Post PC Computing Course – lecture 3

Assigned readings on design approaches (in the Math Library):

- Don Norman's manifesto
- Rob Haitani on how the Palm Pilot was kept simple
  - (optional: Jeff Hawkins interview on Palm)

Proposals received so far:

- Ami Serri-Menkes
- Amir Gonnen, Tamar Bar-El (amigon)
- Boaz Ben-Yaacov, Rafi Chickvashvili (bryan, rafic)
- Claudia Goldman(clag)
- Shay Shemer(shay1)
- Nir Privman (weird1)
- Noa Bar Yosef, Alina Burach (baryosef, alinab)

Today let's complete the list, narrow the selection, and give the projects catchy names…

Semiconductor technology forecasting

- Why can it be predicted?
  - Learning curves
    - Two examples
  - Mead-Conway scaling ('70s and '80s) (actually Dennard)
    - This is the usual reason given, but it's wrong
    - Power, in particular, should have remained constant
      - but is exploding.
  - Industry competition with communication
    - Review SIA roadmap in detail
      - Dimensions of devices and wires
      - Logic speeds
      - Where are the older technologies used?

DRAM Learning Curve

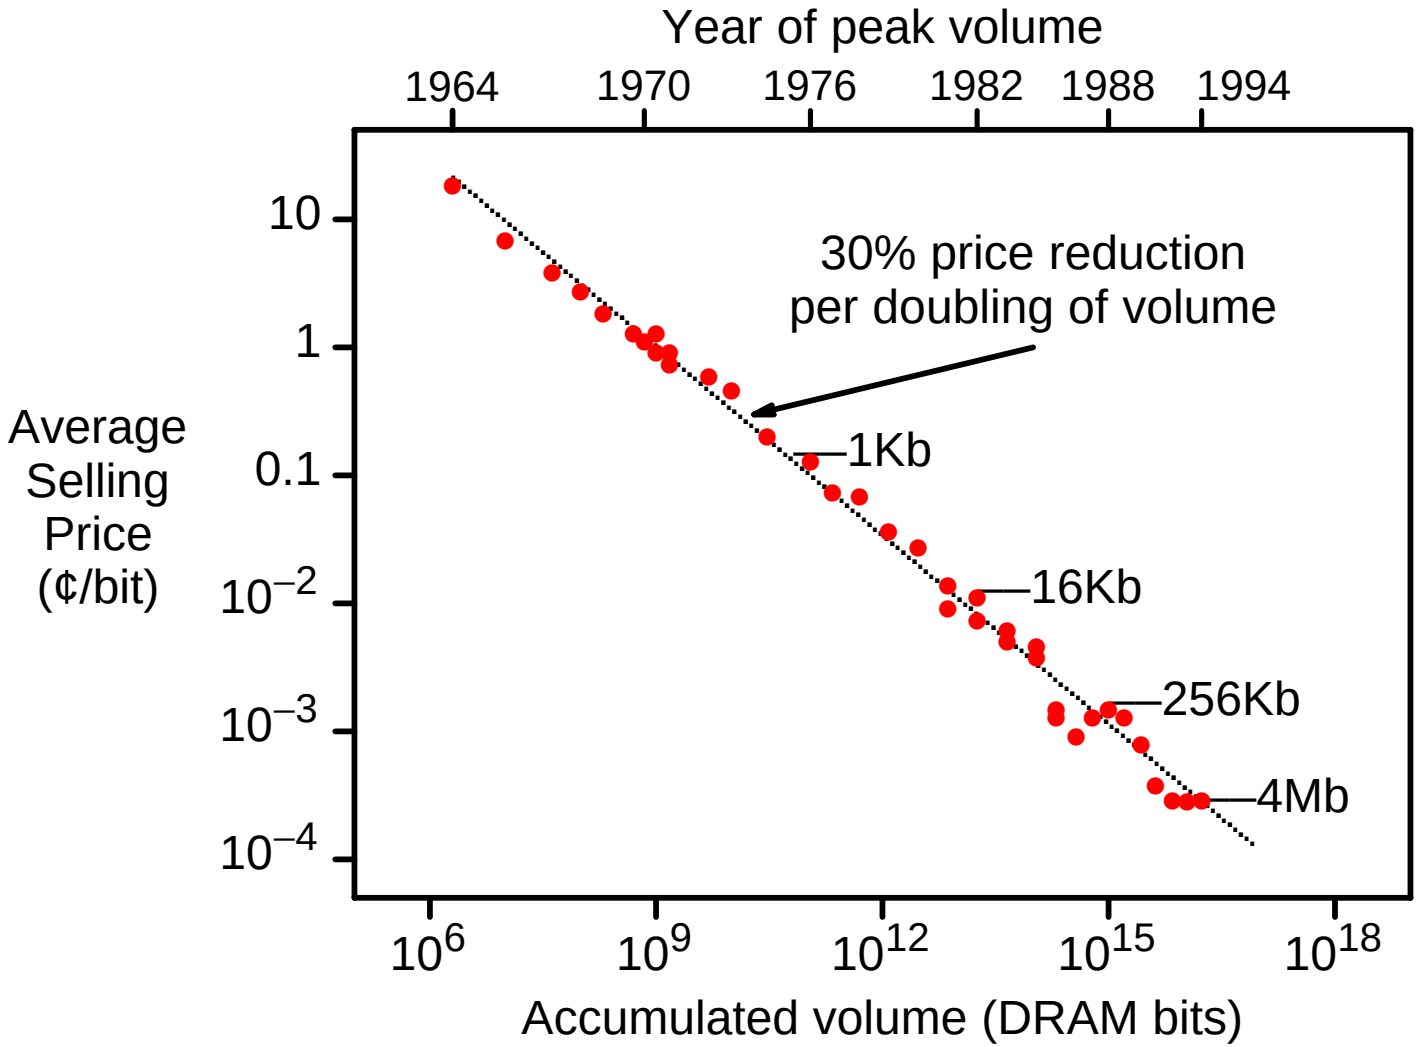

New bits in 1994 ≅ 55% cumulative bits 1970-1993

AT&T Bell Laboratories

/spin.3d/colin/industry/dram_cycle/noyce.vg Jul 27 15:28

Wholesale, loose chips: SIMM streetprice is ~1.25× higher

Broiler Chicken Sales

Learning curve: 1934 - 1975

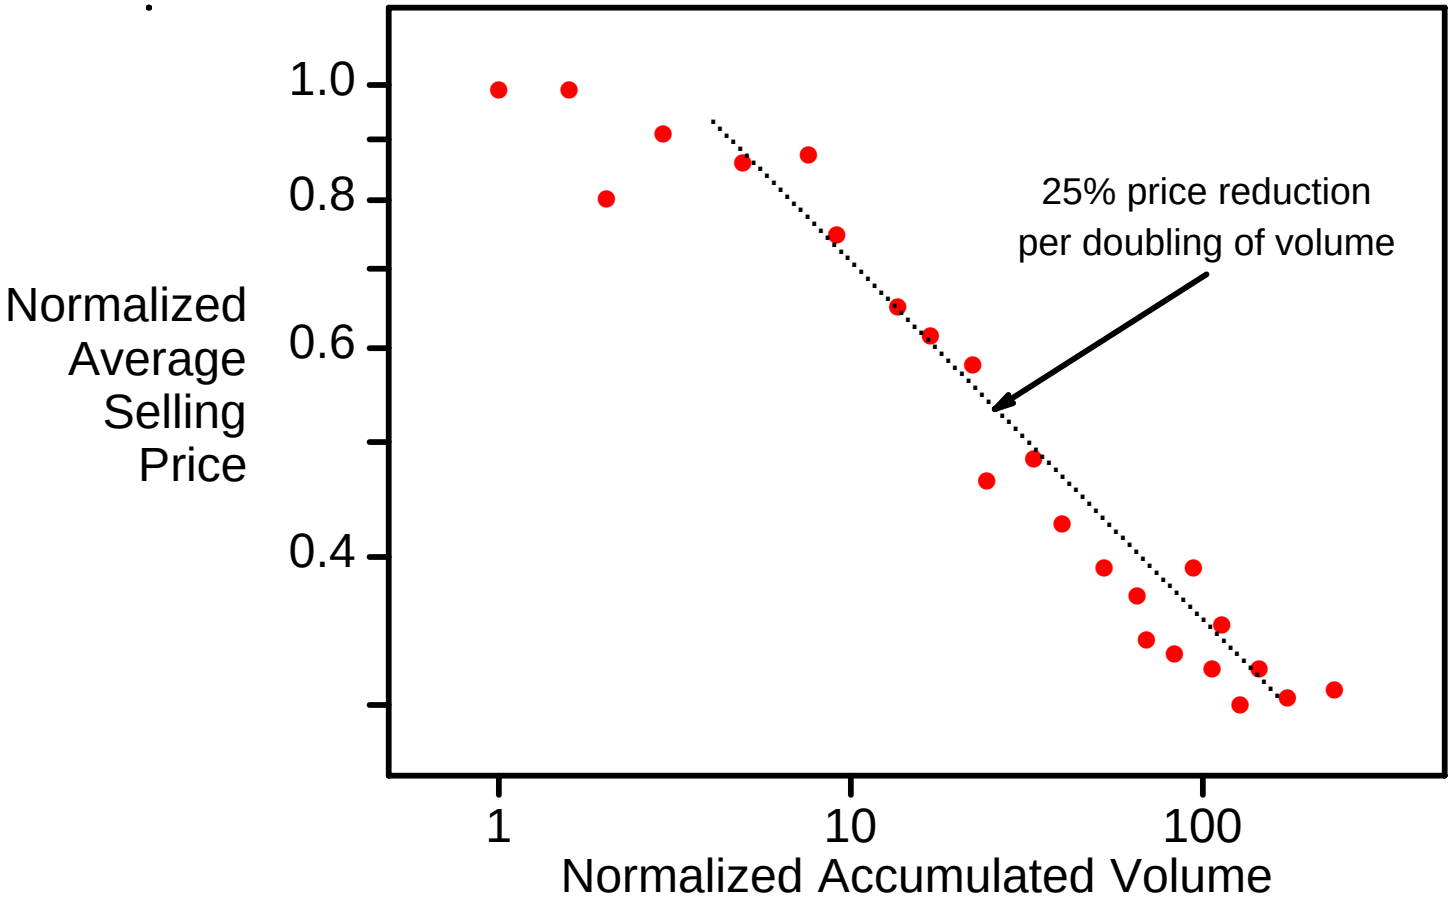

Source: Braxton Assoc., p.41 "Trading places", Clyde Prestowicz, Jr.

AT&T Bell Laboratories

/spin.3d/colin/industry/dram_cycle/chicken.norm.vg Sep 23 13:59

Semiconductor Technology

Power doesn't stay constant under scaling:

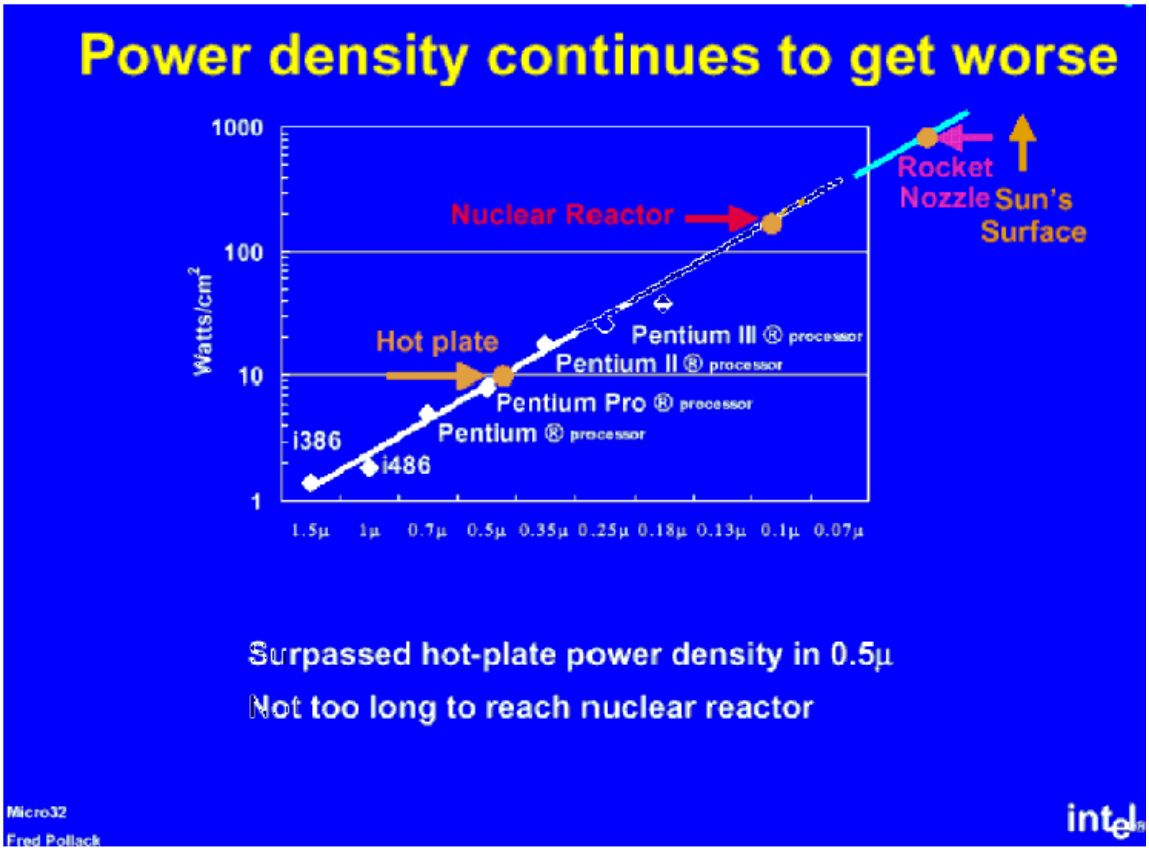

next review details of the SIA/EECA/JEITA/KSIA/TSIA industry roadmap.

Table 2 CMOS Scaling Guidelines for the Next 10 Years

	Late 1980's	1992	1995	1998	2001	2004
Supply Voltage (V)						
High Performance	5	5/3.3	3.3/2.5	2.5/1.8	1.5	1.2
Low Power	—	3.3/2.5	2.5/1.5	1.5/1.2	1.0	1.0
Lithography Resolution (μm)						
General	1.25	0.8	0.5	0.35	0.25	0.18
Gate Level for Short L	—	0.6	0.35	0.25	0.18	0.13
Channel Length (μm)	0.9	0.6/0.45	0.35/0.25	0.2/0.15	0.1	0.07
Gate Insulator Thickness (nm)	23	15/12	9/7	6/5	3.5	2.5
Relative Density	1.0	2.5	6.3	12.8	25	48
Relative Speed						
High Performance	1.0	1.4/2.0	2.7/3.4	4.2/5.1	7.2	9.6
Low Power	—	1.0/1.6	2.0/2.4	3.2/3.5	4.5	7.2
Relative Power/Function						
High Performance	1.0	0.9/0.55	0.47/0.34	0.29/0.18	0.12	0.077
Low Power	—	0.27/0.25	0.20/0.09	0.08/0.056	0.036	0.041
Relative Power/Unit Area						
High Performance	1.0	2.25/1.38	3.0/2.1	3.7/2.34	3.12	3.70
Low Power	—	0.7/0.63	1.25/0.6	1.02/0.72	0.90	1.97

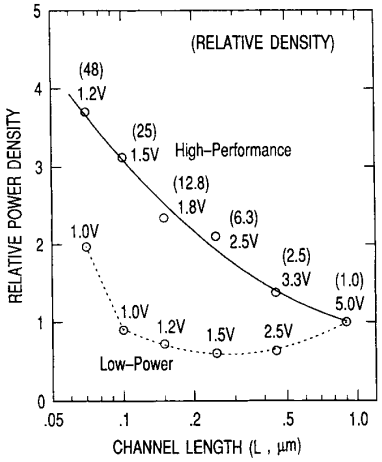

Fig. 7. Relative active power density in CMOS scaled as in Table 2. Relative density is in parentheses.

reciprocal of the square of the relative lithography ground rules for each generation (see Table 2). The upper curve corresponds to the performance driven scenario. In this case, we can see that even with reduced power-supply voltage, the power density increases significantly due to a large increase in the number of devices per unit area. However, one can choose the low power scenario, where the supply voltage is reduced more aggressively, at the expense of the performance (lower curve in Fig. 7). This scenario will be quite attractive for applications where low power and improved power-delay product is the highest priority. As an example, if we compare the 1.5 V operation versus the 2.5 V in 0.25 μm CMOS, there will be over 3.5$\times$ reduction in power (Fig. 7) with only about 30% performance degradation (Fig. 6, $V_t \simeq 0.3$ V).

In the low-power scenario, the power density of the scaled technologies goes down relative to the 1 μm CMOS technology until we reach 1.5 V (Fig. 7). This is due to the fact that in that regime the electric field does not go up substantially and is quite close to the constant electric field scaling as shown in Fig. 4, while the interconnection dimensions (which determine the density) are being scaled down less than the channel lengths. However, beyond the $L = 0.25\,\mu$m at 1.5 V point, the relative power density rises (Fig. 7) as we depart from the CE scaling (Fig. 4) because the lower limit on V_t imposes a lower limit of about 1 V for the power-supply voltage without significantly impacting speed.

A key concern in the low-power scenario is the availability of the complete chip set to make up the systems at reduced supply voltage. However, most of the problems can be overcome by various techniques to mix and match different supply voltages on the board or on the chip. Also, the need for a total low power system solution at low cost should drive the semiconductor industry to even a faster pace for offering various memory and ASIC products at reduced supply voltages.

Another key concern is the susceptibility to soft errors due to alpha particles, which is expected to increase due to reduced voltage and capacitance. This could require improved structures such as SOI to reduce the volume of junction area exposed to alpha particle hits.

D. CMOS Scaling Guideline for the Next Ten Years

The above scaling optimizations and limitations are summarized in Table 2 as a guideline for CMOS scaling over the next ten years for logic applications such as microprocessors. The key design parameters are the power-supply voltage, the lithography resolution, the channel length, and the gate insulator thickness. The more aggressively scaled device and power-supply voltages represent the high-performance and low-power scenarios discussed above. In the columns through the 90's less aggressively scaled device design points are included, representing lower cost options with correspondingly lower speed and higher power dissipation. It is assumed that the wiring pitch scales with the general lithography resolution and the relative density is the reciprocal of the square of the relative