

# Data @ Bare Metal Speed

*The Wisconsin Quickstep Project*

Jignesh M. Patel



THE UNIVERSITY  
*of*  
**WISCONSIN**  
MADISON

Blog: <http://bigfastdata.blogspot.com>

# Good News: Big Data Opportunity

Databases are at the core of the big data revolution

**Databases Rule!**



# Bad News: Big Data Software

What we have delivered looks like a pig with lipstick



Mick Coulas

# Good News: Big Data Opportunity

People are focused on the massive and real growth in this area



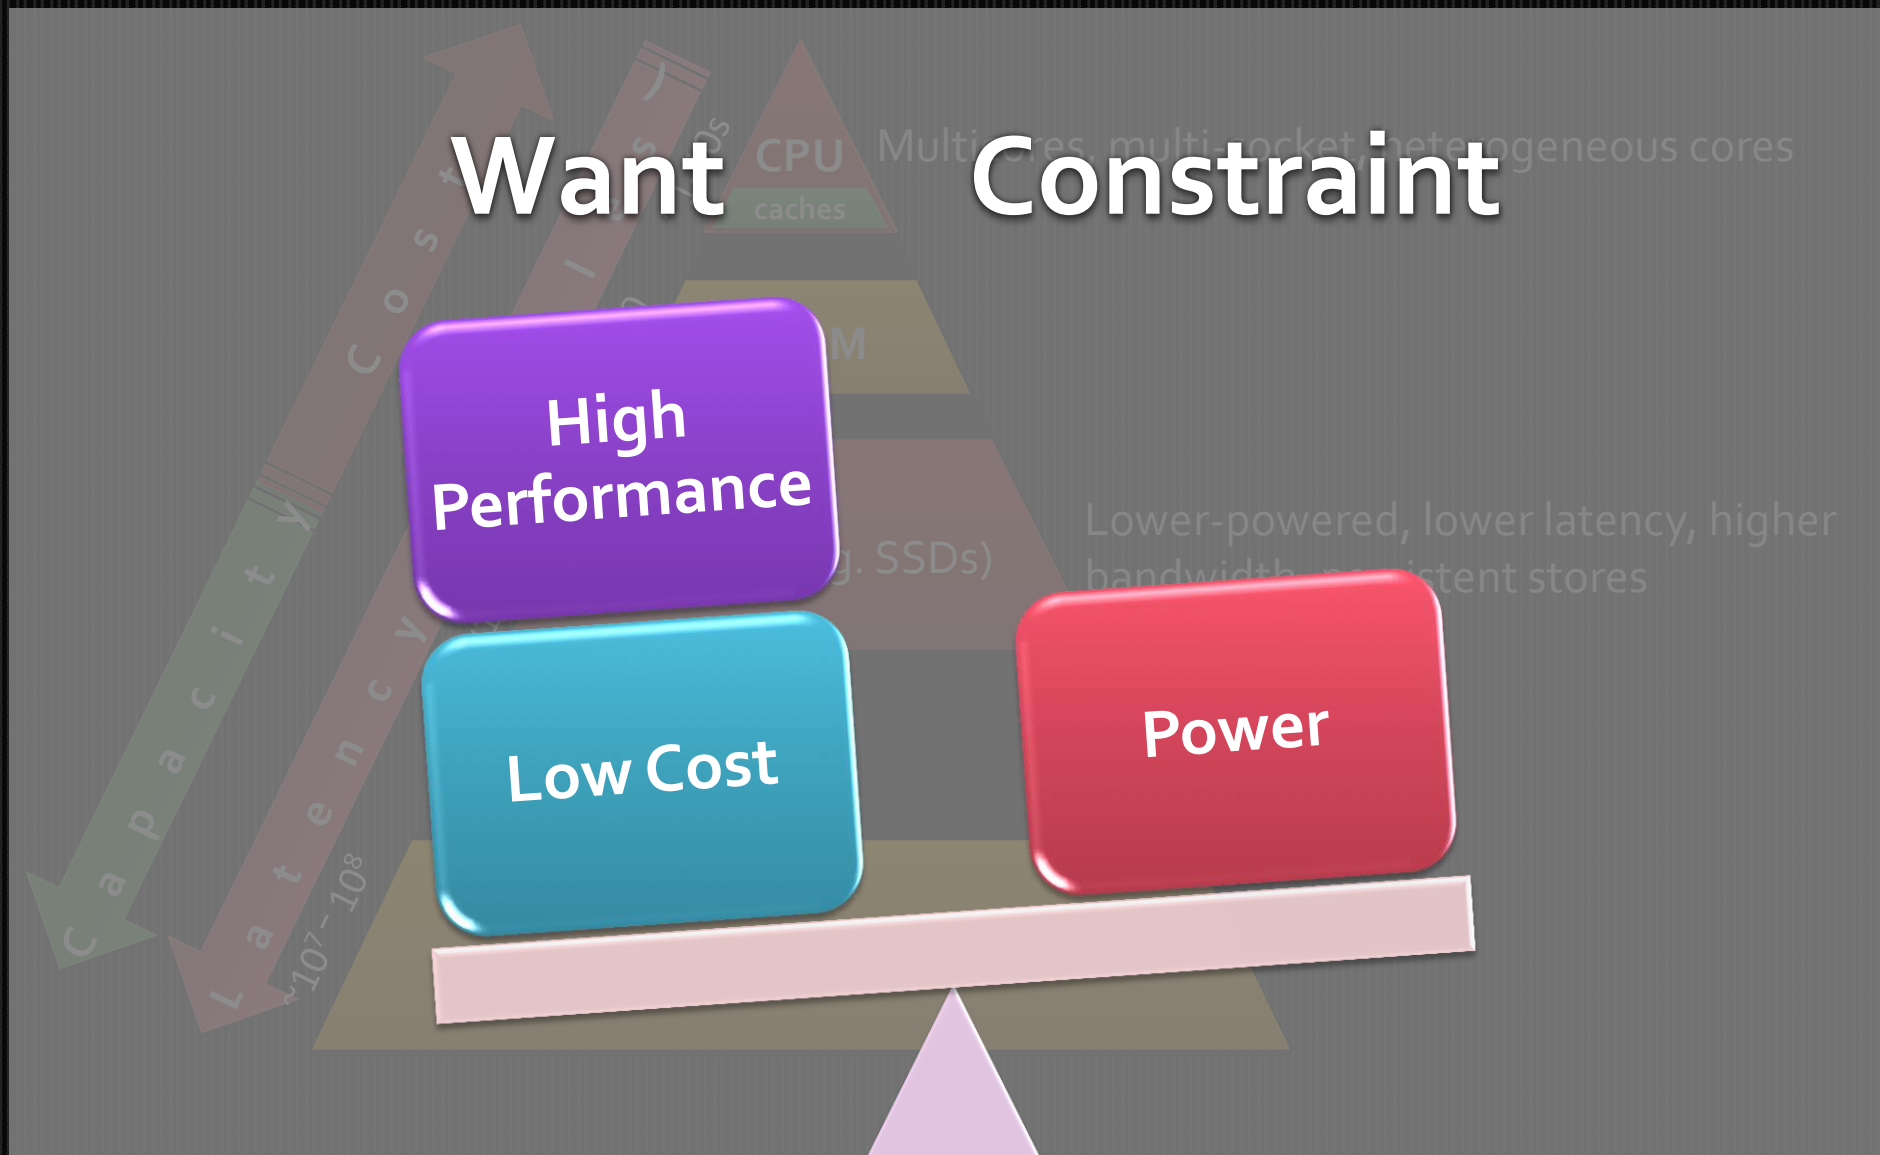
shutterstock

# Bad News: The game is up!

We can't hide for  
much longer  
a disruptive way

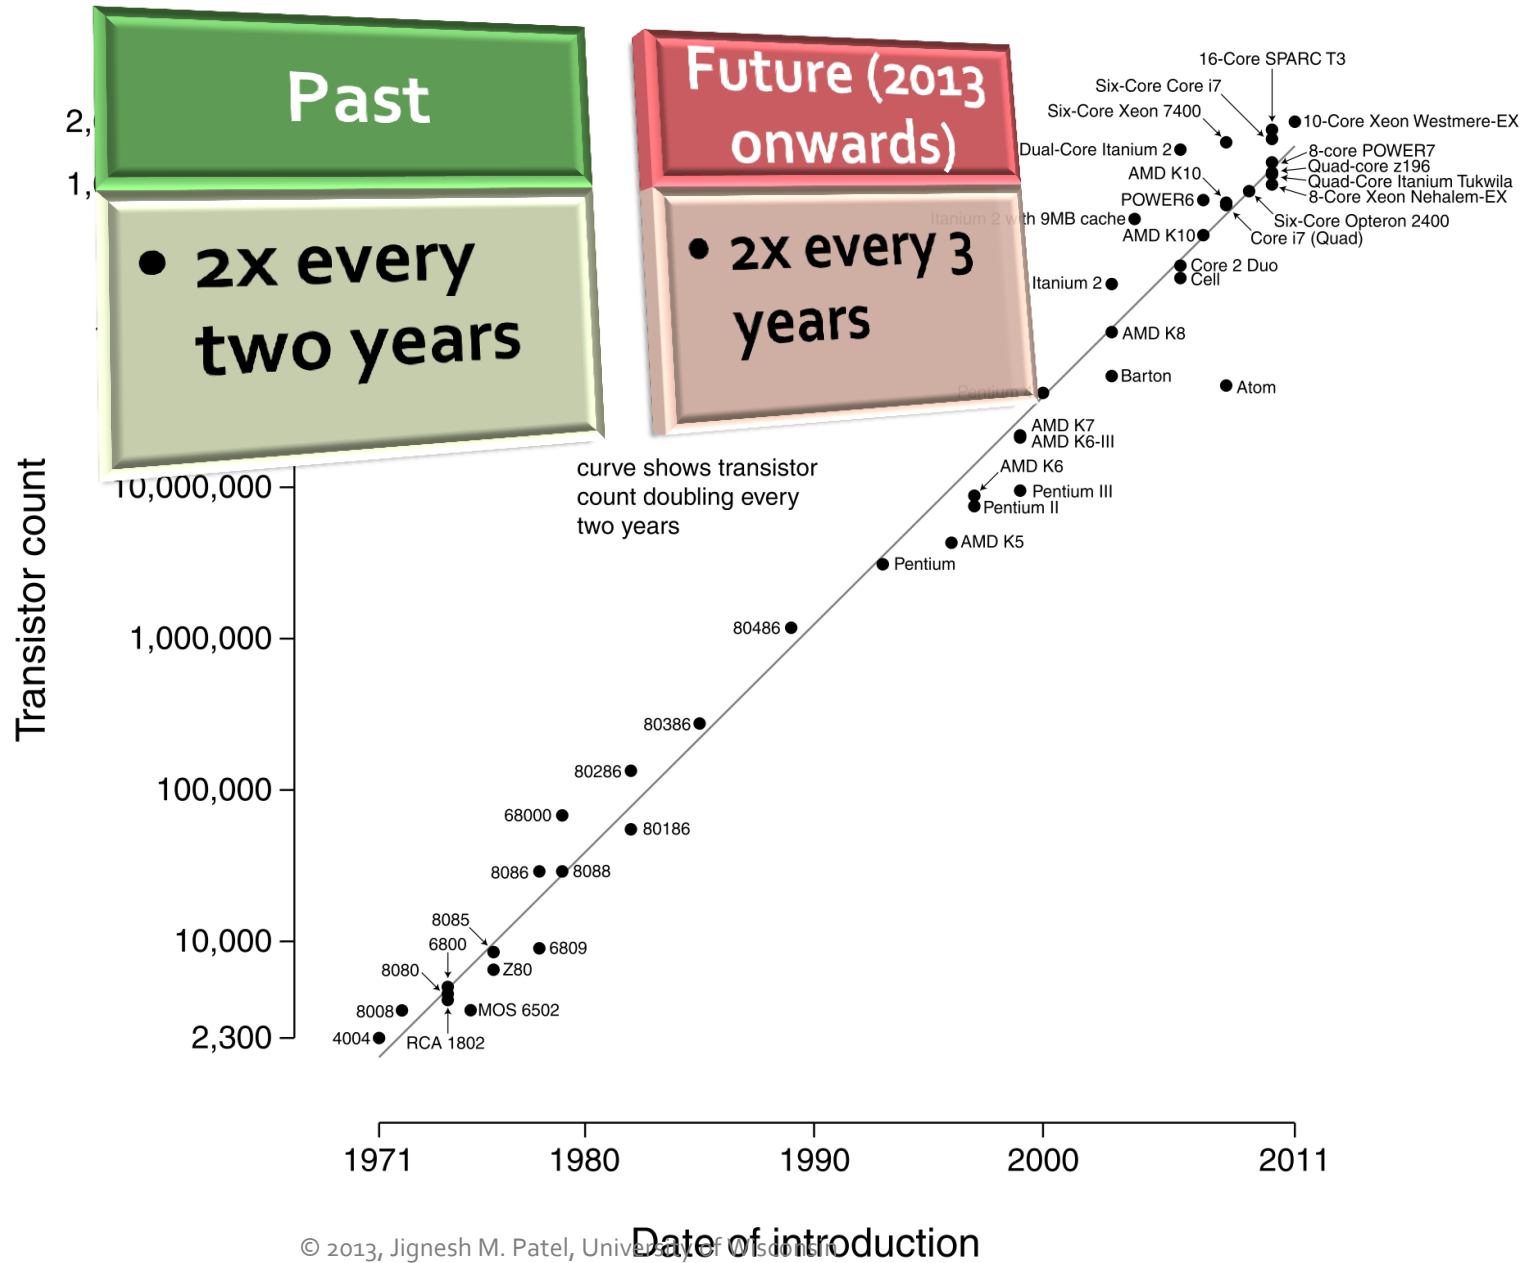


# Disruptive hardware trends

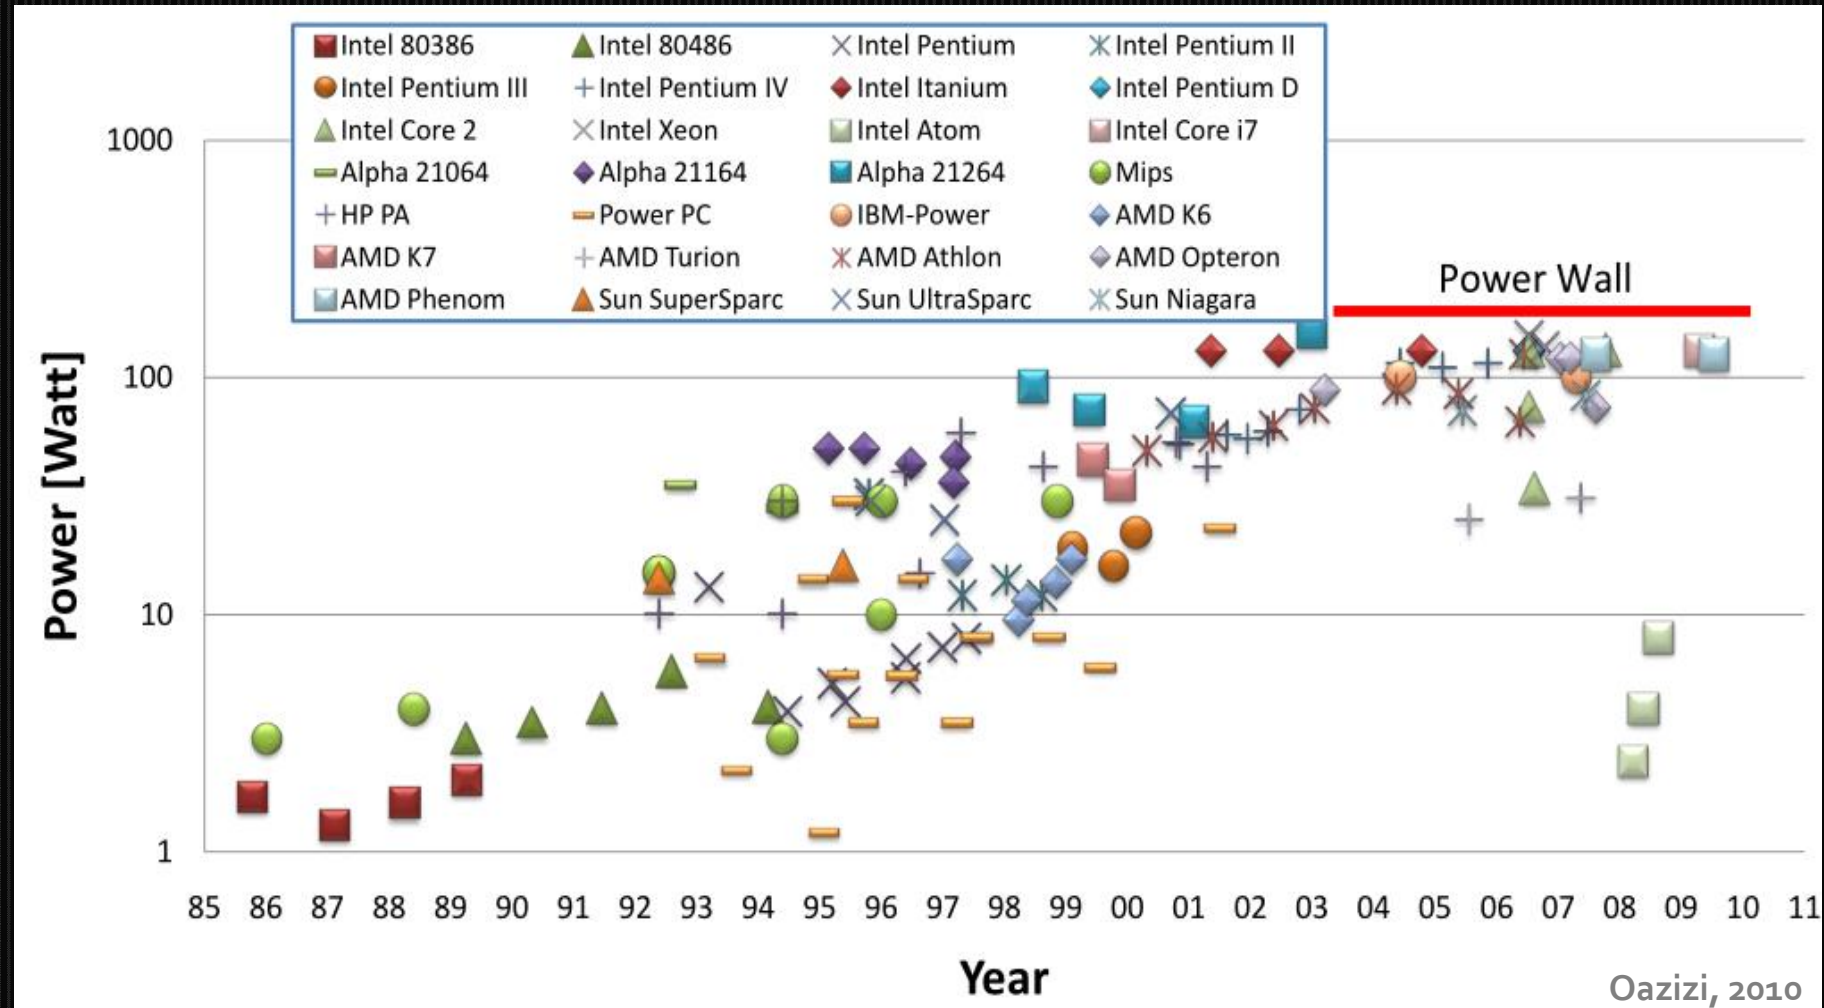


# Microprocessor Transistor Counts 1971-2011 & Moore's Law

Image credit:  
Wikipedia



# Moore's Law continues, but ...





# What's Next for Processors?

- Future processor design?
  1. Keep adding cores  
(~40% per generation)
  2. Heterogeneous cores
  3. Programmable functional units
- But, systems must work within a power budget
- Data growth continues unabated
- **Need to do more with less.**

# Scan: A Key Data Processing Kernel

## What?

- Scan a column of a table applying some predicate

## Why?

- A key primitive in database
- Often most invoked kernel in deep analytic processing

## How?

- Conserve memory bandwidth: **BitWeaving** the data
- Use every bit of data that is brought to the processor efficiently using **intra-cycle parallelism**

# Focus on Column Scan (can be generalized)

*Traditional Row Store*

shipdate	...	discount	quantity
Mar-12-2013		5%	5
Jan-08-2013		2%	4
Apr-29-2013		10%	3
May-14-2013		0%	6
...	...	...	...
Feb-28-2013		5%	0

*Column Store*

shipdate	...	discount	quantity
Mar-12-2013		5%	5
Jan-08-2013		2%	4
Apr-29-2013	...	10%	3
May-14-2013		0%	6
...		...	...
Feb-28-2013		5%	0

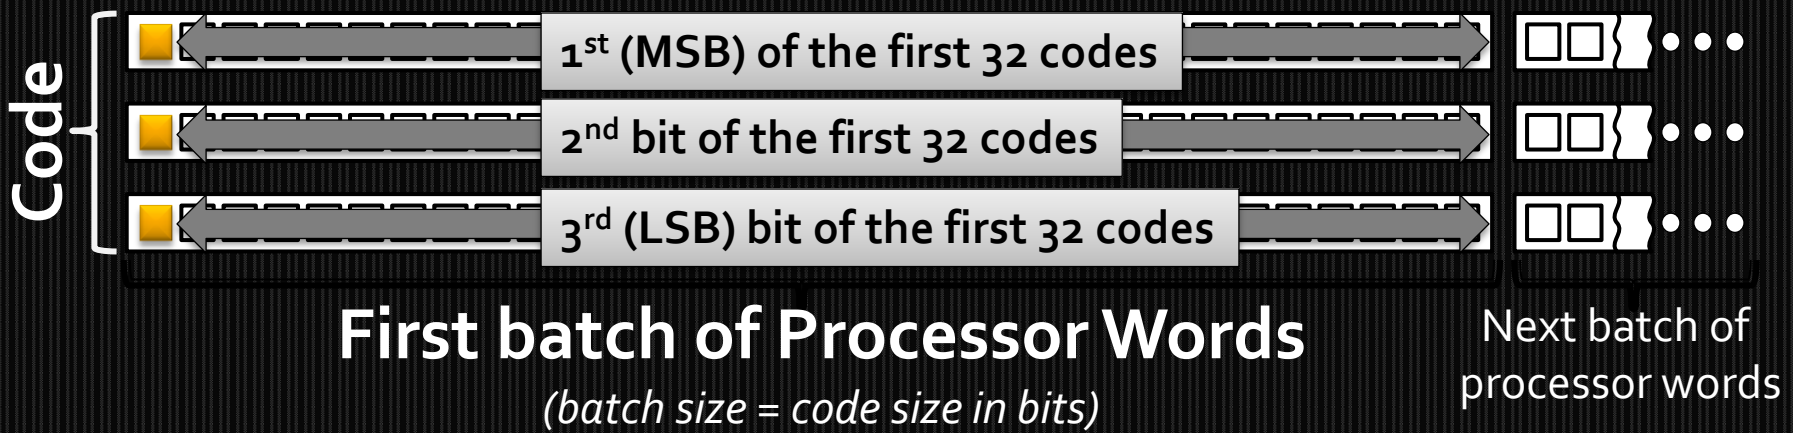
File: n-1      File: n



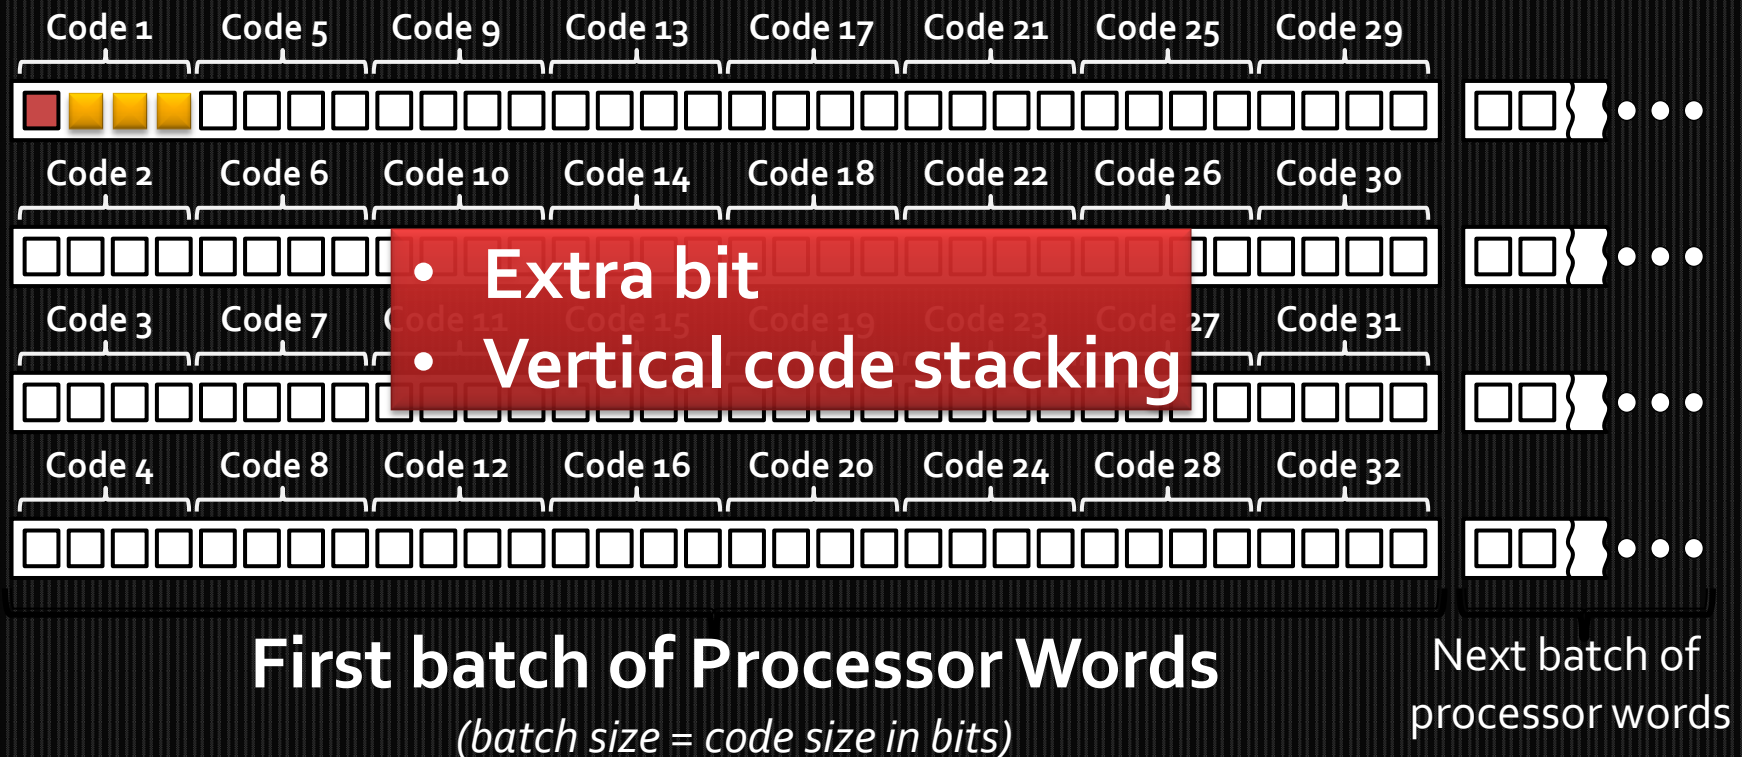
Column Codes: 

5	4	3	6	2	7	1	0
---	---	---	---	---	---	---	---

BitWeaving/V

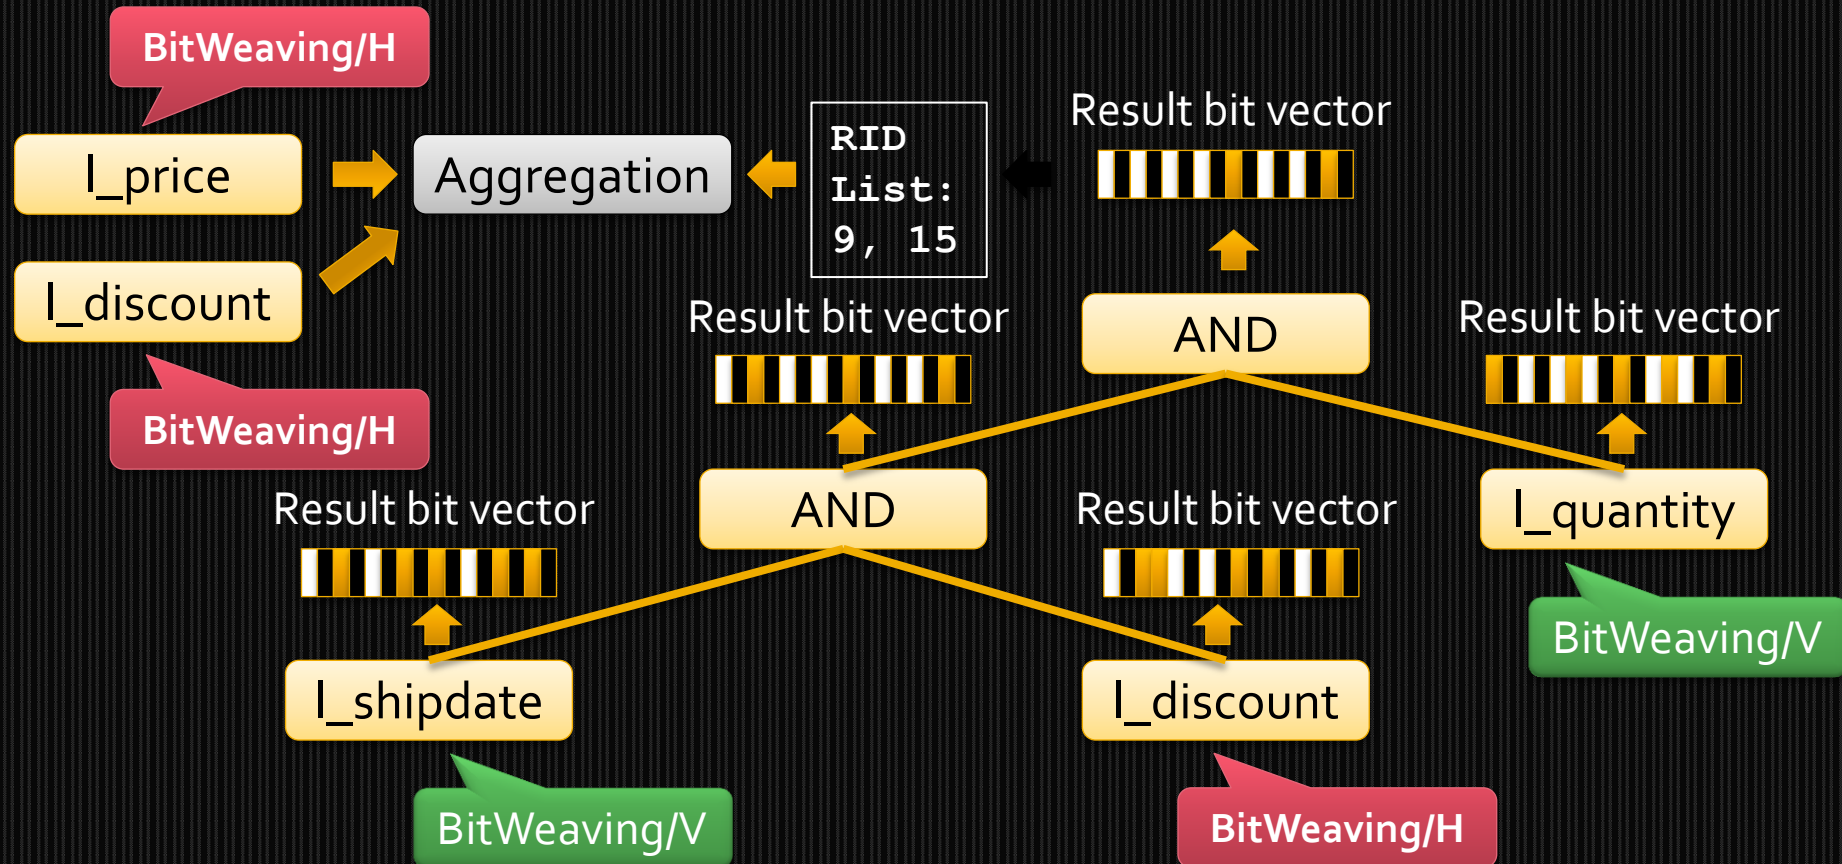



BitWeaving/H



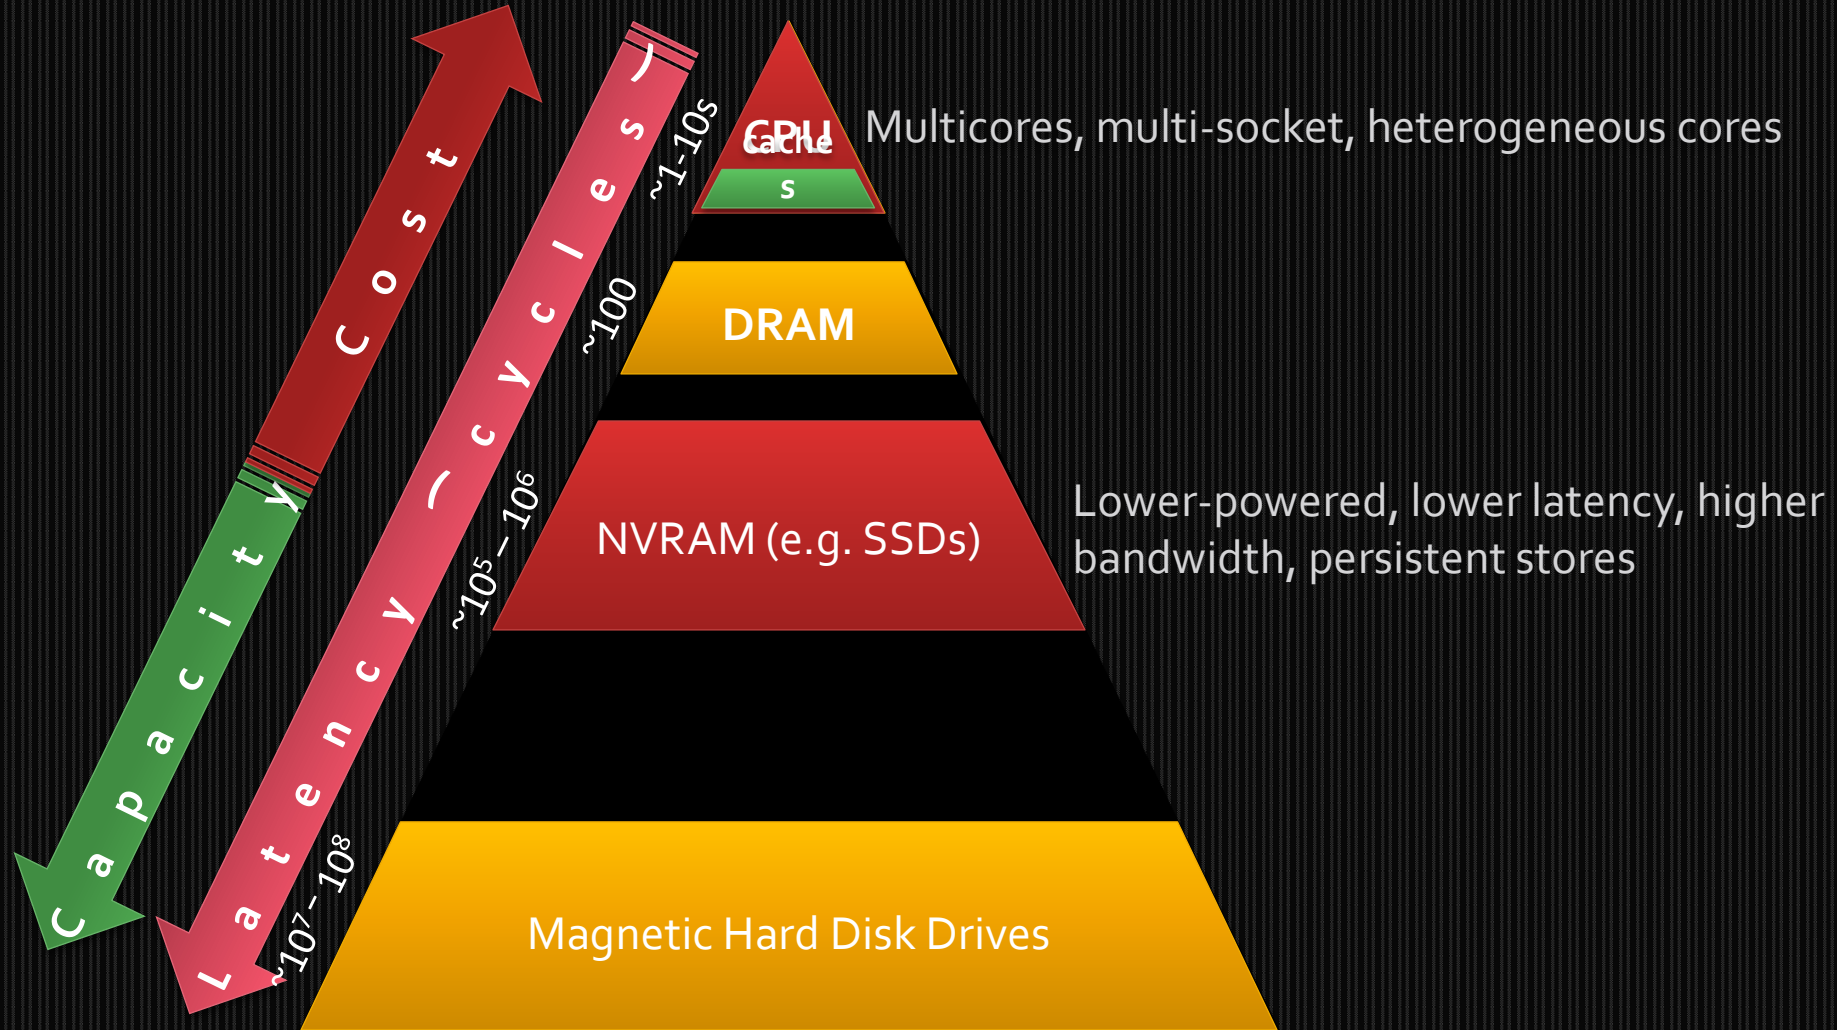
# Framework – Example

```
SELECT SUM(l_discount * l_price) FROM lineitem
WHERE l_shipdate BETWEEN Date AND Date + 1 year
      AND l_discount BETWEEN Discount - 0.01 AND Discount + 0.01
      AND l_quantity < Quantity
```





# Disruptive hardware trends

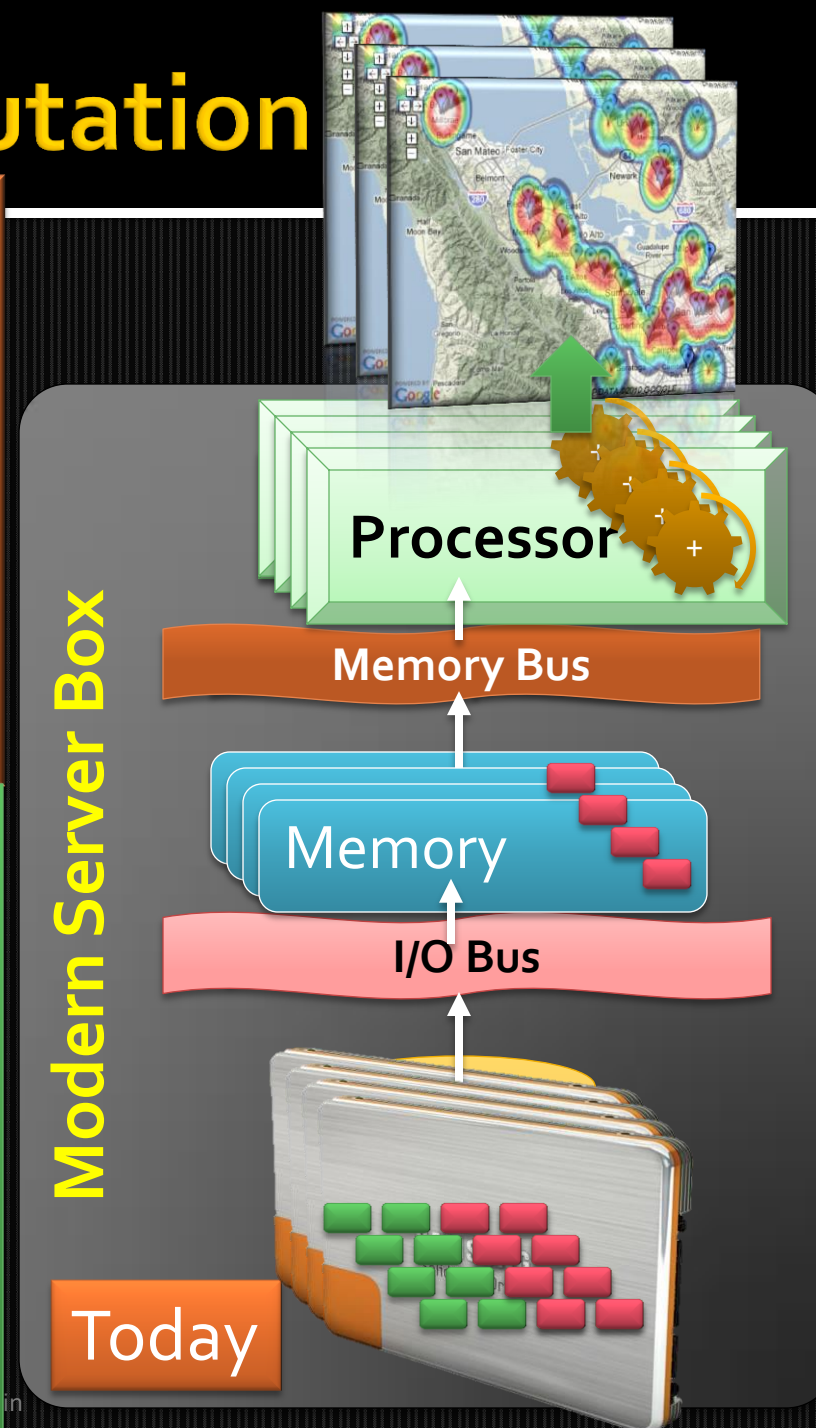


# Data and Computation

## Long Term:

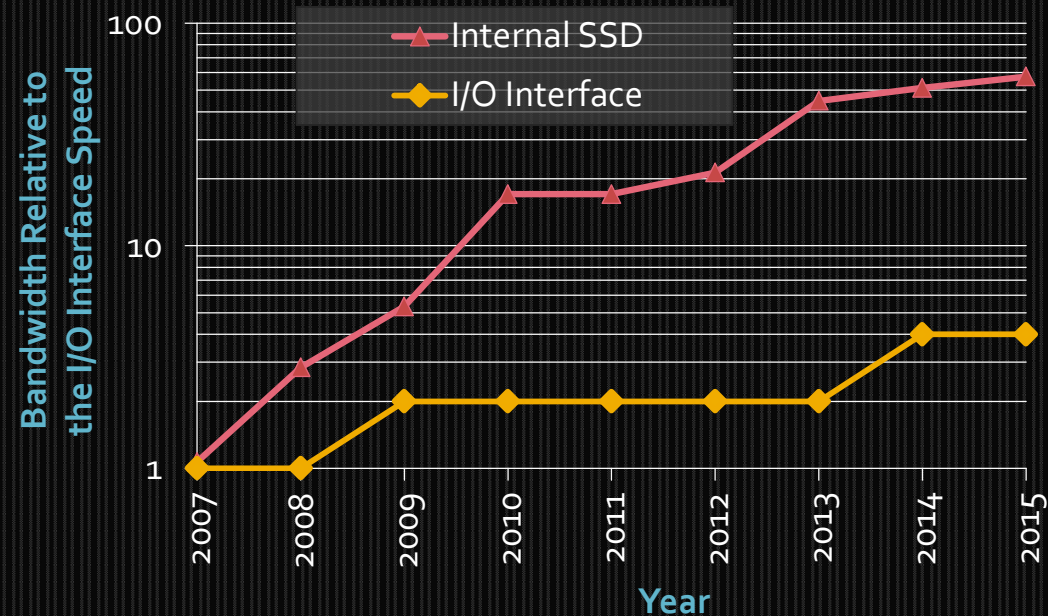
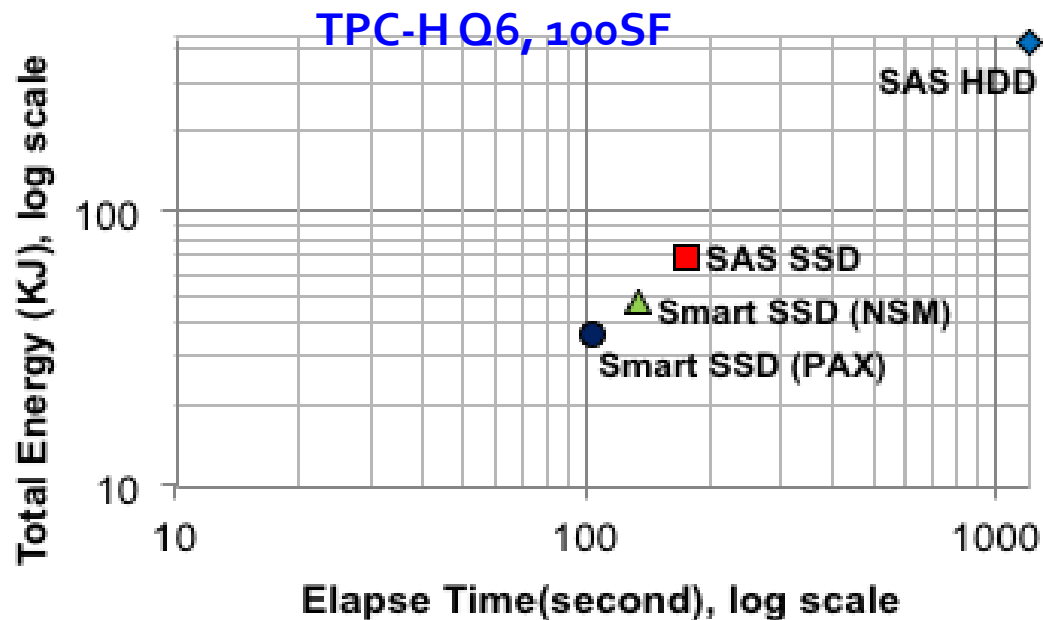
Raw computing and storage costs tends to zero!

The cost is in moving data and powering the circuits/devices





# Example: Flash SSD Architecture



There are similar ways of using hardware creatively, e.g. ASICs or GPUs.

**Basically, need hardware and software synergy!**

# Conclusions

Transformative architectural changes at all levels (CPU, memory subsystem, I/O subsystem) is underway

Need to rethink data processing kernels; e.g. BitWeaving

Need to think of hardware software co-design



# Thanks!



Craig  
Chasseur



Harshad  
Deshmukh



Yinan  
Li



Sangmin  
Shin



Qiang  
Zeng

## The Quickstep Team

*Blog: <http://bigfastdata.blogspot.com>*