



IBM Chipkill Memory

Advanced ECC Memory for the IBM Netfinity 7000 M10

IBM Chipkill Memory (Chipkill), part of the IBM Netfinity X-architecture initiative, will be available during the first quarter of 1999 as an option for the IBM Netfinity 7000 M10 server. With Chipkill, an IBM Netfinity 7000 M10 will be protected from any single memory chip that fails and any number of multi-bit errors from any portion of a single memory chip. This advanced technology derived from IBM large systems will greatly reduce Netfinity 7000 M10 downtime, resulting in a more robust Intel processor-based client-server computing platform.

Background

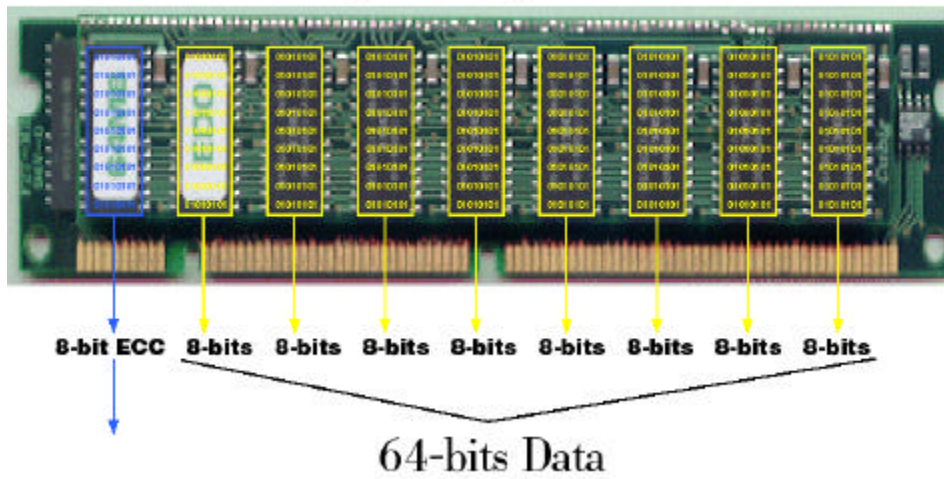
In the early 1990s, most Intel processor-based servers employed parity memory technology. Parity memory can detect errors but cannot correct them. When a memory parity error occurs in a server, the system crashes and all data stored in memory is lost. For example, the server may be processing 32-64MB of data that is in memory but not yet stored on disk. In this instance, a memory parity error could cause a significant amount of information to be lost. (All servers put recent data into memory before the data is stored onto disk.)

In the past, memory parity errors occurred with sufficient frequency to warrant a move by the entire industry to adopt ECC (Error Checking and Correcting) memory. ECC memory can detect and correct a single-bit error and can usually detect, but not correct, double-bit errors. Because multiple bit errors were rare, ECC provided a significant increase in reliability. Today ECC memory technology is standard on almost every server.

To obtain acceptable performance, today's servers require large amounts of memory. This is partly because Intel processor-based servers have increased CPU performance by a factor of about 20 since the early 1990s. However, disk drives have improved performance by only about 5 in the same period. To obtain adequate performance, large amounts of memory are used to isolate fast CPUs from slow disk drives. Because servers require large amounts of memory, customers have also demanded lower memory prices. Low memory costs are obtained by using memory chips that contain as many bits as possible. By using very dense memory chips, the \$/MB ratio has been lowered drastically. To achieve this density, however, a single memory chip usually provides 4 or 8 data bits on each access.

An unrecoverable data loss condition occurs when a memory chip on a memory DIMM fails because the number of bits supplied from each memory chip is greater than ECC can protect. In the following example, each memory chip supplies 8 bits of data. When a memory chip fails, ECC cannot recover the 8 bits of data that are lost. As a result, the server crashes. Data loss from memory chip failures can occur for both 4-bit and 8-bit wide memory technologies because standard ECC only protects against single-bit failures.

Memory Organization



On current high-performance Intel processor-based file servers, these memory chip failures cause a system crash that can result in the permanent loss of several gigabytes of business data! This is a catastrophic failure with no possible data recovery in most file server applications.

Memory failures even create significant downtime for database server applications that protect from data loss by logging each transaction to disk. The reason is that the database recovery process may require many hours to consolidate the database contents with the transaction log updates that have not yet been written to the database. Thus, it can take several hours for a database server to recover from a memory failure before the database can resume business transactions. This database recovery period may result in a loss of up to many hours of business productivity.

Solving Memory Failures

For those systems requiring the highest availability, such as the IBM S/390 class of enterprise servers, the problem of multi-bit or DRAM failures is handled by architectural means in main memory. The memory subsystem design is such that a single chip, no matter what its data width, will not affect more than one bit in any given ECC word. For example, if 4-bit wide DRAMs were in use, each of the 4 bits would feed a different ECC word; that is, a different address of the memory space. Thus, even in the case of an entire memory chip failure, no single ECC word will experience more than one bit of bad data—a situation that can be fixed by the ECC logic, thereby maintaining the fault-tolerance of the memory subsystem.

Thus, IBM mainframe computers are protected against memory chip failures because of an intelligently designed memory subsystem, which organizes each memory module so that the number of bits supplied by each memory chip is equal to the number of bits protected by the ECC logic. Current Intel processor-based servers do not have this capability because market forces demand cheap memory, forcing designers to use very dense "industry standard" memory chips.

IBM engineers have solved this problem for Netfinity servers by placing a Redundant Array of Inexpensive DRAM (RAID) processor chip directly on the memory DIMM. The RAID chip calculates an ECC checksum for the contents of the entire set of chips for each memory access and stores the result in extra memory space on the protected DIMM. Thus, when a memory chip on the DIMM fails, the RAID result can be used to "back up" the lost data, allowing the Netfinity server to continue functioning. This RAID technology is similar to the RAID technology used to protect the contents of

an array of disk drives. We call this memory technology Chip-kill DRAM.

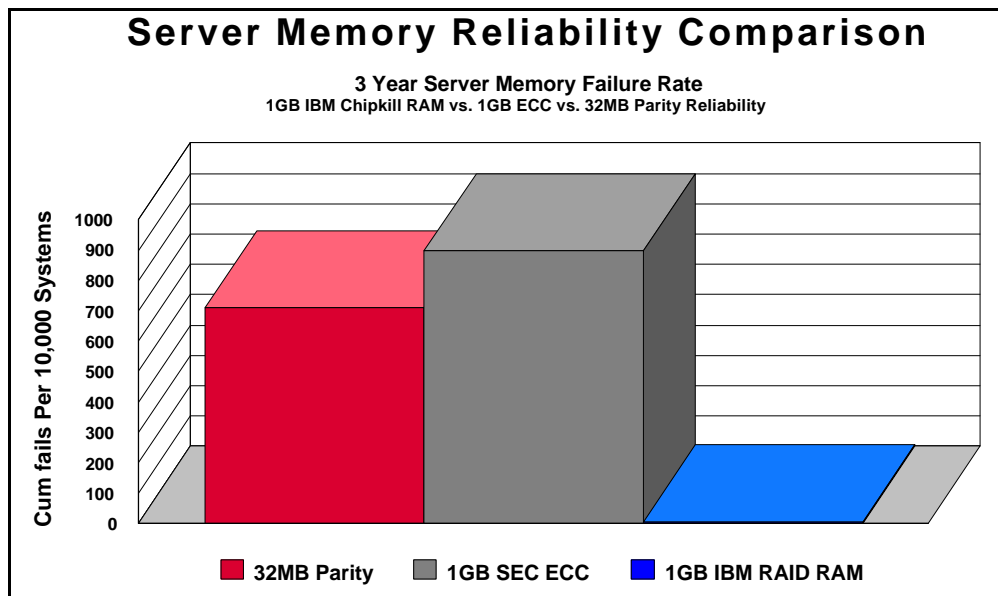
Performance

IBM engineers have designed the ECC Application-Specific Integrated Circuit (ASIC) engine for Chip-kill so that it performs the error correcting calculation without slowing access to the high-performance 50 nsec EDO memory. Thus, IBM Chipkill Memory DIMMS can be installed without modification in a Netfinity Intel Xeon processor-based server without any measurable performance degradation!

Reliability Comparison

The chart shown below illustrates the results of IBM analysis comparing the failure rates of parity, ECC and Chipkill-equipped servers. In summary, the following failure rates were identified:

- The 32MB parity memory-equipped server received 7 failures per 100 servers over 3 years.
- The 1GB ECC memory-equipped server received 9 failures per 100 servers over 3 years.
- **The 4GB Chipkill-equipped server received 6 failures per 10,000 Servers over 3 years!**




Failure rates for 1GB ECC memory are actually higher than failure rates for 32MB parity memory! The reason is that the number of 64Mbit DRAM modules required to supply 1GB of memory has increased to the point that there is a greater potential for a memory chip failure than there is for a single-bit parity error when using 32MB of parity memory. IBM Chipkill Memory technology reduces the failure rate of IBM Netfinity servers to an incredibly low 6 failures in 10,000 systems for a 3-year lifetime!

IBM Chipkill Memory Goes to Mars

The following illustration shows the Chipkill Memory DIMM with the NASA Pathfinder Mars probe. This IBM memory technology flew to Mars as part of the Pathfinder probe.

Enhancing IBM Netfinity Server Reliability



Notice the white RAID chip in the center. This chip performs the ECC calculation for each data access and stores the result in additional DRAM modules on the DIMM. Chipkill technology was used by NASA in the Mars Pathfinder probe. NASA engineers understood that this problem was real and they did not want to invest the billions of dollars to send the Pathfinder probe all the way to Mars only to have the entire mission ruined by a memory chip failure!

Summary

The most straightforward approach to providing Chipkill-correct ECC for Intel processor-based servers is to adopt IBM mainframe techniques and design the memory subsystem in a way that only one I/O bit per chip is used in each of several ECC words. This approach does present design hurdles, such as minimum memory granularity and support for wider chips (as the very common x8 DRAM, but for a system designed from scratch). These hurdles can be overcome. Because Intel processor-based server products are designed utilizing commodity components, the market has expectations of price/performance and frequency of new product introductions that preclude special “proprietary” memory controller designs. For systems designed with industry-standard components, another approach is needed to offer Chipkill capability. The only known alternative method is to deploy a standard memory subsystem framework and populate it with DIMMs that incorporate a retrofittable and plug-compatible on-DIMM Chipkill-correct ECC. This class of product, known as Chipkill Memory from IBM, provides an instantaneous upgrade from an existing ECC memory subsystem to a Chipkill-correct ECC subsystem.

The same market forces that drove yesterday’s 32MB parity memory Intel processor-based servers to include single error correction ECC memory as a checklist item will drive tomorrow’s multi-gigabyte Intel processor-based servers to include Chipkill Memory as a checklist item. In retrospect, this conclusion is not surprising based on the direction set by IBM large-system technology. If one concedes that the mainframe model was a precursor to the distributed, network-centric model of client-server computing, then it seems inevitable that the reliability and fault-tolerance features of the mainframe will soon be demanded in the Intel processor-based servers server arena.

Enhancing IBM Netfinity Server Reliability



© International Business Machines Corporation 1999

IBM Personal Computer Company
Department LO6A
3039 Cornwallis Road
Research Triangle Park NC 27709

Printed in the United States of America

1-99
All rights reserved

For terms and conditions or copies of IBM's limited warranty, call 1 800 772-2227 in the U.S. Limited warranty includes International Warranty Service in those countries where this product is sold by IBM or IBM Business Partners (registration required).

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM reserves the right to change specifications or other product information without notice.

IBM Netfinity servers and PC servers are assembled in the U.S., Great Britain, Japan, Australia and Brazil and are comprised of U.S. and non-U.S. parts.

IBM, Netfinity, X-architecture and Chipkill ECC Memory are trademarks of International Business Machines Corporation in the United States and/or other countries.

Intel, LANDesk, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation.