K-1 Product Summary

Report KC-177, Version 1.0
October 28, 1988

Key Computer Laboratories, Inc.
4245 Technology Drive
Fremont, CA 94538 USA
(415) 770-0169 FAX (415) 770-0493

(Page Intentionally Left Blank)

1.0 Overview

The K-1 product family consists of multi-user, multiprocessor supercomputers. The K-1 family will be a breakthrough in computing power for a very wide base of potential customers because of its high scalar and vector throughputs, large memories, and high I/O capabilities. This document describes the products and services with which Key will establish its position as a world leader in high-performance computing.

1.1 Introduction

Key's products, the K-1 family of supercomputers, are based on an evolutionary (von Neumann) architecture. They can be considered as descendants of current top-of-the-line supercomputers (superpipelined, RISC-like instruction set), but are optimized to overcome the application mismatches that result from existing supercomputer designs. Specifically, Key's K-1 supercomputers are designed to:

- Substantially improve scalar performance, while retaining high vector processing capability, allowing for higher system throughput.
- Support a large address space (48-bits) and allow the effective use of huge physical (32 Gbytes in the initial implementation) and virtual memories.
- Provide huge I/O rates (> 1 Gbytes/second) to balance processing power.
- Support efficient execution of the UNIX operating system.
- Support a high-utilization multiprocessor (eight processors in the initial implementation).
- Have large high-bandwidth cache memories, allowing for the cost-effective support of huge main memories, while retaining fast scalar processing rates.
- Support industry standards such as enhanced UNIX System V, IEEE floating-point, byte addressable architecture, VME and HSC I/O busses.
- Have high compilation speeds (> 60,000 lines per minute) for unoptimized as well as highly optimized code (> 15,000 lines per minute).

These features are implemented with very dense and very fast ECL gate-arrays. which are packaged with an advanced liquid cooling technology that Key has developed. The result is a supercomputer with a 6 nanosecond clock period containing over a million gates, and whose core CPU fits on a single board.

Summaries of the user requirements, system features, and resultant benefits of the K-1 hardware, software, and services are listed in Tables 1, 2, and 3, respectively.

Requirement	K-1 System Feature	User Benefit
High scalar, vector, and system code performance	2.7 GFLOPS system peak; 54 MFLOPS LFK harmonic mean per CPU, 76 MFLOPS 100 x 100 LINPACK double precision FORTRAN per CPU	Balanced performance across a majority of mainstream applications
Highly configurable	1 to 8 CPU multiprocessor, fully field-upgradable from entry to maximum system; 2 to 64 I/O processors	Protects investment as requirements grow, excellent multi-user system
Large memory capacity	512 Mbytes to 32 Gbytes shared among all CPUs	Reduced paging, improved performance on jobs where the application demands lots of data
Large disk capacity	> 1 trillion bytes maximum	Can store very large files and databases
High I/O throughput	> 1.0 Gbyte/sec maximum	Fast access to data sets
Convenient packaging	Air cooled to ambient main system, self-contained water-cooled subsystems	Lower operational costs and higher reliability than other systems
Standards support	IEEE floating-point, VME, HSC I/O busses, FDDI, and others	Compatibility, connectivity with other systems
Price	$1 million to $10 million	Affordable, expandable
Price/performance	4 times that of Cray Y-MP in 1991	Affordable and applicable

Table 1. Hardware Requirements, K-1 Features, and User Benefits

Requirement	K-1 System Feature	User Benefit
UNIX operating system	System V Release 3.2 with BSD extensions passes System V and NBS test suites	Application portability protects software investments
Supercomputing extensions to UNIX	High-performance I/O, multiprocessing support, etc.	High performance deliverable to applications
Standard languages	Ada, FORTRAN, C, Pascal with DEC, IBM, Cray extensions	Application portability and ease of development
High-performance compilers	Common optimizing back-end for all languages; compile speed >60,000 lines per minute	Supercomputer performance for all source languages
Support for standard networking protocols	TCP/IP, NFS, FDDI, Ethernet, DECnet, SNA and others	Connectivity to heterogeneous environments protects user investments
Support for standard graphics protocols and high-speed connections	PHIGS+, X-Windows, HSC, etc.	Use of system as a superior visualization tool
Ability to handle large data sets	Database Management Systems (DBMS) and operating system support for large files and large numbers of files	Better performance and greater insight into task

Table 2. Software Requirements, K-1 Features, and User Benefits

Requirement	Key Feature	User Benefit
Definition of user needs and application expertise	Experienced systems support consultants	Optimal match of problem and Key solutions
Software porting and installation assistance	Scientific Software Group: experienced applications, benchmark, and on-site field support consultants	Smooth transition to Key environment
Performance tuning	Application and industry expertise, experienced field support consultants	Optimal system performance and timely, appropriate system upgrades
Problem "ownership" at all levels from design engineer to field support	"Quality is KEY" program	Rapid resolution of all problems
Software problem management	Product Action Request (PAR) on-line reporting system	Rapid bug resolution
Field support and regular maintenance	Remote diagnostics, several service option levels	Rapid system repair and maximum system uptime
Documentation	Complete, high-quality, and easy to use	First-line problem resolution is easily managed by the end user

Table 3. Support Requirements, Key Features, and User Benefits

2.0 System Description

2.1 Supercomputer Architectural Evolution

Scientific applications contain a great deal of fine-grained or instruction-level parallelism. This means that there are many instructions that can be executed concurrently within basic code blocks or inner loops of scientific application programs. Several architectural approaches have been offered to take advantage of program parallelism. These architectural foundations permit supercomputers to execute many operations simultaneously.

Early supercomputers, such as the Control Data CDC6600 and CDC7600, issued a single instruction per clock cycle, but had a clock rate much faster than the basic instruction times. These systems can issue instructions faster than most basic operations can be executed, exhibiting the basic requirement of a superpipelined architecture. The bottlenecks in early superpipelined machines were the limited number of registers, the issue rate, and the memory latency. As a result, the fast pipelined functional units could not be kept busy on most systems. In order to solve this problem. Seymour Cray added a vector instruction set, a set of vector registers, and a set of backing registers to the basic architecture of the CDC7600, creating the Cray-1 supercomputer. This was a major advance, and allowed the Cray-1, when it was able to use its vector instructions, to keep its functional units operating for a much greater percentage of time than was possible with the earlier machines. However, the Cray-1 faced the same limitations as the CDC7600, and was unable to keep its functional units busy for scalar code.

Another approach to high-performance computing is to build a superscalar machine, or one where the processor can issue more than one instruction per cycle. A current example of this approach is the Multiflow TRACE series of minisupercomputers. which employ a very long instruction word (VLIW) architecture. VLIW machines have been built with slow cycle times, resulting in long latencies for operations, thus limiting speed except for codes with a very large amount of fine-grained instruction-level parallelism. VLIW designs require complex hardware logic. This contributes to the slow cycle times of VLIW systems, resulting in overall slower performance for most applications. Since VLIW machines statically schedule all operations at compile time. the variable access times associated with using cache memories is counter to the VLIW approach. The result on VLIW systems is long access times for memory references. further reducing overall performance on scalar-dominated programs. Code generation for VLIW machines tends to be very difficult, resulting in complex and slow compilers. Slow compilation rates limit the ability of VLIW systems to be used by customers performing large amounts of software development work. The result is that the VLIW machines, like vector systems, will become narrow niche products.

Key's approach is the next step in the evolution of supercomputer architectures. It retains the superpipelined implementation approach with a finely segmented pipeline, but rather than going to a vector instruction set to keep the functional units busy, it issues multiple instructions per clock cycle.

Key's approach solves a number of problems that exist in the Cray machines. It solves the Cray register bottleneck problem by replacing Cray's five register types with 64 uniformly addressed scalar registers. The memory bottleneck is resolved by adding a supercomputer-scale cache memory that allows for rapid access to variables while maintaining a high memory bandwidth. The issue rate bottleneck is resolved by issuing multiple instruction issues in a single cycle. The result is a machine that can achieve near-vector rates without vector instructions, and which is much faster on scalar code than is possible on systems using the Cray vector architecture approach.

2.2 K-1 System Architecture

Key's approach to delivering an extremely fast supercomputer is to build a superpipelined, superscalar implementation of a RISC-like instruction set architecture employing VLSI silicon ECL gate-arrays in an advanced liquid-cooled package. The K-1 central processor has a cycle time of 6 nanoseconds. Two instructions are issued per clock cycle. The system supports an extensive set of high-precision, IEEE-compatible, floating-point instructions as well as a full complement of integer, logical and addressing operations.

The K-1 processor can be divided into four main subsections: the instruction fetch and decode units, the register file, the functional units, and the main memory and I/O systems (see Figure 1). The most central of these is the register file, containing 64 general-purpose registers of 64-bits each. The register file is used to store temporary or permanent data items of all types. Functional units take their inputs from registers or from constants which are part of the instruction. Functional unit results are always stored in registers. Most instructions can specify three independent register addresses. For example, the add instruction adds two registers together and stores the result in a third register.

There are five different types of functional units which process information from the registers: the integer, floating add, floating multiply, floating divide, and load/store units. These functional units are all pipelined, except for divide, and can accept a new operation every clock cycle. The divide functional unit is capable of processing up to four divide operations concurrently.

Key has added several special features to its architecture to support the state-of-the-art optimizations provided by the K-1 compilers. For example, to reduce the effective latency of a memory load instruction, it can be started several cycles before the result is used. In addition, to minimize the effect of the delays normally associated with branches, compilers can move code across branches. The extent to which these optimizations can be done is determined by the existence of features which current architectures generally don't support. Typically, in scalar code, every third or fourth instruction is a branch, which is commonly followed by a memory load instruction. If the legality of the memory reference is dependent on the outcome of the branch (such as a reference through a possibly illegal pointer), then a compiler would not be able to move the load before the branch, resulting in a significant performance reduction. When Key studied this problem, it was found that the addition of an "early load" capability would increase the performance of some loops by more than a factor of three.

Early loads allow load instructions to be moved over branches, even if they generate an illegal memory access. The instruction fault is only generated if an attempt is made to utilize the results of the early load after the branch. This feature is essential to allow compilers to optimize across basic code blocks (the code in between branches) and to effectively use the high issue rate of the K-1 processor.

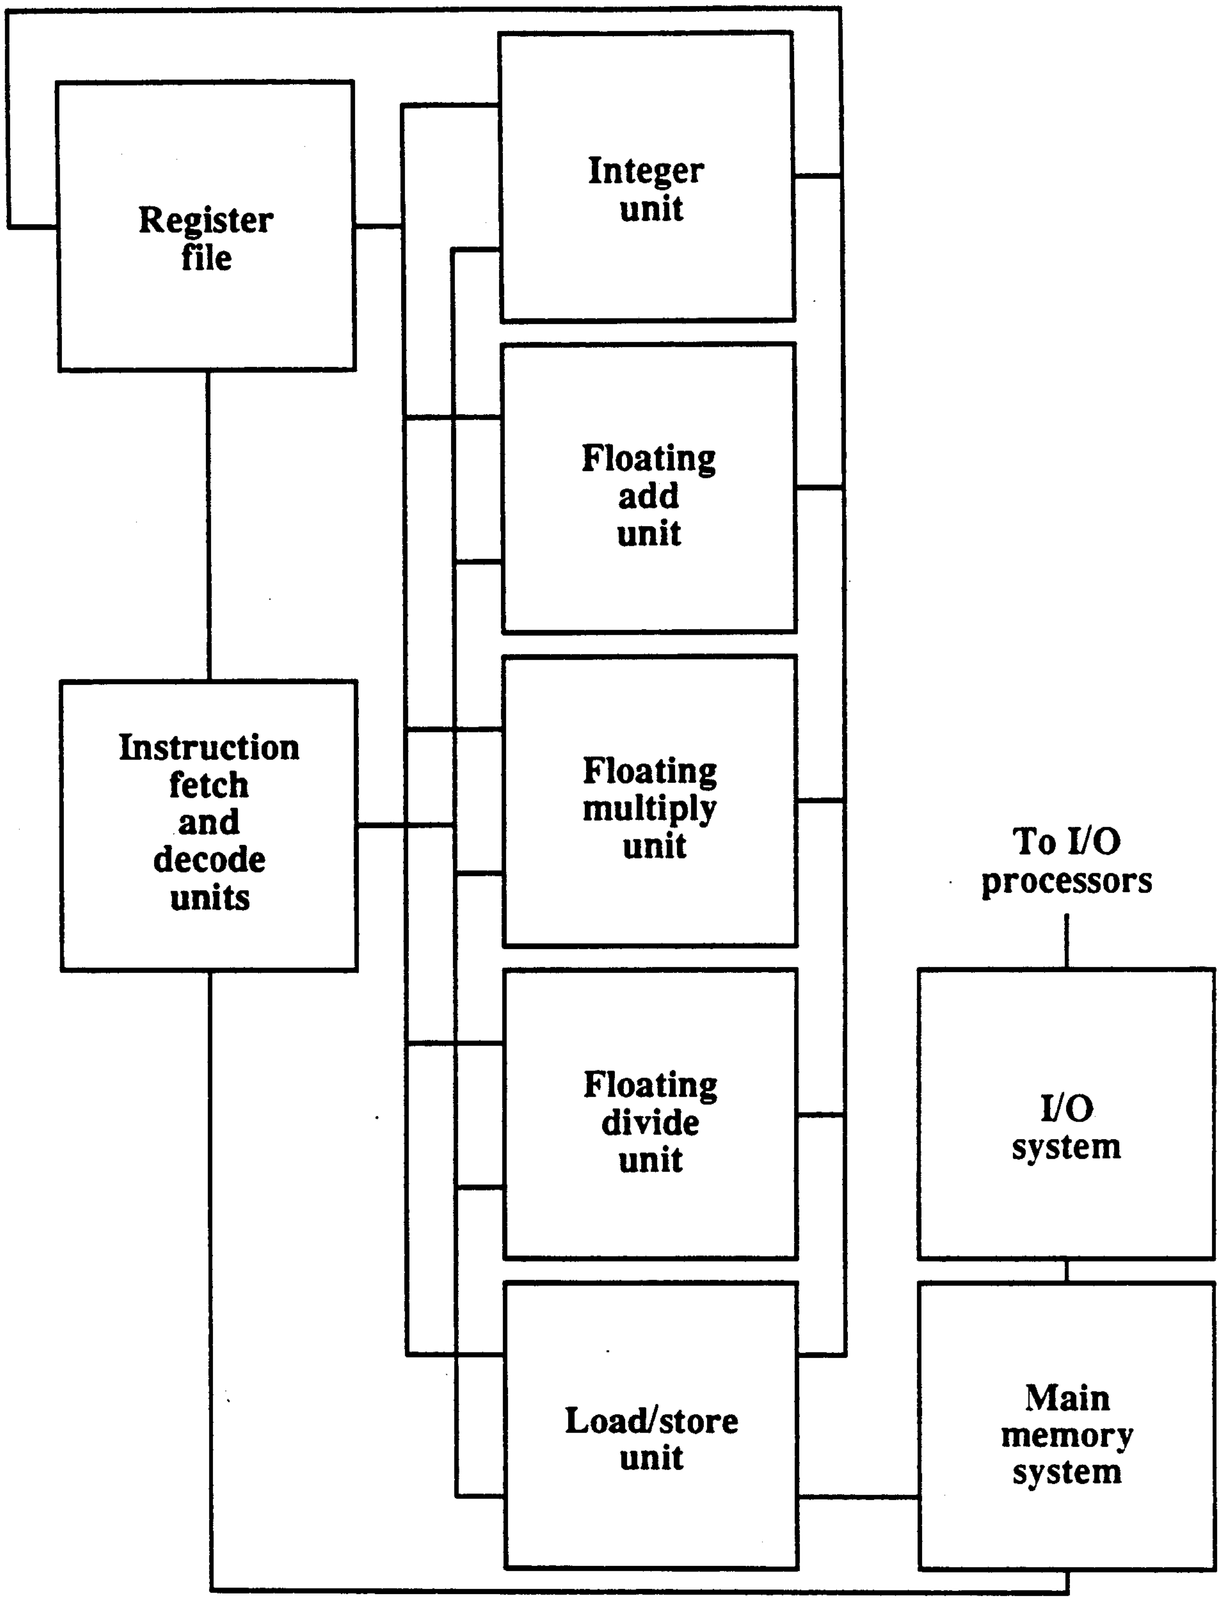

Figure 1. Processor Architecture

Some of the other features which Key has added which are not generally available in current architectures include:

- conditional execution of instructions
- delayed branches (two slots) with conditional execution of delay instructions
- multiple branch flags
- select instruction
- 48-bit memory addressing

The conditional execution of instructions allows branches to be completely eliminated when branching around a small sequence of code. Delayed branches allow the otherwise wasted time after a branch to be used. The instructions in the "delay slots" can be conditionally executed depending on the direction of the branch, maximizing the utilization of the delay slots. Multiple branch flags allow for a number of compare operations to be done in parallel, followed by a sequence of branches. This can greatly speed up the execution of a sequence of test conditions. The select instruction allows for the complete elimination of branches which are used to select between two results. 48-bit memory addressing allows direct referencing of huge physical memories.

With these capabilities, the Key compilers can eliminate many of the bottlenecks that limit current supercomputer scalar performance. New RISC architectures like SPARC (from Sun Microsystems) and MIPS are inappropriate for supercomputing because they lack these features which are so essential for code optimization around branches. The current implementations of these RISC architectures are not highly pipelined and do not issue multiple instructions per cycle, and therefore do not have a great requirement for these features. If, however, RISC implementations are to approach supercomputing levels of performance, they will need to adopt these features in some form. For this reason, and to avoid the limitations of 32-bit addressing, Key has designed a new architecture specifically for scalar supercomputing.

2.3 Memory System

The demand to model larger and larger problems is growing dramatically. Very large memories are necessary to keep the central processors busy and to ensure that optimal system performance does not depend on the latency of disk media for virtual memory management. The K-1 architecture is designed to support the largest engineering and scientific computing needs. It is designed to support a huge, linearly addressed, 48-bit (256 trillion byte) virtual address space, allowing applications to utilize this memory capability easily. The K-1 implementation also supports a huge physical memory, which is shared among all of the processors. The initial implementation is capable of supporting up to 8 Gbytes using 1 Mbit (millions of bits) DRAMS, and up to 32 Gbytes using 4 Mbit DRAMS. Even larger memories will be supported as larger DRAMS become available. The memory controller supports DRAMs and SRAMs of different speeds so that customers can protect their investment in existing hardware as well as upgrade to the highest capacities available. The main memory bandwidth of 3.5 Gbytes/second is achieved through the use of a very wide, 64-byte (512-bit) wide transfer path from the memory system to each processor.

To insulate the system from the access time of main memory, there are two large caches totalling 3 Mbytes for each processor. One cache is 1 Mbyte in size and is used exclusively to hold instructions. The data cache is 2 Mbytes in size. The caches operate transparently with complete cache coherence between processors in the multiprocessor system. The architecture provides instructions for manipulating the caches and for explicitly updating memory. When a processor gets a cache miss, data is transferred in 256 byte lines between the processor and memory. This data is transferred over the main memory bus at 3.5 Gbytes/second. Taking into account the memory access time and overhead associated with processing a cache miss, a individual processor can achieve an effective transfer rate of about 800 Mbytes/second between its cache and main memory. The higher memory rate of the memory bus can be utilized when a number of processors are cache missing at the same time.

2.4 I/O Architecture

The K-1 I/O architecture is designed to provide the very large I/O system bandwidths required for meeting the high system throughput demands of supercomputer applications. Systems with one I/O controller (IOC) provide a peak I/O bandwidth of over 500 Mbytes/second, and systems with two IOCs provide a peak of over 1 Gbytes/second.

The I/O system is designed in a modular fashion, so that it can be configured to support the needs of large and small systems. Multiple I/O channels can support a variety of peripheral devices. These channels can be standard High Speed Channels (HSC) directly connected to a VME bus-based I/O processor (IOP), fiber optic connections with multiple devices, or standard peripheral devices and network connections. Each IOP can support a disk drive subsystem with more than 20 Mbytes/second bandwidth, or a range of lower-speed peripherals such as tape drives and serial data lines. Systems can be configured with from 2 to 64 IOPs, providing a deliverable I/O capability of over 1 Gbyte/second. To maintain a high degree of reliability, it is possible to configure the I/O system with redundant paths between the memory and the I/O processors.

An example of a small system configuration is shown in Figure 2, and of a large system configuration in Figure 3. The configurations available at first customer shipment will support the 100 Mbyte/second copper HSC channel. The fiber optic channel and the 200 Mbyte/second HSC channel will be added at a later date. High-speed peripheral devices, such as graphics frame buffers, can be connected directly to a channel on an I/O Channel Adapter (IOCA) card.

2.5 Packaging Technology

Key's overall system is air-cooled, so the customer need not provide any special cooling and power equipment other than normal computer room air conditioning and 220 volt power. No motor generators or special chillers are required, as they are for current supercomputers like Cray machines. Because of the high power dissipation chips used in the processor, liquid cooling is used within the processor. The processor has an integral heat exchanger which transfers heat to the room air. The I/O system and the memory system are air-cooled. For large configurations, an optional chilled water hookup will reduce the amount of room air-conditioning required.

Heavy use of high-pin-count chips and surface mount technology is made throughout the system. The processor uses a multi-layer (30-40 layer) PC board technology.

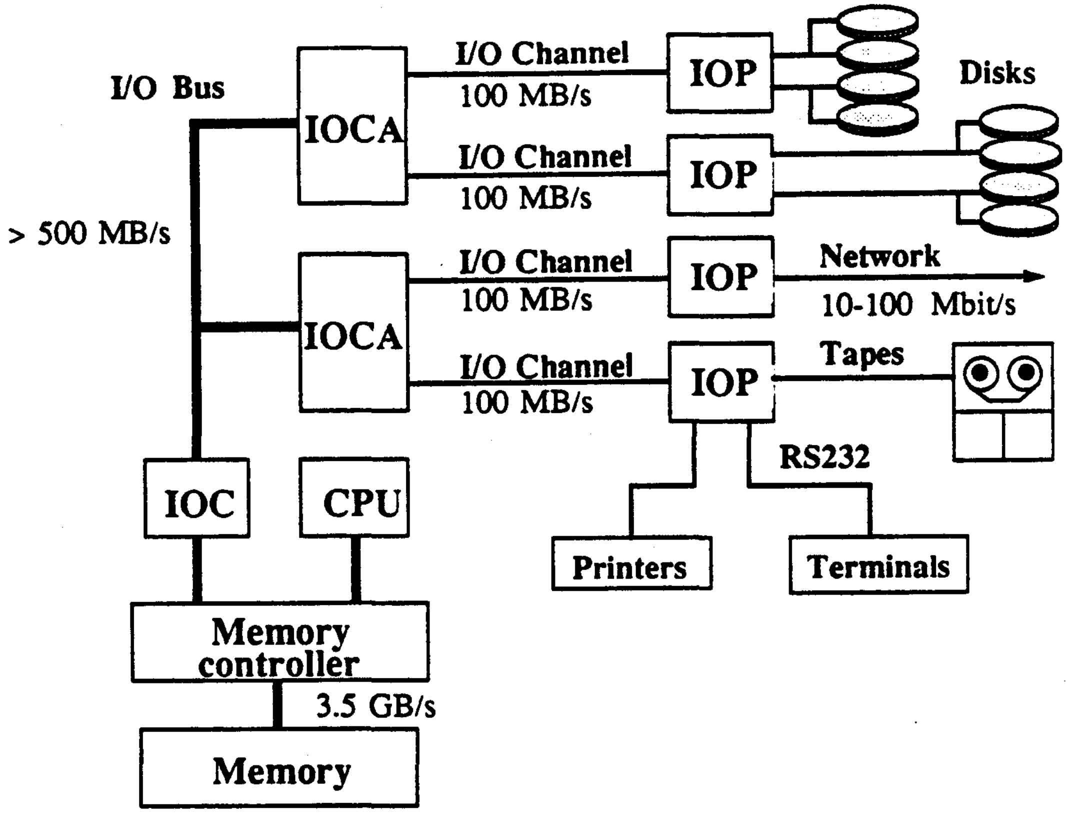

Figure 2. Sample Small I/O System Configuration

Key's packaging complies with the major European (VDE) and United States (UL) regulatory agency requirements to facilitate an early introduction of Key's products to the worldwide marketplace.

2.6 System Reliability

The K-1 system in composed almost entirely of VLSI ECL components, resulting in a very low component count as compared to current supercomputers. The use of advanced gate-array technology permits Key to design chips with 100% scan capability, allowing production of highly testable and reliable CPU, memory, and I/O subsystems. System reliability will be enhanced through a program of quality assurance involving all personnel from design engineers through manufacturing to field maintenance technicians.

2.7 Operating System Software

Key's operating system (KEYNIX) provides a high-performance, standard environment for applications, along with connectivity to engineering workstations and other computers. To meet these requirements, Key has adopted the most important standards and enhanced the internal structure of the operating system to provide high performance.

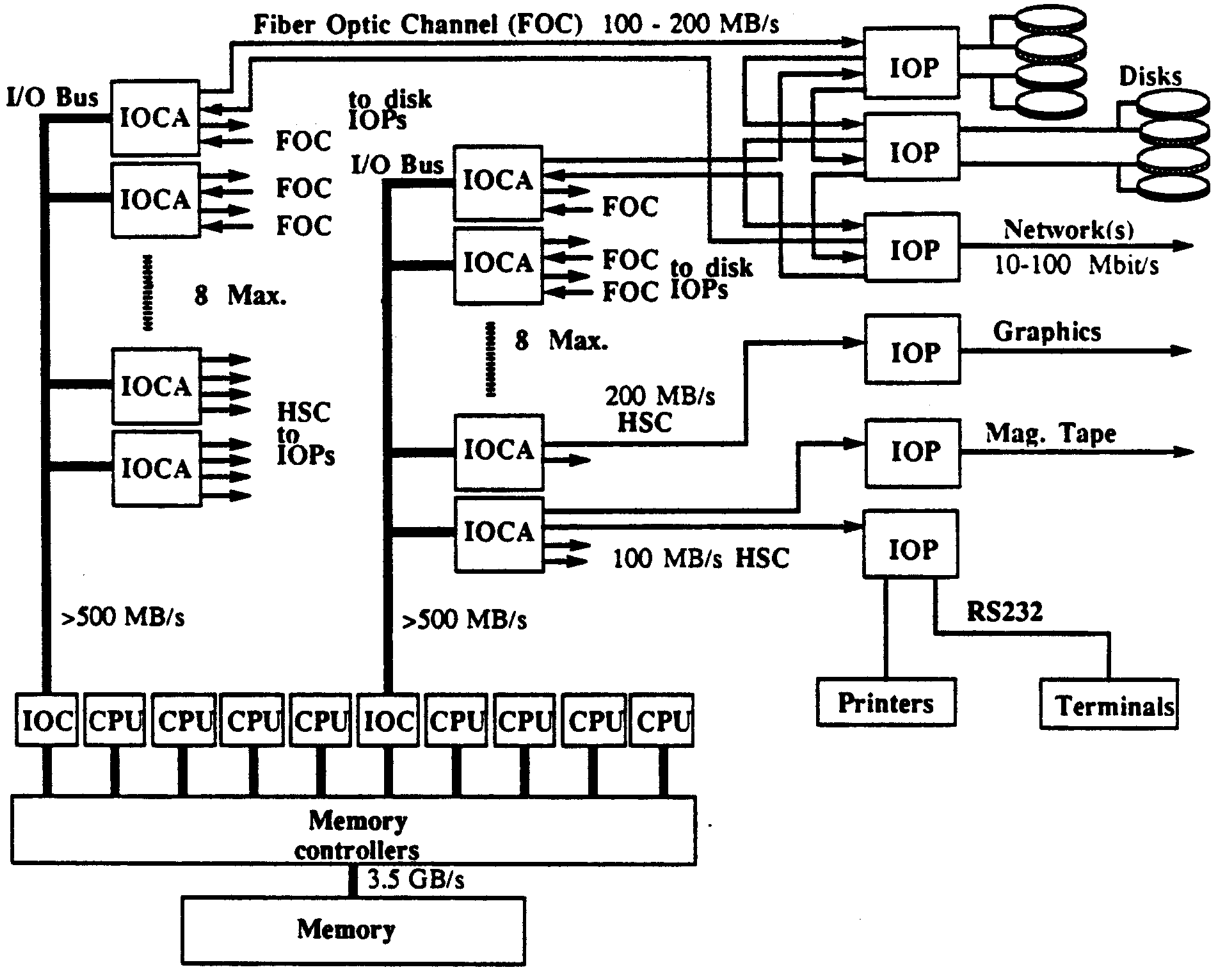

Figure 3. Sample Large I/O System Configuration

Standards are important to technical computer systems users as they ensure software portability across heterogeneous system environments. The KEYNIX operating system is based on AT&T's UNIX System V Release 3.2 (SVR3.2) operating system and passes the System V Verification Suite (SVVS). System V is the most widely accepted UNIX standard in the marketplace today. Nearly every major computer vendor has endorsed a System V-based standard. KEYNIX also provides compatibility with several unofficial and emerging standards such as POSIX by building on top of the System V base.

Supercomputing extensions to UNIX are being developed by standards-setting groups. Key implements many of these facilities for high performance, including asynchronous I/O, suspend/resume, checkpoint/restart, multi-tasking, a batch facility, and job accounting. Additional extensions include the adaptation of the operating system to run on a symmetric multiprocessor, the development of a high-performance extent-based file system, support for very large files that span multiple physical drives, and support for ganged disks. All of these extensions to UNIX greatly enhance deliverable system performance.

2.8 Programming Language Support

Key offers the FORTRAN, C, and Pascal language processors. Other languages, including Ada and LISP, will be developed based on customer requirements. Each supported language is based on its corresponding international or industry standard and contains extensions to ensure compatibility with industry leaders in the mainframe, supercomputer, and minicomputer areas (IBM, Cray, and DEC) and fulfills MIL-STD (Department of Defense) requirements. Extensions also include changes emerging from upcoming standard revisions (e.g., FORTRAN 8X). By maintaining compatibility with language standards, Key ensures that application programs can be transported to the K-1 with minimal effort.

Key complements the performance of the K-1 processor with the best optimizing compilers in the industry. The highest possible run-time performance for all languages is achieved by using a common optimizing back-end. Full optimization is achieved without user intervention in order to reinforce the K-1 as a dependable, no-surprise, balanced system on which software developers need no detailed understanding of the system architecture.

Compilation-time performance of the Key language processors is over 50,000 lines per minute for unoptimized object code and at least 15,000 lines per minute for optimized CPU target programs. This high compilation rate and the highly optimizing compilers positions the K-1 as an excellent software development machine as well as a target system for application execution.

2.9 Networking and Graphics

Several network media are supported by KEYNIX, including Ethernet and FDDI. Special devices such as frame buffers are connected at very high speed to an HSC channel directly to the Key I/O subsystem. Requirements for connection to additional networks will be determined at a later date. DECnet and SNA support are currently under investigation.

Network interprocess communication services are provided by the BSD sockets interface and TCP/IP which have become unofficial standards for network communication among UNIX systems over Ethernet networks. Through these mechanisms, KEYNIX provides remote login and execution services as well as peer-to-peer remote communication. File sharing over a local area network is provided by Key's implementation of Sun Microsystems' Network File System (NFS) and Yellow Pages, which have become de facto standards. The K-1 can share files with a large number of engineering workstations and other computer systems available from many different vendors.

Key fully supports the emerging standards for high-speed fiber optic networking embodied in the FDDI standard specification. Key also supports very high-speed device connections via the HSC interface and FDDI specifications.

Several higher-speed local area networks have been developed by third-party vendors (for example, VectorNet from Scientific Computer Systems and UltraNet from the Ultra Corporation). As of this writing, it is not clear which of these products, if any, will emerge as a "network of choice" for supercomputing. Key will work closely with these and other vendors to ensure the best fit of Key's products with high-speed networks.

In conjunction with high-performance workstation vendors, Key provides the ability to visualize computed results quickly and easily. It is neither necessary nor desirable for Key to supply graphics rendering displays directly. Instead, Key provide a complete supercomputing/visualization solution in cooperation with leading workstation suppliers. This goal is accomplished by using a rendering engine (such as Silicon Graphics') connected to the K-1 via a high-speed fiber optic or HSC channel link. To facilitate the accessibility of high-performance graphics on the K-1 family, Key fully supports PHIGS+, X-Windows version 11, FDDI, and other evolving graphics and networking standards.

2.10 Coexistence with Other Systems

Many K-1 prospects already have made substantial investments in workstations as well as DEC and IBM systems. Both DEC and IBM have made commitments to networking and system software standards. Adherence to standards ensures that Key systems will remain compatible and can easily coexist with these computer systems and that a Key system will be a viable, low-risk augmentation to existing computer sites. All of this protects end-user investment in heterogeneous hardware and application software.

Continued investigation will determine the extent to which DEC compatibility, in the form of DECnet, common utilities, and a VMS operating system shell is required. Key will not undertake to develop these software products, as they are available from third-party vendors.

2.11 Database Management System (DBMS)

Key will offer a database management system to address the unmet need for organizing large amounts of data in a supercomputing environment. Key can attain a position of leadership in database management for two reasons: the K-1 architecture is well-suited to the high I/O and reliability requirements of database-oriented applications, and no current supercomputer vendors provide comprehensive DBMS support. Key will not undertake to develop DBMS software internally, as popular databases are currently available from high-quality suppliers. Key will establish relationships with established vendors rather than start-up firms engaged in supercomputer database development. Support of database packages such as Oracle, Ingres, Unify, Informix, or Sybase would meet this need. Key will also examine the applicability and feasibility of implementing an object-oriented database and FORTRAN bindings for SQL support on the K-1 supercomputer.

2.12 Application Software

Proprietary code developed in-house at customer sites migrate easily to the K-1 because of the familiar programming environment and simplicity of the system architecture. These factors help third-party developers to transport their software to the Key system.

Key's third-party software strategy is twofold: to allow a new class of applications to be developed and used, and to attract the transportation of existing applications in Key's target segments. Key is establishing joint marketing partnerships to identify appropriate hardware/software solutions. In emerging markets like computational chemistry, Key's new class of supercomputers encourage the development of new algorithms and applications.

Through the **Key Partners Application Alliance**, Key supports research and commercial software development groups in order to facilitate the porting and development of application tools. Key has established a **Scientific Software Group** to conduct customer benchmarking and to support application software porting to the K-1. This group is essential in educating K-1 users on porting and optimizing user code. Extensive use will be made of simulators to estimate application performance and to anticipate and resolve problems before K-1 hardware is available. This will result in important third-party codes being ported to the K-1 at or soon after first customer shipments. The group also provides rapid turnaround of customer benchmarks to support the selling process. For initial shipments, which will be less dependent on third-party codes, the Key Scientific Software Group will work directly with customers to help port their in-house codes.

2.13 Support and Services

Recognizing that customer satisfaction relies heavily on a close and continued relationship with prospects and customers, Key offers unparalleled pre- and post-sales support services. Field support covers application and system software as well as hardware. Factory personnel provide extensive backup support, especially through the Scientific Software Group.

The K-1 supercomputer family is designed to allow for efficient and rapid servicing in the field. Each system is supplied with a maintenance processor which can run a comprehensive set of test vectors on the system using the scan paths provided throughout the CPU, memory, and I/O system. This extensive diagnostic capability is revolutionary for high-end supercomputers and is made possible by the use of very high-density chips. The maintenance processor can be accessed remotely by Key factory personnel through telephone/modem connection, so that the exact cause of system problems can be determined and field support representatives will be dispatched with the proper solution in hand. Field replacement is at the board level. Spare boards are kept in field service offices located near major customer centers.

Key will offer two levels of system support services starting in 1992: Basic and Priority service. Both services will be provided for all Key hardware and software products, and will act as the first point of contact for problems with third-party peripherals and applications. Basic service will include: 8AM to 5PM (local time) support coverage from the nearest service center, scheduled preventative maintenance. installation of field change orders, parts and labor, access to the Key Product Action Request (PAR) toll-free "hot line" problem reporting procedure. Priority service will include all of the features of Basic service, plus 24-hour, 7 days a week coverage for remedial service and resident, on-site hardware and software personnel for operational. performance, and application expertise to ensure optimal system usage. It is anticipated that most customers will select Priority service. Key will tailor additional service options to meet the specific requirements of individual customers.

2.14 System Configurations

The K-1 system design permits great flexibility in configuring systems with from one to eight processors to meet specific customer requirements. Although Key focuses its efforts on providing corporate resource supercomputer solutions that offer high performance levels, massive memory capacity, and high I/O bandwidth, a variety of configurations are possible, from expandable entry-level systems to multi-GFLOPS (billions of floating-point operations per second) systems. Sample configurations for entry level, standard, and large K-1 systems are outlined in Table 4.

The K-1 uniprocessor system (K-1/1) delivers a 50% price/performance advantage over anticipated minisupercomputers in 1991. Of course, such an entry level system is just a starting point for eightfold growth within the K-1 family, while minisupercomputer products are at or near the limits of their growth potential. For this reason, larger K-1 configurations cannot be meaningfully compared with minisupercomputers because minisupercomputers will be unable to achieve such high performance rates. Comparisons between the K-1 family and future superminicomputers (such as DEC VAX) and mainframes (IBM) strongly favor Key. This is because such systems cannot be reasonably configured anywhere near Key's performance range, except as clusters or as networks of machines. Key still enjoys a four- to ten- times price performance advantage over superminis and mainframes. Finally, mid- and high-end K-1 systems will deliver a full four times the price/performance of supercomputers such as the Cray Y-MP and Cray-3.

	Entry-Level K-1/1	Medium System K-1/4	Large System K-1/8
CPUs	1	4	8
Peak MFLOPS	333	1,332	2,667
LFK MFLOPS	54	216	432
Memory	0.5 Gbytes	4.0 Gbytes	8.0 Gbytes
Options			
Disks	8 Gbytes	96 Gbytes	336 Gbytes
Disk I/O Rate	24 Mbytes/second	288 Mbytes/second	1 Gbyte/second

Networks, graphics, tape drives, other peripherals, software, etc., are available as additional options.

Table 4. K-1 System Configurations

2.15 System Performance

One of the most respected and widely-quoted benchmarks for supercomputers is the Livermore FORTRAN Kernels (LFK). These kernels, or program loops, are 24 pieces of code extracted from compute-intensive and largely vectorizable scientific applications that are executed at the Lawrence Livermore National Laboratories

(LLNL). The kernels are embedded in a benchmark driver which runs them several times on different data sets, checks for result and timing accuracy, reports execution rates, and summarizes the results with several statistics.

The LFK benchmark suite can provide insights into the sensitivity of the system being tested to various levels of program vectorization. For example, the minimum and maximum rates define the performance range of the system, while the arithmetic mean corresponds to programs characterized by 90% and higher vector content.

An ideal metric to use for comparing real-world application performance across systems is a weighted harmonic mean of the Livermore kernels. Weighting would be dependent upon the percentage of the application that is approximated by each of the LFK programs. The LFK harmonic mean weights each kernel equally. As such, it measures the average rate at which floating-point operations are executed. For this reason, and because most technical applications are not characterized by a high vector content (as recently verified by Dataquest), Key believes that the LFK harmonic mean is a good measurement of real-world, floating-point performance across a range of applications. The performance of the K-1/1 entry system on these kernels, as well as values for current Cray Y-MP and Convex C-210 systems are listed in Table 5. At 54 MFLOPS, the K-1 harmonic mean is 2.3 times that of the Cray Y-MP. Since the K-1 and the Y-MP have the same cycle time, this means Key's architectural innovations provide a great efficiency improvement over Cray's architecture. The K-1 provides far more balanced performance than the Y-MP. The K-1's slowest kernel runs at 5 times the speed of the slowest kernel of the Cray Y-MP.

The K-1 delivers uniformly high performance across a variety of applications that range from low to high levels of vectorization. This is in distinct contrast to Convex, Cray, and other vector systems that deliver their maximally-obtainable performance only on code that has a high vector content. From this, the advantages of the K-1 system for all but the most highly vectorized problems are clear.

Vector processing performance as measured by the 100 by 100 full precision all-FORTRAN LINPACK benchmark is rated at 76 MFLOPS, comparable to that of the Cray Y-MP at 74 and much higher than the Convex C-210 at 10. Clearly, the Key system achieves superior scalar performance and high levels of vector capability.

The combination of high per-processor performance and system expandability yields a supercomputer solution that delivers balanced performance across many applications, for many users, and which can grow in capability with the needs of the installation.

LFK Kernel Benchmark	Key K-1/1†	Y-MP/1‡	Convex C-210
1	188	233	39
2	67	64	7
3	150	215	22
4	112	93	12
5	26	20	2
6	75	17	4
7	166	266	44
8	116	207	26
9	153	223	44
10	53	87	10
11	33	18	3
12	75	105	24
13	44	8	2
14	48	31	4
15	79	7	5
16	25	9	2
17	32	14	3
18	163	157	19
19	28	19	3
20	25	19	2
21	100	154	29
22	47	93	12
23	59	20	4
24	56	5	13
Harmonic Mean	**54**	**23**	**5**

† = estimated; ‡ = scaled up from X/MP

Source: Key Estimates, LLNL

Table 5. K-1 Uniprocessor Performance Comparisons