

Improved Cost Model Puts Pentium at \$180

MPR Cost Model Improved for Small Chips, Large Chips

by Linley Gwennap

When the MPR Cost Model was first introduced (*see 071004.PDF*), it estimated the cost of building 0.8-micron Pentium chips to be \$483 each. Yet today, Intel is selling these chips at a list price of \$418. Is the company taking a loss just to move Pentium into the market faster? Hardly. Our current estimated manufacturing cost for the 0.8-micron Pentium is about \$180, giving Intel a gross margin of 57% for the 60-MHz parts, low by Intel standards but still quite reasonable.

This drop in estimated manufacturing cost is due to both real-world issues and improvements in the Model. As Intel has developed more experience with its 0.8-micron BiCMOS process, defect densities have declined; for a large die like Pentium, a reduction in defects creates a large cost improvement. In the Model, we have made some further across-the-board reductions in defect densities and also tuned the test and packaging areas, making the Model more accurate for low-cost devices.

Several Improvements to Cost Model

Although the Model uses the same basic equations described in the original article, several refinements help it more accurately estimate the cost of certain devices. With the advent of 200-mm (8-inch) wafers, the Model now distinguishes between devices built using larger or smaller (150-mm) wafers. The larger wafers increase wafer cost by about 50% but also increase the

number of die per wafer by 80–100%, yielding an overall cost improvement of about 20%. All of IBM's microprocessors are built on 200-mm lines, as are the latest Intel, MIPS, and SPARC processors.

MIPS processors are unusual in relying on a four-transistor (4T) RAM cell for on-chip cache rather than the larger 6T cell used by other vendors. The 4T cell reduces the size of the die but requires a second poly layer not found in most microprocessor processes. The Model now increases the wafer cost for chips with two poly layers. It also handles processes with a local interconnect layer (*see 080504.PDF*). Based on the number of process steps required, the extra poly layer is counted the same as an extra metal layer, while the local interconnect counts as one-half of a metal layer.

Defect densities in the original model ranged from 1.0 to 1.6 defects per cm^2 . Discussions with various manufacturers indicate that mature processes are operating at around 0.6 defects per cm^2 , with newer processes as high as 1.2. This change has lowered costs for nearly all the products tracked but is most significant for chips larger than 200 mm^2 .

The Model now prices a variety of packages—including TAB, BGA, and CQFP—based on the pin count (*see 071304.PDF*). Multichip modules are more difficult to handle; the estimated costs of the Power2 and Hyper-Sparc MCMs were generated by hand and are slightly less than the combined cost of discrete packaging for all the components.

	WDC 65C02	Intel 386SX	Intel 486DX2	Intel Pentium	Intel Pentium	Motorola 68060	PowerPC 601+	Micro- Sparc	R4400SC 200 MHz
Type of Process	CMOS	CMOS	CMOS	BiCMOS	BiCMOS	CMOS	CMOS	CMOS	CMOS
Process Dimension (drawn)	0.8 μm	1.0 μm	0.8 μm	0.8 μm	0.5 μm	0.5 μm	0.5 μm	0.8 μm	0.35 μm
Number of Metal Layers	2	2	3	3	4	3	5.5	2	3
Wafer Size (mm)	150 mm	150 mm	150 mm	150 mm	200 mm	200 mm	200 mm	150 mm	200 mm
Estimated Total Wafer Cost	\$1,000	\$500	\$1,200	\$1,300	\$3,400	\$2,700	\$4,100	\$1,100	\$4,000
Die Area (mm^2)	7 mm^2	43 mm^2	81 mm^2	294 mm^2	163 mm^2	198 mm^2	74 mm^2	225 mm^2	143 mm^2
Effective Area	31%	85%	85%	85%	85%	62%	97%	65%	34%
Gross Die per Wafer	2587	360	181	40	157	127	372	56	182
Estimated Defects per cm^2	0.6	0.6	0.6	0.7	1.0	1.2	1.2	0.6	1.2
Yield Percentage	99%	81%	68%	25%	32%	30%	47%	46%	59%
Untested Cost per Good Die	\$0.40	\$2.00	\$10	\$129	\$68	\$71	\$24	\$42	\$38
Package Size	44 pin	100 pin	168 pin	273 pin	273 pin	208 pin	304 pin	288 pin	447 pin
Package Type	PQFP	PQFP	PQFP	PGA	PGA	CQFP	CQFP	TAB	PGA
Estimated Package Cost	\$0.20	\$1.60	\$3	\$25	\$25	\$12	\$25	\$8	\$50
Total Test & Assembly Cost	\$0.40	\$1.50	\$4	\$27	\$25	\$10	\$20	\$7	\$28
Total Manufacturing Cost	\$1.00	\$5	\$17	\$181	\$118	\$93	\$69	\$57	\$116

Table 1. Applying the MPR Cost Model to several microprocessors, we see a range of manufacturing costs that is dependent mainly on the die area. For chips with similar die areas, the process complexity (wafer cost) and the package cost are also important.

The Model now varies the cost of wafer testing from \$50 to \$500 per hour, based on the number of signals that must be tested and the frequency of these signals. These two factors determine the price of the wafer-probe tester, which is typically a million-dollar piece of equipment. A low-cost tester can be used for a 20-MHz embedded processor in a 40-pin package, for example, but a much more expensive tester is required for a high-speed Alpha chip running at 200 MHz.

Some Processors Cost As Little As \$1

Table 1 shows how the updated Model applies to a variety of processor chips. The first column shows the 65C02 (see *080903.PDF*), an 8-bit embedded processor with a die size of just 7 mm². This chip is actually a special case because the die yield is so high (99%) that there is no point in doing any wafer test at all, eliminating this cost entirely. The chip, which sells for about \$4, costs about \$1 to build, according to the Model.

Stepping into the 32-bit world, Intel's 386SX is now estimated to cost about \$5 to build. Because Intel's 1.0-micron fabs are fully depreciated, the wafer cost of the 386 is quite low, reflecting only the incremental cost of operating the fab. The chip, which was estimated to cost \$8 under the old model, also benefits from the modifications to the test cost equations.

Comparing the 486DX2 with the 0.8-micron Pentium shows how die size directly impacts both the gross die per wafer and the yield, causing a huge difference in die cost even with a similar wafer cost. Moving Pentium to 0.5-micron BiCMOS (which Intel conservatively calls a 0.6-micron process) reduced the die size, creating a significant improvement in manufacturing cost. This cost will continue to decline as the 0.5-micron process matures, eventually dropping below \$100.

Motorola delivers several 68060 products using the same die. The 68LC060 is not guaranteed to have a working FPU, while the 68EC060 may have a bad FPU or a bad MMU; only the 68060 is fully functional. Assuming that the company can sell all parts with defective MMUs and FPUs as EC or LC versions, the manufacturing cost of all three versions is the same. The Model handles this by subtracting the MMU and FPU from the effective area of the die.

The R4400 uses a similar trick to produce PC, SC, and MC versions from the same die. Its effective area is also reduced by empty areas within the pad ring, due to the pad-limited design. The R4400 also uses redundant columns to correct defects in its large cache arrays, protecting these areas from defects. As a result, the chip has a very small effective area, just 38% of the total die area (based on measuring the die photo).

MicroSparc also has a small effective area, although this is due mainly to large empty areas throughout the die because the design was never compacted.

	Process	Die Area	Package	Est. Cost
Intel 960Sx	1.0μ, 2M	51 mm ²	PQFP-84	\$5
ARM 610	1.0μ, 2M	71 mm ²	PQFP-144	\$9
ARM 710	0.8μ, 2M	46 mm ²	PQFP-144	\$9
Intel 960Jx	0.8μ, 3M	64 mm ²	PQFP-132	\$13
Hitachi SH7604	0.8μ, 2M	82 mm ²	PQFP-144	\$14
Intel 486SX	0.8μ, 3M	72 mm ²	PQFP-196	\$15
Intel 960CA	1.0μ, 2M	137 mm ²	PQFP-196	\$21
AMD 29040	0.7μ, 3M	89 mm ²	PQFP-196	\$29
MIPS R4200	0.6μ, 3M	78 mm ²	PQFP-208	\$29
Intel 486DX2	0.8μ, 3M	81 mm ²	PGA-168	\$31
AMD 486DX2	0.7μ, 3M	89 mm ²	PGA-168	\$33
MIPS R4600	0.64μ, 3M	77 mm ²	MQFP-208	\$36
Intel 960CF	0.8μ, 2M	120 mm ²	PGA-168	\$38
Cyrix 486DX2	0.72μ, 2M	125 mm ²	PGA-168	\$40
Intel DX4	0.5μ, 4M*	77 mm ²	PGA-168	\$40
PowerPC 603	0.65μ, 4M	85 mm ²	CQFP-240	\$47
Motorola 68040	0.8μ, 3M	164 mm ²	PGA-179	\$48
R4400PC/150	0.6μ, 3M	185 mm ²	PGA-179	\$85
HP PA-7100LC	0.8μ, 3M	196 mm ²	PGA-432	\$116
HP PA-71x0	0.8μ, 3M	196 mm ²	PGA-504	\$130
PowerPC 604	0.65μ, 4M	196 mm ²	CQFP-304	\$138
R4400SC/150	0.6μ, 3M	185 mm ²	PGA-447	\$143
MicroSparc-2	0.5μ, 3M	209 mm ²	PGA-321	\$144
Digital 21066	0.6μ, 3.5M	209 mm ²	PGA-287	\$146
Digital 21064A	0.5μ, 4.5M	166 mm ²	PGA-431	\$157
SuperSparc/60	0.6μ, 3.5M*	256 mm ²	PGA-293	\$204

Table 2. According to the MPR Cost Model, some 32-bit processors cost as little as \$5 to build, while the hefty SuperSparc weighs in at \$204. (See Table 1 for other chips.) *Uses BiCMOS process

This chip, with its large die, benefits greatly from the decline in defect density. The original cost estimate of \$121 gave the product a 32% gross margin; the new estimate of \$57 gives it a more comfortable 68% margin.

The 100-MHz PowerPC 601+ demonstrates how a vendor can increase the wafer cost to reduce the per-chip cost. The 601+ uses an expensive 0.5-micron process with five metal layers plus a local interconnect. With tiny gates and lots of routing layers, the process reduces the die size of the 3.1-million transistor chip to just 74 mm², smaller than a 486. This brings the cost down to about \$69, or 10% less than the slower, 0.72-micron 601. This cost could drop another 10% as the new process matures.

Table 2 shows updated manufacturing cost estimates for a sampling of other processors. The i960SA delivers 32-bit RISC performance at a cost of just \$5 using an inexpensive 1.0-micron process, while the hefty SuperSparc costs more than \$200 to build. The R4600 is notable for costing one-fourth as much as processors with similar performance. These results show that good CPU designers must trade off several factors as they attempt to deliver high performance without breaking the manufacturing cost budget. ♦

We are continuously improving the accuracy of our cost model. If you have any comments or suggestions on cost issues, please contact us.