IBM

Research topics
Researcher resources
Research news
Latest news
resources for the press
IBM research worldwide
Did you know?
Employment
Feedback

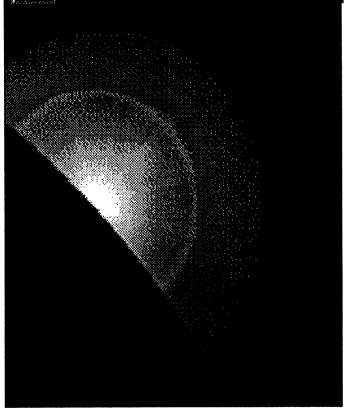

Blue Gene Architecture

IBM's Recipe for Building a Petaflop Computer

Building a supercomputer that can operate at a speed of one petaflop, or one quadrillion floating point operations (FLOPS) per second, requires a radical departure from the design of traditional supercomputers. In fact, at today's rate of progress based on Moore's Law and other measures of advancement, achieving a petaflop-scale supercomputer would take up to 15 years. IBM's goal: less than five.

A FLOP (floating point operation) is an arithmetic operation, such as addition or multiplication, on numbers in scientific notation. It is generally the most complex of basic computer operations.

In order to achieve this groundbreaking speed in such a short time, IBM researchers are experimenting with ways to simplify traditional computer architectures, allowing chips to achieve performance improvements at a much faster pace than the current roadmap for today's chips.

IBM researchers are calling their experimental new design "SMASH," which stands for Simple, Many, And Self-Healing. "Simple" refers to the elemental architecture; "Many" describes Blue Gene's one million processors working in parallel; and "Self-Healing" means greater fault-tolerance and system stability. This experimental architecture should allow for more processors in less space using less power, resulting in more operations overall. Researchers expect this massively parallel system to be capable of eight million simultaneous threads of computation (compared to the maximum of 5000 threads today).

The details on SMASH

Research topics
Researcher resources
Research news
latest news
resources for the press
IBM research worldwide
Did you know?
Employment
Feedback

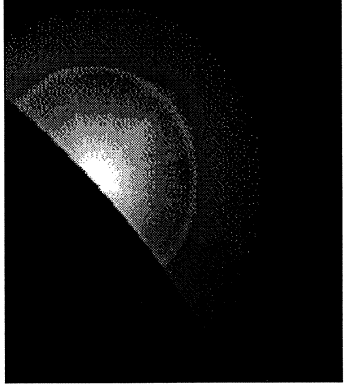

Blue Gene Architecture

IBM's Recipe for Building a Petaflop Computer

Building a supercomputer that can operate at a speed of one petaflop, or one quadrillion floating point operations (FLOPS) per second, requires a radical departure from the design of traditional supercomputers. In fact, at today's rate of progress based on Moore's Law and other measures of advancement, achieving a petaflop-scale supercomputer would take up to 15 years. IBM's goal: less than five.

A FLOP (floating point operation) is an arithmetic operation, such as addition or multiplication, on numbers in scientific notation. It is generally the most complex of basic computer operations.

In order to achieve this groundbreaking speed in such a short time, IBM researchers are experimenting with ways to simplify traditional computer architectures, allowing chips to achieve performance improvements at a much faster pace than the current roadmap for today's chips.

IBM researchers are calling their experimental new design "SMASH," which stands for Simple, Many, And Self-Healing. "Simple" refers to the elemental architecture; "Many" describes Blue Gene's one million processors working in parallel; and "Self-Healing" means greater fault-tolerance and system stability. This experimental architecture should allow for more processors in less space using less power, resulting in more operations overall. Researchers expect this massively parallel system to be capable of eight million simultaneous threads of computation (compared to the maximum of 5000 threads today).

The details on SMASH

This approach will require researchers to explore significant departures from traditional computer design, including:

Embedded memory: The dynamic random access memory (DRAM) will be on the chips, making the memory much more accessible to the processors and radically improving access time and bandwidth. Tightly integrating the logic and memory like this also significantly reduces power requirements.

Minimalist design: By simplifying the overall architecture and using so many processors in parallel, Blue Gene will achieve incredibly fast system performance. Also, because the entire system will be built by replicating one chip 32,000 times, the overall project complexity is greatly reduced.

Multi-threading: Each processor is like a cook preparing eight recipes at once -- the cook starts one dish mixing, then moves to the next recipe, and so on. By the time the cook gets the eighth recipe started, the first is ready for its next step. In computer design, this is called multi-threading, and a primary goal for computer scientists is to keep the "cook" as busy as possible. Blue Gene will have one million of these "cooks," working on eight million concurrent "recipes" or threads.

High communication bandwidth: Blue Gene will have extremely high communication bandwidth among its chips, allowing them to exchange massive amounts of data faster than ever before. With six channels on each chip, each channel sending data at two gigabytes per second, the total communication is 300 terabytes per second. In fact, the aggregate communication bandwidth of Blue Gene is roughly equal to every one of the six billion people in the world operating four ISDN modems simultaneously. If you could harness Blue Gene's bandwidth to download the entire contents of the Internet -- all 100 terabytes -- it would take less than a second.

Self-Healing: The self-healing aspect of Blue Gene's architecture is one of the biggest challenges the research team faces. The hardware -- with its tremendous redundancy in processor and communication paths -- makes this concept feasible since it provides many routes to access a processor, and many processors over which to distribute calculations. Self-management is a huge challenge for

a machine of this scale. IBM's software researchers are exploring advanced technologies for distributed control and recovery to insure that Blue Gene runs continuously.

Small footprint delivers big power

The final machine will have 64 interconnected racks, each rack containing eight boards, each board containing 64 chips, and each 20x20 mm chip comprised of 32 processors. Each processor will operate at one gigaflop, which is one billion operations per second.

Just one of Blue Gene's processor boards (a two-foot by two-foot component) will be capable of two teraflops -- the same raw computational power as the ASCI Blue Pacific supercomputer that IBM installed earlier this year at Lawrence Livermore National Laboratory.

Upon completion, Blue Gene will occupy less than 2,000 square feet, taking up much less space than the 8,000 square feet needed to house the world's largest computer today. At a speed of one petaflop, Blue Gene will be:

- 500 times more powerful than the fastest computer in the world today;
- 1,000 times more powerful than Deep Blue;
- 40 times faster than the sum power of the 40 fastest supercomputers in the world today;
- two million times more powerful than today's top desktop PC.

BACK >>

Research home | IBM home | Order | Privacy | Legal | Contact IBM

IBM Announces $100 Million Research Initiative to build World's Fastest Supercomputer

"Blue Gene" to Tackle Protein Folding Grand Challenge

YORKTOWN HEIGHTS, NY, December 6, 1999 -- IBM today announced a new $100 million exploratory research initiative to build a supercomputer 500 times more powerful than the world?s fastest computers today.

The new computer -- nicknamed "Blue Gene" by IBM researchers -- will be capable of more than one quadrillion operations per second (one petaflop). This level of performance will make Blue Gene 1,000 times more powerful than the Deep Blue machine that beat world chess champion Garry Kasparov in 1997, and about 2 million times more powerful than today's top desktop PCs.

Blue Gene's massive computing power will initially be used to model the folding of human proteins, making this fundamental study of biology the company's first computing "grand challenge" since the Deep Blue experiment. Learning more about how proteins fold is expected to give medical researchers better understanding of diseases, as well as potential cures.

"This is exactly what IBM Research does best -- continuously placing big, aggressive bets on technologies that change the future of computing," said Dr. Paul M. Horn, senior vice president of IBM Research. "In many ways, Deep Blue got a better job today -- if this computer unlocks the mystery of how proteins fold, it will be an important milestone in the future of medicine and healthcare."

Experimental New Architecture Key to Petaflop Performance IBM Research believes a radical new approach to computer design and architecture will allow Blue Gene to achieve petaflop-scale performance in about five years -- one-third of the close to 15 years it would normally take following Moore's Law. The two fastest computers in the world today are part of the ASCI program run by the U.S. Department of Energy, and which were recently tested at about 2 teraflops -- two trillion operations per second each.

"We think a tremendous gain in performance will be made possible by the first major revolution in how computers are built since the mid-1980s," said Dr. Ambuj Goyal, IBM Research's vice president of computer science. "We call this new approach to computer architecture SMASH, which stands for Simple, Many and Self-Healing."

The SMASH architecture differs from existing approaches in three ways:

* It dramatically simplifies the number of instructions carried out by each processor, allowing them to work faster and with significantly lower power and chip surface requirements (the traditional approach is to add complex features to gain performance);

* It will facilitate a massively parallel system capable of more than 8 million simultaneous threads of computation (compared to the maximum of 5000 threads today);

* It will make the computer self-stabilizing and self-healing -- automatically able to overcome failures of individual processors and computing threads.

Blue Gene will consist of more than one million processors, each capable of one billion operations per second (1 gigaflop). Thirty-two of these ultra-fast processors will be placed on a single chip (32 gigaflops). A compact two-foot by two-foot board containing 64 of these chips will be capable of 2 teraflops, making it as powerful as the 8000-square foot ASCI computers.

Eight of these boards will be placed in 6-foot-high racks (16 teraflops), and the final machine (less than 2000 sq. ft.) will consist of 64 racks linked together to achieve the one petaflop performance.

Protein Folding Holds Key to Understanding Basics of Life The scientific community considers protein folding one of the most significant "grand challenges" -- a fundamental problem in science or engineering that has broad economic and scientific impact and whose solution can be advanced only by applying high-performance computing technologies.

Proteins control all cellular processes in the human body. Comprising strings of amino acids that are joined like links of a chain, a protein folds into a highly complex, three-dimensional shape that determines its function. Any change in shape dramatically alters the function of a protein, and even the slightest change in the folding process can turn a desirable protein into a disease.

Better understanding of how proteins fold will give scientists and doctors better insight into diseases and ways to combat them. Pharmaceutical companies could design high-tech prescription drugs customized to the specific needs of individual people. And doctors could respond more rapidly to changes in bacteria and viruses that cause them to become drug-resistant.

"Breakthroughs in computers and information technology are now creating new frontiers in biology," said Horn. "One day, you're going to be able to walk into a doctor's office and have a computer analyze a tissue sample, identify the pathogen that ails you, and then instantly prescribe a treatment best suited to your specific illness and individual genetic makeup."

IBM Research

About 50 scientists from IBM Research's Deep Computing Institute and Computational Biology Group will work on Blue Gene and the protein folding grand challenge. IBM Research is the world's largest research organization dedicated to information technology, with eight labs around the world, including Austin, Beijing, Delhi, Haifa, Tokyo, San Jose, Yorktown Heights (New York), and Zurich. IBM Research works closely with the company's Life Sciences Unit, which was formed earlier this year to deliver information technology services and solutions to advance the pharmaceutical, biotechnology and agri-science industries.

For pictures available for download, go here.

IBM ANNOUNCES $100 MILLION RESEARCH INITIATIVE TO BUILD WORLD'S FASTEST SUPERCOMPUTER

IBM IBM today announced a new $100 million exploratory research initiative to build a supercomputer 500 times more powerful than the world's fastest computers today.

The new computer - nicknamed "Blue Gene" by IBM researchers -will be capable of more than one quadrillion operations per second (one petaflop). IBM Research believes a radical new approach to computer design and architecture will achieve petaflop-scale performance in about five years - less than one third of the time it would normally take following Moore's Law.

Blue Gene's massive computing power will be initially used to model the folding of human proteins, making this fundamental study of biology the company's first computing "grand challenge" since Deep Blue defeated world chess champion Garry Kasparov in 1997. Learning more about how proteins fold is expected to give doctors better understanding of diseases, as well as potential cures.

"This is exactly what IBM Research does best - continuously placing big, aggressive bets on technologies that change the future of computing," said Dr. Paul M. Horn, senior vice president of IBM Research. "What's more, if this computer unlocks the mystery of how proteins fold, it will be an important milestone in the future of medicine and healthcare. We feel like Deep Blue grew up today and got a real job."

The secret behind Blue Gene's performance: "SMASH" architecture

Upon completion, the new IBM Research computer is expected to be capable of more than 1,000,000,000,000,000 operations per second, making it:

* 500 times more powerful than the fastest computers in the world today (two computers that are part of the ASCI programme run by the U.S. Department of Energy, and which were recently tested at about 2 teraflops - two trillion operations per second - each); * 1,000 times more powerful than the Deep Blue machine that beat Garry Kasparov; * 40 times faster than the sum power of the 40 fastest computers in the world today; and * 2 million times more powerful than today's top desktop PCs.

"This tremendous gain in performance will be made possible by the first major revolution in how computers are built since the mid-1980s," said Dr. Ambuj Goyal, IBM Research's vice president of computer science. "We call our new approach to computer architecture SMASH, which stands for Simple, Many and Self-Healing."

The SMASH architecture is radical in three ways:

* It will dramatically simplify the number of instructions carried out by each processor, allowing them to work faster and with significantly lower power requirements (the traditional approach is to make instruction sets more complex to gain performance); * It will facilitate a massively parallel system capable of more than 8 million simultaneous threads of computation (compared to the maximum of 5000 threads today); * It will make the computer self-stabilising and self-healing - automatically able to

overcome failures of individual processors and computing threads;

Blue Gene will be based on IBM's RS/6000 platform and will consist of more than one million processors, each capable of one billion operations per second (1 gigaflop). Thirty-two of these ultra-fast processors will be placed on a single chip (32 gigaflops). A compact two-foot by two-foot board containing 64 of these chips will be capable of 2 teraflops, making it as powerful as the 8000-square foot ASCI computers.

Eight of these boards will be placed in 6-foot-high racks (16 teraflops), and the final machine (less than 2000 sq. ft.) will consist of 64 racks linked together to achieve the one petaflop performance.

Protein folding holds key to understanding basics of life

The scientific community considers protein folding one of the most significant "grand challenges" - a fundamental problem in science or engineering that has broad economic and scientific impact and whose solution can be advanced only by applying high-performance computing technologies.

Proteins control all cellular processes in the human body. Comprising strings of amino acids that are joined like links of a chain, proteins fold into a highly complex, three-dimensional shape that determines their function. Any change in shape dramatically alters the function of a protein, and even the slightest change in the folding process can turn a desirable protein into a disease.

Better understanding of how proteins fold will give scientists and doctors better insight into diseases and ways to combat them. Pharmaceutical companies could design high-tech prescription drugs customised to the specific needs of individual people. And doctors could respond more rapidly to changes in bacteria and viruses that cause them to become drug-resistant.

"The next frontier in biology is being enabled by breakthroughs in computers and information technology," said Horn. "One day, you're going to be able to walk into a doctor's office and have a computer analyse a tissue sample, identify the offending pathogen, and then instantly prescribe a treatment best suited your specific illness and your genetic makeup."

Editors notes

IBM Research

Between 50 and 75 scientists from IBM Research's Deep Computing Institute and Computational Biology Group will work on Blue Gene and the protein folding grand challenge. IBM Research is the world's largest research organisation dedicated to information technology, with more than 3000 scientists and engineers at eight labs around the world, including Austin, Beijing, Delhi, Haifa, Tokyo, San Jose, Yorktown Heights (New York), and Zurich. IBM research works closely with the company's Life Sciences Unit, which was formed earlier this year to deliver information technology services and solutions to advance the pharmaceutical, biotechnology and agri- science industries.

About RS/6000

More than 800,000 IBM RS/6000 systems are in use by over 100,000 commercial and technical customers around the world. The RS/6000 is IBM's family of computers that feature the Reduced Instruction Set Computing (RISC)-based POWER family of microprocessors and AIX, IBM's UNIX operating system. RS/6000 products range in size and capacity from workstations, work group and enterprise servers, to the RS/6000 SP "supercomputer." From businesses working to become more efficient and profitable, to governments and universities seeking to solve the greatest challenges of our time, the RS/6000 supports a wide range of applications and provides the reliability, availability and price/performance that today's information technology managers demand.

For more general information, visit the IBM RS/6000 web site at www.rs6000.ibm.com.

Blue Gene Architecture

IBM's Recipe for Building a Petaflop Computer

Building a supercomputer that can operate at a speed of one petaflop, or one quadrillion floating point operations (FLOPS) per second, requires a radical departure from the design of traditional supercomputers. In fact, at today's rate of progress based on Moore's Law and other measures of advancement, achieving a petaflop-scale supercomputer would take up to 15 years. IBM's goal: less than five.

A FLOP (floating point operation) is an arithmetic operation, such as addition or multiplication, on numbers in scientific notation. It is generally the most complex of basic computer operations.

In order to achieve this groundbreaking speed in such a short time, IBM researchers are experimenting with ways to simplify traditional computer architectures, allowing chips to achieve performance improvements at a much faster pace than the current roadmap for today's chips.

IBM researchers are calling their experimental new design "SMASH," which stands for Simple, Many, And Self-Healing. "Simple" refers to the elemental architecture; "Many" describes Blue Gene's one million processors working in parallel; and "Self-Healing" means greater fault-tolerance and system stability. This experimental architecture should allow for more processors in less space using less power, resulting in more operations overall. Researchers expect this massively parallel system to be capable of eight million simultaneous threads of computation (compared to the maximum of 5000 threads today).

The details on SMASH

This approach will require researchers to explore significant departures from traditional computer design, including:

Embedded memory: The dynamic random access memory (DRAM) will be on the chips, making the memory much more accessible to the processors and radically improving access time and bandwidth. Tightly integrating the logic and memory like this also significantly reduces power requirements.

Minimalist design: By simplifying the overall architecture and using so many processors in parallel, Blue Gene will achieve incredibly fast system performance. Also, because the entire system will be built by replicating one chip 32,000 times, the overall project complexity is greatly reduced.

Multi-threading: Each processor is like a cook preparing eight recipes at once -- the cook starts one dish mixing, then moves to the next recipe, and so on. By the time the cook gets the eighth recipe started, the first is ready for its next step. In computer design, this is called multi-threading, and a primary goal for computer scientists is to keep the "cook" as busy as possible. Blue Gene will have one million of these "cooks," working on eight million concurrent "recipes" or threads.

High communication bandwidth: Blue Gene will have extremely high communication bandwidth among its chips, allowing them to exchange massive amounts of data faster than ever before. With six channels on each chip, each channel sending data at two gigabytes per second, the

total communication is 300 terabytes per second. In fact, the aggregate communication bandwidth of Blue Gene is roughly equal to every one of the six billion people in the world operating four ISDN modems simultaneously. If you could harness Blue Gene's bandwidth to download the entire contents of the Internet -- all 100 terabytes -- it would take less than a second.

Self-Healing: The self-healing aspect of Blue Gene's architecture is one of the biggest challenges the research team faces. The hardware -- with its tremendous redundancy in processor and communication paths -- makes this concept feasible since it provides many routes to access a processor, and many processors over which to distribute calculations. Self-management is a huge challenge for a machine of this scale. IBM's software researchers are exploring advanced technologies for distributed control and recovery to insure that Blue Gene runs continuously.

Small footprint delivers big power

The final machine will have 64 interconnected racks, each rack containing eight boards, each board containing 64 chips, and each 20x20 mm chip comprised of 32 processors. Each processor will operate at one gigaflop, which is one billion operations per second.

Just one of Blue Gene's processor boards (a two-foot by two-foot component) will be capable of two teraflops -- the same raw computational power as the ASCI Blue Pacific supercomputer that IBM installed earlier this year at Lawrence Livermore National Laboratory.

Upon completion, Blue Gene will occupy less than 2,000 square feet, taking up much less space than the 8,000 square feet needed to house the world's largest computer today. At a speed of one petaflop, Blue Gene will be:

* 500 times more powerful than the fastest computer in the world today;

* 1,000 times more powerful than Deep Blue;

* 40 times faster than the sum power of the 40 fastest supercomputers in the world today;

* two million times more powerful than today's top desktop PC.

RS/6000 SP

- Distributed memory, multinode server designed for demanding technical and commercial workloads
- Versatile system running serial, symmetric multiprocessor (SMP) and parallel workloads all managed from a central point-of-control.
- Flexible configurability
  - node types (thin, wide, high)
  - upto 512 nodes per system
- High degree of RAS (reusability, availability, serviceability)

Related Links

- SP specifications and performance
- SP description
- SP description and specs - pdf format
- Parallel Tools and Software
- RS/6000 SP White Papers

RS/6000 SP Comparison Chart

Machine type	**9076**	**9076**	**9076**	**9076**	**9076**
Node Type	332 MHz	332 MHz	375 MHz POWER3 SMP	375 MHz POWER3 SMP	375 MH SMP
Packaging	Thin	Wide	Thin	Wide	High
Microprocessor					
Type	PowerPC 604e	PowerPC 604e	POWER3-II	POWER3-II	POWER
Processors per node	2/4	2/4	2/4	2/4	4/8/12/1
Clock rates	332 MHz	332 MHz	375 MHz	375 MHz	375 MH
Memory					
System(min/max)	256MB/3GB*	256MB/3GB*	256MB/16GB*	256MB/16GB*	1GB/64
L2 Cache	256KB**	256KB**	8MB**	8MB**	8MB**
PCI I/O Capacity					
Slots available(base node)	2 PCI (32-bit)	10 PCI(7 32-bit/3 64-bit)	2 PCI (32-bit)	10 PCI (2 32-bit/8 64-bit)	5 PCI(1 64-bit)
Slots available(per SP Expansion I/O)	N/A	N/A	N/A	N/A	8 PCI (6
Slots available(base node and six SP Expansion I/O Units)	N/A	N/A	N/A	N/A	53 PCI (64-bit)
Internal Disk Capacity					
min/max(base node)	0/36.4GB	0/72.8GB	0/36.4GB	0/109.2GB	0/72.8G
min/max (base node and six SP Expansion I/O Units)	N/A	N/A	N/A	N/A	0/946.4
Benchmarks					
SPECint_base_rate95	245/485	245/485	407/812	407/812	786/156

SPECfp_base_rate95	206/364	206/364	804/1359	804/1359	1670/32
Relative ROLTP Performance	17.9/32.8	17.9/32.8	44.0/67.7	44.0/80.0	81.7/163

***-shared memory**
****-per processor**

9076-500 - short frame available with up to 8 thin nodes or 4 wide nodes
9076-550 - tall frame available with up to 16 thin nodes, 8 wide nodes, 4 high nodes or Expansion I/O Units

Request a quote

If you would like more information about this product or any of IBM's products and solutions, contact IBM.

Back to Large Scale Servers Overview

RS/6000 SP system at a glance

Standard configuration

- Up to 128 nodes per SP system(a) (additional nodes up to 512 by special order)

Node type:	375 MHz POWER3 SMP Thin Node	375 MHz POWER3 SMP Wide Node	375 MHz POWER3 SMP High Node
Processor	375 MHz 2, 4-way POWER3-II	375 MHz 2, 4-way POWER3-II	375 MHz 4,8,12, 16-way POWER3-II
L1 cache	64KB data(a)/32KB instruction(a)	64KB data(a)/32KB instruction(a)	64KB data(a)/32KB instruction(a)
L2 cache	8MB(a)	8MB(a)	8MB(a)
RAM memory	256MB	256MB	1GB
Mem. bus width	128-bit	128-bit	2048-bit
Internal storage	None Required(b,c)	None Required(b,c)	None Required(b,c)
Disk/media bays	Two	Four	Two/Twenty-six(h)
PCI Expansion slots	Two	Ten	Five(f)/Fifty-three(h)
Nodes per frame			
tall/short	Sixteen/eight	Eight/four	Four/*
I/O Bus speeds			
	132MB/second	132 and 264MB/second (triple bus)	132 and 264MB/second (dual bus)
Adapters	Integrated Ultra2 SCSI and Ethernet (10/100 Mbps)(	Integrated Ultra2 SCSI and Ethernet (10/100 Mbps)	Integrated Ultra SCSI and Ethernet (10/100Mbps)
System expansion			
Maximum RAM	16GB	16GB	64GB
Maximum			

internal storage	36.4GB(c)	109.2GB (c)	72.8(c)/946.4GB(h)
SP expansion I/O units	N/A	N/A	Zero through Six
SP Switch and bandwidth (bidirectional)	300MB/second	300MB/second	1GB/second
SP Switch2 and bandwidth (bidirectional)	N/A	N/A	N/A

UNIX® operating system: AIX® Version 4.3.3 (unlimited user license)

System management software

- Parallel System Support Programs for AIX (PSSP) Version 3.1 and 3.2 (pre-installation optional)

System dimensions

- **Tall frame:** 75.8" H x 36.3" W x 51" D (1925 mm x 922 mm x 1295 mm), 971 to 2165 lbs (441 to 984 kg)
- **Short frame:** 49" H x 28" W x 40" D (1245 mm x 711 mm x 1015 mm); 510 to 910 lbs (232 to 414 kg)

Power requirements

- **Tall frame:** 200-240 volts (3 phase) or 380-425 volts (3 phase)
- **Short frame:** 200-240 volts (1-phase)*

Warranty 24x7, On-site for one year (limited) at no additional cost

RS/6000 SP system at a glance

Standard configuration

- Up to 128 nodes per SP system (additional nodes up to 512 by special ord er)

Node type	332 MHz SMP Thin Node	332 MHz SMP Wide Node
Processor	332 MHz 2, 4-way PowerPC 604e	332 MHz 2, 4-way PowerPC 604e
L1 cache	32KB data(a)/32KB instr(a)	32KB data(a)/32KB instr(a).
L2 cache	256KB(a)	256KB(a)
RAM memory	256MB	256MB
Mem. bus width	128-bit	128-bit
Internal storage	None required (b,c)	None required (b,c)

Disk/media bays	Two	Four
PCI Expansion slots	Two	Ten
Nodes per frame tall/short	Sixteen/eight	Eight/four
I/O Bus speeds		
	132MB/second	132 and 264MB/second (triple bus)
Adapters	Integrated SCSI-2 F/W; Ethernet (10 Mbps)	Integrated SCSI-2 F/W and Ethernet (10 Mbps)
System expansion		
Maximum RAM	3GB	3GB
Maximum internal storage	36.4GB(c)	72.8GB(c)
SP expansion I/O units	N/A	N/A
SP Switch bandwidth (bidirectional)	300MB/second	300MB/second
SP Switch2 bandwidth (bidirectional)	N/A	N/A

UNIX operating system: AIX Version 4 (unlimited user license)

System mgt. software:

- PSSP Versions 2 and 3 (pre-installation optional)

System dimensions

- **Tall frame:** 75.8" H x 36.3" W x 51" D (1925 mm x 922 mm x 1295 mm), 971 to 2164 lbs (441 to 984 kg)
- **Short frame:** 49" H x 28" W x 40" D (1245 mm x 711 mm x 1015 mm); 510 to 910 lbs (232 to 414 kg)

Power requirements

- **Tall frame:** 200-400 volts (3 phase) or 380-425 volts (3 phase)
- **Short frame:** 200-240 volts (1-phase)

Warranty: 24x7 On-site for one year (limited) at no additional cost

---

(a) Per processor
(b) Boot from external storage required if no internal storage
(c) Disks must be installed in pairs to support AIX mirroring
(d) 32-bit
(e) Two 32-bit, eight 64-bit
(f) One 32-bit, four 64-bit
(g) Seven 32-bit, three 64-bit
(h) With SP Expansion I/O Units

Performance Table

SPEC and LINPACK Performance (single processor) *

Node	Proc	MHz	SPEC int95	SPEC int_ base95	SPEC fp95	SPEC fp_ base95	LINPACK DP	LINPACK SP	LINPACK TPP
332 SMP	604e	332	14.4	14.0	12.6	12.1	115.7	585	273
375 P3 SMP	P3/II	375	24.4	22.6	50.9	47.1	409.0	--	1236.0
375 P3 SMP	P3/II	375	23.5	21.8	51.3	48.8	424	--	1208

Relative OLTP and SPEC Performance (node)

Node	Proc	MHz	Rel OLTP Perf	SPEC int_ rate95	SPEC int_ base_ rate95	SPEC fp_ rate95	SPEC fp_ base_ rate95	SPEC web96 ops/sec
332 Thin	604e/2	332	17.9	255	245	218	206	--
332 Thin	604e/4	332	32.8	501	485	389	364	6716
332 Wide	604e/2	332	17.9	255	245	218	206	--
332 Wide	604e/4	332	32.8	501	485	389	364	6716
375 P3 Thin	P3/II/2	375	44.0	438	407	844	804	--
375 P3 Thin	P3/II/4	375	67.7	875	812	1382	1359	--
375 P3 Wide	P3/II/2	375	44.0	438	407	844	804	--
375 P3 Wide	P3/II/4	375	80.0	875	812	1382	1359	--
375 P3 High	P3/II/4	375	81.7	845	786	1739	1670	--
375 P3 High	P3/II/8	375	163.7	1684	1569	3418	3290	--
375 P3 High	P3/II/12	375	242.3	2523	2345	4985	4832	--
375 P3 High	P3/II/16	375	319.3	3352	3121	6353	6202	--

IBM Corporation
Marketing Communications
Enterprise Systems Group
Route 100
Somers, NY 10589
Produced in the United States of America
7-00

respective owners.

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this publication in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features, and services available in your area.

Photographs show engineering and design models. Changes may be incorporated in production models.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply, with implementation termed by country.

Information concerning non-IBM products was obtained from the suppliers of these products. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

G221-7001-12

RS/60000 SP POWER3 SMP High Node: Next Generation in Scalable Computing

Authors:

Bob Amos, Product Manager, Sanjay Deshpande - Senior Engineer, Jeff Gruger - Senior Technical Staff Member, Mike Mayfield - Senior Technical Staff Member and Frank O'Connell - Senior Engineer.

Contents

Abstract

IBM's new POWER3 SMP High Node adds a new dimension to the RS/6000 ® SP[tm] product line with an advanced symmetric multiprocessor (SMP) node. Designed with the most advanced Deep Computing and e-business applications in mind, this node features an up to 8-way SMP, with large hardware coherent memory up to 16GB, large 4MB Level 2 cache, extensive I/O connectivity, industry leading 14.2GB/s data cross point switch bandwidth and advanced reliability design. The POWER3 SMP High Node utilizes the POWER3 microprocessor, IBM's newest 64-bit microprocessor. The design integrates the high-bandwidth and floating-point capabilities of its POWER2 predecessor into a fully scalable 64-bit SMP. This node contains the all the design features required to support large SMP nodes for future generations of RS/ 6000 SP.

This paper provides an overview of how node architecture, system technology, and advanced computer design are brought together to produce outstanding high-performance computational capabilities with high reliability.

What follows is a description of the demands of Deep Computing and e-business, followed by a description of the POWER3 SMP High Node design point, highlighting advanced performance features such as memory prefetch. That is followed by a description of the SP I/O Expansion Unit, the POWER3 microprocessor, and an overview of the advanced reliability features designed into these nodes. Next a summary of the expected performance is provided.

Demands of Deep Computing and e-business

Whether the application is engineering analysis such as molecular modeling or crash simulation or business data analysis such as decision support or enterprise resource planning, the new POWER3 SMP High Node provides the power, configurability, reliability and growth to support the most demanding requirements.

- Very large memory and storage to support the data demands of large scientific models, complex decision support, and business modeling.
- A one cycle latency L1 cache and a large low latency L2 cache to provide fast access to data and instructions required for these applications.
- Speculative data prefetching to speed processor access to data.
- Extensive I/O to support connectivity to storage and networks and the I/O bandwidth to required to handle even the most demanding applications.
- Reliability, Availability and Serviceability functions and features to support the most demanding continuous operations demands of e-business and long complex engineering and scientific analysis.
- Configuration flexibility allows processors, memory, I/O and I/O subsystems to be added incrementally as the demand grows.
- Growth for the future. Ability to upgrade to the next generation of processors, memory , and I/O in a simple manner.
- System management and support functions.

POWER3 SMP High Node Overview

The POWER3 SMP High Node complex consists of a compute node (POWER3 SMP High Node) and attached SP Expansion I/O Units. The High Node contains the processors, memory, base I/O and service processor. Up to 6 SP Expansion I/O Units can be optionally attached to the High Node to support I/O demands beyond that provided in the High Node.

The POWER3 SMP High Node is an up to 8W shared memory multiprocessor, utilizing IBM's advanced POWER3 microprocessor. Each microprocessor has a high speed, low latency, 4 MB L2 cache connected to a dedicated high speed bus running at the processor's speed.

The High Node consists of high bandwidth switched address and data paths allowing the interconnection of up to 8 POWER3 processors and of up to 4 I/O elements to two 2-port memory subsystems in a SMP configuration. The processors are organized as nodes (or cards), with 2 processors per node. The I/O elements form a separate node. The processors and I/O elements interface with the system via hardware elements called Node Controllers. The bi-directional paths between the system elements are designed to be point-to-point to allow them to switch at a high rate of 111 MHz.

Figure 1 shows the block diagram of the Central Electronics Complex (CEC) of a POWER3 SMP High Node. It is formed of 1 to 4 processor cards, a separate card carrying I/O elements, up to 4 memory cards, and an active back plane which carries the Address Switch and 2 independent memory controllers, each able to support up to 2 of the memory cards. The following table explains the functionality of each of the chips appearing in Figure 1.

Chip	Function
POWER3	Processor
NCA	Node Controller: Address. Handles address portion of all transactions issued by processors and I/O elements and coordinates data transfers among processors, memory, and I/O elements.
NCD	Node Controller: Data. Contains paths for interprocessor and processor-memory data transfers. Forms a portion of the data switch.
ASX	Address Switch.
Crossbow	Memory Controller. Controls data transfers across paths within the Arrow chips.
Arrow	Switch: Carries data to and from memory and between nodes. Forms a portion of the data switch.
Saber	Portal to I/O subsystems.
SP Fabric Adapter	Portal to inter High Node communication interconnect.

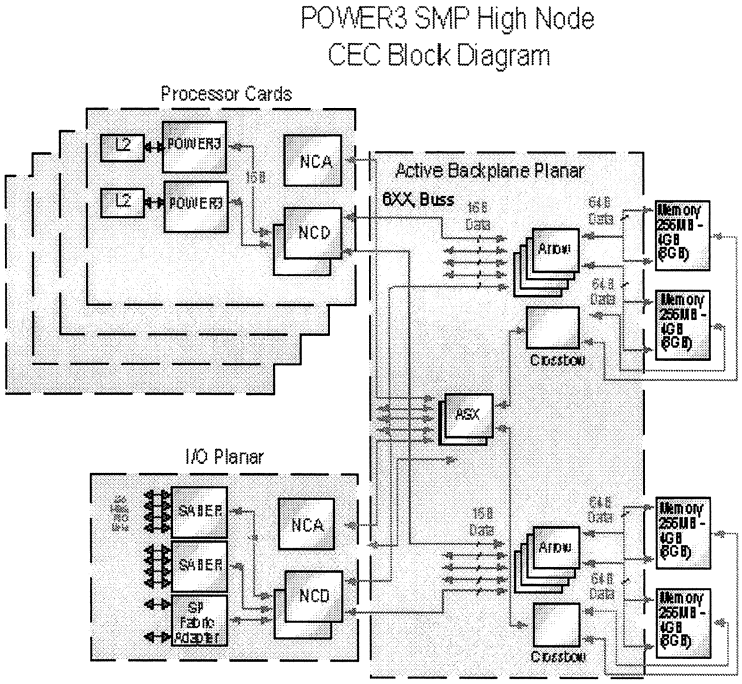

Figure 1.

Each processor card contains two POWER3 microprocessors. Each processor employs 4MB of cache on a private high speed bus running at 222 MHz. Two processors are attached to a set of node controller chips through independent point-to-point address and 16 byte wide data buses running at 111 MHz. Each processor can issue up to one address transaction per cycle to the system.

The node controller chips themselves are attached to the rest of the system via a pair of unidirectional address buses and a pair of 16 byte wide bi-directional data buses also running at 111 MHz. This translates to a data bandwidth of 3.55 GB/s per node. The node controller chips provide extensive buffering of addresses and data to improve system performance. Collectively, they also provide coherency management between all of the processors, memory and I/O using a snoopy bus protocol.

Figure 2 shows a block diagram of the physical hierarchy for system memory. Memory is packaged on up to 4 cards connected to the 4 memory ports in the data cross point switch. Memory cards utilize 128MB SDRAM modules. Each card can contain up to 4GB of memory or up to 16GB for the compute node. 1GB of memory comes in the base system design. Memory can be added in 1GB increments with the

addition of cards and pluggable industry standard DIMMs. This provides a highly configurable and upgradable offering, which will grow with the application requirements.

Each memory subsystem is capable of accepting an address transaction every cycle at 111 MHz. Each memory data port is 64 bytes wide and operates at 55.5 MHz, which translates to a delivered bandwidth of 14.2 GB/s.

Figure 3 shows a block diagram of the I/O subsystem. The figure shows Sabers which serve as portals to disk and communication subsystems, and an SP Fabric Adapter which allows a compute node to connect to the SP Switch to form a multinode system.

POWER3 SMP High Node and SP Expansion I/O Unit
I/O Topology

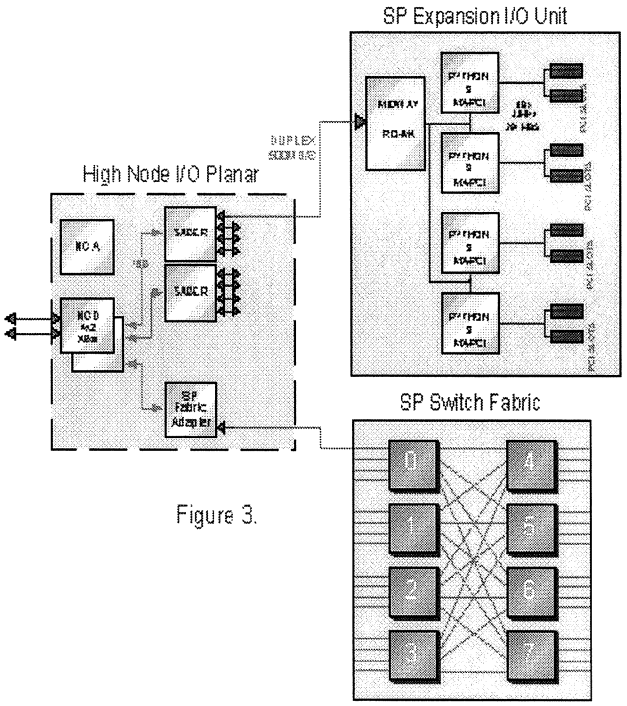

Figure 3.

The High Node I/O consists of four 4-bi6t PCI slots and one 32-bit PCI slots operating at 33 MHz. It provides integrated Ultra SCSI, 10/100 Ethernet® parallel port, and three serial ports. In addition, there are 6 remote I/O (RIO) connections. These connections allow SP Expansion I/O Expansion Units to be connected to the High Node providing incremental growth for applications requiring more I/O connectivity. The RIO ports operate full duplex at 250MB/s each direction with cables up to 15 meters long. This provides outstanding I/O bandwidth to each RIO expansion node. With 6 SP Expansion I/O Units connected, a user can obtain a node complex with 53 PCI slots and 26 DASD bays. This combination of High Node and SP Expansion I/O Units into a node complex provides the highest I/O configuration yet on the RS/6000 SP system.

The POWER3 SMP High Node is designed with growth in mind. The node controller chips will accommodate up to 4 processors per set allowing future growth to a 16-way SMP. System busses are designed to operate up to 125 MHz providing the extra memory bandwidth required for the future. Memory is designed to accommodate increases in DIMM size to allow growth of the system memory in the future.

Upgrades to future High Nodes will be as simple as replacing the pluggable CPU cards.

SP Expansion I/O Unit Overview

The SP Expansion I/O Unit is offered in a thin SP node format. It contains eight 64-bit PCI slots running at 33 MHz and 4 "hot swappable" DASD bays that will accept either SSA or SCSI drives in a variety of sizes.

The node is designed for reliability and easy service. It supports N+1 power and cooling, "hot swappable" mirrored disks drives, redundant links to the POWER3 SMP High Node, and a node supervisor which monitors the health of the system. As an added bonus "hot plug" PCI will be added to these units in the future to support concurrent service of the I/O adapters.

Up to 6 expansion units can be attached to the High Node offering a tremendous compute complex with extensive I/O connectivity and performance.

A large variety of PCI adapters are available for the POWER3 SMP High Node complex, which allow extensive network and storage connectivity and performance. The user may also select from a large variety of IBM's external storage subsystems, which can be attached to this node complex.

POWER3 Microprocessor

The POWER3 microprocessor continues the POWER2 architecture tradition of bringing real solutions to IBM RS/6000 customers' high-performance compute needs, while adding 64-bit addressability, double-word integer operations, and SMP support. To satisfy compute intensive requirements, the POWER3 design contains a highly superscalar core which comprises eight execution units capable of sustaining an execution rate of four instructions per cycle, a 32KB instruction cache, 64KB data cache, and high bandwidth, independent interfaces to the L2 cache and system memory. A block diagram of the POWER3 processor is show in Figure 4.

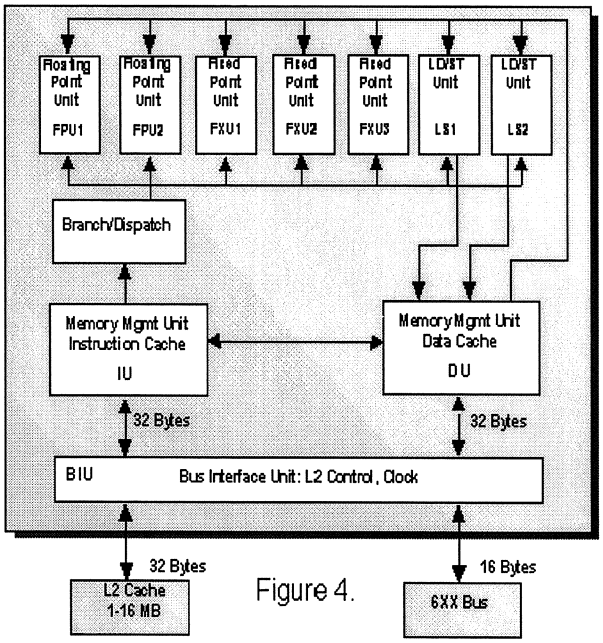

Figure 4.

Two of the three fixed-point units (FXUs) are single cycle execution for the bulk of the integer arithmetic instructions. The third unit executes the multi-cycle integer instructions such as multiply and divide.

The two floating-point units (FPUs) are fully independent, each containing dedicated hardware for square root and divide routines as well as fused multiply-add instruction execution. The FPUs are fully pipelined with three cycle latency, single cycle throughput.

Two load store units provide the data to sustain four floating-point operations per cycle. A 16-entry store queue buffer prevents stores from stalling the machine while loads are being performed. Loads are also executed speculatively, improving data throughput.

The branch execution unit employs dynamic branch prediction, with four pending predicted branches supported. The branch target address cache contains 256 entries (128 by 2 way associative), and the branch history table has 2048 entries.

The instructions are speculatively executed with a

unique register renaming scheme that involves a total of 64 virtual rename registers (32 fixed and 32 floating-point), and a total of 40 physical rename registers actually implemented (16 fixed point and 24 floating-point).

The on-board BIU contains the interface logic supporting up to 16Mbytes L2, 6XX system bus protocols, and dedicated hardware to reduce latency to memory.

Containing 15 million transistors, the POWER3 processor die is shown in Figure 5. It is manufactured in IBM's 0.25 micron hybrid CMOS 6S2 technology, with five levels of interconnect metallurgy.

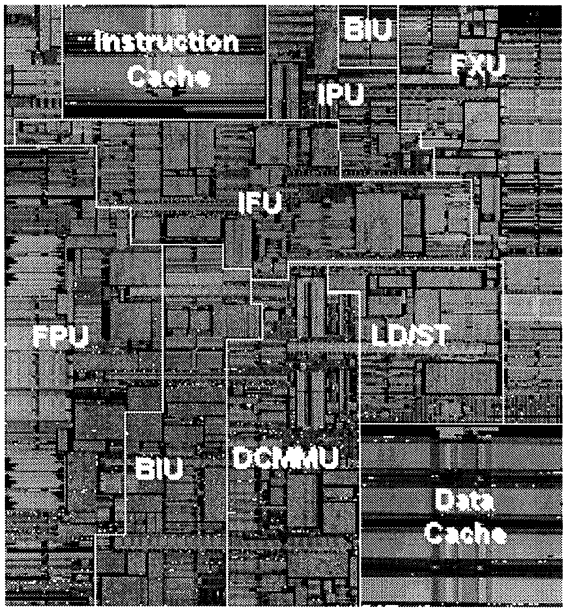

Figure 5.

To ensure that potentially needed data and instructions are available to keep the core from stalling, the POWER3 processor designers invested in two key latency reduction techniques.

First, all caches are non-blocking. The instruction cache supports two outstanding misses, and the data cache supports up to four.

Second, the POWER3 processor implements sequential instruction and data access detection algorithms in hardware, which permit the prefetch of cache lines to closer levels of the memory hierarchy. The POWER3 processor prefetches up to four separate data streams with a depth of two lines for each stream. Prefetching data significantly reduces the apparent memory latency and improves the data bandwidth, thus

increasing performance and processor utilization. Instructions are prefetched into the L1 cache up to one sequential line ahead of the line currently being accessed on the predicted path, which often eliminates delays with instruction fetches from memory.

Reliability, Availability and Serviceability Features

The POWER3 SMP High Node and SP Expansion I/O Units were designed with high RAS capabilities in mind. From memory to I/O to processors this node represents significant advancements in system RAS.

Memory and L2 cache are protected with single bit correct, double bit detect ECC. L1 cache is parity protected. Memory incorporates scrubbing and supports continued operation with a full memory chip failure as further protection.

Both the POWER3 SMP High Nodes and the SP Expansion I/O Unit incorporate N+1 power and N+1 cooling. The Expansion Units also provide redundant links to the High Node.

Disk mirroring is standard on both the High Node and the Expansion Unit. It provides redundant storage allowing continued operation in the presence of a disk failure.

The Expansion Unit provides "hot swappable" capability for both SCSI and SSA disks, which enables maintenance concurrent with node operation. This reduces down time for maintenance.

CPU and Memory RepeatGuard capability is provided. This function checks for excessive soft or hard fails at boot time and deconfigures a faulty memory bank or processor for deferred repairs. The soft fail threshold is user selectable.

Performance

The POWER3 SMP High Node's superior performance characteristics are the result of clearly articulated design objectives gleaned from in-depth knowledge of challenging customer applications, and years of experience in designing the RS/6000 family of POWER and POWER2 processors. Consistent with its POWER and POWER2 heritage, the POWER3 SMP High Node distinguishes itself from its competitors in its ability to deliver main memory bandwidth, a crucial attribute to many applications on the computing frontier. These applications can achieve a quantum leap in performance on the High Node and, in turn, provide new scientific insights and competitive advantages to their users.

The POWER3 SMP High Node excels in data delivery; indeed this was its most significant and most challenging design objective. The most sophisticated features of its design are all integral pieces of its data delivery system: dual load/store execution units, an interleaved data cache, multiple-outstanding cache-miss support, wide data paths to memory and its L2 cache, hardware data prefetch, node controllers, and a non-blocking data cross point switch.

Figure 6 shows the history of POWER, POWER2, and POWER3 STREAM performance to date, including the POWER3 SMP High Node, in the STREAM benchmark, a standard measurement of sustained memory bandwidth. Note that the POWER3 SMP High Node provides a significant jump in performance over the POWER3 SMP Thin and Wide Nodes, delivering an average bandwidth to each of its eight processors a rate that exceeds all processors and systems in its class. As in the past, this new RS/6000 system once again sets the industry standard for sustained memory bandwidth in systems of its class.

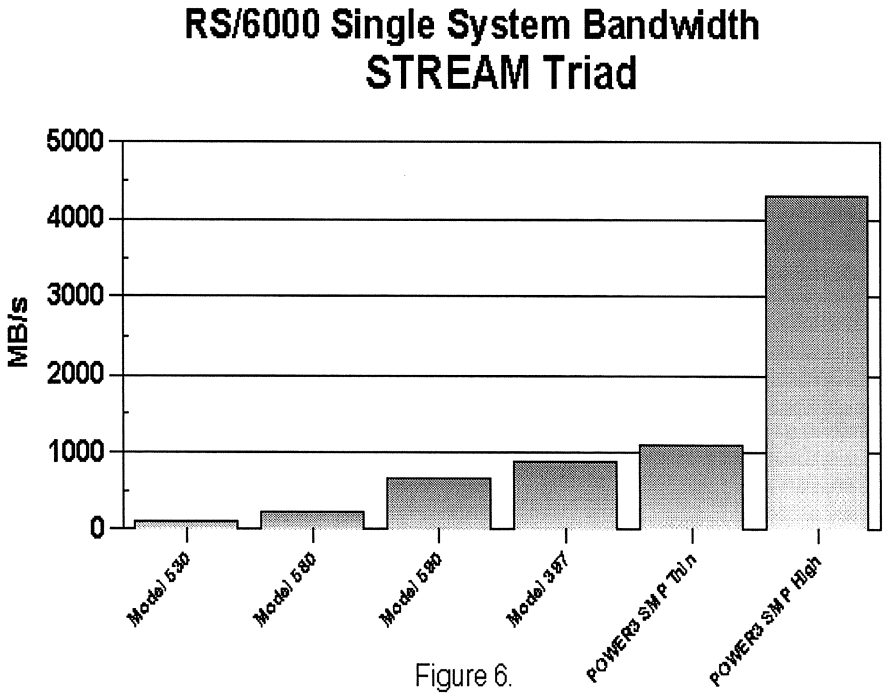

Figure 6.

Complimenting the POWER3 SMP High Node's excellent memory bandwidth are floating-point and fixed-point execution units that can sustain remarkable rates of

computation. Figure 7 shows the LINPACK 1000 performance for the POWER3 High Node and a sample of its product pedigree. This benchmark measures the average performance to factorize and solve a dense matrix of rank 1000 and is representative of large amounts of work found in many important applications. Just as with sustained bandwidth, the High Node greatly improves upon a distinguished record of floating-point performance.

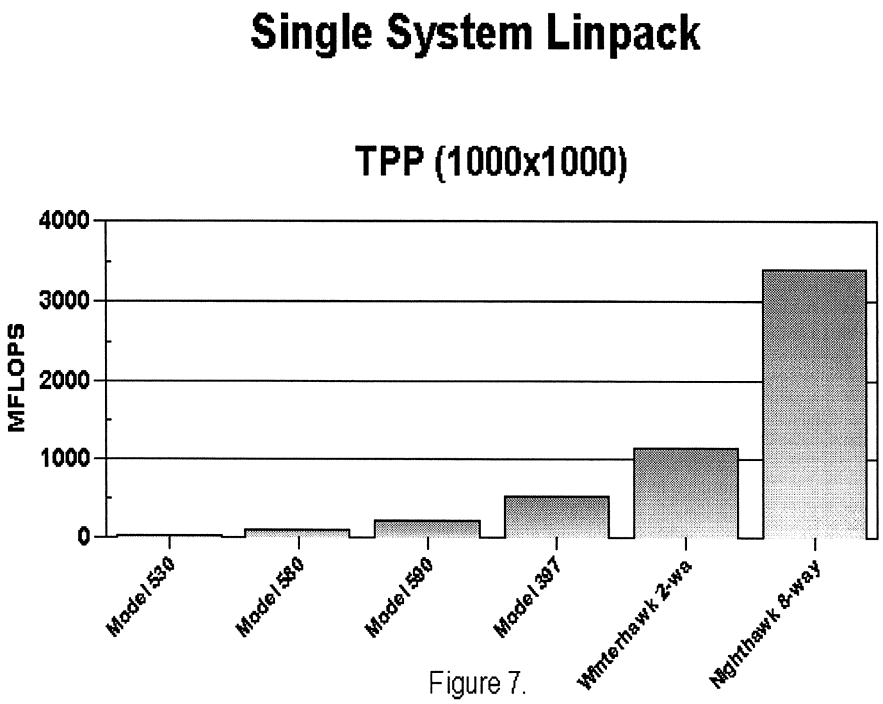

Figure 7.

With its intelligent and powerful design, the POWER3 SMP High Node provides the base for a new family of High Node products.

Summary

The POWER3 SMP High Node and the SP Expansion I/O Unit represent an exciting new offering for the RS/6000 SP product. It introduces large SMP with large memory and extensive I/O. These nodes are designed for user configurability and grown. They are the first in a line of large SMPs.

The POWER3 SMP High Node and the SP Expansion I/O Unit are offered as nodes within the RS/6000 SP system and are also offered in a new family of servers the Model T70,

intended for engineering and scientific applications.

Biographies

Bob Amos is a Product Manager. Sanjay Deshpande is a Senior Engineer. Jeff Gruger is a Senior Technical Staff Member. Mike Mayfield is a Senior Technical Staff Member. Frank O'Connell is a Senior Engineer. All authors are members of IBM Server Group in Austin, Texas.

Any performance data contained in this document was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Notes on Benchmarks and Values

The benchmarks and values shown here were derived using particular, well configured, development-level computer systems. Unless otherwise indicated for a system, the values were derived using 32-bit applications and external cache, if external cache is supported on the system. All benchmark values are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the following on the Web:

TPC http://www.tpc.org

GPC http://www.spec.org/gpc

SPEC http://www.spec.org

Pro/E http://www.proe.com

Linpack http://www.netlib.no/netlib/benchmark/performance.ps

Notesbench Mail http://www.notesbench.org

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX 4.3, and AIX XL FORTRAN V6.1. with optimization where the compilers were used in the benchmark tests.

The Linpack benchmark reflects the performance of the microprocessor, memory architecture, and compiler of the tested system.

- LINPACK TPP (Toward Peak Performance) - n=1,000 is the array size. The results are measured in MFLOPS.

Linpak Benchmarks from:
http://performance.netlib.org/performance/html/PDSreports.html

STREAM is a program which J. McCalpin of University of Virginia developed and measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels. The results reported in this paper are the fastest TRIAD program using a uniprocessor machine. STREAM Benchmark from:
http://www.cs.virginia.edu/stream/

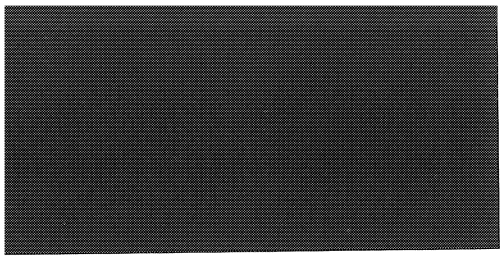

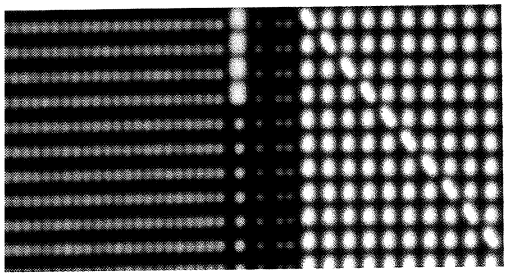

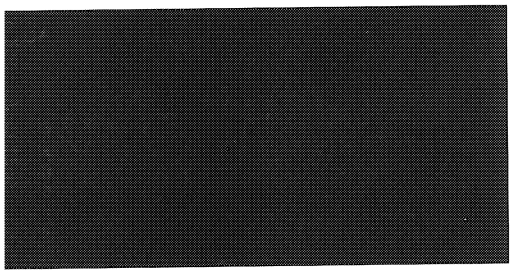

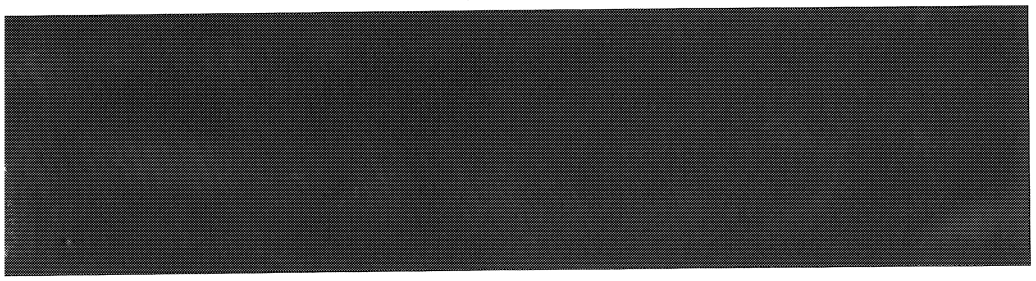

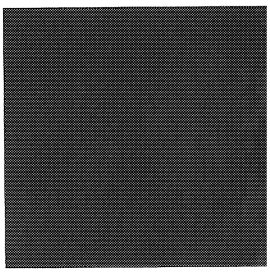

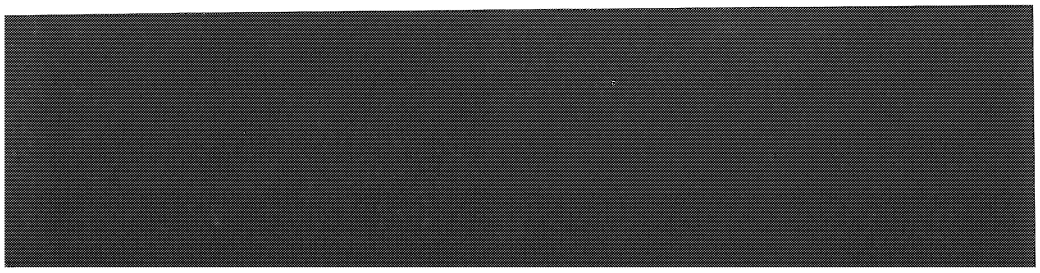

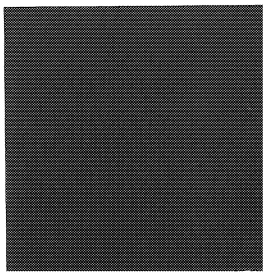

IBM

IBM Enterprise Systems Connection Architecture

The path for growth

System resources increase without impairing operations unduly

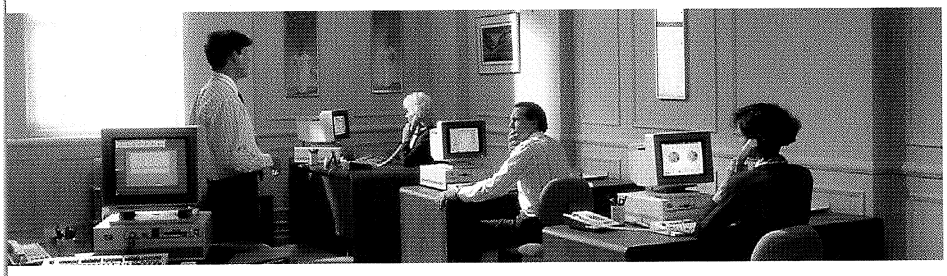

Overcoming limitations with ESCON

Data centre managers are struggling with increasingly difficult problems. As the demands for system availability and growth change dynamically, the constraints of networking technology stay the same. That leaves data centre managers to contend with:

- The weight, bulk and limited range of bus-and-tag copper cabling, even though their computer rooms are at bursting point
- The maintenance of individual and static connections between each processor channel and each I/O control device, even during continual reconfiguration requirements
- The need to exchange data among computing systems that use different protocols, even when connectivity is costly or difficult
- The mandate to protect investment in existing hardware and software, while at the same time seeking new application solutions to meet business needs.

What data centre managers need are immediate solutions. Solutions that allow their enterprise data centre resources to grow, while causing minimal disruption to users. Solutions that help reduce the costs and complexities of managing resources dynamically and automatically, while ensuring continuous availability. And solutions that provide a migration path to new applications and future technologies while maximising return on existing investments.

Fibre optics allow completely new configurations and eliminate many limitations of copper cabling

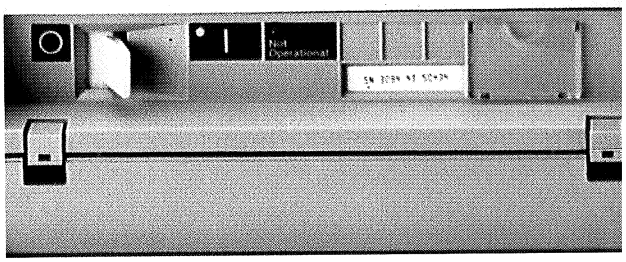

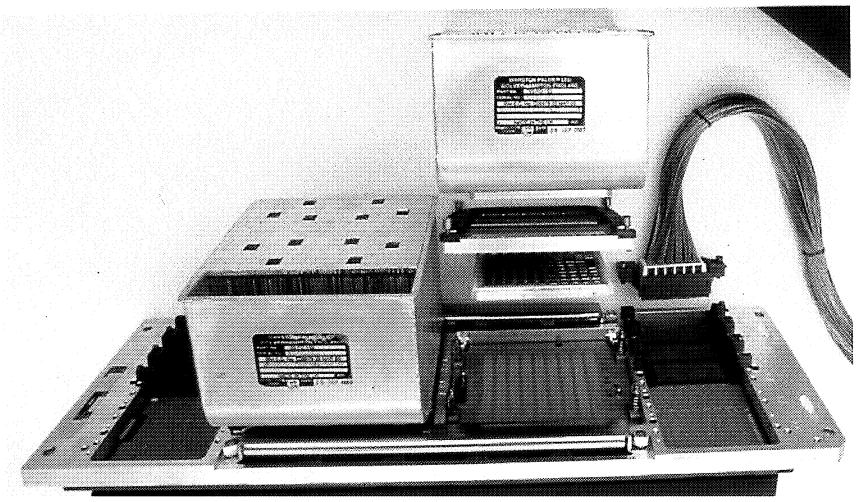

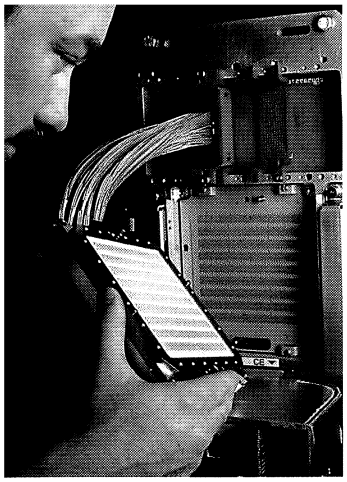

Bringing new technology within reach

IBM Enterprise Systems Connection Architecture (ESCON) offers comprehensive solutions to the data centre manager's problems. IBM ESCON eliminates the limitations of bulky cabling, affording more flexibility for unhampered growth. It helps control spiralling staff costs and safeguards investments in existing I/O resources. And, just as important, it provides new opportunities to broaden the scope of your data centre – and manage it better.

The IBM ESCON offering currently comprises:

- Processors with optical fibre serial channels – Enterprise Systems Architecture/390 (ESA/390) serial channels on the new IBM Enterprise System/9000 (ES/9000)* processors, and on IBM Enterprise System/3090* Model J processors; distances between channels and control units have been increased from 122 metres to 3 kilometres as a direct link, up to 9 kilometres with ESCON converters described below
- Optical interface control units – IBM 3490 Magnetic Tape Subsystem and IBM 3990 Storage Control for attachment to ESA/390 serial channel via adapters
- Dynamic connectivity products – IBM ES Connection Directors (ESCD) for dynamically variable interconnections among multiple optical interface channels and control units
- Network interconnection products – IBM 3172 Interconnection Controller and IBM 3174 Establishment Controller for access to other common or standard communications protocols.
- Time synchronisation and remote control – IBM Sysplex Timer for synchronising time-of-day clocks among processors and ESCON Monitor, a central facility for controlling power sequencing, monitoring environmental conditions and initiating event processing.

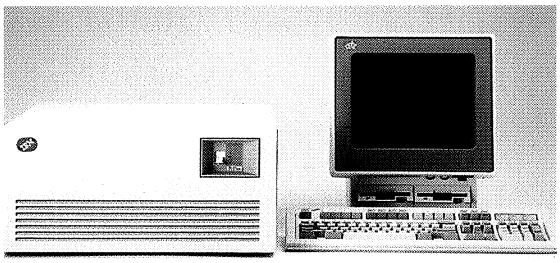

The new ESCON devices bring the advantage of optical fibre serial channels within reach

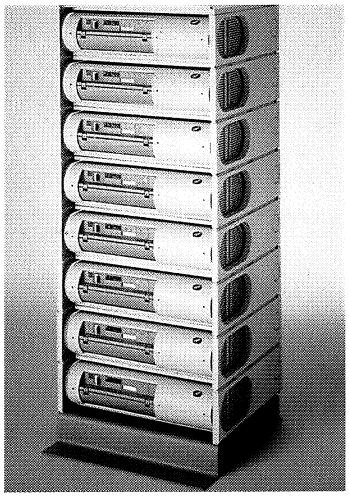

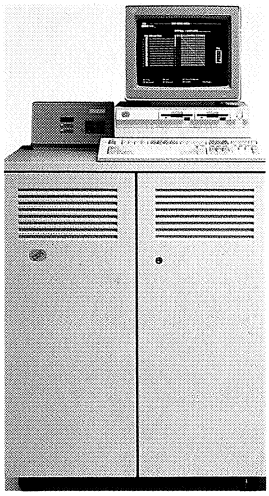

- Software – IBM ESA system control program (operating system) required in each processor, IBM ESCON Manager (ESCM) for environments containing two or more ESCDs, IBM ESCON Supervisor (ESCS) and IBM ESCON Fault Isolation
- Fibre optic products and services – jumper cables, distribution panels, connectors and IBM services to assist in installation
- IBM Enterprise Systems Connection Converters (ESCC) which provide migration paths for parallel processor channels and I/O control units.

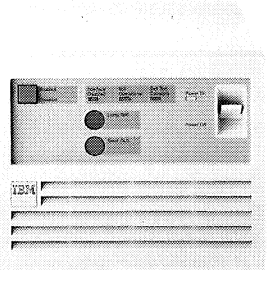

Extend the range of your data centre without taxing your machine room

Gain cabling flexibility and space in your computer complex
IBM ESCON provides a new foundation for data centre growth. Built on the principles of serial fibre optic technology, ESCON introduces a new I/O interface – a fibre optic channel that supports a range of dynamic connectivity products and service offerings.

The fibre optic channel comprises a light-emitting diode (LED) light source, a flexible silica glass filament as transmission medium, and transmitters and receivers that convert light to and from an electrical signal. With its smaller size and weight and its greater physical flexibility over copper cabling, fibre optic cable dramatically reduces the amount of space required for computer room cabling.

Increase logical connectivity with dynamic switching
An integral part of ESCON is the ESCON Director (ESCD), an any-to-any switching device that provides dynamic point-to-point switching between fibre optic channels and attached control devices. The new any-to-any topology allows the same connectivity to be achieved with fewer channel links, thus saving channel hardware.

The ESCDs receive and transmit data so you can greatly increase the distance – up to 9 kilometres – between your host computers and remote I/O units, depending on device type. The ESCD Model 2 is a floor-standing unit that comes with 28 channel ports as standard – and you can upgrade easily in 4-port increments to a total of 60 ports. For networks with fewer connectivity requirements, the table-top ESCD Model 1 offers many of the Model 2's features in an 8-port configuration that's upgradeable to 16 ports.

Effectively monitor, control and automate data centre connections
ESCON implementation enables ESCDs to establish and remove dynamic connections between host computers and control units. IBM ESCON Manager (ESCM) is a software tool that monitors and controls ESCD connections – even for ESCDs located remotely from the ESCM host. It provides a central point of control to help manage enterprise data processing resources that are interconnected by ESCDs.

IBM ESCM:
- Provides host control and coordination of one or more ESCDs
- Automates planning and operating tasks

- Allows network reconfiguration to be non-disruptive
- Helps towards continuous availability
- Facilitates automated network operations when used with IBM NetView* network management software.

Easily migrate from copper cabling to fibre
ESCON lets you take advantage of the benefits of the serial optical environment gradually. With IBM ESCON Converters (ESCCs), you can migrate to fibre as your needs change. The ESCC Model 1 enables your existing control units and their I/O devices with System/370 interfaces to attach to new IBM ES/9000 or ES/3090 Model J via ESCON channels.

And the ESCC Model 2 lets you attach IBM 3990 Model 2 and 3 Storage Control Units with ESCON adaptors to existing processors with parallel channels.

Gain power and performance as you upgrade with serial optical technology

Boost processing power with new system platforms
IBM ESCON products work with the new IBM ES/9000, and most ES/3090 J models. These processors have optional serial channels so you can begin using ESCON products immediately.

Move up to a new I/O interface architecture
ESCON is supported by IBM Enterprise Systems Architecture (ESA), which spans all the ES/9000 processors to provide extended addressing, expanded storage and a new I/O subsystem architecture for System/370 users. The I/O architecture supports a non-synchronous protocol and serial data transfer – giving faster and more efficient data transfers between the processor and I/O device.

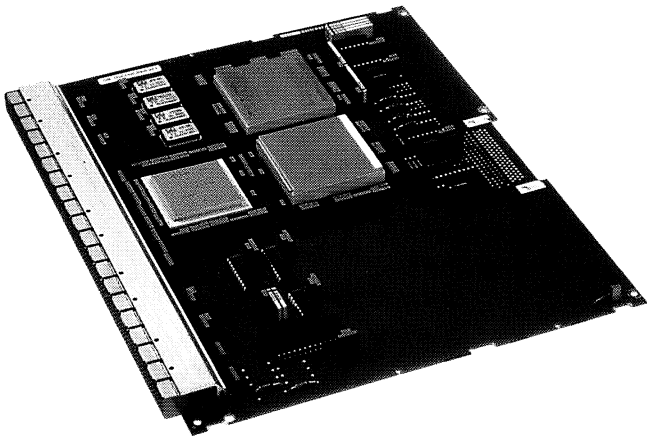

Use fewer channels to serve a greater number of control units
With ESCON, you can attach one ES/3090 Model J or ES/9000 serial channel to multiple I/O device control units – or one control unit to multiple serial optical channels. That way, you can achieve the required connectivity with far fewer links and much less cable – and you'll gain valuable space under the machine room floor.

Manage processors from a single point of control
To manage several IBM processors and systems in the ESCON environment, IBM developed the concept of the Sysplex (SYStem comPLEX). With its innovative Sysplex Timer, Sysplex provides a means to synchronise processor system clocks with an external time source. So you get consistent time-stamping in transaction processing applications and in distributed environments. And with the Cross System Coupling Facility (XCF), Sysplex provides a foundation for managing multisystem communications by effectively linking separate systems.

Add connection capabilities without sacrificing productivity
ESCON helps you manage changes, even when you add users or devices, or increase the volume of information processing. Flexible reconfiguration capabilities let you expand your data centre without affecting the productivity of your users.

I/O connections can be made without a power disruption, so new hardware can be added immediately. What's more, backup components can be integrated into your facility without requiring you to move units or cables, so installation time can be greatly reduced. That makes ESCON highly cost-effective and helps you improve the overall efficiency and usefulness of your network.

Extend your network to multivendor environments
ESCON provides an interface between ESA/390 or System/370 systems and other industry standard protocols. With the IBM 3172 Interconnect Controller, channel connection to a wide range of IBM and other industry-standard Local Area Networks (LANs) is possible.

The 3172 also allows data processing centres to be linked by high-speed communications lines with interconnection of remote mainframes over wide area networks.

Extend the scope of your data centre with advanced applications and services

Customise applications to meet your business objectives
ESCON can establish a platform for developing advanced applications, such as distributed printing and client/server processing together with operational enhancements such as disaster recovery. Its fibre optic channels and dynamic connectivity provide the speed and efficiency necessary to support even the most demanding applications, like high-speed image processing.

Implement network solutions quickly and efficiently
IBM Customer Engineering (CE) offers the knowledge and experience to help you implement ESCON quickly, and efficiently. CE offers connectivity services that range from network consultation and design to turnkey installation and net-work management.

Find out more
If you would like more information on how IBM ESCON can extend the reach and scope of your data centre, contact your local IBM marketing representative.

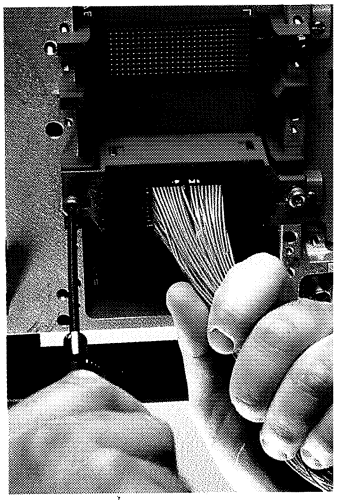

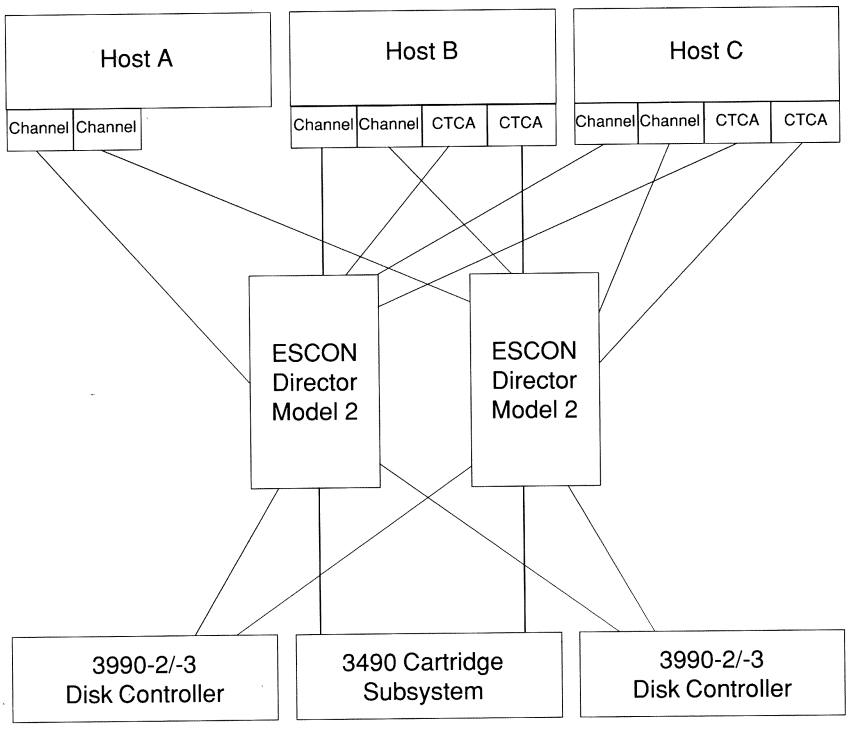

GU20-0086

IBM United Kingdom Limited
National Enquiry Centre
389 Chiswick High Road
London W4 4AL
Telephone 081-747 0747

Registered in England: No. 741598
Registered Office: PO Box 41
North Harbour, Portsmouth
Hampshire, PO6 3AU

Printed in England by
Cedar Colour Limited
Chandlers Ford, Hants

IBM

IBM

Enterprise System/9000™

Tailoring Technology for Leading Edge Solutions

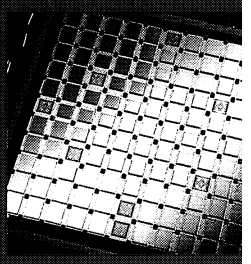

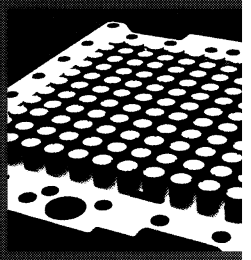

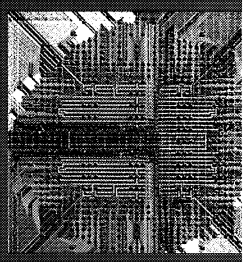

ES/9000™ family of processors

The broadest range of system solutions in the industry

Price competitive solutions

The extraordinary array of processor price and performance alternatives in ES/9000 systems is made possible by IBM's unique achievement in applying a wide range of semiconductor technology, creative packaging techniques, and innovative cooling methods to the implementation of the industry's largest family of compatible shared systems.

S/390 central electronics complex

In addition to supporting world class functions for data management (Data In Memory, Data Facility Storage Management Subsystem, dynamically allocated channels, ES Connection Architecture, global networking), all IBM ES/9000 processors, from smallest to largest, share certain hardware implementation characteristics.

All ES/9000 systems include multiple instruction execution units supporting arithmetic and logical operations, I/O activity, system management functions, and other tasks.

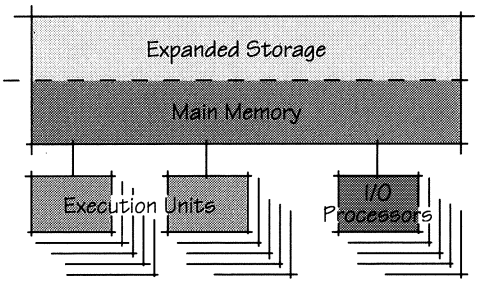

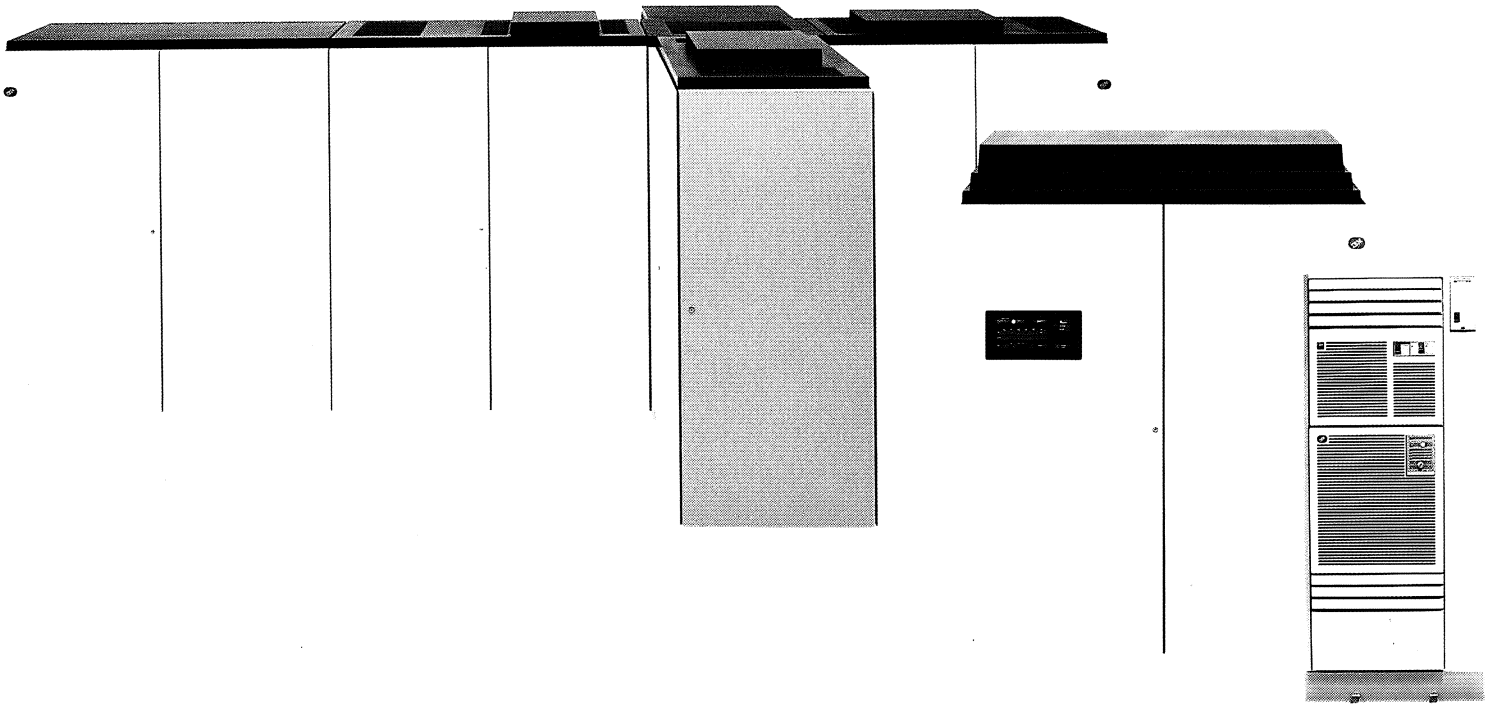

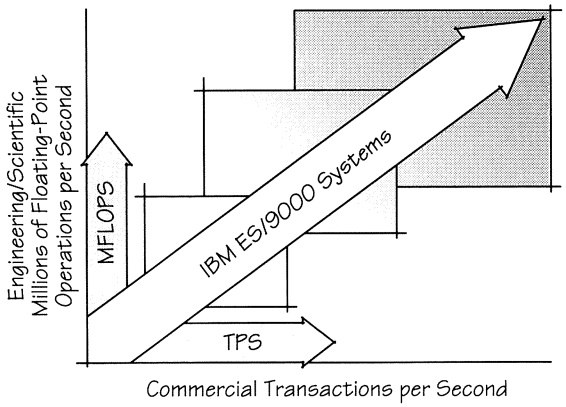

PERFORMANCE FOR DATA MANAGEMENT AND ANALYSIS
Performance measurements on the new ES/9000 systems confirm their outstanding capability both in commercial and engineering/scientific performance. IBM ES/9000 systems perform equally well in *diverse* workload environments characterized by heavy demands for I/O processing and high intensity numeric operations. ES/9000 systems offer extensive granularity, with many options for distinctive solutions: packaging, processor speeds, storage size, channel capacity, and vector facilities.

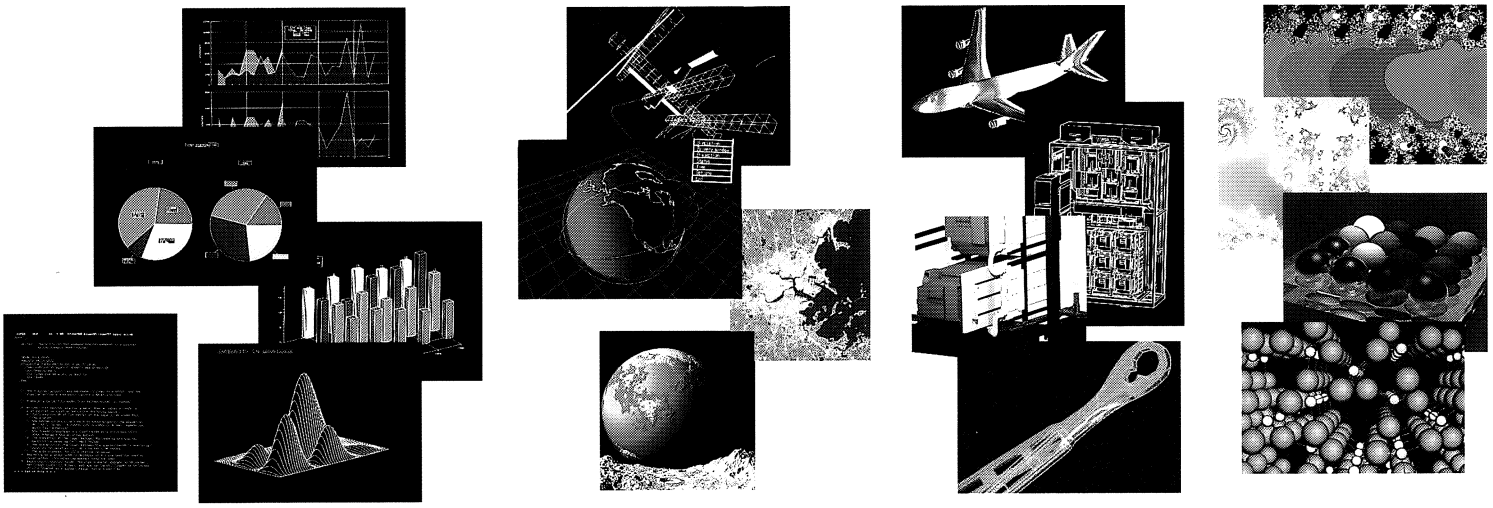

Enterprise Systems Architecture/390™

A standard foundation for custom solutions

IBM ESA/390™defines a comprehensive framework of leading industry mainframe system facilities - allowing developers to build products tailored to specific user requirements.

Today, IBM S/390™supports enriched functions enabling solutions for some of the major technical challenges of the 1990s: network integration and management, data management, application integration, and security.

ESA/390 provides a single architectural platform empowering dynamic new hardware technologies and software facilities in powerful processing systems. New processors, new functions like Sysplex Timer and Integrated Cryptographic Feature, and new architectures for Enterprise Systems Connection and multi-systems management are defined in the scope of ESA/390 solutions.

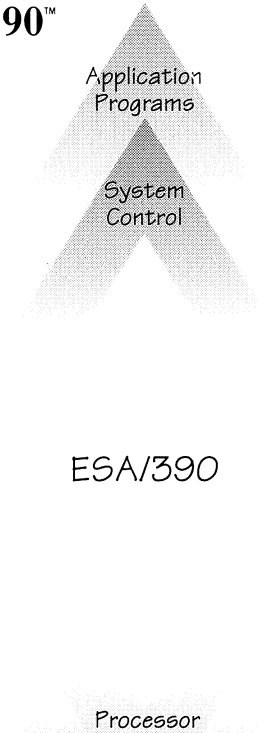

SYSTEM/390: ENHANCING AN INVESTMENT IN S/370 APPLICATION SOLUTIONS

Supporting the range of S/370 compatible application solutions (for S/370, SAA,™ and AIX™), S/390 systems can participate in global networks using efficient instruction processing (for both data management and data analysis tasks) and IBM's broadest bandwidth I/O subsystems.

Operating systems on ES/9000 models feature tailored support for host interactive computing, transaction processing, data management, and mixed workloads.

ESA/390 systems are sized to satisfy user requirements: from small work group systems to the largest data centers managing global networks. The base architecture is the same. The implementation may vary. All are built to a similar standard of excellence.

A wide range of technology, packaging, and upgrade alternatives have been exercised by ES/9000 development engineers, scaled against the full spectrum of end user service level requirements and cost considerations. Systems and upgrade paths have been customized for current and emerging user needs.

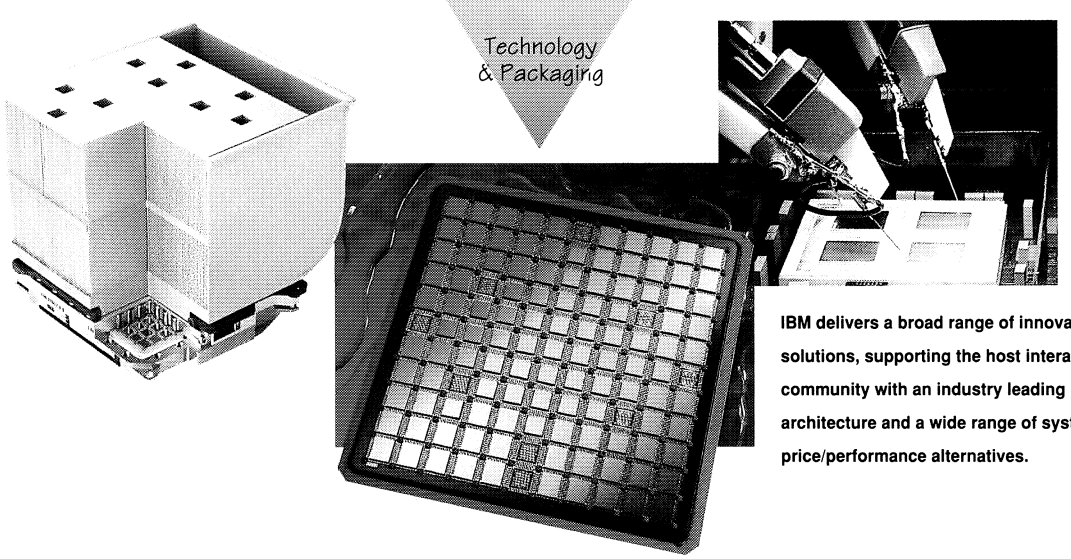

IBM delivers a broad range of innovative technology solutions, supporting the host interactive user community with an industry leading mainframe architecture and a wide range of system price/performance alternatives.

IBM locations throughout Europe are making important contributions to the technical excellence of the new ES/9000. Advanced processes and fine-tuned expertise are applied to building processors, logic chips, control units and telecommunications devices, and to designing advanced operating systems and software - key elements of this leading edge solution.

Among the facilities supporting the ES/9000 are Boeblingen, Mainz and Sindelfingen in Germany; Montpellier, La Gaude, Bordeaux and Essonnes in France; Valencia, Spain; Hursley, England; and Jarfalla, Sweden.

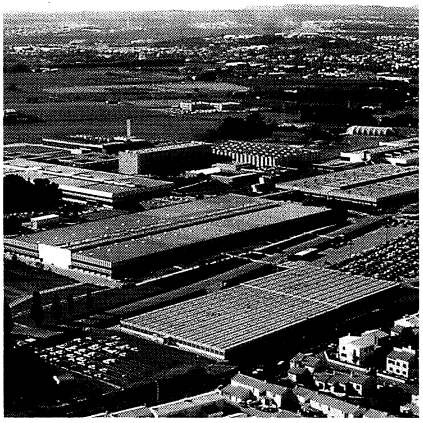

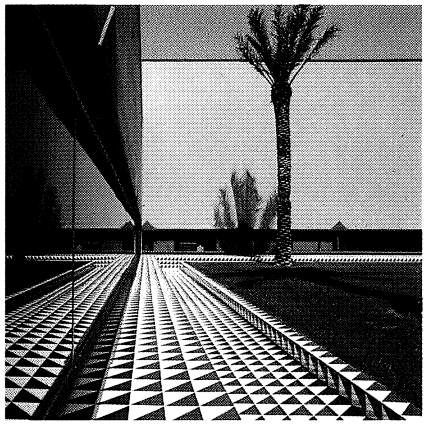

Significant development input is contributed by the Boeblingen laboratory, while Montpellier and Valencia are the primary manufacturing locations for Europe.

IBM

Front Cover:
The cover design includes an assortment of IBM-developed advanced technology components. Using customized circuitry and highly versatile hardware packaging, IBM ES/9000 systems include extensive capacity for technical upgrade and extended system life.

G511-1395-0

Printed in Denmark by Bonde's

IBM Eurocoordination SA
au capital de 2.700.000 F
Siège social Tour Pascal
22 Route de la Demi-Lune
92075 Puteaux
RCS Nanterre B 304 538 192
France

IBM World Trade
Americas Group
Town of Mount Pleasant
Route 9 North Tarrytown
New York 10591 USA

IBM World Trade Asia
Pacific Group
IBM Kamiyacho Building
3-9 Toranomon 4-chome
Minato-ku
Tokyo 105 Japan

IBM

IBM Enterprise System/9000

Introducing a new computing family

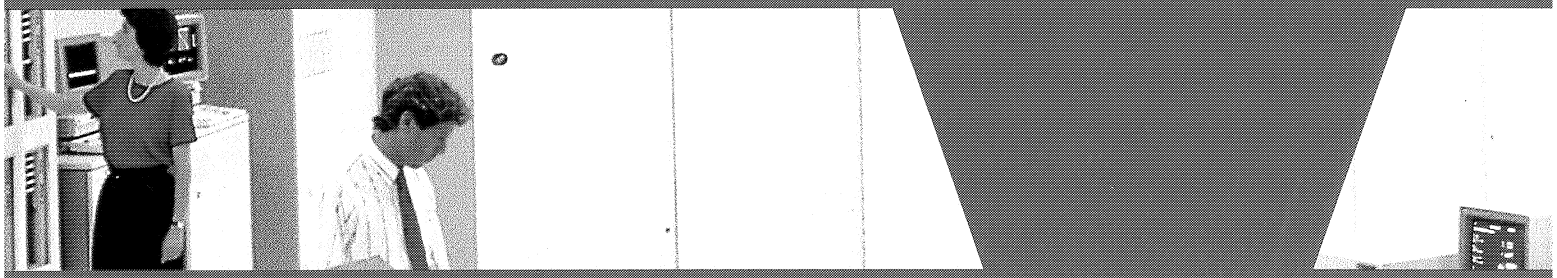

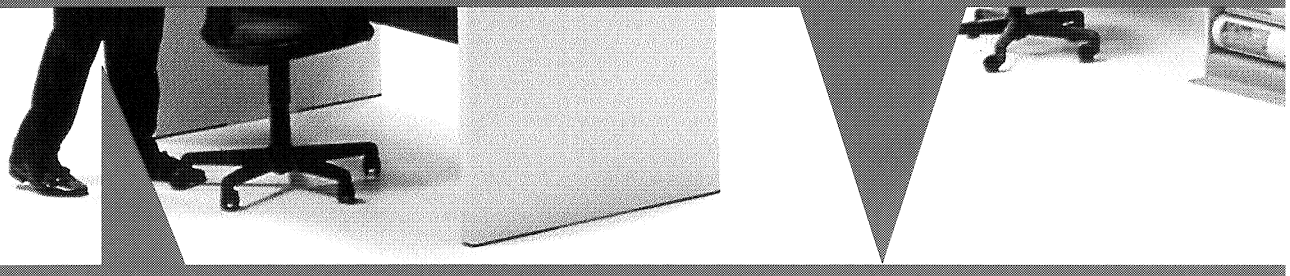

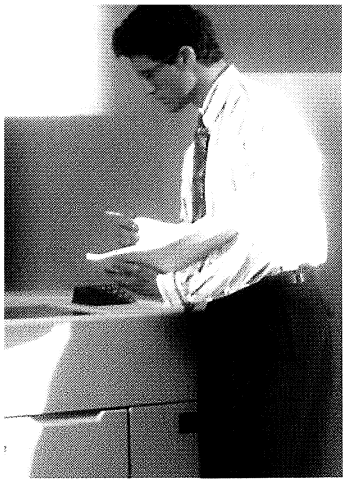

A new computing family offering a wide range of systems

The IBM Enterprise System/9000 (ES/9000)* family is based on the newest proven technologies, and offers a comprehensive computing range to meet today's needs for high power and availability.

The new IBM ES/9000 processor models are part of IBM's System/390 computing platform, which includes enhanced system hardware and new operating system software. The powerful ES/9000 family incorporates innovations such as four-megabit chips and high-speed channel technology to generate new possibilities in computing. This technology complements broadened system and network management capabilities, and enables strategic applications to meet your specific business requirements.

The ES/9000 family's 18 models – successors to the proven IBM Enterprise System families ES/9370, ES/4381, and ES/3090* – provide a solid foundation for increased computing performance and function. The family's performance range spans a 100-fold growth from the compact, rack-mounted Model 120 to the powerful Model 900.

System/390 introduces new platforms and functions, including a new channel subsystem architecture, IBM Enterprise System Connection Architecture (ESCON), which implements a fibre-optic channel connection permitting new application opportunities in some cases, and shared applications in others.

All models in this powerful family share the following common architecture, features and functions:

- IBM Enterprise Systems Architecture/390 (ESA/390)
- Processor Resource/Systems Manager (PR/SM)
- ESCON architecture
- Multisystem management support
- Common operating systems
- Enhanced price performance.

This rich set of common functions offers a solid foundation for your enterprise solutions today, as well as for tomorrow's growth requirements.

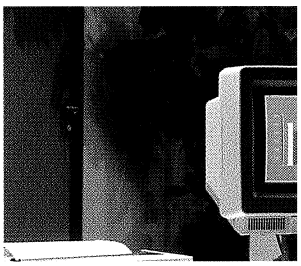

The demands of customers for high-performance computing are met by the new functions of the IBM ES/9000 family

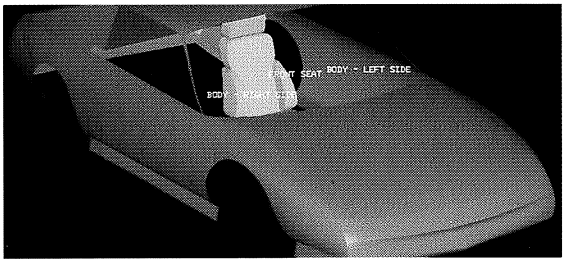

. . . the engineer develops a vehicle design . . .

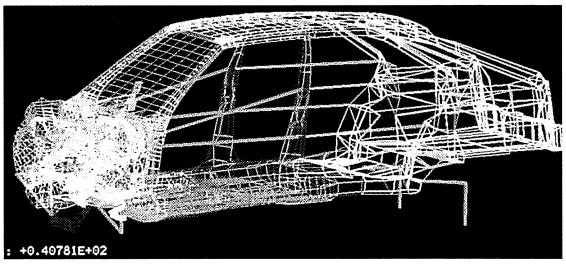

. . . the protection of passengers is assessed with the aid of crash test simulation . . .

ESA/390: gateway to enterprise-wide applications

The Enterprise Systems Architecture/390 (ESA/390), which underlies the ES/9000 family, offers a broad range of capacity as well as growth potential. To operate your system, IBM has enriched and enhanced the established System/370 operating systems with the introduction of Virtual Storage Extended/Enterprise System Architecture (VSE/ESA), Virtual Machine/Enterprise System Architecture* (VM/ESA) and the enhancement to Multiple Virtual Storage/ Enterprise System Architecture (MVS/ESA). These three ESA/390 systems operate on all System/390 processors.

The newly announced VSE/ESA operating system will offer improved support for increased real and virtual memory, providing more application flexibility. This operating system's strong affinity with MVS/ESA helps to make application development and use easier.

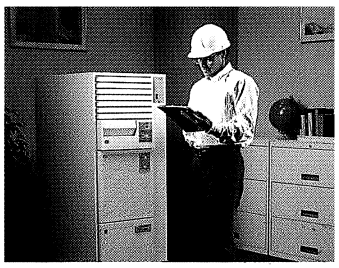

The VM/ESA operating system enhances and makes uniform three current IBM VM products VM/SP, VM/SP HPO and VM/XA. Building on their capabilities, VM/ESA allows a single VM to be run across the ES/9000 family. This uniformity is complemented by new levels of performance and throughput as well as by improved systems management and user productivity.

A new version of the MVS/ESA operating system adds dramatic new function over previous versions, including:

- Enhanced management of multiple systems (Sysplex)
- The ability to add or remove dynamically I/O resources from systems in use (Dynamic Reconfiguration Management)
- Single point of control
- Enablement of Systems Application Architecture (SAA)*-based distributed and cooperative processing
- Increased system availability

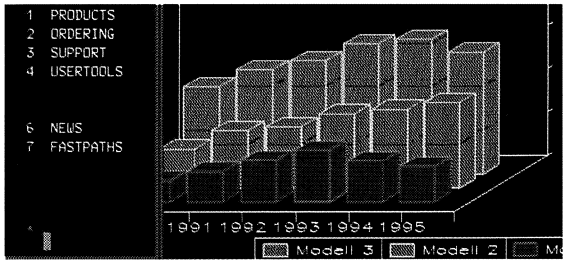

. . . market analysis assesses potential sales . . .

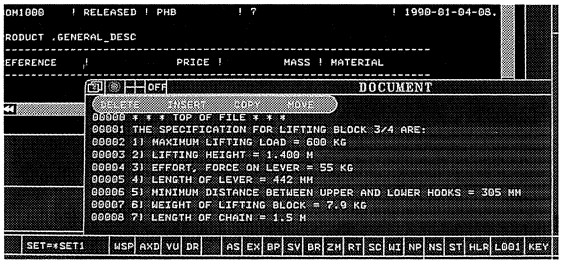

. . . and the production programme is planned

Logical partitioning for configuration flexibility and economy

The ability to partition a physical processor into logical images enables your enterprise to consolidate workloads, enhancing the flexibility of your processor configuration. You can also use logical partitioning to create migration platforms, allowing regular workloads to run concurrently with your test or conversion efforts. This hardware feature, known as Processor Resource/Systems Manager (PR/SM) is now standard on all ES/9000 models. PR/SM allows VSE/ESA to operate on all ES/9000 processors, and it also enables you to run multiple preferred guest systems under VM/ESA, VM/XA and VM/SP.

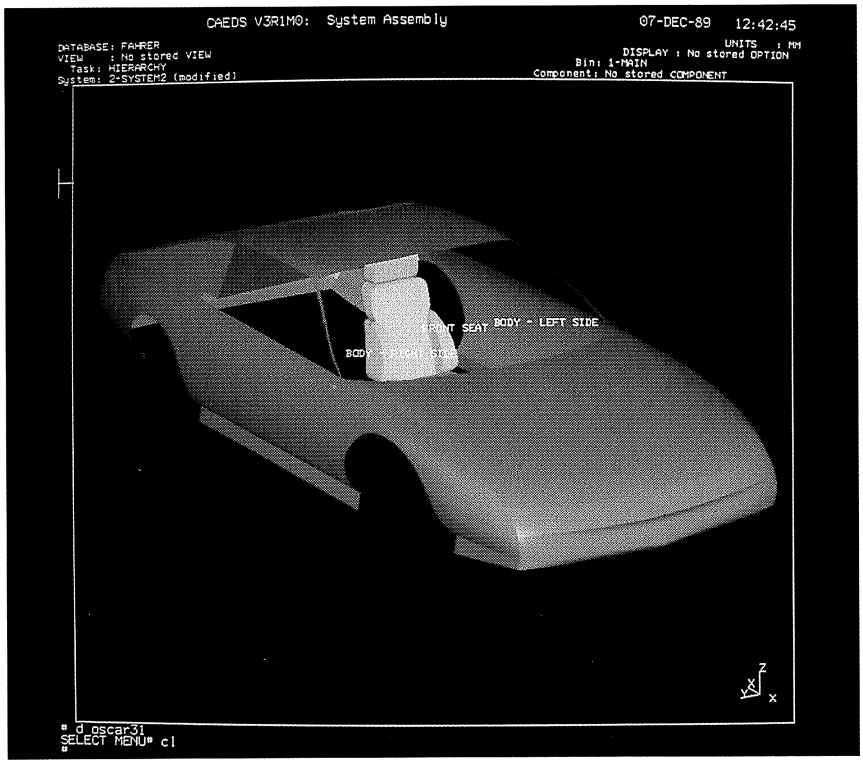

The integration of CAEDS* in enterprise communications allows the engineer to take factors such as standards, parts lists and costs into consideration when constructing a new model of a car

For applications with high availability requirements, PR/SM can be used in conjunction with MVS/ESA Version 4 and the OPC/ESA program product to provide automated restart of a failed production partition. Manual delays can be eliminated or greatly reduced, providing significant potential availability improvements.

Connections to facilitate communication within - and beyond - the enterprise

The IBM ES/9000 family enables your enterprise to communicate through standard protocols with both workstations and other host systems, linking your various enterprise workgroups and facilitating communication with your business partners. You can, for example, electronically distribute your product catalogue to your dealers or distributors.

The available protocols complementing your System/390's networking capability include TCP/IP, OSI, and IBM 3172 controller support for Ethernet (parallel channel environment), as well as IBM Token-Ring local area network protocols. In addition, multi-vendor devices can be attached to host processors by means of the Fibre Distributed Data Interface.

Complementing your enhanced networking capability, the IBM NetView* product family increases your network control through enhanced automated procedures and flexible dynamic network configuration capabilities. As a result, you can increase system availability by reducing the likelihood of system service disruptions.

Such system products and functions provide growth path alternatives through distributed systems, and address the cooperative processing needs of individual work environments.

All System/390 family models share a common channel implementation and offer both serial and parallel channel capabilities. As a result, a wide range of IBM I/O devices can be channel-attached to System/390 systems, and phased migration can be achieved.

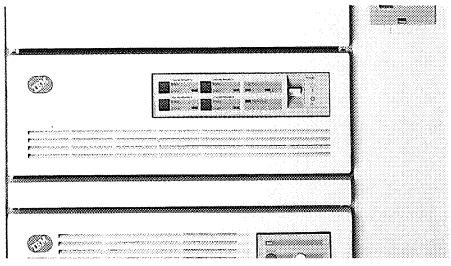

Advanced channel architecture

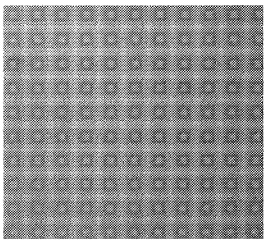

The new ESCON channels, which transmit data at up to 10 MB/sec., extend the boundaries of the traditional information processing centre to give you far greater flexibility in systems configuration. The new ESCON architecture, a new medium and storage interface, and a broad range of products and services support our ESCON channels. These new serial channels significantly increase the distance across which devices can be connected.

Dynamically and at high speed, a new ESCON Director (ESCD) switching device connects any host to its dependent devices or to other hosts. Together with the new fibre-optic cabling, the ESCD enables you to reduce cable bulk and physical line requirements significantly, simplifying system planning and installation as well as eliminating many space constraints.

The ESCDs enable you to choose where you wish to locate your resources – where space within your office complex or work environment permits – up to distances of 9 kilometres (5.4 miles) with two ESCDs. The current limit is 122 metres (400 feet). To accommodate these longer distances, the ESCD can be remotely managed via host-system software called the ESCON Manager (ESCM).

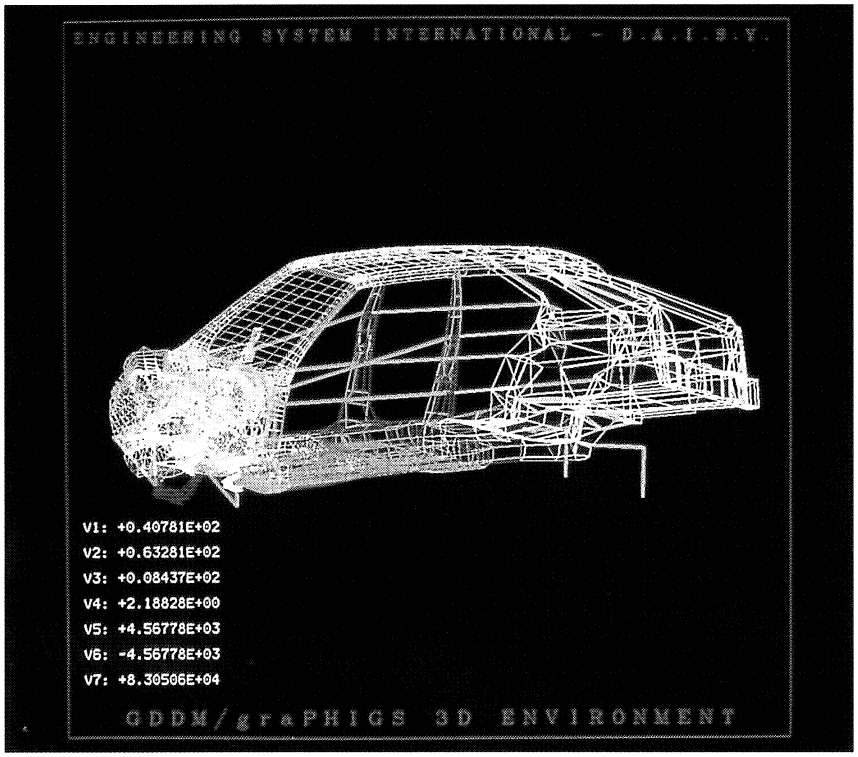

As early as in the design phase, the effectiveness of passenger protection can be assessed using simulation

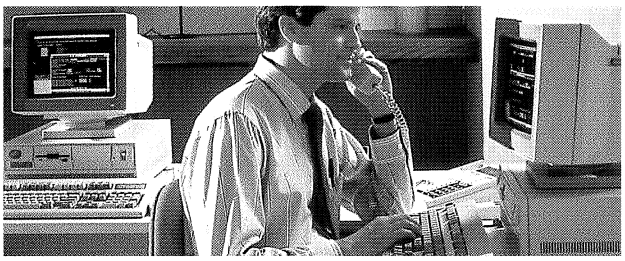

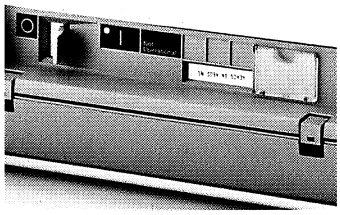

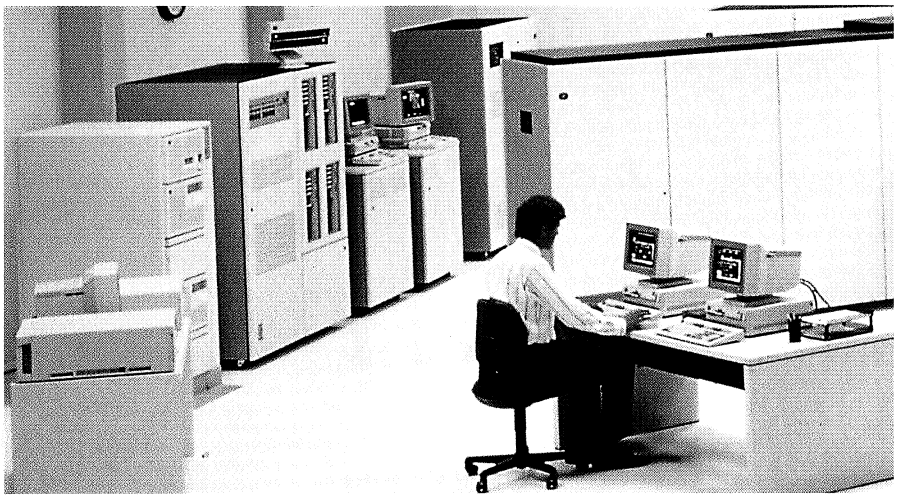

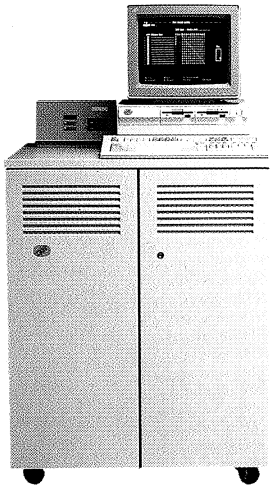

ESCON's fibre optic technology breaks through traditional computing centre barriers, opening up a wealth of new possibilities for configurations

Fibre-optic technology, the physical medium used in this advanced method of channel connectivity, allows data to be transmitted much faster than today's 4.5 MB/sec. data rate – up to a maximum of 10 MB/sec. In addition, you can locate remote functions and disaster-backup equipment far more flexibly than before. ESCON connections also enable you to dynamically add or remove I/O resources while your system remains available to users.

Co-existence and migration are simplified between ESCON and existing System/370 environments by using IBM ESCON Converters (ESCC). The ESCC Model 1 allows you to attach a serial channel to a non-serial control unit, enabling orderly conversion while protecting your existing investment. The ESCC Model 2 permits an ESCON-capable 3990 Model 2 or 3 Storage Control Unit to attach to a System/370 parallel channel.

IBM's latest storage devices, the IBM 3490 Magnetic Tape Subsystem and 3990 Model 2 and 3 Storage Control Units can all be upgraded to operate with ESCON adapters. Two new models of the IBM 3174 Establishment Controller are also capable of attaching to ESCON channels. ESCON channel attachment is also available to the 3172 Interconnect Controller, providing a remote channel-to-channel communication function.

Working together with MVS/ESA SP Version 4.2 Dynamic Reconfiguration Management, the new ESCON channels can significantly reduce the system outages formerly encountered when systems were reconfigured or new devices added.

Additional computing support for a broad range of needs

IBM makes available vector computing options on selected models, and additional systems software to address specific application environments and requirements.

ESCON Configuration Example

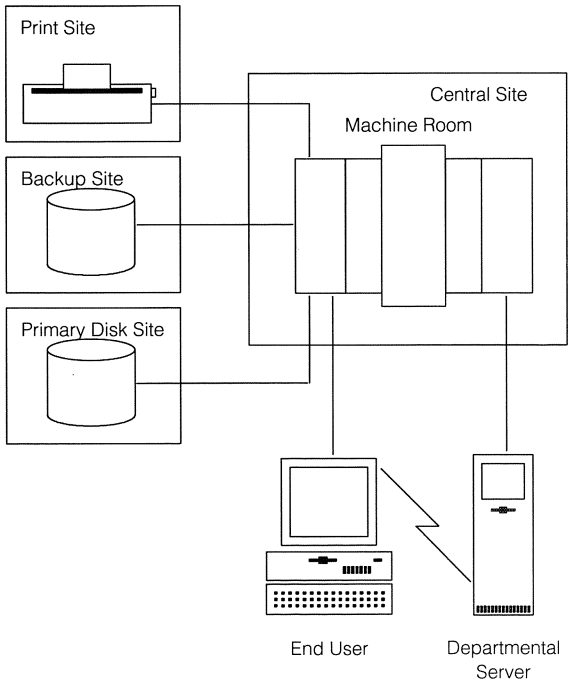

Only common architectures can provide the right basis for trend-setting and innovative future applications

Vector Facility available for more processor models

Information systems users who employ mathematical modelling in their computing applications often seek added processing power for improved response time. This is true in business areas as diverse as financial modelling, finite-element computing and simulation.

For example a manufacturer may use a computer-aided design (CAD) application to develop a model and simulate its operation. The IBM ES/9000 family includes an optional integrated vector capability on selected models, which results in faster feedback on test results. Such rapid feedback can allow engineers to increase their productivity, as well as improve product quality, by using simulation to redesign critical parts when necessary.

Centralised management of local and remote systems

The ES/9000 family models take steps toward your information system requirements for 'lights out' operation and effective control over multiple concurrent systems. The family models offer features and functions for remote location management and coordinated timing for multiple systems. NetView Version 2, IBM's platform for automation, and related systems management products enable you to automate your enterprise and manage multiple systems from a single point of control.

Selected system models can now be remotely power-controlled and environmentally monitored with the ESCON Monitor System (ESCMS), eliminating the need to activate and deactivate power manually at a remote location. Instead, the power can be turned on and off automatically or manually from remote locations. Once a remote system is powered on, the Target System Control Facility can be used in conjunction with NetView to initialise and control multiple target systems remotely from a single point of control.

Multisystem configurations often require that internal time-of-day clock settings be closely coordinated. Now, a cluster of selected systems can use the IBM Sysplex Timer to ensure that accurate time reference data is shared.

New security options protect sensitive information assets

Practically every enterprise needs to protect the sensitive information it transmits and stores. The water-cooled models of the ES/9000 family provide the optional Integrated Cryptographic Feature. Implemented in tamper-resistant hardware, this new security option allows sensitive information to be encrypted by the sending processor and decrypted by the receiving processor. This optional feature provides high performance and can be used in conjunction with the IBM Resource Access Control Facility (RACF) security platform.

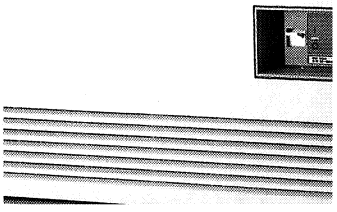

This cryptographic option improves productivity in a cost-effective manner by reducing the complexity of security administration. The result is a more secure distribution capability for sensitive information.

You can also use other IBM-provided security functions to address the need for access control or software 'virus' detection. Whichever IBM security option you choose, you can extend the range and reach of your business information services without compromising the integrity or security of your vital business information assets.

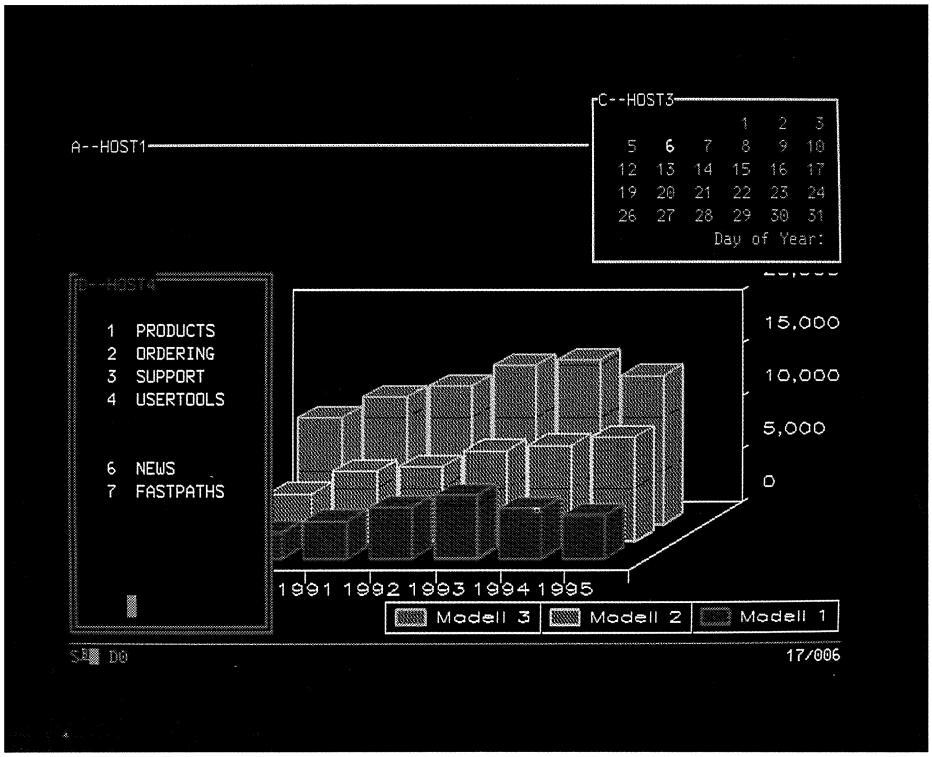

The sales plan displays the results of market analysis into windows support parallel work processes and allow data to be copied from window to window

Environmental efficiency to enhance availability
To help address your environmental concerns, all ES/9000 family models operate on a common 50/60 Hertz power supply.

Allowing for protection of your vital business information, battery- powered backup is provided for models 120 to 170 of the ES/9000 family. If the external power sourse is interrupted, the battery backup facility provides power to the central processor until external power is restored (for a maximum of five minutes).

System availability is improved by a newly designed power system on models 330, 340 and 500 through to 900 of the ES/9000 family. This enhancement provides an additional power capability to reduce system interruptions due to power supply failures.

System software alternatives for added processing power.
The following system software products use or extend part or all of the IBM ES/9000 family functions.

The IBM enhanced AIX* operating system meets user needs for an open system that supports industry standards. Running as a guest under VM, including VM/ESA, AIX/370 version 1, release 2 is supported on Model 190 upwards. ES/9000 increases the number of AIX users who can be supported, and makes additional storage, memory and processing power available. AIX provides consistent user and application interfaces, as well as specific features and functions that build on the advantages of the ES/9000 system capabilities.

Distributed Processing Program Executive/370 (DPPX/370) offers a strong, economical platform for unattended, distributed transaction processing solutions. A fully integrated system, DPPX/370 distributes functions needed to operate and manage remote production environments efficiently. DPPX/370 is supported on Models 120 through 170 of the ES/9000 family. It also operates on a logically partitioned processor.

Transaction Processing Facility
The IBM Transaction Processing Facility (TPF), the high-availability, high-performance system, supports real-time, transaction-driven applications. Its specialised architecture can optimise system efficiency, reliability and responsiveness for data communications and data base processing. IBM TPF Version 3 Release 1 is supported on ES/9000 family Models 190 and above.

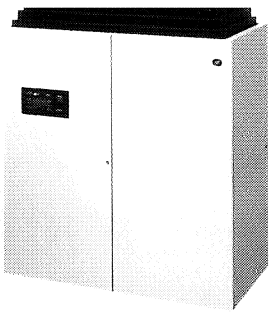

MUMPS
The ANSI-standard Massachusetts General Hospital Utility Multi-Programming System (MUMPS) is implemented in the high-level programming language of IBM MUMPS/VM, a versatile programming system that allows users to develop base application programs. IBM MUMPS/VM also includes a comprehensive data base management facility and a flexible operating system. IBM MUMPS/VM is supported on all Models of the ES/9000 family.

Current operating systems
Support for non-ESA operating systems is dependent on model. For specific release levels and model support, check with your IBM marketing team.

Technology for the '90s - and beyond
IBM technological leadership rests on the belief that customer satisfaction is the basis for success. The ES/9000 family delivers, in a single architecture, the broadest performance range available in the information processing industry. Underlying the ES/9000 family's design are multiple functional elements. These performance-optimised units work independently and in parallel for efficient operation.

Other technological elements in the family include:
- High-density packaging for fewer components
- Thermal Conduction Modules (TCMs)
- Bipolar and CMOS chips
- One-megabit and four-megabit storage chips
- Dynamic reconfiguration capabilities
- Remote service facility
- Extensive error-checking and correction enhancements
- Automatic problem-analysis invocation

Versatile models to accommodate your business growth
The ES/9000 family encompasses 18 models, including powerful, economical entry systems as well as high-performance ones. The growth path within the family includes both rack-and frame-mounted, air-cooled and water-cooled models:
- Air-cooled rack-mounted models 120, 130, 150 and 170
- Air-cooled frame models 190, 210, 260, 320, 440 and 480
- Water-cooled models 330, 340 and 500 to 900

Air-cooled, rack-mounted models
The four powerful, economical entry systems in the ES/9000 family are based on a new complementary metal oxide semiconductor (CMOS) technology offering high density (from 40,000 to 80,000 circuits per chip) and low power consumption. The increased chip density reduces the number of chips required comparable to current systems, such as the IBM ES/9370. This increased density enables greater functionality in rack-mounted systems, as well as improved reliability.

The circuit switching time on the chips in these models has been reduced to less than a nanosecond. This allows the cycle time to be reduced, creating an impressive speed improvement over the capability of the IBM ES/9370.

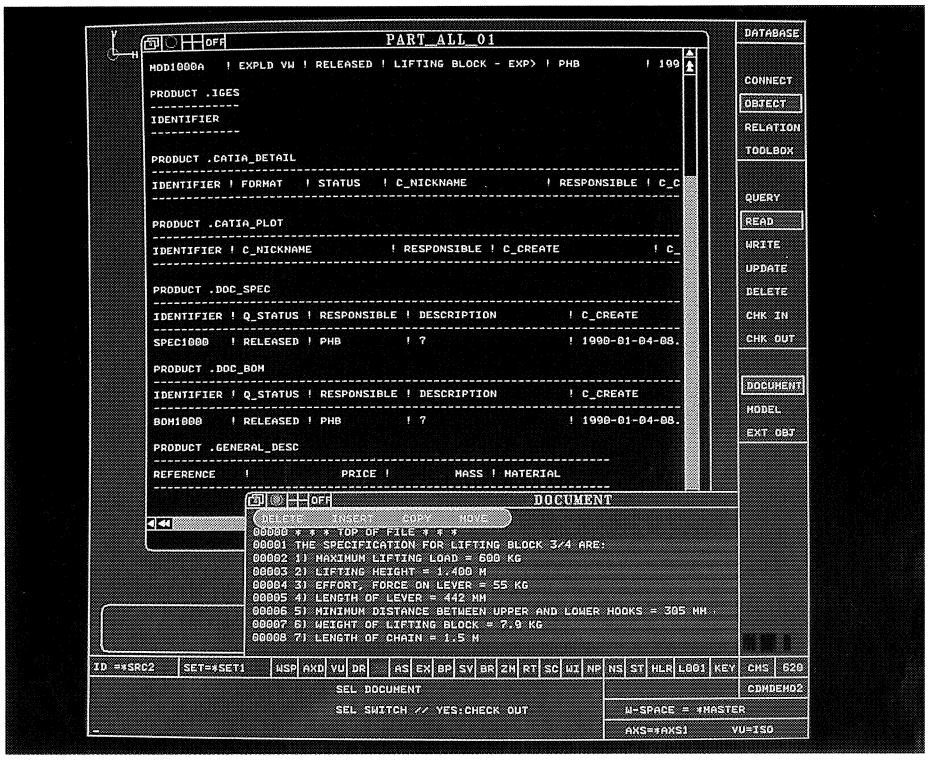

The planning of production capacity incorporates capacity scheduling and optimum use of production facilities

Rack-mounted ES/9000 models contain a compact processor and channel design. This compactness provides new reliability and serviceability and eases your upgrade from IBM ES/9370 processors. A two-sided circuit board used in the rack-mounted design reduces the number of cables needed by integrating the hardware.

Air-cooled frame models
The six ES/9000 air-cooled frame models provide performance, function and technology that was previously available from IBM only in water-cooled processors. These models pack the power of a small IBM ES/3090 mainframe into a footprint the size of an IBM ES/4381 processor.

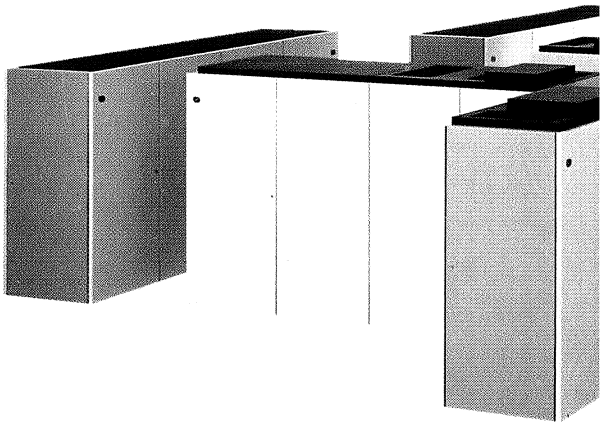

The functions of IBM's ES/9000 family of processors

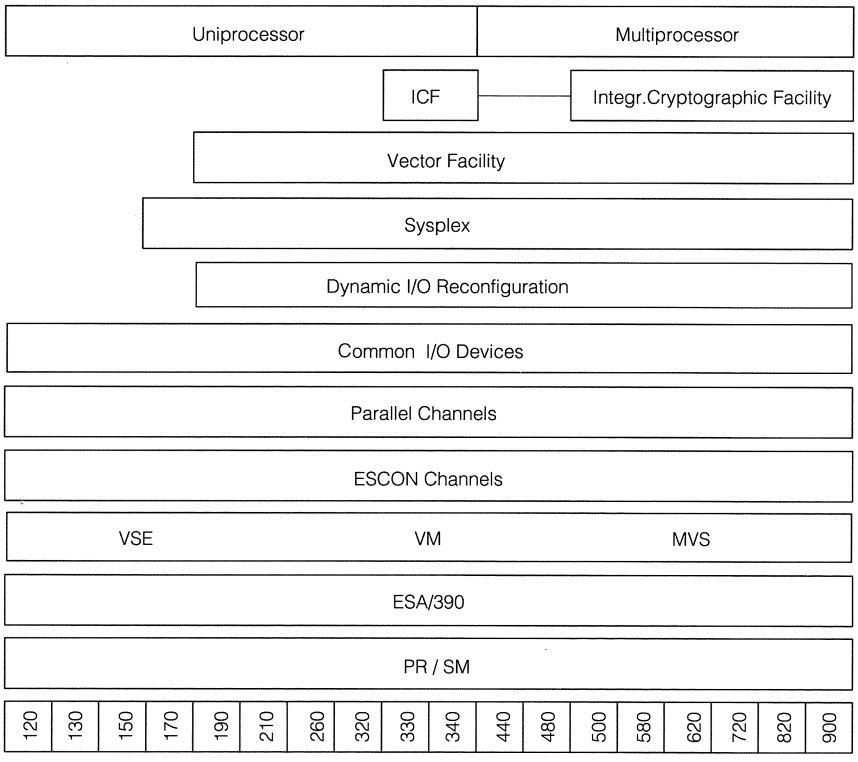

Demanding mathematical applications need systems with extra power in order to keep run times short

These models use leading-edge CMOS and bipolar technology, which provide a high-function Differential Current Switch (DCS) circuit design, and 128K CMOS Static Random Access Memory and an enhanced air-cooled TCM. The DCS circuit design was chosen for the air-cooled models because it provides high performance at a lower power rating. This is a critical requirement in a high-speed air-cooled machine. The main advantage DCS provides over emitter-coupled logic (ECL) is the ability to operate the logic with approximately one-third the signal voltage, which allows for substantial power reduction.

High-performance CMOS chips are used for the high-speed buffer, central storage and expanded storage. The CMOS chip requires only one-sixth of the power and provides twice the density of bipolar chips. This leads to significant system performance gains while reducing the total number of TCMs required. The ES/9000 air-cooled frame models utilise four TCMs to provide the same function as 21 TCMs provided for the ES/3090.

This new air-cooled TCM incorporates several design enhancements over previous TCMs. Three key design changes have been incorporated which result in improved heat dissipation to allow air cooling. These changes are:

- The cooling media has been changed from helium to a mineral oil-like fluid
- Larger copper pistons are used to improve cooling
- Aluminium finned heat sinks are mounted directly over the TCM.

(Continued on back page)

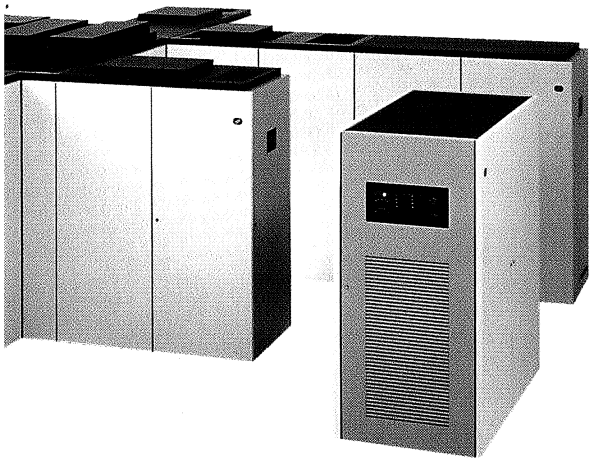

Dynamic growth: a comparison of the performance of IBM ES/9000 models with the IBM ES/9370, the IBM ES/4381 and the IBM ES/3090

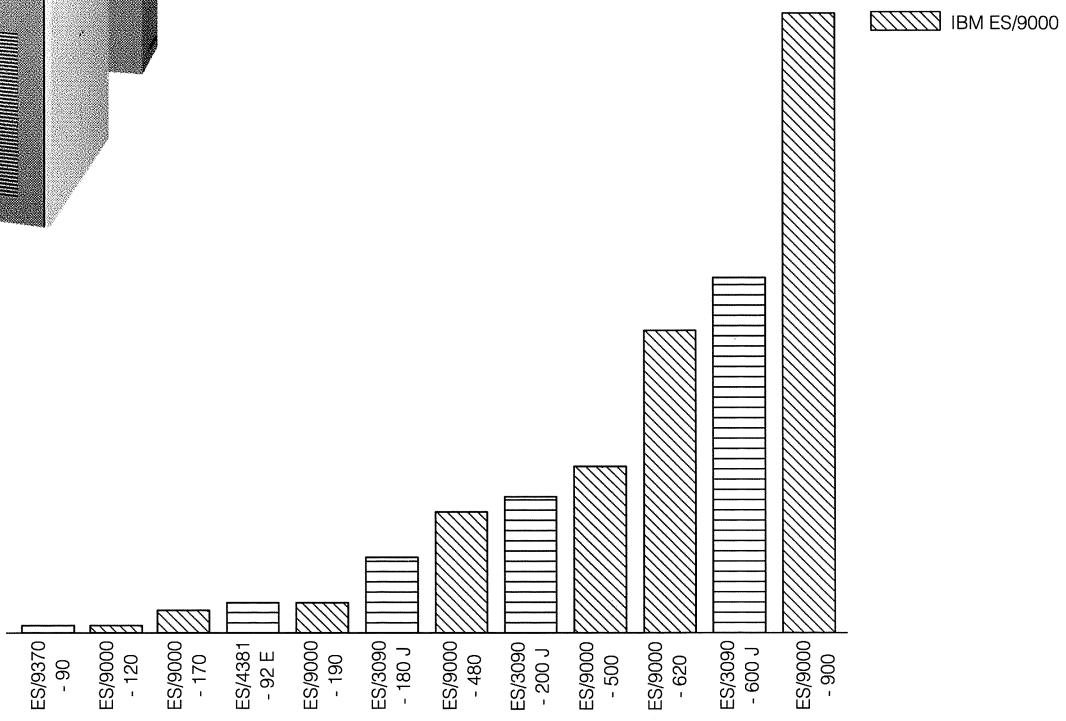

Water-cooled frame models
All water-cooled models of the ES/9000 family have enhanced power systems which improve serviceability and availability. Certain models of the ES/9000 family use new bipolar chip logic and high speed buffer storage providing the high-switching speeds needed for top performance. These chips employ both ECL and DCS circuit designs. Used together on a chip, these two logic designs match speed requirements with appropriate logic circuits. New high-speed buffer storage chips are faster and denser than ones previously used. These advances in chip technology provide improved computing power and availability as well as significantly reduced cycle times.

Newly redesigned TCMs incorporate several changes from earlier versions. Made of glass ceramic, a new substrate allows signals to run faster than current substrates. To accommodate denser chip packaging and increased heat dissipation, a high-conduction thermal design provides significant improvements in cooling. At the same time, the increased circuit density enables a reduction in the total number of TCMs required. Each board holds six TCMs. A central processor contains four TCMs, allowing the Vector Facility option or the Integrated Cryptographic Feature to be placed on the same board as the central processor.

Increased reliability, availability, and serviceability
Architecture, design, packaging and technology together benefit overall processing in several ways. Distance reductions made possible by the denser chips contribute to better performance and higher availability. The packaging's 'replaceable technology' approach enables IBM to offer new technology or function through a TCM exchange, providing you with additional growth alternatives for your strategic computing needs.

Find out more
If you would like more information about the System/390, the new ES/9000 processor models and the common architecture and application platforms, please contact your IBM marketing representative.

BY APPOINTMENT TO HER MAJESTY THE QUEEN SUPPLIERS OF ELECTRIC AND ELECTRONIC TYPEWRITERS IBM UNITED KINGDOM LIMITED PORTSMOUTH

BY APPOINTMENT TO HRH THE PRINCE OF WALES SUPPLIERS OF INFORMATION SYSTEMS IBM UNITED KINGDOM LIMITED PORTSMOUTH

1986 THE QUEEN'S AWARD FOR EXPORT ACHIEVEMENT

1988 THE QUEEN'S AWARD FOR TECHNOLOGICAL ACHIEVEMENT

IBM United Kingdom Limited
National Enquiry Centre
389 Chiswick High Road
London W4 4AL
Telephone 081-747 0747

Registered in England: No. 741598
Registered Office: PO Box 41
North Harbour, Portsmouth
Hampshire, PO6 3AU

Printed in England by
Cedar Colour Limited
Chandlers Ford, Hants

Enterprise System/9000 Models 330, 340, 500, 580, 620, 720, 820 and 900

• *Designed to generate new solutions to business needs*

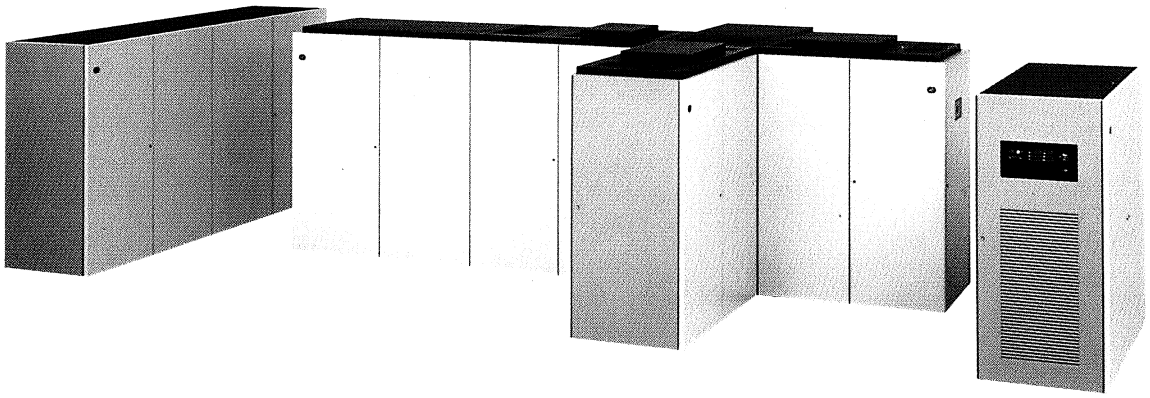

The IBM Enterprise System/9000 (ES/9000) provides the most extensive computing range ever offered within a single processor family. Based on IBM's newest technologies and proven systems architecture, this powerful processor family provides over 100-fold growth from the smallest rackmounted systems to top-of-the-line, general-purpose computers.

The result is a new era in computing that redefines management of the computing environment and lifts computing power to new levels of performance, designed to generate new and better solutions to specific business needs.

Consistent family platform
Based on IBM's Enterprise Systems Architecture/390 (ESA/390), the IBM ES/9000 family offers users a common platform across all of its 18 processor models. This unprecedented potential for growth along clearly defined paths provides exceptional granularity in computing power.

Using state-of-the-art technology and design, the ES/9000 processors introduce the advantages of Enterprise Systems Architecture/390 (ESA/390) to these operating system environments: VSE/ESA, VM/ESA and MVS/ESA. The powerful new processor family also extends such benefits as 31-bit addressing, larger virtual storage addressing and dynamic channel subsystems — making use of the increased central storage and expanded storage capacities that these processors provide.

IBM

In addition, the ES/9000 family offers an array of powerful capabilities. These include vector processing, cryptography, an Enterprise Systems Connection Channel Architecture (ESCON) implemention that includes a channel-to-channel capability, and parallel channels with data rates of up to 4.5 MB/sec. Advanced features and functions such as these are designed to break new ground in application opportunities.

Eight water-cooled models

The IBM ES/9000 family includes eight water-cooled models: the 330, 340, 500, 580, 620, 720, 820 and 900.

Family upgrade flexibility*

As the need for added processing capability increases, upgrade paths are available from existing Enterprise System/3090 Model J and the new Enterprise System/3090-9000T (ES/3090-9000T) models into selected watercooled models of the ES/9000 family. And upgrade paths are available within the ES/9000 family to satisfy additional processing needs as well.

Technology and design leadership

The ES/9000 family Models 330 through 720 contain the proven technology and packaging of the Thermal Conduction Module (TCM). The Models 820 and 900 use TCM packaging that has been significantly enhanced.

Design enhancements for ES/9000 Models 820 and 900 include a second-level high-speed buffer, a split first-level high-speed buffer and improved interprocessor communication. These technological and design enhancements are key elements underpinning the heightened speed and performance of the top-of-the-line family models.

**Upgrades from 620 to 820 and from 720 to 900 involve substantial change to the customer's existing processor. It is the customer's responsibility to determine the tax and accounting handling of processor installations and upgrades within the customer's own organisation.*

IBM ES/900 processor support units

Model	Processor controller element	Power and coolant distrib. unit	Display stations	Modem
330	9022	1	2-5[1]	1
340	9022	1	2-5[1]	1
500	9022	1	2-5[1]	1
580	9022	1	2-5[1]	1
620	9022	2	3-6[1]	2
720	9022	2	3-6[1]	2
820	9022	2	3-6[1]	2
900	9022	2	3-6[1]	2

[1] *3206 Model 100*

IBM ES/9000 hardware features*

ESA/390 Architecture	S
PR/SM	S
ESCON channels	O[1] S[2]
4.5MB parallel channels	O[2] S[1]
Sysplex Timer	O
Vector Facility	O
Integrated Cryptographic Feature	O
SIE Assist	S
DB2 Sort Enhancement	S
VM data spaces	S
Dynamic Reconfiguration Management	S
Enhanced power system	S
Console integration	S
Integrated I/O features	—
Integrated communications subsystems	—

[1] *Models 330,340,500-720*
[2] *Models 820, 900*
* *Specific software levels may be required*

S = Standard feature
O = Optional feature
— = Not applicable

Highlights

The water-cooled ES/9000 processor models provide these advanced functions:

Logical partitioning. The IBM Processor Resource/Systems Manager (PR/SM), a standard function on ES/9000 processors, allows different operating systems to run concurrently in separate logical partitions (LPARs), with a high degree of isolation. In LPAR mode, up to seven logical partitions can be simultaneously active, and up to 14 in physically partitioned configurations. LPAR mode supports System/390 and System/370 architectures.

IBM ES/9000 software support

IML Mode:	ESA/390	LPAR
VSE		
VSE/ESA V1.1	●	●
VSE/SP V4.1.2	—	●
VSE/SP V3.2.2	—	●
VM		
VM/ESA		
370 feature	—	●
ESA feature Rel 1.0	●	●
ESA feature Rel 1.1	●	●
VM/XA SP R2.1	●	●
VM/HPO R5, R6	—	●
VM/SP R5, R6	—	●
MVS		
MVS/ESA SP V4.1.0, V4.2.0	●	●
MVS/SP V3.1.0e, V3.1.3	●	●
MVS/SP V2.2.0, V2.2.3	●	●
MVS/SP V1.3.5	—	●
AIX (under VM)		
AIX/370 V1.2	●[1]	●[2]
TPF		
TPF 3.1	●[4]	●[3]

[1] *VM/ESA only*
[2] *Any supported VM level*
[3] *Single system environment*
[4] *820 & 900 in physical partition mode*
340, 500-900 in loosely coupled environment
— = Not applicable

Note:
A growing number of enabled applications are available in the areas of cooperative processing, performance monitoring, connectivity, industry specific solutions and others. See your IBM representative for current information.

Dynamic Reconfiguration Management. This new function works together with the appropriate level of the MVS/ESA operating system to allow changes to the current system I/O configuration. Channels, control units and devices can be added or removed without the need for a power on, reset and IPL — an important contribution to increased system availability.

[1] *64 up to 128; 128 up to 256;*
[2] *64 up to 128; 128 up to 256 per side*
[3] *128 up to 256; 256 up to 512 per side*
[4] *64 up to 128; 128 up to 256; 256 up to 512*
[5] *64 up to 256; 256 up to 512; 512 up to 2,096*
[6] *256 up to 512; 512 up to 1,024; 1,024 up to 4,096 per side*
[7] *Mutually exclusive with Vector Facility*
[8] *Mutually exclusive with one Vector Facility*
[9] *Maximum one Cryptographic Feature per side and mutally exclusive with one Vector Facility on that side*
†*per side*

Integrated Vector Facility. The optional Vector Facility provides a specialised execution element to process vectorised programs. Many vectorised application programs are available, encompassing a broad range of engineering and scientific solutions.

Integrated Cryptographic Feature. The new optional Cryptographic Feature is designed to strengthen your data processing installation and network security. This facility supports encryption, decryption, message authentication, personal authentication and key management.

Multisystem management. The ES/9000 processor family takes important steps toward key goals, such as 'lights-out' operation and increased control of multiple system configurations.

Connection options, enhanced with the use of fibre optics, now permit communication links between data centres, as well as between systems within a data centre. In addition, a remote data centre location can be powered on and off and environmentally monitored using the IBM ESCON Monitor System product.

Within a data centre, the ESCON capability allows a SYStems-ComPLEX (Sysplex) configuration, enabling workload balancing and backup capability between systems. Multisystem configurations requiring exceptionally accurate time-of-day coordination can now use the IBM Sysplex Timer feature to assure accuracy when timing interconnected systems operations.

Environmental changes. The ES/9000 family Models 330 to 900 come with a 50/60 Hertz (Hz) power source and a two-inch water feed between the chiller and the coolant distribution unit. These changes improve power utilisation and heat transfer performance of the new models.

IBM ES/9000 processor options

	Total channels		Parallel channels			ESCON channels		
Model	Min.	Max	Min.	Max.	Incr.	Min.	Max.	Incr.
330	16	64	16	32	16	0	32	16
340	16	64	16	32	16	0	32	16
500	32	64	32	64	16	0	32	16
580	32	64	32	64	16	0	32	16
620	64	128	64	128	16†	0	64	16†
720	64	128	64	128	16†	0	64	16†
820	128	256	0	96	16†	32	256	16†
900	128	256	0	96	16†	32	256	16†

	Processor storage (MB)		Central storage (MB)			Expanded storage (MB)		
Model	Min.	Max.	Min.	Max.	Incr.	Min.	Max.	Incr.
330	32	640	32	128	32	0	512	64[4]
340	32	1,152	32	128	32	0	1,024	64[5]
500	64	2,304	64	256	64[1]	0	2,048	64[5]
580	64	2,304	64	256	64[1]	0	2,048	64[5]
620	128	4,608	128	512	64[2]	0	4,096	64[5]
720	128	4,608	128	512	64[2]	0	4,096	64[5]
820	256	9,216	256	1,024	128[3]	0	8,192	256[6]
900	512	9,216	512	1,024	256†	0	8,192	256[6]

	Vector Facility			Integrated Cryptographic Feature			Logical partitions
Model	Min.	Max.	Incr.	Min.	Max.	Incr.	Max.
330	0	1	1	0	1[7]	1	7
340	0	1	1	0	1[7]	1	7
500	0	2	1	0	1[8]	1	7
580	0	3	1	0	1[8]	1	7
620	0	4	1	0	2[9]	1	7/14
720	0	6	1	0	2[9]	1	7/14
820	0	4	1	0	2[9]	1	7/14
900	0	6	1	0	2[9]	1	7/14

IBM ES/9000 upgrade performance comparison (ITR)

Model				SEAP (NIC)	
To	From	VM/XA SP 2.1[1]	MVS/ESA[1]	Scalar[1]	Vector[1]
330	15T	1.2	1.2-1.6	1.0-1.4	1.1-1.3
330	17T	1.0-1.1	1.1-1.2	1.0-1.1	1.0-1.1
500	180J	1.8-2.0	1.9	2.0	1.9-2.0
500	18T	1.8-2.0	1.9	2.0	1.9-2.0
500	330	2.0-2.1	2.0-2.2	2.0-2.1	2.0-2.1
500	340	1.8-2.0	1.9	2.0	1.9-2.0
580	200J	1.5-1.6	1.4	1.5	1.3-1.5
580	500	1.5-1.6	1.4	1.5	1.3-1.5
620	200J	2.0-2.1	1.8-1.9	1.9-2.0	1.7-2.0
620	280J	2.0-2.1	1.8-1.9	2.0	1.8-2.0
620	28T	2.0-2.1	1.8-1.9	2.0	1.8-2.0
620	500	2.0-2.1	1.8-1.9	1.9-2.0	1.7-2.0
620	580	1.3	1.3	1.3	1.2-1.3
720	300J	1.9	1.7-1.9	1.9-2.0	1.6-2.0
720	580	1.9	1.7-1.9	1.9-2.0	1.6-2.0
720	400J	1.4	1.3-1.4	1.5	1.3-1.5
720	620	1.4	1.3-1.4	1.5	1.3-1.5
720	500J	1.2	1.1-1.2	1.2	1.1-1.2
720	600J	1.0	1.0	1.0	1.0

[1] *Performance is in Internal Throughput Rate (ITR) ratio, based on measurements and projections using IBM benchmark workloads, using MVS/ESA SP V3.1.3 and VM/XA SP2.1.*

Improved availability. In response to user requirements for continuous operations, the ES/9000 processors provide for:

- Concurrent repair of an offline central processor, while remaining central processors continue to function
- Replacement of critical air-moving devices on the processor and of motors and pumps on the coolant distribution unit
- Dynamic Reconfiguration Management to allow less disruptive configuration changes
- Enhanced power system to provide continued system operation for most power supply failures
- Removal of the Model 3089 motor generator to eliminate a potential point of failure.

IBM ES/9000 physical characteristics*

	330		340		500		580	
	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.
Acoustics, Bels		7.8		7.8		7.8		7.8
Power consumption, 50/60Hz, KVA	32.1	38.4	32.1	38.8	40.6	49.4	48.1	57.4
Heat output, KBTU/hr								
To water	56.0	68.9	56.0	68.9	76.1	95.9	95.6	120.5
To air	18.8	40.6	18.8	40.6	35.5	44.0	37.9	47.1
Total	74.8	109.5	74.8	109.5	111.6	139.9	133.5	167.6
Floor space								
Sq. feet	82.4	88.4	82.4	99.1	82.4	99.1	93.3	99.1
Sq. metres	7.7	8.2	7.7	9.2	7.7	9.2	8.7	9.2
Including service clearance								
Sq. feet	440.7	461.7	440.7	497.1	440.7	497.1	476.3	497.1
Sq. metres	40.9	42.9	40.9	46.2	40.9	46.2	44.3	46.2
Approximate weight								
Lbs.	10985	12780	10985	12780	11925	13710	13085	13710
Kg.	4983	5797	4983	5797	5409	6219	5935	6219

	620		720		820		900	
	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.
Acoustics, Bels		8.1		8.1		8.1		8.1
Power consumption, 50/60Hz, KVA	77.2	92.8	92.2	111.2	108.6	138.6	126.6	165.8
Heat output, KBTU/hr								
To water	152.2	191.8	191.2	241.0	224.6	305.2	272.4	377.4
To air	58.0	75.0	62.8	81.2	63.4	85.4	73.0	97.0
Total	210.2	266.8	254.0	322.2	288.0	390.6	345.4	474.4
Floor space								
Sq. feet	152.6	186.9	178.0	186.9	159.3	159.3	181.2	181.2
Sq. metres	14.2	17.4	16.5	17.4	14.8	14.8	16.8	16.8
Including service clearance								
Sq. feet	720.0	834.1	791.4	834.1	728.1	728.1	799.5	799.5
Sq. metres	66.9	77.5	73.5	77.5	67.6	67.6	74.3	74.3
Approximate weight								
Lbs.	22295	24625	24625	27635	23179	25203	25823	27847
Kg.	10113	11170	11170	12535	10514	11432	11455	12631

**Specifications are subject to change without notice*

Published by IBM United Kingdom Limited for distribution only in the United Kingdom.

Registered in England No. 741598.

Registered Office:
PO Box 41, North Harbour
Portsmouth PO6 3AU

GU20-0093

Enterprise System/9000 Models 820 and 900

• *Designed to generate new solutions to business needs*

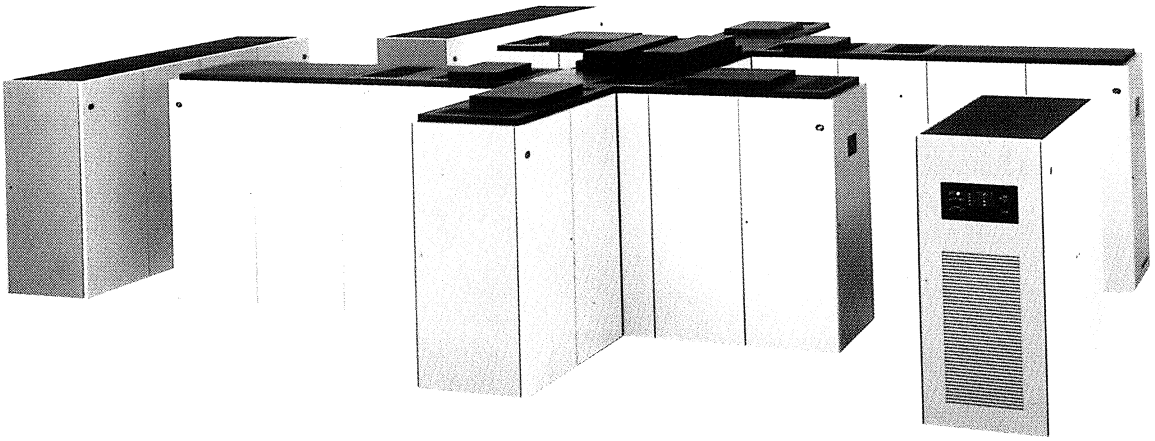

The IBM Enterprise System/9000 (ES/9000) provides the most extensive computing range ever offered within a single processor family. Based on IBM's newest technologies and proven systems architecture, this powerful processor family provides over 100-fold growth from the smallest rack-mounted systems to top-of-the-line, general-purpose computers.

The result is a new era in computing that redefines management of the computing environment and lifts computing power to new level of performance, designed to generate new and better solutions to specific business needs.

Consistent family platform

Based on IBM's Enterprise Systems Architecture/390 (ESA/390), the IBM ES/9000 family presents users with a common platform across all of its 18 processor models. What's more this unprecedented potential for growth along clearly defined paths offers exceptional granularity in computing power.

Using state-of-the-art technology and design, the ES/9000 processors introduce computing advantages of ESA/390 to these operating system environments: VS/ESA, VM/ESA and MVS/ESA. The powerful new processor family also extends such benefits as 31-bit addressing, larger virtual storage addressing and dynamic channel subsystems — making use of the increased central storage and expanded storage capacities that these processors provide.

IBM

In addition, the ES/9000 models offer an array of powerful capabilities. These include vector processing, cryptography, an Enterprise Systems Connection Channel Architecture (ESCON) implementation that includes channel-to-channel capability, and parallel channels with data rates of up to 4.5 MB/sec. Advanced features and functions such as these break new ground in application opportunities.

High-end, high-performance models

The ES/9000 family Models 820 and 900 incorporate enhancements in technology and design that underscore IBM's leadership in processing power and function.

Upgrade flexibility*

As the need for added processing capability increases, ES/9000 Model 620 and 720 users can take advantage of clear-cut upgrade paths to move up to the high-end Models 820 and 900. It's also an easy upgrade for the Model 820 user who wants to advance to the top-of-the-line Model 900.

Technology and design leadership

The high-end Models 820 and 900 incorporate a redesigned Thermal Conduction Module (TCM) containing almost twice the number of circuits and double the number of chip connections available on earlier TCMs. A new glass ceramic substrate allows signals to travel faster, and a new cooling design dissipates the heat.

The increased packaging density reduces the number of TCMs required for certain functions. Redesigned circuit boards hold six TCMs, compared to nine previously required on boards of comparable size.

Additional ES/9000 design enhancements include a second level high-speed buffer, a split first-level high speed buffer and improved interprocessor communications. These technological and design enhancements are key elements underpinning the heightened speed and performance of the ES/9000 family models.

**Upgrades from 620 to 820 and from 720 to 900 involve substantial change to the customer's existing processor. It is the customer's responsibility to determine the tax and accounting handling of processor installations and upgrades within the customer's own organisation.*

IBM ES/9000 processor support units

Model	Processor controller element	Power and coolant distrib. unit	Display stations	Modem
820	9022	2	3-6[1]	2
900	9022	2	3-6[1]	2

[1] *3206 Model 100*

IBM ES/9000 hardware features*

ESA/390 Architecture	S
PR/SM	S
ESCON channels	S
4.5MB pasrallel channels	O
Sysplex Timer	O
Vector Facility	O
Integrated Cryptographic Feature	O
SIE Assist	S
DB2 Sort Enhancement	S
VM data spaces	S
Dynamic Reconfiguration Management	S
Enhanced power system	S
Console integration	S
Integrated I/O features	—
Integrated communications subsystems	—

** Specific software levels may be required*

S = Standard feature
O = Optional feature
— = Not applicable

IBM ES/9000 software support

IML Mode:	ESA/390	LPAR
VSE		
VSE/ESA V11	●	●
VSE/SP V4.1.2	—	●
VSE/SP V3.2.2	—	●
VM		
VM/ESA		
370 feature	—	●
ESA feature Rel 1.0	●	●
ESA feature Rel 1.1	●	●
VM/XA SP R2.1	●	●
VM/HPO R5, R6	—	●
VM/SP R5, R6	—	●
MVS		
MVS/ESA SP V4.1.0, V4.2.0	●	●
MVS/SP V3.1.0e, V3.1.3	●	●
MVS/SP V2.2.0,V2.2.3	●	●
MVS/SP V1.3.5	—	●
AIX (under VM)		
AIX/370 V1.2	●[1]	●[2]
TPF		
TPF 3.1	●[4]	●[3]

[1] *VM/ESA only*
[2] *Any supported VM level*
[3] *Single system environment*
[4] *820 & 900 in physical partition mode*

Note:
A growing number of enabled applications are available in the areas of cooperative processing, performance monitoring, connectivity, industry specific solutions and others. See your IBM representative for current information.

Highlights

The powerful ES/9000 family Models 820 and 900 provide these advanced functions:

Logical organisation. Data flow is more efficient in the ES/9000 Models 820 and 900. That's because data requests can now be routed through multiple elements, reducing contention in storage-level access. In these models, the two system control elements link directly to each processor in the complex, and each resides on its own board with central storage. The new interconnect communication element controls the channels, internal buses to the system control elements and expanded storage.

Logical partitioning. The IBM Processor Resource/Systems Manager (PR/SM), a standard function on the ES/9000 processor family, allows different operating systems to run concurrently in separate logical partitions with a high degree of isolation. In logically partitioned (LPAR) mode, up to seven logical partitions can be active simultaneously (up to 14 in physically partitioned configuration). LPAR mode supports System/390 and System/370 architectures.

Dynamic Reconfiguration Management. This new function works together with the appropriate level of the MVS/ESA operating system to allow changes to the current system I/O configuration. Channels, control units and devices can be added or removed without the need for a Power On Reset and IPL – an important contribution to increased system availability.

Integrated Vector Facility. The optional Vector Facility provides a specialised execution element to process vectorised programmes. Many vectorised application programmes are available, encompassing a broad range of engineering and scientific solutions.

Integrated Cryptographic Feature. The new optional Cryptographic Feature is designed to strengthen your data processing installation and network security. This facility supports encryption, decryption, message authentication, personal authentication and key management.

Multisystem management. The ES/9000 processor family takes important steps toward key goals, such as 'lights out' operation and increased control of multiple system configurations. Enhanced connectivity options using fibre optics now permit connections between data centres as well as between systems within a data centre. In addition, a remote data centre location can be powered on and off and environmentally monitored using the IBM ESCON Monitor System product.

Within a data centre, the IBM ESCON capability allows a SYStems-comPLEX (Sysplex) configuration, enabling workload balancing and backup capability between systems. Multi-system configurations requiring exceptionally accurate time-of-day coordination can now use the IBM Sysplex Timer feature to assure accuracy when timing interconnected systems operations.

Environmental changes. The ES/9000 family Models 820 and 900 require a 50/60 Hertz (Hz) power source and a two-inch water feed between the chiller and the coolant distribution unit. These changes eliminate a potential point of failure and improve heat transfer performance for these new models.

Improved availability. In response to user requirements for continuous operations, the following features result in increased systems availability:

- Concurrent repair of an offline central processor, while remaining central processors continue to function
- Replacement of critical air-moving devices on the processor and of motors and pumps on the coolant distribution unit
- Dynamic Reconfiguration Management, allowing less disruptive configuration changes
- Enhanced power system to provide continued system operation for most power supply failures
- Removal of the Model 3089 motor generator to eliminate a potential point of failure.
- Significant parts reduction made possible by advanced chip technology.

IBM ES/9000 processor options

	Total channels		Parallel channels			ESCON channels		
Model	Min.	Max	Min.	Max.	Incr.	Min.	Max.	Incr.
820	128	256	0	96	16†	32	256	16†
900	128	256	0	96	16†	32	256	16†

	Processor storage (MB)		Central storage (MB)			Expanded storage (MB)		
Model	Min.	Max.	Min.	Max.	Incr.	Min.	Max.	Incr.
820	256	9.216	256	1.024	128[1]	0	8.192	256[2]
900	512	9.216	512	1.024	256†	0	8.192	256[2]

	Vector Facility			Integrated Cryptographic Feature			Logical partitions
Model	Min.	Max.	Incr.	Min.	Max.	incr.	Max.
820	0	4	1	0	2[3]	1	7/14
900	0	6	1	0	2[3]	1	7/14

[1] *128 up to 256; 256 up to 512 per side*
[2] *256 up to 512; 512 up to 1,024; 1,024 up to 4,096 per side*
[3] *Maximum one Cryptographic Feature per side and mutally exclusive with one Vector Facility on that side*
†*per side*

IBM ES/9000 physical characteristics*

	820		900	
	Min.	Max.	Min.	Max.
Acoustics, Bels		8.1		8.1
Power consumption, 50/60Hz, KVA	108.6	138.6	126.6	165.8
Heat output, KBTU/hr				
To water	224.6	305.2	272.4	377.4
To air	63.4	85.4	73.0	97.0
Total	288.0	390.6	345.4	474.4
Floor space				
Sq. feet	159.3	159.3	181.2	181.2
Sq. metres	14.8	14.8	16.8	16.8
Including service clearance				
Sq. feet	728.1	728.1	799.5	799.5
Sq. metres	67.6	67.6	74.3	74.3
Approximate weight				
Lbs.	23179	25203	25823	27847
Kg.	10514	11432	11455	12631

**Specifications are subject to change without notice*

IBM ES/9000 upgrade performance comparison (ITR)

Model		Commercial	SEAP (NIC)	
To	From	work-loads[1]	Scalar[1]	Vector[1]
720	600J	1.0	1.0	1.0
820	620	Up to 1.9	Up to 2.7	Up to 2.8
900	720	Up to 1.9	Up to 2.7	Up to 2.8

[1] *Performance is in Internal Throughput Rate (ITR) ratio, based on measurements and projections using IBM benchmark workloads, using MVS/SP V3.1.3 and VM/XA SP 2.1.*

Published by IBM United Kingdom Limited for distribution only in the United Kingdom.

Registered in England No. 741598.

Registered Office:
PO Box 41, North Harbour
Portsmouth PO6 3AU

GU20-0092

Enterprise System/9000 Models 120, 130, 150 and 170

• *Designed to generate new solutions to business needs*

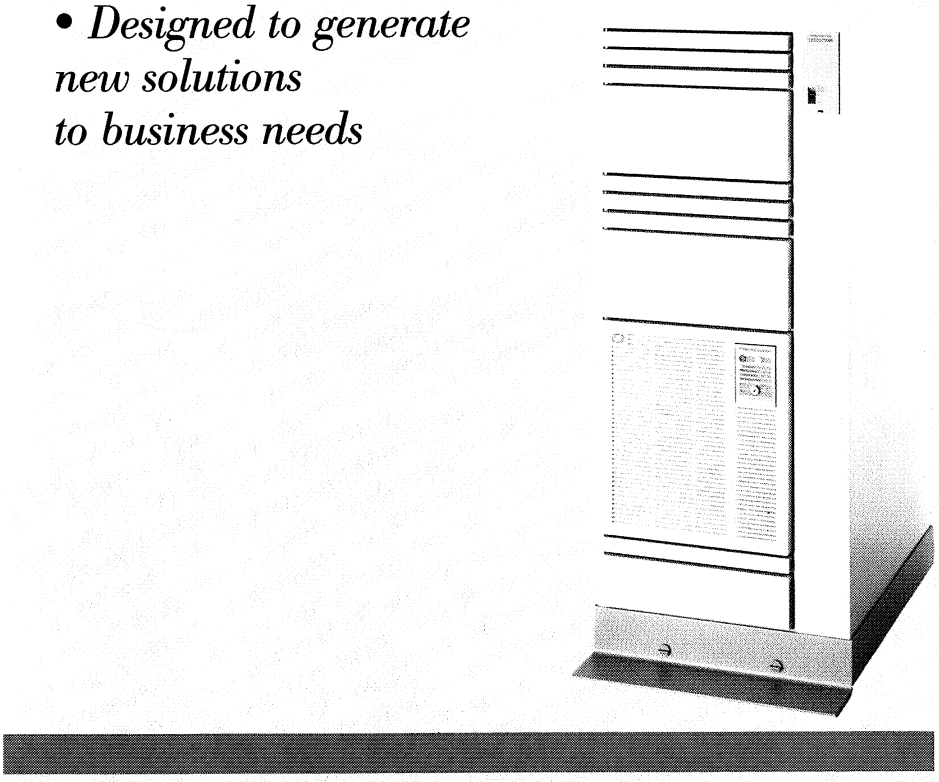

The IBM Enterprise System/9000™ (ES/9000™) provides the most extensive computing range ever offered within a single processor family. Based on IBM's newest technologies and proven systems architecture, this powerful processor family provides over 150-fold growth from the smallest rack-mounted systems to top-of-the-line, general-purpose computers. The result is a new era in computing that redefines management of the computing environment and lifts computing power to new level of performance, designed to generate new and better solutions to specific business needs.

Consistent family platform
Based on IBM's Enterprise Systems Architecture/390™ (ESA/390™), the IBM ES/9000 family provides a common platform across all of its 18 processor models. This unprecedented potential for growth along clearly defined paths offers exceptional granularity in computing power.

Using state-of-the-art technology and design, the ES/9000 processors introduce the advantages of ESA/390; VM/ESA, VSE/ESA and MVS/ESA. The powerful new processor family also extends such benefits as 31-bit addressing, larger virtual storage addressing and dynamic channel subsystems — making use of the increased central storage and expanded storage capacities these processors provide.

In addition, the ES/9000 family models offer an array of powerful capabilities. These include a fibre-optic channel implementation of Enterprise Systems Connection Architecture™ (ESCON™), which provides channel-to-channel capability, parallel channels with data rates of up to 4.5 MB/sec, and serial channels with up to 10 MB/sec. Advanced features and functions such as these are designed to break new ground in application opportunities.

IBM

Four air-cooled, rack-mounted processor models

The ES/9000 family includes four rack-mounted processors — Models 120, 130, 150 and 170. Growth within this group is straightforward and easy. The Enterprise System/9370™* models — including the 50, 60, and 80 and 90 — can be upgraded to the new models. And you can migrate from the IBM ES/4381,™ as well.

Entry into the ES/9000 family offers a new dimension in computing benefits and positions your enterprise for unprecedented computing growth.

ES/9000 models support most existing rack-mounted devices, as well as new rack-mounted, direct-access storage devices (DASD), tape and communication units. This extensive support helps protect your investment in current input/output (I/O) devices.

Advanced function, comparable to that of current high-end systems, is built into the rack-mounted processors through new design enhancements that include:
*Enterprise Systems Architecture/390
*4.5 MB/sec. parallel channels
*IBM's new ESCON channels
*VM operating system assists
*Larger processor storage
*Expanded storage
*Processor Resource/Systems Manager™ (PR/SM™)

The latest advances in state-of-the-art technology and packaging are hallmarks of the ES/9000 air-cooled, rack-mounted processors. These include enhanced system reliability, increased availability and simplified servicing, all resulting in reduced environmental requirements.

Highlights

The ES/9000 family of processors offers a checklist of advanced functions. This list includes:

IBM ES/9000 processor support units

Model	Processor controller element	Power and coolant distrib. unit	Display stations	Tape streamer	Modem
120	PS/2 Model 70[1]	—	1[2]	1	1
130	PS/2 Model 70[1]	—	1[2]	1	1
150	PS/2 Model 70[1]	—	1[2]	1	1
170	PS/2 Model 70[1]	—	1[2]	1	1

[1] *Shipped preconfigured with the system*
[2] *Alternate and remote consoles are available*

IBM ES/9000 hardware features*

ESA/390 Architecture	S
PR/SM	S
ESCON channels	O
4.5MB pasrallel channels	O
Sysplex Timer	O
VM data spaces	S
Console integration	S
Integrated I/O features	O
Integrated communications subsystems	O
Rack-mounted MCCU	S
Battery backup	S

** Specific software levels may be required*

S = Standard feature
O = Optional feature

IBM ES/9000 software support

	Models 120-170		
IML Mode:	ESA/390	LPAR	S/370
VSE			
VSE/ESA V1.1	●	●	●
VSE/SP V4.1.2	—	●	●
VSE/SP V3.2.2	—	●	●
VM			
VM/ESA			
370 feature	—	●	●
ESA feature Rel 1.0	●	●	—
ESA feature Rel 1.1	●	●	—
VM/XA SP R2.1	—	●	—
VM/HPO R5, R6	—	—	—
VM/SP R5, R6	—	●	●
MVS			
MVS/SP/ESA V4.1.0, V4.2.0	●	●	—
MVS/SP V3.1.0e, V3.1.3	●	●	—
MVS/SP V2.2.0, V2.2.3	—	●	●
MVS/SP V1.3.5, V1.3.6	—	—	—
AIX (under VM)			
AIX/370 V1.2	—	—	—
TPF			
TPF 3.1	—	—	—
DPPX			
DPPX/370 V1.1, V1.2	—	●	●
DBX/370 V1.3	●	●	●
MUMPS/VM	—	●	●

— = Not applicable

Note:
A growing number of enabled applications are available in the areas of cooperative processing, performance monitoring, connectivity, industry specific solutions and others. See your IBM representative for current information.

ESA/390. The rack-mounted, air-cooled processors support these ESA/390 operating systems: MVS/ESA, VM/ESA and VSE/ESA — in addition to current operating systems.

PR/SM. PR/SM, a standard function on the processor family, allows different operating systems to run concurrently. The rack-mounted models can run in up to four logical partitions with a high degree of isolation. This capability allows jobs in separate environments — such as production, test and migration — to be run concurrently on one physical system.

Flexible storage. Up to 256 megabytes (MB) of processor storage are available on the ES/9000 rack-mounted models. All 256MB can be configured as central storage; or up to 240MB configured as expanded storage at IML (Initial Machine Load) time.

Channel choices. Both parallel and ESCON channels are available on all rack-mounted processors. ESCON channels offer a data transfer rate of up to 10MB/sec over distances of up to 9 kilometres (using two ESCON directors) in VM/ESA and MVS/ESA operating system environments. The parallel channels support data rates of up to 4.5MB/sec. The combination of parallel and ESCON channels allows users to attach IBM's fastest I/O devices as well as most other I/O devices in use today.

ES/9000 upgrade performance comparison (ITR)

Model			
From	To	Ramp/C VSE 4.1[1]	Compute Intensive Linpac
9370-50	130	2.1	4.8
9370-60	130	2.7	5.5
120	130	1.8	1.8
9370-50	150	2.9	6.9
9370-60	150	3.9	7.9
9370-80	150	2.4	5.3
9370-90	150	1.8	4.0
130	150	1.4	1.4
9370-90	170	2.3	4.6
150	170	1.3	1.2
4381-91E	130	0.6	1.7
4381-91E	150	0.9	2.4
4381-91E	170	1.1	2.8

[1] *Performance is in Internal Throughput Rate (ITR) ratio, based on measurements and projections using IBM benchmark workloads.*

Integrated I/O. Up to six integrated bus adapters are available. These adapters can attach existing and new rack-mounted I/O. They operate at speeds of up to 6.5 MB/sec. And they can also be used to attach a maximum of 16 integrated I/O controllers. Integrated I/O, parallel and ESCON channels can all be conveniently combined in the same processor.

VM enhancements for improved system performance. Several special-purpose VM enhancements are also offered to improve system performance of the four air-cooled, rack-mounted processors.

Multisystem management. The new IBM ES/9000 processors take important steps toward key goals, such as "lights out" operation and increased control of configurations with multiple systems. Enterprise Systems Connection Architecture now permits connections between systems within a data centre as well as between different data centres.

A remotely located data centre, for example, can be powered on and off and monitored from a central location.

Within a data centre, the ESCON options enable workload balancing and a backup capability between systems. Multisystem configurations that require accurate time-of-day coordination can use the Sysplex Timer feature for accurate time reference to interconnected systems.

To support the installation IBM Customer Engineering are able to offer a Software Preload Service.

IBM ES/9000 processor options

	Total channels		Parallel channels			ESCON channels			Integrated I/O buses		
Model	Min.	Max.	Min.	Max.	Incr.	Min.	Max.	Incr.	Min.	Max.	Incr.
120	0	12	0	12	1 or 3	0	12	1 or 3	0	4	2
130	0	12	0	12	1 or 3	0	12	1 or 3	0	4	2
150	0	12	0	12	1 or 3	0	12	1 or 3	0	6	2
170	0	24	0	24	1 or 3	0	24	1 or 3	0	6	2

	Processor storage (MB)			Central storage (MB)			Expanded storage (MB)		
Model	Min.	Max.	Incr.	Min.	Max.	Incr.	Min.	Max.	Incr.
120	16	256[1]	16[2]	16	256	*	0	240	*
130	16	256[1]	16[2]	16	256	*	0	240	*
150	16	256[1]	16[2]	16	256	*	0	240	*
170	32	256[1]	32[2]	16	256	*	0	240	*

[1] *128 if integrated I/O are installed*
[2] *16 up to 32; 32 up to 128; 64 up to 256*
**Granularity of central storage and expanded storage at system initialisation is model dependent.*

IBM ES/9000 physical characteristics*

	120		130		150		170	
	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.
Acoustics, Bels	6.2	6.5	6.2	6.5	6.2	6.5	6.2	6.5
Power consumption, 50/60Hz, KVA	0.6	1.4	0.6	1.4	0.6	1.4	0.6	1.4
Heat output, KBTU/hr Total air	1.7	3.8	1.7	3.8	1.7	3.8	1.7	3.8
Floor space**								
Sq. feet		6.48		6.48		6.48		6.48
Sq. metres		0.6		0.6		0.6		0.6
Including service clearance**								
Sq. feet		86		86		86		86
Sq. metres		6		6		6		6
Approximate weight**								
Lbs.		187		187		187		187
Kg.		85		85		85		85

**Specifications subject to change without notice.*
***Processor rack only*

Published by IBM United Kingdom Limited for distribution only in the United Kingdom.

Registered in England No. 741598.

Registered Office:
PO Box 41, North Harbour
Portsmouth PO6 3AU

GU20-0089

Enterprise System/9000 Models 190, 210, 260, 320, 440 and 480

• *Designed to generate new solutions to business needs*

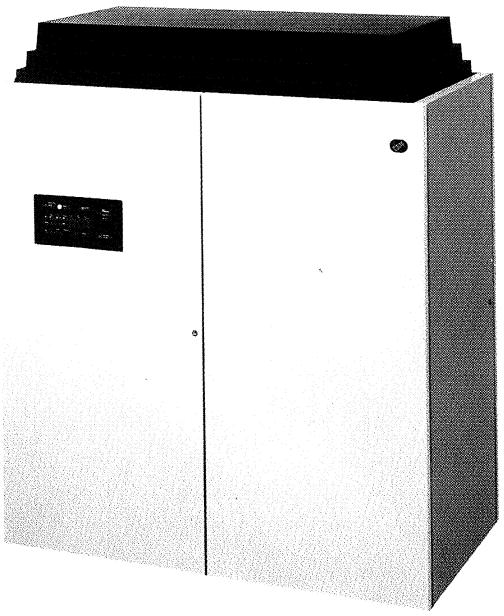

The IBM Enterprise System/9000 (ES/9000) provides the most extensive computing range ever offered within a single processor family. Based on IBM's newest technologies and proven systems architecture, this powerful processor family provides over 100-fold growth from the smallest rack-mounted systems to top-of-the-line, general-purpose computers.

The result is a new era in computing that redefines management of the computing environment and lifts computing power to a new level of performance, designed to generate new and better solutions to specific business needs.

IBM

Consistent family platform

Based on IBM's Enterprise Systems Architecture/390 (ESA/390), the IBM ES/9000 family provides a common platform across all of its 18 processor models. This unprecedented potential for growth along clearly defined paths offers exceptional granularity in computing power.

Using state-of-the-art technology and design, the ES/9000 processors introduce the advantages of ESA/390 to these operating system environments: VSE/ESA, VM/ESA and MVS/ESA. The powerful new processor family also extends such benefits as 31-bit addressing, larger virtual storage addressing and dynamic channel subsystems — making use of the increased Central Storage and Expanded Storage capacities that these processors provide.

In addition, the ES/9000 family offers an array of powerful capabilities. These include vector processing, Common Cryptograph Architecture, a fibre-optic channel implementation of Enterprise Systems Connection Architecture (ESCON) which provides channel-to-channel capability, and parallel channels with data rates of up to 4.5 MB/sec. Advanced features and functions such as these are designed to break new ground in application opportunities.

Six air-cooled, frame processor models*

The ES/9000 family includes six air-cooled frame processors — Models 190, 210, 260 and 320 are uniprocessors; Models 440 and 480 are dyadic processors. Growth within this group is straightforward and easy. These models offer a significant improvement in internal throughput over previous IBM air-cooled processors. For example, Model 480 internal throughput is over four times that of the ES/4381 Model 92E. In addition, advanced functions previously limited to the IBM Enterprise System/3090 (ES/3090) processors are now available on these ES/9000 models.

**Growth from Model 170 to Models 210 and 260 involves substantial change to the customer's existing system; converting from a rack-mounted to a frame processor includes the replacement of the central processor complex (CPC).*

IBM ES/9000 processor support units

Model	Processor controller element	Power and coolant distrib. unit	IOSP Display stations	Tape streamer	Modem
190	I/O Support Processor[1]	—	1-5	1	1
210	I/O Support Processor[1]	—	1-5	1	1
260	I/O Support Processor[1]	—	1-5	1	1
320	I/O Support Processor[1]	—	1-5	1	1
440	I/O Support Processor[1]	—	1-5	1	1
480	I/O Support Processor[1]	—	1-5	1	1

[1] *Shipped preconfigured with the system*

IBM ES/9000 hardware features*

ESA/390 Architecture	S
PR/SM	S
ESCON channels	O
4.5MB parallel channels	S
Sysplex Timer	O
Vector Facility	O
Integrated Cryptographic Feature	—
SIE Assist VM data spaces	S
DB2 Sort Enhancement	S
Dynamic Reconfiguration Management	S
Enhanced power system	—
Console integration	S
Integrated I/O features	—
Integrated communications subsystems	—
Rack-mounted MCCU	—
Battery backup	—

** Specific software levels may be required*

S = Standard feature
O = Optional feature
— = Not applicable

IBM ES/9000 software support

IML Mode:	ESA/390	LPAR
VSE		
VSE/ESA V1.1	●	●
VSE/SP V4.1.2	—	●
VSE/SP V3.2.2	—	●
VM		
VM/ESA		
370 feature	—	●
ESA feature Rel 1.0	●	●
ESA feature Rel 1.1	●	●
VM/XA SP R2.1	●	●
VM/HPO R5, R6	—	●
VM/SP R5, R6	—	●
MVS		
MVS/ESA SP V4.1.0, V4.2.0	●	●
MVS/SP V3.1.0e, V3.1.3	●	●
MVS/SP V2.2.0, V2.2.3	●	●
MVS/SP V1.3.5, V1.3.6	—	●
AIX (under VM)		
AIX/370 V1.2	●[1]	●[2]
TPF		
TPF 3.1	●	●[3]

[1] *VM/ESA only*
[2] *Any supported VM level*
[3] *Single system environment*
**Uniprocessor only*
— = Not applicable

Note:
A growing number of enabled applications are available in the areas of cooperative processing, performance monitoring, connectivity, industry specific solutions and others. See your IBM representative for current information.

Entry into the ES/9000 family offers a new dimension of computing potential, and positions your enterprise for unprecedented computing growth.

The new design enhancements include:

- Enterprise Systems Architecture/390
- 4.5 MB/sec. parallel channels
- IBM ESCON channel architecture
- VM special-purpose enhancements
- Larger processor storage
- Expanded Storage
- Processor Resource/Systems Manager (PR/SM)
- Vector facility
- Dyadic processors

IBM System/370 users can take advantage of these enhancements by migrating to the new ES/9000 processors.

The latest advances in state-of-the-art technology and packaging are hallmarks of the ES/9000 air-cooled, frame processors. These include enhanced system reliability, increased availability and simplified servicing, all resulting in reduced environmental requirements.

Highlights

The ES/9000 family of processors offers a checklist of advanced functions. This list includes:

ESA/390. Users of the ES/9000 air-cooled frame models can take full advantage of all MVS/ESA operating system capabilities. VSE/ESA and VM/ESA operating systems are also supported, as are current operating systems.

IBM ES/9000 physical characteristics*

	190		210		260		320		440		480	
	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.	Min.	Max.
Acoustics, Bels		7.2		7.2		7.2		7.2		7.2		7.4
Power consumption, 50/60Hz, KVA		7.4		7.6		7.9		8.0		11.6		11.6
Heat output, KBTU/hr												
Total air	14.1	18.5	14.1	19.0	14.6	19.5	14.6	19.8	21.2	28.7	21.2	28.7
Floor space												
Sq. feet	14.7	14.7	14.7	24.1	14.7	24.1	14.7	24.1	14.7	24.1	14.7	24.1
Sq. metres	1.37	1.37	1.37	2.24	1.37	2.24	1.37	2.24	1.37	2.24	1.37	2.24
Including service clearance												
Sq. feet	96.6	96.6	96.6	125.2	96.6	125.2	96.6	125.2	96.6	125.2	96.6	125.2
Sq. metres	9.0	9.0	9.0	11.5	9.0	11.5	9.0	11.5	9.0	11.5	9.0	11.5
Approximate weight												
Lbs.	1865	2000	1865	2765	1865	2765	1865	2765	2000	2900	2000	2900
Kg.	839	906	839	1248	839	1248	839	1248	906	1315	906	1315

**Specifications are subject to change without notice.*

PR/SM. PR/SM, a standard function on all ES/9000 processors, allows different operating systems to run concurrently. The air-cooled, frame models can run in up to seven logical partitions, with a high degree of isolation. PR/SM allows jobs in separate environments — such as production, test and migration — to be run concurrently on one physical system.

Flexible storage. The frame models offer up to one gigabyte (GB) of processor storage. A minimum of 32MB and a maximum of 256MB can be configured as Central Storage; or up to 992MB configured as Expanded Storage at IML (Initial Machine Load) time (480MB max on Model 190).

Channel choices. Both parallel and ESCON channels are available on all six frame models. ESCON channels provide data transfer rates of up to 10 MB/sec over distances of up to 3 kilometres in all ESA/390 environments, and 9 kilometres (using two ESCON directors) in VM/ESA and MVS/ESA environments. The parallel channels support data rates of up to 4.5 MB/sec. The combination of parallel and ESCON channels allows users to attach IBM's fastest I/O devices as well as most other current I/O devices.

Special-purpose enhancements. The air-cooled frame models offer special-purpose enhancements for the VM/ESA operating system as well as a DB2 sort enhancement. These enhancements contribute to improved system and application performance.

Multisystem management. The new ES/9000 processor family takes important steps toward key goals, such as "lights out" operation and improved control of multiple-system configurations. Enterprise Systems Connection Architecture now permits connections between systems within a data centre as well as between different data centres.

IBM ES/9000 processor options

	Total channels		Parallel channels			ESCON channels		
Model	Min.	Max	Min.	Max.	Incr.	Min.	Max.	Incr.
190	8	32	8	24	4	0	20	4
210	8	48	8	48	4	0	36	4
260	12	48	12	48	4	0	36	4
320	12	48	12	48	4	0	36	4
440	12	48	12	48	4	0	36	4
480	12	48	12	48	4	0	36	4

	Processor storage (MB)			Central storage (MB)			Expanded storage (MB)		
Model	Min.	Max.	Incr.	Min.	Max.	Incr.	Min.	Max.	Incr.
190	64	512	64[3]	32	128	*	0	480	*
210	64	1,024	64[1]	32	256	*	0	992	*
260	64	1,024	64[1]	32	256	*	0	992	*
320	64	1,024	64[1]	32	256	*	0	992	*
440	128	1,024	128[2]	32	256	*	0	992	*
480	128	1,024	128[2]	32	256	*	0	992	*

	Vector Facility				Logical partitions
Model	Min.	Max.	Incr.	Integrated Cryptographic Feature	Max.
190	0	1	1	—	7
210	0	1	1	—	7
260	0	1	1	—	7
320	0	1	1	—	7
440	0	2	1	—	7
480	0	2	1	—	7

[1] *64 up to 128; 128 up to 256; 256 up to 512, 512 up to 1,024*
[2] *128 up to 256; 256 up to 512; 512 up to 1,024*
[3] *64 up 128; 128 up to 256; 256 up to 512*
* *Granularity of central storage and expanded storage at system initialisation is model dependent.*
— = *Not applicable*

A remotely located data centre, for example, can be powered on and off and environmentally monitored from a central location.

Within a data centre, the ES/9000 family allows for a SYStems-comPLEX (sysplex) configuration, permitting workload balancing and a backup capability between systems.

Multisystem configurations that require exceptionally accurate time-of-day coordination can now use the Sysplex Timer feature to assure accuracy when timing interconnected systems operations.

Vector Facility. Businesses now implement numerically intensive applications to solve a broad range of business and technical requirements. For many applications, the Vector Facility can increase application throughput in comparison with scalar performance. IBM offers sources for many vectorised application programs, encompassing a broad range of engineering and scientific solutions. This valuable function is now available on the four uniprocessors and the two dyadic processors.

Dyadic models. The new dyadic design of the Models 440 and 480, implemented through the system control element, provides a new level of performance for IBM air-cooled processors.

Design advantage. An independent I/O processor, incorporated in all models, handles all channel-related activity, freeing the central processor for other tasks and improving overall system throughput.

IBM ES/9000 upgrade performance comparison (ITR)

Model				SEAP (NIC)	
To	From	VM/XA SP 2.1[1]	MVS/ESA[1]	Scalar[1]	Vector[1]
210	190	1.2-1.3	1.4-1.6	1.0-1.6	1.1-1.8
260	190	1.5-1.6	1.7-2.2	1.1-2.1	1.2-2.2
260	210	1.2	1.2-1.4	1.0-1.3	1.0-1.3
320	210	1.4	1.3-1.9	1.0-1.6	1.1-1.5
320	260	1.2	1.1-1.4	1.0-1.3	1.0-1.3
440	320	1.6	1.3-1.7	1.8-1.9	1.6-1.9
440	260	1.9-2.0	1.8-1.9	2.0-2.4	1.8-2.0
480	260	2.3	2.2-2.6	2.0-2.2	2.0-2.1
480	320	1.9-2.0	1.8-1.9	2.0	1.8-2.0
480	440	1.2	1.1-1.4	1.0-1.1	1.0-1.1

[1] *Performance is in Internal Throughput Rate (ITR) ratio, based on measurements and projections using IBM benchmark workloads, using MVS/SP V3.1.3 and VM/XA SP 2.1.*

Dynamic Reconfiguration Management. The new Dynamic Reconfiguration Management capability in the air-cooled, frame models works together with the appropriate level of the MVS/ESA operating system to allow dynamic reconfiguration of the system I/O. Channels, control units and devices can be added or removed without the need for Power On Reset and IPL (initial program load). This makes an important contribution to improved system availability.

Published by IBM United Kingdom Limited for distribution only in the United Kingdom.

Registered in England No. 741598.

Registered Office:
PO Box 41, North Harbour
Portsmouth PO6 3AU

GU20-0091

IBM

IBM System/390

Empowering your enterprise

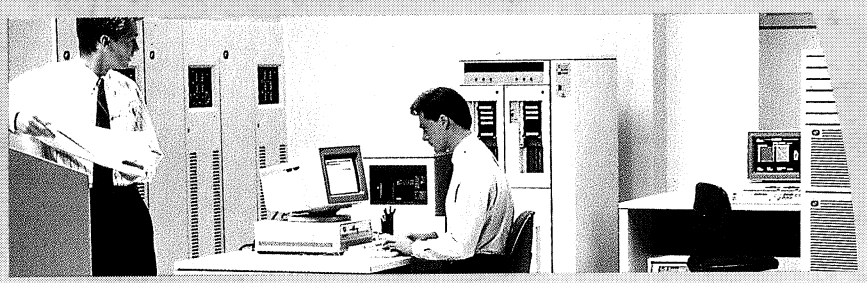

Turning business realities into advantages

Wherever your enterprise operates ... whatever that business may be ... your organisation probably faces worldwide opportunities and competitive pressures. In the past decade alone, business opportunities and challenges have escalated and become global in scope. As a result, you may be among the the executives who routinely need to:

- Make decisions about engaging in new market initiatives, strategic partnerships, joint ventures, mergers or acquisitions.
- Enable your enterprise to respond to new opportunities by helping your diverse individual users, workgroups and systems to exchange vital information easily and more quickly.

Information processing should reflect an enterprise's operating conditions

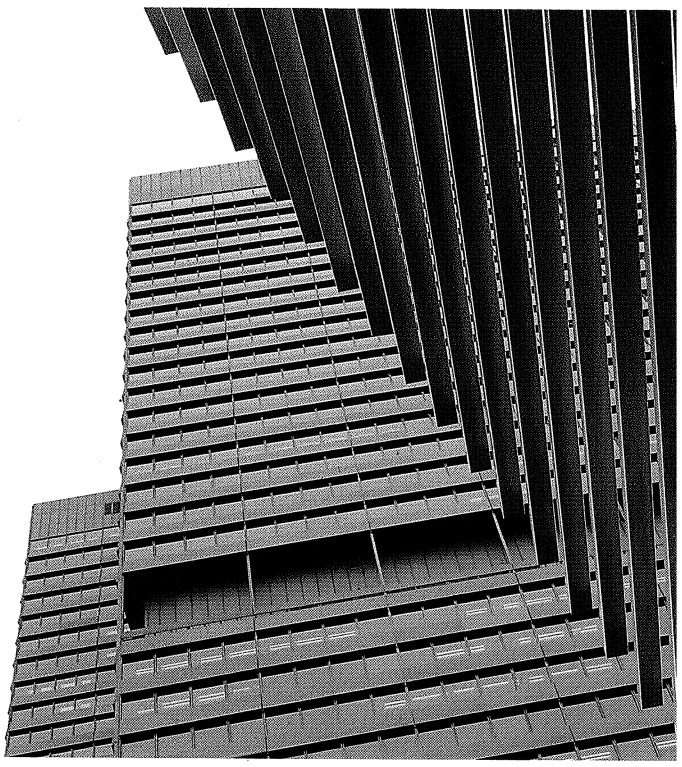

The concept behind the IBM ES/9000 family enables information processing systems to be fully integrated

- Permit suppliers, subsidiaries and field offices to exchange information with your head office – and often with one another, as well.
- Control costs.
- Deploy personnel for maximum effectiveness.

An opportunity to transform your enterprise
Issues such as these represent an implicit opportunity to transform and streamline your enterprise.

Often, such transformation is welcome – especially when fine-tuning your information infrastructure can enhance tomorrow's success by enabling you to:

- Get the information needed to make rapid decisions and respond quickly.
- Simplify access to resources, facilities and systems.
- Protect against the increased corporate security risks that accompany information exchange within and beyond enterprise boundaries.
- More easily manage multiple central, remote and backup computing sites and complicated communications technologies.
- Exploit competitive-edge application and networking opportunities without impeding your business growth by exceeding your resource capacity.

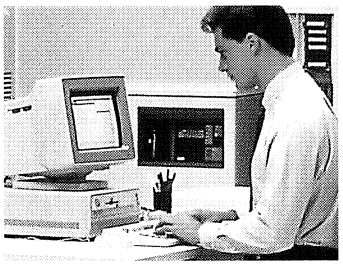

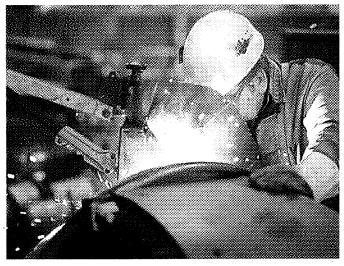

Information processing should be used uniformly as a tool throughout the enterprise

Powering up your enterprise in the new era

The excellent balance between price and performance makes entry into the mainframe class an attractive proposition

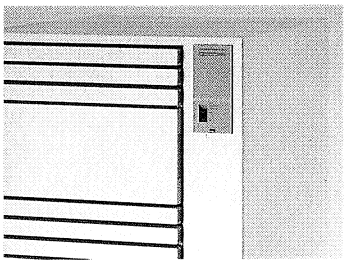

A widespread call for increased access and control

Executives worldwide tell us that barriers to corporate success must be removed. They want fast, reliable and easy access to information within their enterprises. They also seek the freedom to share data among manufacturing, marketing, finance, head office, suppliers and customers. And they want to manage and control the diverse systems throughout their enterprises.

IBM listened carefully as business leaders around the globe told us their concerns, needs and requirements. From them, we formed the foundation for a new era in computing.

With the advent of IBM's new era in computing you can equip your enterprise with the power to increase productivity, reduce costs, develop global markets, automate appropriate procedures, shorten development cycles and respond to customers more quickly and effectively. This new era encompasses:

- The IBM System/390*, a new family of growth-oriented, high-performance computing systems and software
- Broader, more comprehensive information access in a secure environment
- Innovative technology supporting your business vision and enhancing your current information assets
- Increased managerial control over diverse network and system elements and functions.

A solid foundation for increased access and growth

IBM System/390 provides the power through a range of new systems and products, as well as communications routers and switching devices, networking solutions, advanced applications and a wealth of ancillary services. Separately in some instances, together in others, these offerings provide a solid foundation. And this foundation enables your applications and networked systems (non-IBM as well as IBM) to grow in sophistication and accessibility as you strategically deploy them to achieve such enterprise-wide goals as increased market share, improved profit margins, higher quality, greater customer satisfaction and heightened competitiveness.

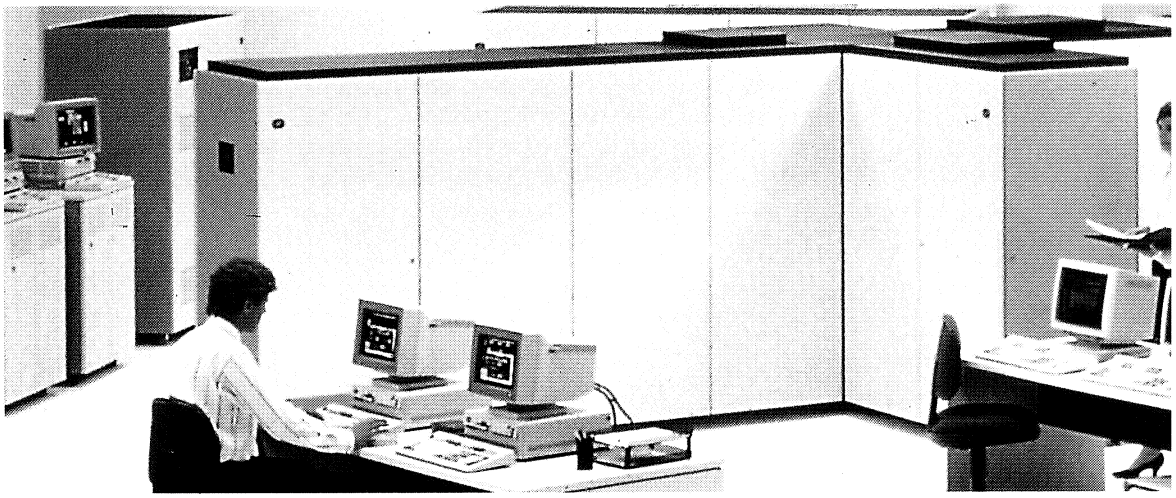

The following examples illustrate how our innovative products and architectures can advance your business vision – and help you improve your company's performance:

- Retail buyers can gather the latest sales trend information from multiple worldwide sites through our new systems, software and networking solutions.
- International bankers can encrypt sensitive electronic funds transfers for maximum confidentiality and security.
- Government employees can consolidate and graphically depict budget authorisations and actual expenditures on an agency or departmental basis by using our decision-support offerings and integrated office systems.
- A food company's management team can analyse information on field inventory and brand movement on a daily basis.
- Manufacturers can give their distributors online catalogue access through our new systems and networking solutions.
- Educational systems can be accessed by home - or office-based users for interactive skills tutoring in languages or mathematics.
- Clothing importers can use our new systems and IBM business partner software to track imported finished goods and letters of credit, and to exchange vital information with their international suppliers by electronic data interchange.

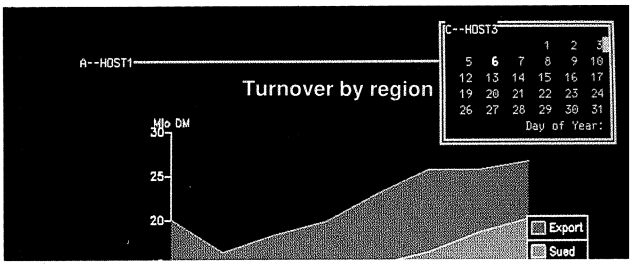

Far-reaching technological innovations lend the traditional computing centre a whole new image

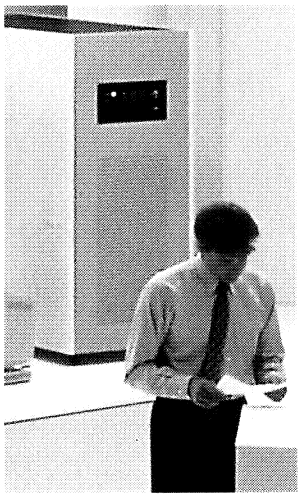

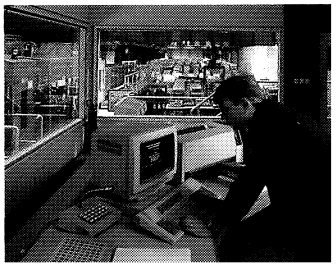

Designed for the '90s

Powerful workstations reach a new level of performance through integration in high-speed networks

Innovative systems and services to advance your business vision

Scenarios such as these can be multiplied many times, depending on country, industry and business strategy. In all cases, customer value drives our efforts to help you realise your enterprise goals. No matter what your industry or business plan may be, here are some of the ways that we can help you empower your enterprise:

Firstly, there is the growth potential and performance power of the new IBM System/390 computing family. System/390 includes the new Enterprise System/9000 (ES/9000)* processor models, as well as enhanced system hardware and new operating system software.

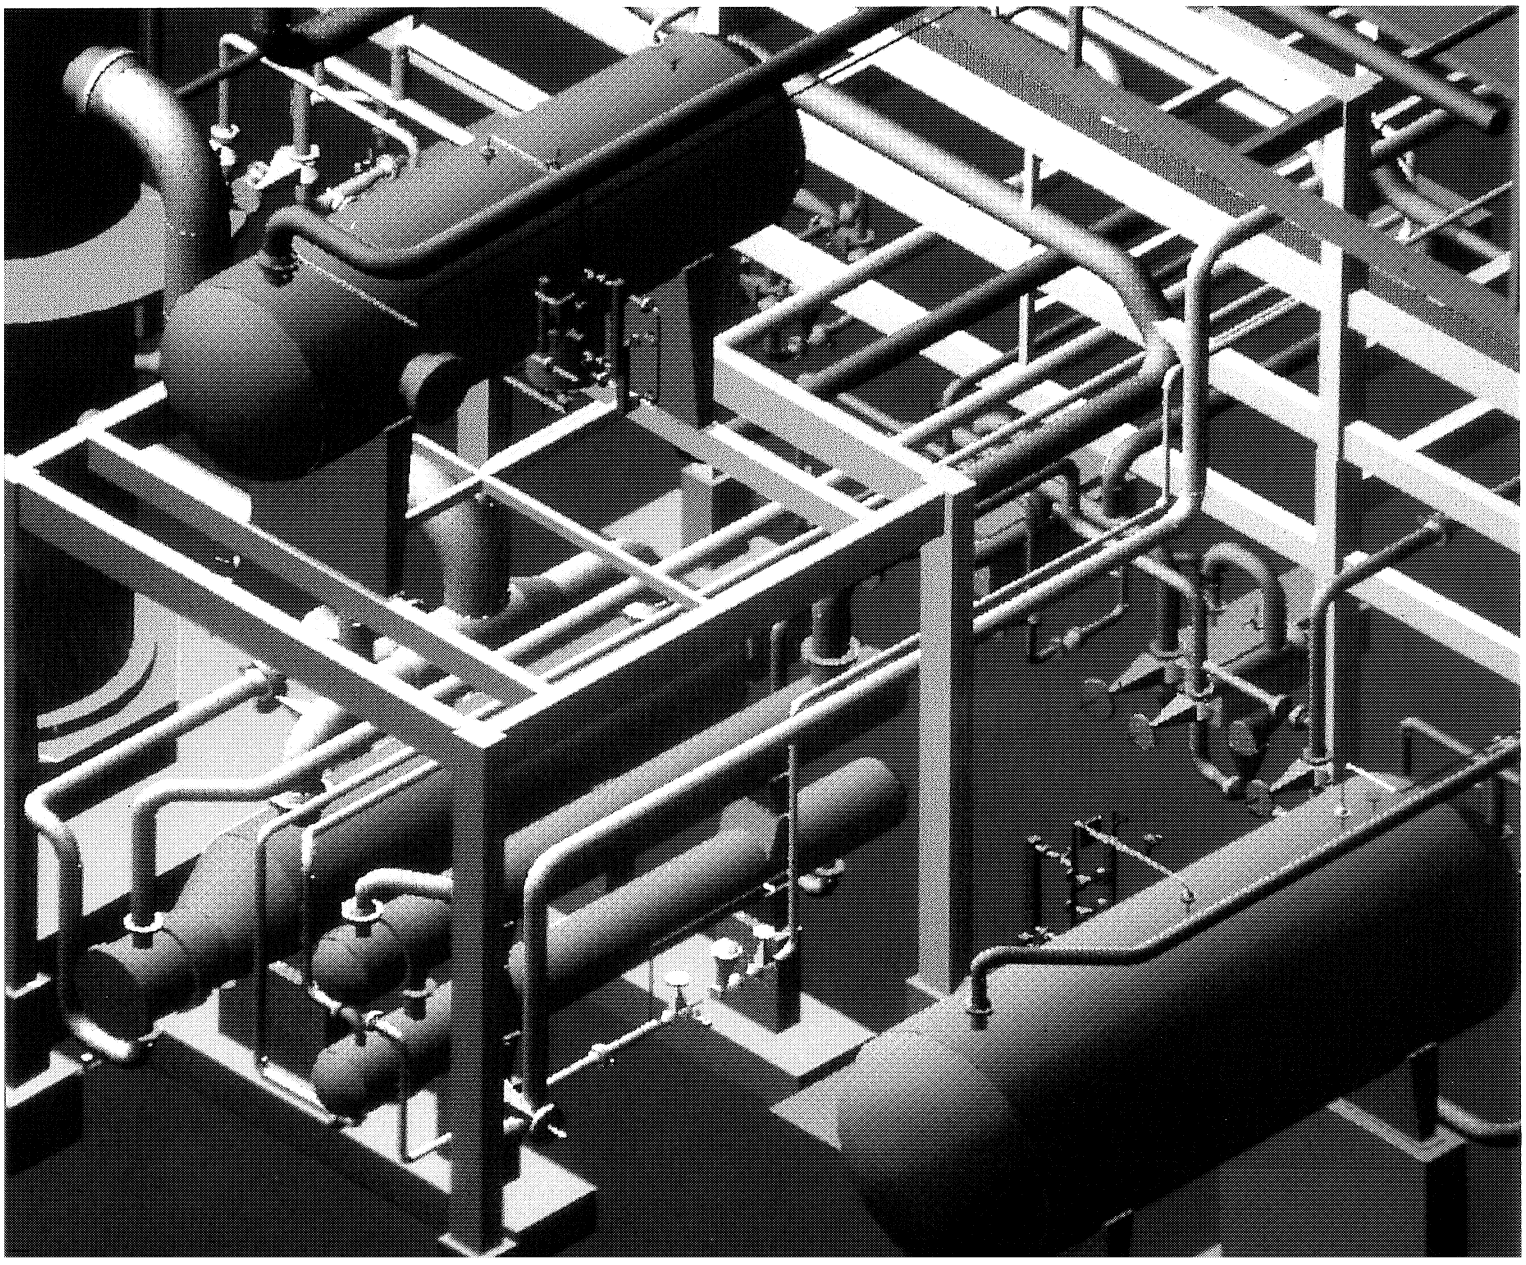

System/390's broad range of processor models offers high performance and significant growth potential, enriching your existing investment in systems, data, applications and personnel.

Your employees can use the System/390 family to gain wider, more comprehensive information access, within and beyond enterprise boundaries. This gain is the result of our enriching and streamlining three operating systems – MVS, VM and VSE – and using common communication protocols and network architectures in the design of our new systems.

You can also gain a more comprehensive managerial reach over your network performance and resources. This capability derives from network solutions that integrate diverse local area networks and offer comprehensive graphic views of the resources they contain ... new system management directions ... and varied security options.

Designed for the '90s, these powerful IBM computing solutions incorporate many valuable technological innovations:

- A unique integrated cryptographic feature
- Fast, lightweight fibre optic cabling
- Environmentally efficient systems.

Our products integrate the systems that process your enterprise information with other office systems, departmental systems, client-server LANS, workstations, workgroups and departments. For example,our Enterprise Systems Architecture /390 (ESA /390)* employs logical and physical interconnections that enable designers working at varied geographical locations to gain faster access to electronically dispersed engineering changes. As a result, they can increase productivity and shorten the product development cycle.

Accommodate cooperative and distributed processing

The new IBM ES/9000 processor models and operating system software also accommodate cooperative and distributed processing. As a result, our new systems move you closer to the day when any individual computing system or workstation in your enterprise can interact with any other – even if you have hundreds of local area networks, each with hundreds of workstations.

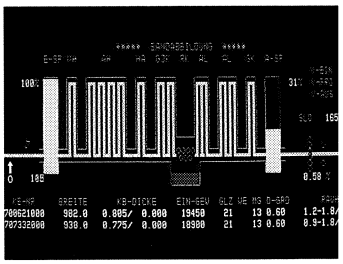

Your information systems organisation can use standard IBM platforms to develop enterprise-wide applications and link appropriate enterprise personnel with your outside clients and suppliers. And workgroups in one enterprise department or division can share information with those in other departments or divisions, regardless of whether they are on the other side of the site, across the country or in a remote corner of the globe.

It no longer matters if terminals belong to different departments or are spread across numerous sites at home and abroad

Key computing attributes of the empowered enterprise

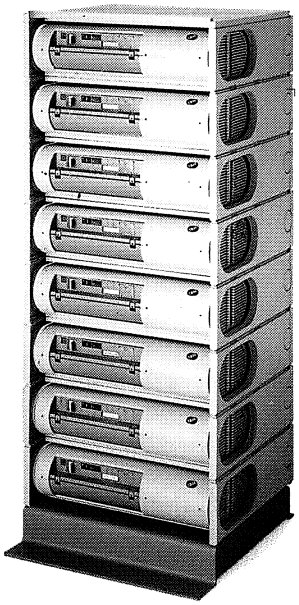

The possibilities available through fibre optics mean a system's I/O devices no longer have to be concentrated in one place

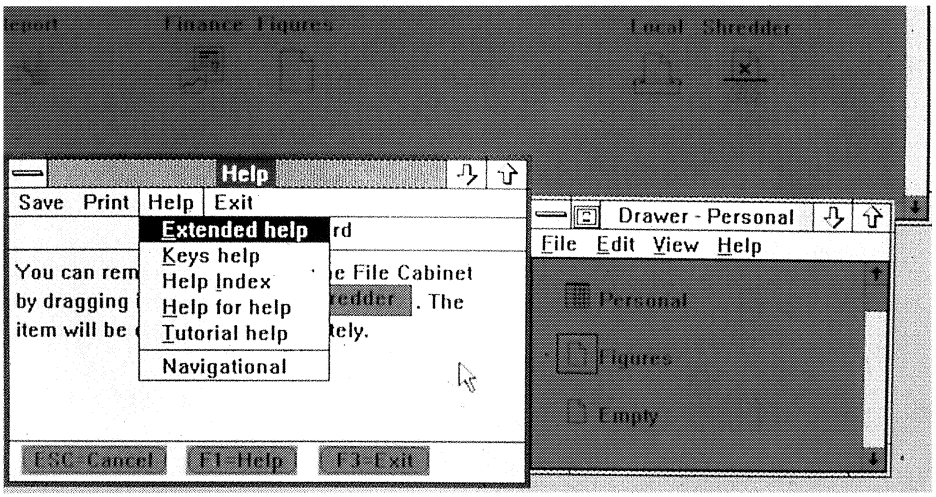

Operational flexibility
Your IBM System/390-empowered enterprise can accommodate changes far more flexibly than ever before. Applications can be shifted from one large host system to another ... devices such as disks and tapes can be changed with minimal disruption to regular operations ... and computing resources can be located wherever space permits.

The fibre optic technology underlying channel connections in our systems and controllers requires far less room, an appealing alternative to the limits imposed by the heavy, bulky cabling that has long served the industry. Our new channel technology also permits data transmission over distances of up to 9 kilometers (5.4 miles) – a significant increase over the former 122 meter (400 feet) standard. As a result, you can now locate system processors in one part of your business complex, and devices for those same systems can occupy space in another building or complex. You gain further flexibility because you can access data from a variety of systems and locations.

IBM's new systems also provide your enterprise with a cooperative processing facility, coupling the benefits of personal computing with the performance, power and control of host-based computing. Cooperative processing allows applications to be split between workstations and hosts.

Increased security options to protect your enterprise
You want your information to be accessible; yet, the more critical that information is to the operation, the more sensitive it becomes. The System/390-empowered enterprise can provide the degree of security protection you require, at workstation, host and network levels, combining ease of use with system-managed security.

Broader, simpler system and network management
Our new products are designed to respond to your requirement for broader, simpler, enterprise-wide system and network management. If improved operational management, lower costs, and increased system availability and responsiveness are high on your agenda, our new products and services can meet your needs.

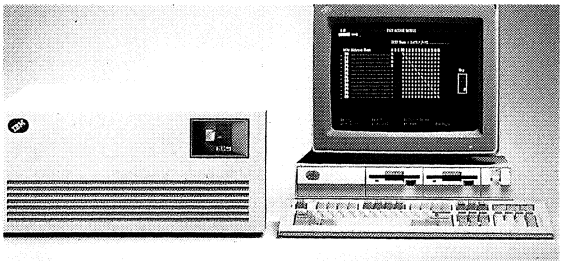

Powerful system growth and expansion
The 18 models in the IBM Enterprise System/9000 (ES/9000) family offer outstanding room for growth. This family includes compact, rack-mounted models, as well as both air-cooled and water-cooled, frame-mounted models, all sharing a common architecture.

IBM's expanded VSE/ESA software offering for the ES/9000 family accommodates fast business growth, while providing the option to migrate your enterprise more easily into the MVS/ESA operating system environment. The ES/9000 family offers many upgrade options from your existing systems, enhancing your IBM computing investment.

Advanced applications
Information processing has moved from automating 'back-office' applications such as accounting ... to front-office applications such as on-line reservations systems ... to advanced applications that can give you a competitive edge, like image processing and computer-integrated manufacturing.

Image processing for increased productivity
One leading bank has described the potential of image processing in the financial community as "the most significant banking development since the computer". Today, banks can capture cheque images at the rate of 40 per second. And now, IBM is making such applications more widely available by basing them on standard system platforms.

As a result, image processing is gaining popularity in other industries and practices as well. Advanced image processing can include photos, engineering drawings or even medical images, as well as cheques and paper. This advanced application area offers the potential to increase transaction volumes handled by individual employees by as much as 30 percent.

Expanded capability for technical applications
Executives are sometimes surprised to learn that technical computing procedures are as useful in commercial applications, such as financial forecasting and production scheduling, as they are in complex scientific applications. Such applications often require the processing of vast quantities of stored information. Because the ES/9000 systems also offer fast storage access, they can act as both compute- and file-servers, working with high-speed networks and powerful programmable workstations.

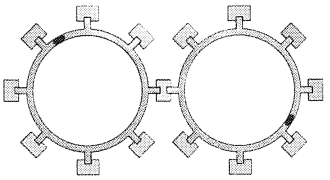

The new ES/9000 family offers significantly improved scalar capabilities and, on most models,a vector facility as well. This increases high-speed computational capabilities for more numerically intensive scientific and technical applications. Our largest ES/9000 model can be expected to offer from 2.0 to 2.8 times the computing capability of our current largest system with the Vector Facility.

A well-designed concept for the protection of data provides a high degree of security

Achieve your goals using our solutions

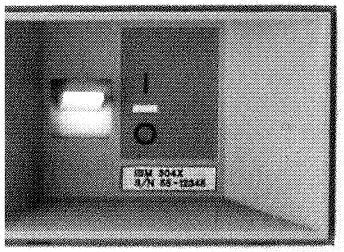

Bottom-line benefits
Together, IBM's ES/9000 processor family, efficient channel architecture, powerful operating system software, diverse security offerings, communications and networking products and advanced applications constitute a complete system approach to satisfying the information systems challenges inherent in supporting your business goals. With these products, you can realise the bottom-line benefits of such cost-saving innovations as the implementation of fibre optic technology.

For example, one company estimates that it can save millions of pounds because this technology does away with their need to construct a new data centre ... and it can save millions more by increasing systems availability at all levels. In addition, these products reflect IBM's continuing commitment to high-quality solutions that answer your needs and concerns.

Applications in the areas of investment and production planning, chemical engineering, structural analysis and simulation require extra performance

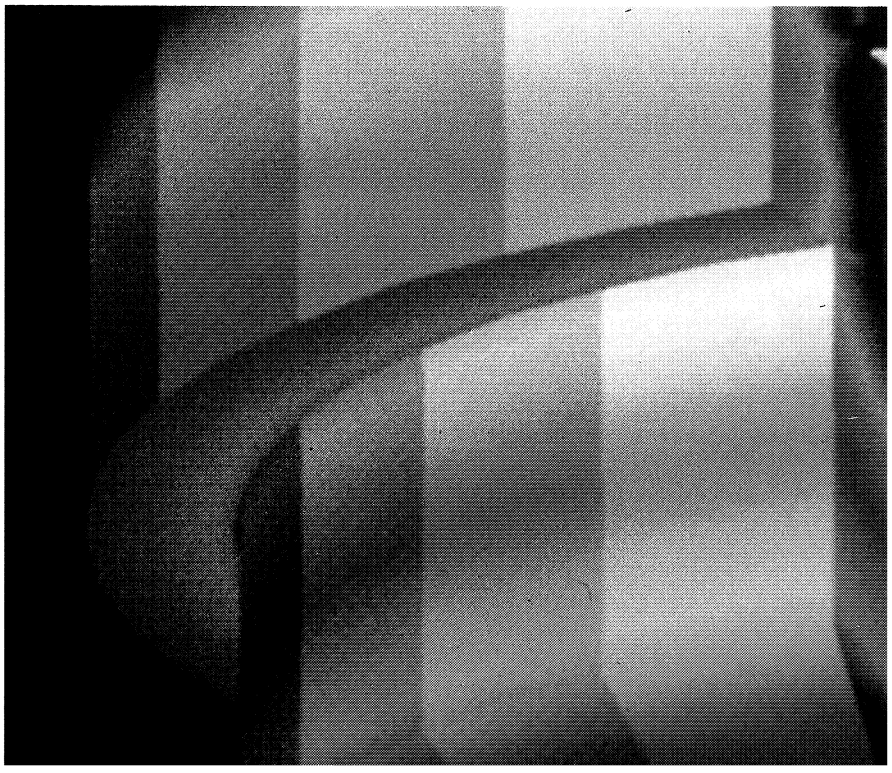

Worldwide support for products and services
IBM's leading-edge solutions and service offerings span the total computing spectrum, including home computing, powerful workstations, numerically intensive and image-intensive processing, powerful communications and storage offerings, and more.

We service hardware and software systems almost anywhere in the world. And our products are backed by our comprehensive support programmes. These include consulting, capacity planning, network planning, site preparation, education, software and hardware conversion, disaster recovery, hardware and software maintenance and facilities management.

And as you choose specific IBM offerings to provide solutions for your enterprise, IBM can fulfill your financing needs – whether funding is required locally or around the world.

Your IBM representative will be happy to give you details about how the System/390 family and related products and services in systems, software, communications and applications can empower your enterprise from top to bottom.

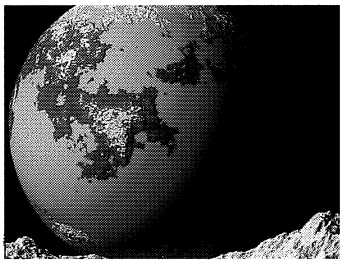

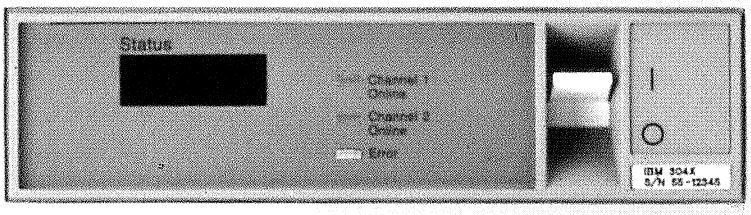

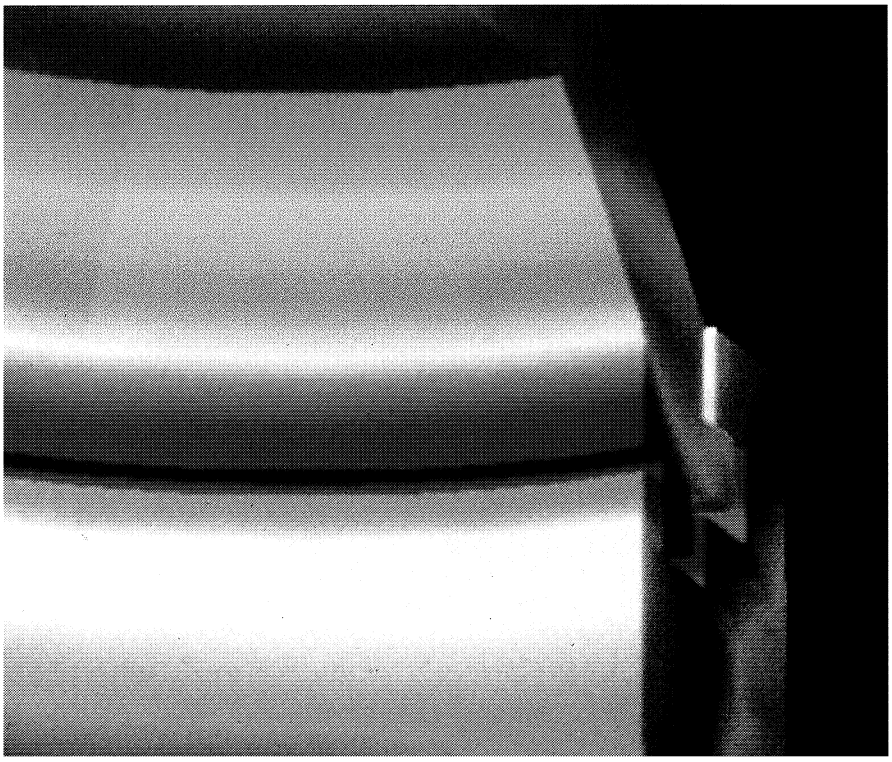

Solutions and services covering the whole spectrum of information processing

GU20-0082

IBM United Kingdom Limited
National Enquiry Centre
389 Chiswick High Road
London W4 4AL
Telephone 081-747 0747

Registered in England: No. 741598
Registered Office: PO Box 41
North Harbour, Portsmouth
Hampshire, PO6 3AU

Printed in England by
Cedar Colour Limited
Chandlers Ford, Hants

Growth through innovation

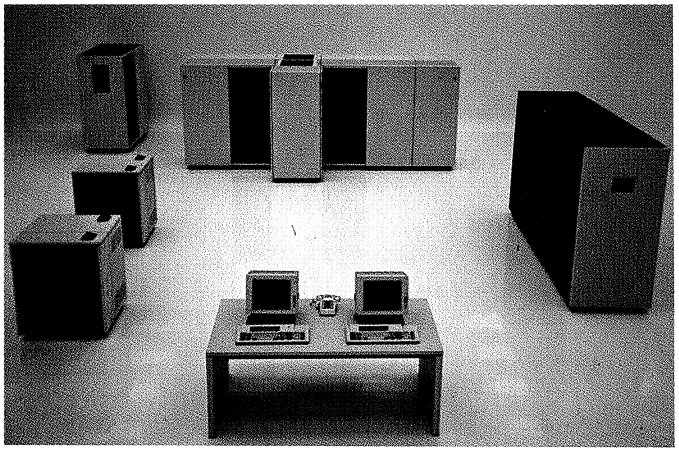

Growth through innovation

The IBM 308X processors, from the Model CX to the Model QX, represent the broadest field-upgradable range in the large systems industry today. The addition of two new processors, the IBM 3090 Models 200 and 400, increases the options upwards to meet the needs of our faster growing large customers such as those who have a need for growth beyond the 3084, or those with large engineering and scientific workloads, or those needing a very large VM system.

This growth is not, however, achieved at the expense of compatibility. The 308X and 3090 processors share common architecture, packaging and technology.

The capabilities of the 308X and 3090 processors are further enhanced by developments in systems software, storage products, and assists.

The same new software supports all 308X and 3090 processors; new facilities include the SIE and SORT Assists which, provide improved performance.

Like the 308X, the 3090 provides support for both the standard and extended capability 3380 disk subsystem, as well as much improved interconnect capabilities with the IBM 3044 Fibre Optic Channel Extender Link.

Customers can continue to install and upgrade through the 308X range, as the common foundation of Thermal Conduction Module technology and the IBM Extended Architecture offers continuity across both the 308X and 3090. Customers can benefit from a broader range of growth options, the enhanced software and the innovative assists. Those customers who need further growth can install 3090 processors when necessary.

Building on the base

The IBM 308X processors, designed with the TCM and 370-XA architectural innovations, form the cornerstone of IBM's large system offerings. IBM is continuing to expand on this foundation with the 3090 processors and latest releases of the Extended Architecture (XA) software.

Extended Architecture is the foundation for IBM's large and intermediate processors, providing the base for IBM large system's growth. It is further exploited by the latest releases of MVS/XA which support the Extended Recovery Facility (XRF), and also the newly announced VM/XA Systems Facility.

MVS/XA Extended Recovery Facility (XRF) gives the capability of providing a typically non-disruptive service for end users with XRF supported terminals. It enables an IMS system running on an active processor to have automatically initiated switches to an alternate system on a different processor in the event of an outage of the active system, and in many cases the IMS end user may be unaware of the disruption.

VM/XA Systems Facility significantly improves the peformance of a preferred guest running under VM/XA Systems Facility.

These new features, along with other enhancements, are available on all the 308X and 3090 processors as shown in the diagram. This underlines IBM's commitment to its large systems customers, by providing growth through evolution and coexistence.

Extended Architecture – the foundation for growth

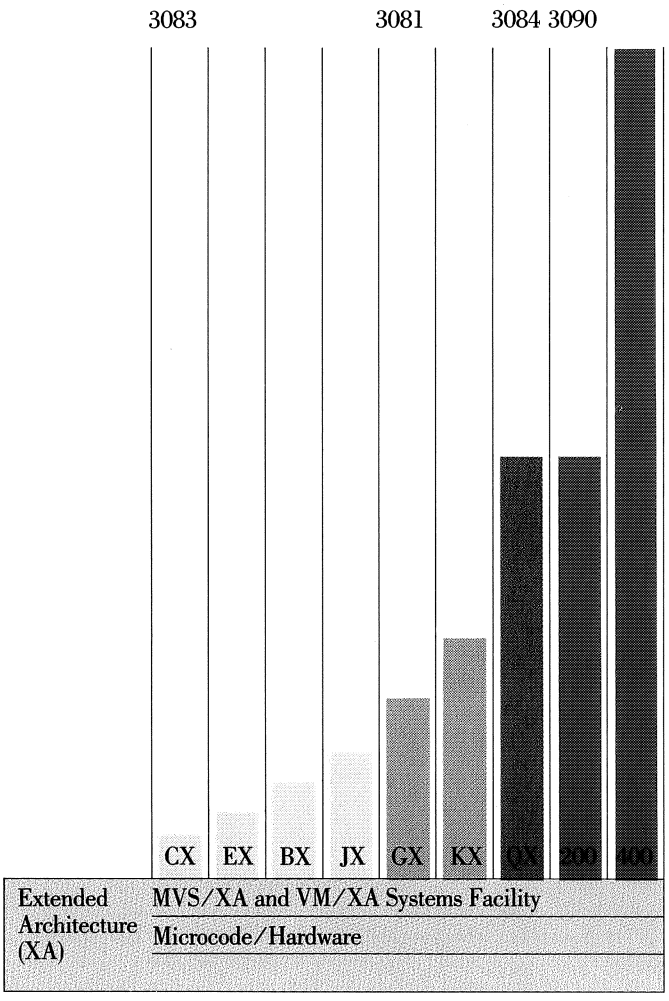

The Thermal Conduction Module (TCM) – technology for the 80s

The packaging of circuitry in a processor is one of the most significant contributors to processor speed, availability, and price/performance. The Thermal Conduction Module (TCM) is IBM's exclusive packaging innovation, a technology introduced in 1981 on the IBM 3081 Model D Processor. Since then, it has formed the basis for the enhanced 308X models announced in 1984 and now for the IBM 3090 processors.

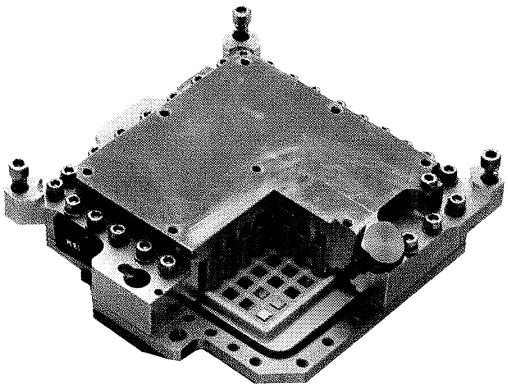

TCM technology implementation

1981 **308X**

1984 **308X X Models**

1985 **3090**

Growth through evolution and coexistence

The IBM 308X and 3090 processors offer the same assists and support the same new software.

Hardware

In addition to the common TCM-based technology, the 3090 and 308X processors support recent hardware announcements such as:

- The IBM 3044 Fibre Optic Channel Extender Link
- The extended capacity IBM 3380 disk models, with improved capacity and performance
- The enhanced IBM 3880 Model 23 Cache Controller.

Software

The performance improvements offered by the VM/XA Systems Facility are provided by the SIE Assist and software enhancements. SIE, without which the VM/XA Systems Facility cannot work, is further enhanced by the SIE Assist available on 308X and 3090 processors.

The DFSORT program product, which uses the new SORT Assist, has also been enhanced. This offers a major performance improvement compared to IBM SORT programs currently available.

Summary

All the 308X and 3090 processors are supported by the new releases of MVS and VM.

The compatibility of the 308X and 3090 processors as described above allows customers to decide on their own growth rates, budgets, financing options and so on, rather than any consideration of technology or availability of function.

IBM 308X

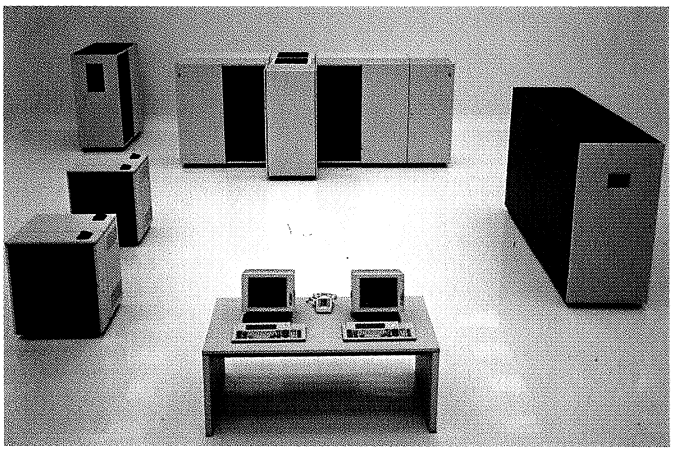

IBM 3090

IBM 3090 Processor complex

Technical information

- **Performance**
  - IBM 3090 Model 200: 1.7 – 1.9 x IBM 3081 Model KX MVS/XA – Commercial
  - IBM 3090 Model 200: 1.9 – 2.9 x IBM 3081 Model KX MVS/XA – Engineering and Scientific
  - IBM 3090 Model 400: Performance (IER) is expected to be in the range 1.7 – 1.9 x IBM 3090 Model 200.
- **System Components**
  The IBM 3090 processors are available in two models: the dyadic Model 200 and the four-way Model 400. The Model 200 can be field upgraded to the Model 400.

  The 3090 processor complex comprises the following units:
  - IBM 3090 Processor Unit
  - IBM 3089-3 Power Units (or equivalent)
  - IBM 3097 Power and Coolant Distribution Unit
  - IBM 3092 Processor Controller
  - IBM 3370A2 Disk Storage Devices
  - IBM 3180 Models 140 and 145 Display Stations.

Units are distributed as follows in the two processor complexes:

IBM 3090 Model 200	
Processor Unit	2 central processors
Central storage	64 Mb
Expanded storage (optional)	64 or 128 Mb
Channels	32,40,48
IBM 3092-1 Processor Controller	1
IBM 3097 Power and Coolant Distribution Unit	1
IBM 3089-3 Power Unit	2
IBM 3370A2 Disk Storage Devices	2
IBM 3180 Model 145 Display Station	2

IBM 3090 Model 400	
Processor Unit	4 central processors
Central storage	128 Mb
Expanded storage (optional)	128 or 256 Mb
Channels	64,80,96
IBM 3092-2 Processor Controller	1
IBM 3097 Power and Coolant Distribution Unit	2
IBM 3089-3 Power Unit	4
IBM 3370A2 Disk Storage Devices	2
IBM 3180 Model 145 Display Station	3

Expanded storage provides near equal performance to central storage at a reduced cost. Expanded storage and central storage are both managed by the operating system and are transparent to application programs. The implementation of larger storage reduces paging requirements thereby providing a basis for increased systems throughput.

(1 Mb = 1,048,576 bytes)

By Appointment to
HRH The Prince of Wales
Suppliers of Typewriters
and Word Processing Equipment
IBM United Kingdom Limited
Portsmouth

IBM United Kingdom Limited
389 Chiswick High Road
London W4 4AL
Telephone: 01-995 1441

Registered in England: No. 741598
Registered Office: PO Box 41
North Harbour, Portsmouth
Hampshire PO6 3AU

Printed in England by
Gros Monti Limited
Ashford, Kent

An Investment For Growth

An Investment for Growth

The acquisition of a new computer system is a significant investment of your company's financial and human resources. Like all investments, it must be evaluated carefully in terms of return, flexibility and risk. IBM has designed the 308X Processor family to provide the optimum investment for your current and future needs. This brochure introduces you to the members of the family and their innovative design, technology and architecture.

Meeting Today's Need

Within the IBM 308X family there are processors with the power to satisfy a wide range of computing requirements. Equally at home in commercial, engineering, scientific or academic environments, 308X Processors are designed for both batch and interactive applications—from a few hundred to several thousand interactive or online users. With thirty-one models in the range—offering an extensive choice of processing power and memory capacity—your initial investment can be closely matched to today's needs.

Satisfying Tomorrow's Requirements

The growth of computer applications is seen by many as an inevitable requirement for improved productivity—the key to success in a competitive world. But predicting the rate of growth is far from easy. A major aspect of the IBM 308X design is the ability to upgrade an installed system on site, as and when the need occurs. Even the smallest IBM 308X Processor may be upgraded in convenient steps to provide close to ten times its original processing power. This growth path can protect your investment against frequent hardware and software replacements.

The IBM 308X Family

There are three basic model types in the IBM 308X family ranging, in increasing power, from the 3083 uniprocessors through the 3081 two-way multiprocessors to the 3084 four-way multiprocessor. A unique design concept allows each model to be upgraded to the next, usually in a few hours and, except for a 3081 to 3084 upgrade, no increase in floor space.

The IBM 3083 Processor

Within the IBM 308X family, the IBM 3083 provides an ideal entry point into large scale computing. The 3083 has a uniprocessor design with one central processor for instruction execution, plus independent elements for memory and channel operations. Variations in the degree of internal overlap and sophistication of processor organisation provide different internal instruction execution rates for each of the four basic models. Other options include from 8 to 32 megabytes of memory and 8 to 24 input/output channels, depending on the model.

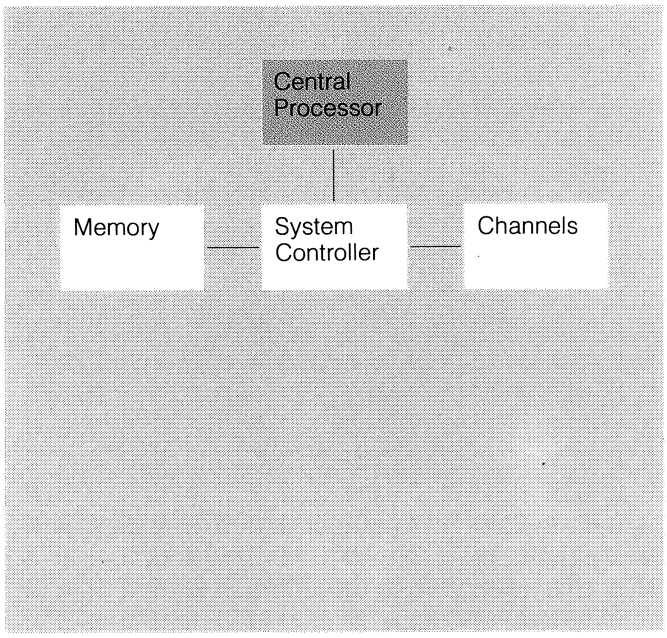

Innovative Technology

The 'building block' for all IBM 308X Processors is the unique IBM-developed Thermal Conduction Module (TCM) shown here in cut-away form. Each TCM can house over 90,000 circuits in a 125mm (4.9 in.) square. Connections between circuits are made via a multi-layer ceramic substrate containing up to 33 layers of wiring, virtually eliminating the need for discrete cables. Combined with a sophisticated cooling system which ensures that each component is maintained at the optimum temperature, the fully integrated packaging of the TCM makes a significant contribution to both reliability and performance.

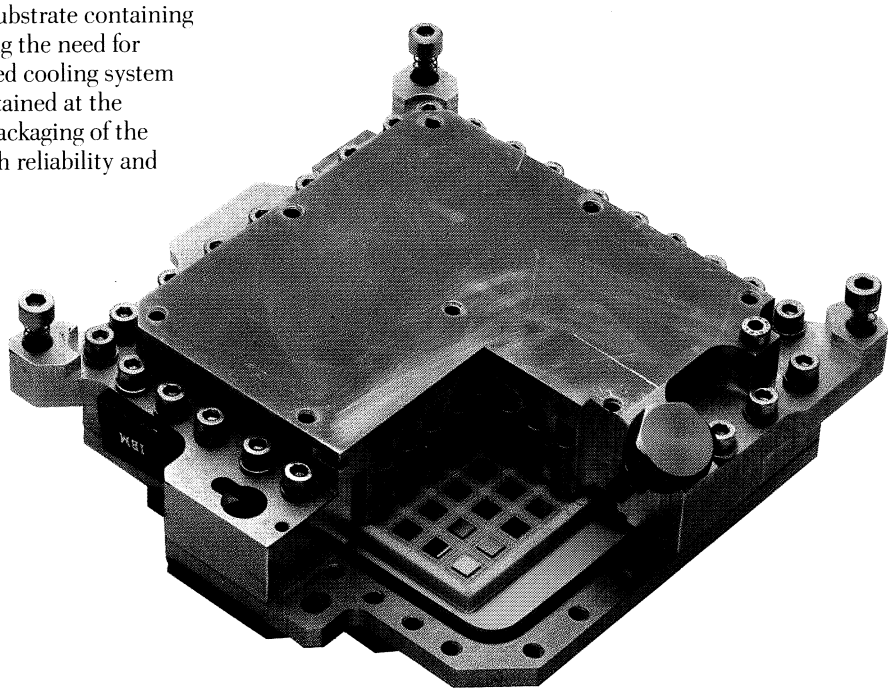

The IBM 3081 Processor

The addition of a second central processor and its associated support systems upgrades a 3083 uniprocessor to a 3081 dyadic processor. The term 'dyadic' reflects a particular design of two-way multiprocessor where the two central processors share common memory and channel subsystems via a system controller. The dyadic design offers high performance with enhanced availability at a lower cost than traditional multiprocessor implementations. The two basic models of the 3081 offer a choice of internal execution rates and, consistent with the higher level of performance, the maximum memory size option is increased to 64 megabytes.

The IBM 3084 Processor

At the head of the IBM 308X family is the 3084 four-way multiprocessor. The upgrade from a 3081 to a 3084 is accomplished by the installation of a second dyadic processor, tightly coupled to the first via the system controllers. As all elements of the processor are now fully duplicated, the 3084 can achieve very high levels of availability, with the possibility of running one half as a dyadic processor concurrent with maintenance operations on the other. Equally, the 3084 may be operated as two separate dyadic processors, each running its own control program. Processor balance is maintained with additional 96 and 128 megabyte memory options and 48 channels as standard.

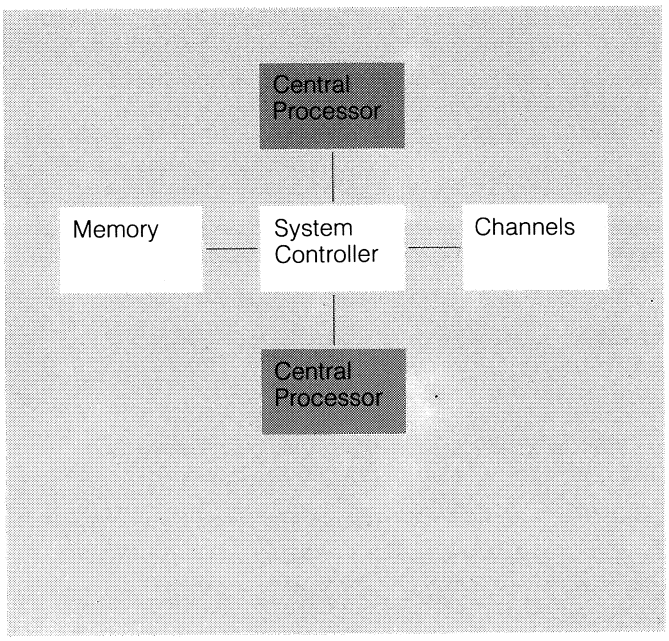

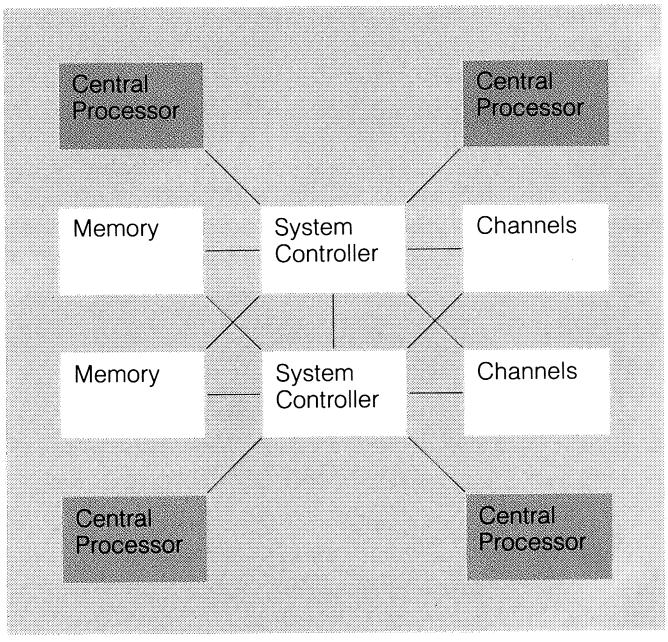

An Architecture For Growth

The ability to exploit the well established IBM System/370 Architecture and IBM's latest Extended Architecture (XA) is fundamental to the design of the IBM 308X family. Since the early 70s, System/370 Architecture has allowed IBM's customers to develop their applications and expand their systems in an evolutionary and compatible fashion. IBM Extended Architecture builds on this base by providing even greater function and availability while protecting your current investment in application software. Existing and new applications can run concurrently within IBM's Extended Architecture, which is designed to take IBM computing through the 1980s and beyond.

The IBM 308X family

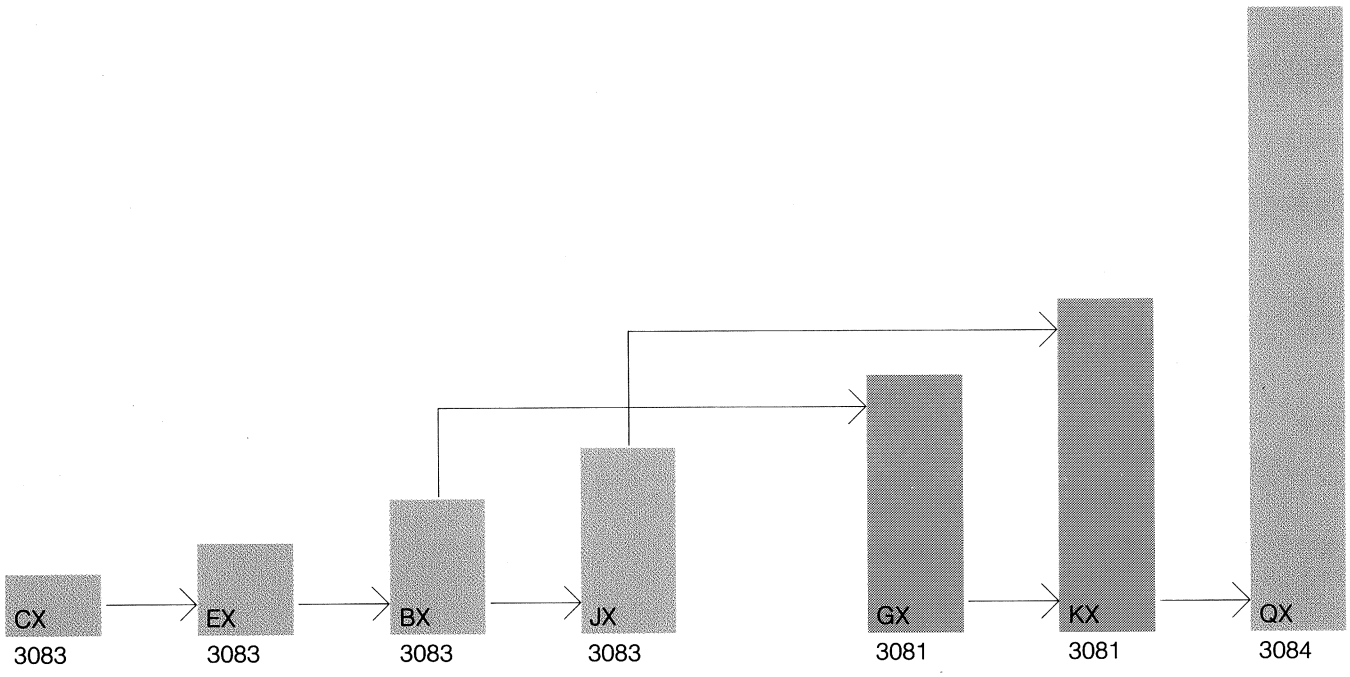

Main storage and channel options

Model		Main storage in megabytes								Number of channels			
		8	16	24	32	48	64	96	128	8	16	24	48
3083	CX	●	●	●	●					●	●		
	EX	●	●	●	●					●	●		
	BX	●	●	●	●					●	●	●	
	JX	●	●	●	●					●	●	●	
3081	GX		●	●	●	●	●				●	●	
	KX		●	●	●	●	●				●	●	
3084	QX				●	●	●	●	●				●

By Appointment to
HRH The Prince of Wales
Suppliers of Typewriters
and Word Processing Equipment
IBM United Kingdom Limited
Portsmouth

IBM United Kingdom Limited
389 Chiswick High Road
London W4 4AL
Telephone: 01-995 1441

Registered in England: No. 741598
Registered Office: PO Box 41,
North Harbour, Portsmouth
Hampshire PO6 3AU

This publication is for general guidance only. Please contact your IBM sales Representative for further information

Printed in England by
Gros Monti Limited
Ashford, Kent

IBM 3090
Processor Unit Models 200 and 400

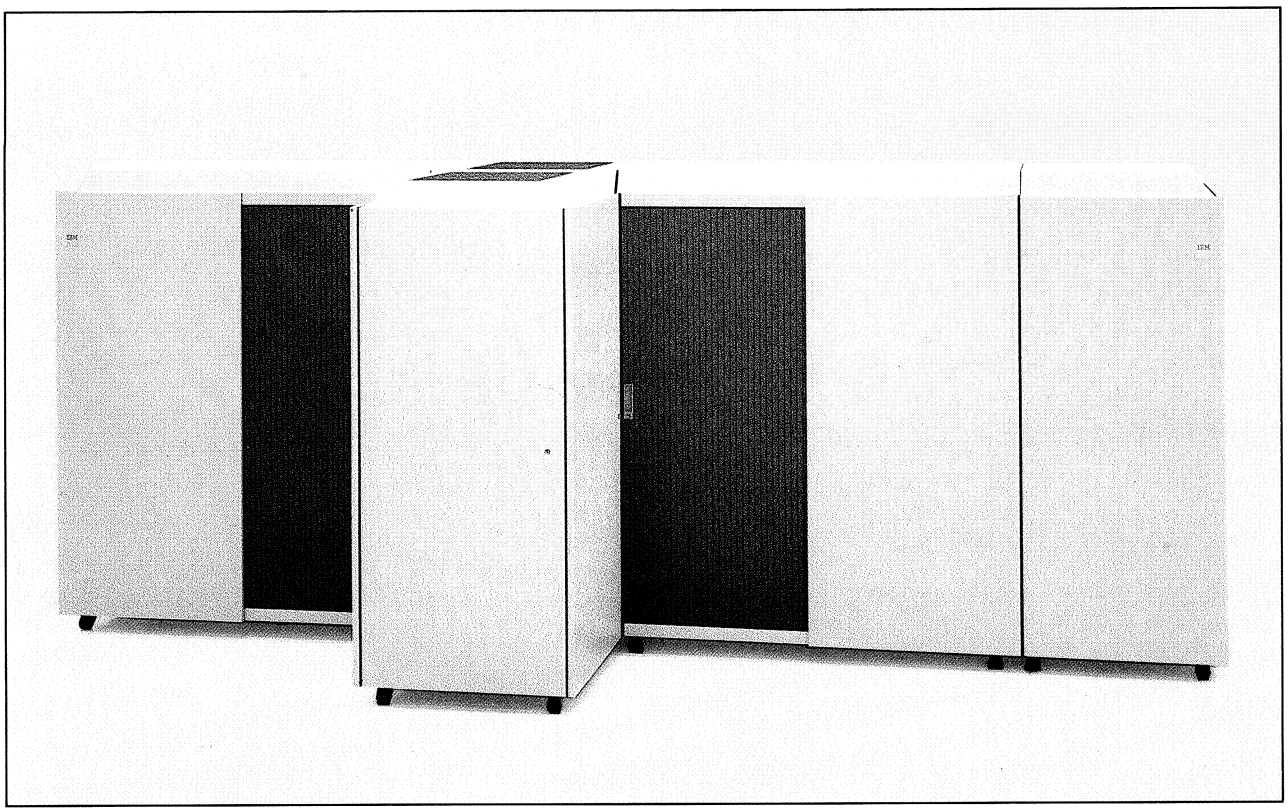

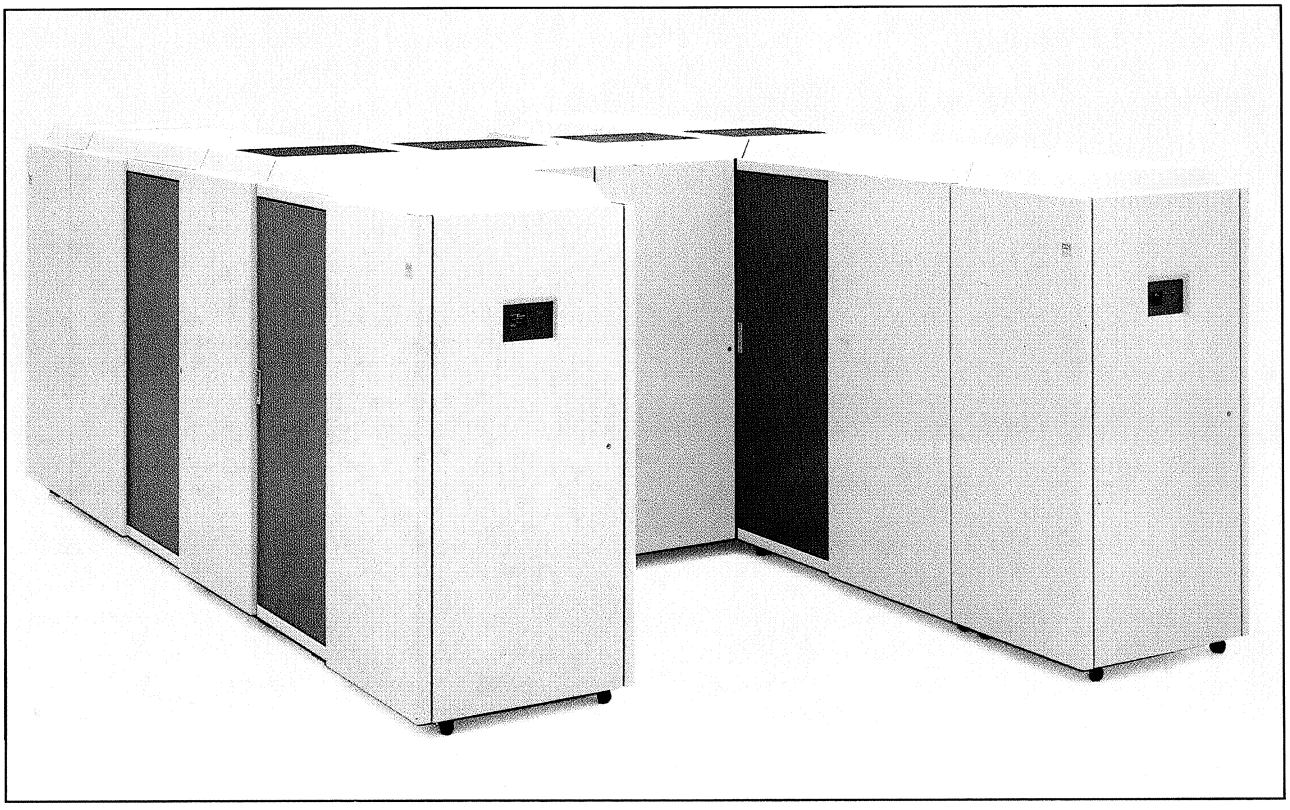

IBM

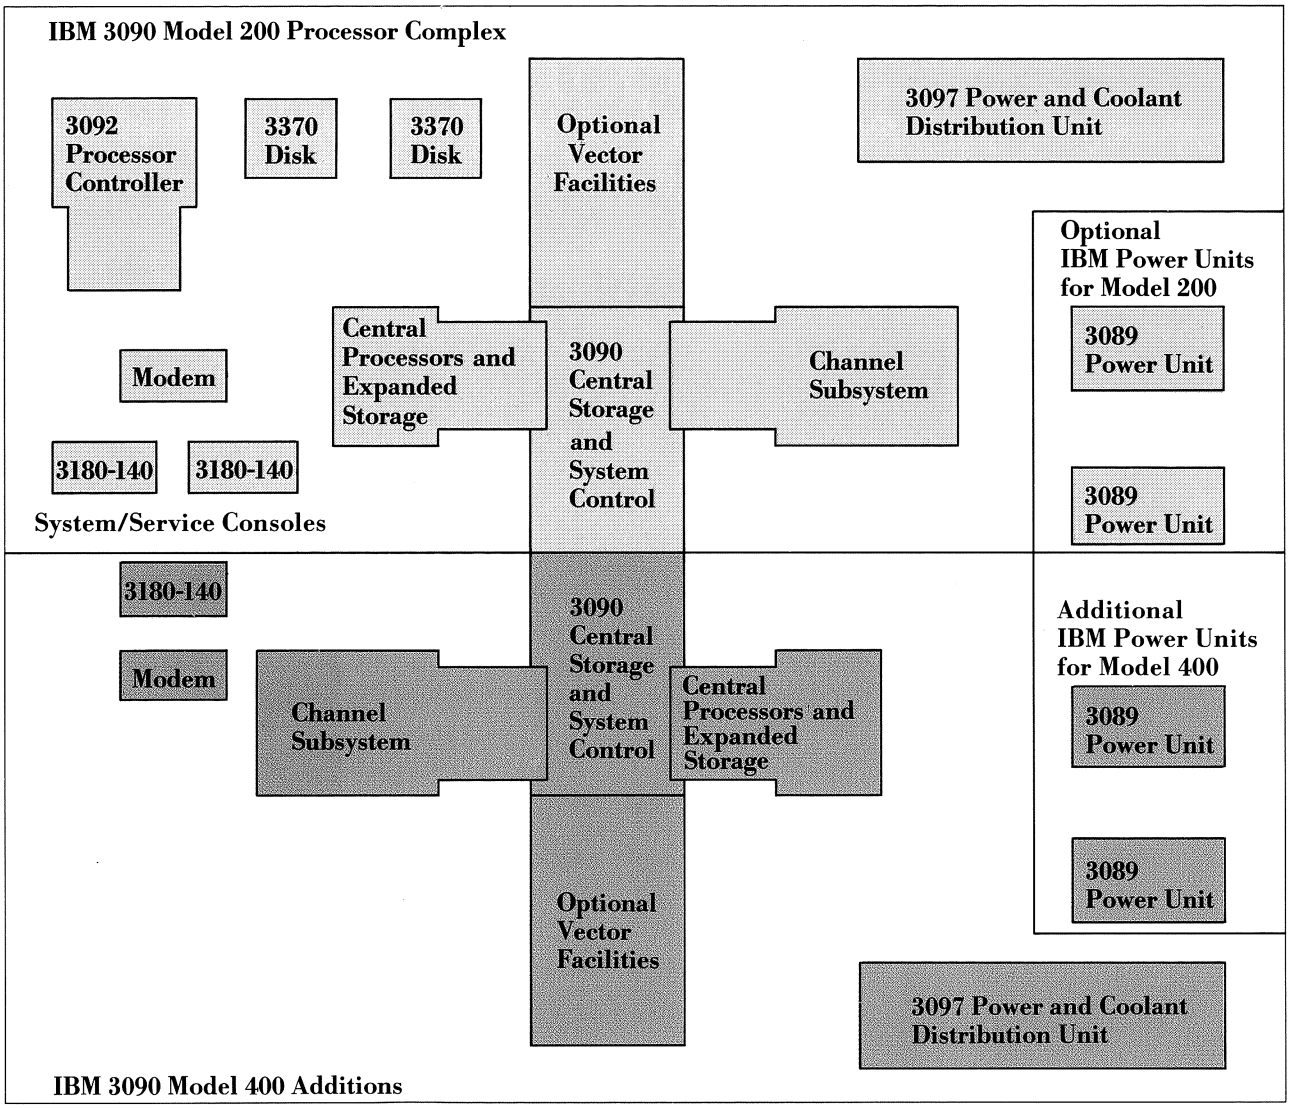

IBM 3090 Processor Configuration

An IBM 3090 Model 200 Complex comprises:

- One 3090 Model 200 Processor Unit with two central processors and one or two optional Vector Facilities.
- One 3092 Model 1 Processor Controller.
- Two 3370 Model A2 Direct Access Storage Devices.
- Two 3180 Model 140 Display Stations as system/service consoles.
- One 3864 Model 2 Modem or a 4800 bits per second (bps) switched network modem supporting an autocall/autoanswer feature.
- One 3097 Model 1 Power and Coolant Distribution Unit.
- Two 3089 Model 3 Power Units or equivalent 400 Hz power source.

An IBM 3090 Model 400 Complex is obtained by field-upgrading a Model 200. It comprises:

- One 3090 Model 400 Processor Unit with four central processors and one, two, three or four optional Vector Facilities.
- One 3092 Model 2 Processor Controller.
- Two 3370 Model A2 Direct Access Storage Devices.
- Three 3180 Model 140 Display Stations as system/service consoles.
- Two 3864 Model 2 Modems or two 4800 bps switched network modems supporting an autocall/autoanswer feature.
- Two 3097 Model 1 Power and Coolant Distribution Units.
- Four 3089 Model 3 Power Units or equivalent 400 Hz power source.

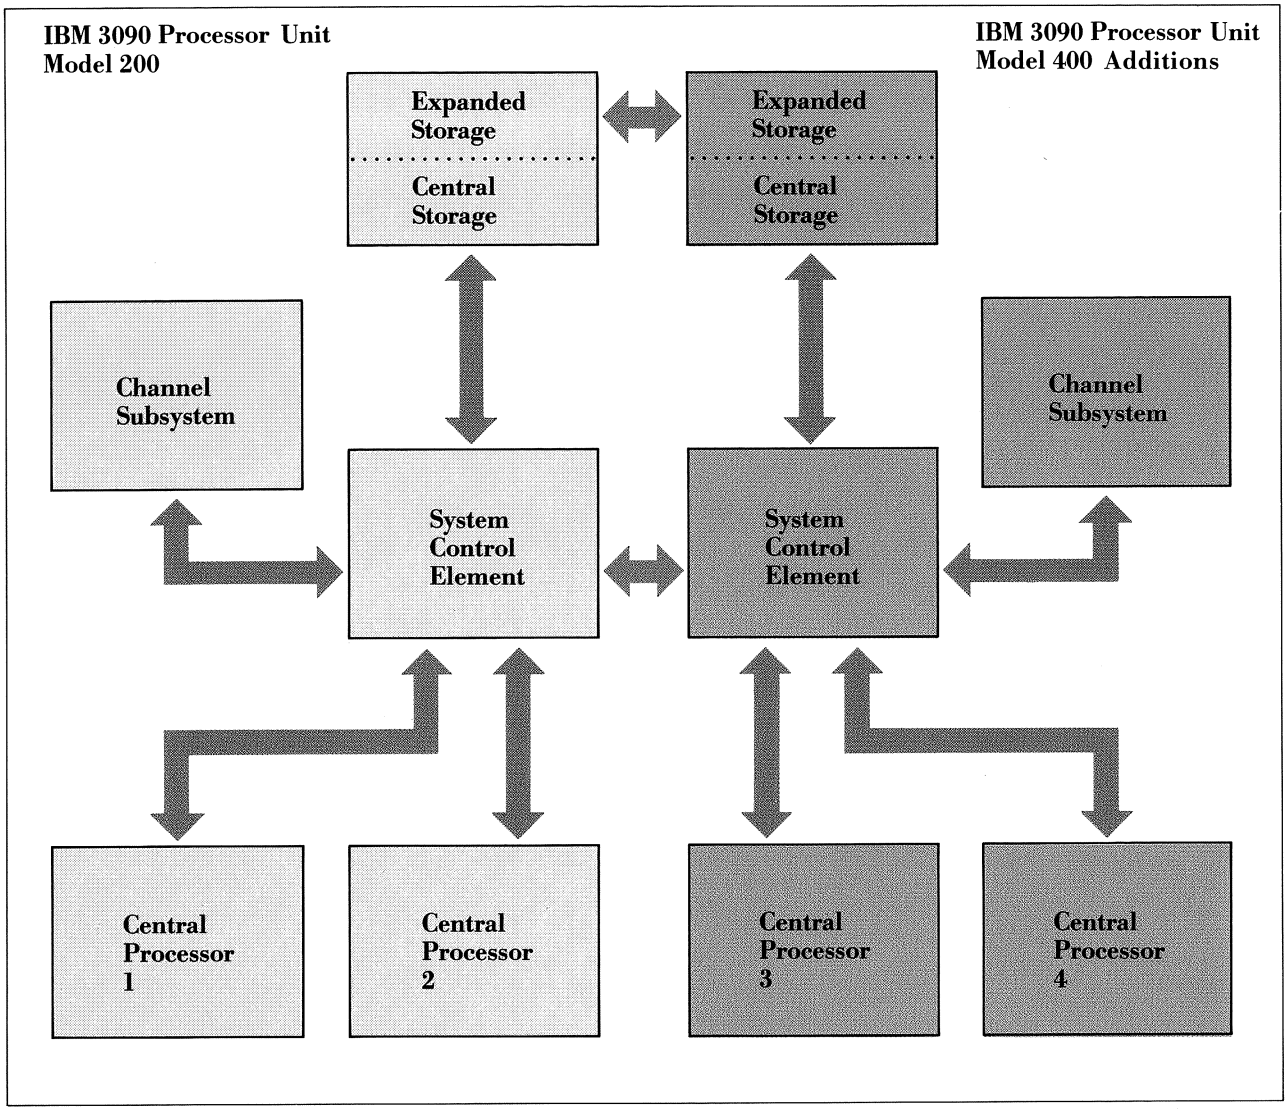

IBM 3090 Processor Unit Design

The IBM 3090 Model 200 and Model 400 Processor Units perform the data processing functions of the IBM 3090 Processor Complex and are the most powerful computers available from IBM.

All IBM 3090 Processor Units are composed of four main elements:

- Shared Central and (optional) Expanded Storage.
- System Control Element.
- Channel Subsystem.
- Central Processor.

The IBM 3090 Model 200 is a dyadic unit, which has:

- Two Central Processors.
- One of each of the other elements.
- It operates as a single logical system.

The IBM 3090 Model 400 is a four-way unit, which has:

- Four Central Processors.
- Two of each of the other elements.
- It can operate as:
  - A single logical system, in single image mode,
  - Two logical dyadic systems, in partitioned mode.

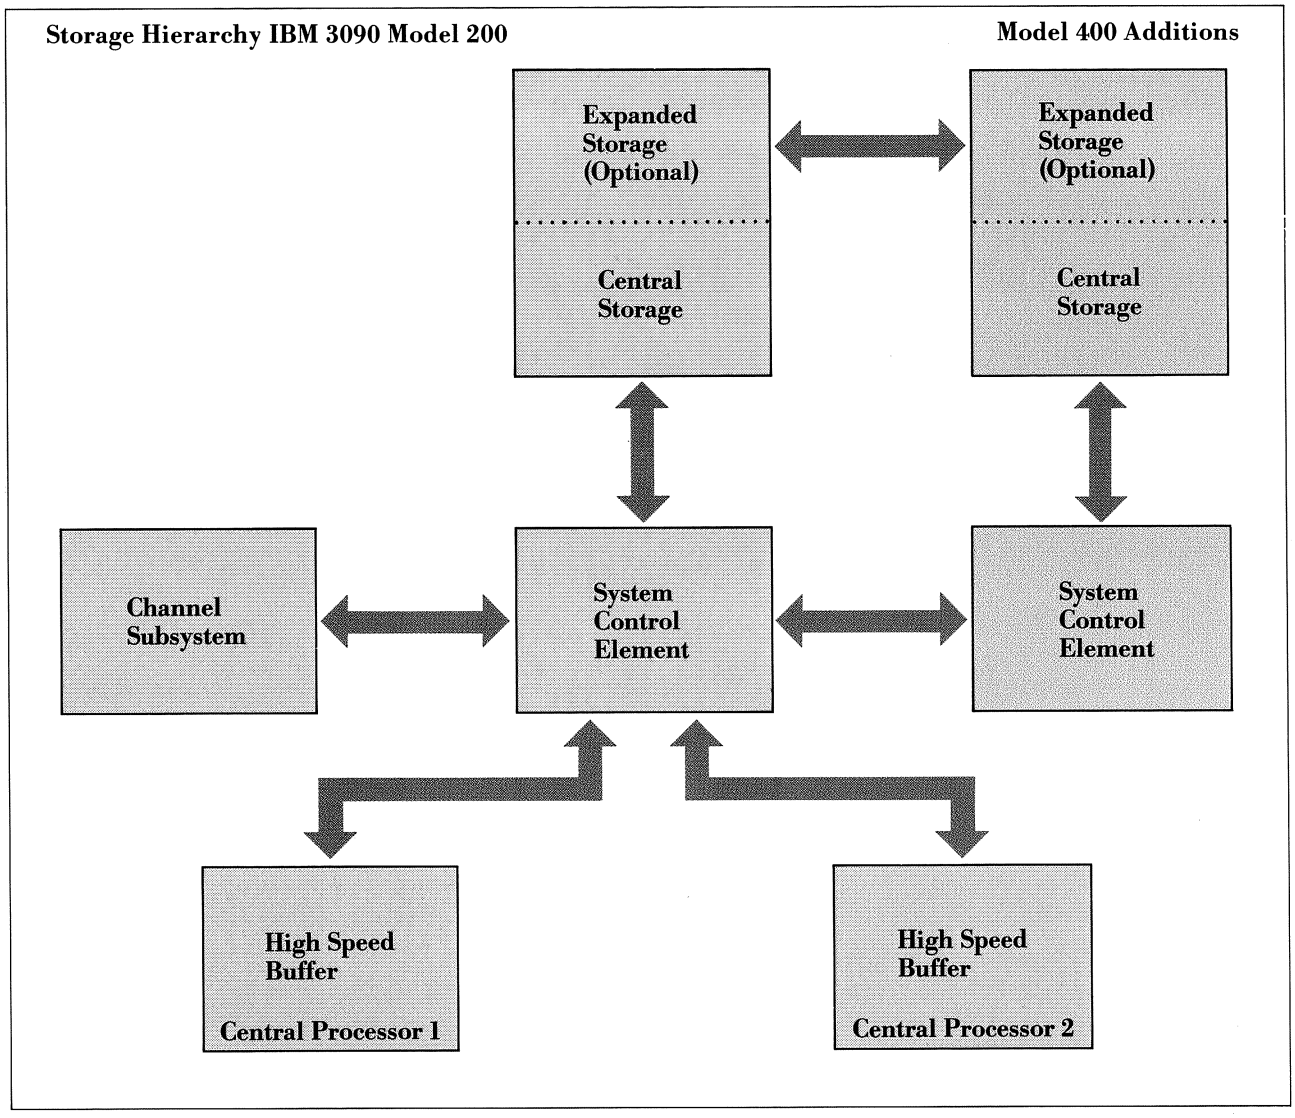
Storage Hierarchy IBM 3090 Model 200
Model 400 Additions
Expanded Storage (Optional)
Central Storage
Expanded Storage (Optional)
Central Storage
Channel Subsystem
System Control Element
System Control Element
High Speed Buffer
Central Processor 1
High Speed Buffer
Central Processor 2

Storage Hierarchy

The 3090 Processor Unit has a three-level storage hierarchy consisting of High Speed Buffer, Central Storage and optional Expanded Storage, managed by the System Control Element and Buffer Control Element.

System Control Element:

- Contains logic for data storage and retrieval for the processor complex.
- Provides the communication path between the processor complex and Central Storage.
- Communicates with and controls optional Expanded Storage.
- Performs error checking and correction.

High Speed Buffer/Buffer Control Element:

- Provides independent storage access function for each Central Processor.
- Contains a 64K-byte High Speed Buffer, Buffer Directory, Translation Look-aside Buffer (TLB) and Dynamic Address Translation (DAT) hardware.
- The High Speed Buffer is used to store active data and instructions and can be accessed by its Central Processor via a 64-bit data path.
- The High Speed Buffer is transparent to programs that are being executed.
- The buffer directory contains the real addresses in Central Storage of data in the High Speed Buffer.
- The TLB saves the translated real addresses of pages in Central Storage which have been referenced, avoiding the need for re-translation if subsequently accessed.
- DAT hardware performs high-speed translation of virtual to real addresses for storing in the TLB.
- Buffer-to-buffer data paths contribute to total processor complex performance.

Central Storage:

- Central Storage is shared by the Central Processors in the 3090 Processor Unit:
  - Two Central Processors in the Model 200 share 64 Mb.
  - Four Central Processors in the Model 400 share 128 Mb.
- A Hardware System Area is reserved for system information, including microcode, control information for Input/Output (I/O) devices, message buffers, tables and trace information.
- All single bit errors are detected and corrected. All double bit errors are detected and some are corrected.
- Dynamic page de-allocation permits failing page-frames to be de-allocated in 4K-byte increments under system control program control for deferred maintenance.
- Key-controlled Storage Protection provides both Store and Fetch protection, preventing unauthorised access.

Expanded Storage:

- Expanded Storage provides optional low-cost, large-capacity solid state storage for paging using IBM-developed 288K-bit chips:
  - 64 Mb or 128 Mb for the Model 200,
  - 128 Mb or 256 Mb for the Model 400.
- Transfer is exclusively to and from Central Storage in 4K-byte pages performed by the System Control Element at system control program request.
- Expanded Storage is transparent to user programs.
- All single and double bit errors are corrected. All triple bit errors are detected and some are corrected.
- Expanded Storage reduces the paging and swapping load on channels and I/O devices and improves system response and performance.

IBM 3090 Central Processor

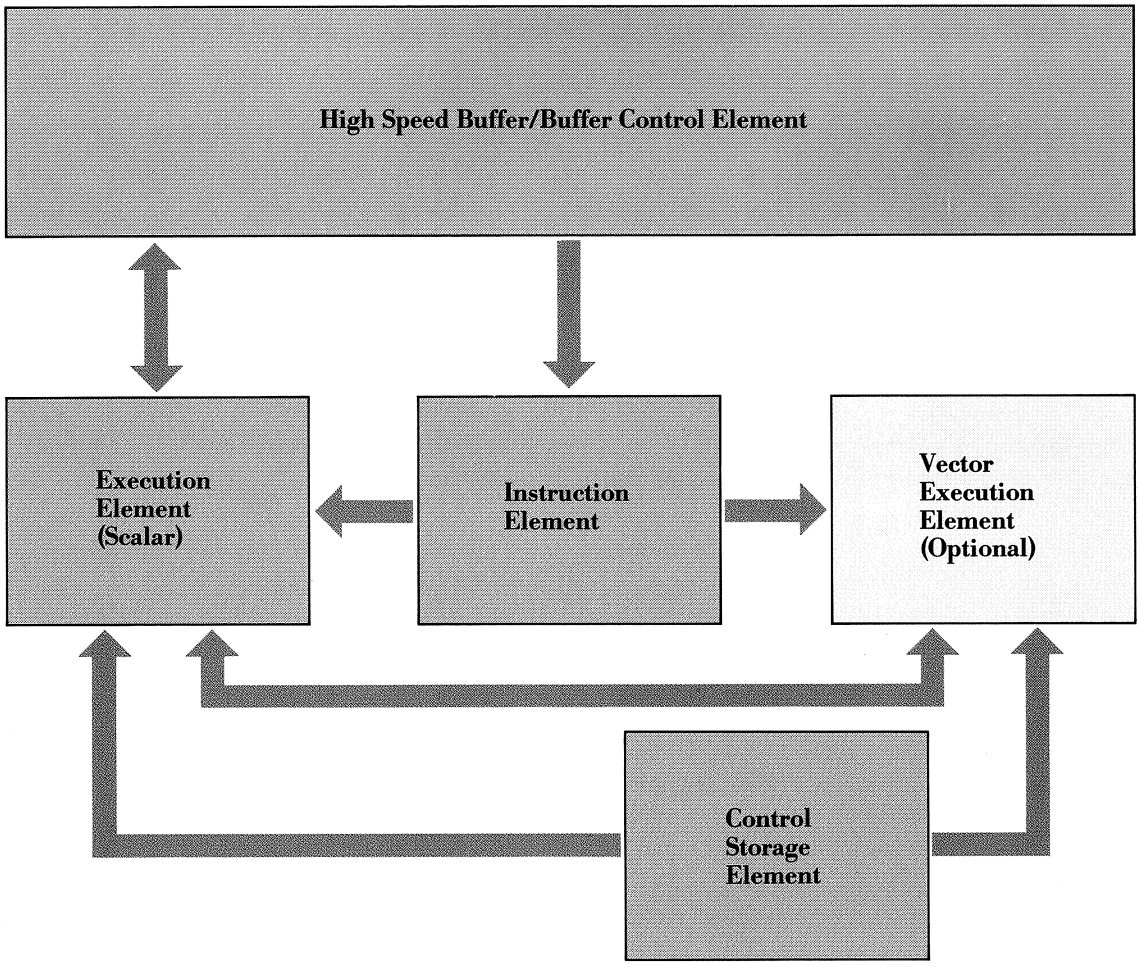

Central Processor

High performance processor design:

- 18.5 nanosecond cycle time.
- 64K-byte High Speed Buffer for fast access to instructions and data.
- Highly overlapped design permits Instruction Element to decode instructions in parallel with scalar or vector instruction execution.
- Multi-processor design permits two Central Processors in the Model 200 and four in the Model 400 to operate simultaneously.

Supports both System/370 and System/370 Extended Architecture:

- System/370 Architecture permits up to 16 Mb virtual storage per address space.
- System/370 Extended Architecture permits 2048 Mb virtual storage per address space and Bimodal Addressing for co-existence of System/370 programs.
- General purpose instruction set of 212 scalar instructions.
- 52 scalar floating point instructions.
- Short (32 bit), long (64 bit) and extended (128 bit) precision scalar floating point support.
- 16 general purpose registers (32 bits wide).
- Four 64-bit floating point registers, which can be paired for 128-bit operations.

Microcode assists for added performance:

- Control Storage Element stores instructions for microcoded operations and assists.
- Start Interpretive Execution (SIE) Assist improves performance for preferred guests running under VM/XA Systems Facility:
  · MVS/XA performs at 90-92% of native mode,
  · MVS/370 performs at 86-88% of native mode.
- Preferred Machine Assist (PMA) enables MVS/370 guests to run at approx. 95% of native speed under the VM/SP High Performance Option.
- SORT assist improves performance of the DFSORT program product.

Features for outstanding scalar performance:

- High speed multiply feature implemented in hardware.
- Fast floating point add/subtract hardware.
- Special loop control circuitry.
- 64-bit wide data paths.

Optional Vector Facility:

- Provides complementary high performance for computations that may be vectorised.
- Executes vector arithmetic and logical operations on up to 128 sets of operands with a single instruction.
- Pipelined design permits completion of up to two operations per cycle.
- 171 vector instructions allow arithmetic and logical operations with binary, short (32-bit) and long (64-bit) precision floating point vectors.
- 16 vector registers can each store 128 32-bit elements and can be paired for 64-bit operations.
- Addressing can be performed using fixed stride values, indirect element selection or under mask control.
- Vector Activity Register, supported by the System Measurement Facility (SMF), permits analysis of vectorisation levels achieved.

IBM 3090 Channel Subsystem

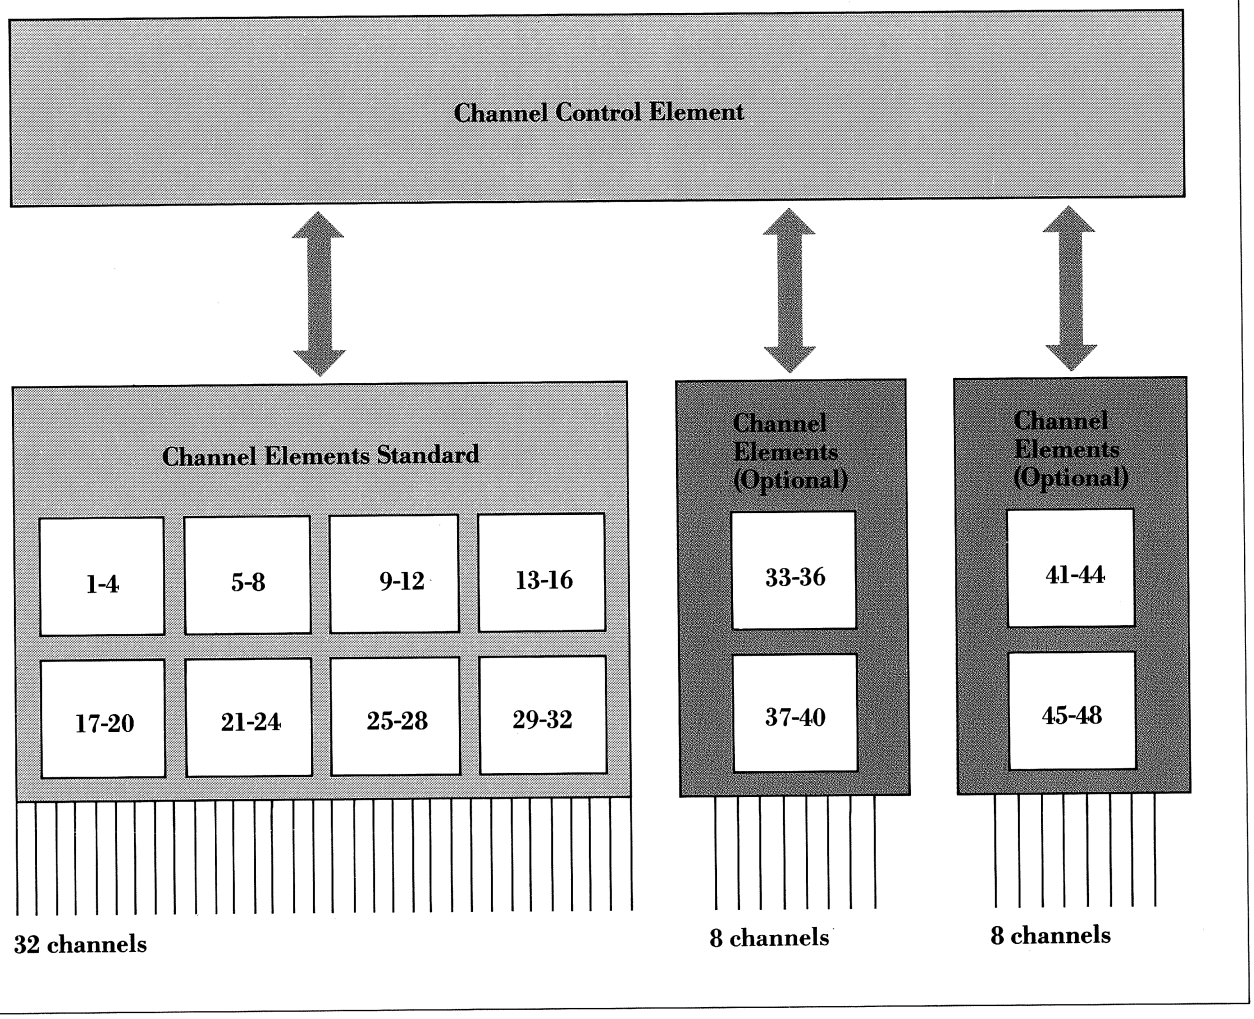

Channel Subsystem

- The Channel Subsystem performs I/O operations for the processor complex.
- The Model 200 has one Channel Subsystem, the Model 400 has two subsystems.
- Each Channel Control Element communicates with its associated System Control Element to access Central Storage and manages up to 12 Channel Elements.
- Each Channel Element operates four channels, which can each support up to eight control units.
- All channels configured as block multiplexer channels can operate at 3 megabytes per second in data streaming mode and 1.5 megabytes per second in interlocked mode.
- Both data streaming and non-data streaming devices may be attached to any block multiplexer channel on an intermixed basis.
- Optionally, up to four channels per Channel Control Element can be initialised to operate as byte multiplexers. This permits:
  - Up to four byte multiplexers per Model 200,
  - Up to eight byte multiplexers per Model 400.

XA Mode Operation:

- Any processor may initiate an operation with any I/O device and process any I/O interrupt using any channel path.
- Up to four different logical paths to a single device may be used.
- In single-image mode the two Channel Subsystems of a 3090 Model 400 co-ordinate activity and appear as one dynamic Channel Subsystem to the system control program.

System/370 Mode Operation:

- Channels may be grouped into logical channel sets with up to 32 channels to a set, and one set assignable to a Central Processor.
  - MVS/370 supports a maximum of 16 channels per channel set.
  - VM/SP High Performance Option supports up to 32 channels per channel set.
- Channel Set Switching permits a change of channel set assignment in the event of a Central Processor failure, so that I/O access can be maintained.

Channel Configuration:

- The Model 200 has 32 standard channels with two optional increments of eight, giving a maximum of 48 channels.
- The Model 400 has 64 standard channels with two optional increments of 16, giving a maximum of 96 channels.
- The channel, control unit and device configurations are defined to the Channel Subsystem by the I/O Configuration Dataset (IOCDS) selected at system initialisation.
- The IOCDS is created by the I/O Configuration Program and stored on the 3370 Direct Access Storage Devices attached to the 3092 Processor Controller.
- This permits transition between XA and System/370 modes whenever required by the installation.

Performance

- Commercial, Compute Intensive Scalar and Vector throughput figures are based on 3090 Model 200 internal throughput rate ratio measurements.
- Parallel processing turnaround figures are based on:
  - 3090 Model 200: external scalar job turnaround improvement ratio measurements.
  - 3090 Model 400: estimated external scalar job turnaround improvement ratios derived from measurements on an IBM 3084-Q four-way processor complex.
- The IBM 3090 Model 400 is expected to execute instructions in the range of 1.7 to 1.9 times that of an IBM 3090 Model 200. Actual performance measurements are scheduled for 1Q86.

Commercial Workloads

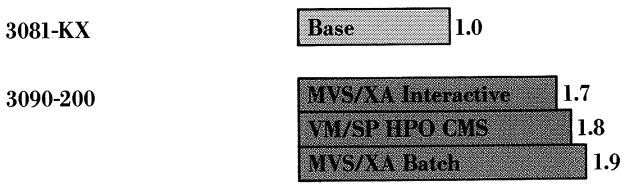

Compute Intensive Workloads

- Measured under MVS/XA

Scalar Processing

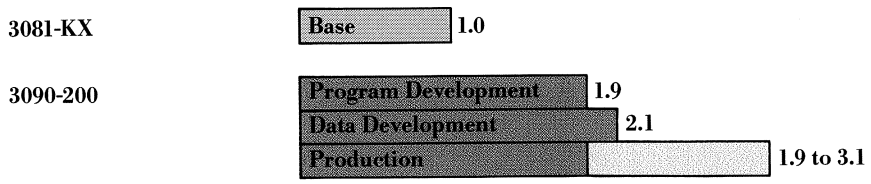

Vector Processing

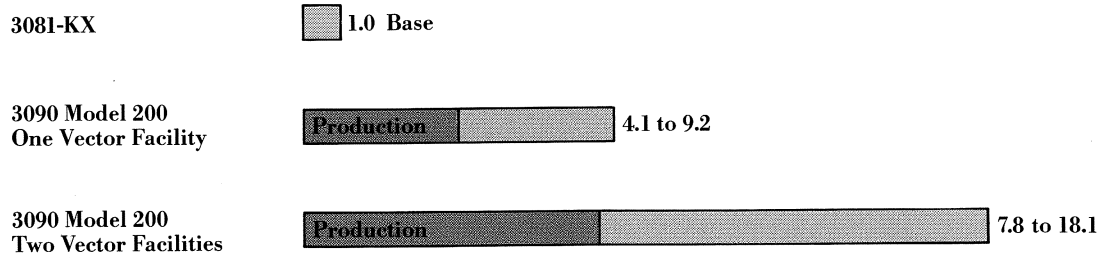

Parallel Processing

- Using IBM VS FORTRAN Program Multitasking Facility

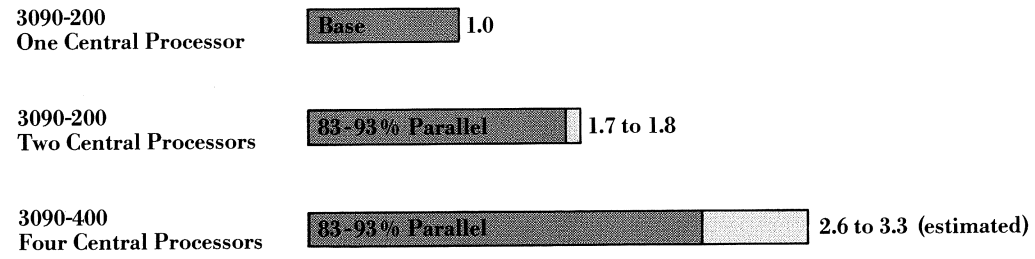

The measurement data is unique to the jobs evaluated and does not necessarily represent the actual performance improvement that any user will experience with his application.

Variables affecting performance include hardware configuration, system control program and job specific parameters such as problem size, solution technique and data array dimensions.

Operating Modes, System Control Program Support and Features for 3090 Processor Unit Model 200.

Operating modes	System / 370 XA		System / 370		
System Control Programs	MVS / SP 2.1.3	VM / XA SF	MVS / SP 1.3.5	VM / SP HPO 3.6	ACP / TPF V 2.2
Assists	SORT	SIE, SORT	No	VMA, PMA	No
Expanded Storage	Yes	Guest	No	Yes	Yes
Vector Facility	MVS / SP 2.1.3 VFE	No	No	VM / SP HPO 4.2	No

Physical Characteristics:

ITEM	UNITS	3081-KX	3090-200	3090-400
Channels	Quantity	24	40	80
Storage	Mb	32	64	128
Vector Facility	Quantity	n / a	0 / 2	0 / 4
Power Consumption	KVA	21.1	46.7/53.6	88.8/102.6
Heat Output: to water	KBTU/hr	39.1	66.5/80.1	133.1/161.4
to air		19.2	47.8/49.8	84.1/86.2
Total		58.3	114.3/129.9	217.2/247.6
Space	Sq. metres	5.72	7.4 / 8.6	14.3 / 16.0
	Sq. feet	61.55	79.9/92.2	153.8/171.7
(including service clearances)	Sq. metres	35.2	41.1/44.9	70.1/75.9
	Sq. feet	378.7	442.8/483.7	754.5/817.0
Operating Environment	Centigrade	16 to 29	16 to 29	16 to 29
	Fahrenheit	60 to 85	60 to 85	60 to 85
Relative Humidity	%	20 to 80	20 to 80	20 to 80
Maximum Wet Bulb	Centigrade	23	23	23
	Fahrenheit	73	73	73
Weight	Kg	4525	5920/6645	11045/12497
	Lb	9950	13050/14650	24350/27550

Figures for power consumption, heat output, space requirements and weight include 3087/3097 Power and Coolant Distribution Unit(s) and 3082/3092 Processor Controller as well as the 3081/3090 Processor Unit configured as indicated above.

Publications

Title	Reference
IBM Problem Analysis and Remote Support Reference Guide	S229-3299
IBM System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufacturer Information	GA22-6974
IBM System/370 Principles of Operation	GA22-7000
IBM System/370 System Summary Processors	GA22-7001
IBM System/370 Input/Output Configurator	GA22-7002
IBM System/370 Extended Architecture Principles of Operation	SA22-7085
IBM System/370 Extended Architecture Interpretive Execution	SA22-7095
IBM 3090 Channel Characteristics and Configuration Guide	SA22-7120
IBM 3090 Functional Characteristics	SA22-7121
IBM Input/Output Device Summary	GA32-0039
IBM General Information Manual: Installation Manual Physical Planning	GC22-7072
IBM System/370 3090 Processor Complex Installation Manual Physical Planning	GC22-7074
IBM 3090 Processor Complex I/O Configuration Program User's Guide and Reference	SC38-0038
IBM Physical Planning Template - 3090 Processor Complex Metric Scale	GX22-7107
IBM Physical Planning Template - 3090 Processor Complex English Scale	GX22-7108

IBM Eurocoordination
S.A. au capital de 2.700.000 F
Siège social : Tour Pascal
22, Route de la Demi-Lune
92075 Puteaux
R.C.S. Nanterre B 304 538 192
France

IBM World Trade
Americas/Far East Corporation
Town of Mount Pleasant
Route 9, North Tarrytown
New York, 10591 U.S.A.

G 511-0133-0 *Printed in Denmark by Lassen Offset*

IBM 3090
Processor Complex Support Units

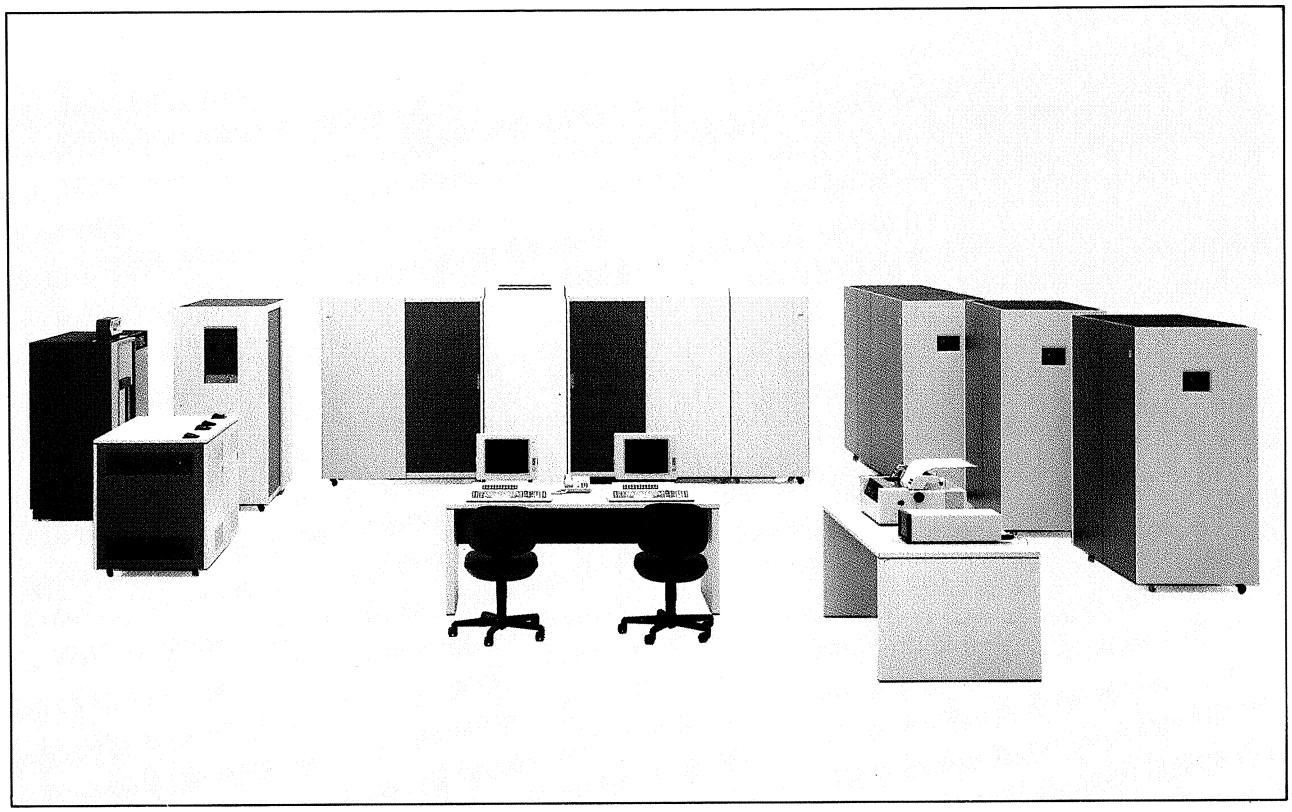

IBM

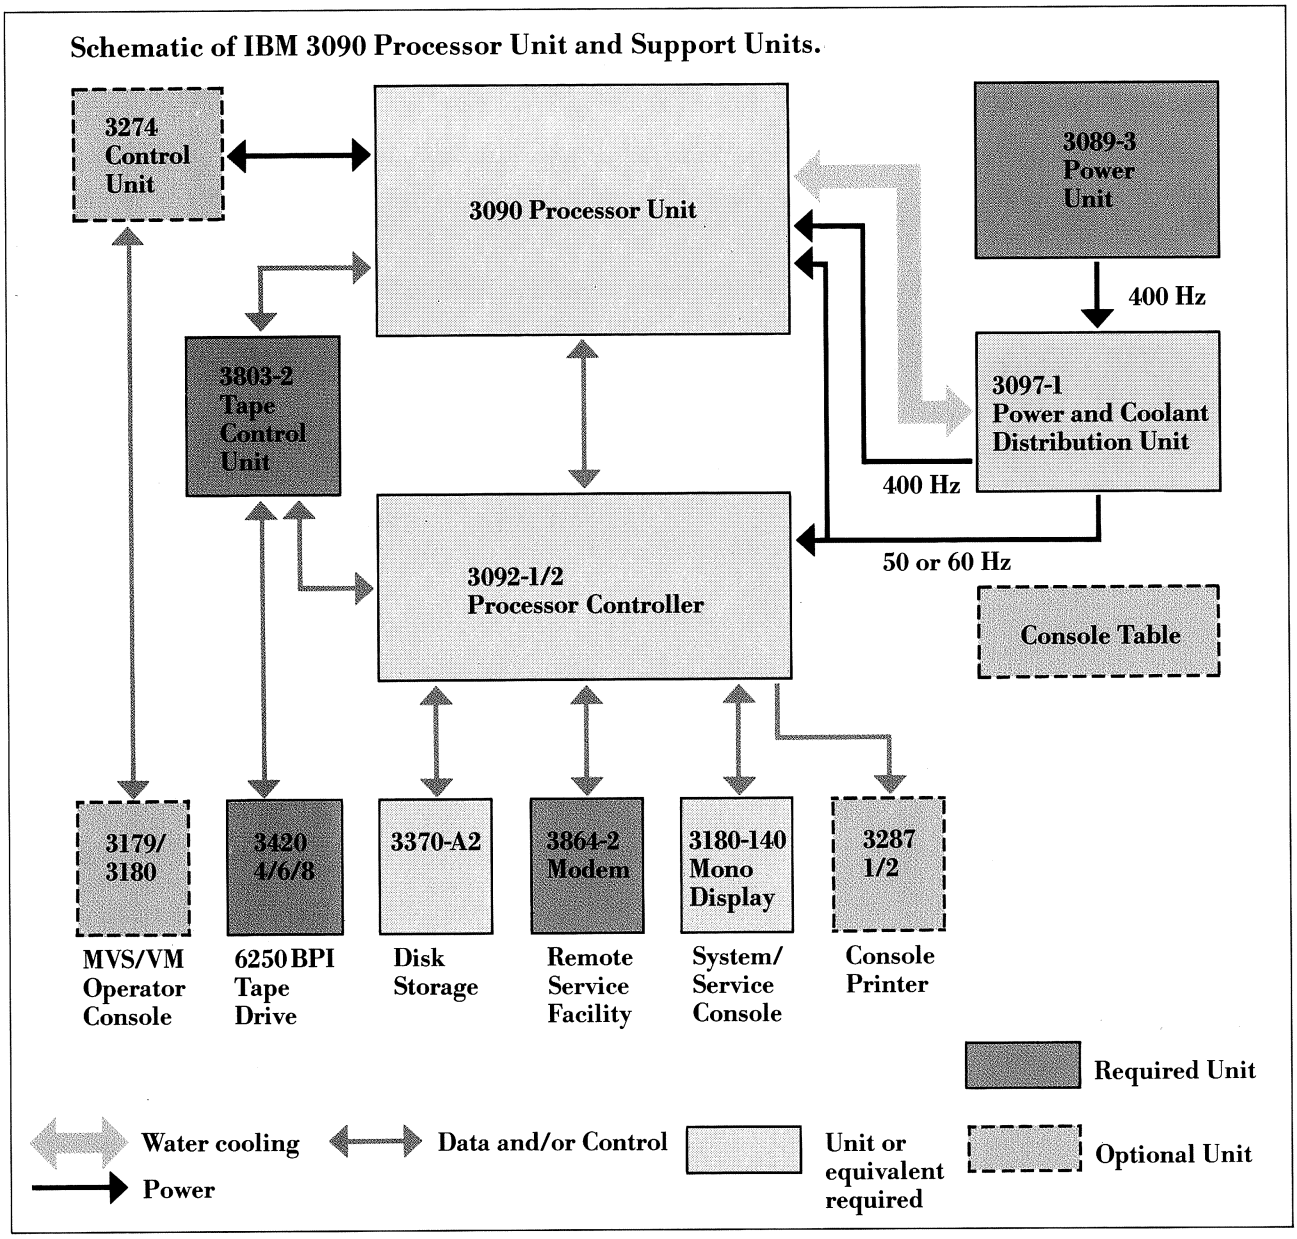

IBM 3090 Processor Complex Support Units

The IBM 3090 Processor Complex is the most powerful IBM computing system available. The IBM 3090 Processor Unit performs the data processing functions of the system. It is supported by a number of required or optional units, which are described in this publication.

The quantity required of each support unit depends on the model of the IBM 3090 Processor Unit. Outline configuration details are provided in the IBM 3090 Processor Unit publication G511-0133. Further information may be obtained from the IBM publications listed at the end of this publication.

Required Unit Types:

- IBM 3089 Power Unit (or equivalent).
- IBM 3092 Processor Controller.
- IBM 3097 Power and Coolant Distribution Unit.
- IBM 3180 Model 140 Display Station as system/ service console.
- IBM 3370 Direct Access Storage Device.
- IBM 3864 Modem (or equivalent).
- Access to IBM 3420 Tape Drive, 6250 bits per inch, (or equivalent).

Optional Units:

- IBM 3179/3180 Display Station (or equivalent) channel-attached operator console for communication with system control program.
- IBM 3287 Printer.
- IBM Console Table.

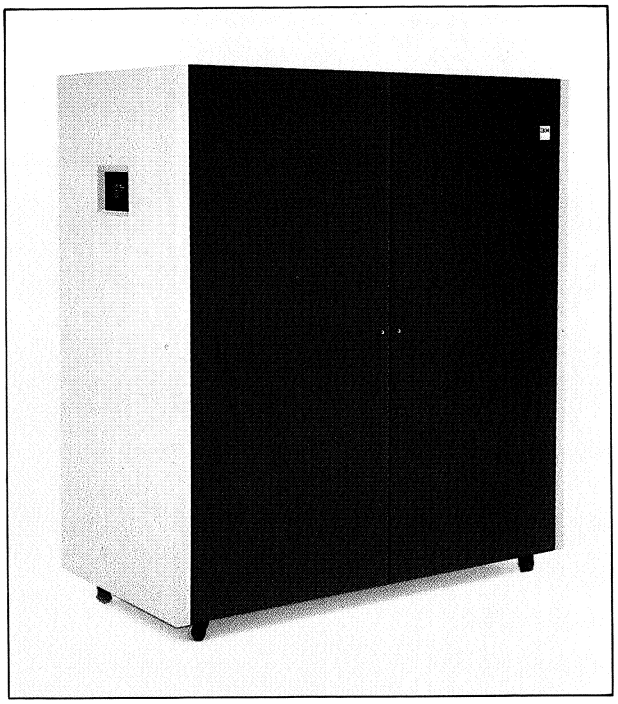

IBM 3089 Model 3 Power Unit

The 3089 contains a motor generator and it, or an equivalent device, provides 400 Hz power to the IBM 3097 Power and Coolant Distribution Unit from which it is distributed to the different frames of the IBM 3090 Processor Unit.

Highlights include:
- Designed for location in the computer room, with the same colour options and low noise characteristics as the IBM 3090 Processor Unit,
- Can be installed by IBM together with the rest of the 3090 Processor Complex, simplifying initial installation.
- Can be maintained by IBM, helping to assure the high availability of the 3090 system, which depends on the correct functioning of the power source for the IBM 3090 Processor Unit.

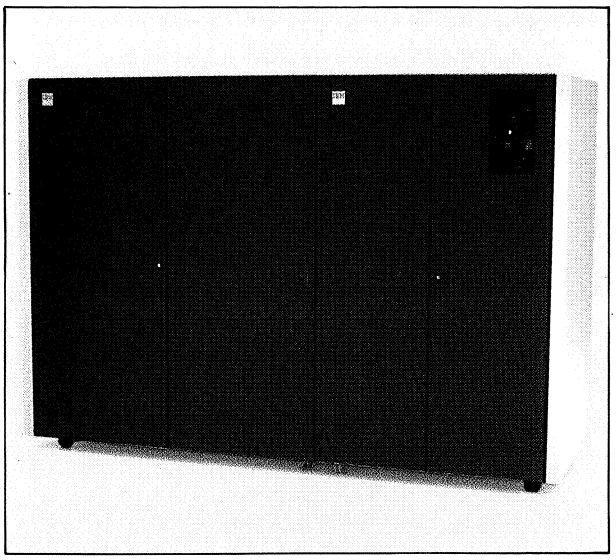

IBM 3097 Model 1 Power and Coolant Distribution Unit

The 3097 provides the combined functions of power distribution to the 3090 Processor Complex and water cooling to the IBM 3090 Processor Unit.

Power Distribution:

- Obtains power from the IBM 3089 Model 3 Power Unit or equivalent 400 Hz power source.
- Distributes 400 Hz power to the IBM 3090 Processor Unit and 50 or 60 Hz power to the IBM 3092 Processor Controller and the IBM 3090 Processor Unit.
- Can power on and off up to 64 control units and, if configured with the optional Power Sequence Control feature, up to 128 control units.

Coolant Distribution:

- Controls the temperature and flow rates of an IBM-supplied closed-loop distilled water cooling system used to cool the Thermal Conduction Modules in the IBM 3090 Processor Unit.
- Pumps the water in the closed-loop system through a heat exchanger in the 3097 where customer-supplied chilled water from an external source absorbs the heat.
- Contains a duplicate pump for the closed-loop cooling system which automatically takes over from the operating pump if it fails.

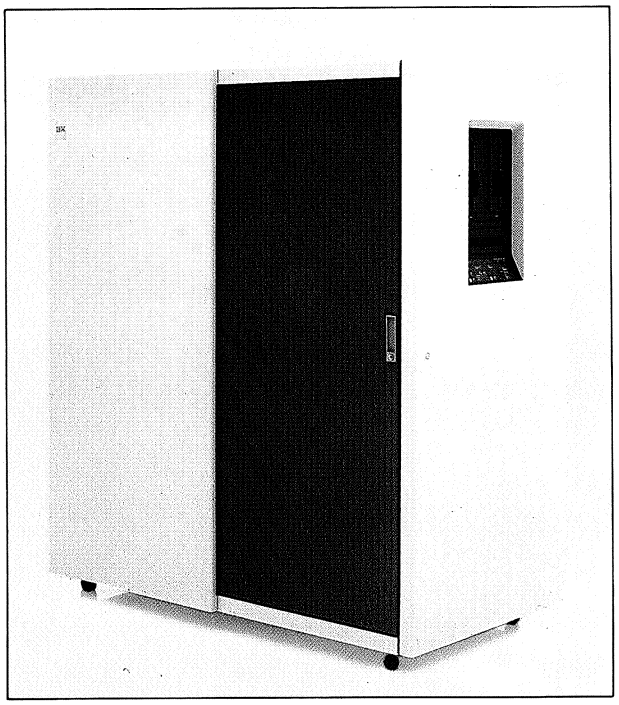

IBM 3092 Processor Controller

The IBM 3092 Processor Controller Models 1 and 2 monitor and control the status of all physical units within a 3090 Processor Complex. Because their correct functioning is essential for 3090 operation, both models of the 3092 contain duplicate processors and are supported by duplicate 3370 Direct Access Storage Devices, permitting the Processor Controller to continue operation after a component failure.

Both models of 3092 perform similar functions. These are:
Device Attachment for:
- Required and optional 3180 system and service consoles.
- 3370 Direct Access Storage Devices.
- 3864 Modem or equivalent.
- Optional 3287 Console Printer.

System Initialisation and Control:

- Power on and power off of all 3090 Processor Complex units and all control units which are under 3097 power sequence control.
- Initial Microcode Load (IML) and Initial Program Load (IPL) of the IBM 3092 Processor Controller and IML of the IBM 3090 Processor Unit.
- Validation of error-free data locations in Central Storage and recording of failing storage locations.
- System configuration for System/370 or System/370 Extended Architecture mode of operation.

Monitoring and Error Logging:

- Collection of information for up to 24 system activity displays including central processor and channel utilisation.
- Monitoring of voltage levels and coolant flow rate.
- Logging of errors, failures and status information relating to Processor Complex and I/O problem determination.

Error Recovery, Fault Diagnosis and Service Support:

- Automatic error recovery where possible, such as the activation of the duplicate water pump in the 3097.
- Automatic identification of the failing field replaceable unit, timeshared with 3092 and 3090 operation.
- Operator alert for non-recoverable or frequent failures.
- Interactive customer Problem Analysis for problem determination and immediate recovery when feasible.
- Reconfiguration of Central Storage, Central Processors and online channels to bypass failures allowed.
- Remote Service Facility to permit immediate fault diagnosis by remote IBM specialists using a dial-up link.

IBM 3864 Model 2 Modem

The IBM 3864 Model 2 Modem with an Autocall Unit feature, or an equivalent device, is required to permit access by the IBM 3092 Processor Controller to IBM's Remote Service Facility (RSF) which is used for remote service updating as well as remote fault diagnosis.

- The 3864 permits data transmission at 4800 bits per second on a switched telephone network.
- Autocall Unit feature permits automatic dialling to or from the IBM remote service location.
- No incoming or outgoing are calls permitted without customer authorisation and manual instead of automatic dial up can be used if required.
- Details of all calls are recorded and each call is listed on a Remote Service Facility (RSF) log index which can be displayed.
- Once the link is established, the RSF can run diagnostics remotely.

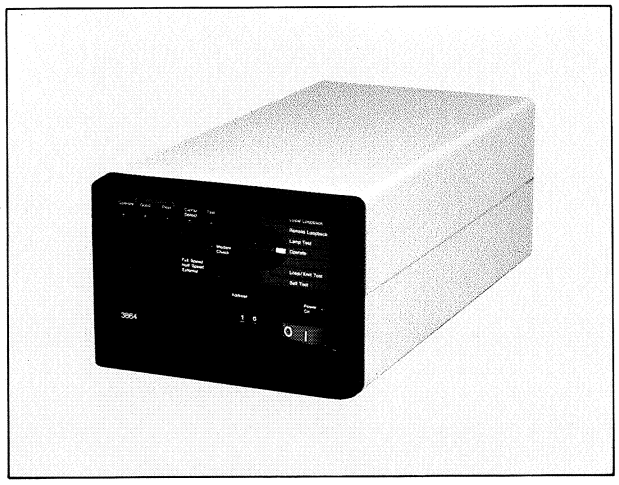

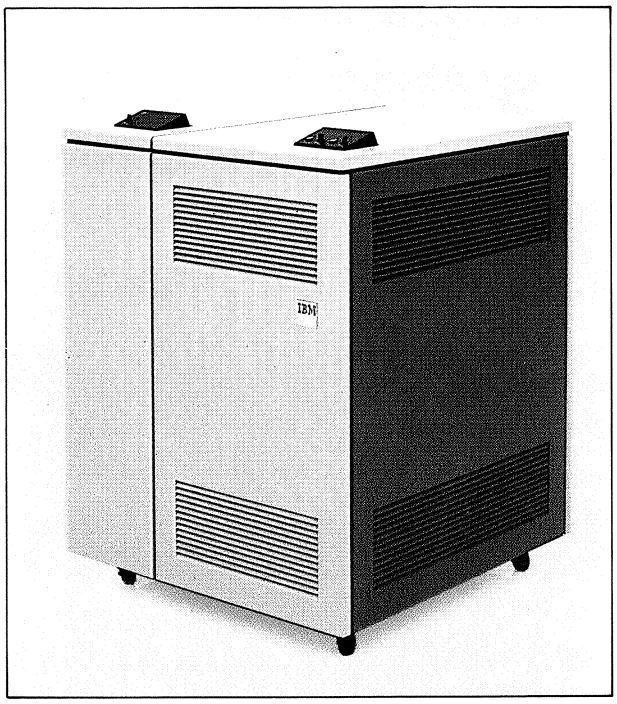

IBM 3370 Direct Access Storage Device

Two IBM 3370 Model A02 Direct Access Storage Devices are used to provide storage for the IBM 3092 Processor Controller:

- Each 3370 has its own control unit and can store up to 730 Megabytes.
- Each 3370 is connected, by means of a string switch feature, to both processors in the 3092.
- The 3370s are used to store microcode for the 3092 and 3090 together with system Input/Output configuration information, diagnostics and the error logs used by service personnel.
- In normal single image operation, one 3370 is active and the other provides backup. Essential information is duplicated on both 3370s so that the backup 3370 can take over at any time.

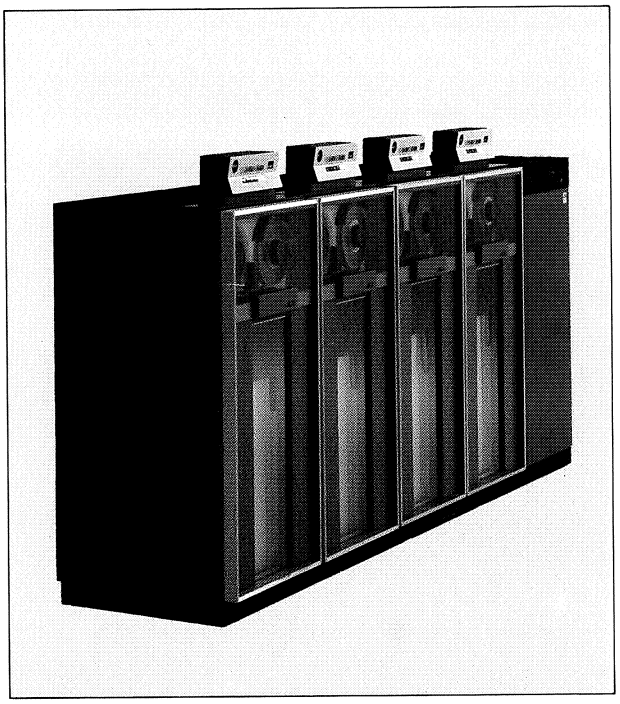

IBM 3420 Model 4/6/8 Tape Drive

Access is required by the IBM 3092 Processor Controller to a tape drive to install microcode and diagnostics, to apply engineering changes and to reload the IBM 3370 Direct Access Storage Devices, if necessary. This is provided by:

- Access to a channel path of an IBM 3803 Model 2 Tape Control Unit (or equivalent) and an IBM 3420 Model 4, 6 or 8 Tape Drive (or equivalent 6250 bits per inch tape drive).
- Location of the tape drive within 200 feet of the 3092 when a 3420 Model 4 or 6 is used or within 72 feet when a Model 8 is used.

IBM 3180 Model 140 Display Station

The IBM 3180 Model 140 Display Station is a monochrome unit used to perform multiple console roles in a 3090 Processor Complex.

- the IBM 3092 Processor Controller can support up to five physical 3180 consoles, required and optional. The quantity of required consoles depends on the model of the IBM 3090 Processor Unit.
- Each 3180 can be used logically as a system console, service support console, programming support console, system console monitor or service support console monitor.
- Any physical display attached to the 3092 can be given multiple logical console functions, which may be re-assigned as required.
- The service support console must be located within 10 metres (33 feet) of the 3092. Other consoles can be located up to 1500 metres (4921 feet) away.

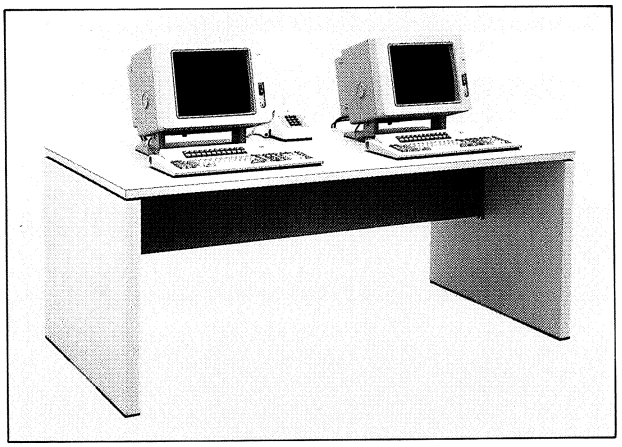

Optional Units

A number of optional units may be attached to the 3090 Processor Complex. These include:

IBM 3179 and IBM 3180 Displays

According to user needs, IBM 3179 Colour and IBM 3180 Model 1 monochrome displays, or other supported devices, may be channel-attached via an IBM 3274 Control Unit for use as MVS or VM operator consoles.

IBM 3287 Printer

One optional console printer can be attached to the IBM 3092 Processor Controller. Two models are available for copying console display images:
- IBM 3287 Model 1, 80 characters per second,
- IBM 3287 Model 2, 120 characters per second.

Console Table

The IBM console table provides an optional operator workstation for one or two operators and their display consoles.

Physical Characteristics

Product		3089 Power Unit	3097 Power and Coolant Distribution Unit	3092 Processor Controller	3370 Direct Access Storage Device	3180 Display Station	Console Table
Height:							
	Metres	1.790	1.790	1.790	1.000	0.370-0.520	0.720
	Inches	70.5	70.5	70.5	39.5	14.5-20.5	28.5
Length:							
	Metres	1.640	3.165	1.615	0.775	0.435	1.780
	Inches	64.5	124	63.5	30.5	17	70
Depth:							
	Metres	0.815	0.815	0.815	0.815	0.360	0.900
	Inches	32	32	32	32	14.25	35.5
Space:							
	Sq M	1.34	2.58	1.32	0.63	0.16	1.60
	Sq Ft	14.33	27.56	14.11	6.78	1.68	17.26
Weight:							
	Kg	1075	1395	770	260	22.25	-
	Lb	2370	3080	1700	580	49	-
Heat Output:							
— To Air							
	Watts	8000	1050	1750	720	130	-
	BTU/Hr	27300	3550	6000	2500	450	-
— To Water							
	Watts	-	1650	-	-	-	-
	BTU/Hr	-	6000	-	-	-	-
Airflow:							
	M3/Min	11.50	3.70	34.41	5.50	Convection	-
	CFM	400	130	1220	190	Convection	-
Power:							
	Watts	38900	*	*	720	110	-
	KVA	43	*	*	1.3	0.2	-

Requirements for operating environment temperature, relative humidity and maximum wet bulb are similar to those for the 3090 Processor Unit which are supplied in the companion publication reference G511-0133

*Included with 3090 Processor Unit power consumption. See publication reference G511-0133.

Publications

Title	Reference
IBM 3180 Display Station Introduction and Pre-installation Planning Manual	GA21-9465
IBM 3270 Information Display System Features Description	GA23-0113
IBM 3370 Direct Access Storage Description	GA26-1657
IBM 3270 Information Display System Configurator	GA27-2849
IBM 3270 Information Display System Installation Physical Planning	GA27-2787
IBM 3278 Display Station Operator's Guide	GA27-2890
IBM 3287 Printer Models 1 and 2 Operator's Guide	GA27-3150
IBM 3287 Printer Component Description	GA27-3153
IBM 3863, 3864 and 3865 Modems, Introduction and Site Preparation Guide	GA27-3200
IBM 3803-2/3420 Magnetic Tape Subsystem Description	GA32-0021
IBM 3279 Display Station Operator's Guide	GA33-3057

Information on the IBM 3089, IBM 3092 and IBM 3097 is to be found in the IBM 3090 Processor Unit publications listed in the companion publication Reference G511-0133

IBM Eurocoordination
S.A. au capital de 2.700.000 F
Siège social : Tour Pascal
22, Route de la Demi-Lune
92075 Puteaux
R.C.S. Nanterre B 304 538 192
France

IBM World Trade
Americas/Far East Corporation
Town of Mount Pleasant
Route 9, North Tarrytown
New York, 10591 U.S.A.

G 511-0134-0 *Printed in Denmark by Lassen Offset*

IBM 3090
Vector Facility and Systems Software

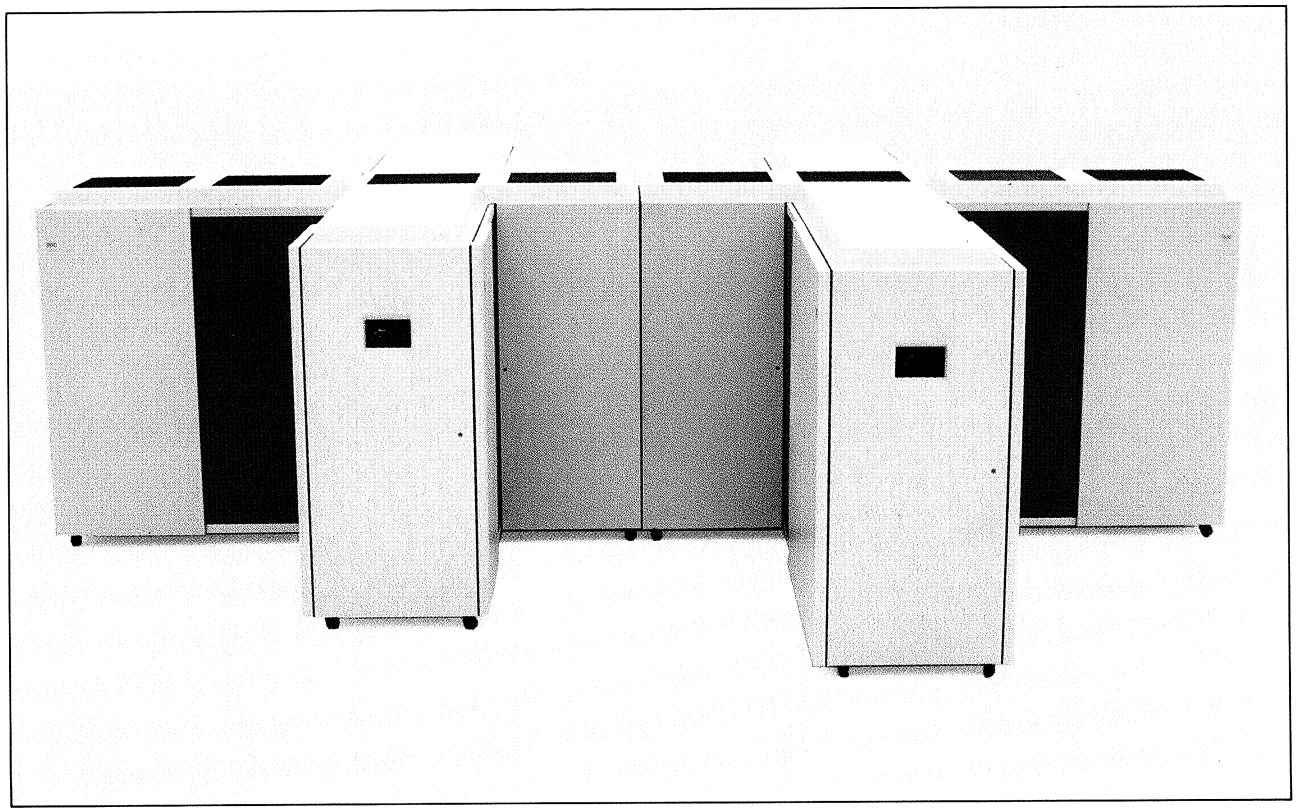

IBM

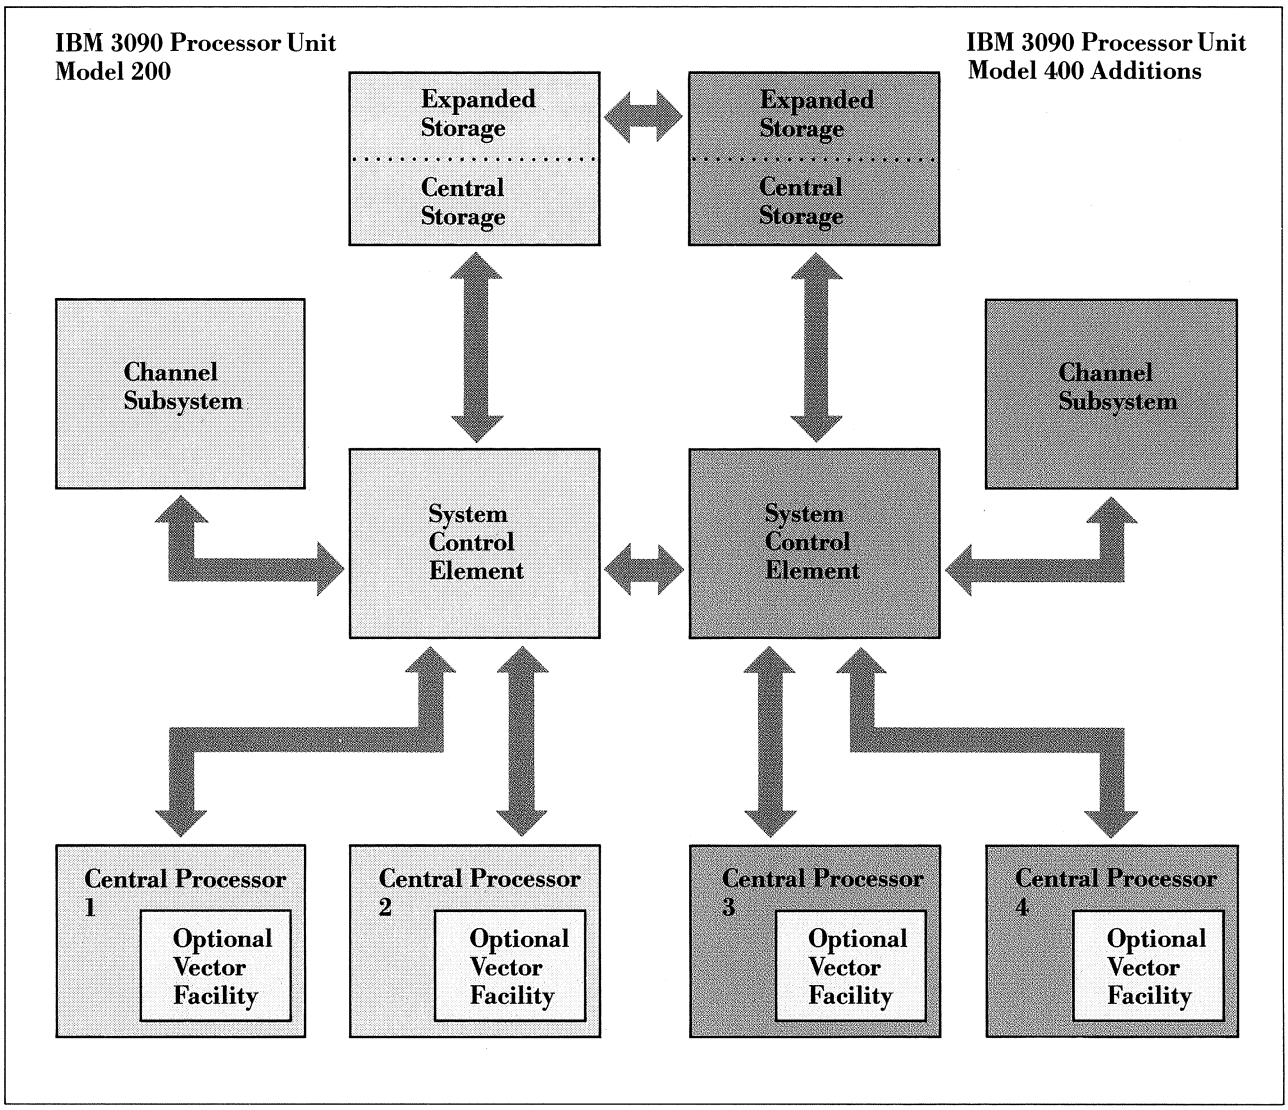

IBM 3090 Vector Facility

The IBM 3090 Vector Facility is an optional field or plant installable feature which adds a powerful vector processing capability to a Central Processor in an IBM 3090 Processor Unit.

The IBM 3090 Processor Unit Models 200 and 400 with one or more Vector Facilities and associated systems software provide users with an integrated solution for scalar, vector and parallel processing that will meet both compute-intensive and general-purpose application requirements.

The Vector Facility is installed in an additional frame, integrated into the IBM 3090 Processor Unit. Availability of up to four Vector Facilities in a 3090 permits users to tailor the configuration to their requirements and handle future growth.

Vector Facility options are:

- IBM 3090 Model 200
  - One or two Vector Facilities.
  - Operation under MVS/XA or VM/SP High Performance Option (HPO).
- IBM 3090 Model 400
  - One, two, three or four Vector Facilities.
  - Operation under MVS/XA in single image mode, and under MVS/XA or VM/SP HPO in partitioned mode.

This publication describes the IBM 3090 Vector Facility and associated systems software. For further information on the IBM 3090 Processor Unit and IBM 3090 Processor Complex Support Units, see publications reference G511-0133 and G511-0134 respectively.

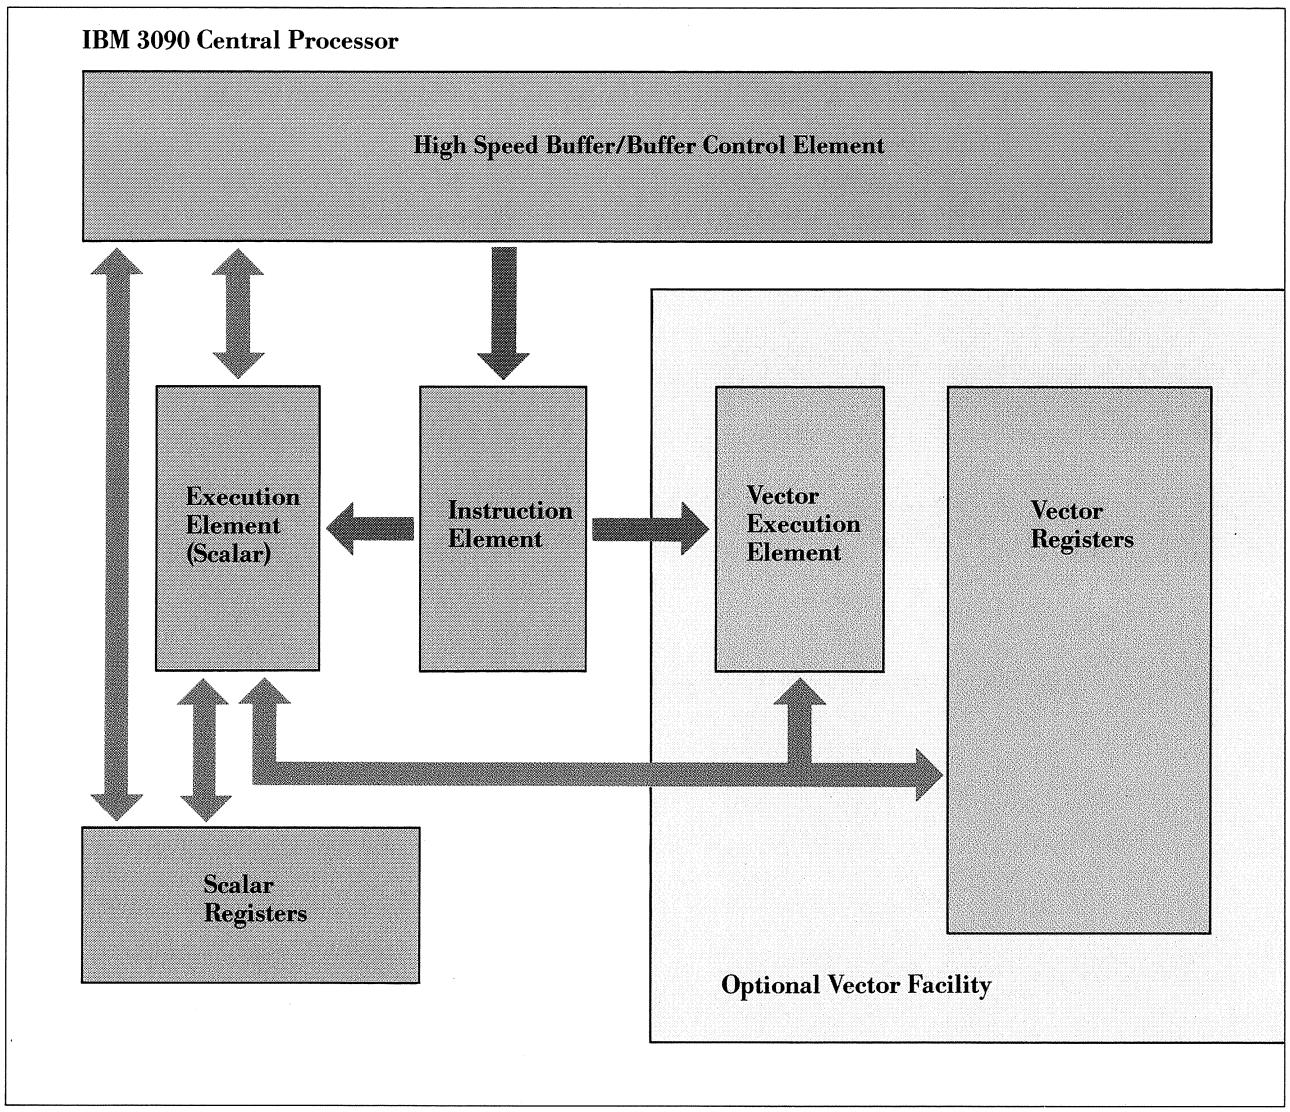

IBM 3090 Vector Facility Feature

Comprehensive vector processing with modular growth options:

- The 3090 Vector Facility provides extra hardware to enable an IBM 3090 Central Processor to execute an extensive repertoire of 171 vector instructions.
- The feature is housed in an additional frame which forms part of the 3090 Processor Unit.
  - The first feature consists of a frame, a Thermal Conduction Module (TCM) board and three TCMs.
  - The second feature consists of three more TCMs.
  - On a 3090 Model 400, two additional features provide another frame, TCM board and TCMs on the "B" side of the Processor Unit.
- The technology is similar to that employed in other parts of a 3090 Central Processor, incorporating high-speed Emitter Coupled Logic (ECL) chips within water-cooled TCMs.

Integrated design permits scalar and vector operations in the same processor:

- Logically each Vector Facility is an integral part of one 3090 Central Processor.
- Instructions are decoded by the same Instruction Element which serves the scalar Execution Element.
- In addition to the Vector Execution Element, each Vector Facility has 8K bytes of vector register storage usable as:
  - 16 vector registers containing 128 32-bit elements.
  - 8 vector registers, by means of pairing, containing 128 64-bit elements.
- Binary as well as 32-bit (short) and 64-bit (long) precision floating point operands are permitted.

Flexible addressing options for efficient data manupulation:

- Three-address instruction formats minimise movement of information in registers and improve operand re-use.
- Input operands may be:
  - Vector register and vector register.
  - Vector register and storage.
  - Vector register and scalar.
  - Storage and scalar.
- Vector results are placed in vector registers, scalar results are placed in scalar registers.
- Source or target vector elements in storage are accessed in one of three ways:
  - Starting address modified by stride values (regular addressing).
  - Starting address modified by a vector of element selection numbers (indirect element selection).
  - Vector register and main storage elements selected under mask control using a 128-bit vector mask register.
- Automatic address updating permits sectioning of vectors longer than 128 elements.

Full use of 3090 storage hierarchy and virtual storage capabilities:

- Data and instructions in storage are normally fetched from the Central Processor's 64K-byte High Speed Buffer, via a 64-bit data path.
- During vector processing, virtual storage permits the addressing of up to 2 048 million bytes in System/370 Extended Architecture mode and up to 16 million bytes in System/370 mode, per address space.

Heavily overlapped Vector Execution Element for high performance:

- Pipelined arithmetic and logical units in the Vector Execution Element are able to produce a 32-bit or 64-bit sum, difference or product each cycle.
- Compound instructions permit two operations to be performed each cycle.
- The 3090 Vector Facility is designed with a vector/scalar ratio of between 3 and 5.
- A vector activity register permits collection of statistics on Vector Facility usage by the System Measurement Facility, so that program vectorisation can be assessed.

Outstanding scalar, vector and parallel processing facilities integrated in one system:

- For the scalar portions of vectorised jobs, as well as jobs not suitable for vectorisation, the 3090 provides exceptional scalar performance.
- Scalar as well as vector object code can be obtained from the same FORTRAN source code using the IBM VS FORTRAN Version 2 Compiler.
- The scalar object code can be tested before installation of a Vector Facility or used in other computer centres not equipped with a Vector Facility.
- Where a high degree of parallelism is possible all the scalar or vector execution elements in a 3090 Processing Unit can be employed simultaneously to execute the same program, using the IBM VS FORTRAN Program Multitasking Facility in an MVS/XA environment.
- These capabilities make the 3090 an outstanding system for nearly all compute-intensive applications.

Performance

The IBM Vector Facility has been measured on a 3090 Model 200 operating under MVS/XA. Internal Throughput Rate (ITR) ratios relative to the Model 200 in scalar mode and to a 3081-KX have been obtained for a variety of vectorisable applications. In all cases the measurements shown here were made using one initiator, giving internal throughput for a single central processor.

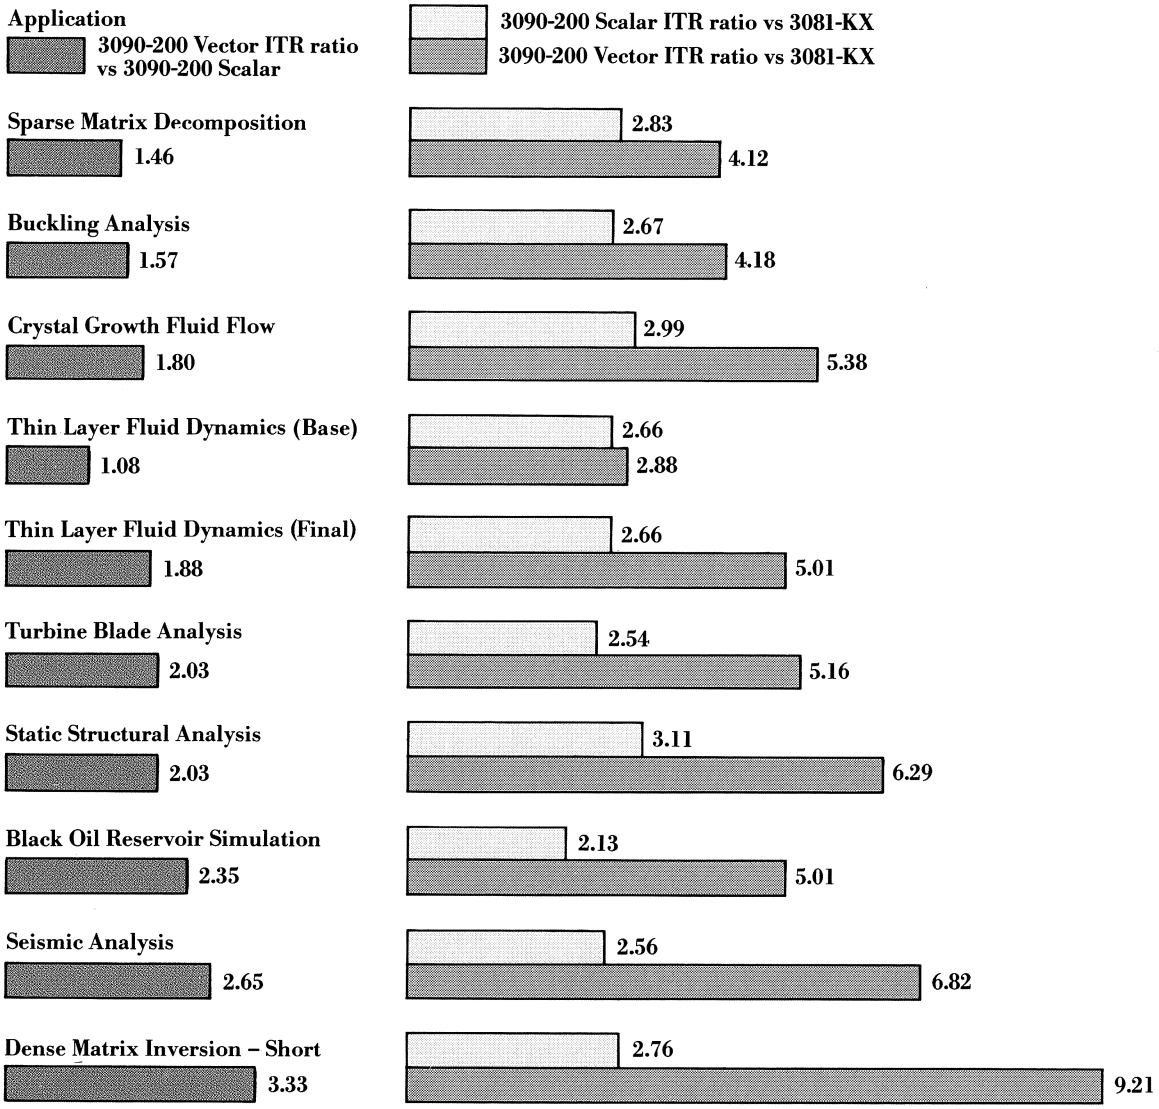

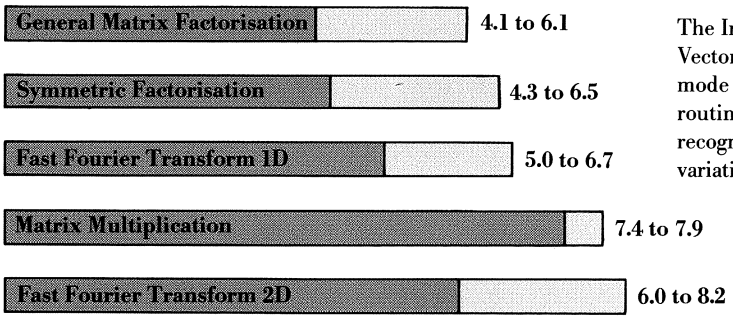

IBM Engineering and Scientific Subroutine Library (ESSL).

The Internal Throughput Rate (ITR) of the 3090 Model 200 with Vector Facility relative to the Model 200 operating in scalar mode has been measured, comparing five out of the 95 ESSL routines with scalar routines obtained from sources generally recognised as efficient. The range of ITR ratios reflects variations in order, vector length and array size.

The measurement data is unique to the jobs evaluated and does not necessarily represent the actual performance improvement that any user will experience with his application.

Variables affecting performance include hardware configuration, system control program and job specific parameters such as problem size, solution technique and data array dimensions.

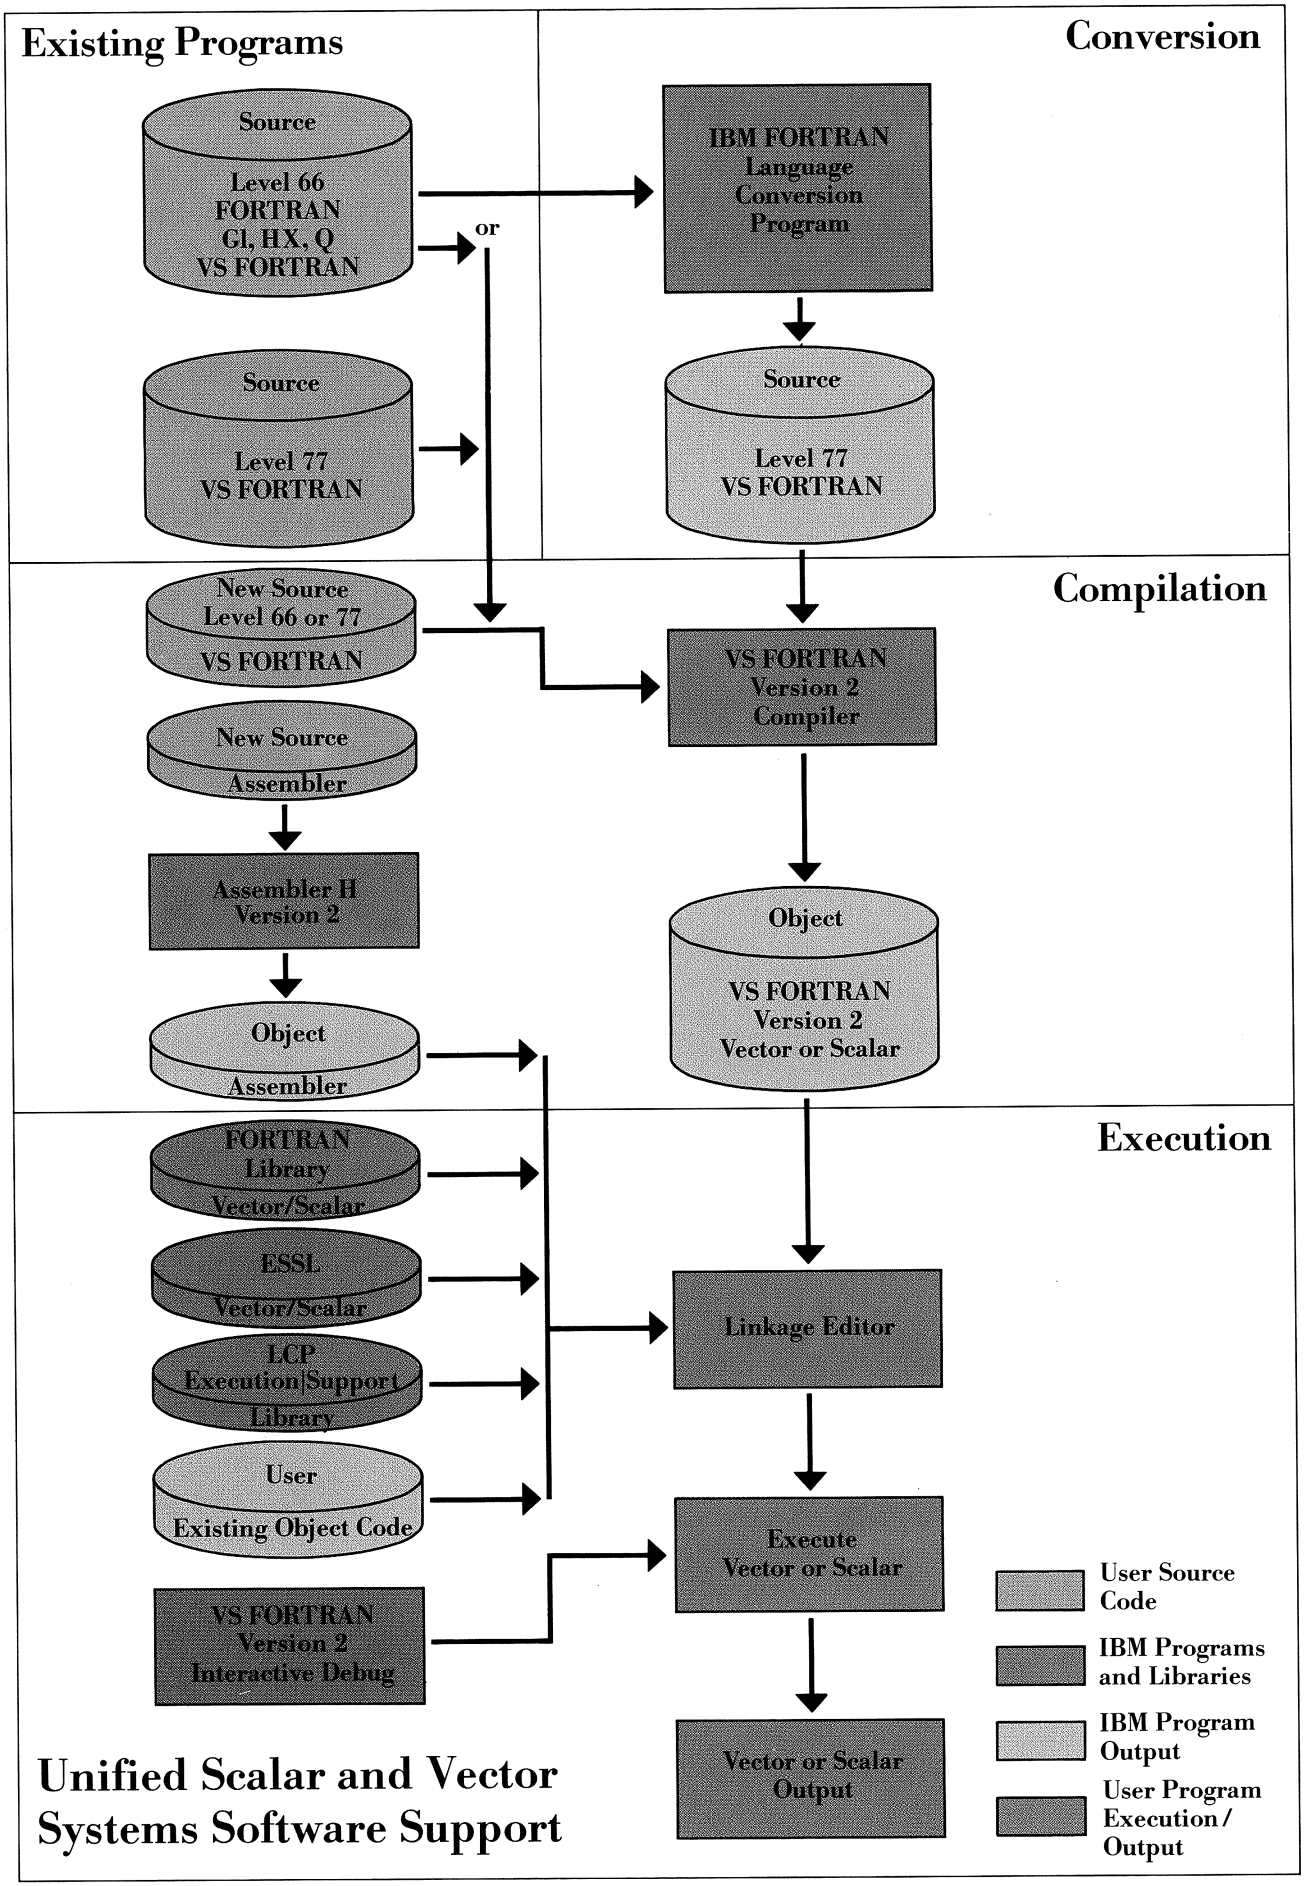
Existing Programs
Conversion
Source
Level 66
FORTRAN
Gl, HX, Q
VS FORTRAN
or
IBM FORTRAN
Language
Conversion
Program
Source
Level 77
VS FORTRAN
Source
Level 77
VS FORTRAN
Compilation
New Source
Level 66 or 77
VS FORTRAN
VS FORTRAN
Version 2
Compiler
New Source
Assembler
Assembler H
Version 2
Object
VS FORTRAN
Version 2
Vector or Scalar
Object
Assembler
Execution
FORTRAN
Library
Vector/Scalar
ESSL
Vector/Scalar
LCP
Execution Support
Library
User
Existing Object Code
VS FORTRAN
Version 2
Interactive Debug
Linkage Editor
Execute
Vector or Scalar
Vector or Scalar
Output
User Source
Code
IBM Programs
and Libraries
IBM Program
Output
User Program
Execution /
Output
Unified Scalar and Vector
Systems Software Support

Systems Software

The systems software support for the IBM 3090 Vector Facility comprises:

— System Control Programs
Compile and Execute:
· MVS/XA.
· VM/SP High Performance Option.

Compile only:
· MVS/370.
· VM/SP.

— Languages and Libraries:
· Assembler H.
· VS FORTRAN Extended including Library, Multitasking Facility and Interactive Debug.
· Engineering and Scientific Subroutine Library.

— Other Programs:
· IBM FORTRAN Language Conversion Program.
· Vector Processing Subsystem/Vector Facility for IBM 3838 simulation.

System Control Program Support

— For vector program execution, the IBM 3090 Vector Facility is supported:
· In S/370 Extended Architecture mode by MVS/SP 2.1.3 Vector Facility Enhancement (VFE).
· In System/370 mode by VM/SP High Performance Option Release 4.2.

— For vector program compilation, using IBM VS FORTRAN Extended or IBM Assembler H Version 2, Release 1 VFE, support is provided in two additional System/370 mode environments:
· MVS/SP Version 1, all releases.
· VM/SP Release 3 or later, with or without the High Performance Option.

Languages and Libraries

VS FORTRAN Version 2, Release 1

VS FORTRAN Version 2 consists of:
- VS FORTRAN Version 2 Compiler.
- FORTRAN Library including Multitasking Facility.
- Interactive Debug.

VS FORTRAN Version 2 has all the function, usability and performance available in:
- VS FORTRAN Version 1 Release 4.1 Compiler and Library including the VS FORTRAN Program Multitasking Facility.
- VS FORTRAN Interactive Debug Release 2.0

It also provides:
- Exploitation of the IBM 3090 Vector Facility.
- New mathematical routines yielding better accuracy and performance.
- Enhancement of the VS FORTRAN Program Multitasking Facility to support vector tasks under MVS/XA.
- Fully integrated debugging support including extended full screen support with windowing and animation.
- Support of the Engineering and Scientific Subroutine Library and the IBM FORTRAN Language Conversion Program.

VS FORTRAN Version 2 Compiler

Re-entrant design with exploitation of 31-bit addressing in XA-mode.

- Input to the Compiler can be:
  - Existing or new programs written in IBM FORTRAN language level 66 or 77 source code.
  - FORTRAN programs which have been converted by the IBM FORTRAN Language Conversion Program.

- Compile options permit creation of highly optimised object code for:
  - Scalar execution on System/370, 43XX or 30XX processors.
  - Vector execution on IBM 3090 Processor Units equipped with one or more Vector Facilities.
  - Parallel execution of scalar and/or vector code on multiple central processors.

- Vectorisation, when requested, is automatic and does not require modification of original (scalar) code:
  - Source statements appropriate for vectorisation are identified.
  - A selection is made from the identified statements so as to optimise performance.
  - Highly optimised vector object code is compiled from the statements chosen for vectorisation.
  - Highly optimised scalar code is compiled from the remaining unvectorised statements.

- An optional Vector Report provides a summary of the vectorisation performed:
  - Indicates which loops have been vectorised.
  - Assists tuning of source code.

- Compiler provides excellent performance in terms of:
  - Compile time.
  - Scalar execution time.
  - Vector execution time.

VS FORTRAN Version 2 Library

A library of mathematical functions for scalar and vector execution that uses new algorithms to provide significantly improved accuracy compared to the VS FORTRAN Version 1 Library. Scalar performance is also improved.

The library includes highly-tuned vector intrinsic functions for optimum performance with the IBM 3090 Vector Facility for:
- Short precision.
- Short precision complex.
- Long precision.
- Long precision complex.

Other notable functions include:
- Scalar and vector versions that provide mathematically equivalent results.
- Functions compatible with object code produced by IBM FORTRAN G1, HX, Q and VS FORTRAN.
- Support for programs converted using the IBM FORTRAN Language Conversion Program.

VS FORTRAN Version 2 Multitasking Facility

Provides a set of routines in the VS FORTRAN Version 2 Library to permit a single VS FORTRAN application program to use all the processors of the host system simultaneously:
- Can be used for all multi-engine processors under MVS on which VS FORTRAN Version 2 operates.
- Applies equally to scalar and vector object code.
- Programmer is required only to insert FORTRAN calls to the Multitasking Facility at the beginning and end of program sections which can benefit from parallel execution.
- Existing as well as new programs can easily be adapted to use the Multitasking Facility.
- Measured performance gains for programs with 83 % to 93 % parallelism showed turnaround improvements of:
  - 1.7 to 1.8 times on a 2-way processor unit.
  - 2.6 to 3.3 times on a 4-way processor unit.
- User program turnaround depends on workload characteristics and system configuration.

VS FORTRAN Version 2 Interactive Debug

Provides all the function in VS FORTRAN Interactive Debug Release 2, plus new facilities:
- Batch debug capability.
- Extended full screen support including:
  - Windowing to permit display of source code alongside debug execution output.
  - Animation to highlight statement being executed.
  - New cursor-oriented commands for improved usability.
  - Individual statement and program unit timing information to aid program optimisation.
- No need for recompilation before repeated execution of the code being debugged.
- Interactive version requires Interactive System Productivity Facility (ISPF).

Languages and Libraries

Assembler H Version 2

Assembler H Version 2 Release 1 VFE accepts and compiles Assembler source programs which include the 171 vector instructions supported by the IBM 3090 Vector Facility.

The Engineering and Scientific Subroutine Library is written in Assembler, where necessary, for optimum performance.

Users may write critical routines of compute-intensive vectorised applications in Assembler to obtain maximum throughput.

Engineering and Scientific Subroutine Library - ESSL

The Engineering and Scientific Subroutine Library provides a high-performance set of mathematical subroutines which are optimised to exploit the characteristics of the IBM 3090 Processor Unit and the Vector Facility:

- Subroutines can be used in VS FORTRAN or Assembler programs under MVS/XA and VM/SP HPO.
- Scalar versions of all the subroutines are provided for development and testing and can be run in any system environment supported by the VS FORTRAN Version 2 Library.
- Accuracy is comparable to that available with leading state-of-the-art scalar algorithms.
- Primary applications supported by ESSL are:
  - Structural analysis.
  - Fluid dynamics.
  - Signal processing.
  - Process simulation.
  - Circuit and device analysis.
  - Reservoir modelling.

Short and long precision versions of almost all subroutines are provided for the following types of mathematical computation:

- Linear algebra: 36 vector-scalar and 14 matrix-vector subroutines.
- Matrix operations: six subroutines for matrix addition, subtraction and multiplication.
- Simultaneous linear algebraic equations: 16 subroutines solve linear systems of equations for a real general matrix, a real positive definite symmetric matrix, a real general band matrix, and a real positive definite symmetric band matrix.
- Eigensystems analysis: 4 eigensystems analysis subroutines find the eigenvalues and eigenvectors of a real general matrix and a real symmetric matrix.
- Signal processing: nine subroutines are provided to compute Fourier transforms, convolution and correlation. Eight general routines compute Ith zero crossing, polynomial evaluation, quadratic interpolation, and time varying recursive filter.
- Random number generation: two uniform random number generation routines are available for general use.

Other Programs

IBM FORTRAN Language Conversion Program - LCP

The IBM FORTRAN Language Conversion Program converts IBM FORTRAN language level 66 source language statements to the correct language level 77 equivalents automatically, to ease standardisation on language level 77.

IBM language level 66 compilers include FORTRAN G1, HX and Q as well as VS FORTRAN.

- The Language Control Program (LCP) provides users with:
  - Language level 77 functions.
  - 31-bit addressing under MVS/XA.
  - A single-compiler installation with simplicity of maintenance, uniformity of results and portability of programs.
  - Installation at the latest ANSI level.

- Operating environments are:
  - MVS/XA.
  - MVS/370.
  - VM/SP, with or without the High Performance Option.

- Source statements produced by the program can be compiled by:
  - VS FORTRAN Version 1 Release 3 or later
  - VS FORTRAN Version 2

- LCP execution support library routines are contained in the libraries of:
  - VS FORTRAN Version 1 Release 4.1 or later, as modified by the Small Programming Enhancement for the LCP.
  - VS FORTRAN Version 2

Vector Processing Subsystem/ Vector Facility-VPSS/VF

The Vector Processing Subsystem/Vector Facility (VPSS/VF) is a simulator of the IBM 3838 Array Processor that allows users to benefit from the IBM 3090 Vector Facility.

- Scalar operation is permitted on a 4381 or 30XX Processor Unit under MVS/XA.
- Vector operation is permitted on a 3090 with Vector Facility under MVS/XA and supported by the Engineering and Scientific Subroutine Library and VS FORTRAN Version 2 Library.
- VPSS/VF allows code written for the 3838 to be executed on the host.
- VPSS/VF dynamically allocates a virtual 3838 and automatically operates in vector mode if a 3090 Vector Facility is installed.
- Real and Virtual 3838s are supported on a 3090 by VPSS/VF permitting users to balance the workload between the host and available 3838s.
- VPSS/VF is designed to be transparent to the user:
  - No modification of code is required.
  - Results produced by a simulated 3838 are mathematically equivalent to those of a real 3838.
- Performance measurements of VPSS/VF in vector mode are scheduled for the fourth quarter of 1985.

Publications

Languages

Title	Reference
VS FORTRAN Version 2	
General Information Manual	GC26-4219
Licensed Program Specifications	GC26-4225
Installation and Customisation Guide	SC26-4224
Language and Library Reference	SC26-4221
Programming Guide	SC26-4222
Interactive Debugging Guide	GC26-4223
Diagnosis	SY26-3998
Reference Summary	SX26-3751
Assembler H Version 2	
General Information Manual	GC26-4035
Licensed Program Specifications	GC26-4029
Installation Manual	SC26-4030
Application Programming Guide	SC26-4036
Application Programming Language Reference	GC26-4037

Other Programs

Title	Reference
Engineering and Scientific Subroutine Library	
General Information Manual	GC23-0182
Licensed Program Specifications	GC23-0181
Installation Manual	SC23-0183
Guide and Reference Manual	SC23-0184
IBM FORTRAN Language Conversion Program	
General Information Manual	GC23-0154
Users Guide	SC23-0145
Installation Guide	SC23-0143
Licensed Program Specifications	GC23-0152
Vector Processing Subsystem/Vector Facility	
Licensed Program Specifications	GC23-0626
Installation and Customization Guide	SC23-0627
Licensed Progam Summary	GC23-0625

IBM Eurocoordination
S.A. au capital de 2.700.000 F
Siège social : Tour Pascal
22, Route de la Demi-Lune
92075 Puteaux
R.C.S. Nanterre B 304 538 192
France

IBM World Trade
Americas/Far East Corporation
Town of Mount Pleasant
Route 9, North Tarrytown
New York, 10591 U.S.A.

G 511-0135-0 *Printed in Denmark by Lassen Offset*

IBM 3090
A New Perspective for Engineers and Scientists

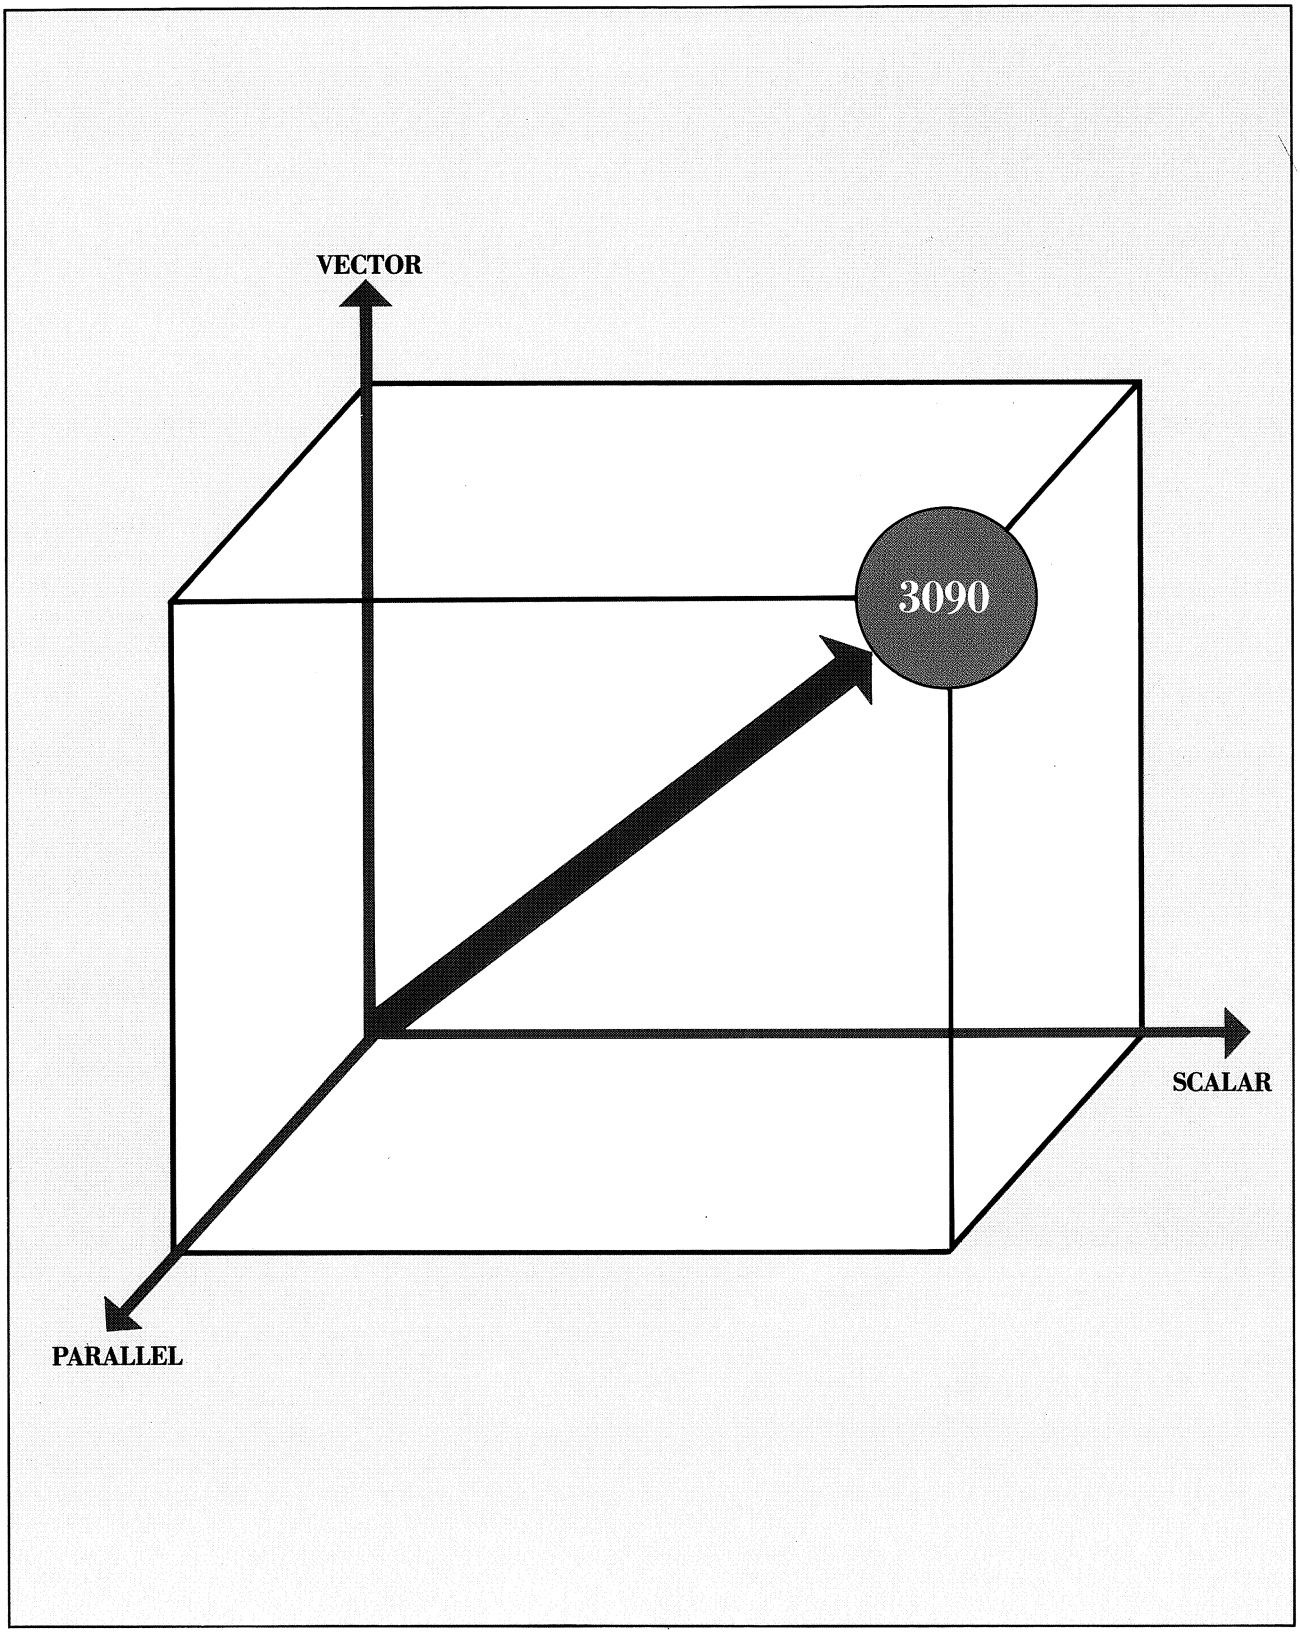

IBM

IBM 3090 - A New Perspective for Engineers and Scientists

From its conception the IBM 3090 has been designed for you, the engineer and scientist, to satisfy your need for greater computational power. In developing the 3090, IBM has provided scalar, vector and parallel computing capabilities to enable users to optimise individual job turnaround and total systems performance.

A solid scalar foundation

Virtually all programs require scalar processing which is why the IBM 3090 has built into it specialised hardware attributes to make it one of the fastest scalar computers available today.

Optional vector capability

Onto this scalar foundation IBM has added an optional vector processing capability, giving to users with highly iterative code (FORTRAN programs with a high proportion of DO loops) the opportunity to obtain a major performance improvement at a low incremental price.

Parallel processing via multitasking

For both scalar and vector applications where many independent calculations are involved, the IBM 3090's multi-processor design (with up to four central processors in the Model 400) coupled with the IBM VS FORTRAN Program Multitasking Facility, easily permits you to split a single program into multiple tasks. The tasks are run in parallel on all the central processors, producing major turnaround improvements for some long running programs. These can be even more significant if the programs are vectorised.

A single integrated system

All this computational function and performance is provided in one system under the control of a single system control program. Scalar, vector and parallel computation can be intermixed at will and the same data files can be shared by different programs or users as required, physically stored once on the 3090's storage devices.

A single compiler

The same powerful IBM VS FORTRAN Version 2 Compiler provides highly optimised scalar or vector code. Vectorisation is performed automatically when selected. In both cases, IBM-supplied subroutines provide state-of-the-art accuracy and performance.

Scalar and vector packages available

Many IBM and non-IBM application packages are available to help you implement the systems you need on the IBM 3090, including non-IBM vectorised code for structural analysis, fluid dynamics, dynamic simulation and mathematical subroutines. The latest European IBM catalogue of scalar programs available from non-IBM sources lists over 250 products.

Options for now and the future

The IBM 3090's multi-processor design, providing up to four central processors and four vector facilities, gives you a wide range of choice to match to your processing requirements today and to handle your growth in the years ahead.

Availability you can count on

Using innovative Thermal Conduction Module (TCM) packaging technology, a memory design with powerful automatic error correction and an online Remote Service Facility for maintenance, the IBM 3090 is designed to surpass IBM 308X computers which have a world-wide average of better than 99.5 % availability.

Protection for your investment

Operating under IBM's System/370 Extended Architecture and System/370 Architecture, the 3090 continues a tradition, which goes back over 20 years, of providing compatibility with previous IBM large computers; permitting you to invest in IBM hardware and software with confidence for the future.

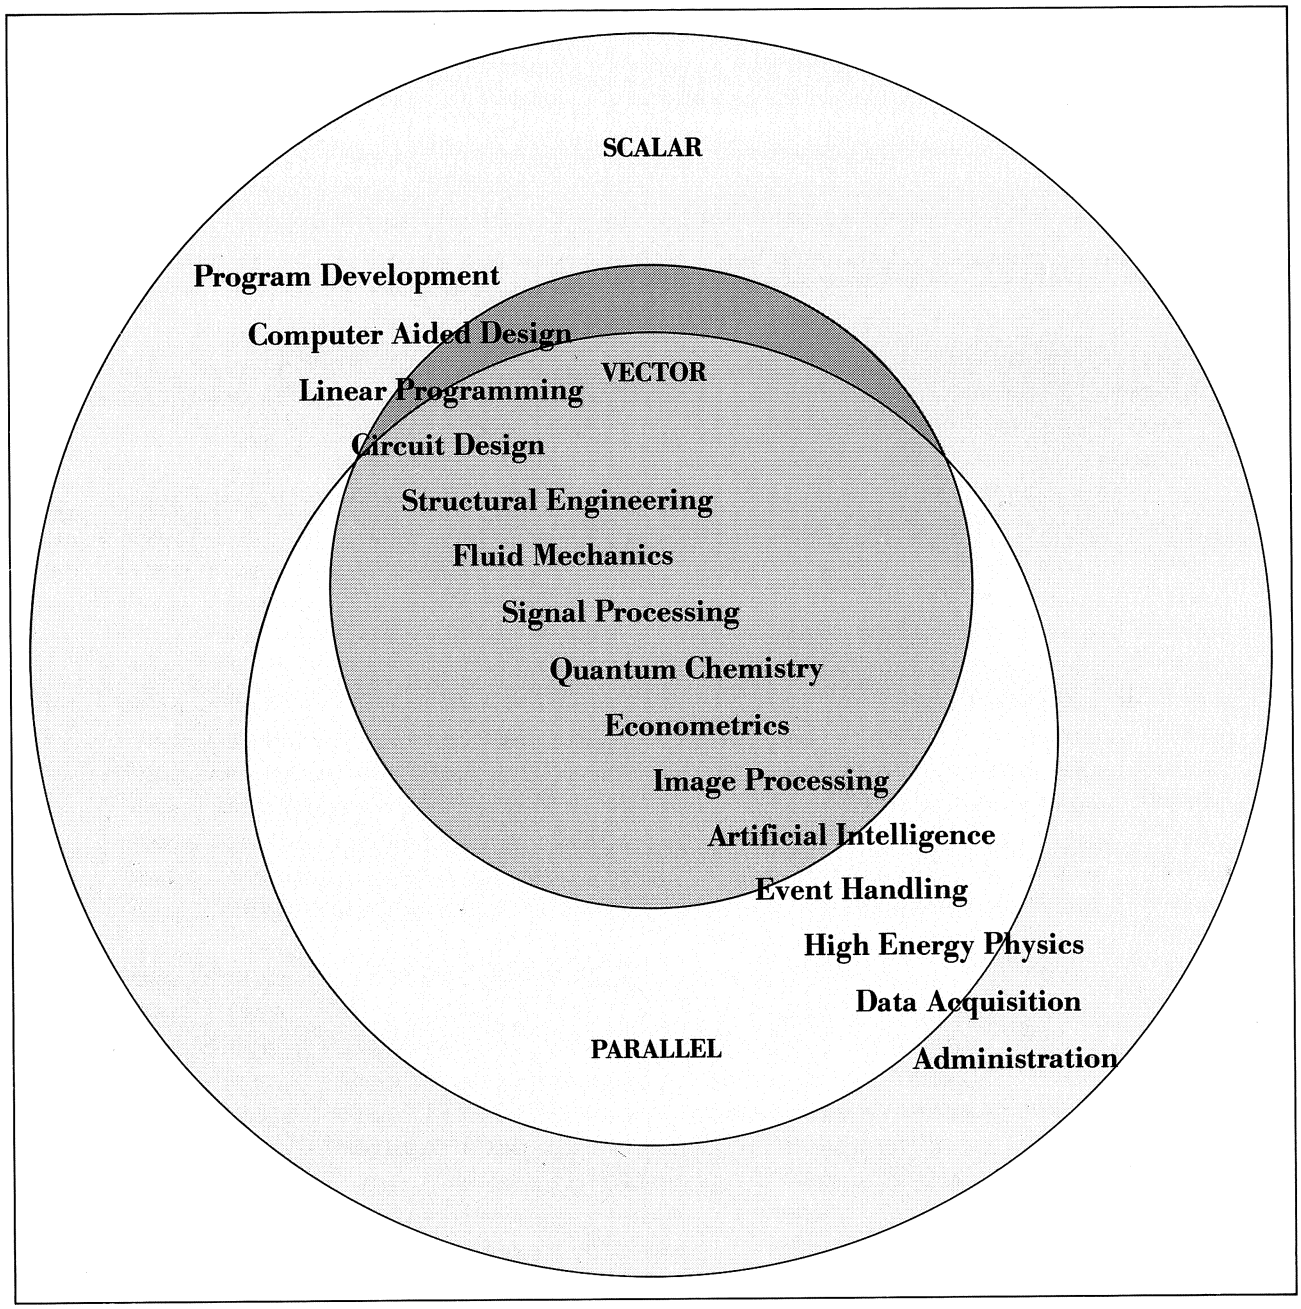

Sample Applications by Processing Domain

IBM 3090 - Designed for the Engineer and Scientist

Computers increasingly vital

Whether your field is engineering design or geophysical exploration, government research or higher education, computers are being used more and more to solve problems and to develop new products.

As computer price/performance continues to improve, so does the sophistication of the computer techniques used by the engineering and scientific community.

In manufacturing industries, cost and competitive pressures are causing engineers to supplement wind-tunnel testing and vehicle crash trials with computer simulation. Elsewhere computers provide the only acceptable solution. For example, in oil reservoir modelling, estuary simulation and astrophysics, direct physical experimentation is either not feasible or inaccurate; in seismic processing and defence applications, the best attainable results justify the investment; and in circuit design and solid modelling, only computers can handle the complexity.

The IBM solution

The IBM 3090 has been designed to provide the engineering and scientific computing solution to your growth requirements: it combines in a single integrated system all the facilities needed to obtain outstanding performance in all application areas.

A major opportunity

Two computing techniques provide engineers and scientists with a major opportunity to improve job turnaround, examine more alternatives, use more refined models or embark on new areas of study. These techniques, parallel and vector processing, offer performance gains which far outstrip those available just from developments in computer electronics. Both are available with the IBM 3090.

Two complementary approaches

In most computers, one instruction operates on only one data pair and one program executes in only one processor at any given moment. This is traditional scalar processing.

Vector processing permits a single computer instruction to perform the same operation on many data pairs in swift succession. Parallel processing permits a single program to be split into independent elements which are executed by separate processors in a multi-processing computer simultaneously.

Each of these techniques used on its own can provide significant performance improvements. When combined on an IBM 3090, even greater performance improvements are possible.

Scalar performance is fundamental

Good scalar performance is essential, not only for those programs which cannot be vectorised, but also for those which can. Even vectorised programs can contain significant sections of scalar code. That is why IBM started by designing the 3090 to be one of the fastest scalar computers available today.

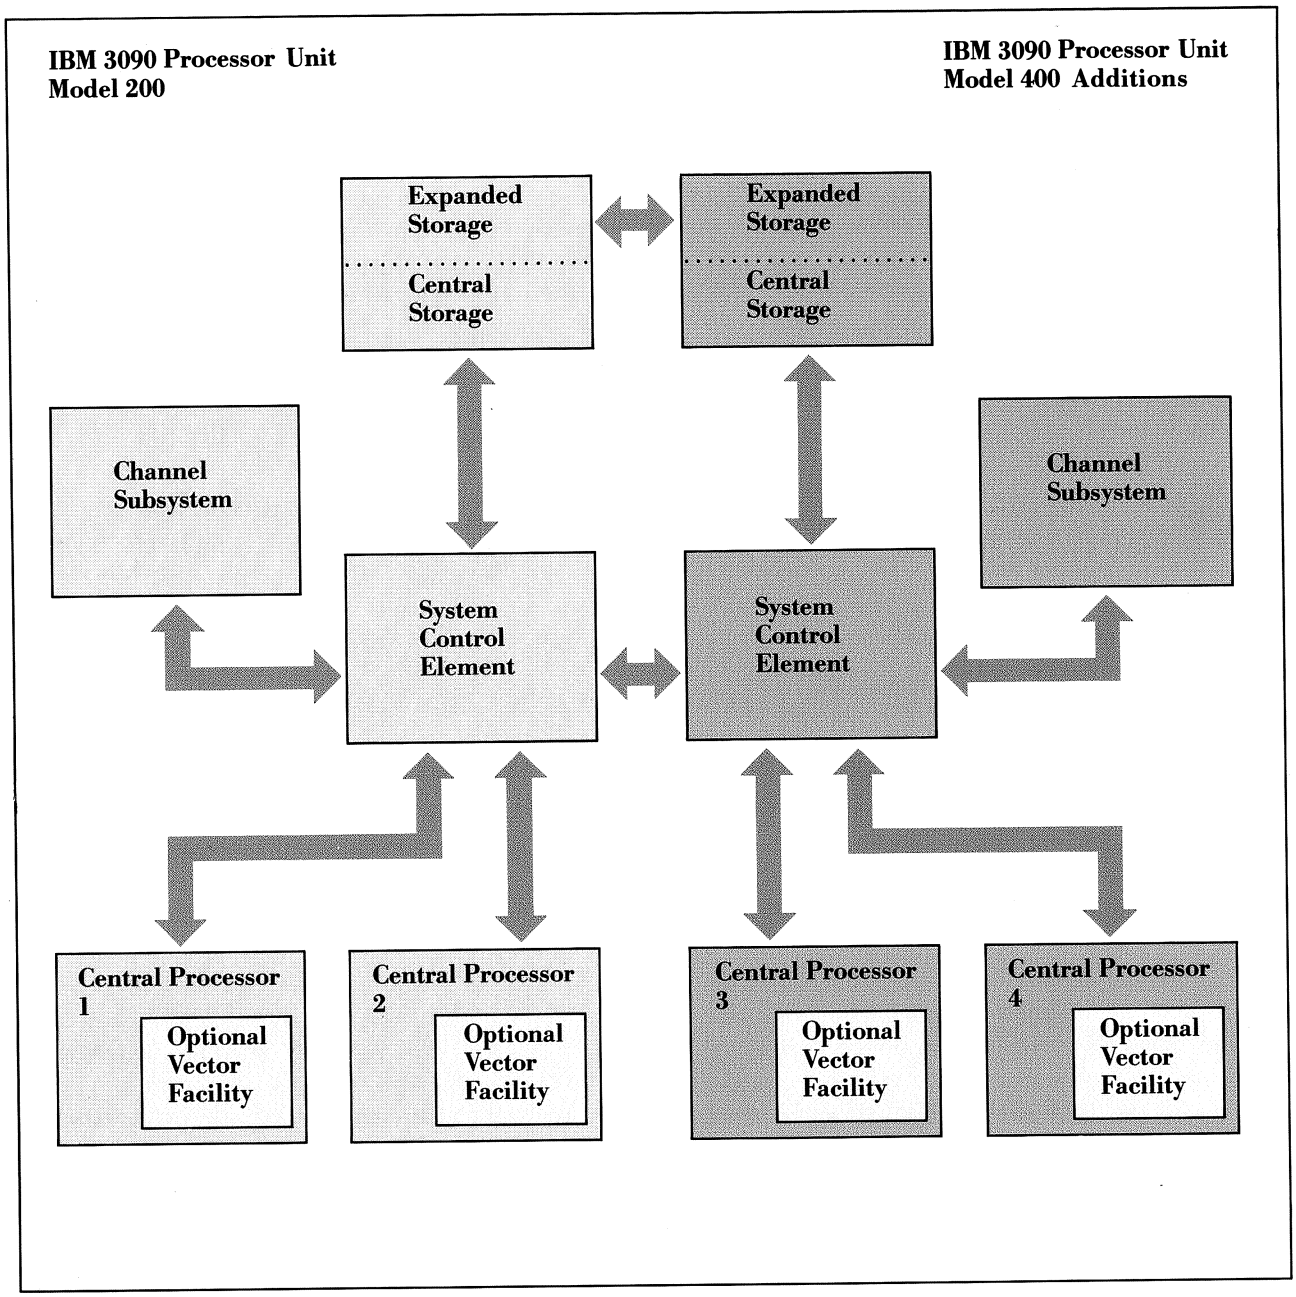

IBM 3090 Processor Unit Design

IBM 3090 - One of the world's fastest Scalar Computers

Engineering, scientific and commercial capability

The IBM 3090 provides outstanding scalar performance for compute-intensive applications through a combination of architecture, technology and design which optimises the processor unit's performance for integer and floating point calculations. As with previous System/370 and System/370 Extended Architecture processors, the IBM 3090 retains a comprehensive commercial processing capability.*

Extensive range of computing options

The general purpose instruction set of the IBM 3090 contains 212 scalar instructions including 52 scalar floating point instructions. These provide the capability of performing 32-bit integer and 32-bit (short), 64-bit (long) and 128-bit (extended) precision floating point operations. The number range supported by the IBM 3090 hardware is from 10^{+76} to 10^{-77} which can be extended by IBM software from 10^{+9861} to 10^{-9865}. The floating point formats permit a precision of from 6 to 7 decimal digits operating with short precision, to 32 to 33 decimal digits operating with extended precision. Error checking is used throughout the IBM 3090 processor complex and its attached IBM input, output and storage devices.

Central processors optimised for fast engineering/scientific computation

A separate channel subsystem for handling Input/Output (I/O) and a system control element for handling accesses to shared central and expanded storage leaves the central processors in an IBM 3090 free to concentrate on program execution. Execution of compute-intensive programs is enhanced by special hardware in all the central processors. This includes a hardware multiplier, optimised add/subtract logic and special loop control circuitry to improve execution performance. Overlapped instruction and execution elements within each central processor provide a highly efficient environment for scalar computation.

Large, fast storage for high performance

IBM System/370 Extended Architecture and the storage hierarchy of the IBM 3090 combine to give the engineer and scientist a liberal programming environment under MVS/XA. Each program may occupy a virtual address space of up to 256 Megawords (Mw) or 2048 Megabytes (Mb). Within the 3090 up to 48 Mw (384 Mb) of real electronic storage provide fast access to data and minimise paging to I/O devices. Use of a three-level storage hierarchy, which is transparent to the programmer, minimises cost and maximises performance. Each central processor has its own 8 Kw (64 Kb) high speed buffer by which most requests for instructions and data are satisfied. Computational routines have a high chance of remaining in the high speed buffer during execution with resulting high performance.

Balanced and responsive I/O subsystem

64-bit data paths are used throughout the IBM 3090 Processor Unit. Up to 96 channels, each of which can operate at up to 3 Mb per second, provide balanced I/O performance for the powerful IBM 3090 central processors. Operating in System/370 Extended Architecture mode, up to four channel paths may be assigned to each I/O device and any path may be selected for each I/O operation. This permits a high level of channel utilisation without loss of performance and can ensure that fast response is obtained from heavily utilised disk files such as program libraries and data bases.

Exceptional performance independently confirmed

This all adds up to a computer with exceptional scalar performance which has been confirmed by independent measurements. The MacNeal-Schwendler Corporation provides a booklet entitled "Time Estimation and Problem Execution" with MSC/NASTRAN.® The June 1985 edition shows that the IBM 3090 executes the MSC/NASTRAN® multiply-add loop faster than any other non-vector computer listed in the publication.**

* If you are interested in the commercial aspects of the IBM 3090, please contact your Marketing Representative to obtain the latest IBM 3090 brochure for commercial users.

® NASTRAN is a registered trademark of the National Aeronautics and Space Administration. MSC/NASTRAN is an enhanced proprietary version developed by the MacNeal-Schwendler Corporation.

** IBM 3090 Vector Facility performance information is expected to be published in October 1985.

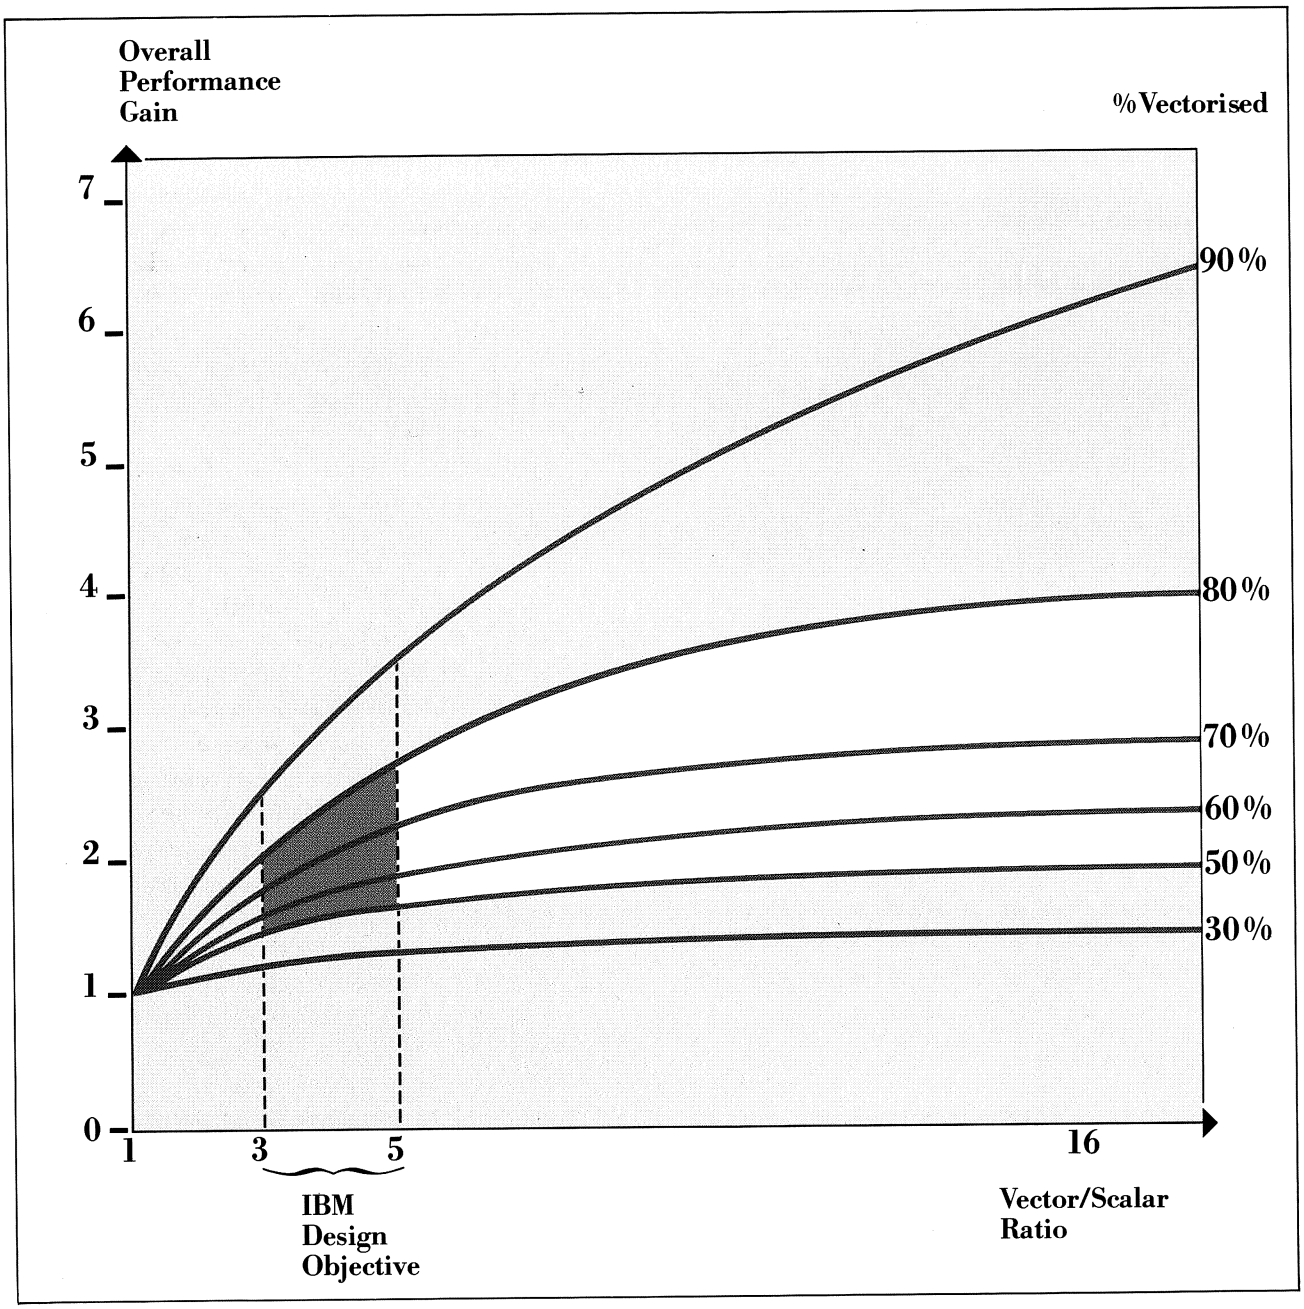

Performance versus Vector/Scalar Ratio

High performance Vector Facilities at low incremental price

Integrated into the IBM 3090

One optional field-installable Vector Facility can be added to each central processor in an IBM 3090 Processor Unit, providing up to four facilities on the largest Model 400. Each Vector Facility forms an integral part of the central processor to which it is attached, providing an attractive performance option with a low incremental price.

Designed to match achievable levels of vectorisation

Even in cases where high levels of vectorisation are theoretically achievable, practical considerations frequently result in achieved levels of vectorisation in the range of 50 % to 80 %. Since the faster execution offered by vector processing only affects the vectorised portion of a program, the law of diminishing returns applies when very fast vector processors are used for programs with low to medium levels of vectorisation. In such cases, scalar performance is equally or more important. The IBM Vector Facility has a design point of three to five times faster than 3090 scalar performance which, as the diagram shows, provides significant improvements in the range of one-and-a-half to three times 3090 scalar performance for typical vectorisable applications.

Powerful vector processing facilities

171 additional instructions are supplied with the Vector Facility, increasing the IBM 3090's instruction set to 383 instructions and providing an exceptionally broad spectrum of machine level capability. Each Vector Facility includes a vector execution element and a set of vector registers. The registers, which contain 128 elements, can be used as sixteen 32-bit registers or eight 64-bit registers for integer and floating point arithmetic. The ability of the registers to hold up to 128 vector elements is expected to be optimum for many applications. When vectors exceed this lenght, automatic sectioning is provided by the IBM VS FORTRAN Version 2 Compiler. Flexible addressing options permit matrices to be processed by row or column taking each element in turn or randomly.

Up to two operations per cycle

The vector execution element is pipelined permitting one result to be produced every cycle. When compound instructions are performed, two operations are executed every cycle. Compound instructions, such as Multiply-and-Add, are designed for commonly-used computations such as dot product and matrix summation.

Easy to implement

Introducing vector processing with the IBM 3090 is very easy. All you have to do is recompile existing language level 66 or 77 IBM FORTRAN programs using the IBM VS FORTRAN Version 2 Compiler, which will result in the automatic generation of vectorised object code. No changes to source code or files are required.
According to user choice, programs can be tested in scalar mode before or after Vector Facility installation. After installation, the scalar code can be used to run the same programs in other locations where a Vector Facility is not available, or as backup. This is because the Vector Facility is an integral part of the 3090 and current IBM VS FORTRAN and Assembler compilers have been extended to support the additional vector instructions.

Cost effective processing options

Introducing vector processing by adding Vector Facilities to an installed 3090 and installing a new 3090 with Vector Facilities are cost effective alternatives. In both cases, mixed scalar and vector as well as commercial work can be run on the same 3090. (Vectorised programs are automatically executed only on central processors which have vector capability.) In this way, the considerable asset that the 3090 vector processing computer represents can be fully utilised, even at the outset when few vectorised programs have been implemented.

Because existing applications will run faster when vectorised, additional capacity can become available for scalar processing once a Vector Facility is installed, thus postponing the need for a more expensive model upgrade or additional system.

3090 Processing Environment	Approximate relative job turnaround time
One Central Processor	
Scalar	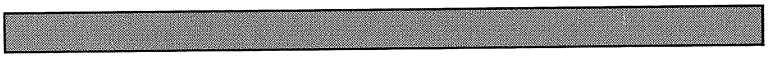
Vector Facility (VF)	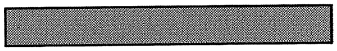
Two Central Processors	
Two-Way Parallel - Scalar	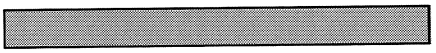
Two-Way Parallel - Two VF	
Four Central Processors	
Four-Way Parallel - Scalar	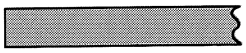
Four-Way Parallel - Four VF	

- Approximate relative job turnaround based on actual IBM 3090 scalar and vector measurements, with parallel processing performance extrapolated from IBM 3090 two-way and IBM 3084-Q four-way scalar processing measurements.
- The Black Oil Simulation Model is the application used as a basis for this performance comparison. Complete measurement data is scheduled for the first quarter of 1986.

IBM 3090 Performance Opportunities

Parallel Processing for Scalar and Vectorised Applications

Breaking through the technology barrier

Electronic technology is a key factor affecting the rate at which individual processors can be made to run faster. Multiple processors are a means of obtaining greater performance: using multi-processing to obtain greater aggregate performance and parallel processing to obtain greater program performance.

In multi-processing a single control program allocates many programs to be executed on the processors making up the multi-processing complex. A given program is executed on only one processor.

In parallel processing a single program may be divided into separate, independent tasks which are then executed concurrently on the different processors; providing the opportunity for significant program turnaround improvements, depending on the degree of parallelism in the program, the number of parallel processors available and the efficiency of the multi-processing design. Parallel processing is of particular interest for critical long-running jobs.

IBM multi-processing design approach

IBM has long recognised the potential limitations of a single processor computer design. IBM's first multi-processor, the System/360 Model 67, was introduced in 1965. Today MVS/XA provides for up to 16-way multi-processing. With experience stretching back 20 years, IBM's range of 3090 dyadic and four-way computers embody the design expertise needed to provide highly effective parallel and multi-processing.

Easy parallel processing

FORTRAN programmers can easily obtain the benefits of parallel processing using the Multi-tasking Facility in IBM's VS FORTRAN Version 1 and Version 2 Program Libraries. Only FORTRAN calls are required, to indicate the beginning and end of program sections which can be executed independently. During execution, the MVS system control program uses its multitasking capability to assign these sections to different processors and control the synchronisation requested by the programmer.

Parallel processing is not limited to programs which operate on vectors. Any scalar program where a significant amount of independent computation is required can also benefit. Some applications are more suitable for parallel processing than vector processing.

Combined vector and parallel processing provides best job turnaround

With both vector and parallel processing facilities available on the IBM 3090, compute-intensive users have the opportunity to obtain significant improvements in job turnaround. Performance gains greater than three-fold have been measured using each technique. IBM 3090 measurements showing the benefits of both techniques together are scheduled for the first quarter of 1986.

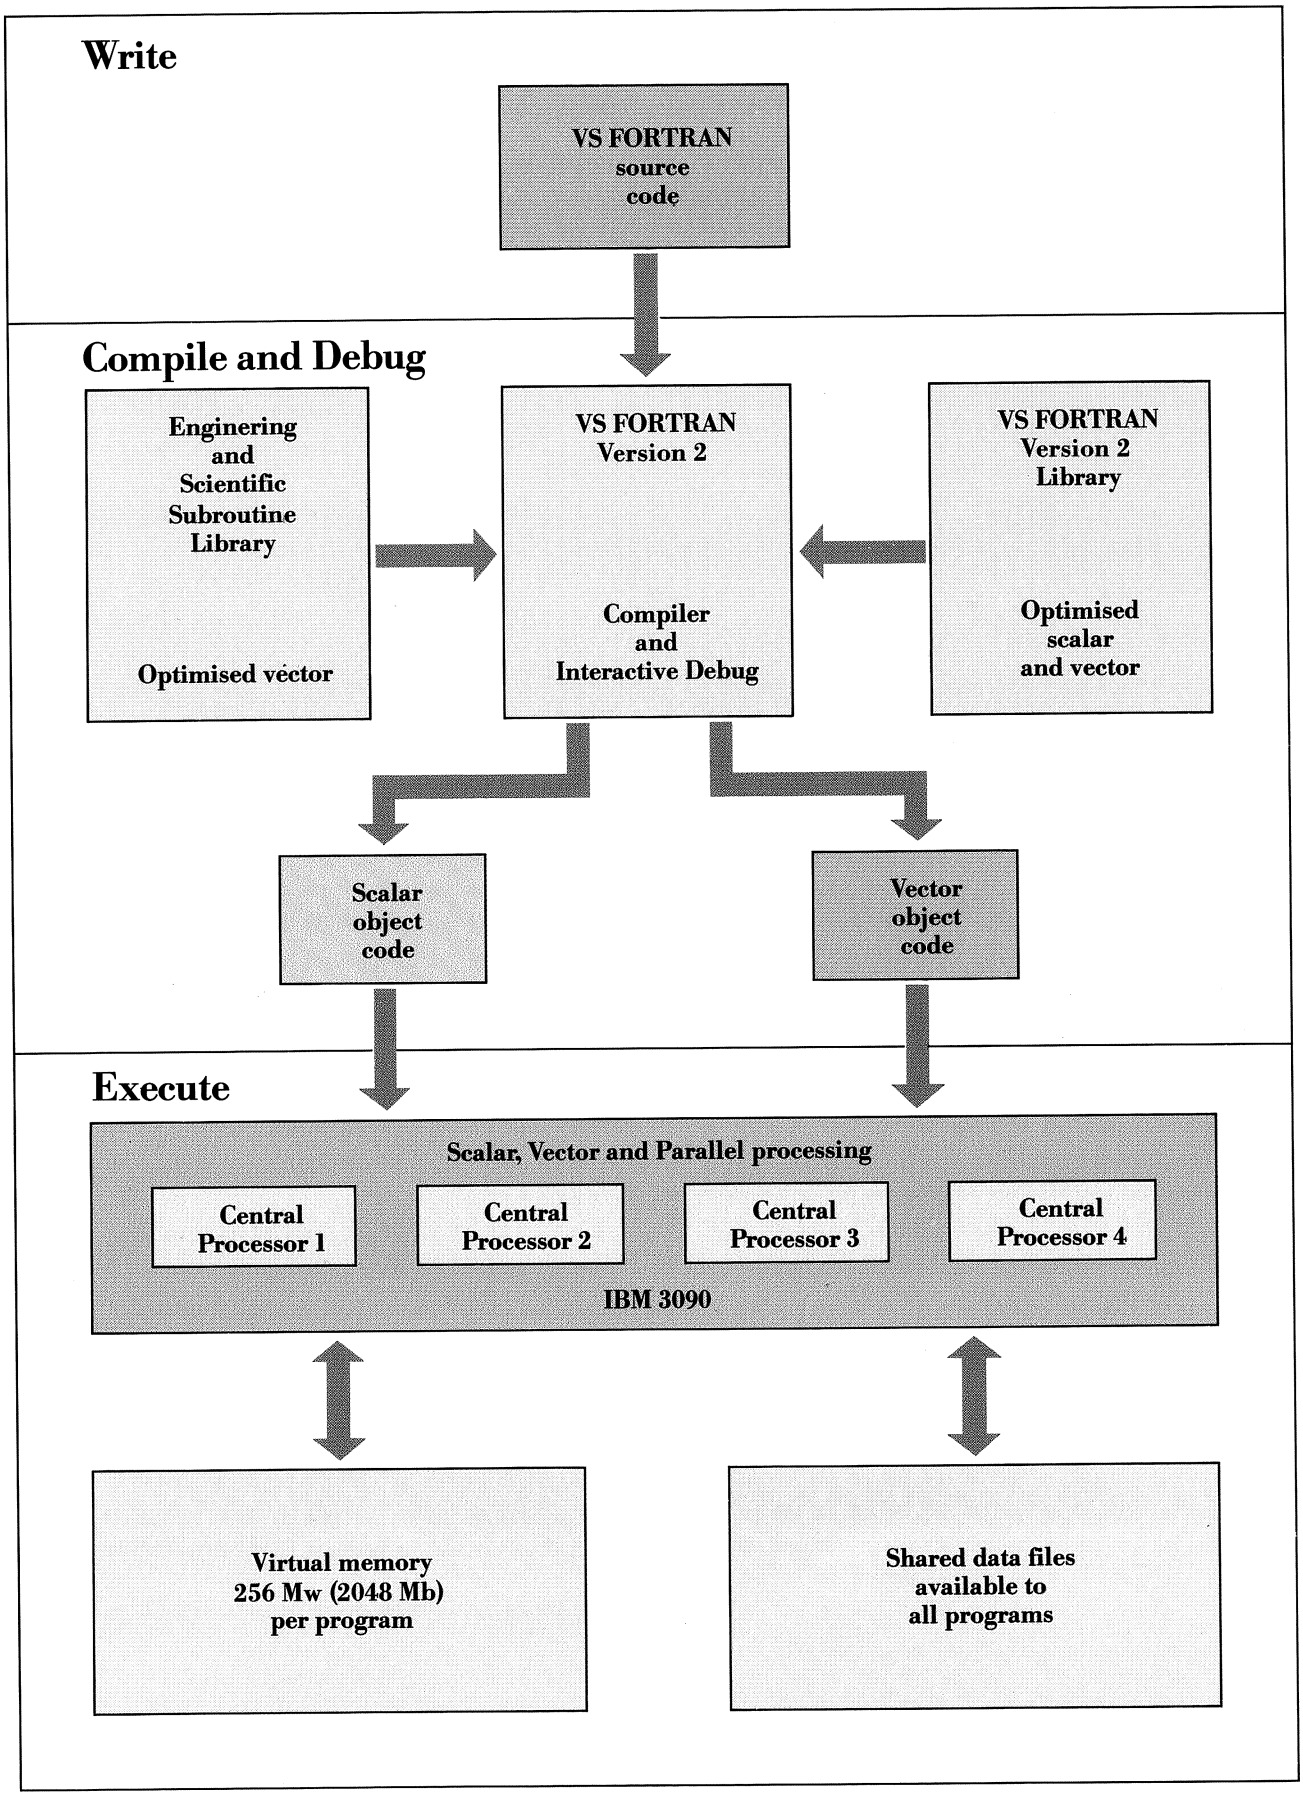

IBM 3090 - Unified Hardware and Unified Software

Comprehensive Systems Software Support

A unified approach

The IBM 3090 provides hardware for scalar processing, vector processing and parallel or multi-processing in a single processor complex. Likewise IBM's system software support utilises a unified approach: one compiler provides FORTRAN object code for scalar or vector processing, in a parallel processing environment if required. Two system control programs, MVS/XA and VM/SP High Performance Option, support scalar and vector processing simultaneously in the same IBM 3090 processor complex.

MVS/XA provides four-way and parallel processing

MVS/XA exploits the IBM 3090's hardware and provides a liberal environment for end users and programmers. This includes a virtual storage capacity of up to 256 Mw (2048 Mb) for each program. In addition MVS/XA supports all four processors of the largest IBM 3090 Model 400 as part of a single logical system, minimising control program overheads and allowing FORTRAN users to perform four-way parallel processing with the VS FORTRAN Program Multitasking Facility.

Simple end-user interaction with VM

The VM/SP High Performance Option (VM) provides a virtual machine environment. Together with the Conversational Monitor System (CMS), it permits each end user to access the IBM 3090 system independently, with the perception that he or she is the sole user. Working in conjunction with VM/CMS, the Engineering/Scientific Support System (E/S^3) is designed to enable the engineer and scientist to perform program development and execution in a simple manner, through consistent task-oriented dialogues which require a minimum of system knowledge. VM also provides guest System Control Program (SCP) support, not only for other IBM SCPs but also for IX/370, IBM's implementation of UNIX®.

® UNIX is a registered trademark of AT&T Bell Laboratories.

Scalar, vector and parallel processing support with VS FORTRAN

IBM's latest FORTRAN compiler is designed to enable the engineer and scientist to exploit the scalar, vector and parallel processing capabilities of the IBM 3090 easily. In a single product, VS FORTRAN Version 2, IBM provides state-of-the-art compilation to produce highly-optimised object code for scalar or vector execution, using the same source code. Language level 66 and 77 source code is accepted by VS FORTRAN Version 2 and by VS FORTRAN Version 1, the currently available scalar version of the compiler. Both compilers include the FORTRAN Program Multitasking Facility, permitting parallel execution of programs.

Extensive subroutine libraries provide highly optimised code.

To provide the engineer and scientist with efficient code for the execution of widely-used mathematical functions, there are FORTRAN library products for both scalar and vector processing. The VS FORTRAN Version 2 Library contains highly-tuned intrinsic functions whose latest scalar and vector versions provide enhanced accuracy using recent algorithm developments.

The Engineering and Scientific Subroutine Library (ESSL) provides a wide range of application-oriented subroutines optimised for vector execution. In addition, ESSL contains duplicate routines written in scalar code to facilitate program development and testing on 3090s without a Vector Facility or other IBM computers. ESSL's wide range of 95 subroutines supports such applications as structural analysis, fluid dynamics, signal processing, process simulation, circuit and device analysis and reservoir modelling.

Debug facilities for rapid program implementation

The enhanced Interactive Debug provided with VS FORTRAN Version 2 provides state-of-the-art facilities, including full screen support with windowing and animation. Using the individual source statement and program unit timing information provided, FORTRAN programmers can not only debug but also tune their programs in a productive way, while executing on the IBM 3090 in scalar mode or actually using the Vector Facility.

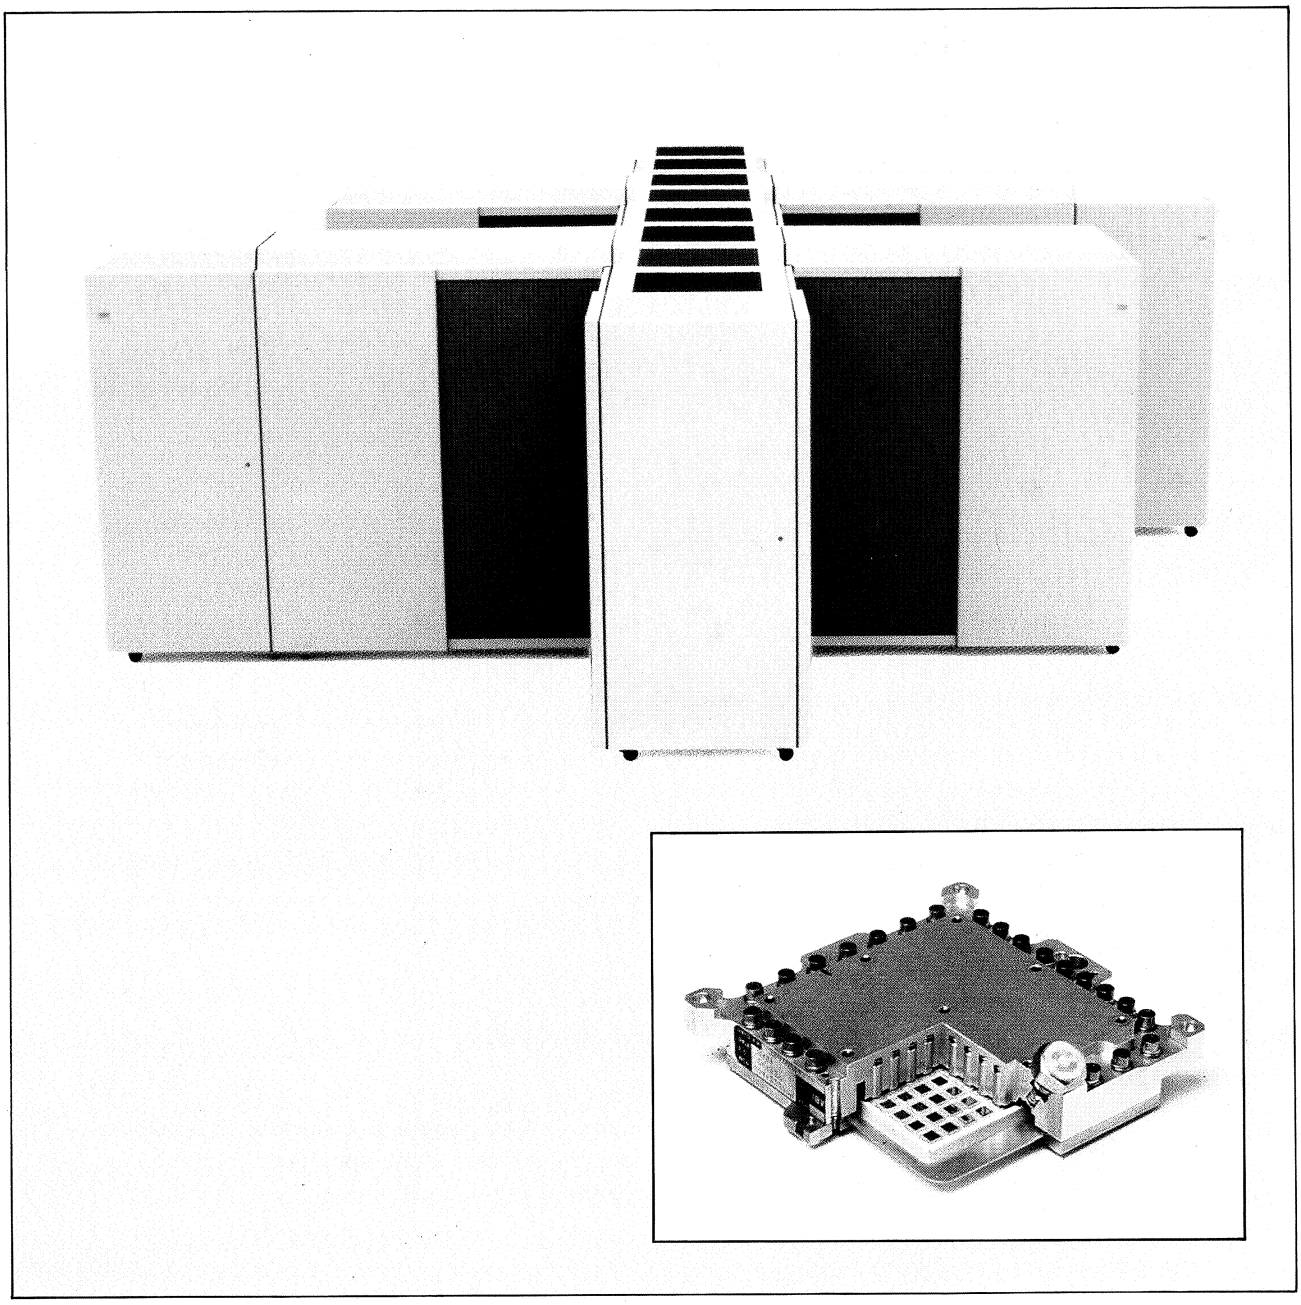

IBM 3090 Model 400 with four Vector Facilities and Thermal Conduction Module (inset)

IBM 3090 - A modular, high-quality European product

Flexible performance options

The IBM 3090 provides a wide range of growth and processing options to permit you to tailor your initial configuration to your requirements and expand it later as required. Scalar processing power can be increased by field-upgrading the Model 200 dyadic processor to a four-way Model 400 which is expected to provide a 70% to 90% increase in performance. Vector processing power can be added in four increments. Each Vector Facility can increase the performance of its associated central processor one-and-a-half to three-fold. Memory options from 8 Mw (64 Mb) to 48 Mw (348 Mb) and channel options from 32 to 96 channels are also available.

High availability

Users need high availability with a computer as powerful as the IBM 3090. That makes IBM's determination to provide better availability with each successive generation of products very important for engineers and scientists contemplating the purchase of an IBM 3090. The world-wide average availability of IBM's 308X computers is better than 99.5% and this provides the benchmark for the 3090.

Proven technology for reliability

The high-speed Emitter Coupled Logic (ECL) chips, which are used to create the circuits of the IBM 3090 and its optional Vector Facilities are connected to their multilayer ceramic substrates by IBM-developed controlled-collapse chip connections. This technique, now over 20 years old, has stood the test of time and provides a durable, trouble-free bond. Interchip connections are established via printed conductors embedded in the multiple layers of the IBM-developed ceramic substrate, kiln-fired to form a single solid block. The substrate is mounted in an IBM Thermal Conduction Module (TCM) which provides a stable operating environment via water cooling, while protecting the chips and connections from corrosion in a sealed helium gas filled space.

The TCMs are plugged into TCM boards via zero-insertion-force spring connectors which permit easy removal of TCMs for upgrades, engineering changes or repairs while ensuring effective connection to the boards. The multilayer epoxy/glassfibre printed-circuit boards provide reliable connections to other TCMs on the same board. Connections between boards and other functional units of the processor unit are made with short multiwire cables.

Recovery and service support

The IBM 3092 Processor Controller, which contains duplexed components, forms part of a 3090 processor complex and performs monitoring, error logging, reconfiguration and diagnostic functions. It works with the 3090 to achieve automatic error recovery whenever possible. In the case of an unrecoverable failure, sections of shared central and expanded storage, complete central processors and channels can be taken off line; often without interruption to processing.

Built-in diagnostics allow the 3092 to perform error diagnostics in parallel with 3090 operation. The operator is alerted when a serious problem occurs and, using interactive customer Problem Analysis, can often restart the system. In the case of a solid 3090 failure, the 3092 can transmit fault and diagnostic information to a remote service centre, where specialists can exercise the 3090 remotely and diagnose the problem. Frequently this diagnosis can be performed fast enough for the required spare part to be delivered by the time the IBM customer engineer arrives.

Built in Europe

The IBM 3090 reflects IBM's policy of manufacturing products in Europe. For example: the memory chips, TCM substrates and TCM boards are made in Germany. The system consoles are built in the United Kingdom. The logic chips are made in France, where final assembly and testing of the IBM 3090, its TCMs and TCM boards also takes place. Supplying these IBM plants are many European companies whose contributions range from the keytops for the system consoles to the complete frames of the processor units. All of which demonstrates IBM's positive and continuously developing contribution to the European economy.

Research and Development in IBM

Each year, IBM makes significant investments in fundamental and applied research and development, related to all aspects of information processing and its application in today's world. This work is carried out in research and development facilities throughout the world, and results in new products, new manufacturing processes, and new applications of computing. The development and dissemination of much of this work is shared with scientific and industrial communities to promote the further application and understanding of advanced technology.

IBM Development Laboratories in the United States

Austin: Personal computers and low-cost workstations.
Boca Raton: Low-cost personal computers, small systems, modular products, specialised applications.
Boulder: Copiers and printers.
Burlington: Semi-conductor components.
Charlotte: Printers and financial systems.
Danbury: Scientific instrumentation and application programs.
East Fishkill: Semi-conductor components and ceramic substrates.
Endicott: Printers, small and medium range systems and software, ceramic substrates, printed circuits and circuit boards.
Gaithersburg: Telecommunications systems, air traffic control, global telemetry systems, special systems, complex systems software.
Houston: Space shuttle programs, Space Mission Control Headquarters programming, in-flight software.
Kingston: Distributed systems, technical and scientific processors, data entry systems, graphics systems.
Lexington: Typewriters, printers, keyboards.
Manassas: Submarine systems, signal processing.
Owego: Aerospace processors and systems, command and control systems, electro-magnetic systems.
Poughkeepsie: Large processors, operating systems and software.
Raleigh: Telecommunications and application programs.
Rochester: Medium range systems and software, low-cost mass memory.
San Jose: Disk storage systems.
Santa Teresa: Software and programming aids.
Tucson: Magnetic tape systems, disk storage units and printers.
Westlake: Space programs.

Canada

Toronto: National language adaptation and programming support.

Japan

Fujisawa: Handling of uncoded data, interactive display systems.

IBM Research Laboratories

Yorktown Heights (USA): Fundamental, technological and applied research in the fields of physics, mathematics, and semi-conductors.
San Jose (USA): Fundamental, technological and applied research into physics, computer science and data storage systems.
Zurich (Switzerland): Fundamental and technological research into solids, communications and computer science.

IBM Development Laboratories in Europe

France

Corbeil-Essonnes (near Paris): Very Large Scale Integration (VLSI) logic and high speed memory semi-conductor components.
La Gaude (near Nice): Telecommunications systems, signal processing, communications controllers, data communications, data networks.

Germany

Boeblingen (near Stuttgart): Medium range computer systems and processors, operating systems, high speed printers, Very Large Scale Integration (VLSI) semi-conductor components, finance systems, man-machine communications, scientific pilot projects.
Sindelfingen (near Stuttgart): Software development for computer-aided publishing, information retrieval, videotex, industrial, financial and cross-industry applications, networking.

Italy

Rome: Software development for project management, linear programming, production planning, application development productivity, distributed systems support, text management.

United Kingdom

Hursley (near Winchester): Graphic display products and programming, communications programming, low-cost disk storage, advanced display technology.

IBM Scientific Centres in Europe

IBM maintains a strong link with academic and scientific organisations in Europe, to explore the use of new data processing techniques, as applied to scientific problems. Centres in Paris, Rome, Heidelberg, Madrid, Pisa and Winchester carry out work in areas such as applied mathematics, computer science, expert systems, digital image processing, earth and environmental sciences, medical applications, speech and text processing.
The scientific computing facility in Rome, called the European Centre for Scientific/Engineering Computing, assists European scientists in carrying out advanced research in such areas as quantum chemistry, fluid dynamics, high energy physics and biology. In addition the centre has facilities for practical experimentation with parallel processing which are being extended in 1986 to include vector processing.

IBM Eurocoordination
S.A. au capital de 2.700.000 F
Siège social : Tour Pascal
22, Route de la Demi-Lune
92075 Puteaux
R.C.S. Nanterre B 304 538 192
France

IBM World Trade
Americas/Far East Corporation
Town of Mount Pleasant
Route 9, North Tarrytown
New York, 10591 U.S.A.

G511-0136-0 Printed in Denmark by Fair-Print AS, Roskilde