Big- \mathcal{O} Ain't What it Used to Be

unix weapons school

CC3.0 share-alike attribution copyright © 2013 nick black

Asymptotic notation review I

Asymptotic analysis gives us a means of speaking of arbitrarily large growth, independently of arbitrarily (but finitely) large costs not associated with problem size.

Notation	Name	Definition	Introduced
$f(n) \in \mathcal{O}(g(n))$	Big O	$\exists k > 0, \exists n_0, \forall n$	Paul Bachmann
		$n > n_0 \implies f(n) \le g(n) * k$	(1894)
$f(n) \in o(g(n))$	Small O	$\forall k > 0, \exists n_0, \forall n$	Edmund Landau
		$n > n_0 \implies f(n) \le g(n) * k$	(1909)
$f(n) \in \Theta(g(n))$	Big Theta	$\exists k_1 > 0, \exists k_2 > 0, \exists n_0, \forall n$	Donald Knuth
		$n > n_0 \implies g(n) * k_1 \le f(n),$	(1976)
		$f(n) \le g(n) \ast k_2$	
$f(n) \in \omega(g(n))$	Small Omega	$\forall k > 0, \exists n_0, \forall n$	Donald Knuth
		$n > n_0 \implies f(n) \ge g(n) * k$	(1976)
$f(n) \in \overline{\Omega(g(n))}$	Big Omega	$\exists k > 0, \exists n_0, \forall n$	Donald Knuth
		$n > n_0 \implies f(n) > g(n) * k$	(1976)

Advances in (finite) computing technology can only reduce these ignored costs. Wrap the earth with your register file, and still there will be numbers so large that their addition is $\Theta(n)$.

Class (all $c_* > 1$)	Name	Example
$\mathcal{O}(1)$	Constant	Is word-sized unsigned int n a power of 2?
$\mathcal{O}(\lg_{c_1} \lg_{c_2} n)$	Double-log	Interpolative search on uniform distribution
$\mathcal{O}(\lg_c n)$	Logarithmic	Binary search
$\mathcal{O}(n \lg n) = \mathcal{O}(\lg n!)$	Linearithmic	FFT
$\mathcal{O}(n^c)$	Polynomial	Primality testing
$\mathcal{O}(c^n)$	Exponential	Brute-force Boolean equivalence
$\mathcal{O}(n!)$	Factorial	Unrestricted permutations of a poset
$\mathcal{O}(c_1^{c_2^n})$	Double-exp	Presburger arithmetic decision best case

Speaking of still faster growth rates¹ (hyper-exponential, \mathcal{A}) is mostly zoology.

¹Check out "fast-growing hierarchies" and the LöbWainer hierarchy. CS4803UWS at the Georgia Institute of Technology Big- \mathcal{O} Ain't What it Used to Be Algorithmic choices can dominate performance, especially at scale. By the definition of Big O, it should also be obvious that an asymptotically superior algorithm can be slower for small inputs².

That said, no one's going to think implementing a routing table with a linked list is a good idea.

Furthermore, asymptotic analysis speaks of performance as problem size grows. It doesn't speak of real-time. It doesn't speak of bounded memories. We rarely speak of piecewise asymptotics.

But, by all means, do ensure you're not doing linear searches on sorted data etc.

²We will see that small inputs can be surprisingly large. CS4803UWS at the Georgia Institute of Technology Big- \mathcal{O} Ain't What it Used to Be

Naive square (nXn X nXn) matrix multiplication is $\Theta(n^3)$.

$$C = AB \implies C_{ij} = \sum_{m=1}^{k} A_{im} B_{mj}$$
 (1)

Counting the explicit additions and multiplications, there are precisely $2n^3$ "operations".

Fused multiply-add

IEEE 754-2008 floating point support requires FMA, fused multiply-add. Let rn() denote a rounding operation. Typically, a multiply-add chain requires two instructions, and rounds twice:

$$MAC(A, B, C) = rn(rn(A * B) + C)$$
(2)

Fused multiply-add rounds only once, preserving the fully precise product in an internal register:

$$FMA(A, B, C) = rn(A * B + C)$$
(3)

AMD's FMA4 (Bulldozer) implements a fully general SIMD FP FMA. Intel's FMA3 (Haswell, as part of AVX2; also in AMD's Piledriver) implements a destructive SIMD FP FMA³. The throughput and latency are equivalent to standard single SIMD FP adds and multiplies. NVIDIA's Fermi likewise introduced a full-throughput FMA.

There are now precisely n^3 "operations".

³AMD's XOP further implements an SIMD integer FMA.

Wide issue



Haswell can issue and retire 2 VFMADD* instructions per cycle.

There are now precisely $n^3/2$ "operations"⁴.

⁴Assuming that two operations are available every cycle.

SIMD



AVX uses the 16 256-bit YMM registers. There are 8 32-bit IEEE 754-2008 single-precision values in a 256-bit input.

There are now precisely $n^3/16$ "operations"⁵.

⁵Assuming that values are usable in 256-bit chunks.

Multicore



Has well will likely debut in a quadcore physical package. There are now precisely $n^3/64$ "operations"^{*a*}.

 $^{a}\mathrm{Ignoring}$ communication costs, and assuming perfect parallelism.

CS4803UWS at the Georgia Institute of Technology

Memory accesses

- $\ \, \bullet \ \, \forall i,0\leq i\leq n: {\rm Read\ row\ } i {\rm\ from\ } A {\rm\ into\ } a$
- **2** $\forall j, 0 \leq j \leq n$: Read column j from B into b, read C_{ij} into c
- $\forall k, 0 \le k \le n : \text{Store } a_k * b_k + c \text{ into } C_{ij}$
- n^2 loads from A
- **2** n^3 loads from B
- **3** n^2 loads from C
- **4** n^2 stores to C

Counting the loads and stores, there are precisely $n^3 + 3n^2$ memory accesses. There are now precisely $\frac{65n^3}{64} + 3n^2$ "operations".

Arithmetic intensity $\lim_{n \to \infty} \frac{\text{Arithmetic}}{\text{Memory}} = 2$

Hong and Kung proved in 1981 that any schedule of conventional matrix multiplication must transfer $\Omega(\frac{n^3}{\sqrt{Z}}), Z < \frac{n^2}{6}$ words between slow and fast memory. *Tiling* is the optimal strategy.

Of course, AVX's VMOVAPS moves 256 bits, or 8 32-bit single precision floating point words, at a time. And there's 4 cores. With two load/store pipes each. So that's $\Omega_{HK}/64^6$.

Of course, we're not going to be able to pack two VFMADDPS and two VMOVAPS instructions into every 16B/c I\$ fetch⁷.

How does this interact with register banking? Multilevel caching? TLBs? Page cache? Prefetching? DRAM banking? Multilevel disk? Logical cores? Other physical cores? NUMA?

⁶Assuming that the values are located in aligned, contiguous 256-bite chunks in memory. Wait...we can use VMOVUPS if they're unsuitably aligned.

 $^{^7\}mathrm{VEX}\text{-encoded}$ VMOVAPS tends to run $\tilde{}5$ by tes.

Argh

FFFFFF

FFFFF

FFFFFFF

FFFF

CS4803 Spring 2010 Lab 3—All $\mathcal{O}(n^3)$



ATLAS Unleashed—All $\mathcal{O}(n^3)$



In the pure systems space, \mathcal{O} makes still less sense. What's \mathcal{O} of multithreaded file lexing?



CS4803UWS at the Georgia Institute of Technology

Big-O Ain't What it Used to Be

Nonetheless, optimization can be very fruitful.



Analysis-driven optimization I

samples	pcnt	kernel function
12022.00		read hpet
1766.00		spin lock irqsave
1269.00	4.7%	acpi os read port
1209.00	4.5%	hpet next event
410.00	1.5%	schedule
280.00	1.0%	do sys poll
273.00	1.0%	sched clock local
265.00	1.0%	fget light
241.00	0.9%	native read tsc
237.00	0.9%	spin lock
231.00	0.9%	spin unlock irgrestor

High-level ("Macroanalysis")

Coarse tools and algorithmic reasoning, e.g.:

- Ensure sufficient task-level parallelism
- Ensure cores aren't overutilized
- Profiler-driven hotspot location
- High-level memory and I/O flow

Analysis-driven optimization II



Complexity finds a way



A fractal having Hausdorff dimension 2, the space-filling Peano curve is used in cache-oblivious algorithms.

Mastering the complex modern design space requires deftly switching between theoretical analysis, guided experiment, and pure exploration.

This class is about doing so.

Insert 1/3

hyn) Taled a 1999(2) (2004) - WVS & Under CSEC230 FRIN 2007 Deather = fosts of 12 winh the 2 or muse in = raws will there Fire rows in 150+1, a(s,:)*0, where we & are (ma) mat 2#2 Nur c(j) in flam Ap : Sit of A ; n cuche Bp: starts D and some particle surp of c(1,1) is f DAL UN COMPAN ali, 1) * 5(k, j) Bin cala or with 2* D district A in rach Each phase P2 trading Z CCiD is 1: 11 in 12 it to main memory 15p+1 - 2:27 Set . dx 2D

CS4803UWS at the Georgia Institute of Technology

Insert 2/3

- (n) S 0/3 Mark 4 23 27 Conidir XNO EU X N in 50-No (\cdot, \cdot) News ! No SC mur phases (+) m Verd 0 0

CS4803UWS at the Georgia Institute of Technology

Insert 3/3

54 15 A1 menting me 5 Sim ann Lompita Congl ello. (nn nir inchien + unction gaml FJ 12-20 onthe a transfilm (1)5(1) 5 < Non computely group fasti 0 FT シンシン 50 ONUSIC 5 0 2

CS4803UWS at the Georgia Institute of Technology

- Hong and Kung. "I/O complexity: The red-blue pebble game" (1981).
- Yotov et al. "An experimental comparison of cache-oblivious and cache-conscious programs" (2007).
- Irony et al. "Communication Lower Bounds for Distributed-Memory Matrix Mul" (2004).
- Goto et al. "Anatomy of High-Performance Matrix Multiplication" (2008).
- Kalyanasundaram et al. "Improved Simulation of NTMs" (2011).
- François Le Gall. "Faster Algorithms for Rectangular Matrix Multiplication" (2012).
- Eric Quinnell. "Floating-Point Fused Multiply-Add Architectures" (2007).
- Tom Leighton. "Better Master Theorems for Divide-and-Conquer Recurrences" (1996).
- Intel Instruction Set Extensions Programming Reference