3.7 PowerPC 970 in 130nm and 90nm Technologies

Norman J. Rohrer[1], Miles Canada[1], Erwin Cohen[1], Mat Ringler[1], Mike Mayfield[3], Peter Sandon[1], Paul Kartschoke[1], Jay Heaslip[1], James Allen[3], Peter McCormick[1], Thomas Pflüger[2], Jeff Zimmerman[1], Cedric Lichtenau[2], Tobias Werner[3], Gerard Salem[1], Mike Ross[1], David Appenzeller[1], Dana Thygesen[1]

[1]IBM, Essex Junction, VT
[2]IBM, Boeblingen, Germany
[3]IBM, Austin, TX

A family of 64b PowerPC microprocessors is introduced in 130nm SOI technology with eight layers of copper interconnect and migrated to 90nm SOI technology with ten layers of copper interconnect. The core of the 970 processor comes from the POWER4 design [1]. In 970, SIMD instructions are added, the L2 cache is modified as a dedicated 512kB cache and a balanced encoded high-speed bus is added. The architectural features of the two 970 designs are identical[2] as shown in Fig. 3.7.1.

The 970 implements an instruction set extension to the PowerPC architecture, comprising of 162 SIMD instructions. The implementation includes a new vector register file, consisting of 32 registers of 128b each. These 128b wide vector registers, data paths and execution units support 4-way SIMD single-precision floating point operations, and 4-way integer, 8-way short integer, and 16-way byte wide fixed-point operations. The SIMD extension is implemented in the 970 by integrating two new execution units into the POWER4 core. One unit is the vector ALU, containing simple fixed, complex fixed and floating point vector units. The second unit is the permute unit, which supports arbitrary permutations by byte, as well as operand merge, pack, unpack, and splat operations.

Large portions of the SIMD unit are initially designed using 4b dataflow library elements. A preliminary set of block level physical design and timing rules is developed in order to facilitate floorplanning and pipeline optimization. The final implementation is accomplished using a mixture of custom and semi-custom styles. Placement of the semi-custom logic is controlled by schematic annotations.

The 970 implementation is deeply pipelined, and supports dispatch of up to 5 instructions per cycle and issues up to 1 instruction per cycle to each of its 10 execution units. The Dhrystone 2.1 MIPS rating of the 970 is 2.9DMIPS/MHz. This corresponds to a 5800DMIPS rating for the 2.0GHz processor, and a 7250DMIPS for the 2.5GHz processor. The SPECint2000 rating for the 2.0GHz 970 is 898 and the corresponding SPECfp2000 rating is 1045.

The PowerPC 970 implements two unidirectional, point-to-point, 36b wide, source synchronous busses to communicate to the northbridge chipset, Fig. 3.7.2. These busses constitute the elastic interface[3]. The elastic interface can pipeline data across multiple cycle wire delays between source and sink points. This decouples the bus latency and bandwidth, allowing higher bandwidth for a given bus latency. This functionality is supported by an initialization alignment procedure, run during power-on-reset, which deskews bit arrival times across each of the data bits in both busses, and centers the clock in the signal eye for maximum noise immunity.

The bus can run at speeds up to 1.1GHz (1.1GT/s), but always an integral fraction (1/2, 1/3, 1/4 or 1/6) of the processor frequency. For the 2.0GHz processor, the 1.0GHz bus speed corresponds to a raw bandwidth of 8.0GB/s. Allowing for the multiplexing of address and control signals, and handshaking overhead, this corresponds to 7.1GB/s of useable bandwidth. The single ended line requires an eye opening of 517ps. The 970+ adds power-saving capability with selectable termination. The bus can run with or without an effective 55Ω termination, resulting in a power saving of approximately 2W. Unterminated, the bus can attain 633MT/s.

In the 130nm design, the I/Os operate at the core voltage. To maintain compatibility and high bus speeds, the 90nm design split the voltage plane of the I/O from the core VDD plane. When the I/O voltage is held equal to the core voltage, the 970 and 970+ have equivalent module footprints. The 970 design uses 1417 C4, while the newly designed chip has 1446 C4s. The need to maintain package parasitics and C4 current levels drove a doubling of the C4 density compared to the 130nm design.

The 970 has 190nF of on-chip capacitance for power noise decoupling. 52% of the capacitance is a thin-oxide-accumulation capacitor for space and density optimization. The 90nm oxide cannot use thin-oxide capacitors due to prohibitively large oxide tunneling currents. A third gate oxide is optimized to provide maximum capacitance with tolerable leakage at a thickness of 15A. The optimized oxide provides 131nF of decoupling capacitance to the 970+, allowing supply noise to be contained within 3% of the previous design.

PowerPC 970+ implements the use of an electronic fuse, eFUSE. The fuses will be used for array redundancy of the large caches, electronic chip identification, part number identification and thermal diode calibration. Each processor will hold the test temperature and thermal diode voltage read at ambient wafer testing and from elevated module testing. The eFUSE is a silicided PC bridge that is electromigrated to cause a >10x increase in resistance. Figure 3.7.3 shows the blow devices and a single ended sense head along with micrographs of an unblown and a blown fuse. The blow conditions are 15mA for 200μS using 3.3 to 3.7V with VDD=1.0V. The fuse sensing can be done over a VDD range of 0.8-1.5V.

PowerPC 970+ also implements a frequency and voltage scaling technique called PowerTune[4], for reducing the power during operation. PowerTune allows the processor to switch the frequency from its maximum value to half the maximum value or to quarter of the maximum frequency. At any reduced frequency, the voltage can be lowered to further reduce the total power consumption. At any frequency/voltage combination, the processor can enter a nap mode for additional saving power. The power can be further reduced by entering a deep nap mode which runs at 1/64th of the clock maximum frequency. Deep nap allows the processor to reduce the active switching power and maintain the clocks running at a high enough speed to eliminate the impact of the history effect upon restarting the clocks.

The electrical and physical aspects of the 130nm[5] to 90nm[6] design and technology are shown in Fig. 3.7.4. With a total of 58 million transistors, the 130nm design is 118mm^2, while the 90nm design is 62mm^2. The area scaling from 130nm to 90nm is effectively a 50% reduction including the 6-T SRAM cell size, which is reduced from 2.16μm^2 to 1.06μm^2. At nominal process, the operating frequency is increased from 1.8GHz at 1.45V to 2.5GHz at 1.3V. Figure 3.7.5 shows the typical power for 970 and 970+. The typical power decreased from 57W to 49W, going from 1.8GHz to 2.5GHz, respectively. The switching power decreased by 38% due to a lower voltage and less capacitance, even at the higher frequency. The power due to gate tunneling decreased by 16% due to a lower voltage, smaller area and removal of the decoupling capacitors from the thin oxide devices. The power due to the subthreshold current increased due to a leakier device that outpaced the reduction in voltage and the reduction in device width. This power can be further reduced to 31W using frequency scaling and to 15W using both frequency and voltage scaling as shown at f/4 or 625MHz in Fig. 3.7.5.

References:

[1] D.C. Bossen, et al., "Power4 System Design for High Reliability," *IEEE Micro*, vol. 22, no. 2, pp. 16-24, Mar. 2002.
[2] P. Sandon,"PowerPC 970: First in a New Family of 64-bit High Performance PowerPC Processors," *Microprocessor Forum*, Oct. 2002.
[3] F. Ferriaolo, "Synchronous Wave-Pipelined Interface," *Hot Chips*, Aug. 1999.
[4] Cedric Lichtenau et al., "PowerTune: Advanced Frequency and Power Scaling on PowerPC 970+," *ISSCC Dig. Tech. Papers,* pp. 356-357, Feb. 2004.
[5] E. Leobandung et al., "Scalability of SOI Technology into 0.13μm 1.2V CMOS Generation," *IEEE Electron Devices Meeting Digest*, pp. 403-406, 1998.
[6] M. Khare et al., "A High Performance 90nm SOI Technology with 0.992mm 6T-SRAM Cell," *IEEE Electron Devices Meeting Digest*, pp. 407-410, Dec. 2002.

 0-7803-8267-6/04

- **Instruction pipe**
  - 64KB L1 Inst cache, direct mapped
  - 8 instructions fetch / cycle
- **Branch prediction**
  - Highly accurate dynamic prediction
- **Dispatch, issue**
  - 1 group (4 + branch) / cycle
  - Over 200 instructions in flight
- **Data pipe**
  - 32 KB L1 Data cache, 2-way SA
  - 32 x 64b GPR, FPR
  - 32 x 128b VRF
  - 512KB L2 cache, 8-way SA
  - 8 data prefetch streams
- Execution
  - 2 Load/store units
  - 2 Fixed point units
  - 2 IEEE floating point units
  - 2 SIMD sub-units
  - Branch unit
  - Condition register unit

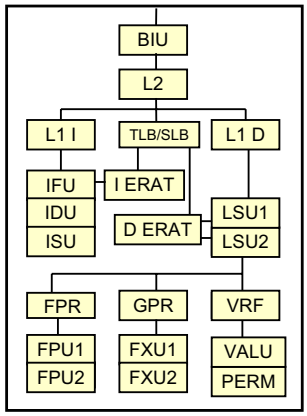

Figure 3.7.1: PowerPC 970 architectural features.

- Features
  - Two unidirectional busses
  - 36-bit read, 36-bit write, encoded to 44b
  - Point-to-point
  - Source synchronous
- Elastic interface
  - Allows multiple cycle wire delays between chips
  - Hardware deskews bit lines at POR
- Bus protocol
  - Address, control and data multiplexing
  - Sideband signals for snoop and ACK
  - Pipelined transactions
  - Out of order data
  - Coherency and sharing via snooping
  - Processor synchronization for SMP
- A 1.0GHz bit rate achieves 7.1GB/s useable bandwidth

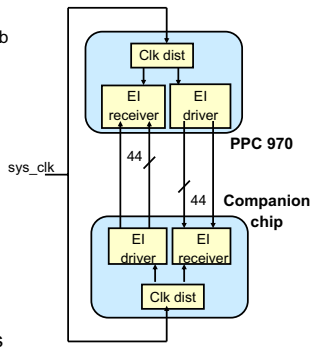

Figure 3.7.2: High bandwidth processor bus.

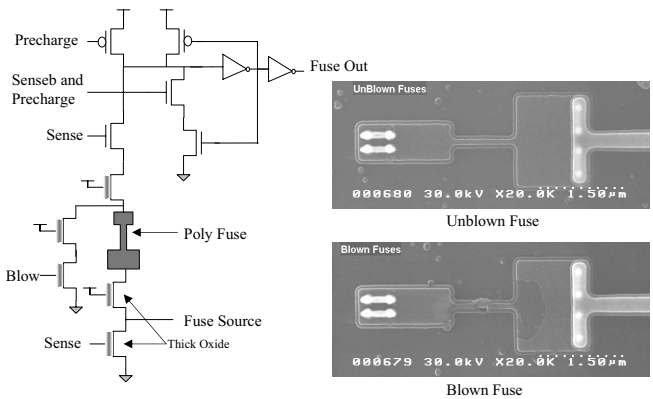

Figure 3.7.3: eFUSE circuit with poly fuses.

	970	970+
Technology	130nm, CMOS w/ SOI	90nm, CMOS w/ SOI
Gate Lpoly	55nm	46nm
Tox	13 & 22 Angstrom	10.5, 15 & 22 Angstrom
NFET/PFET I_{dsat}	917/418 uA/um @ 1.2V	938/400 uA/um @ 1.0V
Metal Levels	8 (5-1x, 2-2x, 1-4x) - FTEOS	10 (5-1x, 3-2x, 2-6x) - FTEOS
Die Size	118mm²	62mm²
FETs	58 Million	
Caches	64KB, instruction cache, w/parity 32KB, data cache, w/parity 512KB, L2 cache, w/ECC	
Voltages	1.2-1.5V for core and I/O VDD	0.9V-1.35V for core VDD 1.6V I/O
Frequency	1.3-2.0GHz	1.3-2.5GHz
Typical Power	66W @ 2.0GHz	50W @ 2.5GHz
Estimated Performance	898 SPECint2000 @ 2.0 GHz 1045 SPECfp2000 @ 2.0 GHz 5800 DMIPS @ 2.0 GHz (2.9DMIPS/MHz)	1082 SPECint2000 @ 2.5 GHz 1361 SPECfp2000 @ 2.5 GHz 7250 DMIPS @ 2.5 GHz (2.9DMIPS/MHz)
Package	25x25mm CBGA 576 pins on 1mm pitch (161 signals)	

Figure 3.7.4: Chip and technology features.

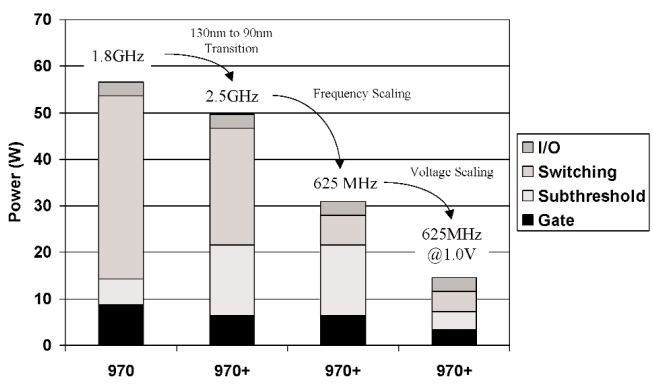

Figure 3.7.5: Typical power for 130nm and 90nm at nominal process and elevated voltage.

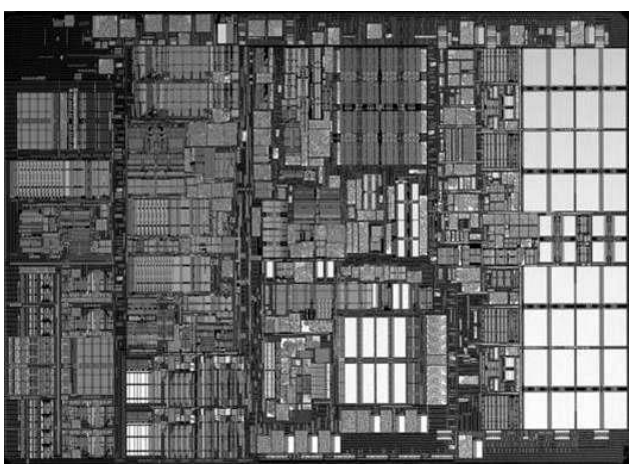

Figure 3.7.6: PowerPC 970 die micrograph.

0-7803-8267-6/04

- **Instruction pipe**
  - 64KB L1 Inst cache, direct mapped
  - 8 instructions fetch / cycle
- **Branch prediction**
  - Highly accurate dynamic prediction
- **Dispatch, issue**
  - 1 group (4 + branch) / cycle
  - Over 200 instructions in flight
- **Data pipe**
  - 32 KB L1 Data cache, 2-way SA
  - 32 x 64b GPR, FPR
  - 32 x 128b VRF
  - 512KB L2 cache, 8-way SA
  - 8 data prefetch streams
- Execution
  - 2 Load/store units
  - 2 Fixed point units
  - 2 IEEE floating point units
  - 2 SIMD sub-units
  - Branch unit
  - Condition register unit

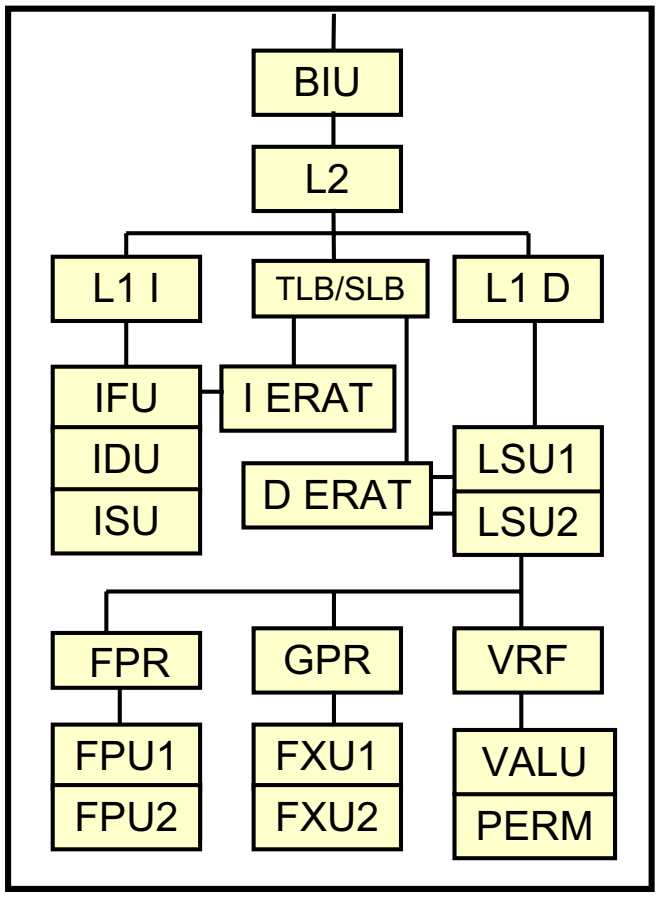

Figure 3.7.1: PowerPC 970 architectural features.

0-7803-8267-6/04

- Features
  - Two unidirectional busses
  - 36-bit read, 36-bit write, encoded to 44b
  - Point-to-point
  - Source synchronous
- Elastic interface
  - Allows multiple cycle wire delays between chips
  - Hardware deskews bit lines at POR
- Bus protocol
  - Address, control and data multiplexing
  - Sideband signals for snoop and ACK
  - Pipelined transactions
  - Out of order data
  - Coherency and sharing via snooping
  - Processor synchronization for SMP
- A 1.0GHz bit rate achieves 7.1GB/s useable bandwidth

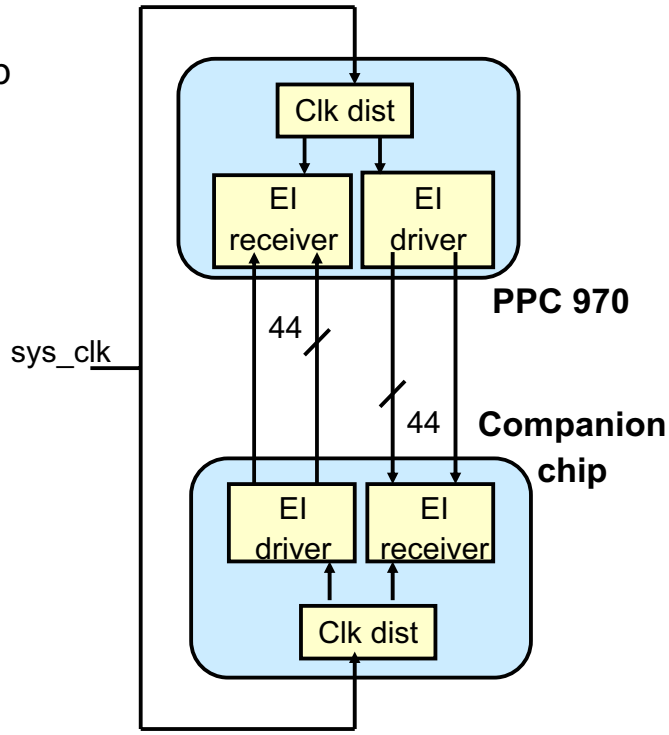

Figure 3.7.2: High bandwidth processor bus.

0-7803-8267-6/04

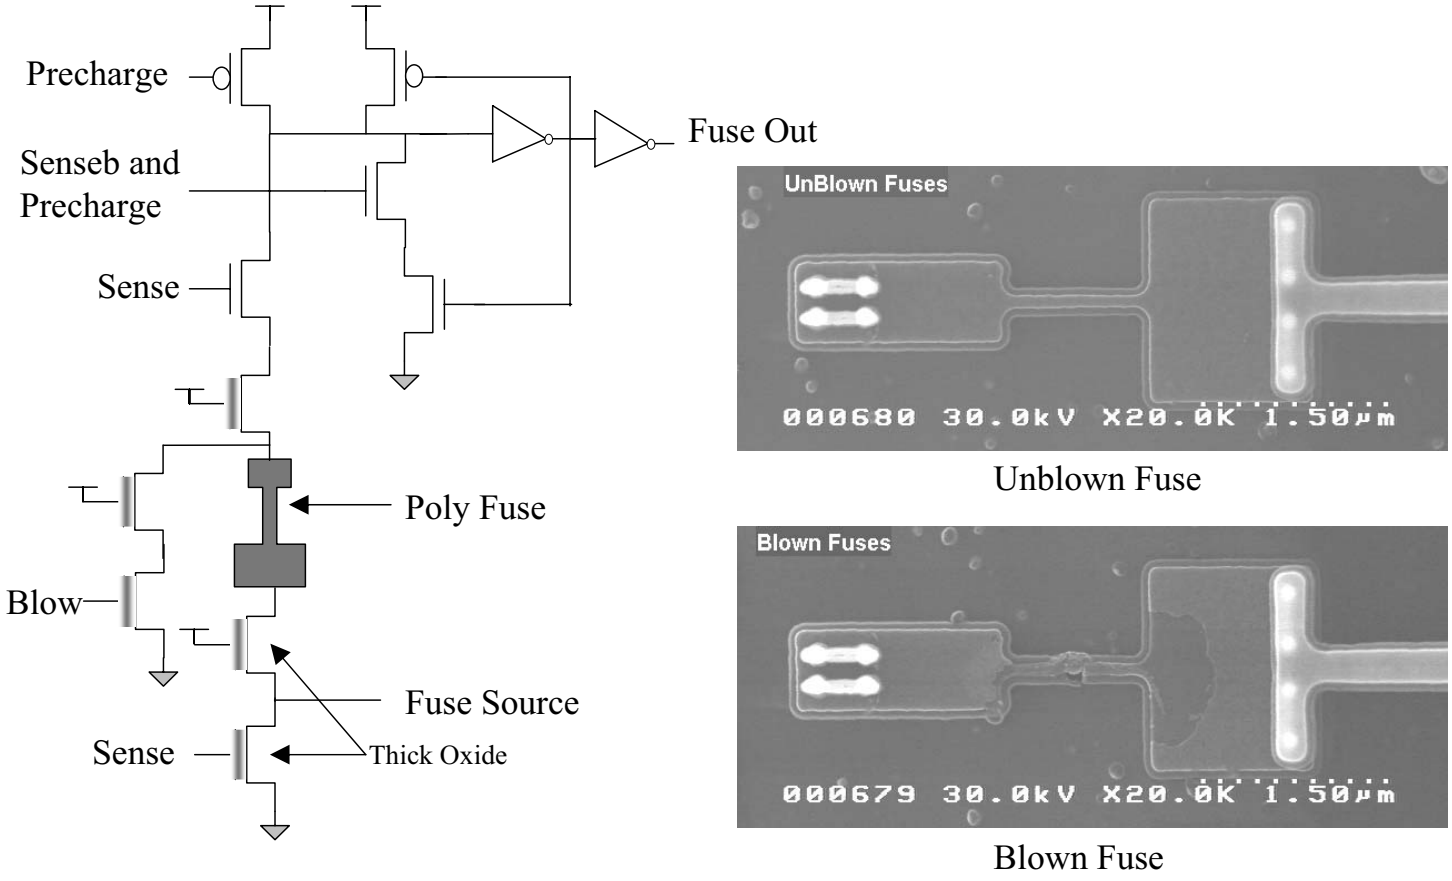

Figure 3.7.3: eFUSE circuit with poly fuses.

	970	970+
Technology	130nm, CMOS w/ SOI	90nm, CMOS w/ SOI
Gate Lpoly	55nm	46nm
Tox	13 & 22 Angstrom	10.5, 15 & 22 Angstrom
NFET/PFET I_{dsat}	917/418 uA/um @ 1.2V	938/400 uA/um @ 1.0V
Metal Levels	8 (5-1x, 2-2x, 1-4x) - FTEOS	10 (5-1x, 3-2x, 2-6x) - FTEOS
Die Size	118mm^2	62mm^2
FETs	58 Million	
Caches	64KB, instruction cache, w/parity 32KB, data cache, w/parity 512KB, L2 cache, w/ECC	
Voltages	1.2-1.5V for core and I/O VDD	0.9V-1.35V for core VDD 1.6V I/O
Frequency	1.3-2.0GHz	1.3-2.5GHz
Typical Power	66W @ 2.0GHz	50W @ 2.5GHz
Estimated Performance	898 SPECint2000 @ 2.0 GHz 1045 SPECfp2000 @ 2.0 GHz 5800 DMIPS @ 2.0 GHz (2.9DMIPS/MHz)	1082 SPECint2000 @ 2.5 GHz 1361 SPECfp2000 @ 2.5 GHz 7250 DMIPS @ 2.5 GHz (2.9DMIPS/MHz)
Package	25x25mm CBGA 576 pins on 1mm pitch (161 signals)	

Figure 3.7.4: Chip and technology features.

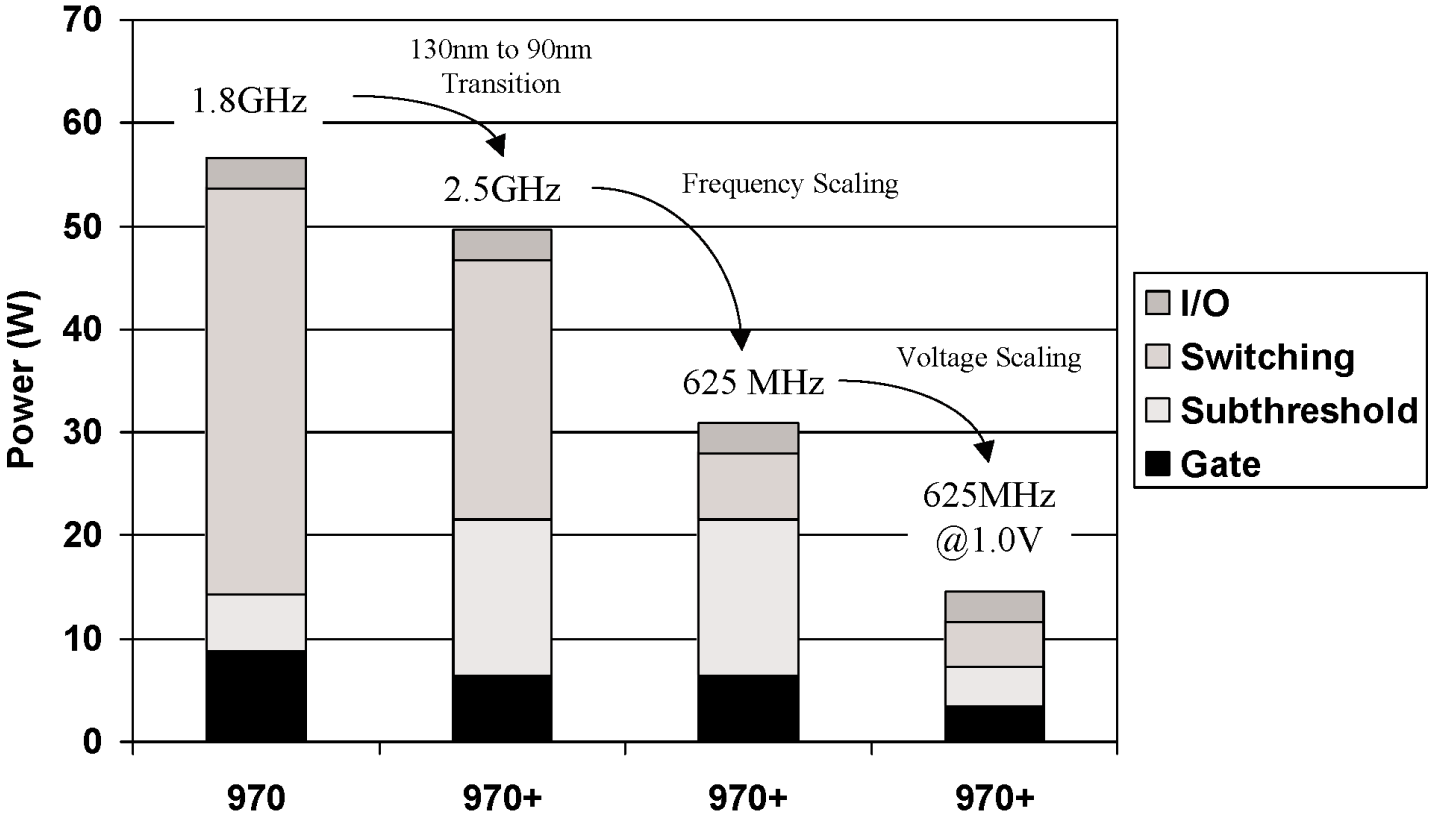

Figure 3.7.5: Typical power for 130nm and 90nm at nominal process and elevated voltage.

0-7803-8267-6/04

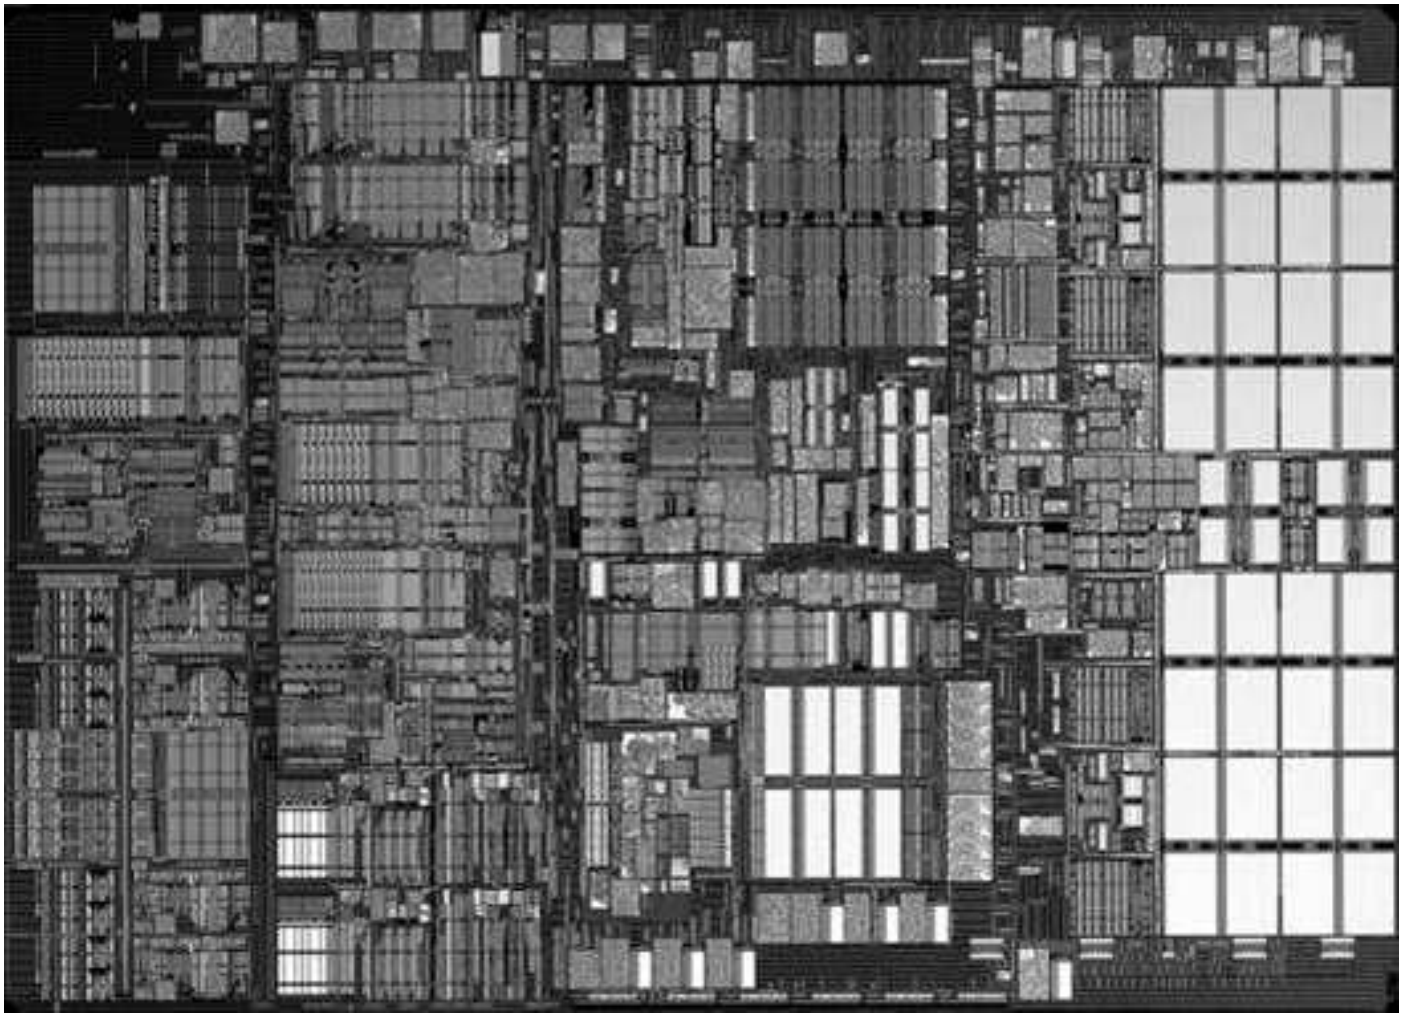

Figure 3.7.6: PowerPC 970 die micrograph.