



OpenEye Scientific Software



GPUs: What is all the fuss about?


Brian (you're-better-off-calling-me-Bruce) Cole

CUP 11




A convincing bar graph

Shape Overlays per Second






A convincing bar graph





A convincing bar graph






- How did hardware get here?
- What is it good at?
- What is it bad at?



Moore's Law

CPU Transistor Counts 1971-2008 & Moore's Law


“I say,
if Gore invented
the Internet,
I invented the
exponential.”
– Gordon Moore




http://en.wikipedia.org/wiki/Moore's_law



Riding Moore's Law






Some Ancient History

	1980 Vax-11/750	Modern CPU	Improvement Since 1980
Clock speed (MHz)	6	3,000	500x
Memory Size (MB)	2	2,000	1,000x
Memory Bandwidth (MB/s)	13	7,000	540x
Memory Latency (ns)	225	70	3x
Memory Latency (cycles)	1.4	210	-150x


Do you c why?



Einstein vs Moore: Ding Ding



VS



speed of light / 3 GHz

Search

[Advanced Search](#)

Web [+ Show options...](#)

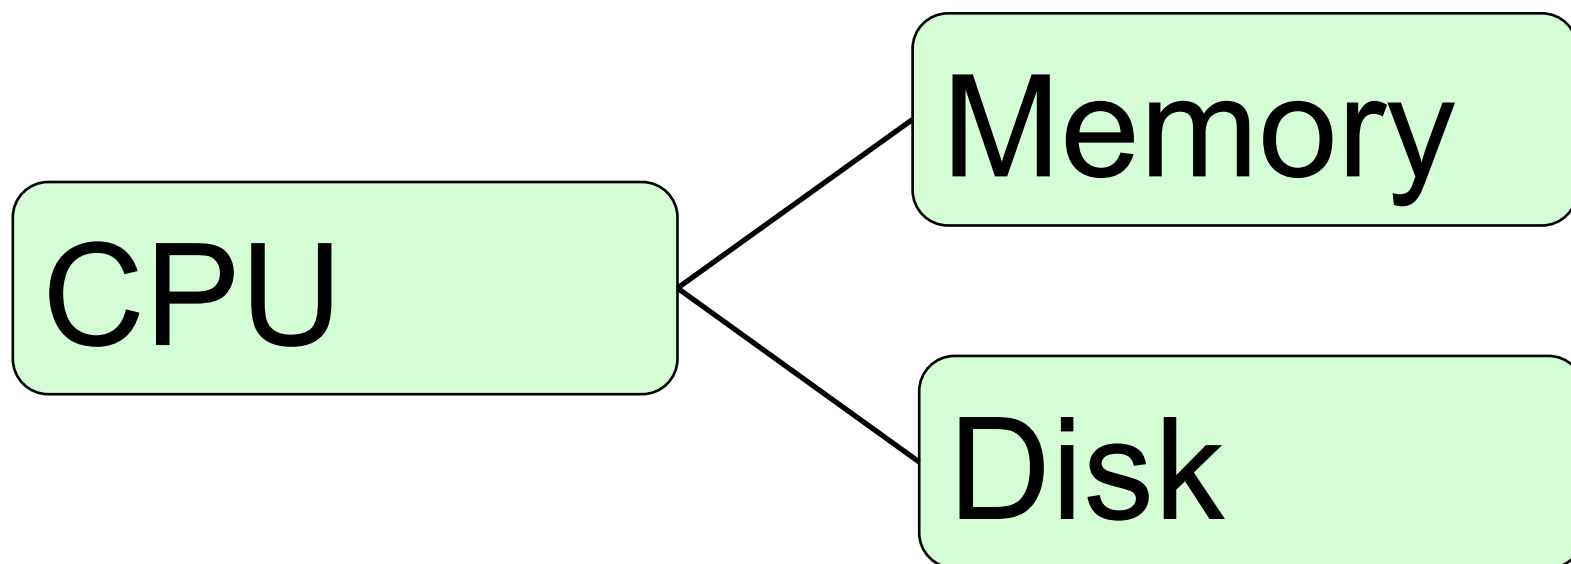
Results 1 - 10 of about 1,170,000 for speed of light / 3 GHz. (0.41 seconds)



the speed of light / (3 gigahertz) = 9.99308193 centimeters

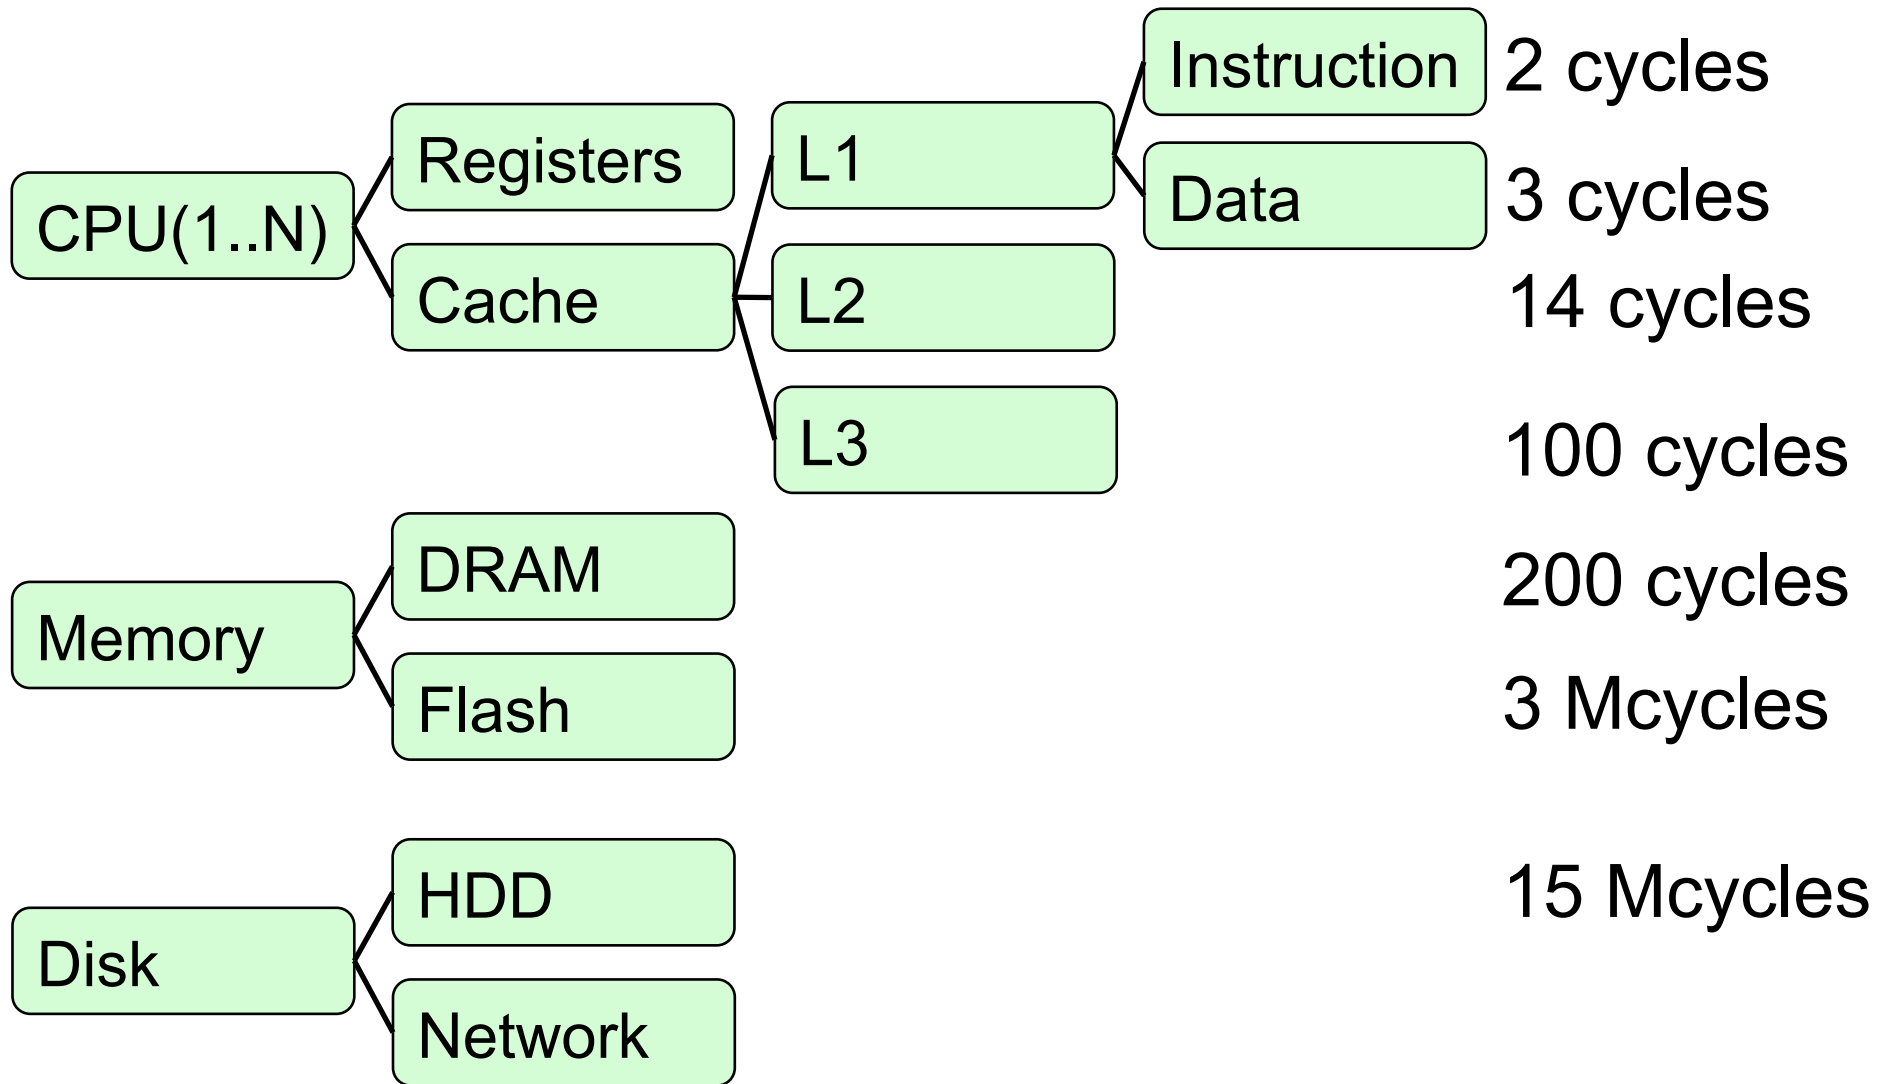
[More about calculator.](#)








Real life (simplified)





$$\text{Bandwidth} \times \text{Latency} = \text{Concurrency}$$





Adding two vectors

```
__kernel void vec_add (__global const float *a,  
                       __global const float *b,  
                       __global float *c)  
{  
    int gid = get_global_id(0);  
    c[gid] = a[gid] + b[gid];  
}
```


Warning: this sucks!



- Registers
 - There's a limited number so usage is fixed at kernel launch
 - Therefore, no recursion, or even a stack to be found
- Local memory
 - Bank conflicts
 - Every work item should ask for a **different** piece of memory
- Constant memory
 - Every work item should ask for the **same** piece of memory
- Global memory
 - Long latency
 - But high bandwidth!



"My job is to apply the formula"



$$\frac{FLOPS}{Bytes} \times \frac{GigaBytes}{Second} = \frac{GigaFLOPS}{Second}$$

Arithmetic Intensity ~~Bus~~ Bandwidth Arithmetic Throughput




Adding two vectors

```
__kernel void vec_add (__global const float *a,  
                      __global const float *b,  
                      __global float *c)  
{  
    int gid = get_global_id(0);  
    c[gid] = a[gid] + b[gid];  
}
```



"If X is less than the cost of a recall, we don't do one."

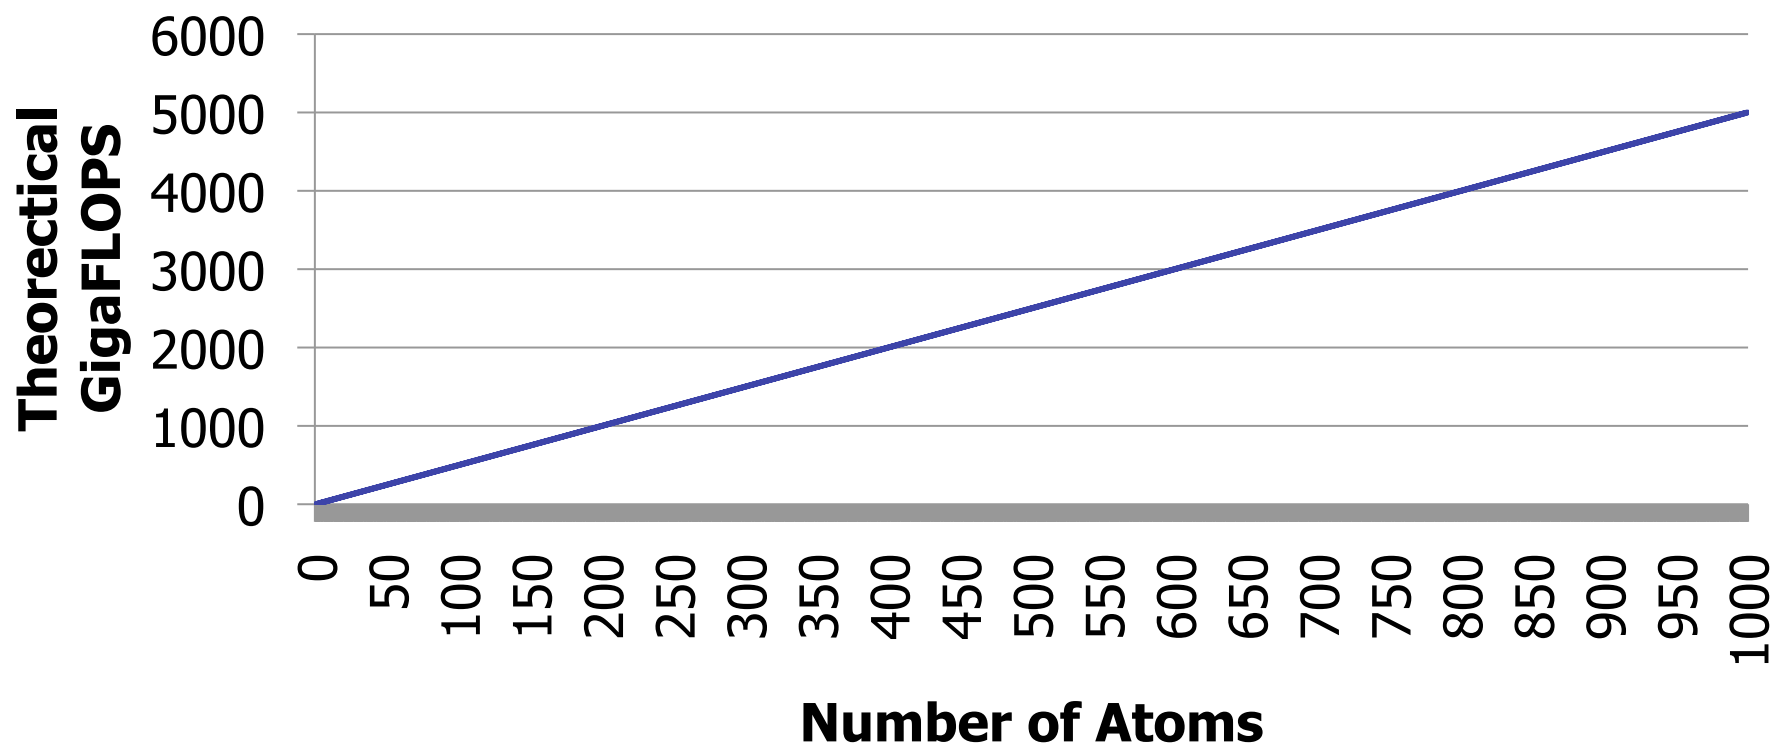


$$\frac{1FLOP}{12Bytes} \times \frac{5GigaBytes}{Second} = \frac{0.4GigaFLOPS}{Second}$$

Theoretical throughput of a Tesla card is 933 GigaFLOPS!



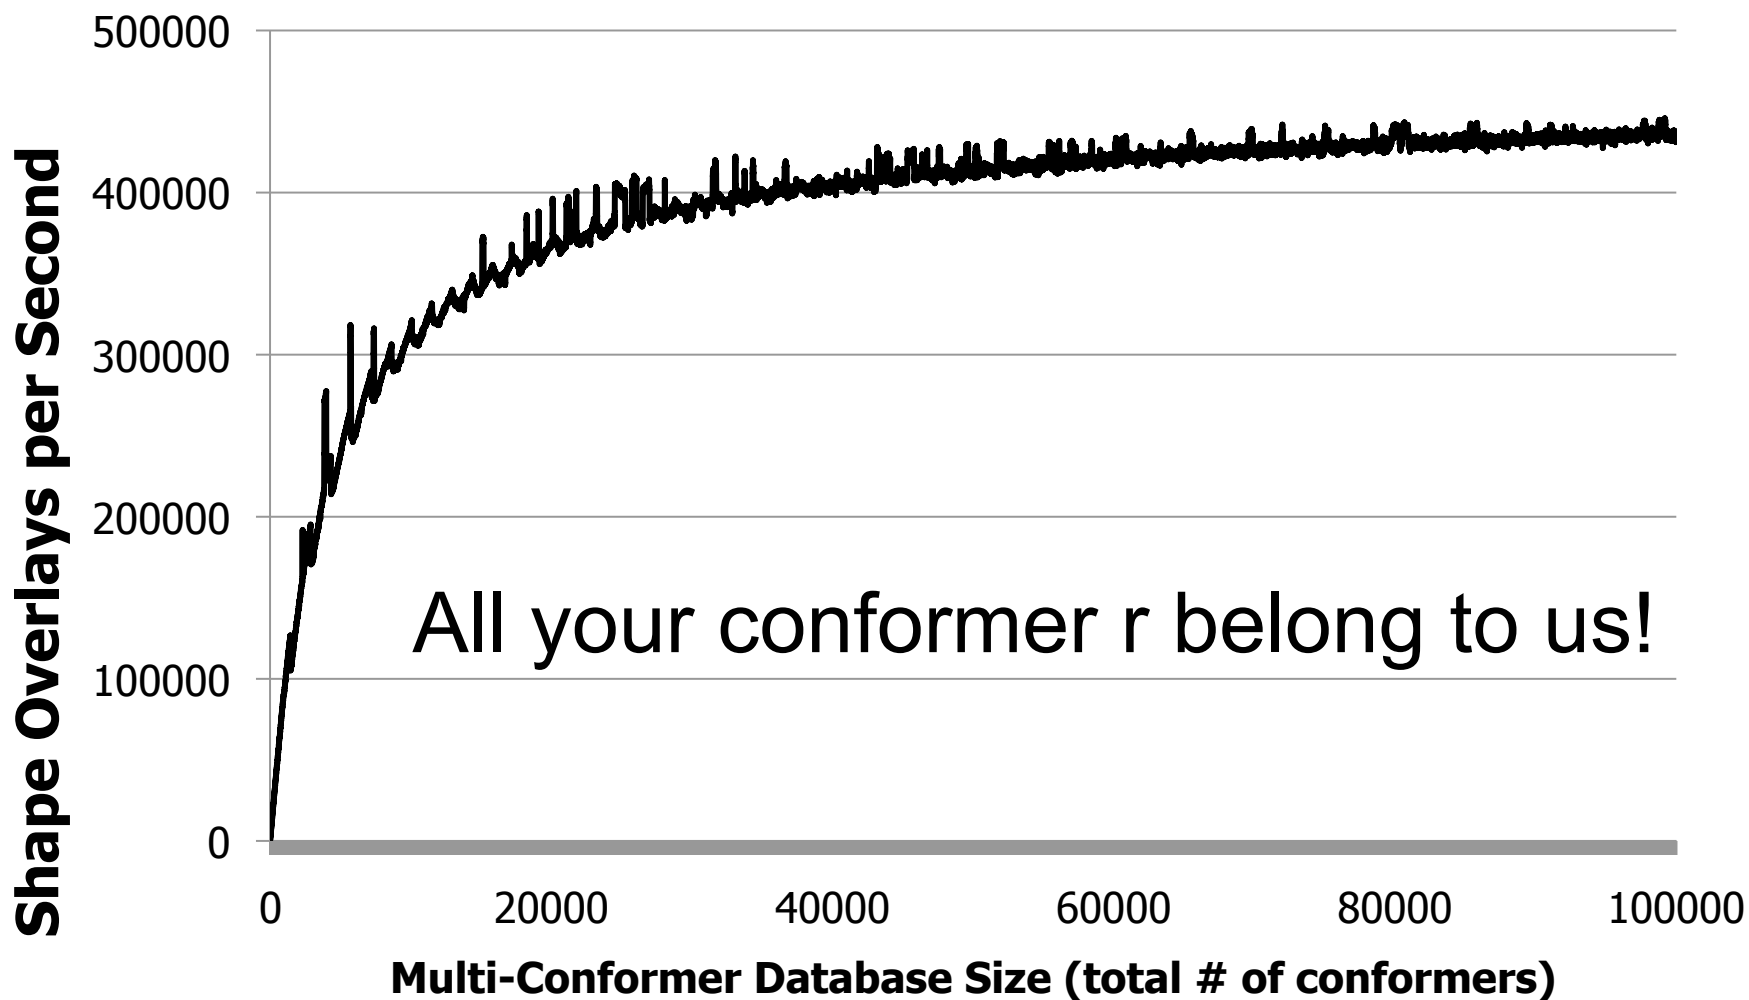
Making N^2 attractive




$$\frac{N^2 \text{ FLOPS}}{N \text{ Bytes}} \times \frac{5 \text{ GigaBytes}}{\text{Second}} = \frac{5N \text{ GigaFLOPS}}{\text{Second}}$$




Moore is better





Turning $O(N)$ into $O(N^2)$

- Determine Connected Components
 - Protein ligand splitting



- Don't expect a GPU OEChem any time soon (or ever)
 - Only very select (i.e. appropriate) algorithms can/will be ported



- Algorithms that port well
 - have **lots** of independent and uniform work
- Algorithms that don't port well
 - do lots of decision making



VIDA 4.0.0 for OpenEye (All)

GPU Search

Search: Submit Query Cancel

Max Hits:

List Window

Name	ID	#
ada-xtal-lig.mol2	1	1
GPU Result 1	5	101
****	4	
ASN_04887305	7	
ASN_08253722	9	
ASN_08255652	11	
ASN_06396361	13	
ASN_06396522	15	
ASN_06482832	17	
ASN_07441561	19	

Style Control List Window

Scripting Window

```
7799285/9616251
7799285/9616251
7799285/9616251

8188076/9616251
8188076/9616251

8487496/9616251
8583150/9616251
9538640/9616251
```

C1CCN(CC1)CCNC(=S)NC[C@H]2CCCCO2

Spreadsheet

	VIDA Name	VIDA ID	ShapeTanimoto
3	ASN_04887...	7	0.8939
4	ASN_08253...	9	0.8832
5	ASN_08255...	11	0.8821

Molecules Proteins Atoms Residues



- OpenEyers
 - El Presidente
 - Bob Tolbert
 - Geoff Skillman
 - Matt Stahl
 - Paul Hawkins
 - Brian Kelley
- NVidia
 - Peng Wang
- Herb Sutter (Microsoft)
 - <http://www.gotw.ca/>
 - <http://herbsutter.wordpress.com/>