

ESE5320: System-on-a-Chip Architecture

Day 17: Oct. 28, 2024
LZW, Associative Maps



Penn ESE5320 Fall 2024 -- DeHon

1

Today

- LZW Compression Algorithm (Part 1)
- Associative Memory
 - Custom (part 2)
 - FPGA (Part 3, time permitting)

Penn ESE5320 Fall 2024 -- DeHon

2

2

Message

- Can adaptively compress data using recurring substrings
- Rich design space for Maps

Penn ESE5320 Fall 2024 -- DeHon

3

3

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 ESE5320 Fall 2024 -- DeHon

4

4

Encoding

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

Penn ESE5320 Fall 2024 -- DeHon

5

5

Algorithm Concept

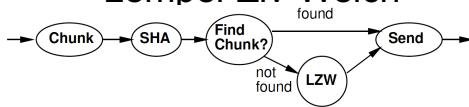
- While data to send
 - Find largest match in window of data sent
 - If length too small (length=1)
 - Send character
 - Else
 - Send $\langle x, y \rangle = \langle \text{match-pos}, \text{length} \rangle$
 - Add data encoded into sent window

Penn ESE5320 Fall 2024 -- DeHon

6

6

Part 1: LZW Compression Lempel-Ziv-Welch



Penn ESE5320 Fall 2024 -- DeHon

7

7

Idea

- Avoid $O(\text{Dictionary-size})$ 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"

Penn ESE5320 Fall 2024 -- DeHon

8

8

Idea

- Avoid $O(\text{Dictionary-size})$ 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

Penn ESE5320 Fall 2024 -- DeHon

9

9

Idea

- Avoid $O(\text{Dictionary-size})$ 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

Penn ESE5320 Fall 2024 -- DeHon

10

10

Idea

- Avoid $O(\text{Dictionary-size})$ 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

Penn ESE5320 Fall 2024 -- DeHon

11

11

Tree Algorithm

Tree Root for each character

- Follow tree according to input until no more match
- Send <name of last tree node>
 - Dictionary entry
 - Named sequentially by insertion
- Extend tree (dictionary) with new character
- Start over with this character

Penn ESE5320 Fall 2024 -- DeHon

12

12

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	???	S	???	???

Penn ESE5320 Fall 2024 -- DeHon 25

25

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	83	S	261	S<-A

Penn ESE5320 Fall 2024 -- DeHon 26

26

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	83	S	261	S<-A
M	?? (different)	A	???	???

Penn ESE5320 Fall 2024 -- DeHon 27

27

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	83	S	261	S<-A
M	<nothing>	A		

Penn ESE5320 Fall 2024 -- DeHon 28

28

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	83	S	261	S<-A
M	<nothing>	A		
M	???	AM (258)	???	???

Penn ESE5320 Fall 2024 -- DeHon 29

29

I AM SAM SAM I AM

Input	Send	Dict	Add Entry	Add What
I	<nothing>			
I	73	I	256	I<-spc
A	32	spc	257	spc<-A
M	65	A	258	A<-M
M	77	M	259	M<-spc
S	32	spc	260	spc<-S
A	83	S	261	S<-A
M	<nothing>	A		
M	258	AM (258)	262	258<-spc

Penn ESE5320 Fall 2024 -- DeHon 30

30

Map

- Later today or Wednesday will talk about higher-level maps
- Version of encode/decode in
 - Geeks-for-Geeks description
 - Decoder we provide
- Based on Maps of strings
 - Simpler to state
 - Likely not $O(1)$ per symbol
 - More expensive implementation in hardware
 - Version here better for hardware (keeps simple)

Penn ESE5320 Fall 2024 -- DeHon

49

49

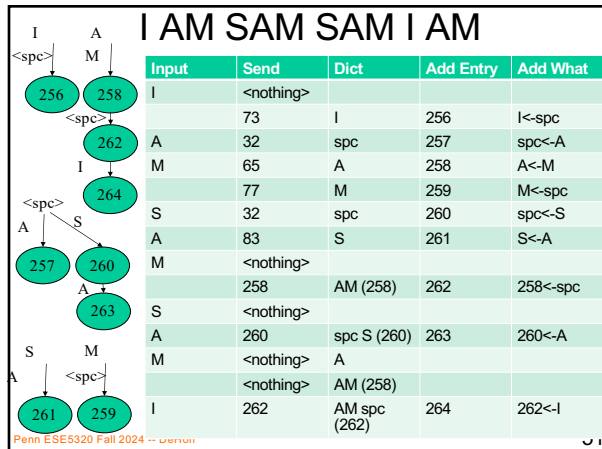
Algorithm with Map

```
curr=0; // pointer into input chunk
nextdict=NUM_SYMBOLS;
// dict initialized symbol i for i<NUM_SYMBOLS
while (curr<chunk_size)
    while(dict.lookup(x+input[curr])!=NONE)
        x=x+input[curr]; curr++;
    send dic.lookup(x);
    dict.insert(x+input[curr],nextdict++);
    x=String(input[curr]); curr++;
```

Penn ESE5320 Fall 2024 -- DeHon

50

50



Penn ESE5320 Fall 2024 -- DeHon

51

51

Compact Memory

- `int encode[SIZE][256];`
- How many entries in this table are not NONE?
 - Maximum number of times execute: `encode[x][input[curr]]=nextdict++;`
 - Maximum number of entries added to encode?
 - Maximum number of NONE entries in encode?

Penn ESE5320 Fall 2024 -- DeHon

52

52

Compact Memory

- `int encode[SIZE][256];`
- Table is very sparse
- If store as associative memory
 - At most SIZE entries
 - (SIZE+256 with first 256 entries for each byte value)
- Look at how to implement associative memories next

Penn ESE5320 Fall 2024 -- DeHon

53

53

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]`
 - 12b node id (1.5B)
 - Need $1.5 \cdot 4096 \cdot 256 = 384 \cdot 4\text{KB}$
 - = 384 BRAMs
 - 1 comparison and lookup per byte
 - (maybe should be SIZE+256)

Penn ESE5320 Fall 2024 -- DeHon

54

54

4K Chunk LZW Search

	BRAMs	Operations/byte
Brute Search	1	4K
Tree with Dense RAM	384	1

36Kb BRAMs on ZU3EG = 216

Complexity

- Reminder from Day 15
 - Optimized implementations tend to be more complex
 - Large problems may force more complexity
 - (e.g. Large window, deal with limited BRAMs on chip)
 - Seeing examples of that today...

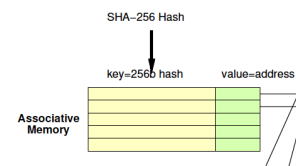
Associative Memories

Part 3

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



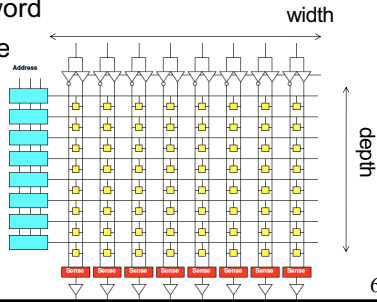
Associative Memories

- Use for deduplication
- Also may use in LZW to reduce BRAMs
 - Just saw
 - **Problem:** Simple 2D tree table requires too many BRAMs
 - **Opportunity:** Tree table sparse

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



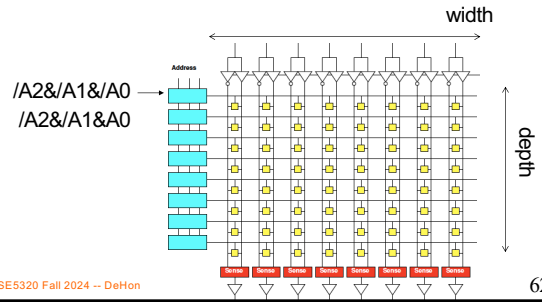
Penn ESE5320 Fall 2024 -- DeHon

61

61

Address Blocks

- Each address match is AND

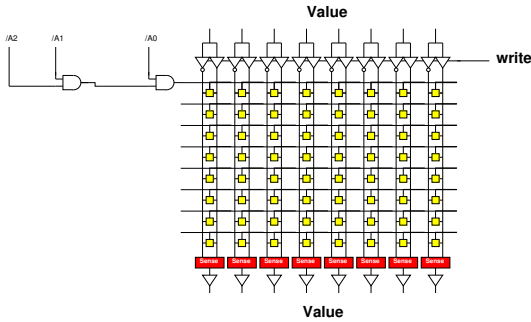


Penn ESE5320 Fall 2024 -- DeHon

62

62

Address Blocks



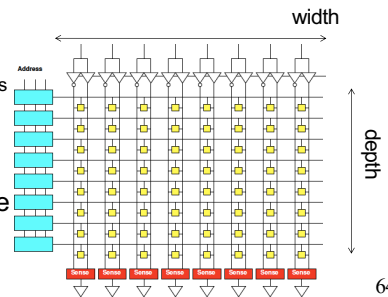
Penn ESE5320 Fall 2024 -- DeHon

63

63

Memory Block Associative

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



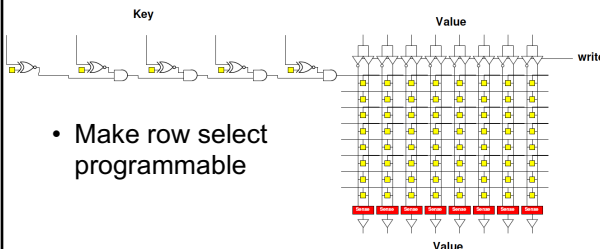
Penn ESE5320 Fall 2024 -- DeHon

64

64

Programmable Key

- Make row select programmable

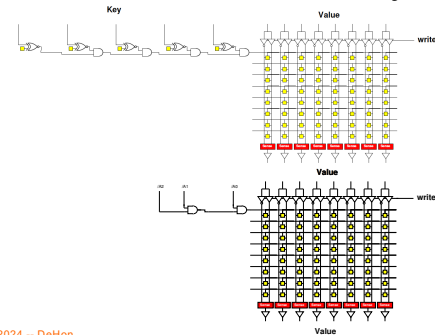


Penn ESE5320 Fall 2024 -- DeHon

65

65

Contrast Assoc. and Dense Memory

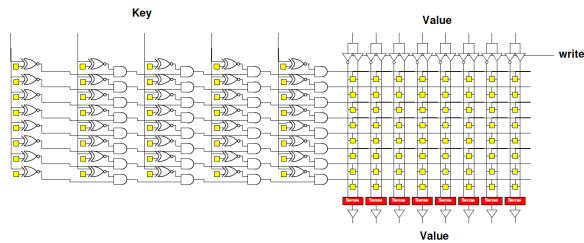


Penn ESE5320 Fall 2024 -- DeHon

66

66

Associative Memory Bank

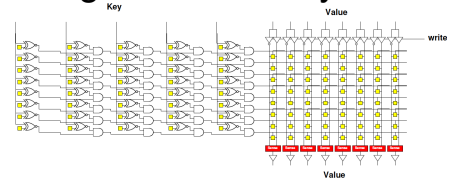


Penn ESE5320 Fall 2024 -- DeHon

67

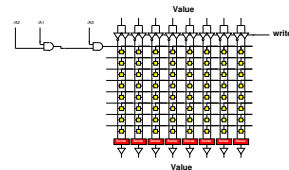
67

Programmable Key



Assoc: 5b key, 8 entries, 8b value
How many memory bits?

Direct: 8 entries, 8b value
How many memory bits?

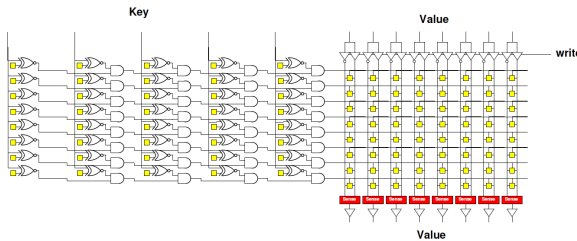


Penn ESE5320 Fall 2024 -- DeHon

68

68

Associative Memory Bank



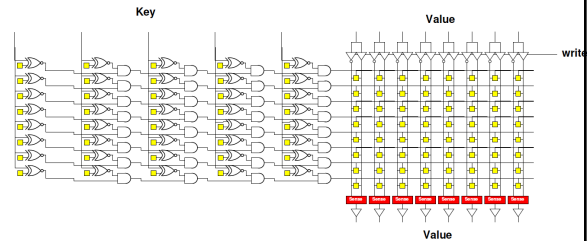
- Memory cells = entries*(keybits+valuebits)

Penn ESE5320 Fall 2024 -- DeHon

69

69

Associative Memory Bank



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

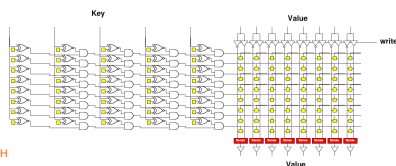
Penn ESE5320 Fall 2024 -- DeHon

70

70

Associate Memory Cost

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



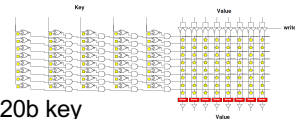
Penn ESE5320 Fall 2024 -- DeH

71

71

Associate Memory Cost

- Physical associative memory for 4KB LZW Chunk tree encode
 - 4K entries
 - 12b (pos) output
 - 12b (pos)+8b (char)=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 ESE5320 Fall 2024 -- DeHon

72

72

Associative Memories FPGA

Part 4
(probably for Day 18)

[Skip wrapup](#)

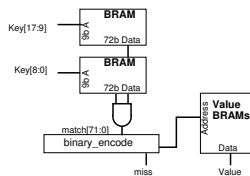
FPGA

- 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

Assoc. Mem from BRAM

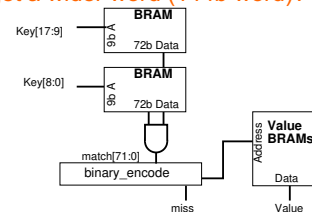
For wider match

- 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

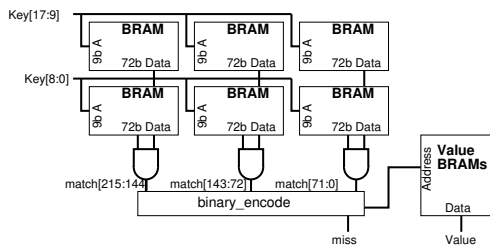


BRAM Associative Memory

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

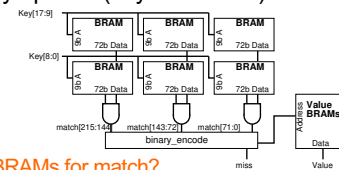


BRAM Associative Memory



Associative Memory Cost

- Match unit
 - Requires 1 BRAM per 9b of key per 72 entries
 - $[\text{keylen}/9\text{b}] \times [\text{entries}/72]$
 - Asymptotically optimal ($\text{keylen} \times \text{entries}$)
 - But large constants
- LZW
 - 4K entries
 - 20b key
 - How many BRAMs for match?



4K LZW Chunk Search: Fully associative

- Match BRAMs:
 - Match key: 20b
 - Entries: 4096
- Value BRAMs:
 - 12b (state [position])
 - 12b x 4096 entries
 - Takes 2 BRAMs

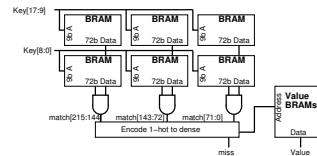
Penn ESE5320 Fall 2024 -- DeHon

79

79

Example Stored Values

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



Penn ESE5320 Fall 2024 -- DeHon

80

80

Memory Contents

Key[17:9] Match BRAM

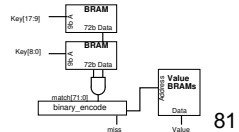
Addr	7	6	5	4	3	2	1	0
0x001	0	0	0	0	0	0	1	1
0x014	0	0	0	0	0	0	0	0
0x0C8	0	0	0	0	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

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



Penn ESE5320 Fall 2024 -- DeHon (only show bottom 8 b; rest 0's)

81

81

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];
    
```

Penn ESE5320 Fall 2024 -- DeHon

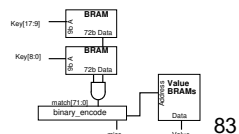
82

82

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



Penn ESE5320 Fall 2024 -- DeHon

83

83

Code Snippet

```

ap_uint<72> key_low[512];
ap_uint<712> 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];
    
```

Penn ESE5320 Fall 2024 -- DeHon

84

84

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 ESE5320 (only show bottom 8 b; rest 0's) 85

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]

match=match_low & match_high;

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?

Penn ESE5320 (only show bottom 8 b; rest 0's) 86

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 ESE5320 Fall 2024 -- DeHon 87

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]

match=match_low & match_high;

addr=binary_encode(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

What addr?

Penn ESE5320 (only show bottom 8 b; rest 0's) 88

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

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

Value BRAM

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

res=value[addr];

What res?

Penn ESE5320 (only show bottom 8 b; rest 0's) 89

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

Penn ESE5320 Fall 2024 -- DeHon 90

Memory Contents

Key[17:9] Match BRAM Value 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

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

Penn ESE5320 Fall 2024 -- DeHon

91

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 ESE5320 Fall 2024 -- DeHon

92

Memory Contents

Key[17:9] Match BRAM Value 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

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

Penn ESE5320 Fall 2024 -- DeHon

93

Memory Contents

Key[17:9] Match BRAM Value 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	1	1	0	0	0
0x0F0	0	0	0	0	0	1	0	0
0x113	0	0	0	0	0	0	0	0

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 ESE5320 Fall 2024 -- DeHon

94

4K Chunk LZW Search

	BRAMs	Operations/byte
Brute Search		1 4K
Tree with Dense RAM	384	1
Tree with Full Assoc	173	1

36Kb BRAMs on ZU3EG = 216

Penn ESE5320 Fall 2024 -- DeHon

95

Checkpoint

- Notice levels of mapping:
 - Prefix Tree algorithm
 - Formulated on a 2D memory
 - Then implemented in assoc. memory
 - (later with Tree ... hash table)

Penn ESE5320 Fall 2024 -- DeHon

96

Big Ideas

- Can adaptively compress data using recurring substrings
 - With constant work for symbol
- Rich design space for engineering associative map solutions

Admin

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