

SPACECRAFT ON-BOARD INFORMATION EXTRACTION COMPUTER (SOBIEC) P. 10

David Eisenman - Deputy Section  
 Manager, Flight Command &  
 Data Management Section  
 NASA Jet Propulsion Laboratory  
 California Institute of Technology  
 4800 Oak Grove Drive  
 Pasadena, CA 91109

Robert E. DeCaro  
 Chief Electronic Systems Engineer  
 Irvine Sensors Corporation  
 3001 Redhill Avenue  
 Costa Mesa, CA 92626

David W. Jurasek  
 Director of Advanced Systems  
 Engineering  
 nCUBE  
 1825 N.W. 167th Place  
 Beaverton, Oregon 97006

**ABSTRACT**

The Jet Propulsion Laboratory is the Technical Monitor on an SBIR Program issued for Irvine Sensors Corporation to develop a highly compact, dual use massively parallel processing node known as SOBIEC. SOBIEC couples 3D memory stacking technology with state of the art parallel processor technology provided by nCUBE. The node contains sufficient network Input/Output to implement up to an order-13 binary hyper-cube. The benefit of this network, is that it scales linearly as more processors are added, and it is a superset of other commonly used interconnect topologies such as: meshes, rings, toroids, and trees. In this manner, a distributed processing network can be easily devised and supported. The SOBIEC node has sufficient memory for most multi-computer applications, and also supports external memory expansion and DMA interfaces. The SOBIEC node is supported by a mature set of software development tools from nCUBE. The nCUBE operating system (OS) provides configuration and operational support for up to 8000 SOBIEC processors in an order-13 binary hypercube or any subset or partition(s) thereof. The OS is UNIX (USL SVR4) compatible, with C, C++, and FORTRAN compilers readily available. A stand-alone development system is also available to support SOBIEC test and integration.

**MISSION REQUIREMENTS**

The general problem of finding optimal techniques for the extraction of scientific information from a wide band data stream has been discussed in depth in a JPL publication by Robert Rice<sup>1</sup>. There, the observation was made, and to a degree quantified, that perhaps the most powerful technique for error-free information extraction is to employ activity and pattern recognition to cue the allocation of digitization and communications resources. In discussions with JPL personnel regarding this technique, the example was given of a Mars explorer spacecraft

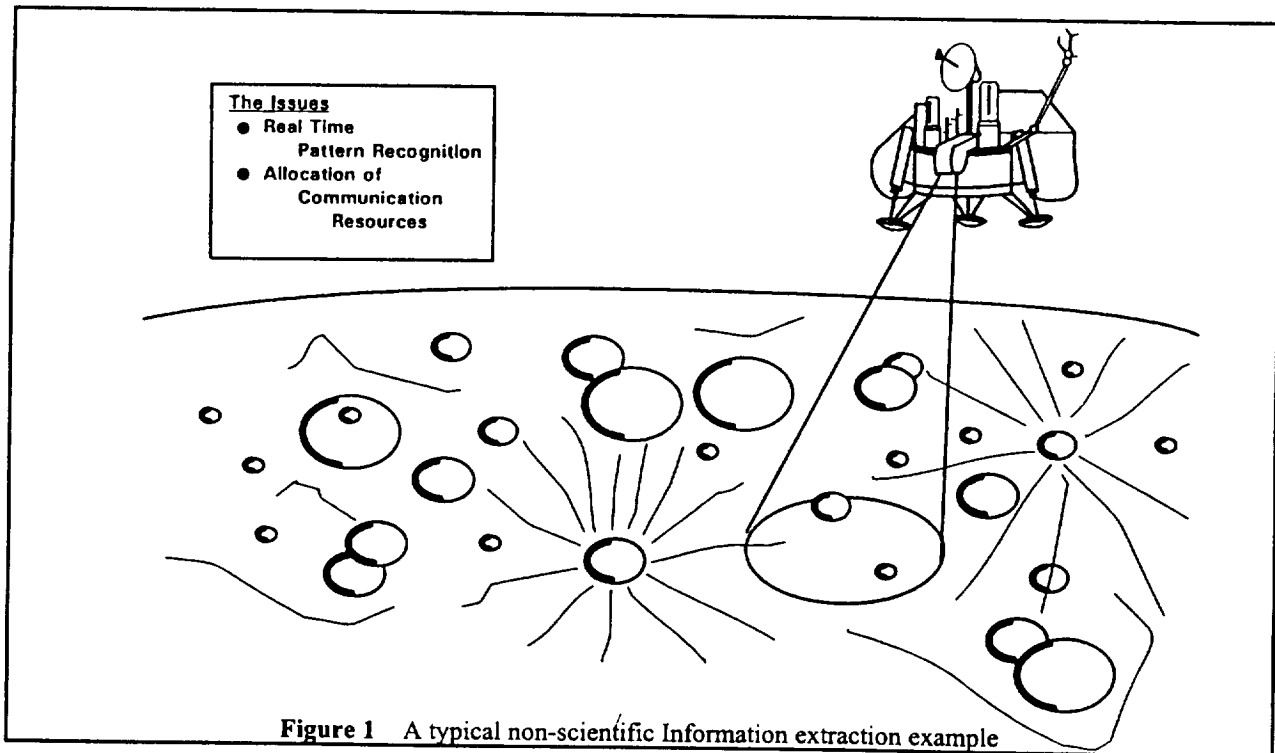
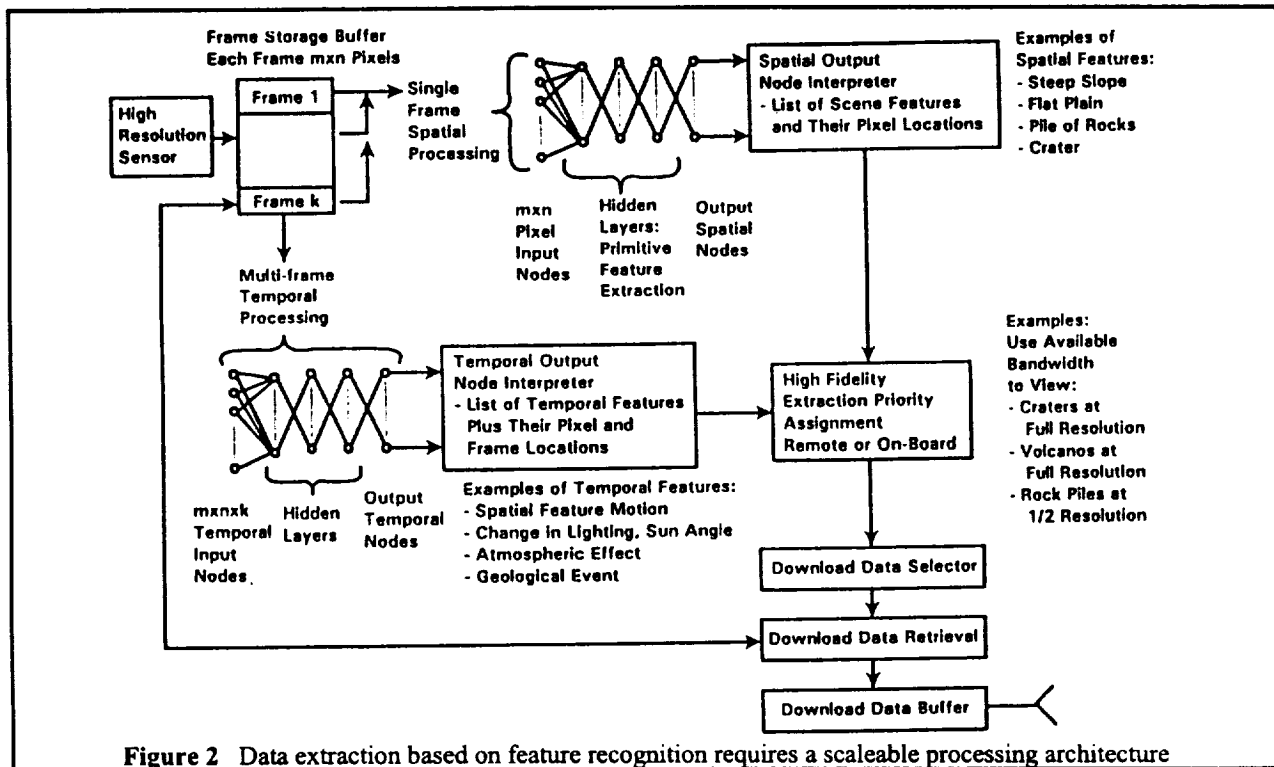


Figure 1 A typical non-scientific Information extraction example

looking for a landing site (Figure 1). In the Mars lander case, it is desirable to avoid areas of high activity and spatial complexity. It is necessary to examine apparent clear areas in very great detail (high spatial and high amplitude resolution) to assure that these areas are indeed clear and flat. In this example, sophisticated feature extraction and pattern recognition capability is important. Comparison of this example to the more obvious one of high fidelity scientific data communication in the face of a limited datalink provides evidence of the generality of the information extraction problem. A general solution to this problem is the Spacecraft on-Board Information Extraction Computer (SOBIEC), a massively parallel, highly interconnected processing system. This effort is funded by a Small Business Innovation & Research (SBIR) contract, monitored by the Jet Propulsion Laboratory (JPL).

Figure 2 shows a concept diagram for using a high-density, parallel processing computer with a large amount of distributed memory to perform feature extraction, which leads to a prioritized downlink of important features at high resolution and optimizes the limited bandwidth communication channel.



This implementation of feature extraction uses parallel processors to emulate neural circuitry performing hierarchical pattern recognition. For any given mission, both hardware utilization (number of nodes and interconnections between nodes) and software (weight definition specific to features require definition. These definitions are left open to make the implementation generic. Nevertheless, it illustrates a potentially powerful method for feature recognition and image extraction, which is also well suited for hardware implementation on a distributed memory parallel processing computer. The SOBIEC massively parallel processing node, developed concurrently between Irvine Sensors Corporation (ISC), nCUBE, and NASA JPL, a highly compact building block, enables compute intensive missions where the processing must be scaled to the application such as the example given, micro-spacecraft and micro-rovers.

### SOBIEC ARCHITECTURE

SOBIEC's electrical architecture, typical of massively parallel processors, is shown in Figure 3. SOBIEC's processor (developed by nCUBE) contains a dynamic RAM controller with 7 bit error detection and correction (EDAC) and fourteen serial communications links. Ten years ago, nCUBE pioneered the field of massively-parallel computing where hundreds or thousands of processors are used to solve large, complex computing problems.

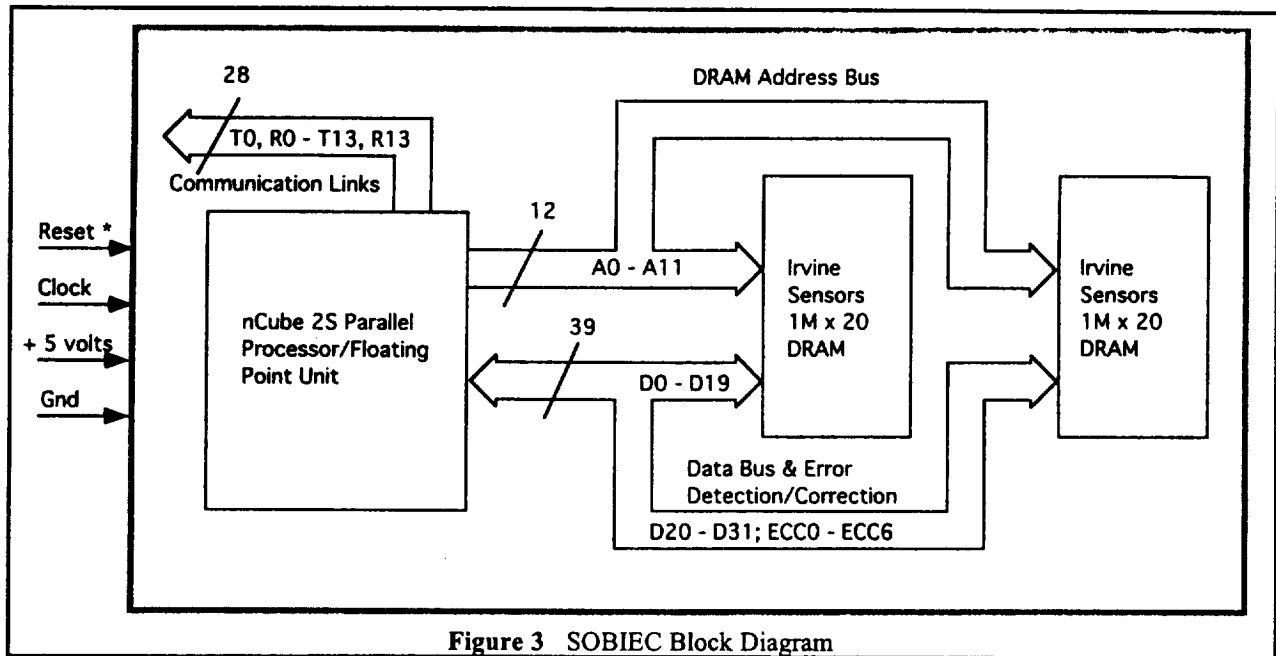


Figure 3 SOBIEC Block Diagram

An important aspect of parallel computers is their ability to continue operating, even when one or more processing elements has failed. Although the failure rate of SOBIEC's highly-integrated processor is extremely low, this "graceful degradation" in computing is vital to mission-critical applications. The relatively low cost of the SOBIEC processor node allows redundancy to be used in a cost-effective manner when needed, while still allowing the flexibility to re-allocate resources to handle peak loads.

### *SIMD vs. MIMD PARALLEL COMPUTERS*

Over the years, many different parallel computer architectures have been proposed and built. Generally, these machines fall into two main categories: Single Instruction, Multiple Data (SIMD) and Multiple Instruction, Multiple Data (MIMD).

SIMD machines are characterized by having one set of instructions (the program) and multiple data sets. In these types of machines, the data is spread over the array of processors, and then all processors execute the same instructions in lockstep. Although this is extremely useful for certain types of applications where the same manipulation of data is applied across a large data set, it has a number of disadvantages for a SOBIEC application. First, it is virtually impossible to have multiple tasks (different applications) on these machines. Secondly, programs on SIMD machines can be difficult to debug because it is difficult to view the data on selected processors as the program is running. Finally, these machines lack flexibility. SOBIEC programs could not "interact" with the data on a processor-by-processor basis, modifying the execution of the program based on the data in an individual processor. These disadvantages proved to be a serious limitation for SOBIEC's mission, and so MIMD processors were investigated.

MIMD machines are characterized by each processor having its own individual instruction set (program) and data set. This distribution of both instructions and data gives these machines a large amount of flexibility. It is very straightforward to divide these machines among multiple tasks, each task taking only as much of the system compute power as is necessary to complete the task in a time-effective manner. Debugging on these machines is aided greatly by the fact that a debug program can be loaded on one or more selected processors, giving the programmer the ability to view the state and data on individual processors or set of processors in the machine. Finally, since separate programs are running on each processor node, these programs need only perform the operations appropriate to the data stored locally. This "extra" processing power can then be effectively used to perform other "background" tasks, increasing the overall flexibility and cost-effectiveness of the machine.

Another distinguishing characteristic of SOBIEC's processing computer is that it is based upon a distributed vs. shared memory architecture. In the shared memory system, all processors have access to common

memory pool, in which instructions and data are stored. Although the shared memory model has the advantage of being familiar to virtually all programmers, it suffers from the serious flaw of being difficult to scale. The limited data bandwidth of the busses that connect all processors to common memory quickly becomes the bottleneck of the system, preventing additional processors from providing the expected increase (linear scaling) in performance. Distributed memory systems (such as SOBIEC) give each processor its own local memory, which gets shared with other processors and the outside world via messages over a communications network. Since these are "private memory arrays", the total memory bandwidth increases linearly as more processors are added. Thus, memory bandwidth is not a limiting factor in the scalability for SOBIEC systems.

SOBIEC's (nCUBE's) communications network is known as a binary hyper-cube. In a binary hyper-cube system, all processors are assigned a binary identification word (the Processor ID). Processors which differ by only one bit in their ID are interconnected with a synchronous, duplex Direct Memory Access (DMA) channel which runs at 2.75 megabytes per second, each direction. Thus, the number of DMA channels on any given processor is the log (base 2) of the maximum number of processors in the system. SOBIEC's processor has 13 DMA ports for array interconnect, allowing up to 8192 processors to be fully interconnected. An additional (fourteenth) DMA port on each processor is used to connect to the outside world via the I/O subsystem, which would consist of additional SOBIEC processors running I/O driver code. The benefit of this network is that it scales linearly as more processors are added, and it is a superset of other commonly-used interconnect topologies such as: meshes, rings, toroids, and trees. This feature makes the SOBIEC processor node truly universal for NASA applications.

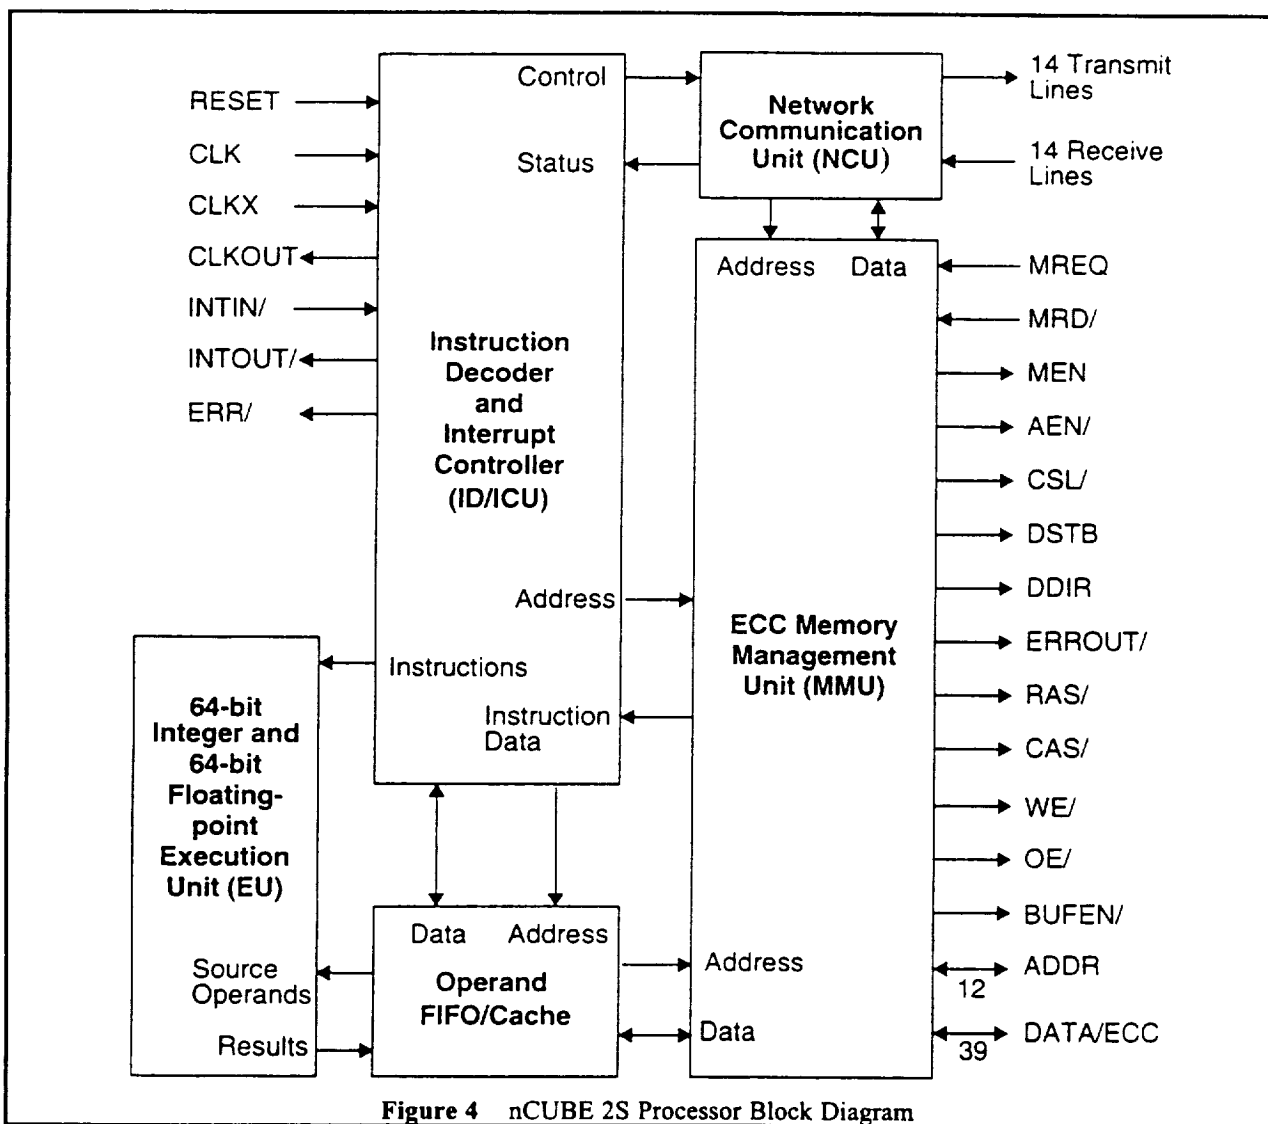


Figure 4 nCUBE 2S Processor Block Diagram

## ***PROCESSOR SELECTION TRADES***

The issue as to whether to use custom, proprietary processor chips, or off-the-shelf commercial processor chips for the SOBIEC program was explored. The advantages of off-the-shelf processors are fairly obvious. These parts are inexpensive, reliable, readily-available, and typically have wide familiarity in the marketplace. But we also found that they also have one large, fairly non-obvious disadvantage: in order to be commercially successful, the manufacturer has had to make these parts quite general purpose and able to fit into a wide variety of systems. This means that for any given system, a fair amount of "glue logic" to make the processor work is required. All that "glue" would make SOBIEC physically larger, slower, less reliable, more power hungry, etc.

nCUBE, in designing the processor used for SOBIEC, chose a more difficult path by designing its own full-custom processor. By designing a minimum parts count "world-class processor", all the disadvantages associated with commercial processors were avoided. The nCUBE MIMD processing element consists of the processor chip and DRAM (Dynamic RAM) chips... nothing else. These elements are physically small, reliable, low power and easy to build and test, making them an ideal choice for SOBIEC. Our experience has also shown that the extra money spent on a relatively low volume custom processor is more than made up for by cost savings due to the lack of glue logic, and simplified build/test of MCMs.

## ***SOBIEC PROCESSOR FEATURES***

Figure 4 is a block diagram of nCUBE's 2S processor used on SOBIEC. The 2S processor consists of five major blocks: Instruction Decoder/Interrupt Controller (ID/ICU), 64-bit Integer/Floating Point Execution Unit (EU), Operand First-In, First-Out (FIFO) Cache, Network Communication Unit (NCU), and ECC Memory Management Unit (MMU). This processor features a 64-bit data path for on-chip integer, floating point processors and dynamic Memory Management unit. The nCUBE 2S Processor is capable of addressing from 1 to 64 Megabytes of memory per node. In addition, the processor supports page mode accesses to increase bandwidth. The MMU directly controls the external DRAM chips with an external 32-bit data bus and 7-bit ECC (Error Check and Correct) bus. This ECC provides SOBIEC the assurance that all single-bit errors will be detected and corrected on-the-fly, while double bit errors will be detected and flagged as such. The extra security provided by this ECC ensures that SOBIEC's space applications will provide correct results, despite the occurrence of such well documented random events as alpha-particle hits.

The processor contains protection logic to protect system software and allow multiple processes per node. A User/Supervisor bit in the Program Status Word (PSW) causes protection faults on certain instruction and restricts access to memory other than the memory assigned to a process.

nCUBE's 2S processor contains an elaborate interprocessor communications network. Fourteen DMA channels, tightly coupled to the MMU enable SOBIEC to share the mission signal processing between on-board processing elements. An additional benefit of this communication network, is in developing a fault tolerant architecture. As stated previously, in the event of a device failure, a properly designed multiprocessor system can be designed to re-allocate resources to achieve a graceful degradation - key to any successful space-based application.

The 2S processor's high integration also results in significant power savings - about 2.5 watts at 20 MHz. Combined with four Megabytes of memory, SOBIEC's total power dissipation is only 6 watts per node. In addition, each node is capable of 3.2 MFLOPs (Million Floating-Point Operations per second), 12 MIPs (Million Instructions per second), and 80 megabytes per second of memory bandwidth.

## ***3D STACKED MEMORIES***

The SOBIEC MPP node achieves its small size by utilizing a pair of 20 megabit DRAM "short stacks" configured as a 1 megabyte x 20 bit word manufactured by Irvine Sensors, Costa Mesa, California. The memory devices are referred to as 3D silicon "short stacks" because the individual memory die are layered on top of each other, similar to a "stack of pancakes". This configuration enables a minimum height component with only a slight overall increase in height (0.060 inches versus 0.025 inches) of the original silicon. The process used to fabricate Irvine Sensors' 3D silicon "short stacks", shown in figure 5 will now be described.

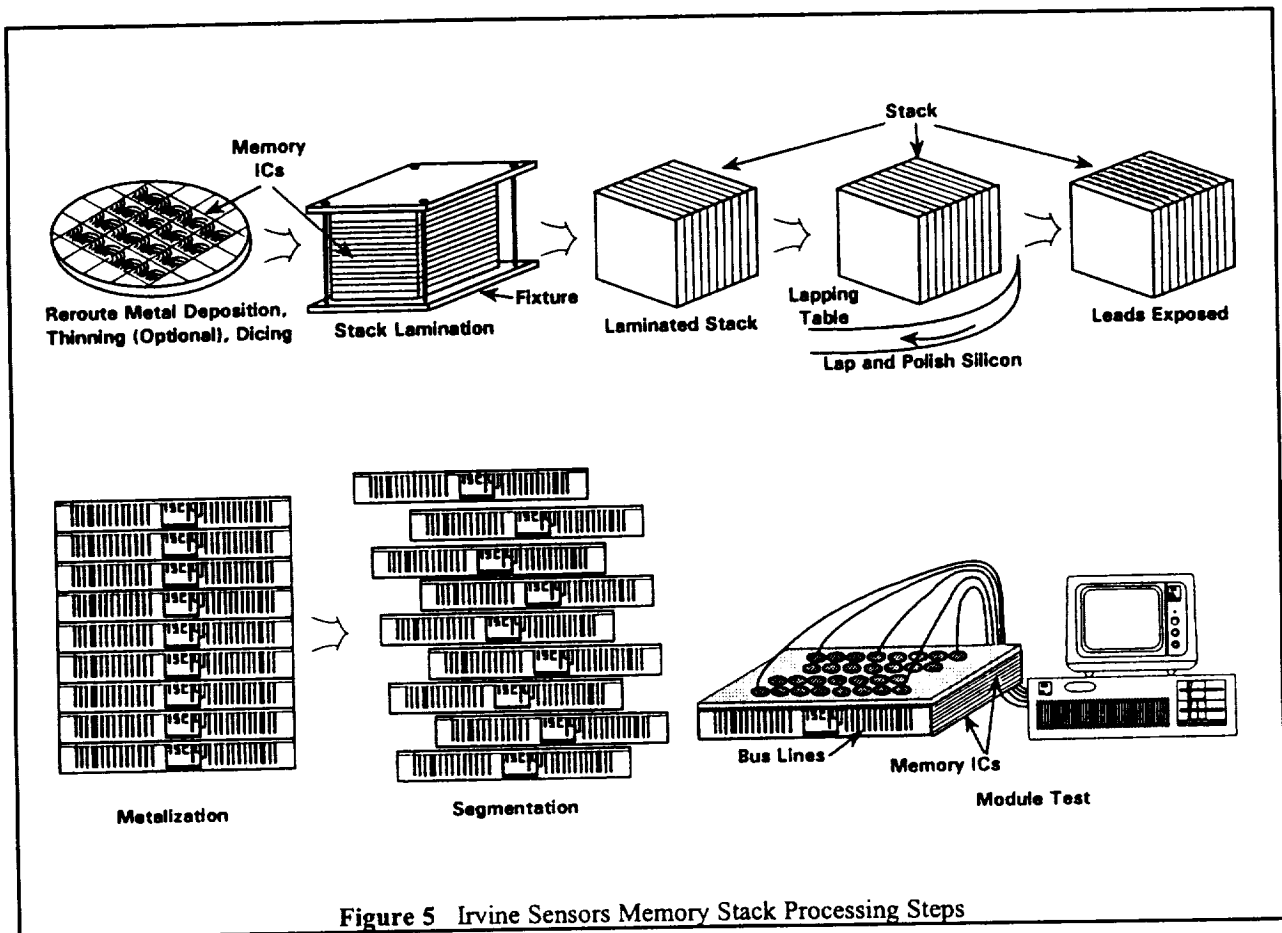


Figure 5 Irvine Sensors Memory Stack Processing Steps

Process Description

3D silicon memory fabrication begins by performing a device lead re-route at the wafer level to bring all the input-output to one side of the die. The wafers are thinned to about 0.010 inches and then sawn to provide individual memory die. The good die are laminated together (to form a cube) using a judicious application of a pair of dielectric adhesives (thermal setting and thermal plastic), special separation ceramic cap chips (top and bottom for eventual separation into short stacks) and lamination fixtures. Following cube lamination, the "cube" is lapped and polished (on the re-routed lead side) to expose the lead conductors. An etching process is applied to the cube for further lead exposure prior to passivation. The cube gets further lapping to again expose the re-routed leads (from the passivation) and then the cube gets a final (bussing) metallization applied to electrically interconnect the devices. To obtain the five layer long-word memories, the cube is heated to its thermal plastic region, and the memory component undergoes final testing.

Technology Maturity

Irvine Sensors has spent over ten years in the development of the stacked memory technology. During that time, 128-layer stacks of 4-mil layer thickness were successfully fabricated and tested to cryogenic temperatures for infrared focal plane applications. Layers as thin as two mils were successfully fabricated. "Short stacks" have survived over 100 cycles of -55°C to +125°C long thermal cycles with no change in interlayer wiring resistance. In addition, 50-layer memory stacks of 15-mil layers have been successfully fabricated, tested and temperature cycled to over 300°C with no change in interlayer wiring resistance. The stability of the technology is so promising that IBM and Irvine Sensors entered into a joint development venture (August '92) to start a memory cubing production line in Burlington, Vermont. This line is projected to be in production of 3D memory products in 1Q94.

### Advantages of Irvine Sensors 3D Memory Short Stacks

Laminated 3D memory, which is processed using thin film technology for interlayer connections, is inherently simple and robust. This technology provides maximum design flexibility and density. The advantages of Irvine Sensors 3D stacked memory are summarized in Table 1.

**Table 1 Advantages of Irvine Sensor's 3D Short Stack Memory Technology**

Attribute	Advantage
Height	<ul style="list-style-type: none"><li>• Stack height is minimized by thinning ICs to a 10 mil thickness</li><li>• Height with two cap chips is 60 mils (four layer memory stack)</li></ul>
Reliability	<ul style="list-style-type: none"><li>• Interconnect with standard semiconductor thin film high reliability techniques</li><li>• External shock &amp; vibration do not effect interconnections due to mass &amp; rigidity of stack</li></ul>
Cost & Availability	<ul style="list-style-type: none"><li>• IBM &amp; Irvine Sensors recently opened a high rate production facility for stacked memories</li><li>• First Product Available 1Q94</li></ul>
Power	<ul style="list-style-type: none"><li>• Capacitance, inductance &amp; RFI susceptibility are reduce with 3D silicon</li></ul>
Speed	<ul style="list-style-type: none"><li>• Speed is optimized due to the reduced capacitance interconnect wiring</li><li>• No more direct or lower impedance interconnects possible than Irvine Sensor's stacked memory technology</li></ul>

### **EARLY SOBIEC PACKAGING CONCEPTS**


The design goals for the SOBIEC program are to develop a high performance, minimum parts count distributed memory massively parallel processing node, that is: low in power, weight, and cost; small sized; and contain sufficient memory for most multi-computer applications. Further system level requirements levied were: low thermal impedance (4°C/Watt Junction to Case) to enable applications with severe temperature extremes, and low conducted noise (EMC). Our early SOBIEC conceptual approach (figure 6) was a 3D silicon architecture that utilized nCUBE's n2S single chip processor as an active substrate for Irvine Sensors' 3D silicon memory devices.

In this approach, a third layer of metallization is added to re-route the processor's DRAM compatible Input/Output to a single stack of ten 4 Mbit DRAMs configured as a 1 Megabyte by 40 bit word. The mechanical interface between the processor and memory, consisted of thermal epoxy and direct wire bonding between the memory stack and processor. Also, this approach required one dimension of the mechanical interface between the processor and memory, to be similar in length to the longest dimension of the stacked memory, in order to provide a stable base for memory attachment. After undergoing a design rule shrink however, nCUBE's processor failed to meet this requirement, and so an alternate approach was sought.

### **FINAL SOBIEC PACKAGE DESIGN**

After several "team" meetings, a pseudo dual cavity 138 pin grid array, alumina package approach (shown in figure 7) was selected for the SOBIEC baseline design. The salient features of this approach are excellent MCM testability, low thermal impedance -- no localized processor heating, no external "glue or ancillary parts," 4 megabytes main memory upgrade able to 16 megabytes (32 megabytes possible with an increase in height of just 0.050 inches!) in the same footprint, small sized -- only 1.2 by 1.2 by 0.31 inches in height, and low electrical noise generation.

The benefits of this approach are: simple implementation of X-Y tiled arrays of processing nodes. Each "node" requires only about one third of the original space. The reduced area is a direct result of Irvine Sensors' 3D silicon memory technology. These memories are located directly under the nCUBE n2S processor, separated by 0.09 inches of alumina. Since these memories are only 0.060 inches thick, and require little more ceramic real estate than the processor itself, a highly compact massively parallel processor node was enabled.



**Figure 6** The Original SOBIEC packaging approach used nCUBE's 2S processor as an "active substrate" for Irvine Sensors' 3D stacked memories

nCUBE's n2S processor is implemented in 1.0  $\mu\text{m}$  CMOS technology, and has modest power requirements. However its high clock frequency and numerous output buffers (address/data, control, error detection and correction) can cause "power surges" as multiple output buffers drive new signal levels simultaneously. Clean on-chip power distribution at (almost) all frequencies is provided by generous "package integral" and "package exterior" voltage bypass capacitors. The package exterior capacitors are a mixture of SMT tantalum and ceramic devices. These devices have been carefully chosen to provide less than 0.1 ohm ESR (equivalent series resistance) from about 2 kilohertz to over several tens of megahertz. Beyond this frequency, SOBIEC's "package integral" capacitance (about 0.001 microfarad parallel plate capacitor) formed by its multiple internal power and ground planes, provides effective noise bypassing to several hundred megahertz. In addition, SOBIEC input power is supplied by 38  $V_{DD}$  and  $V_{SS}$  pins to assist in the supply of low noise power.

### **SOBIEC RELIABILITY**

During the contract period, Irvine Sensors and nCUBE evaluated the reliability of the SOBIEC module in a 55° C environment. The reliability of the SOBIEC processor node included data analysis of the SOBIEC thermal management system using a combination of previous data and SOBIEC package thermal characteristics. The following thermal impedances were used for the evaluation:

Package Thermal Impedance	3°C/Watt	Junction To Case
Short Stack Thermal Impedance	3°C/Watt	Junction To Case
PC Board Thermal Management	10°C	Overall Case To Ambient Temperature Rise



The power dissipations for the SOBIEC electronics are as follows:

Short Stack Memories (0.33 watts x 5)	1.65 watts (each stack)
nCUBE Processor	2.5 watts
Total SOBIEC Power	~ 6 watts

Using this data, SOBIEC's alumina package would elevate in temperature 18°C (3°C/W x 6 W). In addition, the delta temperature of the top-most DRAM die in the short stack would be 5°C (1.65 watts x 3°C/watt). Therefore the junction temperatures of the SOBIEC devices are as follows (for a 55°C environment):

nCube Processor	83°C
Short Stack Memories	88°C

This data translates to an MTBF for the SOBIEC module of about 3 million hours.




Figure 7 SOBIEC's high system level integration is enabled by nCUBE's n2S processor and Irvine Sensors' 3D silicon and MCM technology

### ***SOBIEC OPERATING SYSTEM (OS) AND SOFTWARE SUPPORT***

SOBIEC acceptance by NASA is not only due to its unique packaging, but due to nCUBE's mature development toolset. nCUBE's software--the Parallel Software Environment--provides a familiar UNIX interface with extensions and optimizations to take advantage of SOBIEC's massively parallel hardware. The software includes a micro kernel, libraries, and UNIX utilities for SOBIEC (nCUBE) processors, as well as development and system management software for workstations on the nCUBE's network. Each software component in the Parallel Software Environment has been designed for speed and flexibility--the goals of massively parallel computing.

Running on every nCUBE processor, the nCX operating system manages processes, memory, and interprocessor communication, and supports a UNIX system call interface and POSIX signals -- all in a compact, optimized micro-kernel. Because nCX runs on I/O processors as well as compute processors, programmers can develop custom device drivers in the standard nCUBE programming environment.

nCUBE libraries include standard UNIX libraries, parallelization libraries, math libraries, and graphics libraries. In many cases, a programmer can port an application to an nCUBE 2 supercomputer simply by inserting a few parallelization calls. These calls hide the underlying communication necessary to parallelize an operation, and perform complex operations blazingly fast. A Fast-Fourier Transform or a matrix multiply can be admirably performed with a single call. All the libraries support striped files for faster I/O. Users can run UNIX utilities such as cat and tar on nCUBE processors. These utilities operate on striped files transparently.

The Parallel Software Environment's workstation software includes a set of cross-development tools for writing, compiling, launching, profiling, and debugging parallel programs. The tools use the interfaces and command-line options of tried-and-true UNIX tools. Using the debugging and profiling tools, programmers can step through parallel programs or generate bar graphs of subroutine usage or communication loads. The tools make it possible for UNIX programmers to quickly learn the basics of developing, debugging, and tuning parallel programs. Workstation software also includes user and system administration utilities for monitoring and managing the nCUBE 2 supercomputer. Users can control SOBIEC processes, load multiple programs in complex configurations, or display sample code for a subroutine. System administrators can track system usage with an accounting system, selecting shut down and reboot I/O servers, and manage resources with nQS, nCUBE's batch queuing system.

nCUBE is continuing to develop its software, making parallel processing and parallel I/O faster than ever. Within the year, nCUBE will introduce a new parallel file system that supersedes RAID 5 in performance and reliability. nCUBE is also continuing to develop its networking and database capabilities.

### ***COMMERCIALIZATION POTENTIAL***

The commercialization potential for SOBIEC is enormous. Recently, the commercial market has begun to benefit from the power of massively parallel computers. Parallel computing is taking a dual path to success. The first set of commercial users are strictly interested in the number of processing nodes that can be placed onto a fixed board. In this case, SOBIEC clearly has an edge of it's competition due to it's unique packaging technology. In the second case, commercial users are most interested in matching input/output bandwidth through the use of a parallel configuration of computers. Here again, SOBIEC's small physical size and high degree of interprocessor communications provides a competitive edge over similar technologies.

A commercial application that can immediately benefit from SOBIEC, involves applications requiring large databases such as Oracle. SOBIEC's 2S processor is designed to rapidly process transactions and very complex queries using Oracle. The natural parallelism of information in commercial databases makes them an ideal fit for SOBIEC's massively parallel computing.

### ***CONCLUSIONS***

A highly compact high performance massively parallel processing system has been developed by Irvine Sensors, nCUBE, and NASA JPL, and is in the final stages of integration and test. This production ready design realized significant size, weight, and volume reductions through the judicious application of 2D and 3D silicon technology. This general purpose processing element is packaged in a 138 leaded pin grid package that requires no more board real estate than the original packaged processor itself. The low cost alumina package exhibits excellent thermal and electrical properties and meets all the requirements for a SOBIEC mission. Completing the introduction of this product, is a mature software development system and library to ease the new or experienced user into the work of massively parallel computing.

#### **References:**

- 1 JPL Publication 82-61, End-to-End Imaging Information Rate Advantages of Various Alternative Communications Systems by Robert F. Rice, 9-1-82