

ESE532: System-on-a-Chip Architecture

Day 17: October 30, 2017
Associative Maps, Hash Tables



Today

- Motivation/Reminder
- Hardware Associative Memory
- Software Maps
 - Hash Tables
- Hardware Maps
 - Associative Memories from BRAMs
 - Multi-hash Hash tables from BRAMs

Message


- Rich design space for Maps
- Hash tables are useful tools
- Multiple hashes are better than one

Reminder from Last Time

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

Deduplication Architecture



Associative Memory

- Maps from a key to a value
- Key not necessarily dense
 - Contrast simple RAM

Penn ESE532 Fall 2017 -- DeHon

7

Memory Tree Algorithm

```
curr – pointer into input chunk
// follow tree
y=0;
while(encode[x][input[curr]]!=NONE)
    x=encode[x][c]; y++;
if (y>0)
    send <x,y>
    send input[curr]
    encode[x][input[curr]]=curr
```

Penn ESE532 Fall 2017 -- DeHon

8

Compact Memory

- `int encode[SIZE][256];`
- Table is very sparse
- Store as associative memory
 - At most SIZE entries

Penn ESE532 Fall 2017 -- DeHon

9

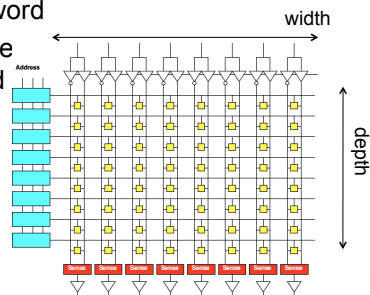
Custom Hardware Associative Memory

Penn ESE532 Fall 2017 -- DeHon

10

Memory Block Review

- Match on address
- Select wordline for a row
- Reads out a word
- Address dense and hardwired
- One row for each $2^{A\text{bits}}$ values

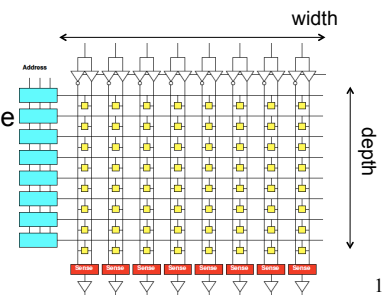


Penn ESE532 Fall 2017 -- DeHon

11

Memory Block Review


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



Penn ESE532 Fall 2017 -- DeHon

12

Programmable Key

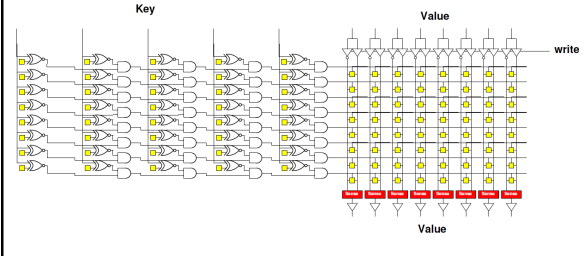


- Make row select programmable

Penn ESE532 Fall 2017 -- DeHon

13


Associative Memory Bank



Penn ESE532 Fall 2017 -- DeHon

14

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 2017 -- DeHon

15

Associate Memory Cost

- More expensive than equal capacity SRAM memory bank
 - Memory cells in decoder
 - No sharing of AND-terms in decoder
 - Need to support write into key



Penn ESE532 Fall 2017 -- DeHo

16

Software Map

Penn ESE532 Fall 2017 -- DeHon

17

Software Map

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

Penn ESE532 Fall 2017 -- DeHon

18

Map Implementations

- How could we implement Map in software?

Penn ESE532 Fall 2017 -- DeHon

19

Preclass 1

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

Penn ESE532 Fall 2017 -- DeHon

20

Tree

- 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 2017 -- DeHon

21

Tree Insert

- Need to maintain balance
- Doable with $O(\log(N))$ insert
 - Tricky
 - See Red-Black Tree

Penn ESE532 Fall 2017 -- DeHon

22

High Performance Map


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

Penn ESE532 Fall 2017 -- DeHon

23

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
 - Great, got direct lookup



Penn ESE532 Fall 2017 -- DeHon

24

Preclass 3a


- Average number of entries per hash when $N > \text{HASH_CAPACITY}$?

Penn ESE532 Fall 2017 -- DeHon

25

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




Penn ESE532 Fall 2017 -- DeHon

26

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 challenge if bucket can get big



Penn ESE532 Fall 2017 -- DeHon

27

Preclass 3b

- Probability of conflict if $N \ll \text{HASH_CAPACITY}$?
- How can we reduce bucket sizes?

Penn ESE532 Fall 2017 -- DeHon

28

Preclass 4

$N=1024$

m→	0	1	2	3	4+
C=1024	0.37				
C=2048					
C=4096					

$$\binom{N}{m} \left(\frac{1}{C}\right)^m \left(1 - \frac{1}{C}\right)^{N-m}$$

Penn ESE532 Fall 2017 -- DeHon

29

Hash

- Can tune hash parameters to control distribution
- Spend more memory → smaller buckets
 - less work finding things in buckets
 - Memory-Time tradeoff
- Still have possibility of large buckets

Penn ESE532 Fall 2017 -- DeHon

30

Perfect Hash

- **Perfect hash:** every key maps to a unique hash value
- Can create perfect hash if know keys in advance

Penn ESE532 Fall 2017 -- DeHon

31

Graph and Edge Model

- Use 2 hashes
- Each hash value is node
- Each entry is edge
 - 2 hashes, connects 2 nodes
- Program output by filling in value for nodes
- add node values modulo capacity to get edge value

```
int G[capacity];
int value[capacity];
h0=hash0(key);
h1=hash1(key);
addr=(G[h0]+G[h1])
%capacity;
res=value[addr];
```


ISP-Hive July 2016

32

Graph and Edge Model

- Use 2 hashes
- Each hash value is node
- Each entry is edge
 - 2 hashes, connects 2 nodes
- Program output by filling in value for nodes
- add node values modulo capacity to get edge value

Consider insert 0xB0 in each case




ISP-Hive July 2016

33

Case 3 often Solvable

Consider insert 0xB0 in each case




ISP-Hive July 2016

34

Case 3 not always solvable

Consider insert 0xB0 in each case




ISP-Hive July 2016

35

Perfect Hash Theory

- If graph is acyclic
 - Can always assign values to nodes to satisfy/produce any edge values
 - Walk graph in order from starting point



ISP-Hive July 2016

36

Perfect Hash Theory

- If graph is acyclic
 - Can always assign values to nodes to satisfy/produce any edge values
- Random graphs with a sufficiently low edge to node ratio are acyclic with high-probability
 - Good hash functions (random) makes graph random

ISP-Hive July 2016

37

Perfect Hash

- If do get cycle
 - Change hash
 - Change modulus
 - (add a bit more capacity)

Penn ESE532 Fall 2017 – DeHon

38

Variants for Hardware

- Modulus can be expensive
 - Unless it is a power of 2
- Set capacity to next power of 2
 - ...and modulus becomes simple
 - Can use xor instead of addition
 - No delay for carry, $O(1)$ combining

Penn ESE532 Fall 2017 – DeHon

39

Perfect Hash

- Can find tools to compute the perfect hash (e.g. gnuperf)

Penn ESE532 Fall 2017 – DeHon

40

Hardware Map

Penn ESE532 Fall 2017 – DeHon

41

FPGA


- Has BRAMs – normal memories, not associative
- 36Kb BRAM
 - 1024x36
- Can be 10b key → 36b value
 - Just using the memory sparsely

Penn ESE532 Fall 2017 – DeHon

42

Assoc. Mem from BRAM

- For wider match
- Cover 10b of key with each BRAM
- Use 36 output bits to indicate if one of 36 entries match
- AND together corresponding entries
- Get 36 match bits
- Re-encode match bits to lookup value



Penn ESE532 Fall 2017 -- DeHon


BRAM Associative Memory

- How would we expand capacity?

Penn ESE532 Fall 2017 -- DeHon

44

BRAM Associative Memory




Penn ESE532 Fall 2017 -- DeHon

45

Associative Memory Cost

- Match unit
 - Requires 1 BRAM per 10b of key per 36 entries
 - $(\text{keylen}/10\text{b}) * (\text{capacity}/36)$




Penn ESE532 Fall 2017 -- DeHon

46

Perfect Hash

- Can do perfect hash very compactly
- Hash to $\log(\text{capacity})$
 - Two different hashes
- Use separate BRAM for each hash
- Xor together values
- Use as dense address to lookup values



Penn ESE532 Fall 2017 -- DeHon

47

Multi Hash

- Perfect hash hints that multiple hashes may be more useful
- E.g.
 - Setup several hashes and memories
 - Compute each hash
 - See if any are free and put there
 - (in software, put in smallest bucket found reduce maximum bucket size)

Penn ESE532 Fall 2017 -- DeHon

48

Multihash

- K hashes
- Probability of conflict in all K tables


$$\left(\frac{N}{C}\right)^K$$

Penn ESE532 Fall 2017 -- DeHon

49

dMHC

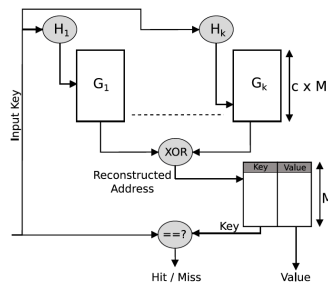
- Dynamic Multi-Hash Cache



Penn ESE532 Fall 2017 -- DeHon

50

Two-Level dMHC




Penn ESE532 Fall 2017 -- DeHon

51

Perfect Hash: Case 3 often Solvable

Consider insert 0xB0 in each case



ISP-Hive July 2016

52

Conflict: victimize and reinsert

Consider insert 0xB0 in each case




- With high likelihood, the item victimized has freedom

ISP-Hive July 2016

53

Perfect Hash: Case 3 not always solvable

Consider insert 0xB0 in each case




ISP-Hive July 2016

54

May not be able to solve

Consider insert 0xB0 in each case



- Due to cycle
- ...or just too long of a chain to repair in limited time

ISP-Hive July 2016

55

On Conflict Reinsert


- Nearly associative – occasionally need to drop an entry
- Makes less than perfect
 - Approximates an associative memory
- Potentially acceptable in some scenarios

Penn ESE532 Fall 2017 – DeHon

56

dMHC Lookup Fast

- Hash (few levels of xor)
- BRAM read
 - 2-3ns
- Xor (one level)
- Can run 1/cycle @200MHz +




Penn ESE532 Fall 2017 – DeHon

57

dMHC Lookups Fast

- Virtex 6 (one generation older than Zynq)



Penn ESE532 Fall 2017 –

58


Hybrid

- If the number of conflicts (things we might drop) are small
 - Maybe just save those in a true associative memory.
- Use dMHC for bulk
- When get (unresolvable) conflict
 - Stick in small associative memory
 - Use BRAM associative memory
- Lookup in both in parallel

Penn ESE532 Fall 2017 – DeHon

59

Hybrid dMHC+Assoc.



Penn ESE532 Fall 2017 – DeHon

60

Design Space

- Area-Time tradeoffs
 - Minimal with sequential search
 - Hash capacity (more memory, less time)
 - Build fully associative – spend area, instead of time
- Area-Quality tradeoffs
 - dMHC tuning c, k
- Pre-computation/specialization opportunities when data known early
 - Perfect hash

Penn ESE532 Fall 2017 – DeHon

61

Big Ideas

- Sparse, near $O(1)$ Map access → Hash Table
- If know set of keys in advance, can build perfect hash
- More hashes, reduce conflicts
- Rich design space for engineering associative map solutions

Penn ESE532 Fall 2017 – DeHon

62

Admin

- **No class Wednesday**
- Go to Horowitz Talk
 - *Life of Breaking and Making: What Happens When You Mix a Driving Curiosity and Episodic Overconfidence*
 - Wu & Chen at 3pm (classtime)
- Olukotun Thursday
 - Scaling Machine Learning performance...
 - Wu & Chen at 3pm
- First project milestone due Friday
 - Including teaming

Penn ESE532 Fall 2017 – DeHon

63