

Parallel Digital Signal Processing: An Emerging Market

Application Report

Mitch Reifel and Daniel Chen
Digital Signal Processing Products — Semiconductor Group

SPRA104
February 1994



IMPORTANT NOTICE

Texas Instruments (TI) reserves the right to make changes to its products or to discontinue any semiconductor product or service without notice, and advises its customers to obtain the latest version of relevant information to verify, before placing orders, that the information being relied on is current.

TI warrants performance of its semiconductor products and related software to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Certain applications using semiconductor products may involve potential risks of death, personal injury, or severe property or environmental damage ("Critical Applications").

TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS.

Inclusion of TI products in such applications is understood to be fully at the risk of the customer. Use of TI products in such applications requires the written approval of an appropriate TI officer. Questions concerning potential risk applications should be directed to TI through a local SC sales office.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards should be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance, customer product design, software performance, or infringement of patents or services described herein. Nor does TI warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used.

Introduction

During the past decade, while CPU performance increased from 5 MIPS in the early 1980s to over 40 MIPS today, applications performance developed exponentially, especially in imaging, graphics, and high-end data processing. This “Malthusian” effect, in conjunction with the “silicon wall”, has created a situation in which application needs have vastly outpaced the ability of single processors to keep up.

This condition inspired rapid development in parallel processing, especially in digital signal processing (DSP). Currently, 75 to 80% of all 32-bit, floating-point DSP applications use multiple processors in their design for several reasons. First, DSP algorithms are inherently suited to task partitioning and, thus, to parallel processing solutions. Second, as the cost of single-chip DSPs decrease, using multiple DSPs in a system becomes increasingly cost effective. Third, the high data throughput, real-time processing capability, and intrinsic on-chip parallelism of DSPs make them especially suitable for multiprocessing systems.

Simply put, parallel processing uses multiple processors working together to solve a single task. Processors can either solve different portions of the same problem simultaneously or work on the same portion of a problem concurrently.

This paper discusses digital signal parallel processing as well as the reasons why DSP and parallel processing have become a natural match:


- Advances in CPU architectures.
- New developments in hardware development tools.
- The emergence of software languages and operating systems for multiprocessing.

This paper looks at solutions from different vendors as well as trends in the industry as a whole.

The Technology Merge

The first practical single-chip DSPs were introduced in the early 1980s. Because of their real-time processing capability, high throughput, and intensive math-processing capability, DSPs began to replace general-purpose processors in many applications. These applications were well suited for real-time processing such as speech processing, telecommunications, and high-speed control. They also pushed DSP to the forefront of technology and created one of the fastest growing markets of the decade (see Figure 1).

Figure 1. Worldwide Single-Chip DSP Market




The DSP market was one of the fastest growing markets of the 1980s. Parallel processing is predicted to follow a similar pattern in the 1990s.

DSPs are now used in a broad range of nontraditional applications, such as graphics, imaging, and servo-control, that were not originally thought of as part of the signal processing domain. Application designers turned to DSPs because their cycle times were faster than those of general-purpose and RISC architectures. By the middle 1980s, however, cycle time improvements in each new generation became smaller.

In the 1990s, processor manufacturers are approaching the physical limitations of silicon and can no longer rely on smaller geometries alone for increasing processor performance for next generation products, as shown in Figure 2.

Figure 2. DSP Performance Evolution




All semiconductor manufacturers are approaching the "silicon wall" and are looking at different multiprocessor solutions to get around the problem.

In the meantime, tasks that were unheard of just a few years ago — such as virtual reality and video recognition — are pushing the envelope of performance requirements. Figure 3 shows the trend with actual designs that use TMS320 DSPs.

Multiprocessing meets these challenges. However, multiprocessing comes in different forms. Some manufacturers gain performance improvements with on-board architectural enhancements, but this technique alone cannot meet every need.

Figure 3. Performance Requirements (Actual TMS320 Designs)



Continued growth in application requirements demands intensive development in processor technology.

On-Chip Vs. Off-Chip Parallel Processing

Parallel processing enhancements can be divided into two broad categories: on-chip and off-chip. On-chip parallelism relies on architectural enhancements for improved performance, while off-chip parallelism incorporates additional processors.

On-Chip Parallel Processing

Architectural enhancements on RISC processors can be grouped into three distinct categories: superpipelining, superscaling, and multi-CPU integration.

Superpipelining

This technique breaks the instruction pipeline into smaller pipeline stages, allowing the CPU to start executing the next instruction before completing the previous one. The processor can run multiple instructions simultaneously, with each instruction being at a different stage of completion.

The main drawbacks of this technique are the increased level of control logic on the processor, difficulty in programming, and difficulty in task switching. Real-time multitasking on a superpipelined processor can become impossible if the pipeline grows too deep.

Processors that use superpipelining are the Intel i860 and the MIPS R4000.

Superscaling

Instead of breaking the pipeline into smaller stages, superscaling creates multiple pipelines within a processor, allowing the CPU to execute multiple instructions simultaneously.

However, when multiple instructions are executed simultaneously, any data dependency between the instructions (such as a conditional branch) increases the complexity of the programming. Programmers must make certain that simultaneously executed instructions don't need the same on-chip resource, or that one executing instruction doesn't need the result of another whose result is not yet available.

Digital's Alpha processor is one example of a CPU that uses superscaling.

Multi-CPU Integration

This technique goes a step further than the preceding techniques and integrates multiple CPUs into a single piece of silicon. The number of processors may vary, depending on chip size, power dissipation, and pin count.

Star Semiconductor's SPROC and the soon-to-be-announced MVP (Multimedia Video Processor) advanced imaging processor from TI implement this technique.

All three of these parallel processing techniques increase processor performance without the need for dramatic cycle time improvement. None of the techniques, however, can achieve the BIPS performance required by today's applications. If an application demands higher performance than on-chip processors can deliver, the solution must be multiple processors.

Off-Chip Parallel Processing

Off-chip parallel processing is not necessarily better — it's inevitable. No single processor, no matter how it is pipelined, how it is scaled, or how many CPUs it has on board, can handle all applications. Recognizing this, manufacturers developed techniques to integrate multiple processors efficiently. Like building blocks, off-chip parallel processors connect easily to form expandable systems of virtually infinite size and variety.

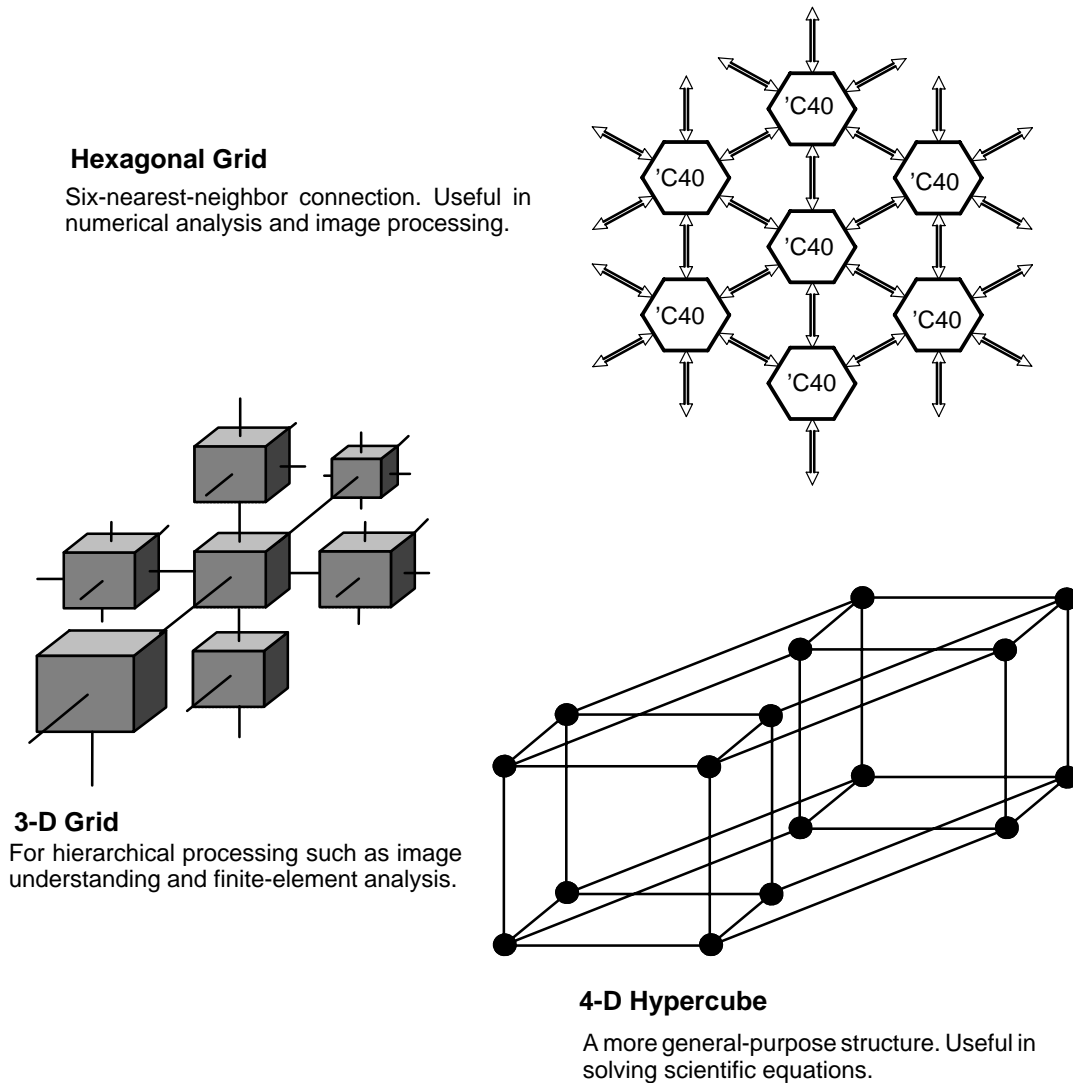
Two processors employ this technique: the Inmos Transputer and the Texas Instruments TMS320C40. Both of these processors also incorporate on-chip parallel processing features to achieve high individual performance. The latest generation Transputer, T9000, uses superpipelining, while the 'C40 uses superscaling. These processors offer both the high performance of the on-chip parallel processing architectural enhancements and the extra features of off-chip expansion.

Off-chip expansion is achieved by connecting multiple processors together with zero glue logic for direct processor-to-processor communication. While methods are different, (TI uses six 8-bit parallel communication ports; Inmos uses four serial links), the concept is the same: connect multiple processors together to create a topology or array of virtually any size to achieve the performance needed by high-end applications (see Figure 4). The communication ports (or links) on the devices are supplemented by parallel memory buses and other support peripherals, allowing designers broad flexibility in designing their systems.

These are some benefits of off-chip parallelism:

- **Expandability** — You can easily add more processors to your system to meet performance requirements.
- **Flexibility** — You can implement a wide array of processor topologies that best fit your application needs. Unlike hardwired multi-CPU integration, off-chip processing can implement everything from 1D pipelines to 4D hypercubes.
- **Upgradability** — With processors that connect like building blocks, systems can be designed in a modular fashion, allowing extra processing power to be added at a later date to meet expanding processing needs.

Figure 4. TMS320C40 System Architectures



The TMS320C40 has six interprocessor communication ports for creating topologies of virtually any size and type.

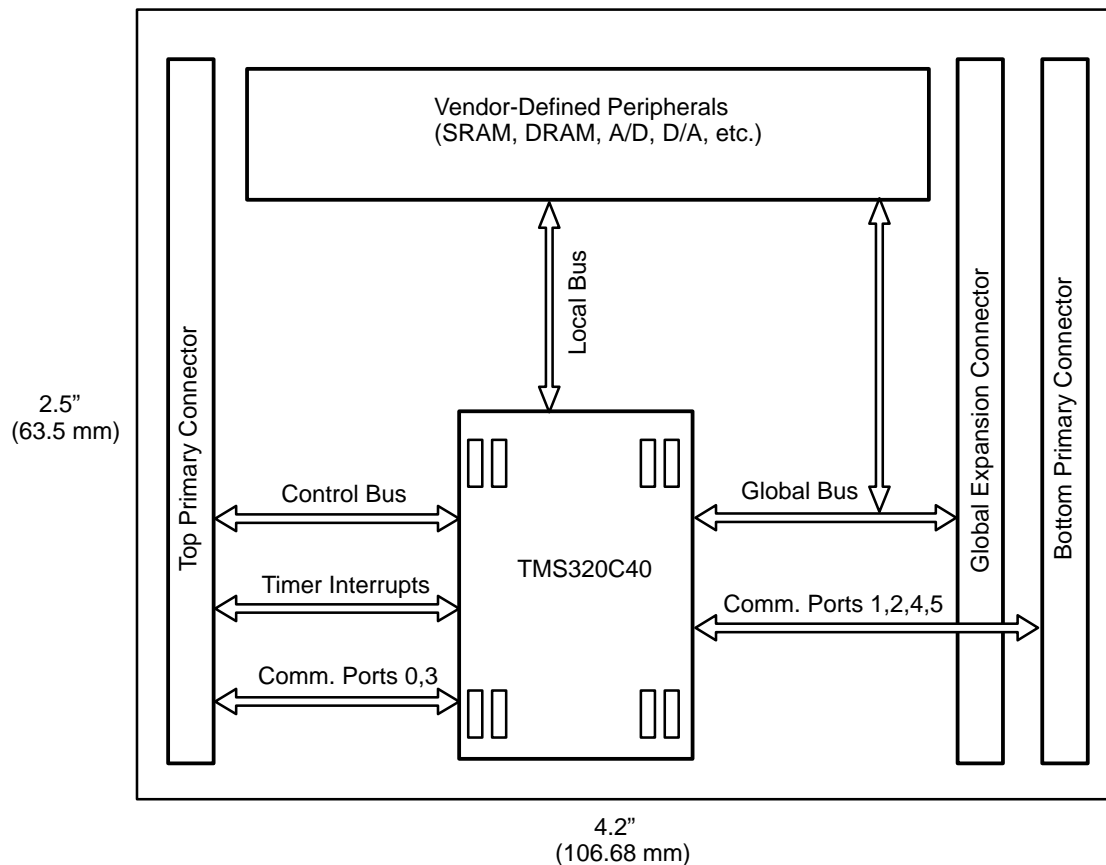
Processing Modules

Upgrading, expanding, and integrating parallel systems is even easier with processing modules than with processors. TRAMs (Transputer Modules) for the Transputer and TIMs (TI Modules) for the 'C40 provide an open standard, easy-to-use approach that saves time.

The TIM-40 and TRAM describe modular building blocks for prototyping and manufacturing parallel-processing systems. Both standards consist of a daughterboard module that can include a parallel processor,

memory, A/D-D/A conversion, and other functions as required. System designs can contain any number of modules, limited only by the amount of room in the system (see Figure 5).

Figure 5. TIM-40 Module



The architecture of the TIM-40 gives you both a standard interface to build parallel processing systems and also the flexibility to add support peripherals and features that best fit your application.

Designs based on the modules can be scaled and upgraded easily as system performance requirements increase. Furthermore, modules used in development activities can be reused in new programs.

The modular approach helps designers enhance system reliability. In a massively parallel system that requires 100 'C40s, TIM-40 modules can reduce the challenge of more than 3,000 pin connections to a task of only 200 daughterboard-to-motherboard connections.

These architectural enhancements make hardware design and integration of multiple processors easy, but they do not address debugging and programming the large parallel systems that result. This is where the 'C40 and Transputer differ. The designers of the 'C40, realizing the problems in debugging large parallel systems, built into the processor features that allow unique multiprocessing debugging capabilities.

Multiprocessing Emulation

Programming and debugging single, serial processors has always been difficult. Programming the enhanced processors in multiprocessor systems is even more difficult. Prior to the availability of the 'C40

and its development tools, developers used tools intended for uniprocessor architectures to design and debug multiprocessor systems. While such tools were satisfactory for their original purpose, they were difficult to use with embedded processors in parallel architectures. Designers used multiple emulators and/or complicated software monitors to debug their parallel systems. These tools provided neither system synchronization, unintrusive real-time operation, or the fine detail required to design and debug embedded parallel processors.

The 'C40 XDS510 in-circuit, scan-based emulator incorporates the same cutting-edge tasks that are used for parallel supercomputing. It supports global starting, stopping, and single-stepping of multiple 'C40s in a target system. It also has the capability to halt all the 'C40s in a system if a single 'C40 hits a breakpoint.

This parallel debug capability of the XDS510 is supported by the on-chip analysis logic designed into the 'C40. The XDS510 can access the analysis module to efficiently debug, monitor, and analyze the internal operation of the device. The analysis module consists of an analysis control block, an analysis input block, and a JTAG test/emulation interface block. The module features program, data, and DMA breakpoints, a program counter trace-back buffer, and a dedicated timer for profiling.

A single XDS510 emulator can perform mouse-driven, windowed debugging of C and assembly language for all the 'C40 processors in your system, regardless of the complexity of the topology. It also determines whether the system load is balanced across the processors.

The TMS320C40 is the only parallel processor that has this emulation and debugging feature.

Software Development

One of the largest problems facing developers of multiprocessing systems is programming. Issues such as program partitioning, load balancing, and program routing present unique difficulties. Various solutions have been offered:

Graphical Programming Languages — Comdisco Systems recently introduced Multiprox — the first graphical programming environment for developing systems that employ multiple TMS320C40s. Multiprox lets you partition a signal flow block diagram into regions for separate processors to execute. Multiprox automatically generates code for each processor, then compiles and downloads the code with all the necessary interprocessor communication. As a result, you can develop algorithms in less time, and the development process is simplified for those who are not parallel processing experts. Topologies of any size and variation can be used with the system.

Operating Systems (OS) — Various operating systems are available to help designers implement realtime multiprocessor systems:

- Helios is a distributed parallel operating system designed to run on multiple-instruction/multiple-data (MIMD) architectures, making it ideal for use in processing modules. After the OS is distributed across the network, each processor runs the Helios nucleus, and they all operate together as a single processing resource. The UNIX-like interface and Posix programming interface allows developers familiar with these environments to program on the 'C40 quickly and easily.
- SPOX offers a hardware development platform and run-time support for real-time systems, thereby simplifying the development of embedded multitasking applications. 'C40 SPOX provides comprehensive sets of parallel DSP operations and includes a high-level software interface that makes it easy to utilize the 'C40's communication ports and DMA coprocessor. SPOX supports both multiprocessing and multitasking applications.

- RTX/MP for the 'C40 is designed for complex distributed systems with large arrays of processors and has support for fault-tolerant systems.

Parallel Programming Languages — Programming languages are emerging to help the programmer implement software across multiple processors. Parallel C for the 'C40 has been introduced by 3L Ltd. Parallel C is a full implementation of C with many additional features that support parallel processing. The compatibility with C allows existing single processor applications to be ported easily and quickly to parallel systems while the parallel processing features facilitate easier network programming and communications. Other languages available on the 'C40 include ANSI C and Ada, both of which come with multiprocessing support.

Emerging Trends

One of the questions usually asked about the flattening in performance of silicon speed is, “What about gallium arsenide?” (also called GaAs). To date, no semiconductor manufacturer has planned mass production of GaAs-based processors, and it will probably be another decade before GaAs processors make it onto the market. When GaAs processors do appear, their performance by itself still won't meet the requirements of the newest applications. Multiprocessing, even with GaAs processors, will be a necessity.

A more imminent trend is multichip modules (MCMs). This is simply an extension of the off-board processing theory that puts multiple processors into a single package, thus requiring smaller pin count and board area than if the processors were used separately. MCMs provide the best of off-chip and on-chip parallel processing. They offer the improved thermal management, power distribution, and signal integrity of signal processors, as well as the flexibility, upgradability, and expansibility of off-chip parallel processing. TI has already announced dual and quad 'C40 MCMs. Even higher integration with new packaging advancements, such as 3-D packaging, are planned.

Conclusion

The inability of single-chip processors to keep up with the expanding needs of emerging applications makes parallel processing potentially one of the most rapidly growing technologies of the 1990s.

On-chip parallelism can improve performance only to a certain degree. Off-chip parallel processing can increase the performance almost infinitely. Three key factors of parallel processing have been identified: interprocessor communication, parallel debugging, and parallel programming. Two processors, the Texas Instruments 'C40 and Inmos Transputer, were discussed. While both processors incorporate features for high-speed processing and off-chip interprocessor communication, only the 'C40 has the on-chip debug capability and the programming tools needed for programming arrays of processors of arbitrary size and complexity.

Bibliography

Peterson, Robert, and John Scoggan, *Electronic Packaging in DSEG*, Texas Instruments Technical Journal, Volume 9. No. 3, May–June 1992.

Simar, Ray, *The TMS320C40 and Its Application Development Environment: A DSP for Parallel Processing*, International Conference on Parallel Processing, Volume 1, p. 149–151.

Weiss, Ray, "Third Generation RISC Processors," *EDN*, March 30, 1992, p. 96–108.

IMPORTANT NOTICE

Texas Instruments (TI) reserves the right to make changes to its products or to discontinue any semiconductor product or service without notice, and advises its customers to obtain the latest version of relevant information to verify, before placing orders, that the information being relied on is current.

TI warrants performance of its semiconductor products and related software to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Certain applications using semiconductor products may involve potential risks of death, personal injury, or severe property or environmental damage ("Critical Applications").

TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS.

Inclusion of TI products in such applications is understood to be fully at the risk of the customer. Use of TI products in such applications requires the written approval of an appropriate TI officer. Questions concerning potential risk applications should be directed to TI through a local SC sales office.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards should be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance, customer product design, software performance, or infringement of patents or services described herein. Nor does TI warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used.