

Cost

cost per unit



Manufacturing costs drop as expertise grows, for that process

- better methods
- better equipment
- less waste (time, materials)

yield



Yield = 1 - waste

- #(devices sellable) versus #(devices produced)
- #(devices sellable) versus (cost to produce them)

E.g. DRAM \Rightarrow price = α cost

80% contract sales to large equipment makers (hidden)

20% open market

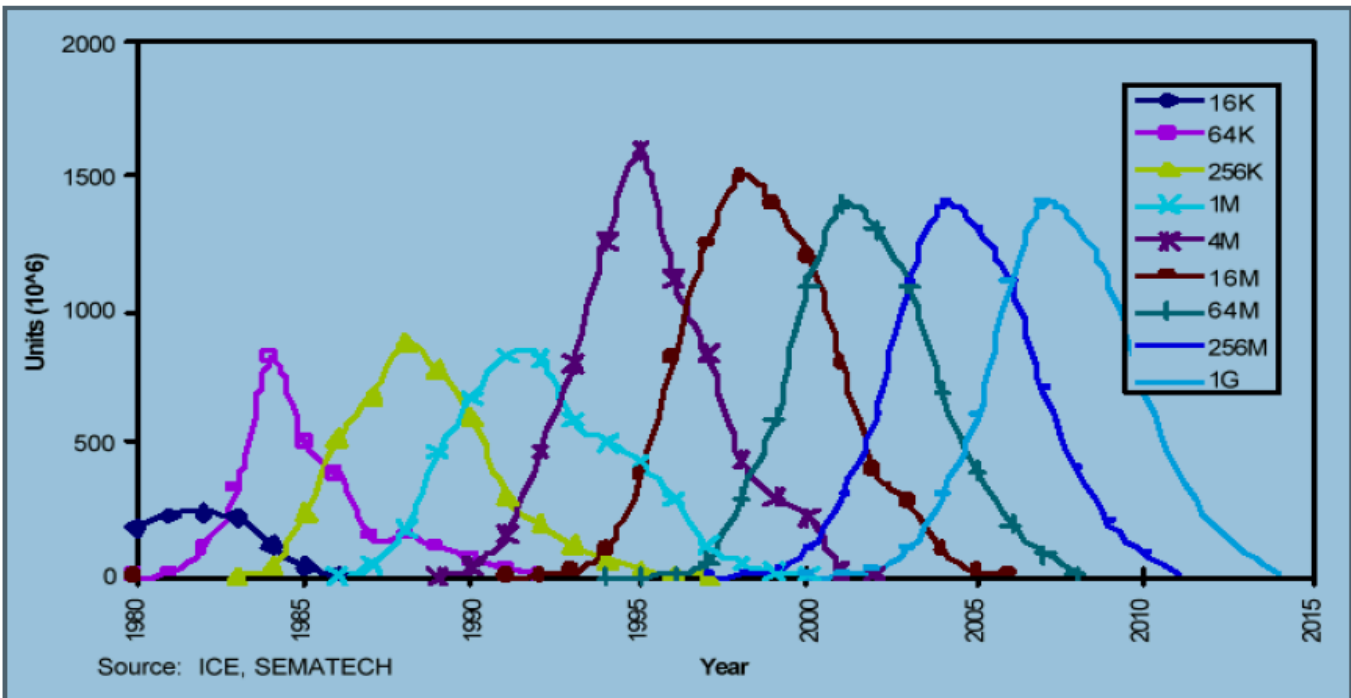
Commodity market

Total Capacity



- lots of vendors
- selling same items

new plant online : \$3B / 3 yr



DRAM Unit Volume by Generation [37]

Invest in largest demand \Rightarrow production cost amortized \Rightarrow larger profit
 hot-new \Rightarrow high price / low volume \Rightarrow old-standard \Rightarrow low price

Costs Drop

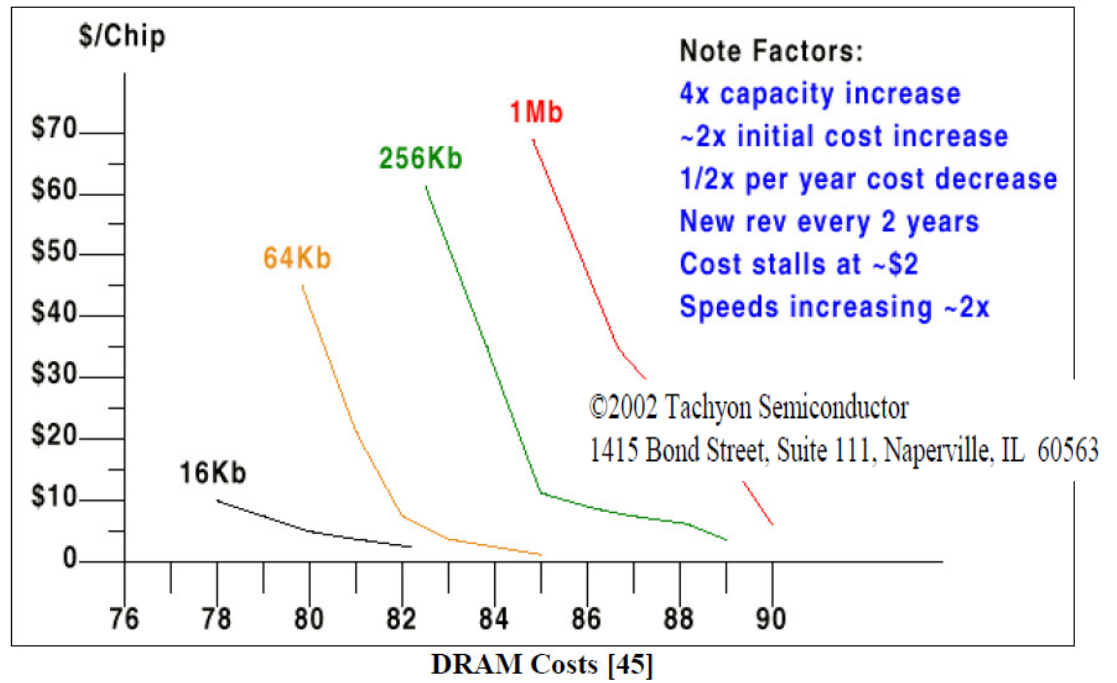
\$50/64

\$65/256

4x capacity
2x speed

Cost per bit
per bit/sec

Drop faster



Changes

Telecom (routers/switches) 20% ↑ ⇒ 50%

latency ↓ vs bandwidth ↑
(PCs)

SRAM

12 ns


\$50/~2 MB

DRAM

40 ns

\$200/GB

Cost



$\frac{1}{2} \times \text{Cost} : 2 \times \text{Volume}$

Volume → supplier competition → lower Cost

⇒ Low-end Market (Price/Performance) ↓

Standardization / Volume =====> market acceptance of innovations

Cloud Pricing

AWS

Combined efficiencies


Description	Type	CU	Original \$ / CU / Hour	Current \$ / CU / Hour	% Reduction	Aug 2006	Oct 2007	May 2008	Oct 2009	Feb 2010	July 2010	Sep 2010	Nov 2010	Nov 2011
Small - "the original"	m1.small	1	\$0.10	\$0.085	15%	\$0.10			\$0.09					
Large	m1.large	4	\$0.10	\$0.085	15%		\$0.40		\$0.34					
Extra Large	m1.xlarge	8	\$0.10	\$0.085	15%		\$0.80		\$0.68					
High-CPU Medium	c1.medium	5	\$0.04	\$0.03	15%			\$0.20	\$0.17					
High-CPU Extra Large	c1.xlarge	20	\$0.04	\$0.03	15%			\$0.80	\$0.68					
High-Memory Double Extra Large	m2.2xlarge	13	\$0.09	0.077	17%				\$1.20			\$1.00		
High-Memory Quad Extra Large	m2.4xlarge	26	\$0.09	0.077	17%				\$2.40			\$2.00		
High Memory Extra Large	m2.xlarge	6.5	\$0.12	0.077	33%					\$0.75				
Cluster Compute	cc1.4xlarge	33.5	\$0.05	\$0.04	19%						\$1.60			
Cluster Compute Eight Extra Large	cc2.8xlarge	88	\$0.03	\$0.03	0%									\$2.40
Micro	t1.micro	0.9	\$0.02	\$0.02	0%							\$0.02		
Cluster GPU Instance	cg1.4xlarge	33.5	\$0.06	\$0.06	0%									\$2.10

Price Reductions

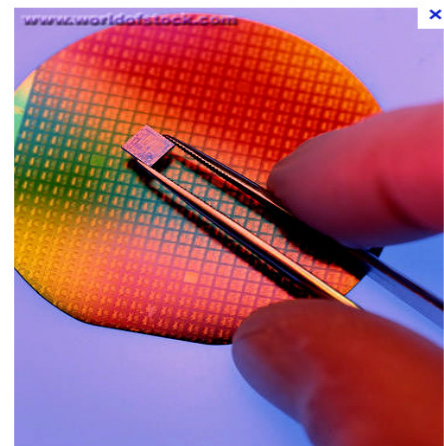
CPUs, Chips, SOC

Si ingots slicing → wafers



masking, etching, doping




dicing →




Circuit testing





Pad Bonding




Pin Packaging




↓ encasing



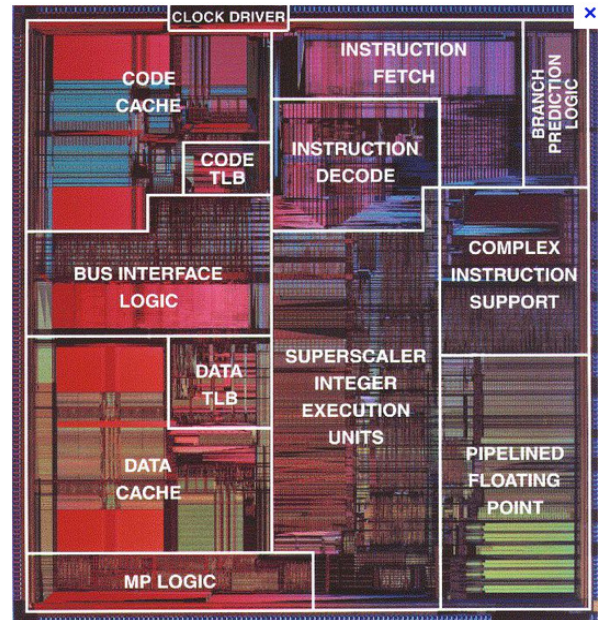
← board printing mounting



what's inside?



i7



P5

$$\text{Cost} = \frac{C_{\text{die}} + C_{\text{Test}_1} + C_{\text{package}} + C_{\text{Test}_2}}{\#(\text{sellable units})}$$


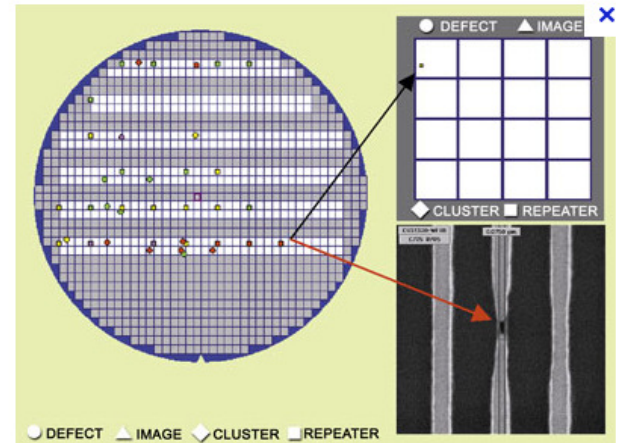
$$C_{\text{die}} = C_{\text{wafer}} / (\# \text{dies})(\text{yield})$$

$$C_{\text{wafer}} \approx \$5000$$

$$\begin{aligned} \#(\text{dies}) &= \left(\frac{\text{Area}_{\text{wafer}}}{\text{Area}_{\text{die}}} \right) - \left(\frac{\text{Circumference}_{\text{wafer}}}{\text{Diagonal}_{\text{die}}} \right) \\ &= \frac{\pi r^2}{A_{\text{die}}} - \frac{2\pi r}{\sqrt{2} \sqrt{A_{\text{die}}}} \end{aligned}$$

$$\text{yield} \approx \frac{\#(\text{good wafers}) / \#(\text{wafers})}{\left[1 + \frac{\#(\text{defects})}{\text{cm}^2} (A_{\text{die}} \text{ cm}^2) \right]^N}$$

$$\text{Curve fitting} \Rightarrow N \in [11.5, 15.5]$$



$\frac{\#(\text{defects})}{\text{cm}^2} \approx 0.04$
 ↪ function of time + volume

300 mm Wafer

$A_{\text{die}} = 2.25 \text{ cm}^2 \Rightarrow 109 \text{ good}$
 $A_{\text{die}} = 1 \text{ cm}^2 \Rightarrow 424 \text{ good}$

P5 Sandy Bridge

20 x 10 mm
 2 cm²
 \$50

embedded CPU, 32b

0.1 cm²
 \$13

printer controller

0.04 cm²
 \$0.1

Die size = $\#(\text{Transistors}) + \#(\text{pins})$ ↑
 +
 Volume ↓
 +
 Customization ↑


} cost

Amortized Costs

Mask = \$1M

⇒ reconfigurable gate arrays

Redundancy, e.g.



Select 3 banks
 1 spare


Warehouse - Scale Costs

Cost_{computing} = Cost_{equipment} / unit Time + Cost_{power} + Cost_{structure} / Time + Cost_# / Time + Cost_{repair}

(60%) (40%)
 (Computers + networks) + (other)
 3 yr

failures, systemic, hierarchic

Service Accomplishment = it works
 Service Interruption = it doesn't



Reliability = time to failure

Mean Time To Failure (MTTF) = Avg (Reliability) = Avg (does)
 Mean Time To Repair (MTTR) = Avg (doesn't)

$$\text{Failure Rate} = 1/\text{MTTF}$$

$$\text{Mean Time Between Failures (MTBF)} = \text{MTTF} + \text{MTTR}$$

$$\text{Availability} = \% \text{ time working} = \frac{\text{MTTF}}{\text{MTBF}}$$

Multi-level
Recovery

- error correction
- redo
- redundancy

E.G., Disk sub-system

10 disks, $\text{MTTF} = 1 \text{ Mhr}$

1 ATA controller, $\text{MTTF} = \frac{1}{2} \text{ Mhr}$

1 power supply, $\text{MTTF} = \frac{1}{5} \text{ Mhr}$

1 fan, $\text{MTTF} = \frac{1}{5} \text{ Mhr}$

1 ATA cable, $\text{MTTF} = 1 \text{ Mhr}$

ASSUME: exponentially dist. failures

- independent of T

- No dependency between components

$$E(\text{* failures in } \Delta T) = \mu \Delta T$$

↖ avg failure rate

$$1 \text{ failure} = \mu \overline{\Delta T}_1$$

$$\text{MTTF} = \overline{\Delta T}_1 = 1/\mu$$

$$\mu = \sum_i \mu_i = \sum_i (1/\text{MTTF}_i)$$

$$= (10 + 2 + 5 + 5 + 1) / \text{Mhr} = 23 / \text{Mhr}$$

$$\text{MTTF} = \frac{10^6 \text{ hr}}{23} \approx \left(\frac{10^2}{25}\right) 10^4 = 40,000 \text{ hr}$$

$$(365 \text{ d})(24 \text{ hr/d}) \approx$$

$$(400)(100/4) = \frac{1}{4}(40,000)$$

$$\Rightarrow 4 \text{ yr.}$$

E.g., 2 power supplies?


Let's add a redundant power supply: if one fails we fix it while the other keeps working. How long before the system fails?

That is, how long before we experience having the back-up fail before we can finish the repair?

Poisson

Bernoulli w/ $p = \lambda dt$

flip coin w/ Prob = $p = \lambda dt$



$$P(m, N, p) = \binom{N}{m} p^m q^{N-m}$$

$$\Rightarrow \frac{e^{-pN} (pN)^m}{m!}$$

let $pN = \lambda$

$$P(m, \lambda) = \frac{e^{-\lambda t} (\lambda t)^m}{m!}, \quad \lambda = \text{rate}$$

$$P(X_n = x) = 1 - e^{-\lambda x}$$

$$E(X) = 1/\lambda$$

1.1 Point Processes

Definition 1.1 A simple point process $\psi = \{t_n : n \geq 1\}$ is a sequence of strictly increasing points

$$0 < t_1 < t_2 < \dots, \quad (1)$$

with $t_n \rightarrow \infty$ as $n \rightarrow \infty$. With $N(0) \stackrel{\text{def}}{=} 0$ we let $N(t)$ denote the number of points that fall in the interval $(0, t]$; $N(t) = \max\{n : t_n \leq t\}$. $\{N(t) : t \geq 0\}$ is called the counting process for ψ . If the t_n are random variables then ψ is called a random point process. We sometimes allow a point t_0 at the origin and define $t_0 \stackrel{\text{def}}{=} 0$. $X_n = t_n - t_{n-1}$, $n \geq 1$, is called the n^{th} interarrival time.

We view t as time and view t_n as the n^{th} arrival time. The word *simple* refers to the fact that we are not allowing more than one arrival to occur at the same time (as is stated precisely in (1)). In many applications there is a "system" to which customers are arriving over time

1.2 Renewal process

A random point process $\psi = \{t_n\}$ for which the interarrival times $\{X_n\}$ form an i.i.d. sequence is called a *renewal process*. t_n is then called the n^{th} *renewal epoch* and $F(x) = P(X \leq x)$ denotes the common interarrival time distribution. The *rate* of the renewal process is defined as $\lambda \stackrel{\text{def}}{=} 1/E(X)$.

Poisson distribution can be derived by taking the appropriate limit of the binomial distribution

$$P(m, N, p) = \frac{N!}{m!(N-m)!} p^m q^{N-m}$$

$$\frac{N!}{(N-m)!} = \frac{N(N-1)\dots(N-m+1)(N-m)!}{(N-m)!} = N^m$$

$$q^{N-m} = (1-p)^{N-m} = 1 - p(N-m) + \frac{p^2(N-m)(N-m-1)}{2!} + \dots \approx 1 - pN + \frac{(pN)^2}{2!} + \dots \approx e^{-pN}$$

$$P(m, N, p) = \frac{N^m}{m!} p^m e^{-pN}$$

$$P(m, \mu) = \frac{e^{-\mu} \mu^m}{m!}$$

1.3 Poisson process

Definition 1.2 A Poisson process at rate λ is a renewal point process in which the interarrival time distribution is exponential with rate λ : interarrival times $\{X_n : n \geq 1\}$ are i.i.d. with common distribution $F(x) = P(X \leq x) = 1 - e^{-\lambda x}$, $x \geq 0$; $E(X) = 1/\lambda$.

$$P(N(t) = k) = e^{-\lambda t} \frac{(\lambda t)^k}{k!}, \quad k \geq 0.$$

*k events
in time t*


Simulating a Poisson process at rate λ up to time T :

1. $t = 0, N = 0$
2. Generate U .
3. $t = t + [-(1/\lambda) \ln(U)]$. If $t > T$, then stop.
4. Set $N = N + 1$ and set $t_N = t$.
5. Go back to 2.

Suppose that for a Poisson process at rate λ , we condition on the event $\{N(t) = 1\}$, the event that exactly one arrival occurred during $(0, t]$. We might conjecture that under such conditioning, t_1 should be uniformly distributed over $(0, t)$. To see that this is in fact so, choose $s \in (0, t)$. Then

$$\begin{aligned} P(t_1 \leq s | N(t) = 1) &= \frac{P(t_1 \leq s, N(t) = 1)}{P(N(t) = 1)} \\ &= \frac{P(N(s) = 1, N(t) - N(s) = 0)}{P(N(t) = 1)} \\ &= \frac{e^{-\lambda s} \lambda s e^{-\lambda(t-s)}}{e^{-\lambda t} \lambda t} \\ &= \frac{s}{t} \end{aligned}$$

Independent, Random faults



individual

VS



biased coin w/

$$\text{Prob}(\text{working}) = p$$

$$\text{Prob}(\text{failed}) = (1-p)$$

coin tosses, independent, constant p

$$\Rightarrow X = \{X_1, X_2, \dots\}$$

$$X_i = \begin{cases} 0, & \text{failure} \\ 1, & \text{working} \end{cases}$$

↑ individual toss outcome

n -trials

$$v \in \{0, 1\}^n$$

$$\text{Prob}(X=v) = p^r (1-p)^{n-r}$$

$$[\text{\#ones}(v) = r]$$

System state $\Rightarrow A(X) \in \{0, 1\}$

$$\Rightarrow \begin{cases} \text{System is working: } A(X) = 1 \\ \text{System has failed: } A(X) = 0 \end{cases}$$

Series System State

$$\text{working}(X) = \begin{cases} 0, & \text{some } X_i = 0 \\ 1, & \text{all } X_i = 1 \end{cases}$$

$$\Rightarrow \begin{cases} \text{Prob}(\text{working}) = p^n \\ \text{Prob}(\text{failed}) = 1 - p^n \end{cases} \quad [\text{\#ones}(X) = n]$$

Parallel System

$$\text{working} = \begin{cases} 0, & \text{all } X_i = 0 \\ 1, & \text{some } X_i = 1 \end{cases}$$

$$\Rightarrow \begin{cases} \text{Prob}(\text{failed}) = (1-p)^n \\ \text{Prob}(\text{working}) = 1 - (1-p)^n \end{cases} \quad [\text{\#zeros}(X) = n]$$

Take n devices, run them for t hrs

k working $X_i = \begin{cases} 1, & \text{device}_i \text{ still working} \\ 0, & \text{device}_i \text{ failed} \end{cases}$

$$k = \sum X_i \quad E(k) = \sum E(X_i) = np$$

$$E(X_i) = 1 \cdot \text{Prob}(X_i=1) + 0 \cdot \text{Prob}(X_i=0) = p$$

Assuming measured $k \approx E(k) \Rightarrow p = k/n$

Assuming t is a unit time interval, let $X = \{X_1, X_2, X_3, \dots, X_n\}$

$X_i = \begin{cases} 1, & \text{device working during period } i \\ 0, & \text{device failed in period } i \end{cases}$

$$\begin{aligned} \text{Prob}(\text{failure at } k) &= \text{Prob}(\text{working for } k \text{ intervals}) \times \text{Prob}(\text{fail in } (k+1) \text{ interval}) \\ &= p^k (1-p) = p^k q \end{aligned}$$

$$E(k) = \sum_{k=1}^{\infty} k p^k q = q \sum_{k=1}^{\infty} k p^k = q p \sum_{k=1}^{\infty} \frac{d}{dp} p^k = q p \frac{d}{dp} \sum_{k=1}^{\infty} p^k = q p \frac{d}{dp} \left(\frac{1}{1-p} \right)$$

$$= q p / (1-p)^2$$

$$= q p / q^2 = p/q$$

$$= \text{MTTF}$$

$$\frac{d}{dp} (1-p)^{-1} = -(1-p)^{-2} \frac{d}{dp} (-p) = (1-p)^{-2}$$

e.g. $p = \frac{1}{2} \Rightarrow \text{MTTF} = (\frac{1}{2}) / (\frac{1}{2}) = 1$
 $q = \frac{1}{2^n} \Rightarrow \text{MTTF} = \frac{1-2^{-n}}{2^{-n}} = 2^n - 1$

2 power supplies

$$\text{Prob}(\text{either fails}) = (1-p^2) = Q$$

$$\begin{aligned} \text{MTTF}_2 &= q/Q = \frac{1-(1-p^2)}{(1-p^2)} = \frac{p^2}{1-p^2} = \frac{p^2}{(1-p)(1+p)} = \frac{p}{q} \frac{p}{(1+p)}, \quad \text{for } p \approx 1 \\ &\approx \frac{p}{q} \left(\frac{1}{2} \right) = \frac{1}{2} \text{MTTF}_1 \end{aligned}$$

Prob(2^{nd} does fail in MTTR)

$$\approx \frac{\text{MTTR}}{\text{MTTF}_1}$$

$$p \approx p \Delta T$$

$$p = \lambda dt$$

$$p = \frac{1}{\text{MTTF}} dt$$

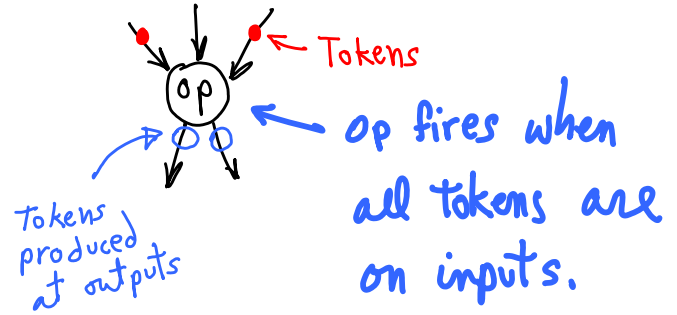
$$p \approx \frac{1}{\text{MTTF}_1} (\text{MTTR})$$

$$P(x) = P_{1-2}(x - \text{MTTR}) P(\text{MTTR})$$

$$E(x) = \frac{1-p}{p} = \frac{1 - \left(\frac{2}{\text{MTTF}} \right) \left(\frac{\text{MTTR}}{\text{MTTF}} \right)}{\left(\frac{2}{\text{MTTF}} \right) \left(\frac{\text{MTTR}}{\text{MTTF}} \right)} \approx \frac{(\text{MTTF})^2}{2 \text{MTTR}}$$

Tagged Dataflow Arch.

Dataflow Graph



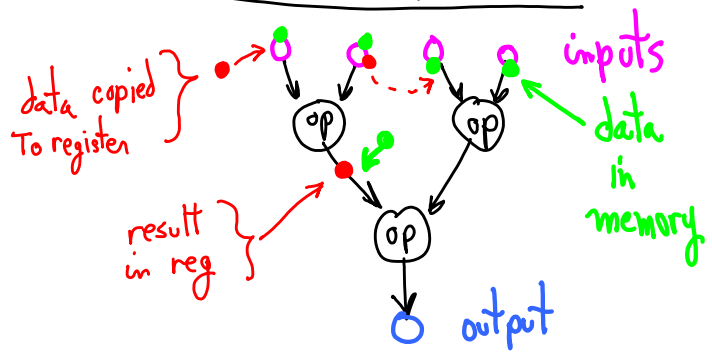
--- Data is copied from **memory** to **registers**:
red pebble on **green**

--- Operation done on register contents

--- **Result** written to **register**:
red pebble on arc

--- **Register reuse**: copy **registers** to **memory**:
green pebble on arc

Data Dependency Graph



Some portion of Data Graph pebbled red or green during computation.
 outputs pebbled green at completion

ID programming (LISP-like)

```

Def vsum A B =
  { C = array (1, n) :
    { for i 1 to n
      Ci ← Ai + Bi
    }
  }
  } In C
  
```

Functional Programming

- Any step can be taken whenever ready
- C allocated in parallel with Loop unrolling
- Function returns before C is filled
- Each C_i waits for A_i and B_i
- Semantics are write once
- Synchronization is implicit
- individual tasks distributed as HW available

Def ip A B

{ s = 0

In {

For i 1 to n

Next s = s + A_i x B_i

finally s


}

Loop unrolling

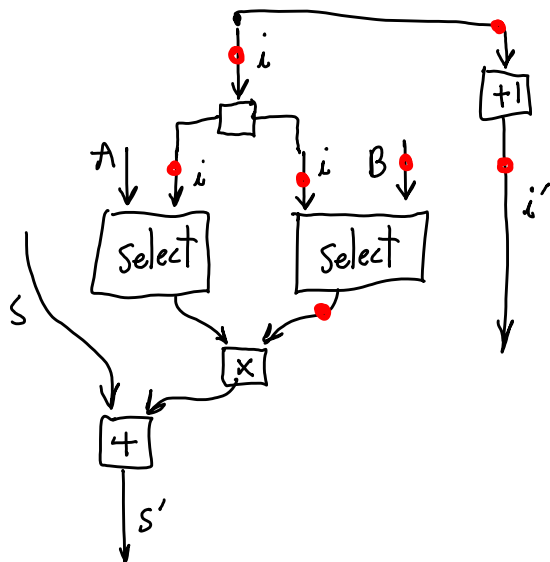
$$s_{i+1} = s_i + A_i + B_i$$

Synchronization?

Write-Once semantics




Dataflow Graph



Tokens "fine" operators


constants



I-structure

mem v A

P	1	2	3	4
A	?	?	?	?
W				



Read token

WRITE token

• ≡ <addr_{data}, READ, addr_{inst}>

• ≡ <addr_{data}, WRITE, value>

→ write data:

Split-phase operation: asynchronous R/W requests sent, asynchronous result returns value.

Good at hiding memory latency?

Good at tolerating high synchronization costs?

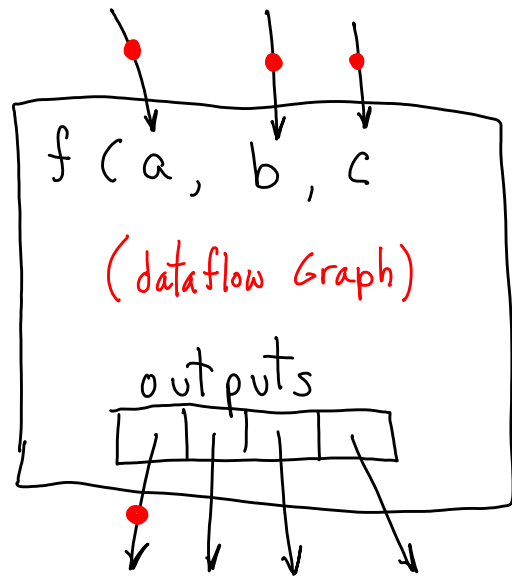
if v = Present → error

if v = Waiting, wakeup list

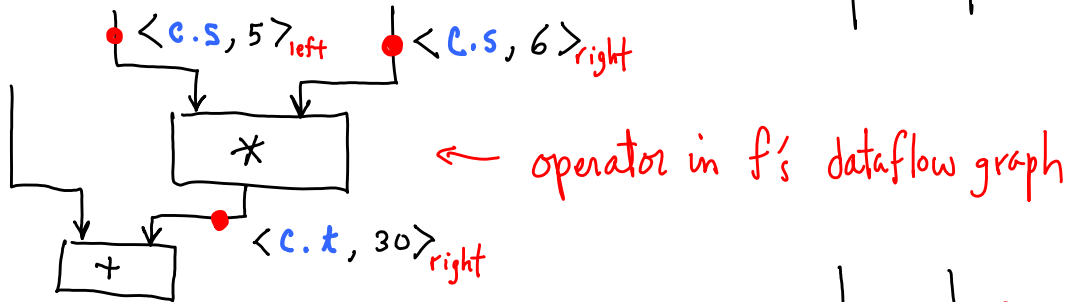
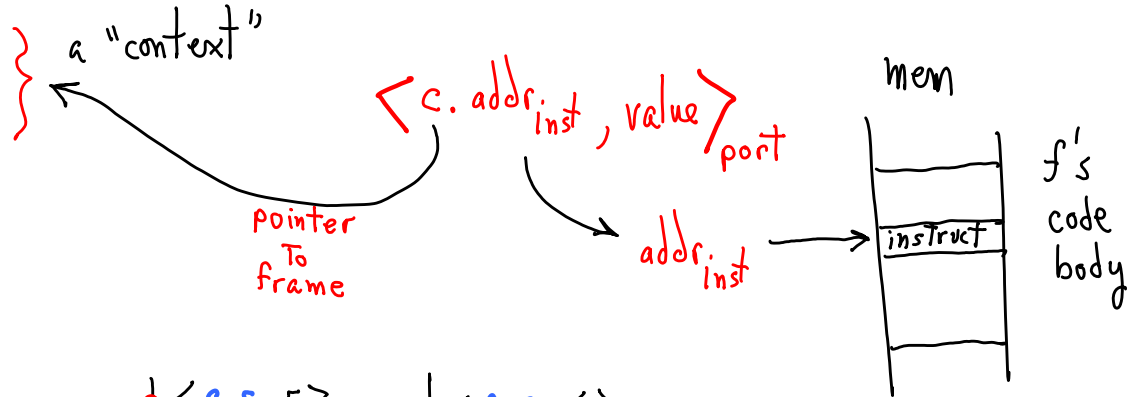

set v = Present

Functions / context


recursive calls	multiples
multiple parallel active calls	f applied in parallel, e.g. loop unrolling



STACK





Firing Rule
Operator fires when tags on inputs match



\Rightarrow multiple tokens w/ different tags may be present


Calls

1. f sets return




Note
"s" is current instruction,
"t" is next instruction

2. Get Callee's context



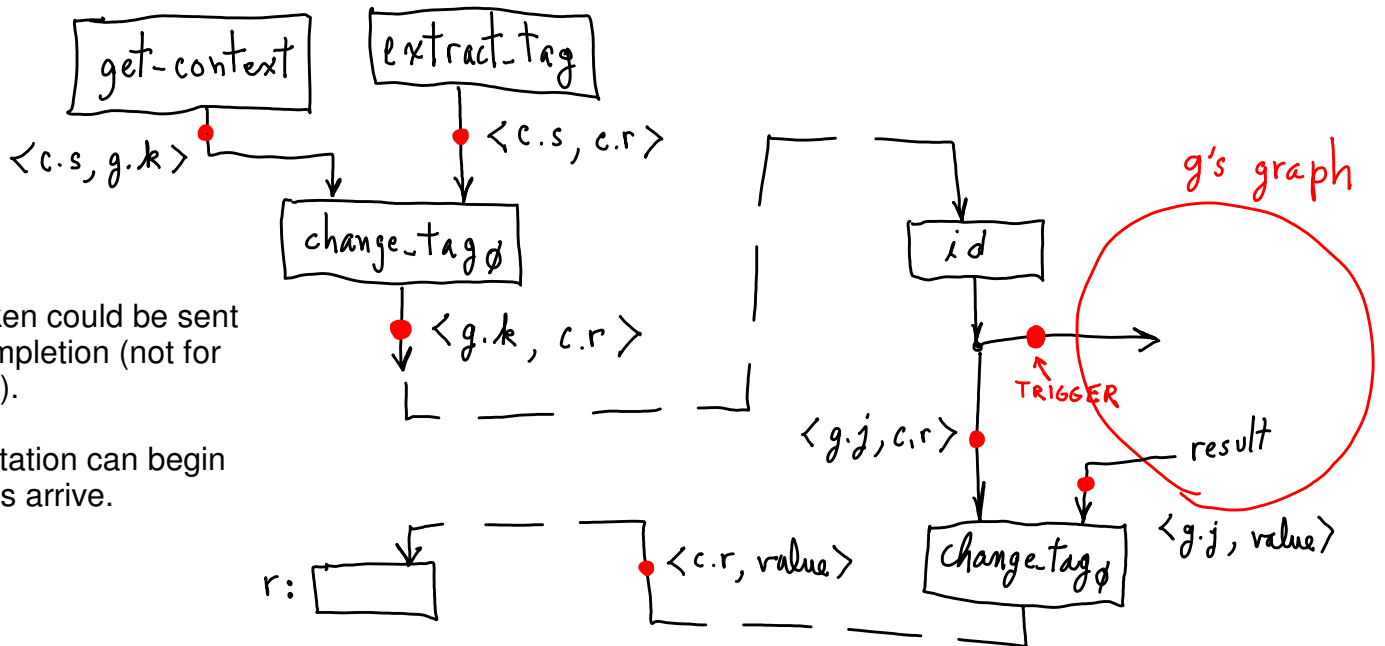
k is entry point for g. context "g" is for this invocation of g().

3. Pass args



Note
Destination of arc depends on left input value.

3. Invoke function

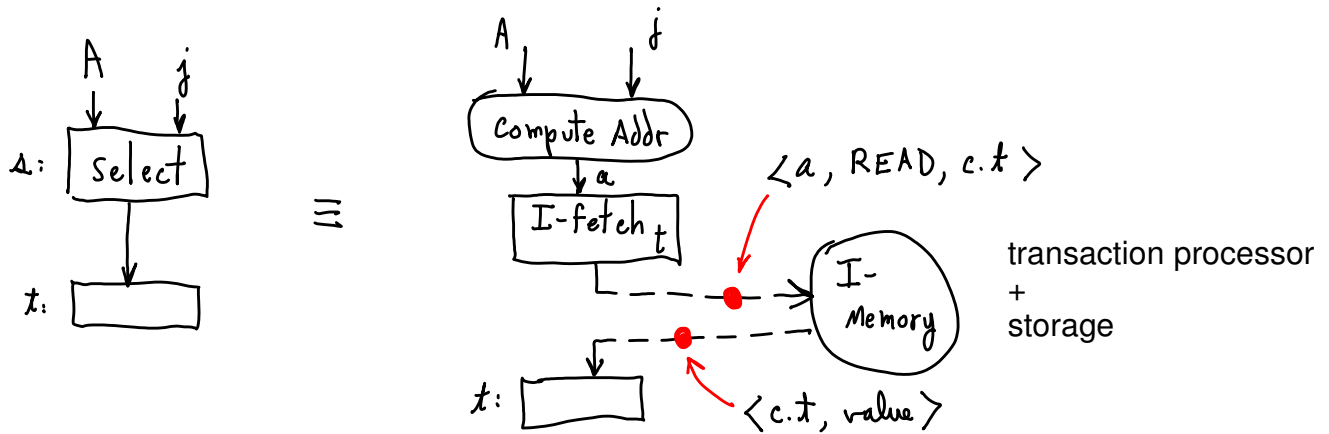


Note:
Return token could be sent before completion (not for g() though).

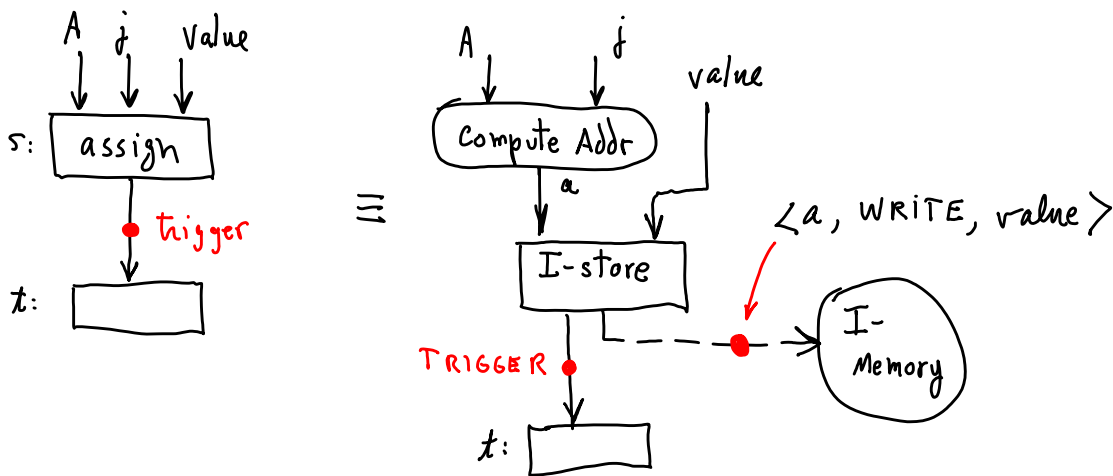
g's computation can begin before args arrive.

I-Structure R/W/allocate

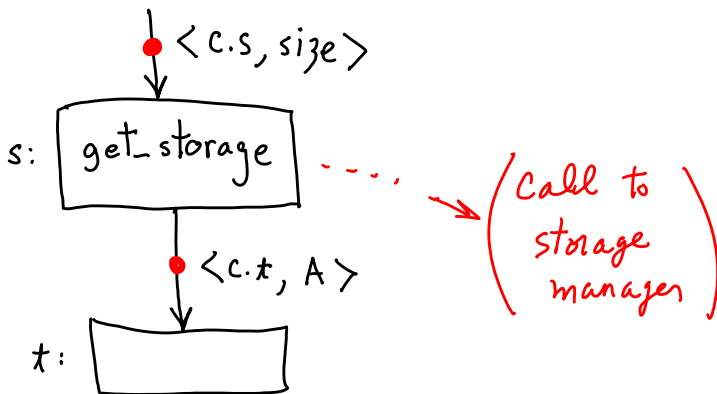
Read



Write



Allocate



Can have other R/W protocols (not Write-Once only)

Can have active memory: operations in memory unit.


De-allocate

Cannot de-allocate while tokens exist in function call for context i.

==> Well-behaved graphs

1. initially, no tokens
2. given one token on every input, one token produced per output
3. after all output tokens produced, graph is empty.

Allocate by code blocks



by function body (same)

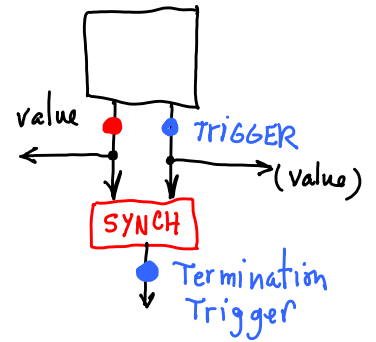
Call = allocate \Rightarrow code to processor
 registers
 context

send trigger to callee (+ args)

Callee sends termination Trigger (+ return value)

Deallocate

conditional output




Functional APPLY

Def vsun A B
 $\{$
 \dots
 $\}$
 In C

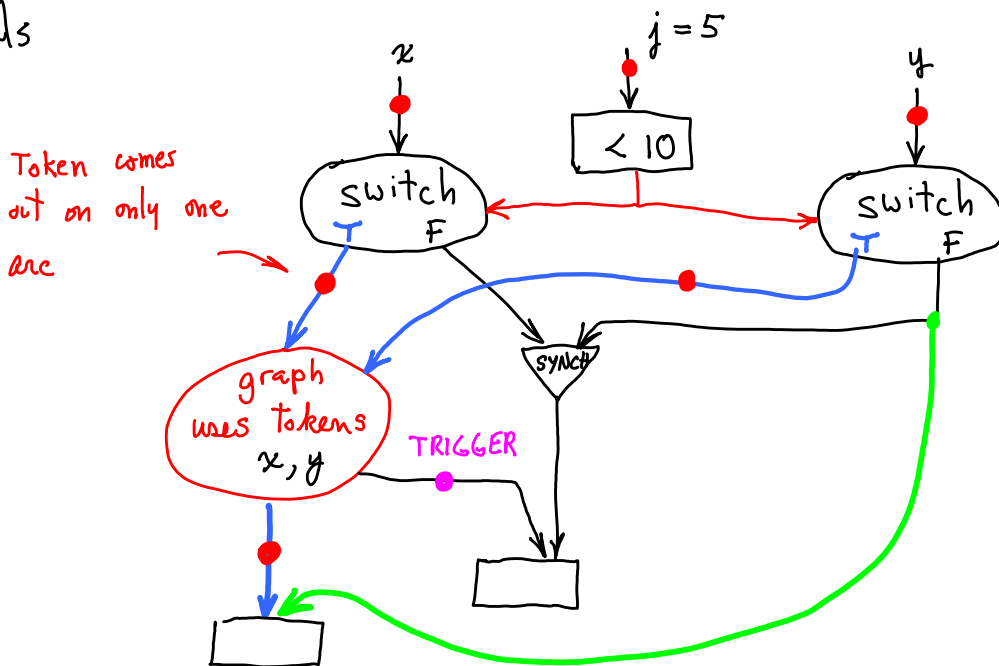
\Rightarrow move_point = vsun A
 move_point B

evaluation occurs when all
 args present



Conditionals

if ($j < 10$)
 $x + y$
 else
 y



Token comes
 out on only one
 arc

Well behaved

Trigger output tokens
 have same address

Value output tokens
 have same address

Triggers only when all
 inputs have arrived
 and output token is
 ready

