

ESE532: System-on-a-Chip Architecture

Day 17: Nov. 1, 2021

LZW, Associative Maps, Hash Tables



Penn ESE532 Fall 2020 -- DeHon

Today

- LZW Compression (Part 1)
- Associative Memory (Part 2)
 - Custom
 - FPGA
- Software Maps
 - Tree (Part 3)
 - Hash Tables (Part 4)
 - (time permit – else next time)

2

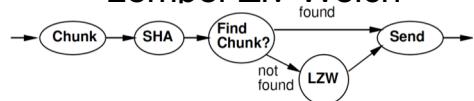
Message

- Rich design space for Maps
- Hash tables are useful tools

Penn ESE532 Fall 2020 -- DeHon

3

Part 1: LZW Compression Lempel-Ziv-Welch



Penn ESE532 Fall 2020 -- DeHon

4

Day 16 Reminder

Idea

- Use data already sent as the dictionary
 - Give short names to things in dictionary
 - Don't need to pre-arrange dictionary
 - Adapt to common phrases/idioms in a particular document

Penn ESE532 Fall 2020 -- DeHon

5

Day 16 Reminder

Encoding

- Greedy simplification
 - Encode by successively selecting the longest match between the head of the remaining string to send and the current window

Penn ESE532 Fall 2020 -- DeHon

6

Algorithm Concept

- While data to send
 - Find largest match in window of data sent
 - If length too small (length=1)
 - Send character
 - Else
 - Send <x,y> = <match-pos,length>
 - Add data encoded into sent window

Day 16: Preclass 4

- How many comparisons per invocation?

```
#define DICT_SIZE 4096
#define LENGTH 256
// clen<=LENGTH
int longest_match(char dict[DICT_SIZE], char candidate[LENGTH], int clen) {
    int best_len=0; best_loc=-1;
    for (int i=0;i<DICT_SIZE-clen;i++) {
        j=0;
        while((candidate[j]==dict[i+j]) & (j<clen)) j++;
        if (j>best_len) {best_len=j; best_loc=i;}
    }
    return((best_loc<>8)|best_len);
}
```

Idea

- Avoid O(Dictionary-size) work
 - Only need to match against positions that start with the character(s) in string to encode
 - Separate dictionary for each?

0	1	2	3	4	5	6	7	8	9	10
I		A	M		S					

Only check position 0 for "starts with I"

Idea

- Avoid O(Dictionary-size) work
 - Only need to match against positions that start with the character(s) in string to encode
 - Separate dictionary for each?
- T-dictionary:
 - Tap, The, Their, Then, There, Tuesday

Idea

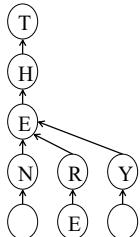
- Avoid O(Dictionary-size) work
 - Only need to match against positions that start with the character(s) in string to encode
 - Separate dictionary for each?
- T-dictionary:
 - Tap, The, Their, Then, There, Tuesday
- If prefix same, why check redundantly?
 - Generalize: Store things with common prefix together
 - Share prefix among substrings
 - Represent all strings as prefix tree
 - Represent all strings as prefix tree

Idea

- Avoid O(Dictionary-size) work
 - Only need to match against positions that start with the character(s) in string to encode
 - Separate dictionary for each?
- If prefix same, why check redundantly?
 - Store things with common prefix together
 - Share prefix among substrings
 - Represent all strings as prefix tree
- Follow prefix trees with **fixed** work per input character

Tree Example

- THEN AND THERE, THEY STOOD...



Penn ESE532 Fall 2020 -- DeHon

13

Tree Algorithm

Tree Root for each character

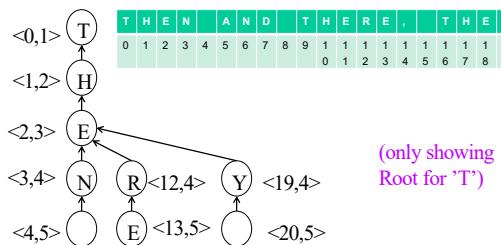
- Follow tree according to input until no more match
- Send <name of last tree node>
 - An <x,y> pair
- Extend tree with new character
- Start over with this character

Penn ESE532 Fall 2020 -- DeHon

14

Tree Example

- Label with <lastpos,len> pair
- THEN AND THERE, THEY STOOD...



Penn ESE532 Fall 2020 -- DeHon

15

Large Memory Implementation

- int encode[SIZE][256];
- Name tree node by position in chunk
 - lastpos
- c is a character
- Encode[lastpos][c] holds the next tree node that extends tree node lastpos by c
 - Or NONE if there is no such tree node

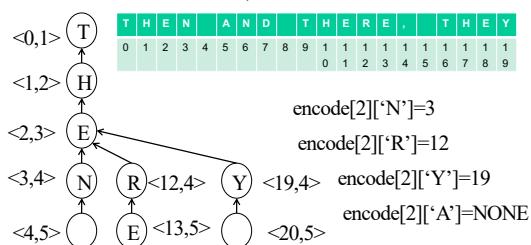
Penn ESE532 Fall 2020 -- DeHon

16

Tree Example

(only showing
Root for 'T')

- Label with <lastpos,len> pair
- THEN AND THERE, THEY STOOD...



Penn ESE532 Fall 2020 -- DeHon

17

Large Table for Tree

Addr	spc	A	...	D	E	...	H	...	N	...	R	S	T	...	Y
-1		5													0
0															1
1															2
2															3
3	4														12
4															19
5															6
6															7
7	8														
8															
9															
10															
12															

Penn E

18

Memory Tree Algorithm

```
curr – pointer into input chunk
// follow tree
y=0; x=-1; // for no match, yet...
while(encode[x][input[curr+y]]!=NONE)
    x=encode[x][input[curr+y]]; y++;
If (y>0)
    send <x,y>
else
    send input[curr+y]
encode[x][input[curr+y]]=curr+y
curr=curr+y+1
```

Penn ESE532 Fall 2020 -- DeHon

19

Memory Tree Algorithm

```
curr – pointer into input chunk
// follow tree
y=0; x=-1; // for no match, yet...
while(encode[x][input[curr+y]]!=NONE)
    x=encode[x][input[curr+y]]; y++;
If (y>0)
    send <x,y>
else
    send input[curr+y]
encode[x][input[curr+y]]=curr+y
curr=curr+y+1
```

Penn ESE532 Fall 2020 -- DeHon

Follow
Tree

20

Memory Tree Algorithm

```
curr – pointer into input chunk
// follow tree
y=0; x=-1; // for no match, yet.
while(encode[x][input[curr+y]]!=NONE)
    x=encode[x][input[curr+y]]; y++;
If (y>0)
    send <x,y>
else
    send input[curr+y]
    encode[x][input[curr+y]]=curr+y
    curr=curr+y+1
```

How much work per character to encode?
Hint:
1) match case
2) not match

Penn ESE532 Fall 2020 -- DeHon

21

Compact Memory

- int encode[SIZE][256];
- How many entries in this table are not NONE?
 - Hint: SIZE is total number of positions.
How many characters process?
How many insertions per character processed?

Penn ESE532 Fall 2020 -- DeHon

22

Compact Memory

- int encode[SIZE][256];
- Table is very sparse
- If store as associative memory
 - At most SIZE entries
- Look at how to implement associative memories next

Penn ESE532 Fall 2020 -- DeHon

23

LZW So Far – 4KB chunks

- Brute Force
 - Needs one byte per byte = 4KB = 1 BRAM
 - DICT_SIZE=4096 comparisons per byte
- Dense memory encode[SIZE][256]
 - Need 2*256B per byte = 512*4KB = 512 BRAMs
 - 1 comparison and lookup per byte

Penn ESE532 Fall 2020 -- DeHon

24

Associative Memories

Part 2

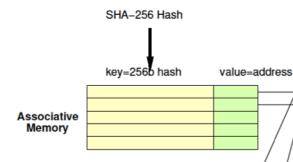
Penn ESE532 Fall 2020 -- DeHon

25

Day 16 Review

Associative Memory

- Maps from a key to a value
- Key not necessarily dense
 - Contrast simple RAM
 - Cannot afford 2^{256} word memory



Penn ESE532 Fall 2020 -- DeHon

Associative Memories

- Use for deduplication
- Also may use in LZW to reduce BRAMs

Penn ESE532 Fall 2020 -- DeHon

27

Day 16 Review

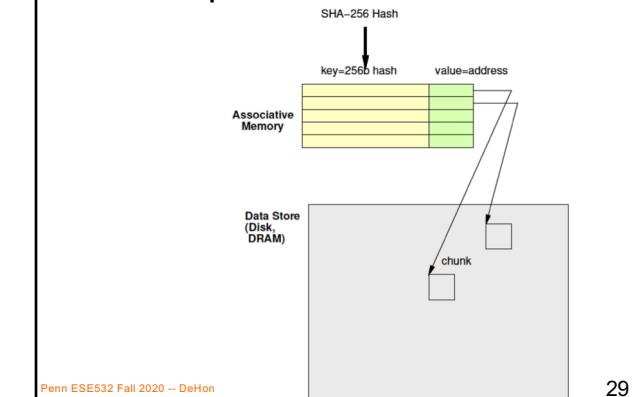
Deduplicate

- Compute chunk hash
- Use chunk hash to lookup known chunks
 - Data already have on disk
 - Data already sent to destination, so destination will know
- If lookup yields a chunk with same hash
 - Check if actually equal (maybe)
- If chunks equal
 - Send (or save) pointer to existing chunk

Penn ESE532 Fall 2020 -- DeHon

28

Deduplication Architecture



Penn ESE532 Fall 2020 -- DeHon

29

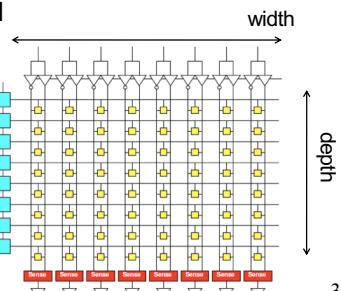
Penn ESE532 Fall 2020 -- DeHon

30

Custom Hardware Associative Memory

Memory Block Review

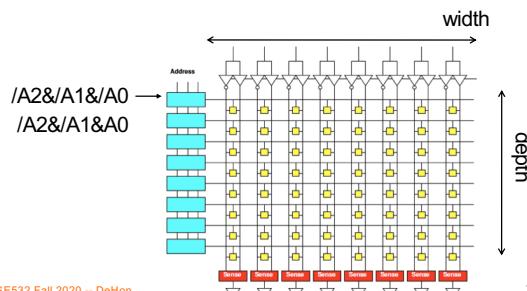
- Match on address
- Select wordline for a row
- Reads out a word
- Address dense and hardwired
- One row for each 2^{Abits} values



31

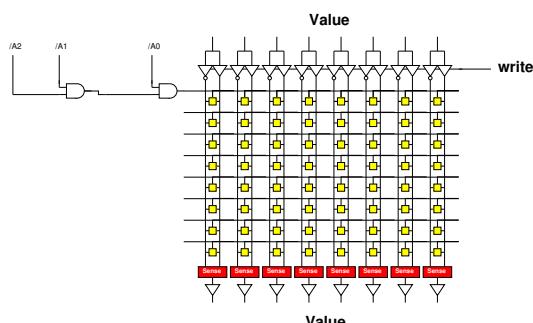
Address Blocks

- Each address match is AND



32

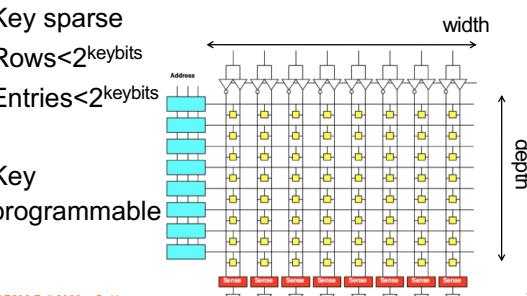
Address Blocks



33

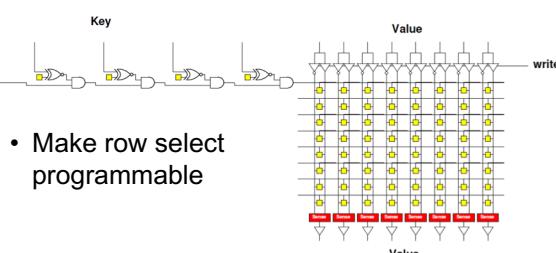
Memory Block Associative

- Want address as key
- Word is value
- Key sparse
- Rows < 2^{keybits}
- Entries < 2^{keybits}
- Key programmable



34

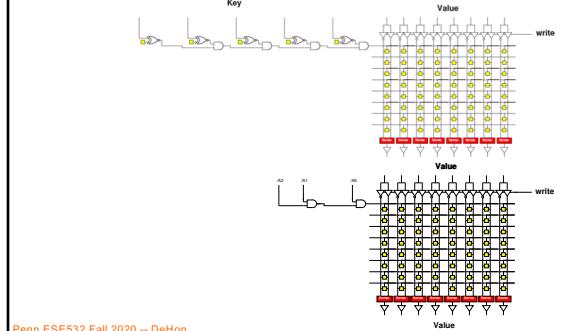
Programmable Key



- Make row select programmable

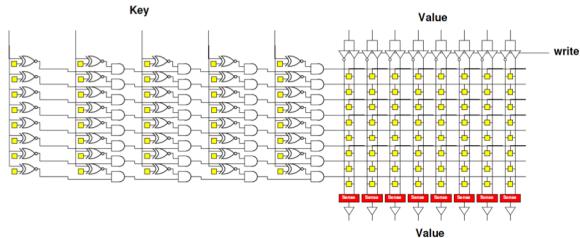
35

Contrast Assoc. and Dense Memory



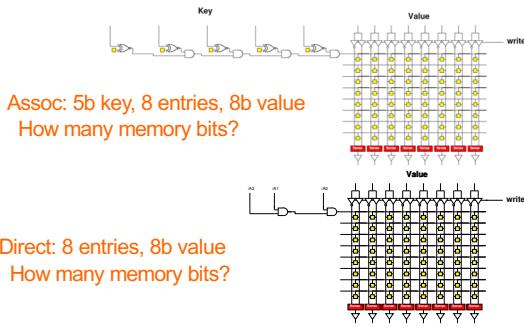
36

Associative Memory Bank



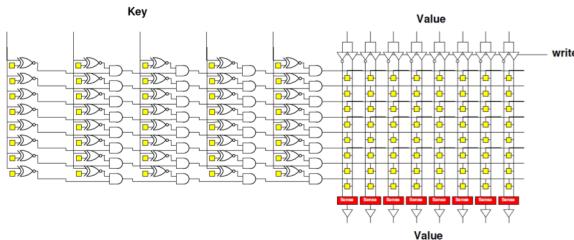
37

Programmable Key



38

Associative Memory Bank

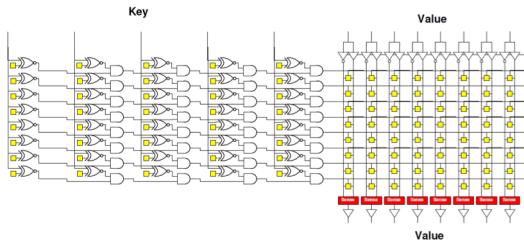


- Memory cells = entries*(keybits+valuebits)

Penn ESE532 Fall 2020 -- DeHon

39

Associative Memory Bank



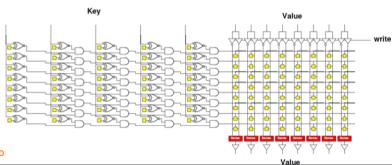
- Will need to be able to write into key
 - Another “fixed” decoder to generate key-word line for programming

Penn ESE532 Fall 2020 -- DeHon

40

Associate Memory Cost

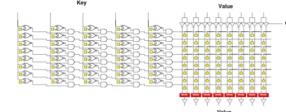
- More expensive than equal capacity SRAM memory bank
 - Memory cells in decoder
 - Need to support write into key



41

Associate Memory Cost

- Physical associative memory for 4KB LZW Chunk tree encode
 - 4K entries
 - 12b output
 - 12b+8b=20b key
- Memory cells assoc.?
- Compare direct 4Kx12 memory (cells)?
 - How much larger is assoc. for same capacity?
- Compare 4096*256 with 12b result for dense LZW case (cells)?
 - How much larger to solve same problem



Penn ESE532 Fall 2020 -- DeHon

42

FPGA

- Has BRAMs – normal memories, not associative
- 36Kb BRAM
– 512x72
- Can be 9b key → 72b value assoc.
– Just using the memory sparsely
- Or interpret as programmable decoder with 72 match lines

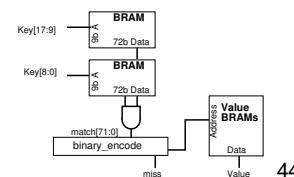
Penn ESE532 Fall 2020 -- DeHon

43

Assoc. Mem from BRAM

For wider match

- Cover 9b of key with each BRAM
- Use 72 output bits to indicate if one of 72 entries match
- AND together corresponding entries
- Get 72 match bits
- Re-encode match bits to lookup value

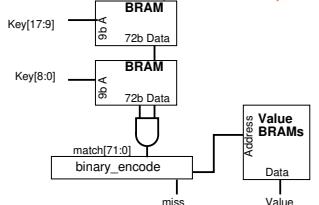


Penn ESE532 Fall 2020 -- DeHon

44

BRAM Associative Memory

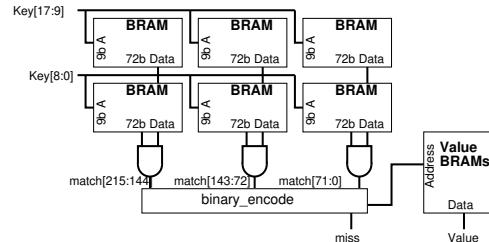
- Previous slide expands match width
- How would we expand capacity?**
– Hint: how get a wider word (144b word)?



Penn ESE532 Fall 2020 -- DeHon

45

BRAM Associative Memory

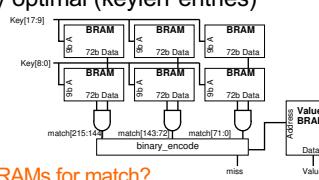


Penn ESE532 Fall 2020 -- DeHon

46

Associative Memory Cost

- Match unit
 - Requires 1 BRAM per 9b of key per 72 entries
 - (keylen/9b) * (entries/72)
 - Asymptotically optimal (keylen*entries)
 - But large constants
- LZW
 - 4K entries
 - 20b key
 - How many BRAMs for match?



Penn ESE532 Fall 2020 -- DeHon

47

4K LZW Chunk Search: Fully associative

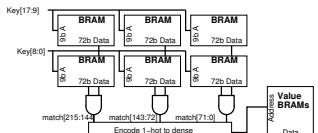
- Match BRAMs:
 - Match key: 20b
 - Entries: 4096
- Value BRAMs:
 - 20b (key) + 12b (state)
 - 32b x 4096 entries

Penn ESE532 Fall 2020 -- DeHon

48

Example Stored Values

Key[17:9]	Key[8:0]	Value
0x001	0x014	0x01
0x001	0x01	0x34
0x0F0	0x014	0xE3
0x0C8	0x113	0xCC



49

Penn ESE532 Fall 2020 -- DeHon

Memory Contents

Key[17:9] Match BRAM

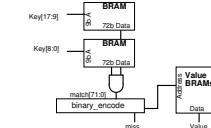
Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

Value BRAM

Addr	Value
0x00	0x01
0x01	0x34
0x02	0xE3
0x03	0xCC
0x04	
0x05	
0x06	

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	0	1	0
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	0	1	0	0



50

Penn ESE532 (only show bottom 8 b; rest 0's)

Code Snippet

```
ap_uint<72> key_low[512];
ap_uint<72> key_high[512];
int value[72];

match_low=key_low[key%512];
match_high=key_high[(key>>9)%512];
match=match_low & match_high;
addr=binary_encode(match);
res=value[addr];
```

51

Penn ESE532 Fall 2020 -- DeHon

Code Snippet

```
ap_uint<72> key_low[512];
ap_uint<72> key_high[512];
int value[72];

match_low=key_low[key%512];
match_high=key_high[(key>>9)%512];
match=match_low & match_high;
addr=binary_encode(match);
res=value[addr];
```

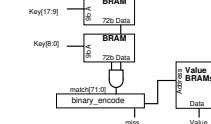
53

Penn ESE532 Fall 2020 -- DeHon

How Lookup Work?

Key[17:9]	Key[8:0]	Value
0x001	0x014	0x01
0x001	0x01	0x34
0x0F0	0x014	0xE3
0x0C8	0x113	0xCC

Lookup 0x214 = 0x001 0x014



52

Memory Contents

Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

match_low=key_low[key%512];
match_high=key_high[(key>>9)%512];

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0

What match_low, match_high?

Penn ESE532 (only show bottom 8 b; rest 0's)

54

Memory Contents

Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

What match?

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0

Penn ESE532 (only show bottom 8 b; rest 0's)

55

What does binary_encode do?

binary_encode(0000000...010000)=0x40

- 10000...0 → 71
- 0000...01 → 0
- 0000...010 → 1
- for (i=0<i<72;i++)
 - If (bit[i]==1) return i
- Return(MISS); // if not find (i.e., all 0's)
- Technicalities – maybe check only one 1

Penn ESE532 Fall 2020 -- DeHon

56

Memory Contents

Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

What addr?

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0

Penn ESE532 (only show bottom 8 b; rest 0's)

57

Memory Contents

Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

Value BRAM

Addr	Value
0x00	0x01
0x01	0x34
0x02	0xE3
0x03	0xCC
0x04	
0x05	
0x06	

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0

res=value[addr];

What res?

Penn ESE532 (only show bottom 8 b; rest 0's)

58

Memory Contents

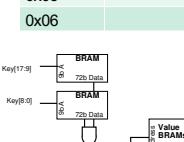
Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	0	1	0	0

Addr	Value
0x00	0x01
0x01	0x34
0x02	0xE3
0x03	0xCC
0x04	
0x05	
0x06	

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0



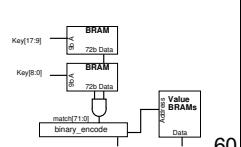
59

Add another entry

match	Key[17:9]	Key[8:0]	Value
0	0x001	0x014	0x01
1	0x001	0x01	0x34
2	0x0F0	0x014	0xE3
3	0x0C8	0x113	0xCC
4	0x0C8	0x01	0x2B

How BRAM contents change to add this entry for 0x19001

Penn ESE532 Fall 2020 -- DeHon



60

Memory Contents

Key[17:9] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	1	1	
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	1	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

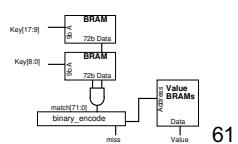
Value BRAM

Addr	Value
0x00	0x01
0x01	0x34
0x02	0xE3
0x03	0xCC
0x04	0x2B
0x05	
0x06	

Key[8:0] Match BRAM

Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	1	0	0	1	0
0x014	0	0	0	0	0	1	0	1
0x0C8	0	0	0	0	0	0	0	0
0x0F0	0	0	0	0	0	0	0	0
0x113	0	0	0	0	1	0	0	0

Penn ESE532 (only show bottom 8 b; rest 0's)



61

Software Map

Part 3

62

4K Chunk LZW Search

	BRAMs	Operations
Brute Search	1	4K
Tree with Dense RAM	512	1
Tree with Full Assoc	175	1

36Kb BRAMs on ZU3EG = 216

Penn ESE532 Fall 2020 -- DeHon

63

Software Map

- Map abstraction
 - void insert(key,value);
 - value lookup(key);
- Will typically have many different implementations

64

Preclass 1

- For a capacity of 4096
- How many memory accesses needed
 - When lookup fail?
 - When lookup succeed (on average)?

Penn ESE532 Fall 2020 -- DeHon

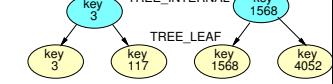
65

Tree Map (Preclass 2)

- Build search tree
- Walk down tree
- For a capacity of 4096, assume balanced...
- How many tree nodes visited
 - When lookup fail?
 - When lookup succeed (on average)?

Penn ESE532 Fall 2020 -- DeHon

66



Tree Map LZW

- Each character requires $\log_2(\text{dict})$ lookups
 - 12 for 4096
- Each internal tree node hold
 - Key (20b for LZW), value (12b), and 2 pointers (12b)
 - 7B
- Total nodes $4K^2$
- Need 14 BRAMs for 4K chunk

Penn ESE532 Fall 2020 -- DeHon

67

Tree Insert

- Need to maintain balance
- Doable with $O(\log(N))$ insert
 - Tricky
 - See Red-Black Tree
 - https://en.wikipedia.org/wiki/Red–black_tree
 - <https://www.geeksforgeeks.org/red-black-tree-set-1-introduction-2/>

Penn ESE532 Fall 2020 -- DeHon

68

4K Chunk LZW Search

	BRAMs	Operations
Brute Search	1	4K
Tree with Dense RAM	512	1
Tree with Full Assoc	175	1
Tree with Tree	14	12

36Kb BRAMs on ZU3EG = 216

Penn ESE532 Fall 2020 -- DeHon

69

Hash Tables

Part 4

Penn ESE532 Fall 2020 -- DeHon

70

High Performance Map

- Would prefer not to search
- Want to do better than $\log_2(N)$ time
- Direct lookup in arrays (memory) is good...

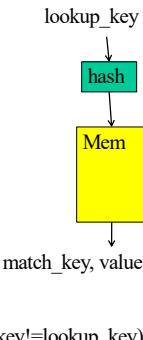
Penn ESE532 Fall 2020 -- DeHon

71

Hash Table

- Attempt to turn into direct lookup
- Compute some function of key
 - A hash
- Perform lookup at that point
- If hash maps a single entry (or no entry)
 - Great, got direct lookup
 - Like sparse table case

Penn ESE532 Fall 2020 -- DeHon



72

Preclass 3a

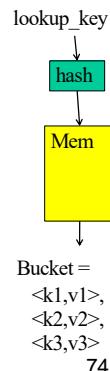
- Average number of entries per hash when $N > \text{HASH_CAPACITY}$?
 - Concrete example
 - $N=4096$
 - $\text{HASH_CAPACITY}=256$

Penn ESE532 Fall 2020 -- DeHon

73

Hash Table

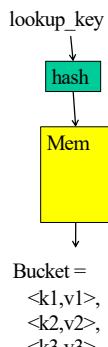
- Attempt to turn into direct lookup
- Compute some function of key
 - A hash
- Perform lookup at that point
- Typically, prepared for several keys to map to same hash → call it a bucket
 - Keep list or tree of things in each bucket



74

Hash Table

- Compute some function of key
 - A hash
- Perform lookup at that point
- Find bucket with small number of entries
 - Searching that bucket easier
 - ...but no absolute guarantee on maximum bucket size



75

Preclass 3b

- Probability of conflict if $N < \text{HASH_CAPACITY}$?
 - Concrete example
 - $N=4096$
 - $\text{HASH_CAPACITY}=409600$
- Impact of HASH_CAPACITY on average bucket size?

Penn ESE532 Fall 2020 -- DeHon

76

Big Ideas

- Rich design space for engineering associative map solutions
- Sparse, near $O(1)$ Map access → Hash Table

Penn ESE532 Fall 2020 -- DeHon

77

Admin

- Feedback (including HW7)
- Reading for Wednesday on Web
- First project milestone due Friday
 - Including teaming

Penn ESE532 Fall 2020 -- DeHon

78