

Cache-efficient Data Structures

Jörn Engel

Lazybastard.org

November 2, 2007

The RAM Myth

Myth:
RAM is *Random* Access Memory

The RAM Myth

Truth:

That used to be the case

Please read Ulrich Drepper's paper at LWN

The RAM Myth

If accessed truly randomly,
DRAM can read one word in 60ns,
giving a bandwidth of 16MW/s

The RAM Myth

RAM has become a high-latency high-bandwidth device


The RAM Myth

just like a hard disk

The RAM Myth

Today's low-latency device is the L1-Cache

The RAM Myth



Data Structure Overview

Machine model:

- 16KiB of L1 data cache
- No L2 cache

Data Structure Overview

Handle 1,000,000 objects

- Insert
- Remove
- Lookup

Data structures are judged by

- Density
- Cache misses per operation

Single-linked list




Single-linked list (unsorted)

	Density	Fan-out	Insert	Remove	Lookup
Sequential	100%	1	0	500,000	500,000
Random	100%	1	0	500,000	500,000
Worst-case	100%	1	0	999,999	999,999

Single-linked list (sorted)

	Density	Fan-out	Insert	Remove	Lookup
Sequential	100%	1	0	999,999	1
Random	100%	1	500,000	500,000	500,000
Worst-case	100%	1	999,999	999,999	999,999

non-embedded Single-linked list



non-embedded Single-linked list

	Density	Fan-out	Insert	Remove	Lookup
Sequential	50%	1	0	999,999	2
Random	50%	1	500,000	500,000	500,000
Worst-case	50%	1	999,999	999,999	999,999


Double-linked list




Double-linked list

	Density	Fan-out	Insert	Remove	Lookup
Sequential	50%	1	0	1	1
Random	50%	1	500,000	500,000	500,000
Worst-case	50%	1	999,999	999,999	999,999

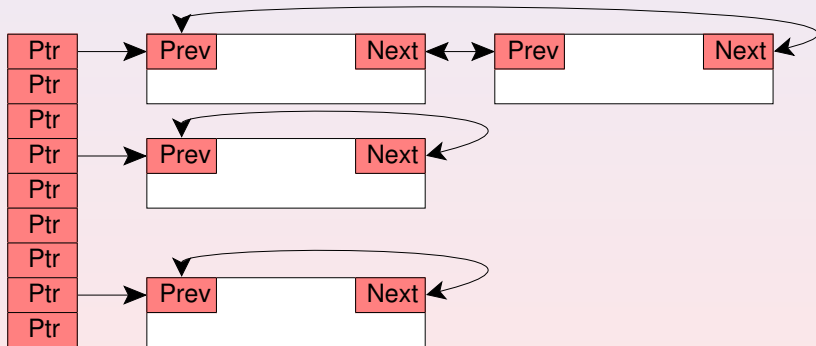
Hash table



Hash table with single-linked list



Hash table with double-linked list



Sparse hash table

	Density	Fan-out	Insert	Remove	Lookup
Sequential	<10%	∞	1	1	2
Random	<10%	∞	1	1	2
Worst-case	<10%	∞	1	1	2

Dense hash table

	Density	Fan-out	Insert	Remove	Lookup
Sequential	50%	mixed	1.5	1.5	2.5
Random	50%	mixed	1.5	1.5	2.5
Worst-case	50%	mixed	1.5	1.5	2.5

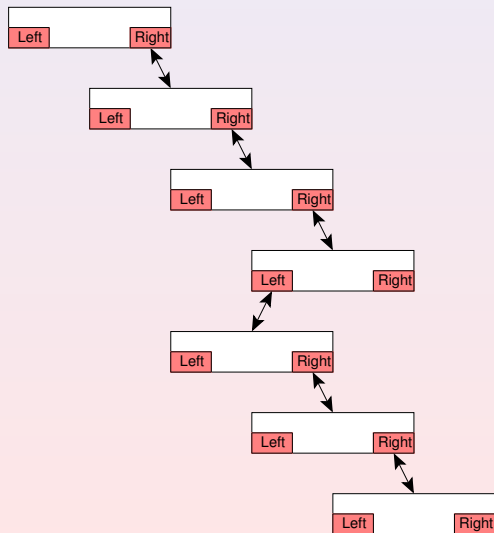
Full hash table

	Density	Fan-out	Insert	Remove	Lookup
Sequential	~100%	mixed	>2	>2	>3
Random	~100%	mixed	>2	>2	>3
Worst-case	~100%	mixed	>2	>2	>3

Memory consumed by Linux kernel hashes

Table	32b/64b <512MiB RAM	32b/64b ≥512MiB RAM
TCP established	96k/192k	384k/768k
TCP bind	64k/128k	256k/512k
IP route cache	128k/256k	512k/1M
Inode-cache	64k/128k	64k/128k
Dentry cache	32k/64k	32k/64k
Total	384k/768k	1248k/2496k


Unbalanced Binary Tree



Red-Black Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	<33%	2	0	1	1
Random	<33%	2	14	14	14
Worst-case	<33%	2	14	14	14


Radix Tree



Radix Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	97%	32	.03	.03	1.03
Random	3%	1-32	3-5	3-5	4-6
Worst-case	3%	1-32	3-5	3-5	4-6

B⁺-Tree



B⁺-Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	24%	8	.13	.13	1.13
Random	36%	12	3-4	3-4	4-5
Worst-case	24%	8	4-5	4-5	5-6

B_#-Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	48%	16	.06	.06	1.06
Random	42%	14	2-4	2-3	3-4
Worst-case	24%	8	4-6	4-5	5-6

B*-Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	32%	11	.09	.09	1.09
Random	40%	13	2-5	2-5	3-5
Worst-case	32%	11	3-6	3-6	4-6

B*#-Tree

	Density	Fan-out	Insert	Remove	Lookup
Sequential	32%	11	.09	.09	1.09
Random	44%	15	2-4	2-4	3-4
Worst-case	32%	11	3-6	3-6	4-6

B*#-Tree (B#-Tree for outer sibling)

	Density	Fan-out	Insert	Remove	Lookup
Sequential	48%	16	.06	.06	1.06
Random	43%	15	2-4	2-4	3-4
Worst-case	30%	11	3-6	3-6	4-6

Judy-Tree



Judy-Tree

? ! @ # \$ %