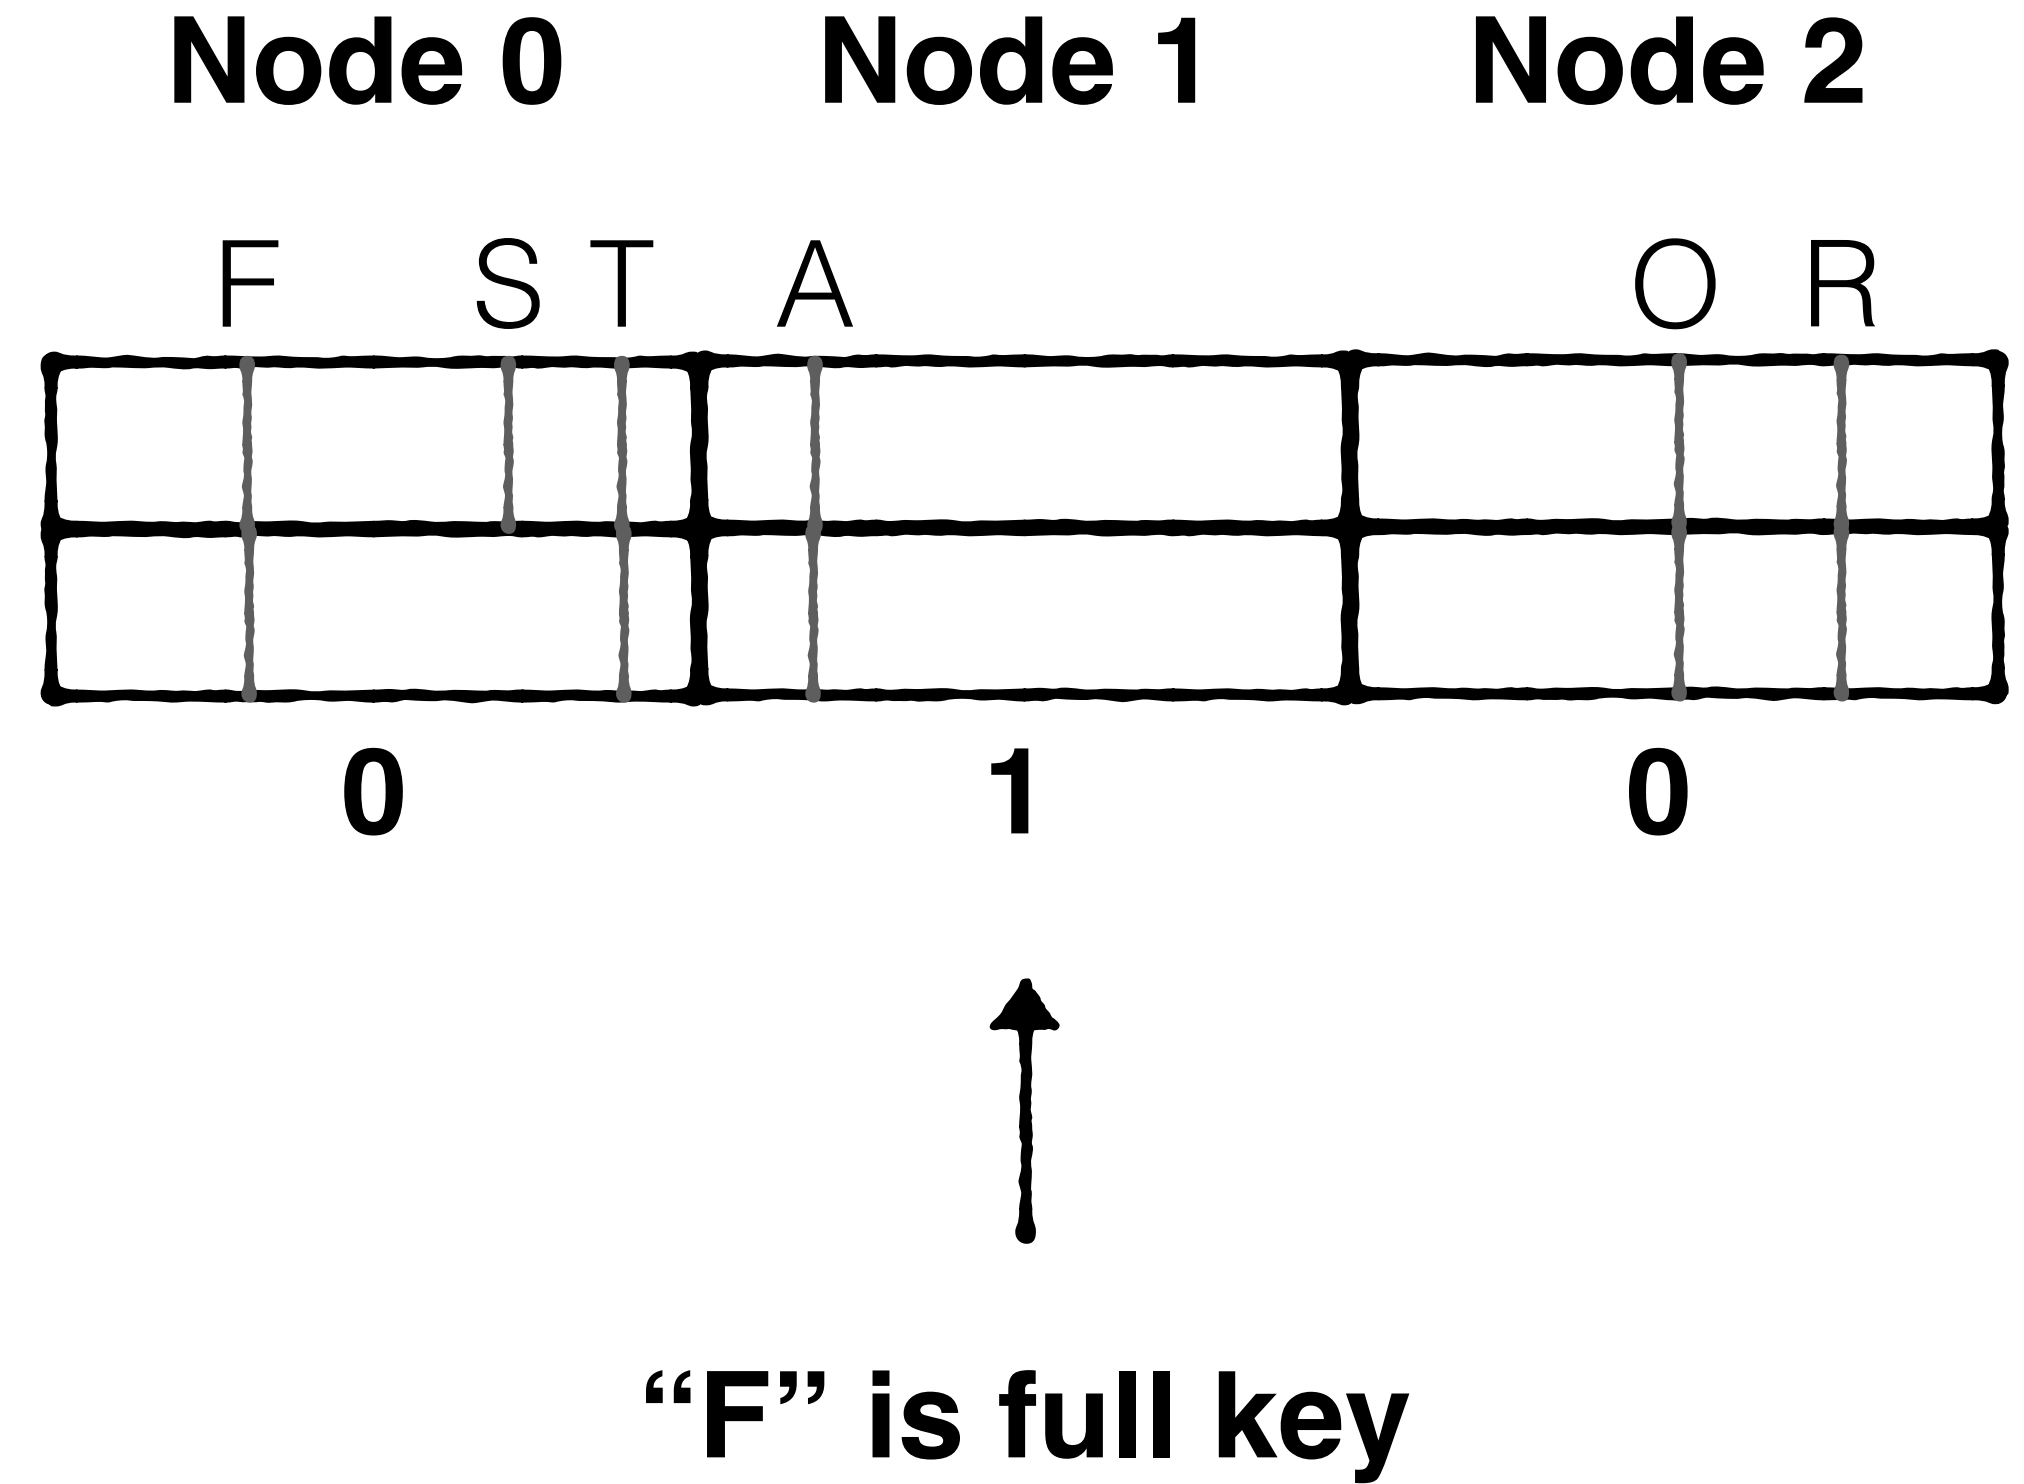



Edges  
Has-child  
**isPrefix**



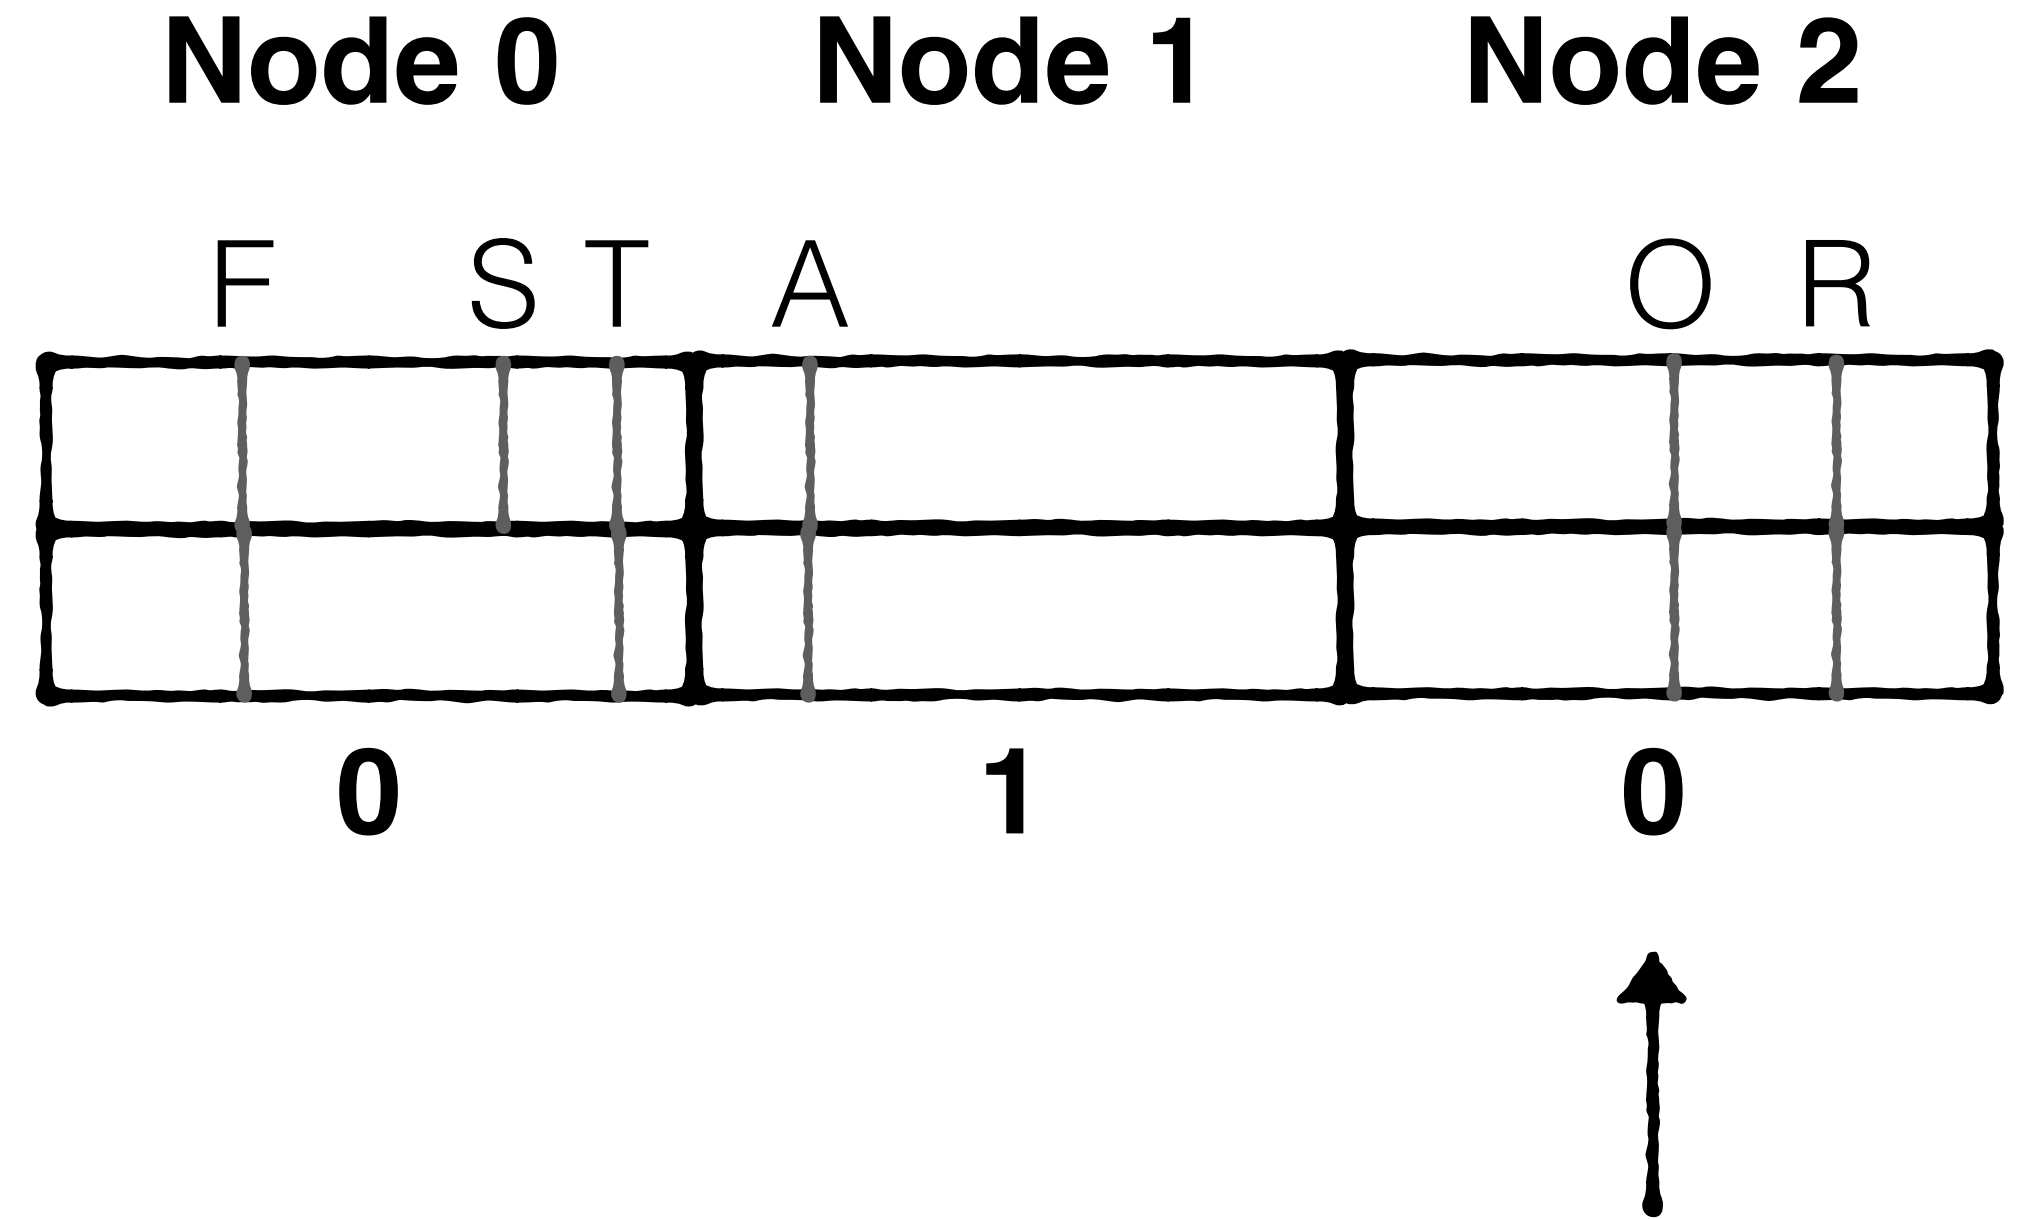


Edges  
Has-child  
**isKey**






Edges  
Has-child  
**isKey**





**“T” is not full key**




Edges  
 Has-child  
 isKey

**Values**





Edges  
Has-child  
isKey




**Values**

**V1**


**V2**


**Stored contiguously in an array**

Does "S" exist?



Edges  
Has-child  
isKey  
Values



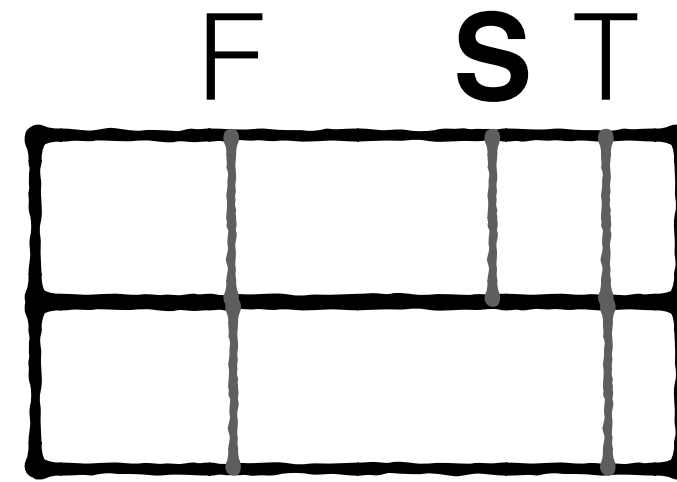


**Edges**  
**Has-child**

isKey

Values


**Does "S" exist?**



0



V1

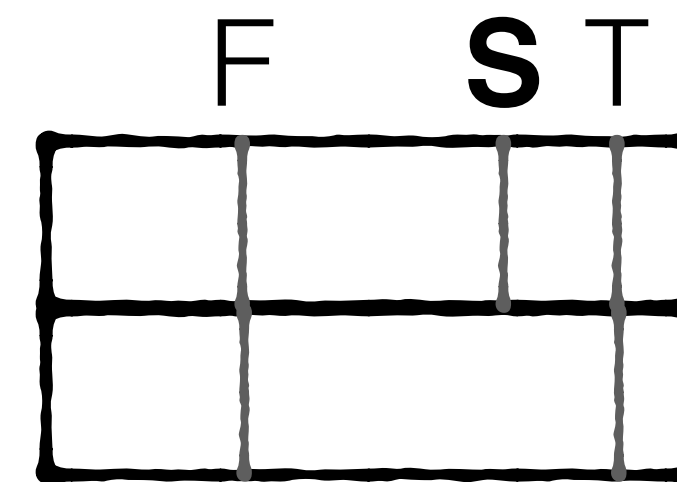


**Edges**  
**Has-child**

isKey

Values

**Does "S" exist?**

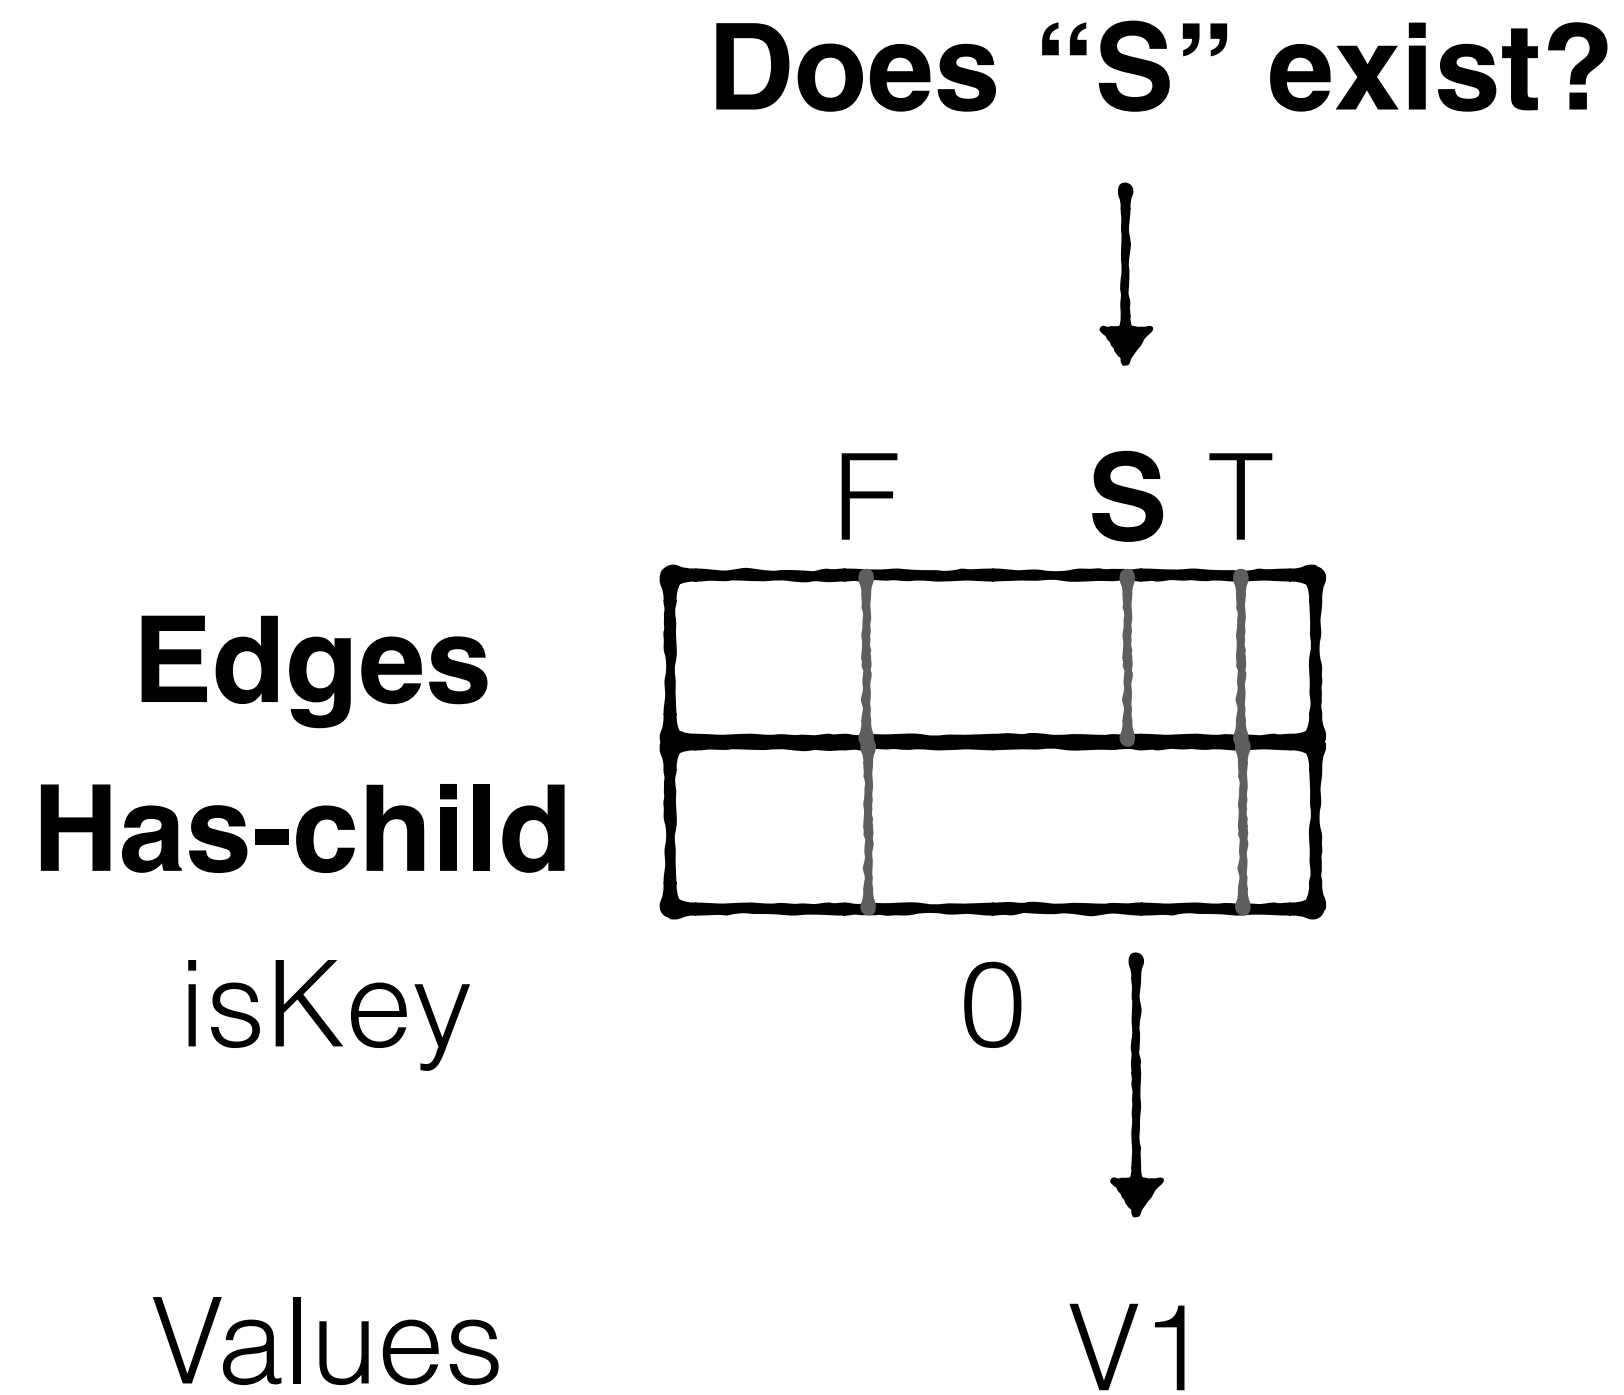



0




V1

**We observe that S exists and has no children, meaning it must be a full key**





We observe that S exists and has no children, meaning it must be a full key

**How to locate its value?**



Edges  
Has-child






**Rank(Edges, S) = 1**





S





**Rank(Has-child, S) = 1**

$$S \text{ value offset} = \mathbf{Rank}(\text{Edges}, S) - \mathbf{Rank}(\text{Has-child}, S) = \mathbf{0}$$





$$S \text{ value offset} = \text{Rank}(\text{Edges}, S) - \text{Rank}(\text{Has-child}, S) = 0$$




**Constant time per each node access :)**

$$S \text{ value offset} = \text{Rank}(\text{Edges}, S) - \text{Rank}(\text{Has-child}, S) = 0$$



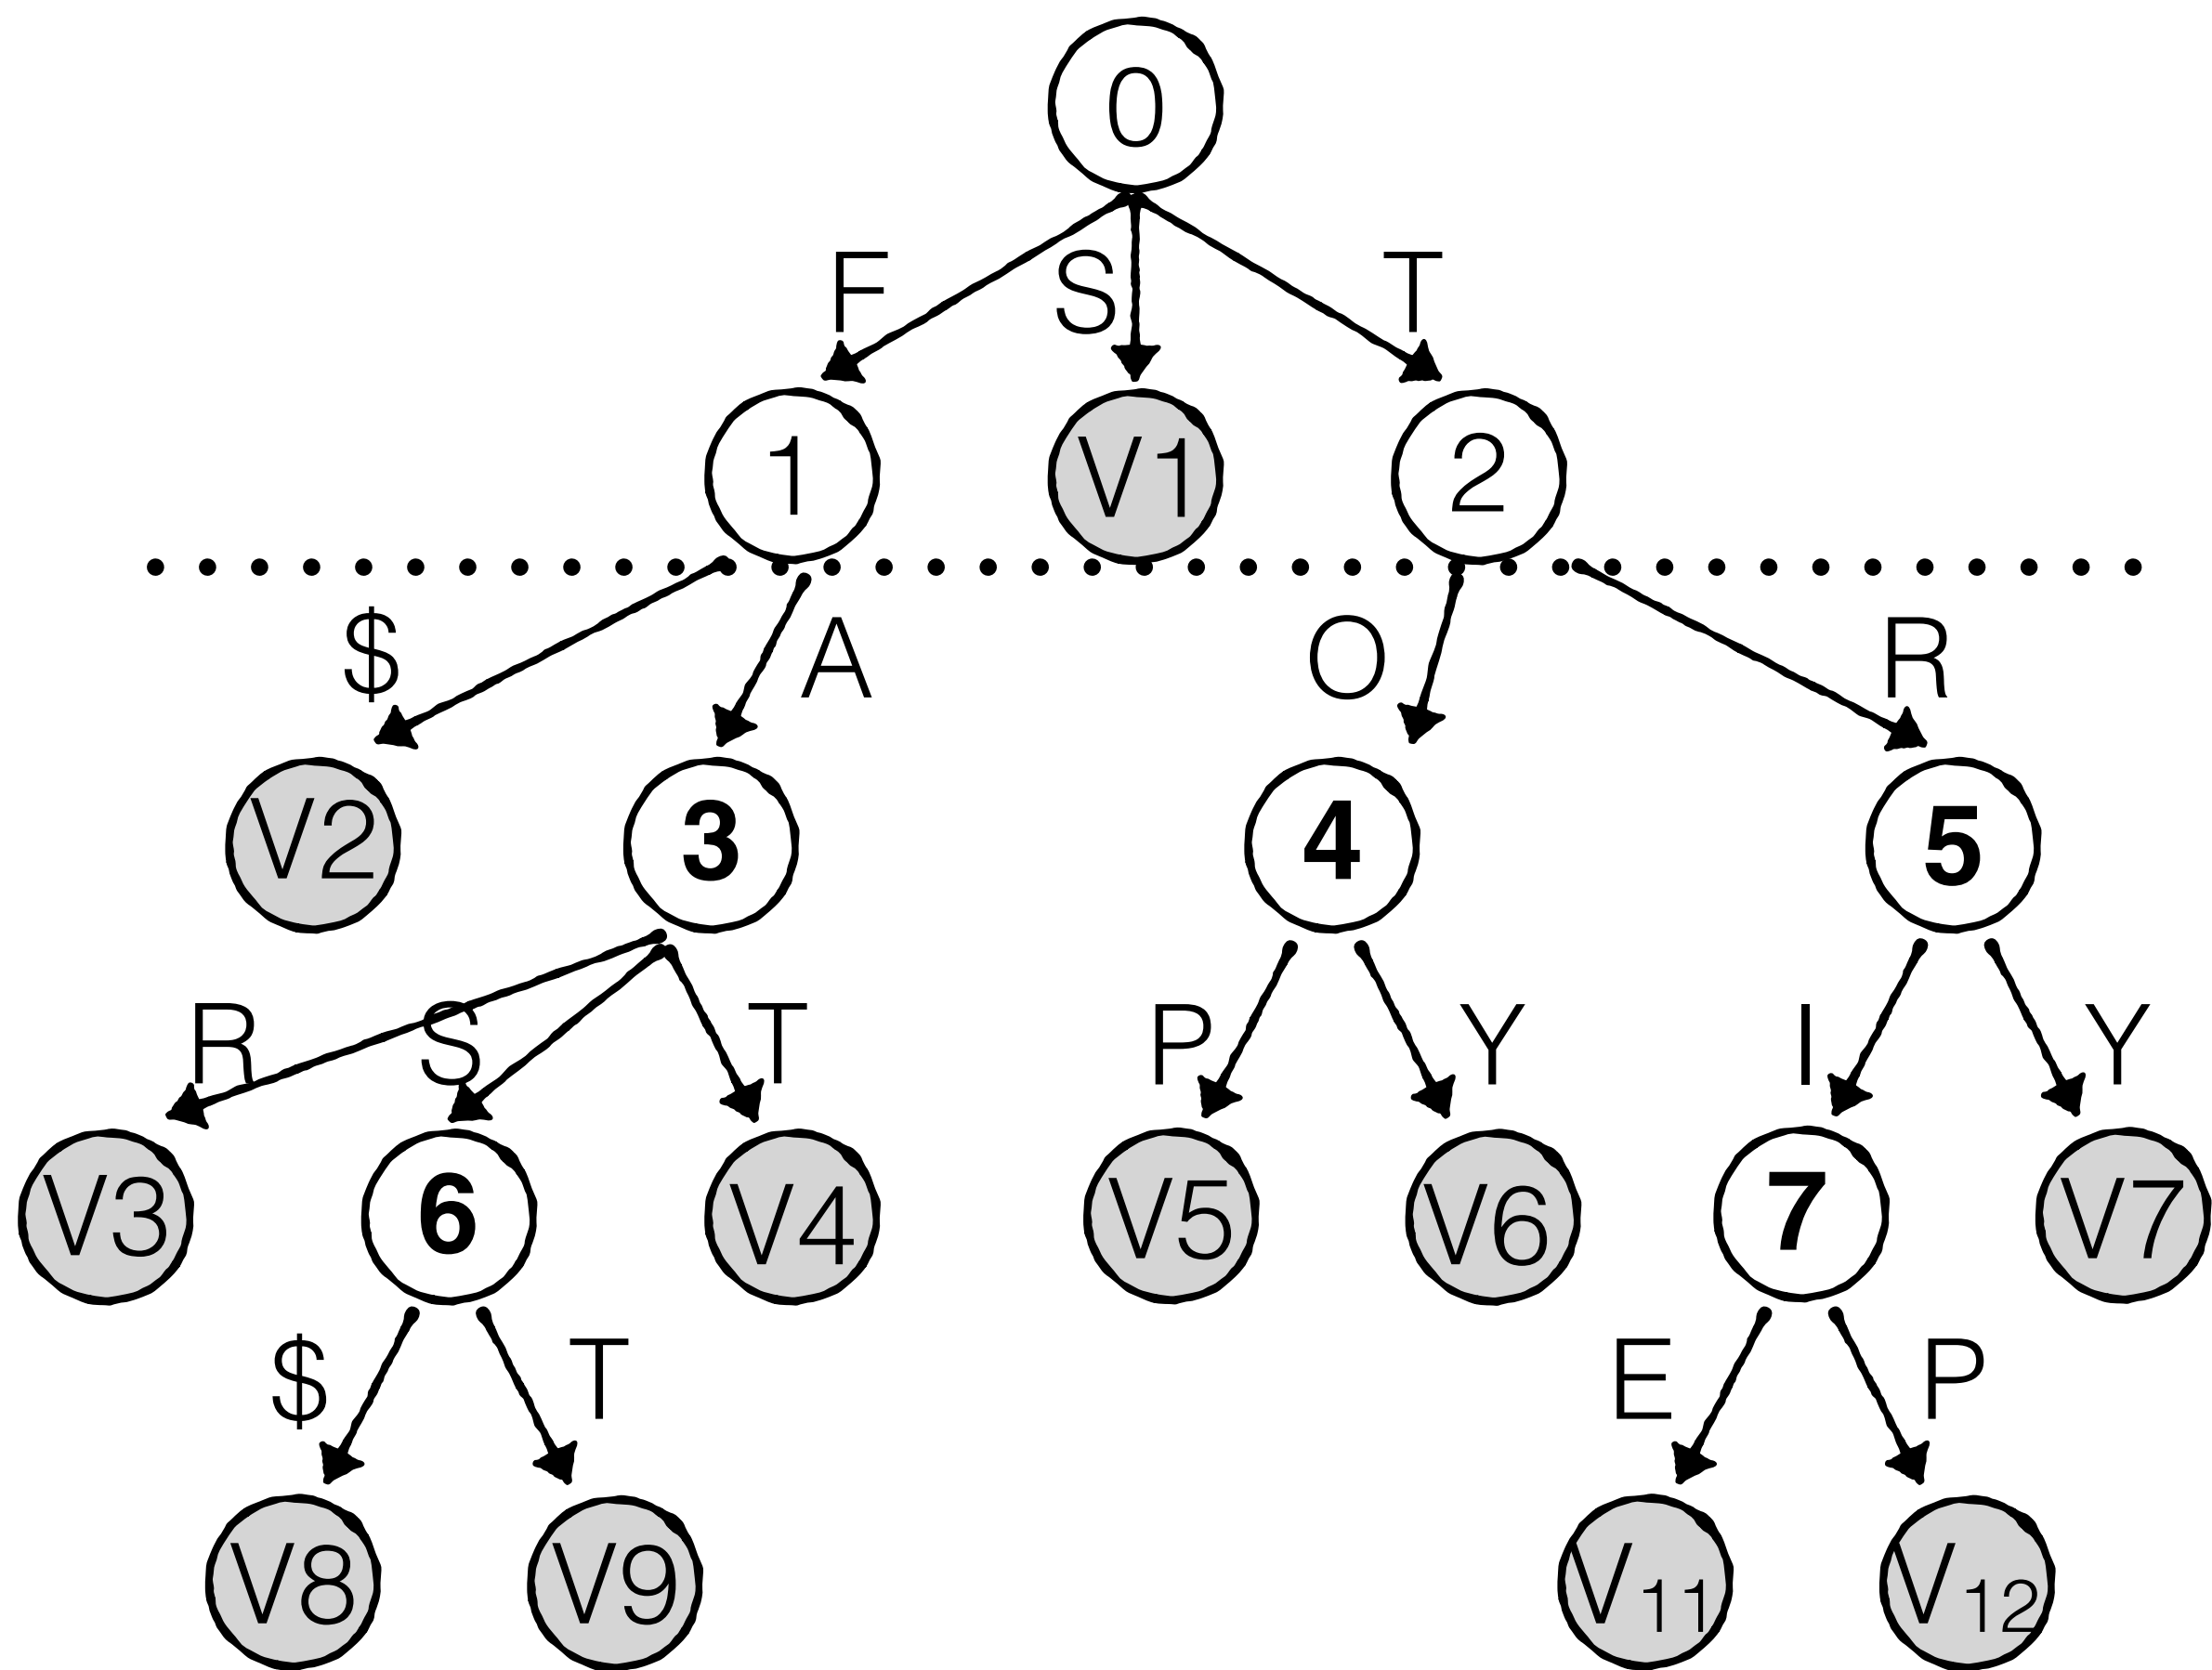
Constant time per each node access :)

**Few bits per entry when base nodes have many edges**




**Internal Nodes in breadth first order**

**3 4 5 6 7**





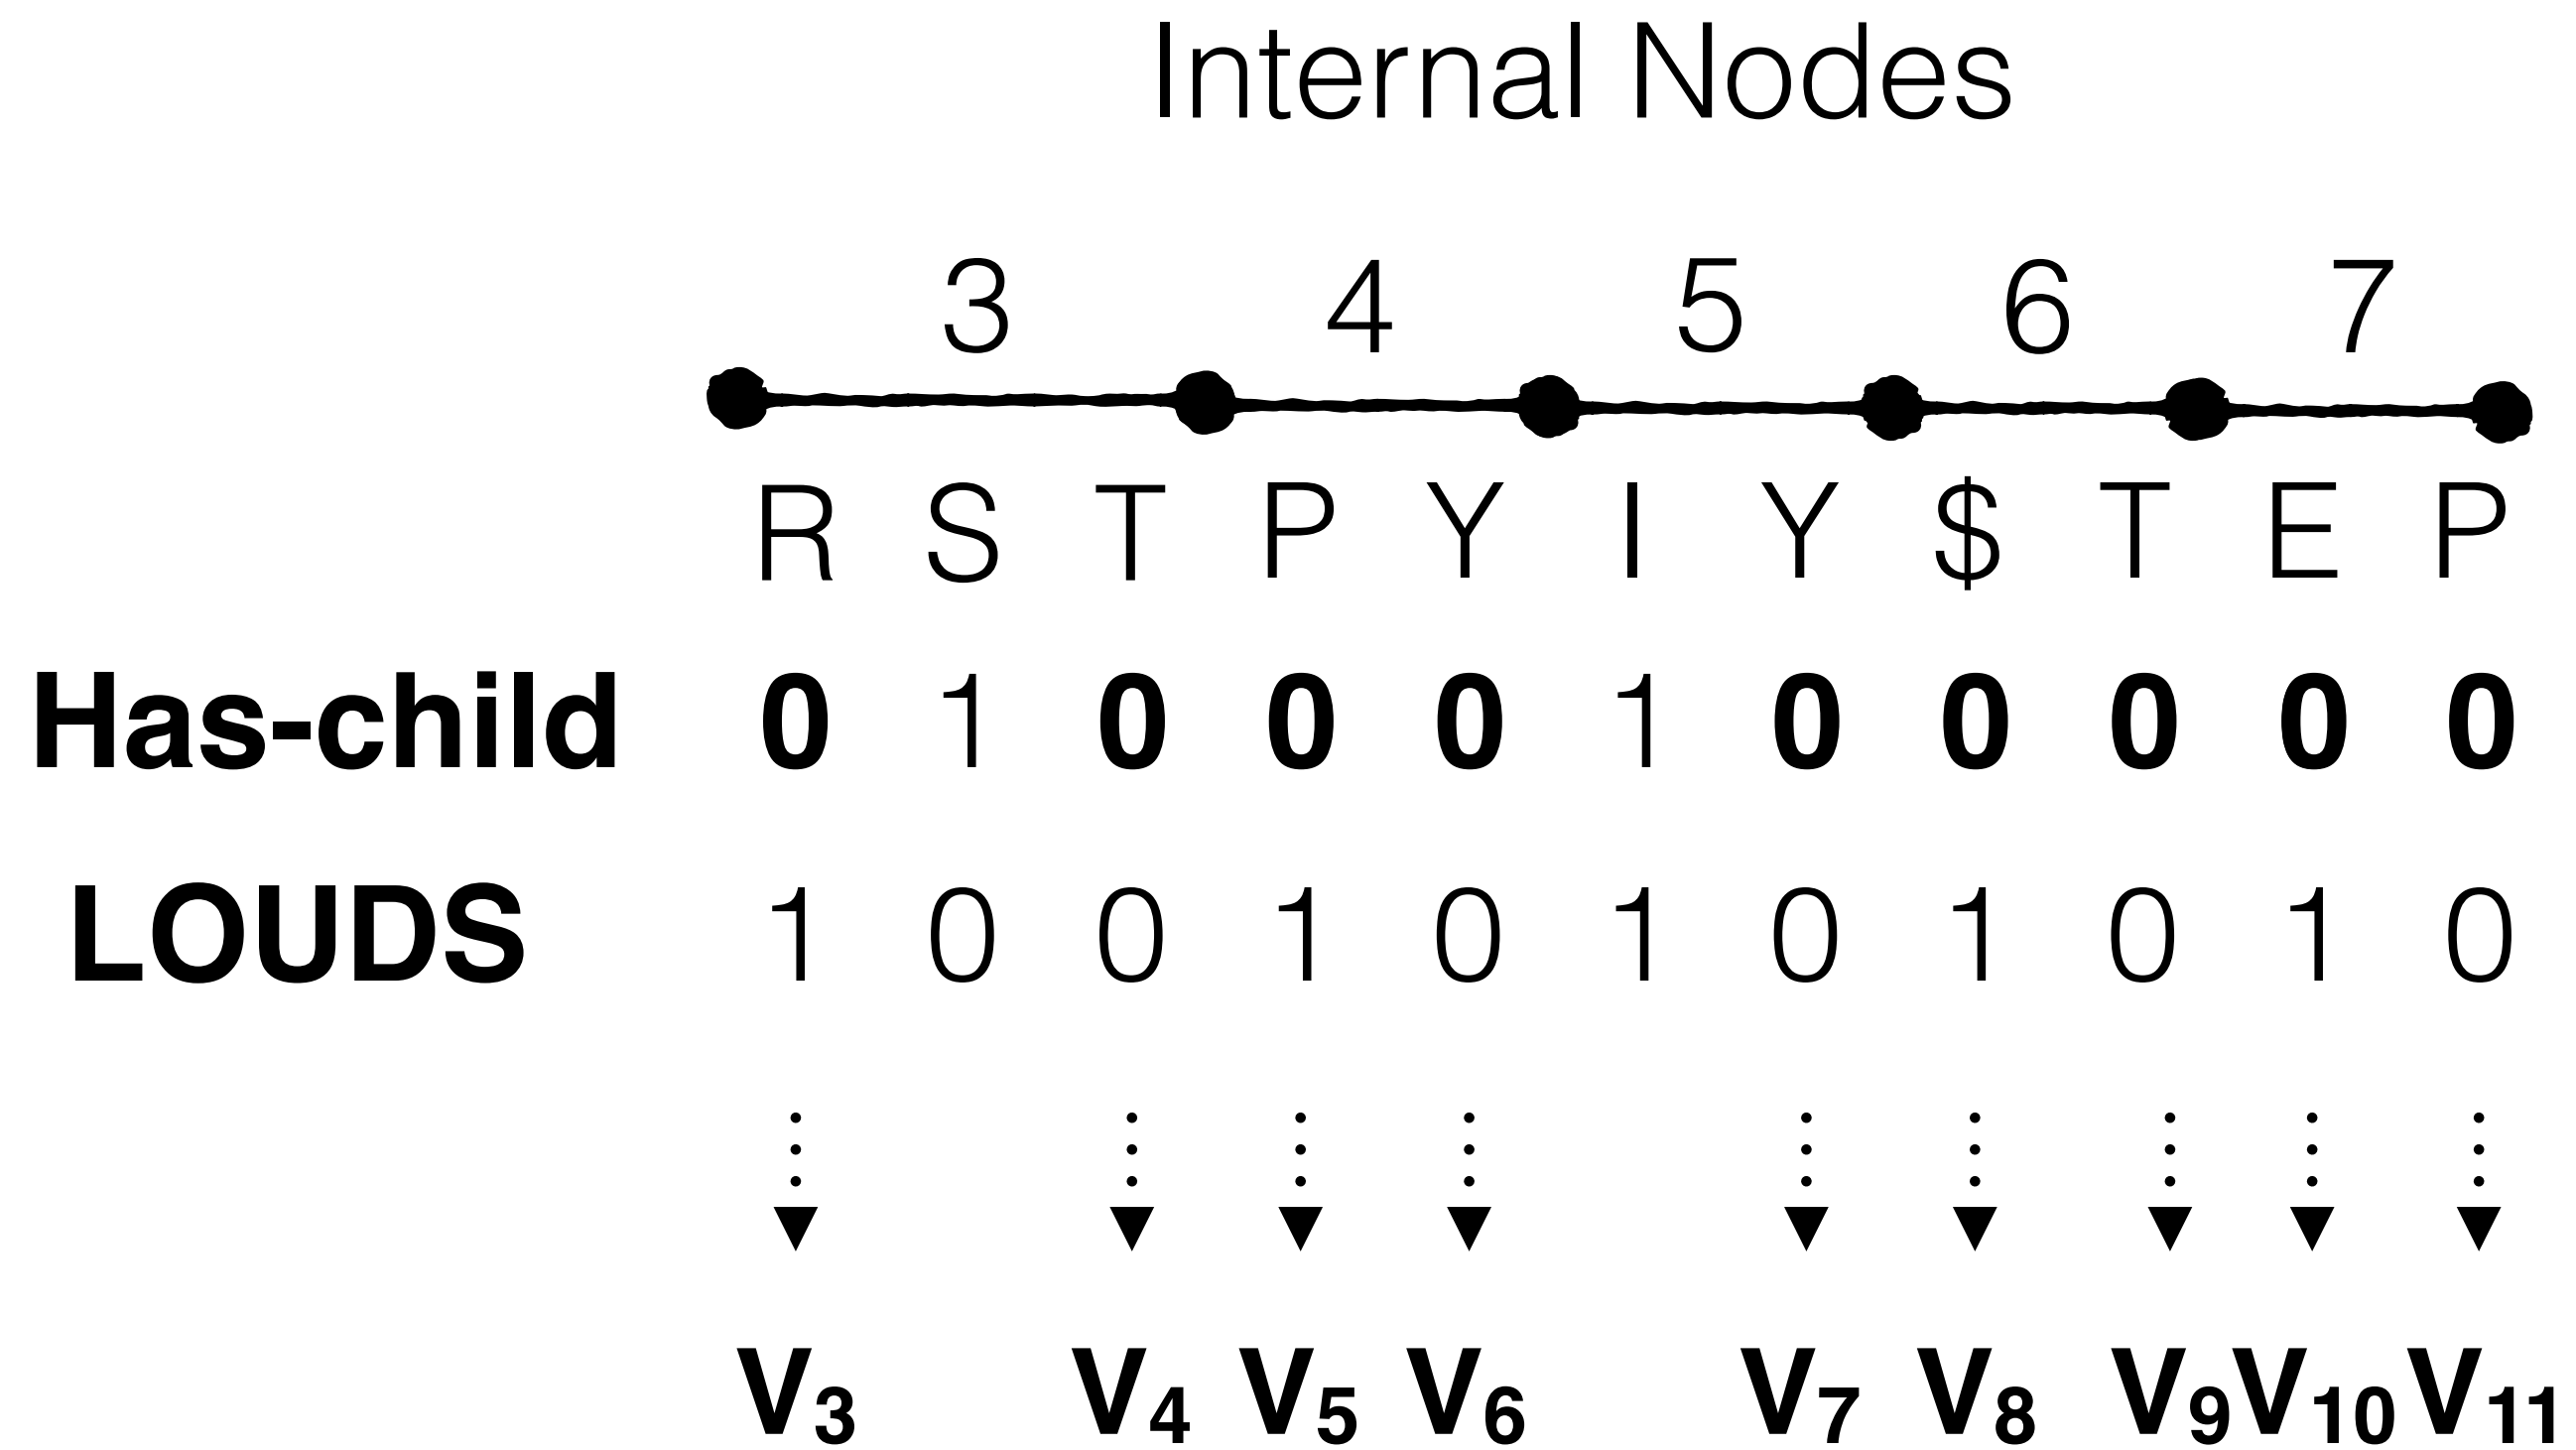

Internal Nodes in breadth first order




**Store children contiguously**











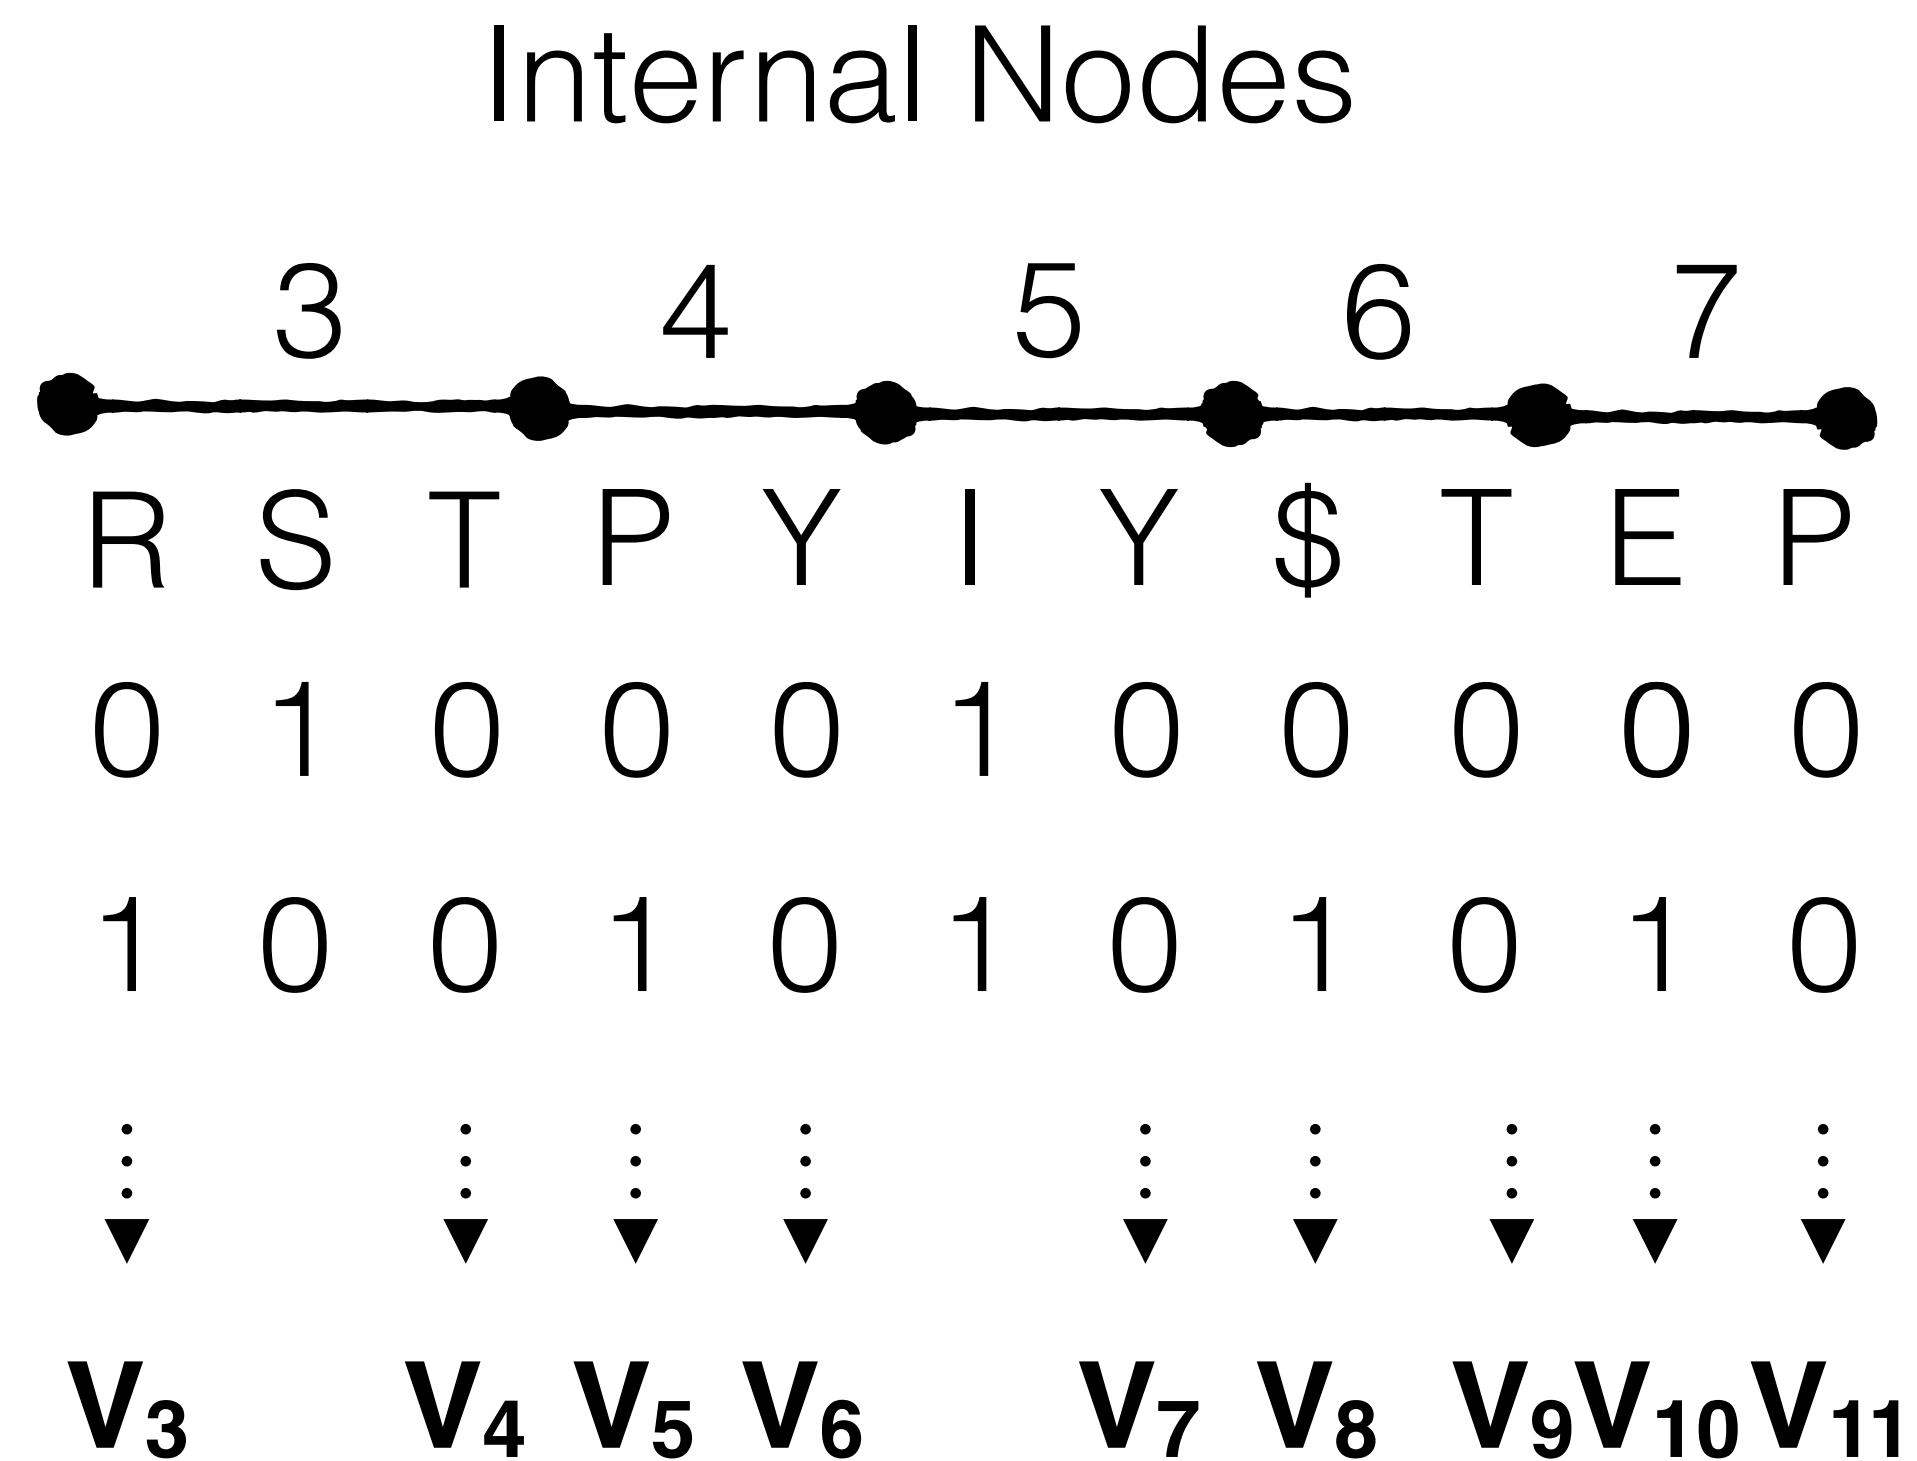
**Has-child  
LOUDS**




**Stored Contiguously**



**Has-child**  
LOUDS




**Each leaf connected to payload**



	F	S	T	A		O	R
Edges							
Has-child							
isKey	0			1		0	

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0

**Whole filter :)**



Edges  
Has-child  
isKey


	F	S	T	A			O	R
Edges								
Has-child								

isKey


	0		1		0
⋮					
↓					
<b>V<sub>1</sub></b>					
↓					
<b>V<sub>2</sub></b>					

Has-child  
LOUDS

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0
⋮											
↓											
<b>V<sub>3</sub></b>		<b>V<sub>4</sub></b>	<b>V<sub>5</sub></b>	<b>V<sub>6</sub></b>		<b>V<sub>7</sub></b>	<b>V<sub>8</sub></b>	<b>V<sub>9</sub></b>	<b>V<sub>10</sub></b>	<b>V<sub>11</sub></b>	




Edges  
Has-child  
isKey



	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0

**V<sub>1</sub> V<sub>2</sub> V<sub>3</sub> V<sub>4</sub> V<sub>5</sub> V<sub>6</sub> V<sub>7</sub> V<sub>8</sub> V<sub>9</sub> V<sub>10</sub> V<sub>11</sub>**



Does "FAS" exist?



F S T A O R


Edges  
Has-child  
isKey


0 1 0

R S T P Y I Y \$ T E P

Has-child 0 1 0 0 0 1 0 0 0 0 0

LOUDS 1 0 0 1 0 1 0 1 0 1 0





**Edges**  
**Has-child**

**FAS**


↓

<b>F</b>	S	T	A			O	R


	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0




	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0



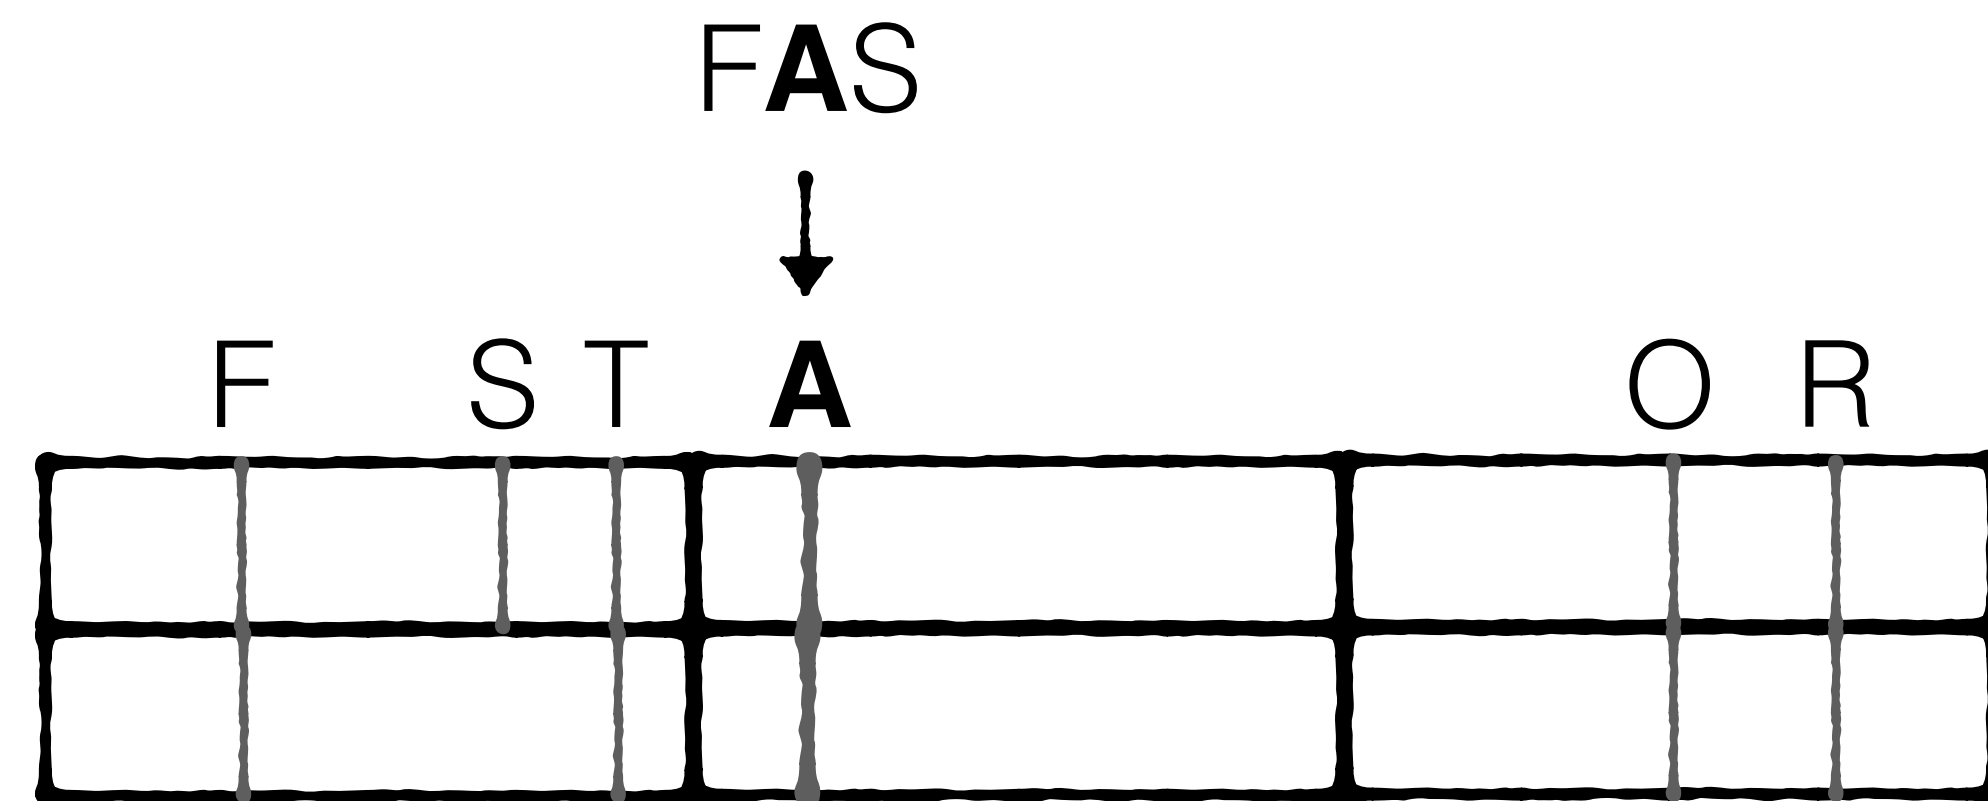
**Edges**  
**Has-child**



	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0




Edges  
Has-child




Rank(A) = 2

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0



**Edges**  
**Has-child**




Rank(A) = 2

**skip**


**skip**

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0





**Edges**  
**Has-child**



Rank(A) = 2      **skip**      **skip**

Has-child


R S T P Y I Y \$ T E P

0 1 0 0 0 1 0 0 0 0 0

**LOUDS**

**1 0 0 1 0 1 0 1 0 1 0**


**Select(0) = 0**



**Edges**  
**Has-child**

	F	S	T	A			O	R

**Node 3**




**R S T P Y I Y \$ T E P**

Has-child 0 1 0 0 0 1 0 0 0 0 0

LOUDS **1** 0 0 1 0 1 0 1 0 1 0

↑  
**FAS**



**Edges**  
**Has-child**

	F	S	T	A		O	R
Edges							
Has-child							

Node 3



R **S** T P Y I Y \$ T E P


Has-child 0 1 0 0 0 1 0 0 0 0 0

LOUDS 1 0 0 1 0 1 0 1 0 1 0



FAS


**Scan**



**Edges**  
**Has-child**

	F	S	T	A		O	R
Edges							
Has-child							



Node 3



	R	<b>S</b>	T	P	Y	I	Y	\$	T	E	P
Has-child	0	<b>1</b>	0	0	0	1	0	0	0	0	0
LOUDS	1	<b>0</b>	0	1	0	1	0	1	0	1	0





**FAS**



Node 3

	R	<b>S</b>	T	P	Y	I	Y	\$	T	E	P
Has-child	0	<b>1</b>	0	0	0	1	0	0	0	0	0
LOUDS	1	<b>0</b>	0	1	0	1	0	1	0	1	0

↑  
FAS




Rank(1) = 0



Has-child	0	1	0	0	0	1	0	0	0	0	0
LOUDS	1	0	0	1	0	1	0	1	0	1	0

↑  
FAS



Edges  
**Has-child**

	F	S	T	A			O	R

Has-child

R **S** T P Y I Y \$ T E P

0 1 0 0 0 1 0 0 0 0 0


LOUDS

1 0 0 1 0 1 0 1 0 1 0



**select(5 - 2) = 7**

↑  
**FAS**




Edges  
**Has-child**

	F	S	T	A		O	R
Edges							
<b>Has-child</b>							

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	<b>0</b>	0	0	0
LOUDS	1	0	0	1	0	1	0	<b>1</b>	0	1	0

↑  
**FAS**




Edges  
**Has-child**

	F	S	T	A	O	R
Edges						
<b>Has-child</b>						

	R	S	T	P	Y	I	Y	\$	T	E	P
Has-child	0	1	0	0	0	1	0	<b>0</b>	0	0	0
LOUDS	1	0	0	1	0	1	0	<b>1</b>	0	1	0

⋮  
**V<sub>8</sub>**

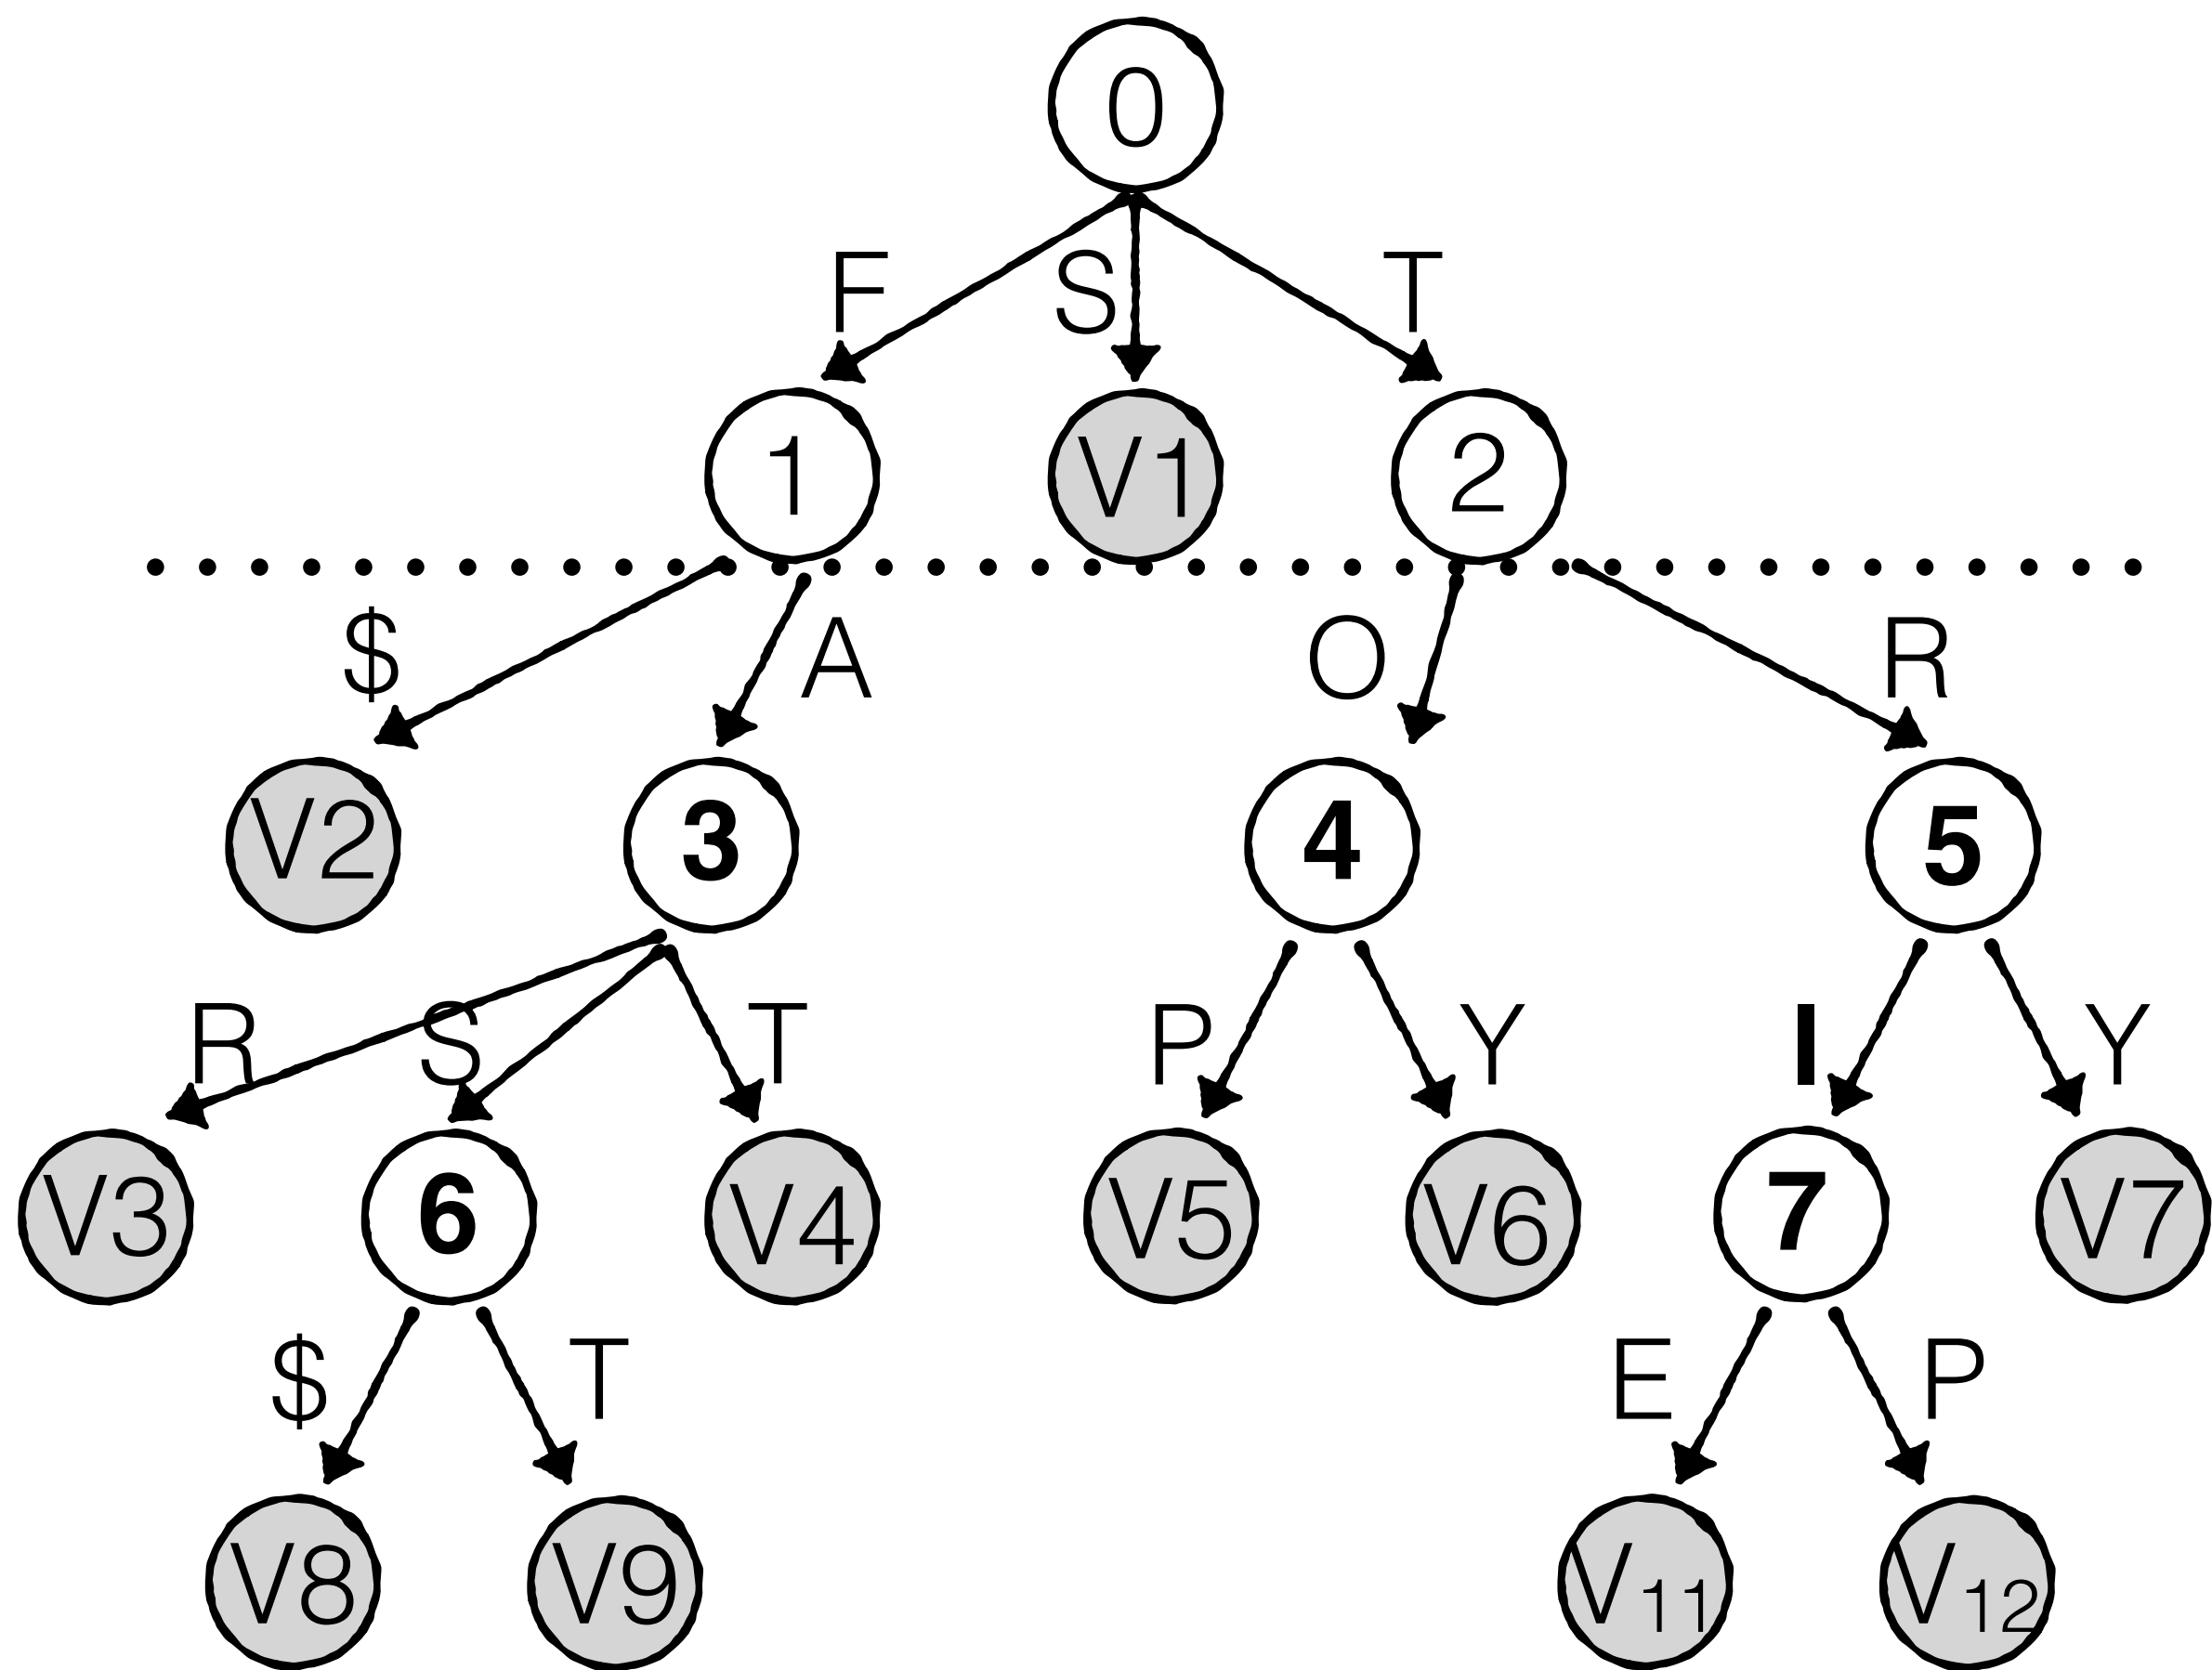


Edges  
Has-child

	F	S	T	A	O	R
Edges						
Has-child						

	R	S	T	P	Y	I	Y	\$	T	E	P
<b>Has-child</b>	0	1	0	0	0	1	0	0	0	0	0
<b>LOUDS</b>	1	0	0	1	0	1	0	1	0	1	0

**Sparse encoding is space efficient for nodes with few edges**



Edges  
Has-child

	F	S	T	A		O	R
Edges							
Has-child							



	<b>R</b>	<b>S</b>	<b>T</b>	P	Y	I	Y	\$	T	E	P
<b>Has-child</b>	0	1	0	0	0	1	0	0	0	0	0
<b>LOUDS</b>	1	0	0	1	0	1	0	1	0	1	0

Sparse encoding is space efficient for nodes with few edges

**But slower as we must scan edges for each node**

# SuRF



**Var-length keys  
& queries**

**Yes**

**FPR guarantee**

**None**

**Query speed**

**$O(L)$**

( $L$  = key length)

**Dynamic**

**No**

SuRF



**Memento**



**Diva**

