

# A Systems-Oriented Microprocessor Bus

National Semiconductor  
Application Note 449  
June 1986



## INTRODUCTION

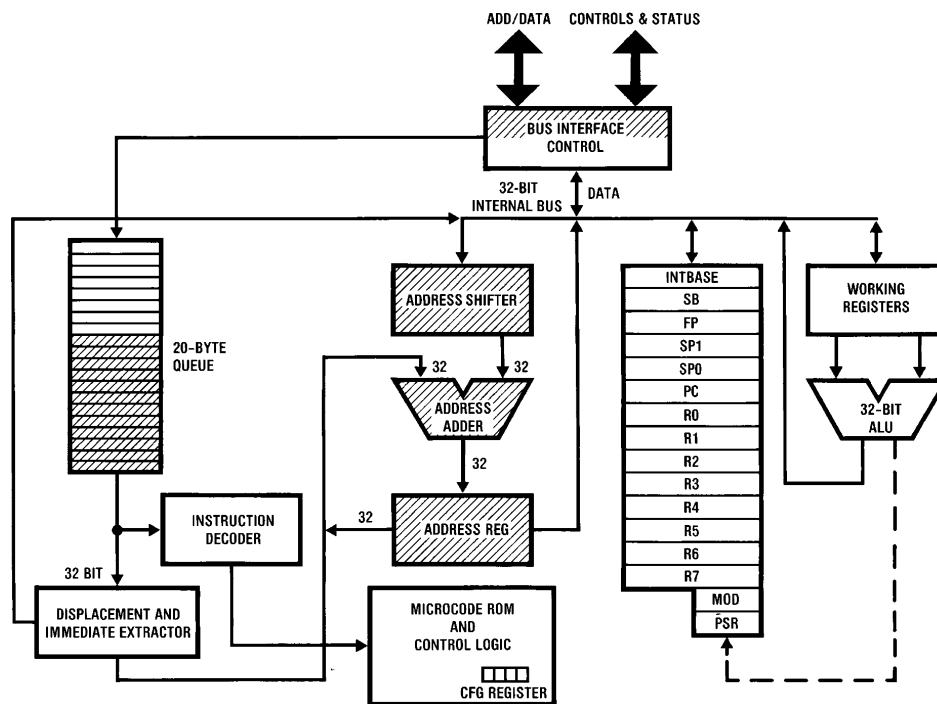
Viewed from the standpoint of hardware throughput, the performance attained by state of the art architectural and technological implementations is a result not only of what gets done per instruction, but also of the available memory speed and bandwidth as perceived by the processor. In addition to the memory parameters, the effects of DMA, processors' contentions and the system bus itself needs to be considered. The clock frequency provided by the process and design techniques is limited by the complexity of architecture and its implementation. This mechanism limiting the performance of the "theoretical" machine running by itself; the actual in-system performance is further limited by the protocols interfacing the machine to system memory and the bandwidth of memory.

Until recently, advances in processes and architectural implementations were the main tools used to improve performance in a uniprocessor or few processors' environment, due largely to the price and size of computing clusters which were considered to be the main resources of the system and the resulting price of communications to memory and I/O. The introduction of computing clusters such as the Series 32000® family, consisting of CPU (Central Processor

Unit), MMU (Memory Management Unit), FPU (Floating Point Unit), TCU (Timing Control Unit), has changed the price/performance criteria to the point where the use of multiple processors has become a viable way to leapfrog the far too slow progress in process and integration capabilities.

Given the appropriate tasks and software, a multiprocessor system can increase its total throughput by a very significant factor compared to the performance that available processes and hardware have acquired for the single processor. Furthermore, this throughput is obtainable incrementally, making the multiprocessing a more flexible and better long range investment for both OEMs and end users.

The NS32332, a new CPU in the Series 32000 family, has addressed the issues of fast system/memory bus protocol and multiprocessing support in addition to internal architectural enhancements. Together with the NS32382 MMU and NS32310 FPA, it becomes the third computing cluster to be released by NSC with a 15 MHz relative performance evaluated at 2-3 times the NS32032 cluster at 10 MHz. The size and price of Series 32000 computing clusters together with the denser packing of other components have combined to



TL/EE/8762-1

**FIGURE 1. NS32332 Internal Architecture.** Shading indicates enhancements to the NS32032 implementation. The principal new elements are a high-speed ALU and address shifter in the address-computation unit, an additional 12 bytes in the instruction prefetch queue (used with burst transfers), and a new high speed external interface.

Series 32000® and TRI-STATE® are registered trademarks of National Semiconductor Corp.

provide multiprocessing performance equivalent to the performance of superminis and mainframes with the added bonus of reliability and low cost maintenance features. System architectures using private caches and global memory for example, are now in a position to contemplate delivering 10 MIPS per board due to the compact size of the new cluster, with maximum performance per box driven by the bandwidth and physical length of the system bus.

**Note:** The MIPS referenced are Series 32000 instruction set MIPS.

### NS32332 ARCHITECTURE

The NS32332 (Figure 1) has been configured as a hardware microarchitectural improvement of the NS32032 CPU, using the relevant parts of its database, automatically ensuring the correctness of the programmer's architectural model. The main NS32032 referenced architectural enhancements relate to the addition of new dedicated addressing hardware (consisting of a high speed ALU, a barrel shifter and an address register), a very efficient increased (20 bytes) instruction prefetch queue, a new system/memory bus interface/protocol, increased efficiency slave processor protocol and finally enhancements of microcode.

The resulting machine is a three stage pipelined microprogrammed CPU. Compared to the NS32032, it has a larger address space (32 bits), faster address calculation, faster execution time and faster memory references causing less "contention per mips" both internally and externally with the system and other processors. The improvements have been made possible using a more advanced process which enabled the implementation of additional circuitry on the chip as well as a higher frequency, the full effect of which is ensured by the zero-waits implementation possible at 15 MHz.


### PHYSICAL DESCRIPTION

Featuring full 32 bit internal and external address/data paths, the 10.1 x 8.7 mm chip contains 80,000 transistor sites and is housed in an 84 pin PGA package. It is implemented in 2.8 micron drawn minimum feature, single metal, single poly NMOS and operates at 15 MHz from one +5V supply dissipating a worst case evaluated power of 3W. The externally generated clock delivers two non-overlapping phases and provides the required system synchronizations via a TTL compatible clock and special circuitry for multiprocessor clock synchronization.

### MEMORY PARAMETERS

The efficiency of the machine/memory protocol determines the maximum frequency at which the machine will run without wait states, the machine transaction rate with memory, and the minimization of contention between processors.

Pipelined machines' performance is impacted by the introduction of wait states due to memory access time (Figure 2), exhibiting sometimes 25% performance loss on the first wait state with decreasing percentages for losses due to subsequent wait states. Non-optimized protocols may negate the achievements of process and circuit technologies by losing the performance advantage to inefficient system interface. For a given machine/memory protocol, the operating frequency and the number of clocks allowed for memory access will determine the required memory access time; conversely, the maximum achievable zero-wait-states frequency, is determined by the available system memory access times and the number of clocks allowed by the protocol between address generation and data setup (memory read cycle).



TL/EE/8762-2

**FIGURE 2. Processor with inefficient system/memory bus. Performance, high without wait states, drops significantly when introduced in a real system. Note the marginal effect clock frequency (f) increase has in the presence of wait states.**

A second parameter is the memory bandwidth required by the machine to support the transaction rate it needs to develop its throughput. Given that a relationship exists between the MIPS rate of the machine and the average number of memory accesses it needs to make per instruction (workload dependent), the available memory bandwidth is perceived to directly impact performance, memory cycle times, protocol efficiency and contention with DMA, cache operations and/or other processors will combine to determine the throughput achievable in a real system.


As will be shown later, the NS32332 cluster has met the above requirements by providing a protocol supportive of zero-waits at top frequency, requiring minimal bandwidth by using burst transactions, achieving in the process a combination of high speed access time and low memory bandwidth.

### MULTIPROCESSING SUPPORT

Extending performance beyond the capabilities of both chip technology and system memory bandwidth, multiprocessing architectures can efficiently harness the power of many processors by providing the appropriate software environment and by increasing the available memory bandwidth. Multiprocessing support consists of providing the features required to support the integrity of program execution across multiple processors and the hardware implementation of memory bandwidth extension. Particularly demanding is the architecture using private caches and global memory. In addition to supporting system interlocks to provide atomic semaphore operations, the computing cluster must ideally provide efficient matching to cache speeds at the maximum cluster frequency. It must obtain zero-wait-states while still ideally providing physical cache addressing, and have low memory bandwidth to minimize the impact of cache operations.

From the point of view of cluster design, these requirements translate to:

1. Providing hardware support for indivisible CPU and MMU operations.
2. Zero-wait-states access timing has to include the MMU translation delay.
3. Reduce memory bandwidth
  - by short minimal protocol
  - interleaved memory techniques
  - pipelined arbitration
  - fast turnaround of I/O TRI-STATE®s



**FIGURE 3. Simplified block diagram of two way set associative cache shows the match and data paths during cache read. MMU internal cache match and data cache match are overlapped. MMU translation drives directories and data RAM in a fall-through mode.**

4. Efficient memory management
  - high translation buffer hit ratio
  - high speed MMU-system memory transactions

#### NS32332 PERFORMANCE AND FEATURES

Featuring 4 Gbyte uniform addressing, the 32-address-bits, 32-data-bits NS32332 is designed to provide at 15 MHz 2–3 times the throughput of the NS32032 at 10 MHz. The performance enhancement is partly due to the 1.5 frequency increase (all of which is utilizable due to zero-waits in realistic 15 MHz systems) and partly due to architectural enhancements (the new address hardware contributes 25%, the bus improvements 15%, the deeper queue 8%, the microcode enhancements contribute 2%). Performance improvement has been approached at the system level via a balanced design where the internal architecture has decreased the “clocks per instruction”, the process and circuits have increased the frequency and the system/memory bus has been designed to provide the bandwidth and access time required for full in-system throughput.


Internally, less clocks per instruction are obtained by two ALUs sharing the work using dedicated buses. The multiple bus Address Unit hardware has improved performance by providing higher speed effective address computation and access to internal registers, a barrel shifter on one input provides array indexing of items sized 1, 2, 4, or 8 bytes with no additional time delay.

The deeper instruction prefetch queue, supported by the burst protocol contributes to reducing the clocks per instruction by minimizing the average instruction access time since instructions can now be fetched not only faster but also causing less contention with data. Similarly, the data path speed has been improved due to the decrease of contention with the instructions path.

The process and circuit techniques are providing a frequency of 15 MHz in a fairly simple chip (80,000 transistor sites) for its instruction set and throughput. The other chips, far simpler than the CPU, are running at the same clock frequency to provide a fully integrated computing cluster.

The NS32332 system/memory bus has been designed to provide the longest time possible for use of RAM (at zero-waits, the better part of two 66 ns clocks, is available between non-multiplexed physical address and read data required). To support the more stringent requirements of multiprocessing caches, the bus design was driven by the need to provide a physically addressed, zero-waits design, not only on the data path but also in generating the cache match signal. In *Figure 3*, the two critical paths are shown on a simplified diagram of a two-way set associative cache. The data access is less demanding as the physical address needs to propagate only via the cache data RAM. The match propagation path has to ascertain that both the internal MMU and cache directories have found a match and still drive the Ready logic with appropriate setup time to stop the transaction if a miss has to be acted on.

In cache read, for example (*Figure 4*), the critical paths of the two way set associative cache have both been supported for a physical cache design using existing components without the need of wait states. Compared to the NS32032, the major enhancements are focused on accepting the Ready signal one half clock later and a new approach for MMU design which has the translation buffer operate in fall-through mode for increased speed, overlapping virtual address propagation and physical address translation, to save the one or two waits required by other processors. The new bus protocol has reduced the clock's count to only four clocks per transaction, MMU translation included, with only two clocks dedicated to actual memory communication, thus decreasing the required memory bandwidth in circuit




**FIGURE 4. NS32332 Cache Read. The protocol at 15 MHz is designed to drive realistic cache memory without wait states.**

implementations where bus switching is performed by AS buffers. Less memory bandwidth sensitivity has resulted. The impact of DMA and processor contention is minimal; with AS buffers two processors may be connected to the same memory bus.

Further reductions of access time and bandwidth requirements are achieved by the provision of a dynamically controllable burst protocol (*Figure 5*), which improves the perceived memory access time by factor 1.6 for interleaved memory systems, nibble mode or static column RAMs. In nibble mode, for example, for inexpensive implementations, the presence of one wait state will be felt only during the first access causing the whole transaction to behave as if the memory was exhibiting only 1/4 wait state. An additional benefit obtained by the burst mode is the reduction of internal contention between the instructions and data transactions. The burst mode is especially efficient for instructions

prefetching, long/non-aligned operands and for bus sizes smaller than 32 bits. The burst protocol involves the use of a Burst Request line driven by the CPU and a Burst Acknowledge line driven by the memory or device addressed. Both lines are dynamically sampled to check for burst continuation. The removal of Burst Acknowledge will dynamically transform the remainder of the transaction into a renewed normal memory access. Data during the burst is moved at 2-clock intervals without regenerating the address; data rate during the burst may reach a maximum of 30 Mbytes/sec, about twice faster than comparable processors. Control of the Ready line may be employed to reduce burst speeds to match the needs of slower memories/buses.

The new bus provides enhanced fault tolerant support via dedicated Bus Retry and Bus Error lines. A dynamically variable data bus width (8, 16, or 32 bits) allows the machine to execute code of varying widths for economy, available system bus data widths or fault tolerant purposes. The burst feature makes bus width reduction be less impacting than normal accessing, for instance, only six rather than eight clocks are required to transfer a 32-bit operand over a 16-bit bus. Bus arbitration is handled at three levels depending upon the nature of the transaction; for interchip communications it is provided by the slave protocol; for chips-to-system communication, by exercising of the CPU Float line and for computing cluster-to-system accesses, by options beginning with the classical but slow Hold/Holdack protocol and ending with a pipelined arbitration scheme to take full advantage of the available memory/system bandwidth. Pipelined arbitration is further supported by early processor BIU transaction status (for determining priorities between processors) and an advanced status strobe line. Prevention of system bus deadlocks is supported by either the bus retry function or cycle hold function, both of which can regenerate the transaction. Interlocked arbitration is made possible via dedicated lines driven by the CPU and MMU chips. The local/system arbiter may then decide if and when to allow the interlocked access using if necessary the bus retry function.



**FIGURE 5. Burst transactions begin with a normal transaction followed by additional data transfers at twice the speed. The BREQ and BACK lines are used to implement the handshake between memory and cluster.**

The slave bus is implemented using the same a/d lines as the system/memory bus. The local-system/memory buffers are used to disconnect the local bus from the system bus during local transactions. Slave transactions are two cycles long and may use the full 32-bit bus width. Compared to the NS32032, the slave overhead has been reduced considerably by internal improvements in the instruction decode/operand move paths as well as by the new 32-bit bus.

The NS32332 protocols' complement supports all of the present and future computing cluster members to provide for easy hardware upgrading and economy of solutions. Switching slave configurations is programmed by the use of configuration registers.

#### **SYSTEM DESIGN STRATEGY**

The computing cluster implementation provides the highest, easiest to purchase throughput per square inch; the small size and price encourage its use in multiprocessor architectures. In single processor configurations, aside from performance, the computer cluster provides the least price and development risk for new products. The Series 32000 family provides continuously improving top performance comput-

ing clusters, tools, compilers and operating systems, covering the most important features of mainframes and minis. The system designer is now free to concentrate on the real software and hardware issues concerning the implementation of system architecture with his efforts spent in the directions which make his system competitive and timely. The up-down compatibility of the architecture across chip generations ensures the continued value of software investments, a benefit carried over to the end user.

#### **SUMMARY**


A first step in the direction of multiprocessing has been taken by the Series 32000 family of computing clusters. With unchanging long term architecture and multiple hardware protocols, the NS32332 cluster provides increased performance without devaluing the software investment, packing 2-3 times the performance of its predecessors in a very small, inexpensive set of chips. As increased cache sizes become possible via denser packing, multiprocessing configurations find themselves in a good position to incrementally deliver throughput previously found only in superminis or mainframes. The NS32332 cluster contributes strongly to the performance, size and price of the new machines and to their longevity.

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.

2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

 <b>National Semiconductor Corporation</b> 1111 West Bardin Road Arlington, TX 76017 Tel: 1(800) 272-9959 Fax: 1(800) 737-7018	<b>National Semiconductor Europe</b> Fax: (+49) 0-180-530 85 86 Email: cnjwge@tevm2.nsc.com Deutsch Tel: (+49) 0-180-530 85 85 English Tel: (+49) 0-180-532 78 32 Français Tel: (+49) 0-180-532 93 58 Italiano Tel: (+49) 0-180-534 16 80	<b>National Semiconductor Hong Kong Ltd.</b> 19th Floor, Straight Block, Ocean Centre, 5 Canton Rd. Tsimshatsui, Kowloon Hong Kong Tel: (852) 2737-1600 Fax: (852) 2736-9960	<b>National Semiconductor Japan Ltd.</b> Tel: 81-043-299-2309 Fax: 81-043-299-2408
---	---	--	--

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.