

**THE TECHNOLOGY BEHIND SIMSTAR™,
AN ALL-NEW SIMULATION MULTIPROCESSOR**

Ronald W. Embley
Electronic Associates, Inc.
West Long Branch, New Jersey 07764/USA

ABSTRACT

The design of the world's fastest and most advanced simulation computer, SIMSTAR, presented many engineering challenges. This paper describes the new technology employed in EAI's SIMSTAR, an all-new simulation multiprocessor. The latest innovations in implementation of both linear and digital integrated circuits, together with new techniques in subsystem circuit design, packaging, and system design make SIMSTAR possible and practical today.

SIMSTAR is a new-generation, parallel simulation multiprocessor designed specifically for analysis of dynamic systems. Consisting of the most advanced linear and discrete circuit technologies, SIMSTAR is a high-performance, automatic computational unit. Capability is provided to accurately model the complex engineering processes, including all their non-linear, discontinuous and stochastic characteristics.

The key innovations, some of which are presented in detail, are:

- An efficient, three-level Connection Matrix (CMOS implementation) and a smart matrix connection search algorithm.
- An efficient Parallel Logic Processing Unit for high-speed sequential and combinational logic generation.
- A 32-bit system-integrated Digital Arithmetic Processor with removable cartridge disc mass storage unit.
- A 16-bit system-integrated Local Control Processor (68000 μ P-based) as an intelligent setup and control interface.
- A memory-mapped (shared-memory) interprocessor digital interface.
- The mathematical computing block concept providing multi-functional capabilities for parallel computing components.
- System-integrated Automated Test Equipment (ATE) with Automated Diagnostics for mathematical computing blocks, connection matrix, and parallel logic processing unit, both static and dynamic.
- An Autobalance System for time and temperature drift correction, ensuring high accuracy in a consistent manner.
- A 3000 Point Solid-State Readout (diagnostic and problem solution) system with an Autoranging ADC and automatic offset correction.
- An Active Ultra-High Quality Ground Preservation System.
- Extended Range (pseudo-floating point) digitally set coefficient units and multipliers.
- Automatic Noise/Oscillation Detection System.
- System-integrated analog signal line translator amplifiers for fidelity connection with external equipment.
- Deglitcher circuits, greatly reducing transients from on-line electronic switching.
- Wide bandwidth, active fixed function generators.
- Digitally set arbitrary function generators with high-speed on-line function data update.
- Compound Operational Amplifiers yielding 60MHz bandwidth PLUS excellent dc. offset and drift characteristics.

AN EFFICIENT THREE-LEVEL CONNECTION MATRIX (CMOS IMPLEMENTATION) AND A SMART MATRIX CONNECTION SEARCH ALGORITHM

A totally automatic interconnection means was required for SIMSTAR to connect over 300 high-performance linear computing devices and related interface channels, all of which could be operating simultaneously. Crosstalk, noise and phase shift errors were all to be minimized in order to provide interconnections electrically transparent to the system.

The resultant approach takes the form of a crossbar switch with high reliability designed in by utilizing ICs employed in telephone switching networks where outstanding MTBF is mandatory.

The Analog Math Block Connection Matrix is a solid-state buffered output switch array which allows any Mathematical Computing Block (MCB) output to be connected to any MCB input or inputs. The matrix also provides switchable analog input and output signal lines for connection to other SIMSTAR consoles and external peripherals and/or equipment. The analog connection matrix is implemented using a three-stage Clos [1] network which greatly reduces the number of switches required in a single-level matrix (Figure 1).

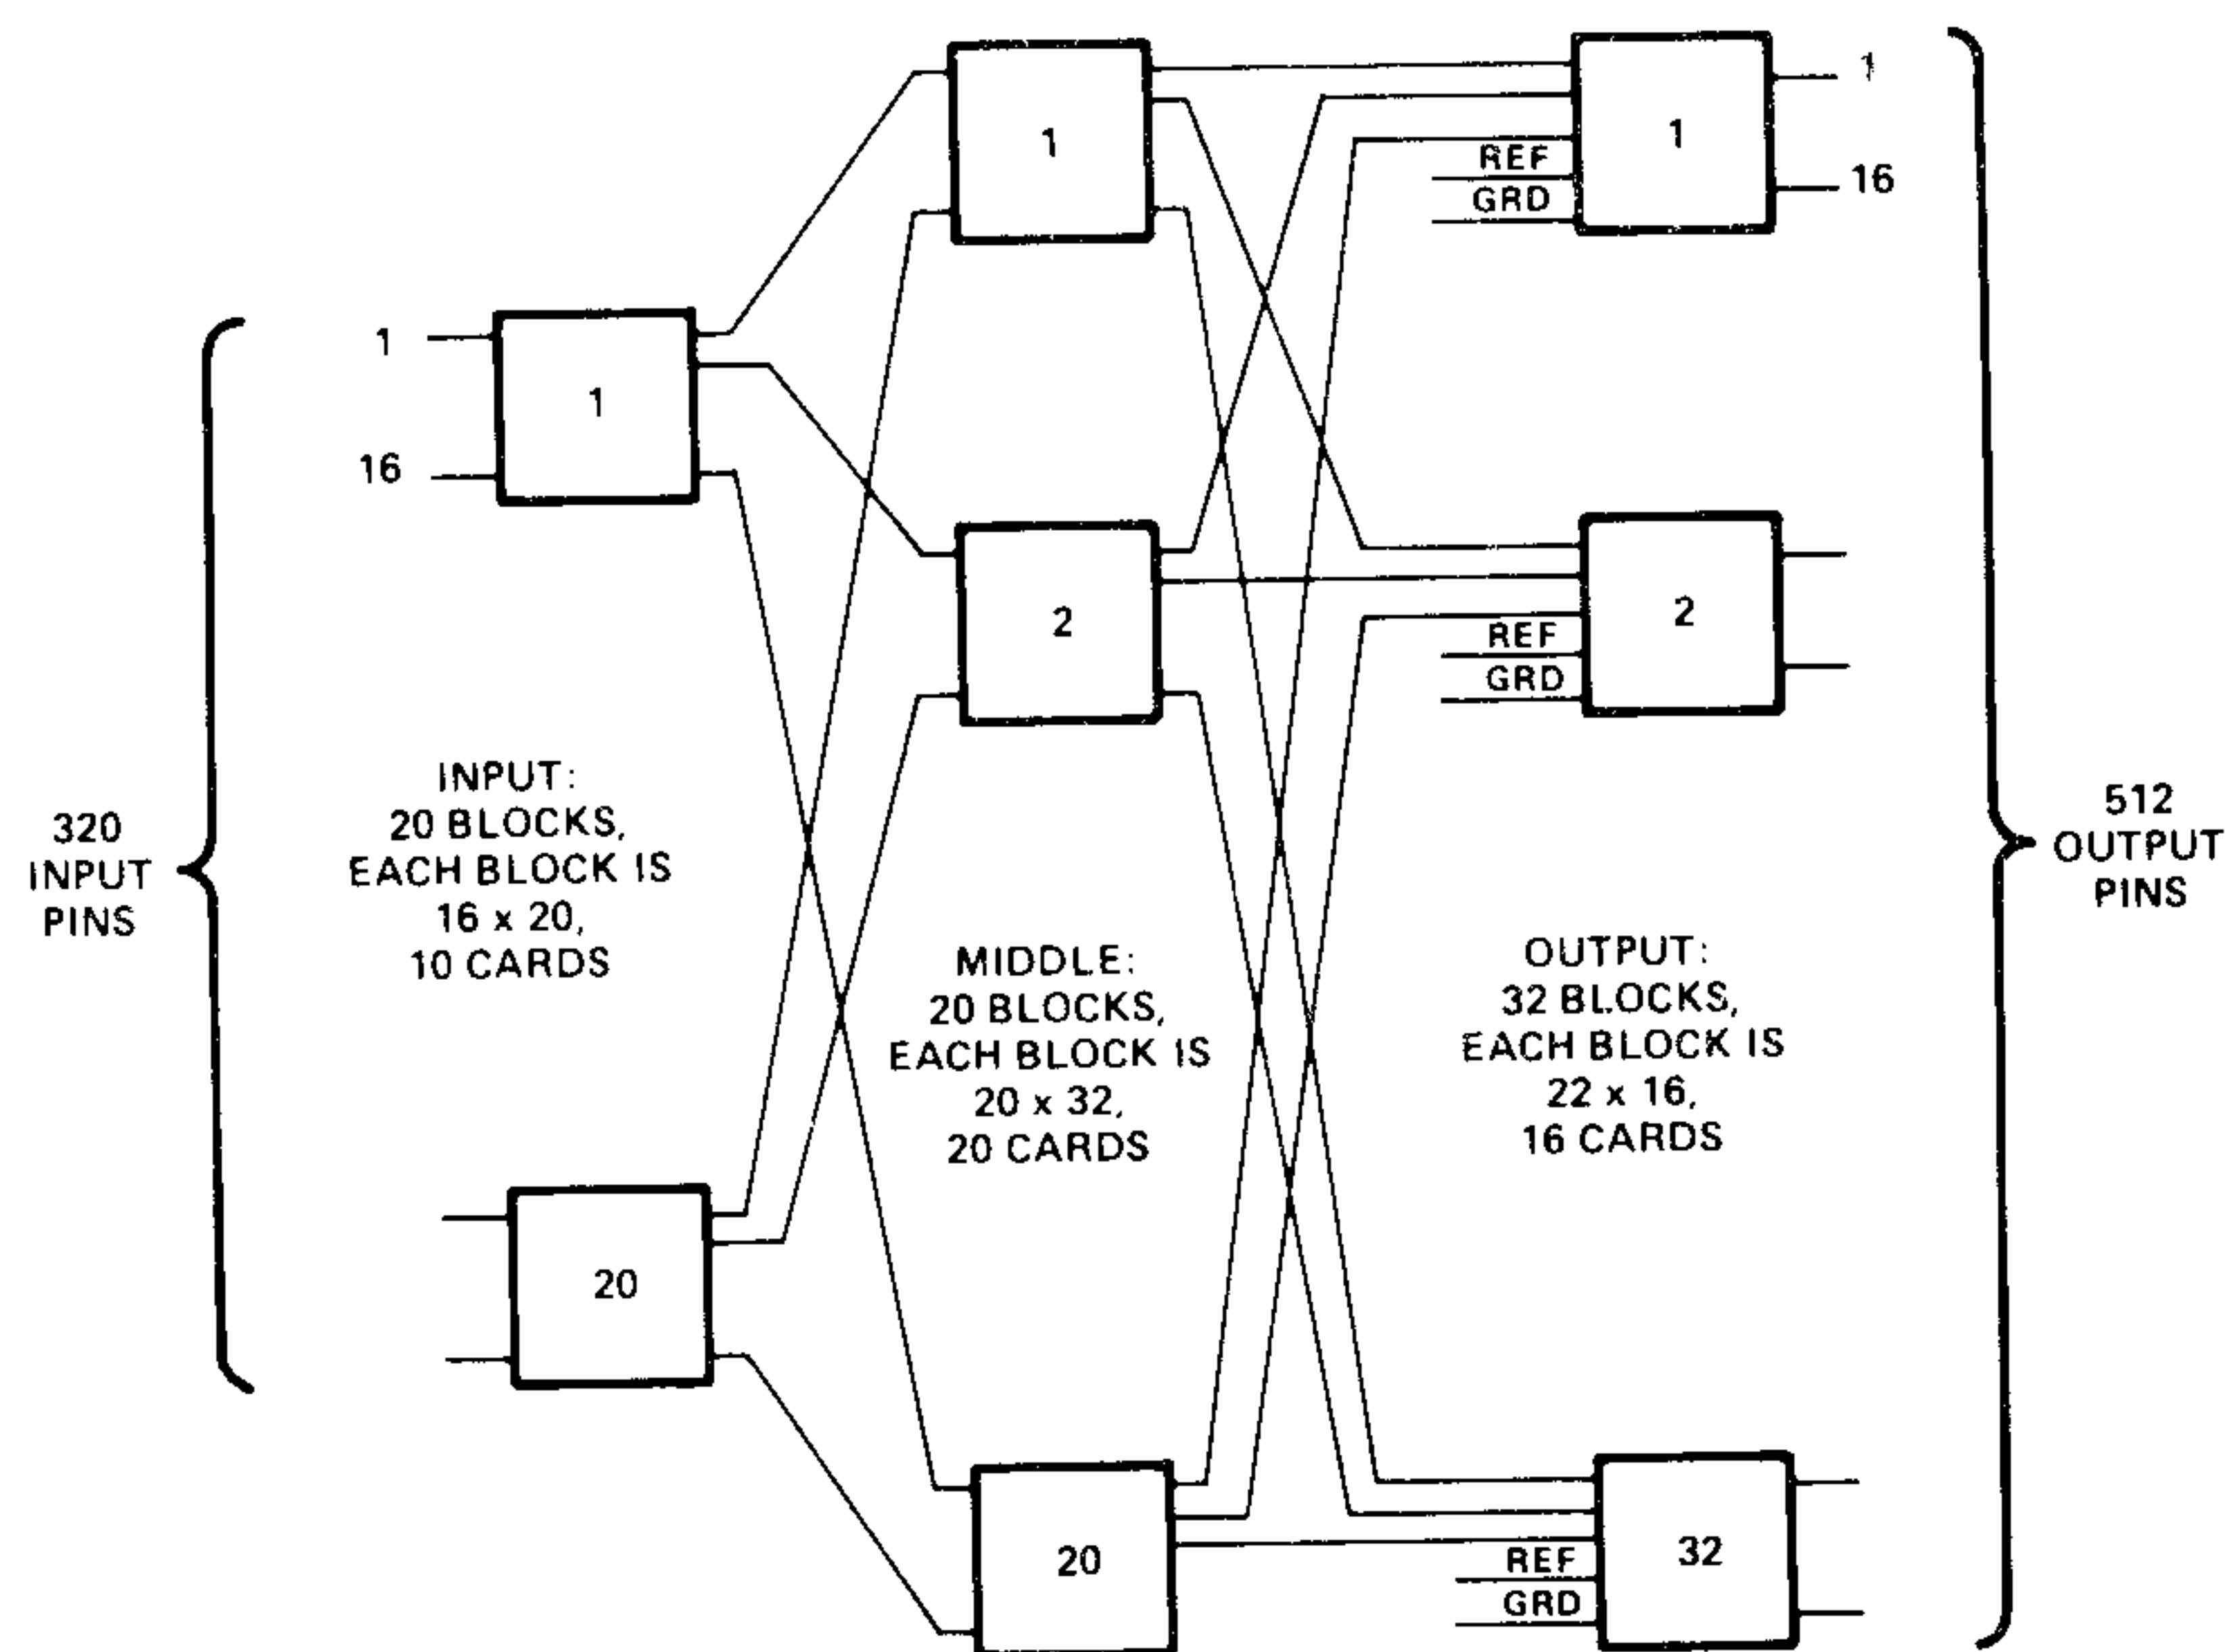


Figure 1 The SIMSTAR 3-Stage Connection Matrix

This is a 320 x 512 matrix which, if implemented as a single-level crossbar switch, would require 163,840 switches. Using the three-level Clos network reduces the number of switches to 29,440. Use of this implementation results in a practical 320 x 512 analog connection matrix packaged on 46 printed circuit boards fitting into two SIMSTAR card files. These cards include an 832 point readout system and 320 pole double-throw diagnostic switch (Figure 2) for automated diagnosis of switch failures.

It is interesting to note that if the straightforward one-level crossbar switch had been implemented, three additional racks would be required by the SIMSTAR Multiprocessor system to house it!

High-density, solid-state CMOS voltage switch integrated circuits are the heart of the Math Block Connection Matrix. The RCA 22100 LSI/CMOS switch chip used combines a 4 x 4 array of crosspoints (transmission gates) with a 4 to 16 line decoder and 16 latch circuits for control memory. The CMOS FET transmission gates are of large geometry providing 75 ohm on impedance resulting in low crosstalk between switches and low phase shift through the three stages of switches, meeting all of SIMSTAR's targeted requirements.

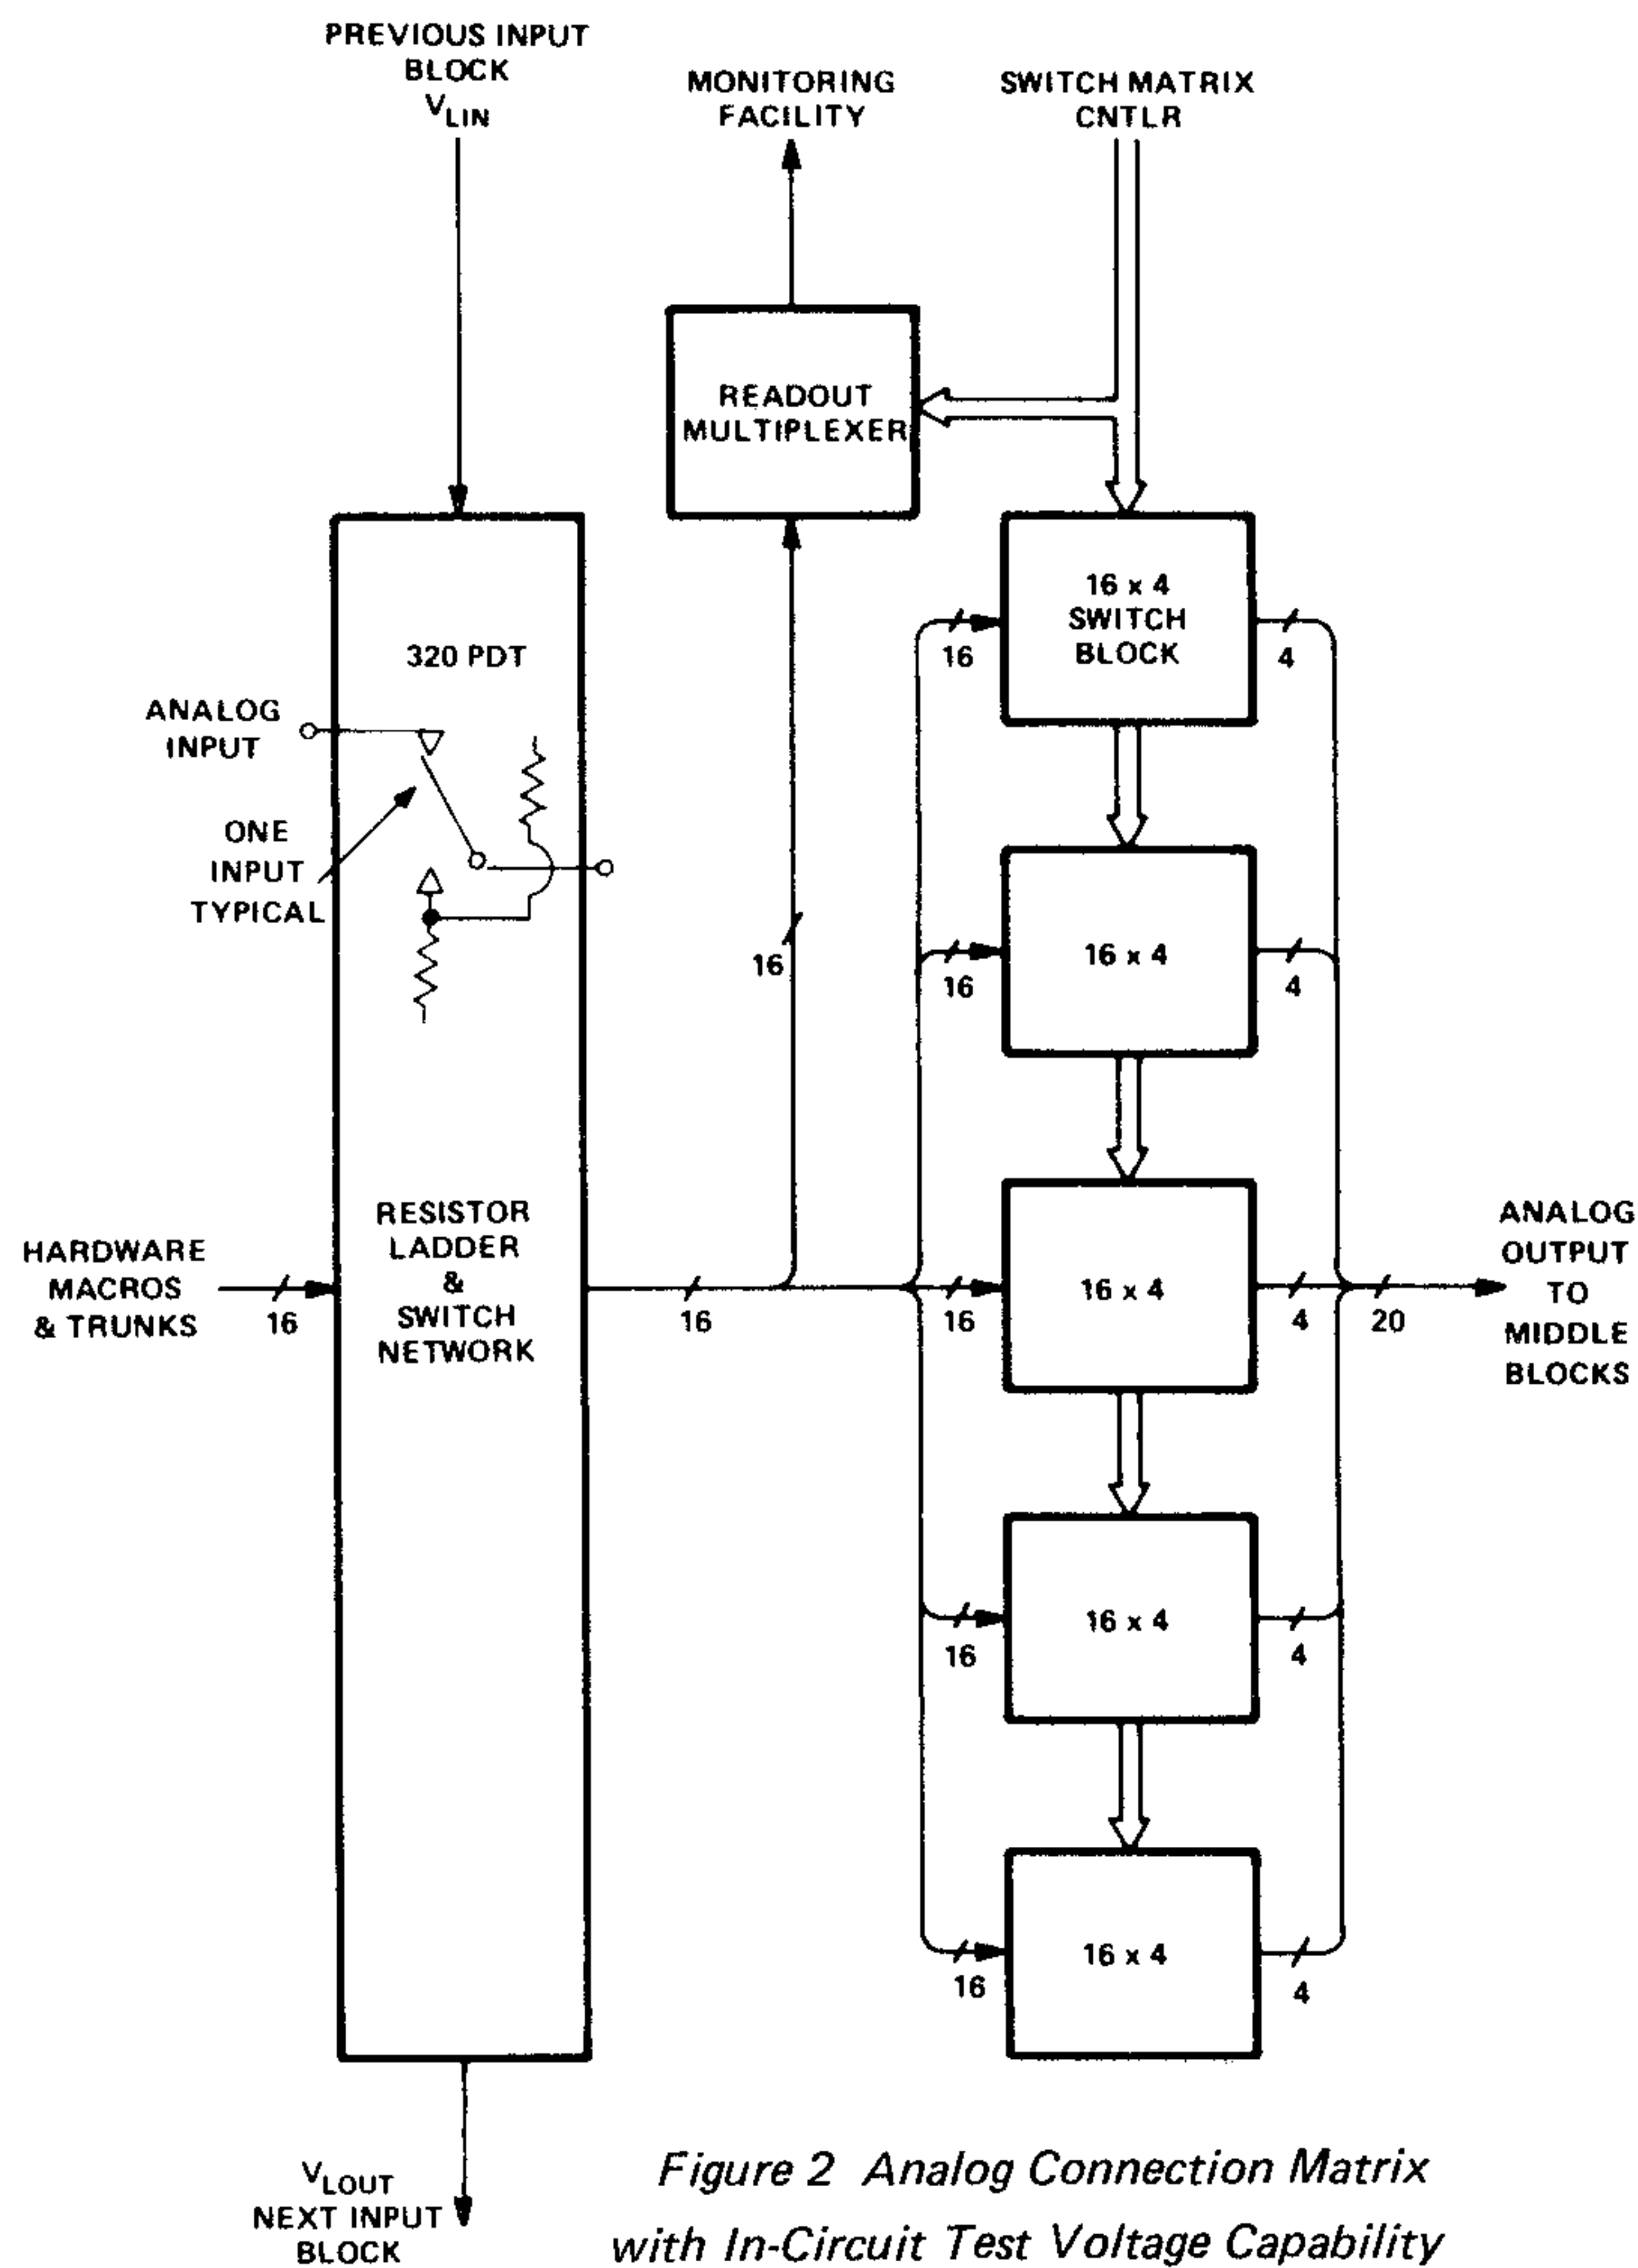


Figure 2 Analog Connection Matrix with In-Circuit Test Voltage Capability

A smart search algorithm has been developed and simulated which enables the optimum middle blocks to be chosen to satisfy a given practical set of input and output connections without signal blocking.

This network routing problem was analyzed and solved as follows:

Given a set of connections to be established on the matrix (in telephone terms, a set of "calls to be set up"), how can one find parallel paths for all the connections without encountering blocking (short circuits)? There is a well-known solution [1] for the case of fanout-free calls, but in the presence of fanout, no efficient algorithm was known [2] [3].

The problem belongs to the class known as NP: Given a proposed set of paths, it is easy to check its validity, but the number of possible candidates rules out an exhaustive search [4]. For SIMSTAR, the number of candidates is $20^{512} \approx 1.3 \times 10^{666}$. The exponent alone shows that the problem is intractable.

The solution required a combination of mathematical and engineering techniques. First, a recursive tree-search algorithm was developed which exhaustively examined all candidates; this worked well for small matrices (e.g., 20×32), but required excessive computing time for larger systems.

In the "engineering" phase, heuristics were developed to eliminate large blocks of candidate solutions at a time [5]. These were evaluated on a set of randomly-generated worst case applications.

The original version of the tree-search found correct paths in only 50% of the test cases, even when allowed to run for over two hours. The final version, with sophisticated heuristics, succeeded in all cases. The median search time was three minutes. Since these were artificially selected worst cases, the time required for real applications can be expected to be substantially less.

AN EFFICIENT PARALLEL LOGIC UNIT (PLU)

A Parallel Logic Unit has been designed which replaces the manually connected parallel logic on predecessor systems. It is a processor which provides real-time monitor and control, sequential logic, combinational logic, external signal line capability, and is high-level language programmable. One hundred sixty input signals can be processed in parallel and the results routed to 320 specified outputs in $1\mu s$.

This high-speed one-bit processor design utilizes a time-sliced approach greatly reducing the number of internal ICs required. Eighty ICs do the job of 1120 in the Boolean Function Generation (BFG). Three hundred twenty Boolean Functions of four variables are generated in each time slice. This results in an economic design, which is completely housed in one SIMSTAR Multiprocessor card file.

The Boolean Function Generator components of the PLU are shown in the block diagram, Figure 3. The memory and 2:1 MUX/registers form a loop where sections of the memory are examined during each of up to fourteen internal time states. This loop is iterated upon to perform the signal switching and Boolean Function Generation necessary for the outputs. The fourteen maximum internal states (less than 100ns each), including one state for synchronizing the inputs and one state for initial condition, comprise one complete cycle. Normally, the first four stages are used to switch the inputs to the fifth stage and the combinational logic is done in the fifth through eleventh stages, the actual number of stages depending on the complexity of the logic to be performed. The last stages are then used to switch the BFG outputs to the proper PLU output. Logic can be performed in any of the stages in case of blockage or need for an expanded function but normally all the logic functions will be created in the fifth through seventh stages, reducing the cycle time to 1.2 microseconds.

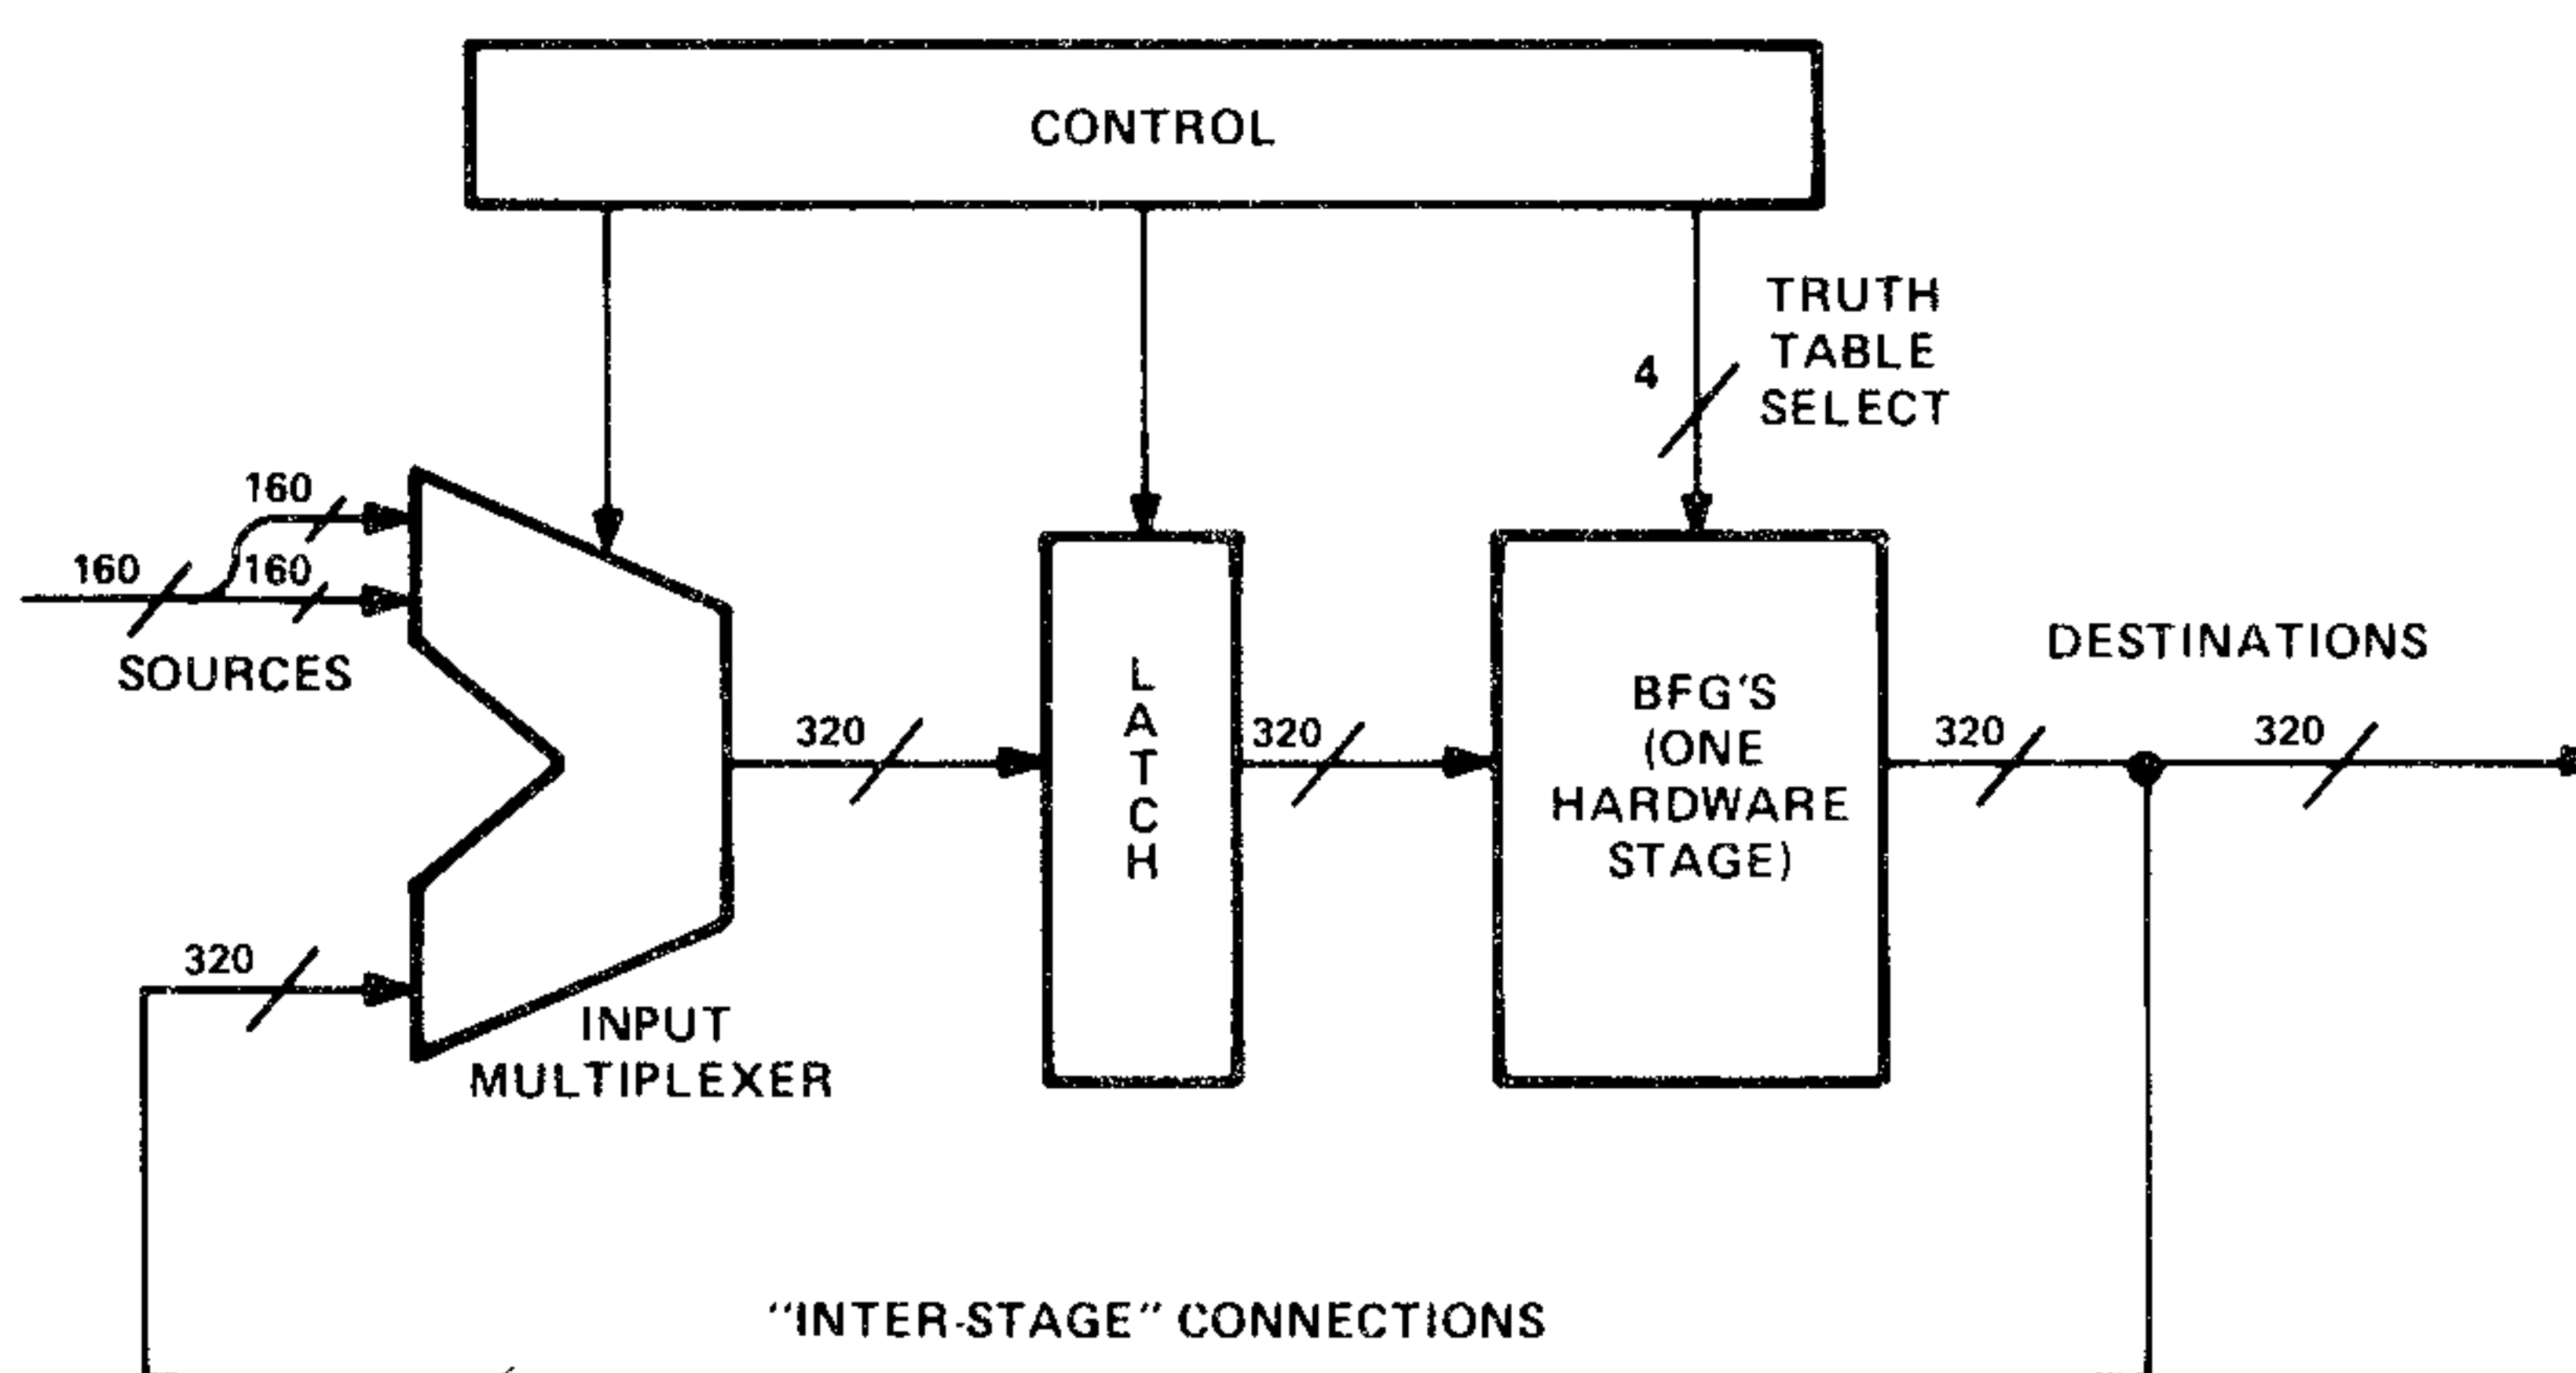


Figure 3 Multi-Stage BFG Network with Reduced Hardware

The resultant PLU subsystem provides the SIMSTAR Multiprocessor system with a seemingly inexhaustible logic processing capability. If needed in a simulation application, the equivalent of over 4000 4-input gates can be represented and updated in less than 2 microseconds.

A SYSTEM-INTEGRATED DIGITAL ARITHMETIC PROCESSOR (DAP)

A state-of-the-art DAP is built into SIMSTAR, providing high-efficiency, economical simulation of the slower processes being simulated and memory-mapped setup and control of the high-speed Parallel Simulation Processor (PSP).

The DAP design provides the following features:

- True 32-bit CPU
- Built-in single and double floating point
- Optional floating point accelerator (1.6-2.2 μs FP)
- Up to two megabytes of memory

- Vectored interrupts
- Shared memory with local control processor and host processor
- Memory management for multiprogramming
- Removable cartridge 80-MByte disc mass storage.

A SYSTEM-INTEGRATED LOCAL CONTROL PROCESSOR (LCP)

The LCP is designed to setup, control and maintain the PSP. The LCP is 68000 μ P-based having a 16-bit data bus and a 24-bit address word. Local memory (256KB) stores all the data sent to the PSP. Complex data transfers from DAP or HOST processor to the PSP are controlled simply by the LCP.

The LCP performs the following PSP functions:

- Initialization
- Macro Inventory Keeping
- Data Format Conversion
- Maintenance - Automated Diagnostics, Autobalance, Temperature and Power Voltage Measurement, Noise and Oscillation Detection
- Problem Solution Readout - Multiplexer Address Selection and ADC Gain Ranging

SYSTEM-INTEGRATED AUTOMATED TEST EQUIPMENT (ATE)

Because of the unique multi-element parallelism of SIMSTAR, a sophisticated testing method was deemed necessary. Maximum possible operational up time along with deterministic performance were the chief design goals. An automated test system has been designed into SIMSTAR providing not only ease of maintenance, but assurance of all key hardware performance specifications in the PSP. Soft, as well as hard failures, such as excessively drifting op-amps, are identified and the faulty unit can be removed from inventory before subtle errors in computation can occur. Faulty units are logged and removed from approved-status inventory automatically. The faulty unit is then automatically and electronically replaced by another of the same type from the approved-status inventory prior to the next problem setup. Efficient diagnostic algorithms have been developed which can pinpoint a single switch failure in the block connection matrix or a single bad bit in a PLU ram. The board and chip are then identified for the maintenance technician. Fault indicator LEDs are provided on all hardware macro computing boards for maintenance convenience and fast recognition during repair.

The test equipment integrated into the SIMSTAR system is:

- Autoranging ADC (26-bit resolution)
- 3000-Point Multiplexed Readout Selector
- Precision Programmable Gain Device
- Precision Error Detection Amplifier
- Precision Sign Changing Amplifier
- Peak Error Detector Amplifier
- Programmable Frequency and Amplitude-Stabilized Oscillator
- Block Connection Matrix used to connect the unit under test (UUT), in the specified test circuit configuration
- Programmable precision voltage sources (16 bits)
- Safe operating temperature and in-tolerance power supply measurement circuits
- Automatic System Power Shutdown Circuit
- A 320-pole, Double-Throw Switch with Precision Resistor Ladder Network

This test equipment hardware, together with comprehensive software diagnostic routines, comprises the automated test system. A block diagram of the automated test system is shown in Figure 4. In this example, a macro unit (UUT) is connected to the diagnostic test unit (DTU) via the analog block connection matrix. The DTU shown in the AC mode provides a sine wave test signal (programmable frequency) to both the UUT input and one input of the error detection amplifier (EDA). The other input of the EDA is connected through the precision sign changing amplifier and gain device to the UUT output via the Block

Connection Matrix. The output of the UUT then is compared with its input, with any sign change or gain magnitude corrected for by the DTU. An error signal is generated at the output of the EDA which is directly proportional to any error or distortion in the UUT. The peak detector in this case captures the total instantaneous dynamic error (TIDE), the vector sum of phase shift error and amplitude error. The peak detector output error signal is then digitized by the ADC, sent to the LCP where it is compared against predetermined specification limits and the result entered in the maintenance log.

With the thoroughness of the resultant hardware/software design, a new level of user confidence in SIMSTAR operation is possible.

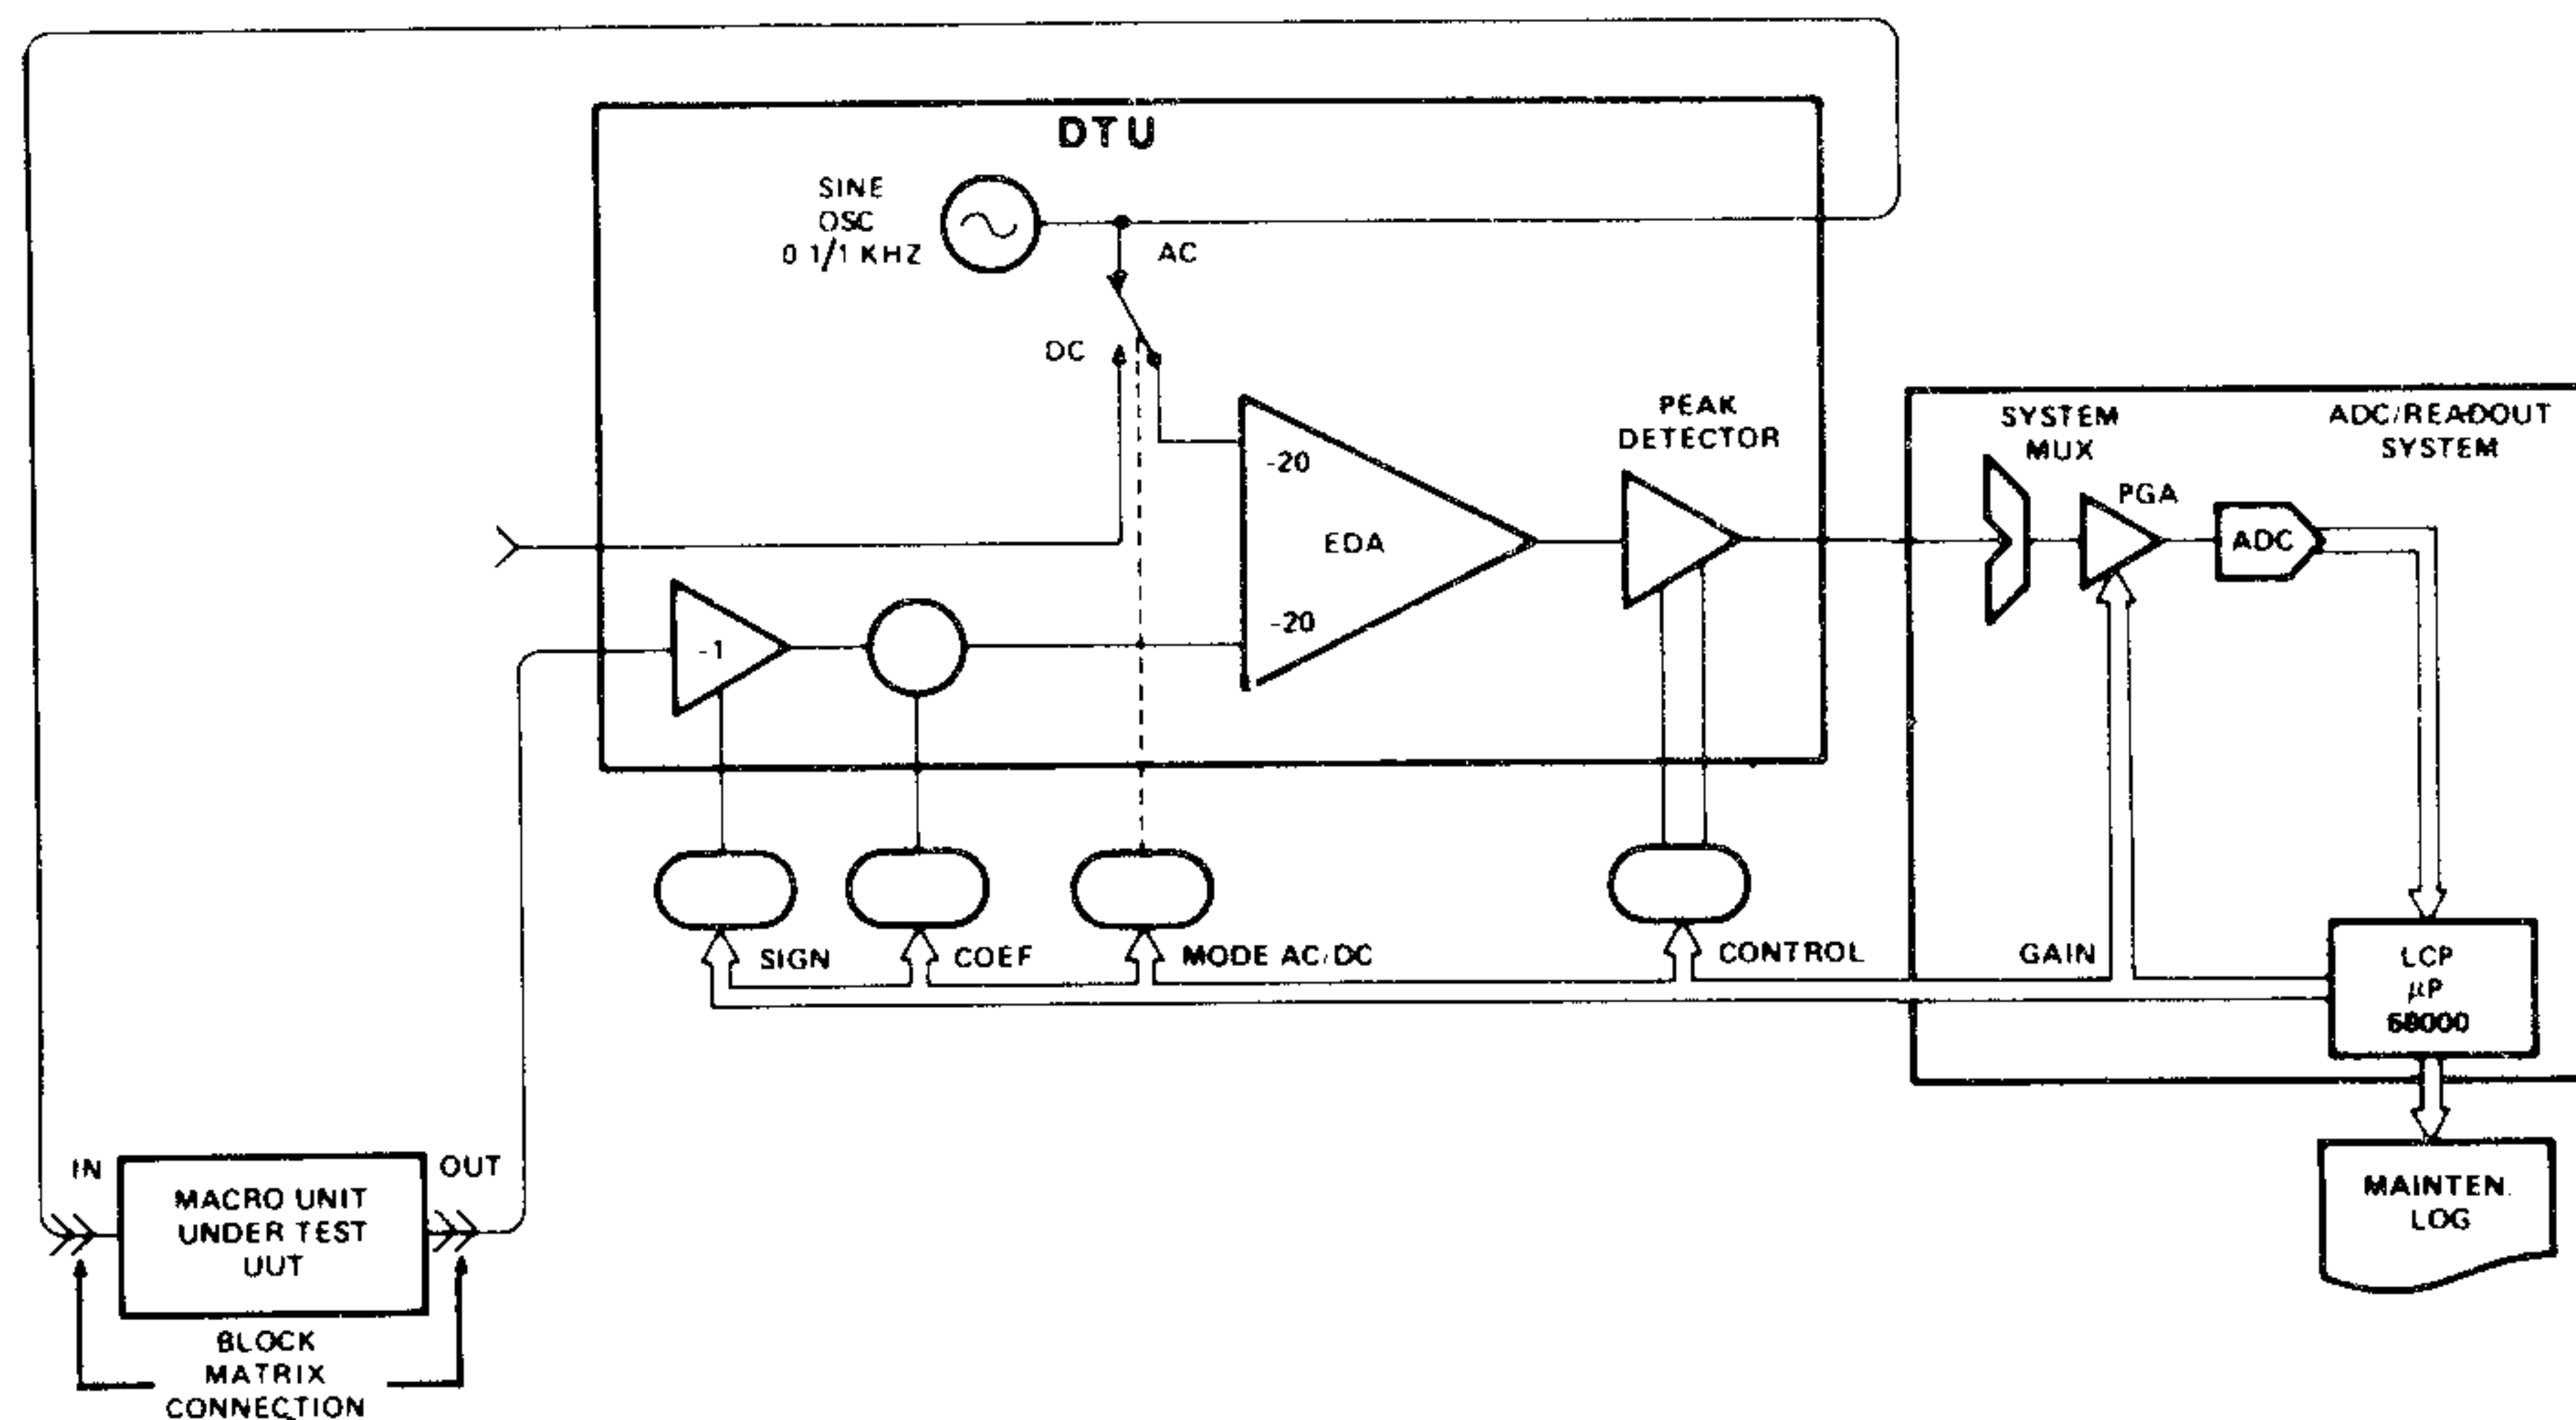


Figure 4 Block Diagram of the Automated Test System (shown in the Macro AC Test Mode (0.1/1 KHZ Tide))

AN AUTOBALANCE SYSTEM FOR TIME AND TEMPERATURE DRIFT CORRECTION

An automatic means for electronically nulling any offset voltage present in critical op-amp circuits has been designed into the SIMSTAR system. This autobalance system completely eliminates the tedious time-consuming task of manual nulling and the simulation inaccuracies caused by drifting op-amps associated with predecessor systems. Not only are all critical amplifiers nulled (to zero ± 5 microvolts) as a function of elapsed time and temperature change, but records are kept on the amount of nulling required for a given amplifier and, if a predetermined limit is reached, that amplifier is flagged as an excessive drifter which can then be replaced before it can effect the accuracy of the simulation.

An autobalance is always performed after a problem load, after a macro configuration change or at the request of the DAP. An autobalance will be performed unless optionally inhibited by the user, after a problem restore, if a change in temperature exceeds 1°C after the last autobalance (temperature is measured via four precise solid-state sensors at ten minute intervals); and if an autobalance has not been executed within the last eight-hour period. An autobalance is never initiated during the RUN mode of SIMSTAR.

The autobalance hardware consists of over 600 8-bit correction DACs, each being connected to critical op-amp circuits as shown in Figure 5. Referring to Figure 5, the simplified autobalance algorithm can be understood:

1. Select UUT output for readout via the readout system.
2. Set all connection matrix outputs to zero.
3. Set correction DAC to zero output.
4. LCP receives and records output offset.
5. Correction DAC is loaded with the corresponding nulling value.
6. Repeat steps 4 and 5 again, output is nulled.

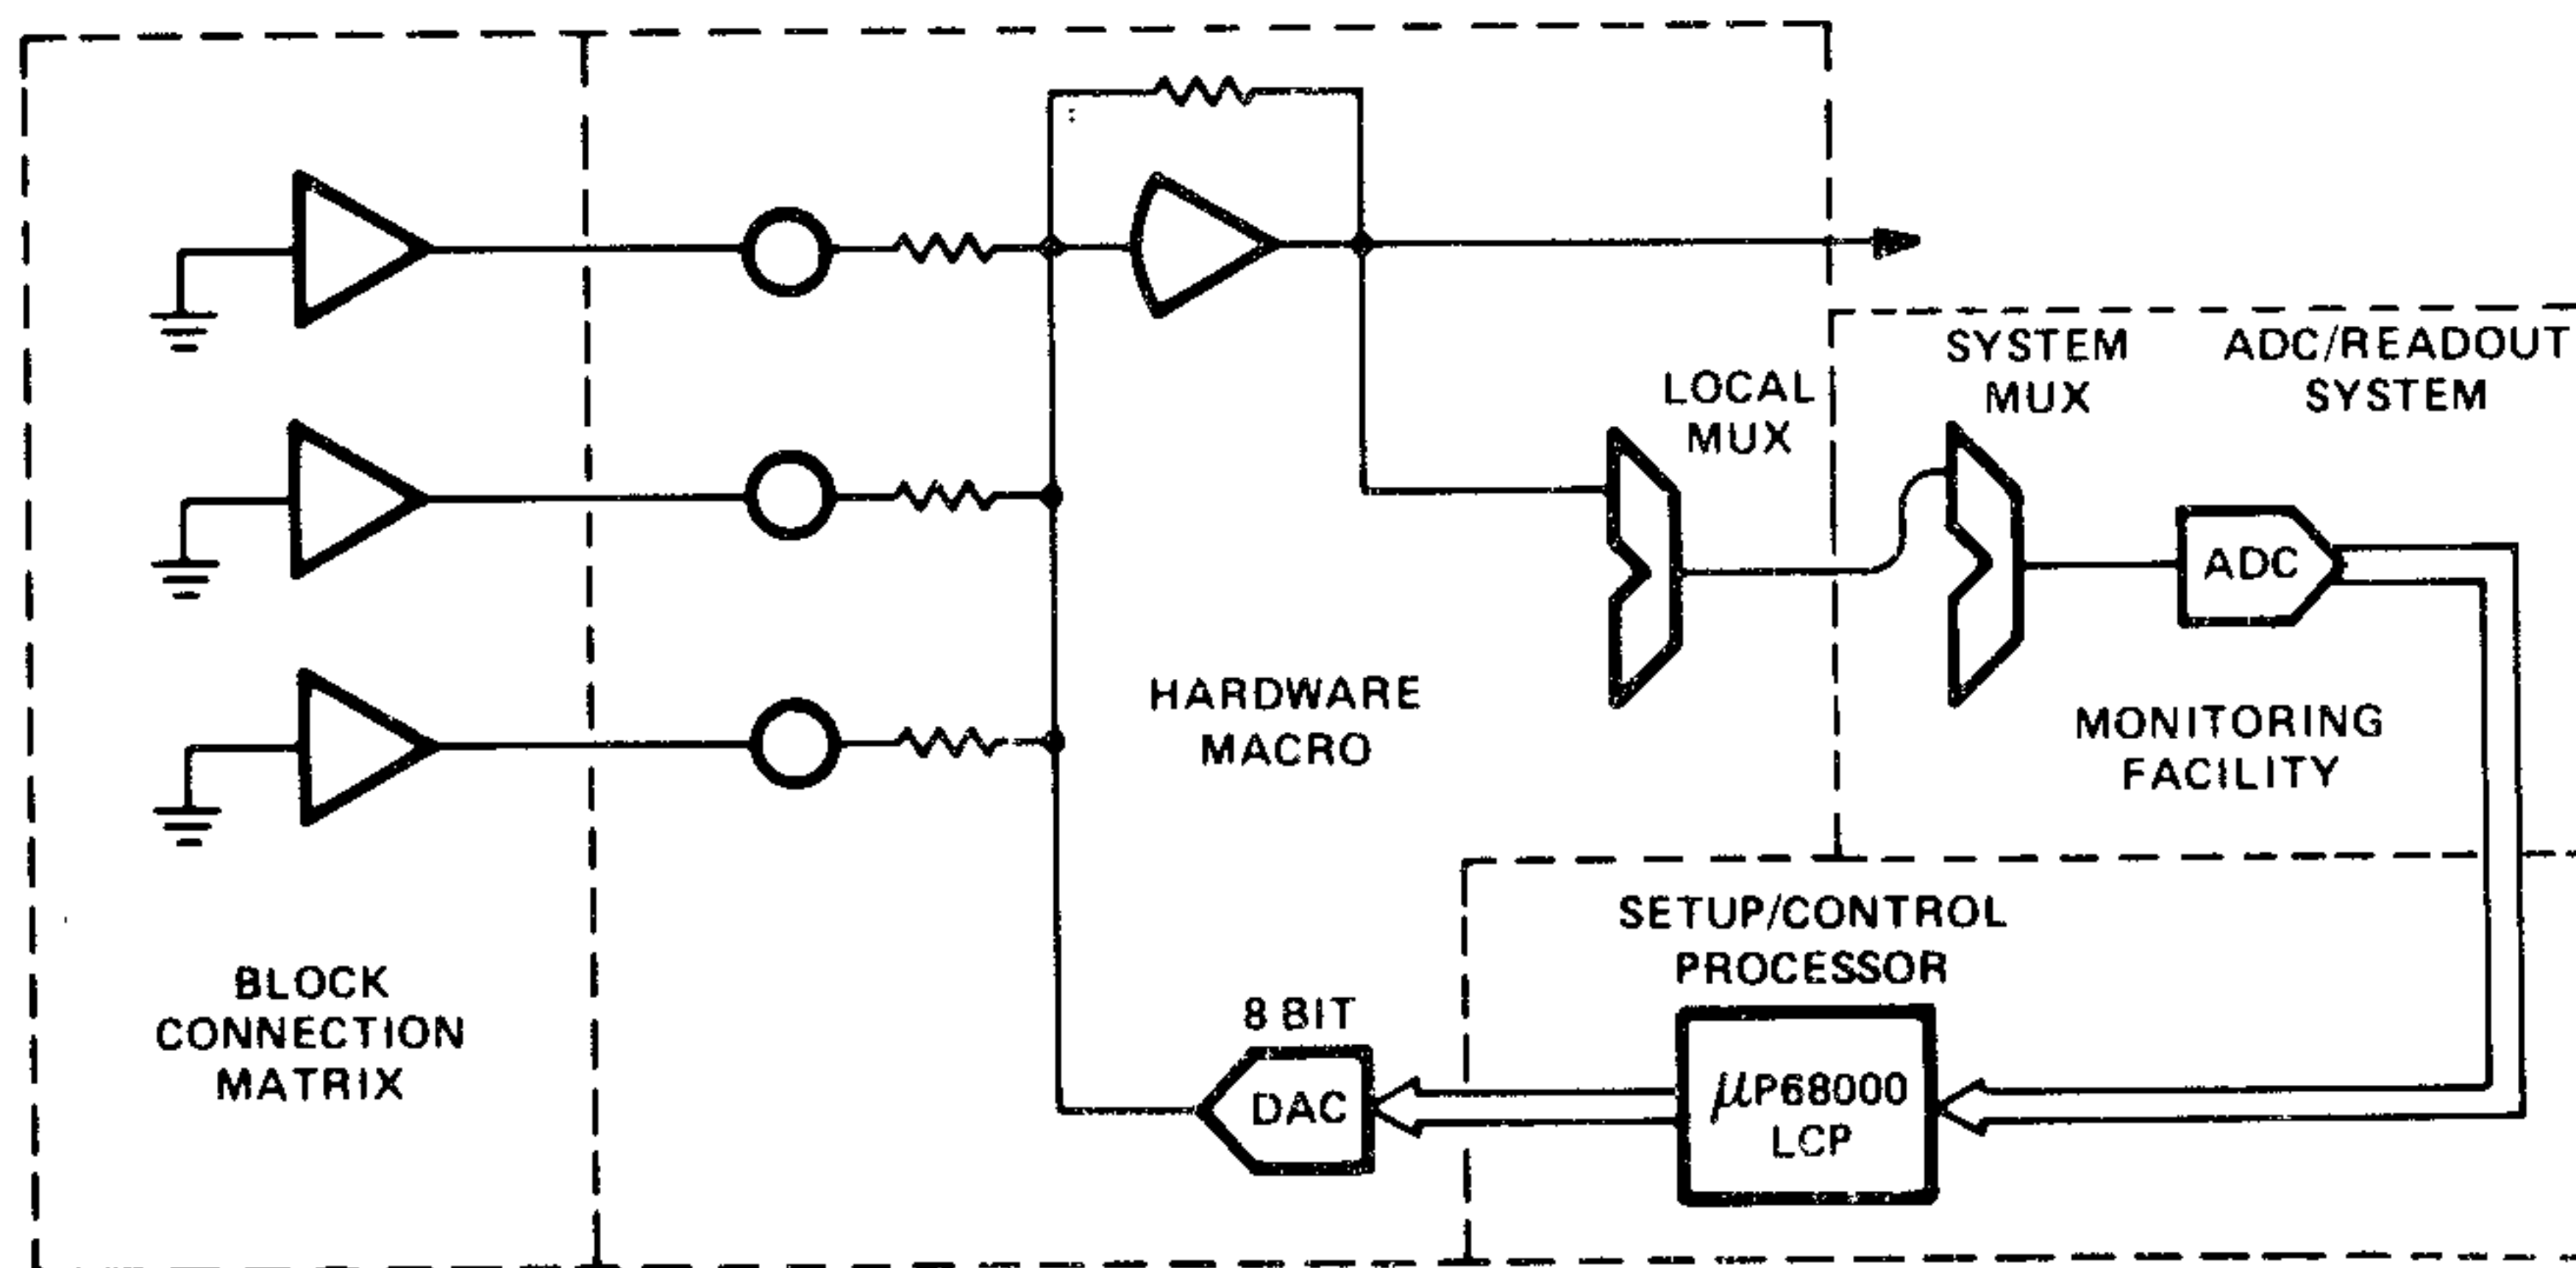


Figure 5 System Block Diagram of the Autobalance System

AN ACTIVE ULTRA-HIGH QUALITY GROUND PRESERVATION SYSTEM

A high-quality (HQ) grounding system was needed for SIMSTAR that would maintain extremely small potential differences (micro-volts) between all of the over 200 analog computing macros in respect to the system central HQ ground point. Figure 6 shows a diagram of the "star point" ground system employed in SIMSTAR which eliminates ground loops, but IR drops in the individual ground distribution wires still produce intolerable DC offsets. A unique solution to this problem has been implemented in SIMSTAR. An active system has been developed which reduces the HQ ground current flowing from each macro unit by four orders of magnitude, hence essentially eliminating the proportional IR drop in the distribution wires (see figure 7a). This definitely minimizes error contributions from ground sources.

The active ground circuit (AGC) (see Figure 7b) is comprised of a low drift op-amp connected in the voltage follower mode to produce a low impedance current source (or sink) whose output potential is maintained at zero voltage (virtual ground). The net result is that the normal tens of milliamps of current flowing to or from a computing component ground terminal is steered to the insensitive $\pm 15V$ busses, while the HQ ground is unaffected.




Figure 6 Star Point Ground System

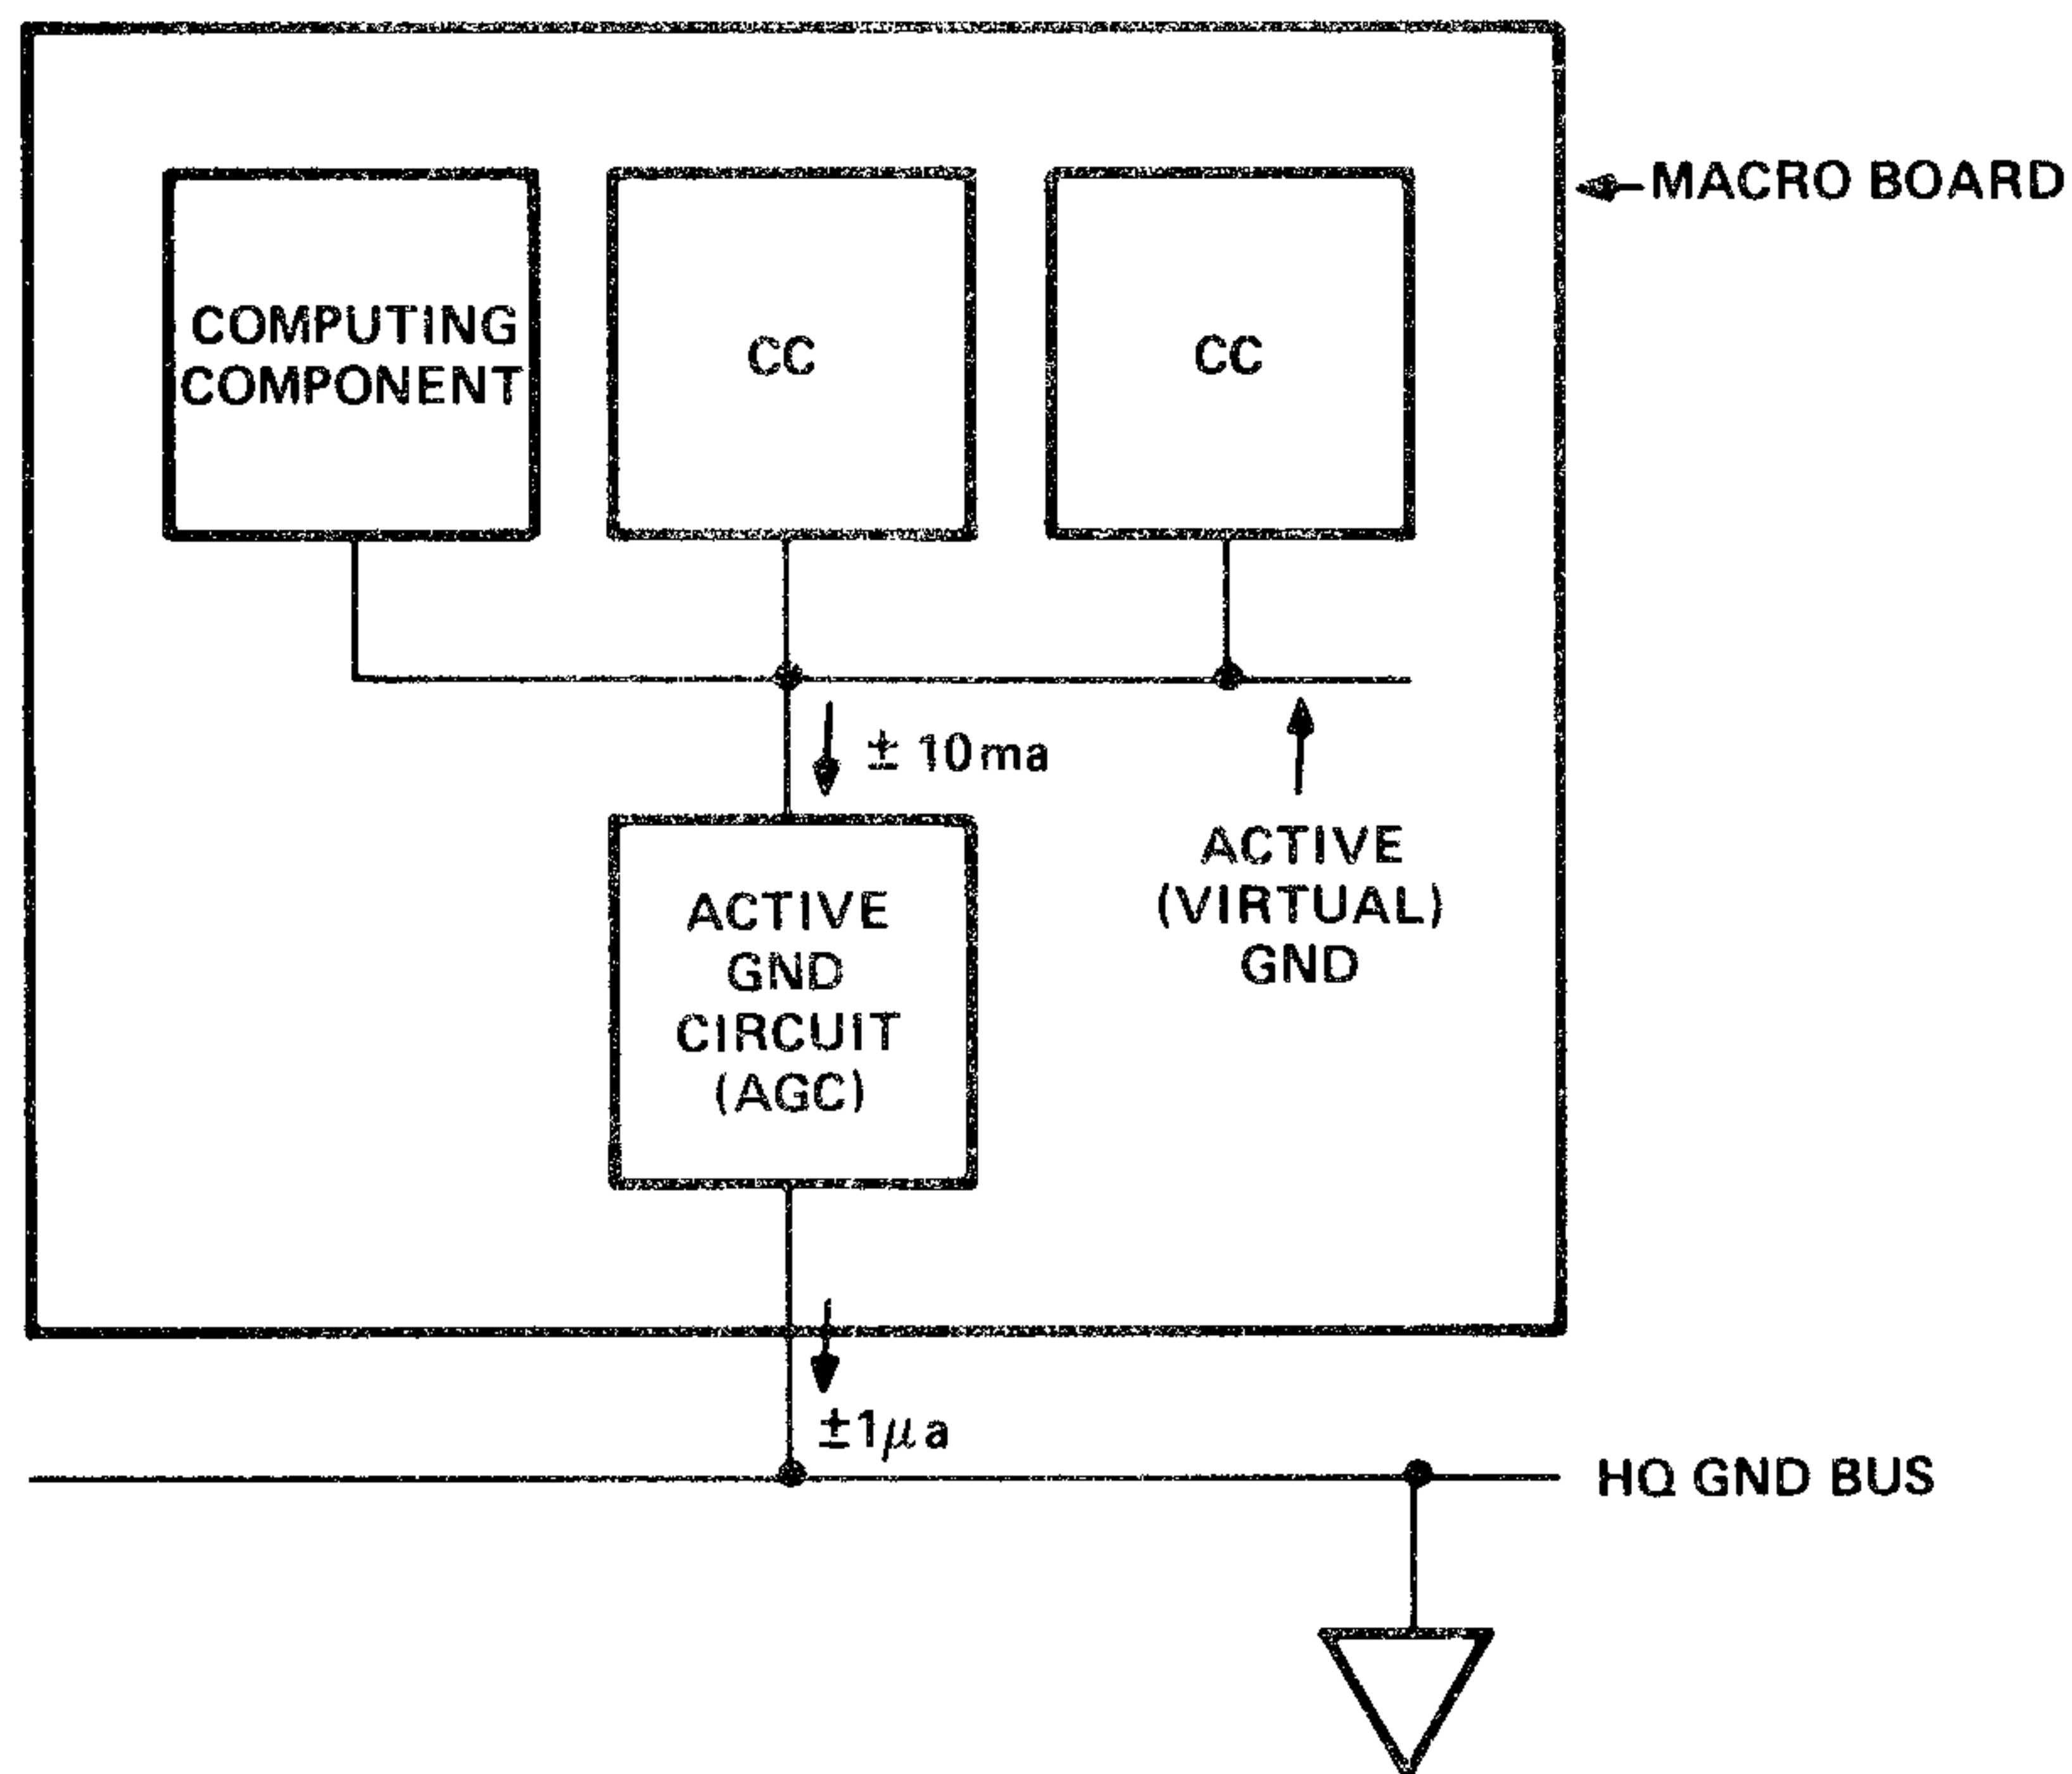


Figure 7a Functional Block Diagram

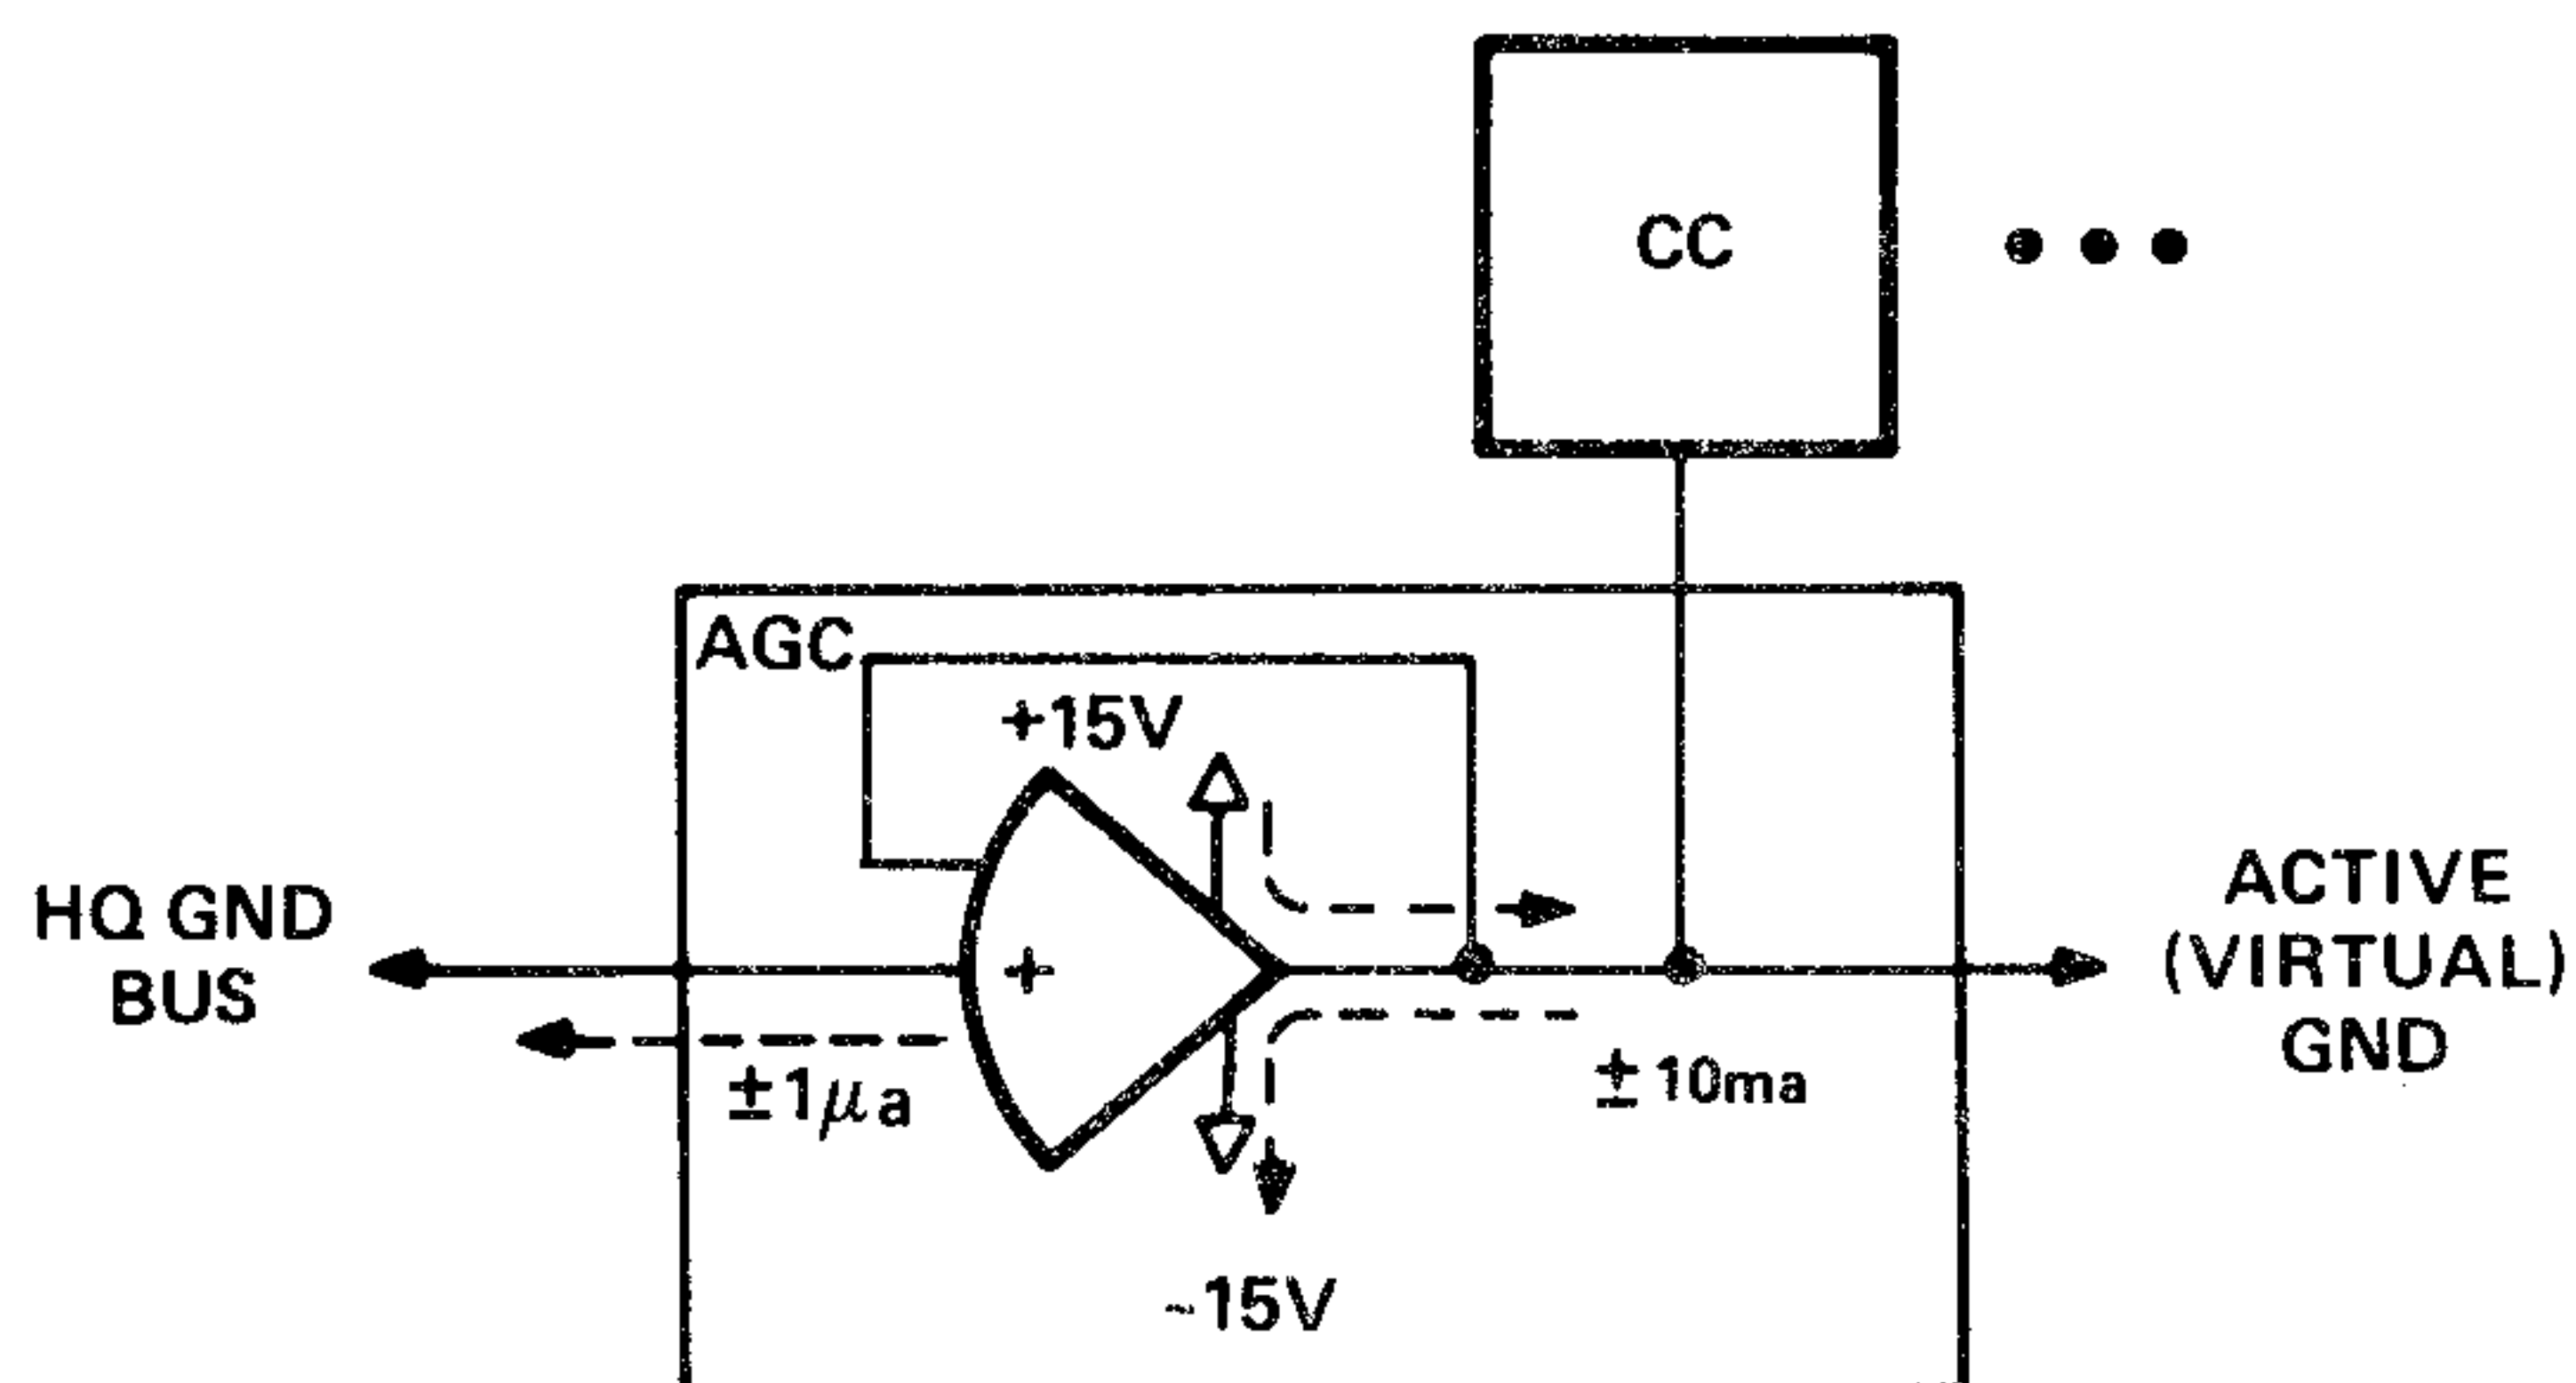


Figure 7b Simplified Schematic of the Active Ground Circuit

Figure 7 Active High Quality Ground Preservation System

EXTENDED RANGE (PSEUDO-FLOATING POINT) DIGITALLY SET COEFFICIENT UNITS AND MULTIPLICATION DEVICES

A factor of ten improvement in the useable range of SIMSTAR over predecessor systems has been realized by incorporating automatic gain changing (local rescaling) circuits within each digitally set coefficient unit (DSCU) and analog multiplication devices.

The extended range DSCU design (see Figure 8) provides both improvements in accuracy and resolution for small coefficient values (below 1/4). The DSCU functions, together with the LCP, to utilize the high order bits (most accurate portion) of the MDAC even for small coefficient settings. This autoranging or autoscaling technique results in an effective overall range of 18 bits (+ SIGN) with 16- and 18-bit resolution below settings of 1/4 and 1/16, respectively. Accuracy also improves by factors of 4 and 16, respectively, tending towards a percent of output error characteristic.

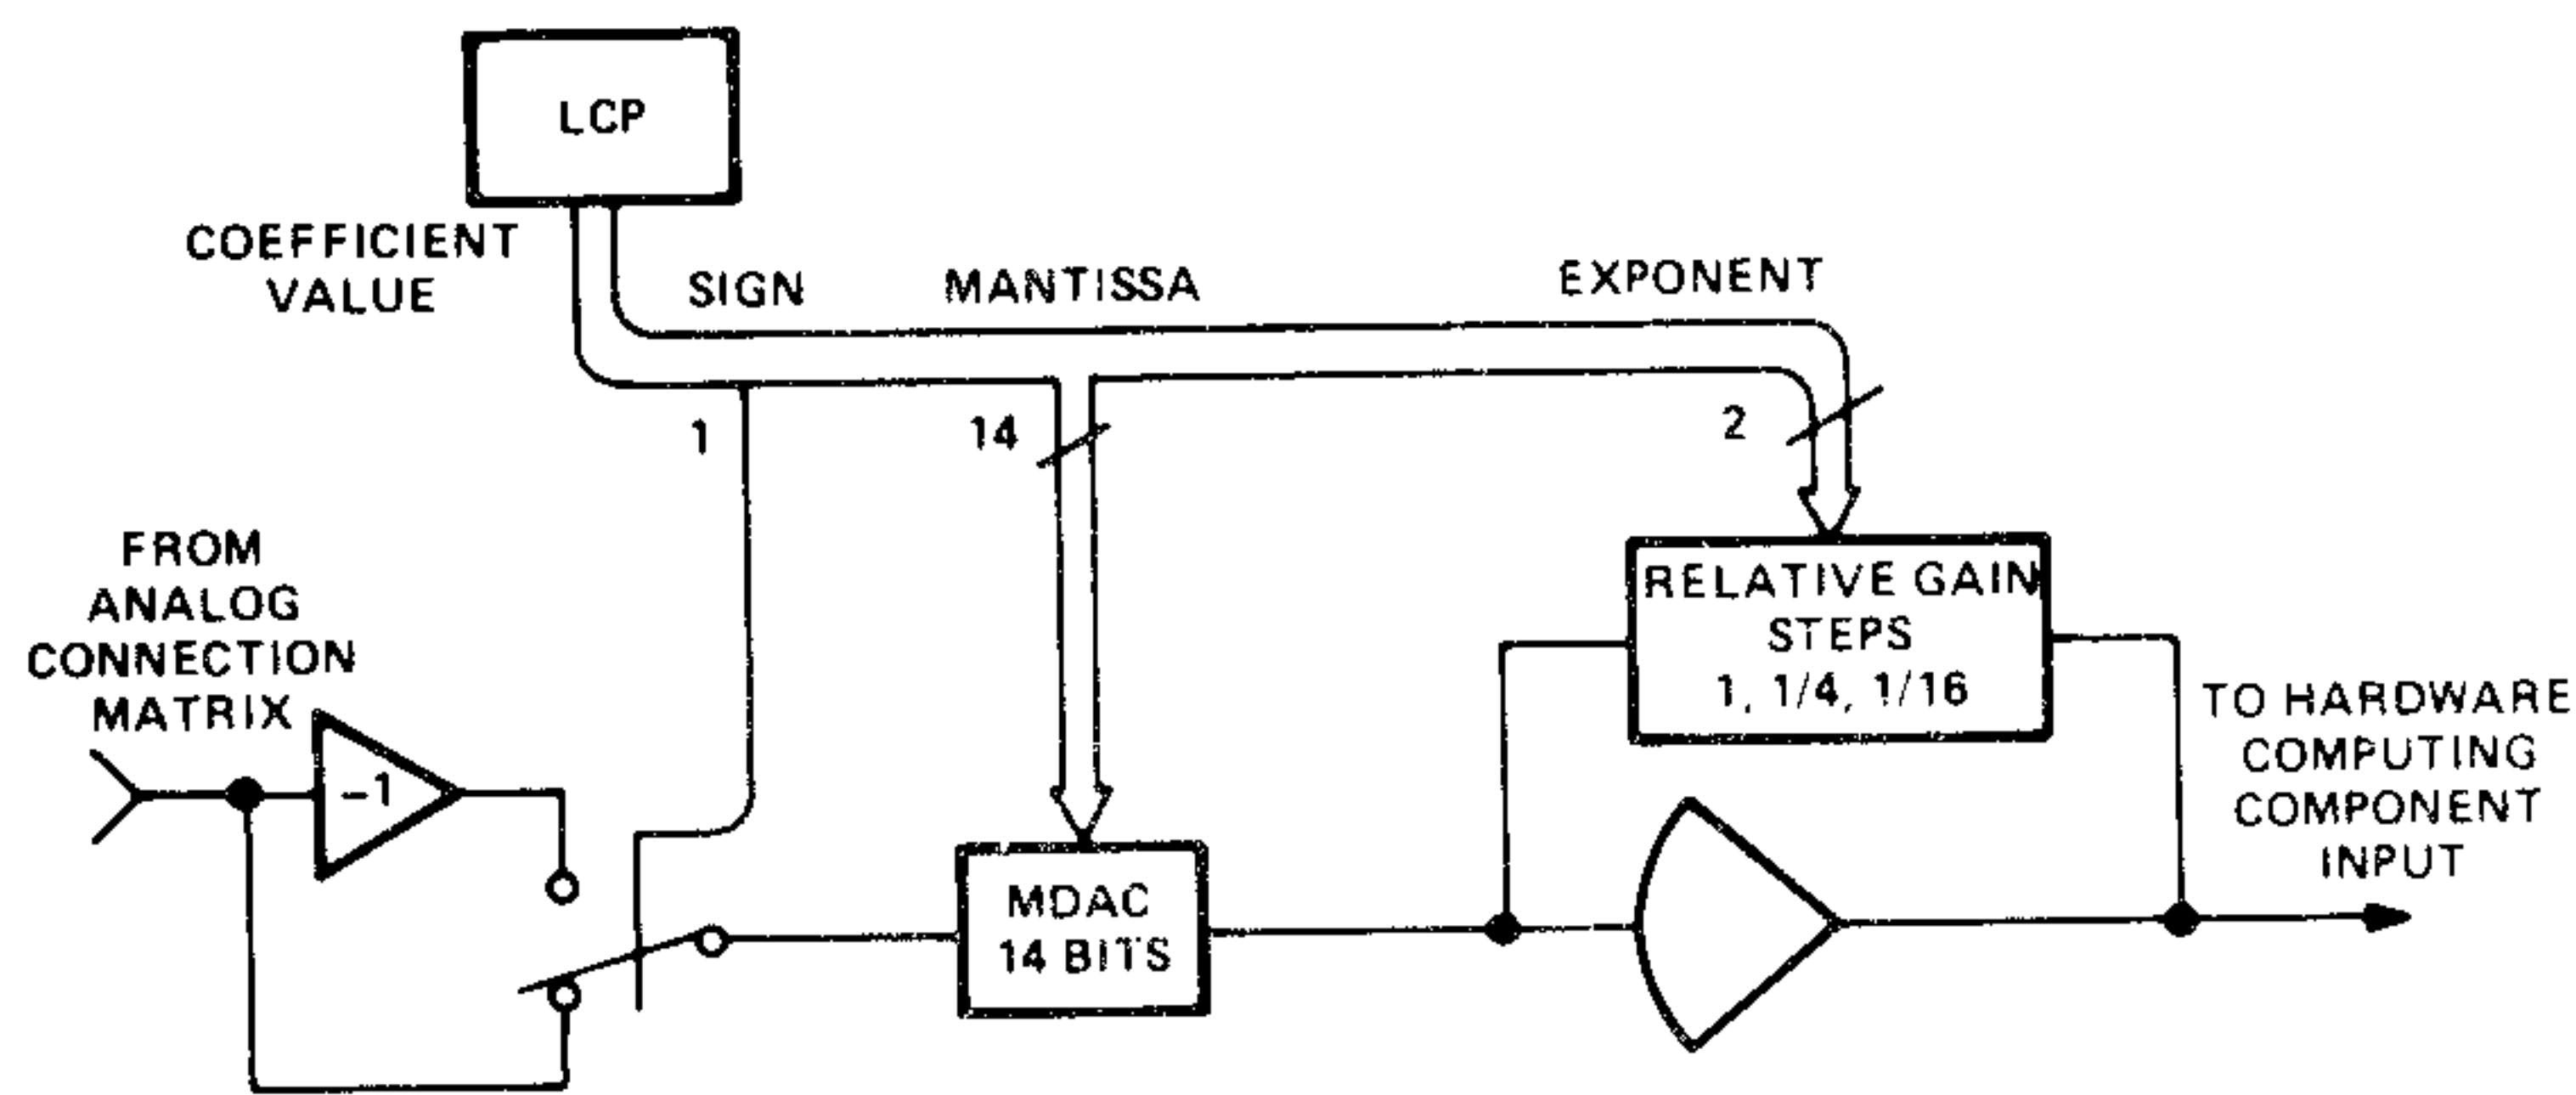


Figure 8 Simplified Block Diagram Schematic of the Extended Range DSCU (% of Output Error)

The LCP functions as follows with the DSCU to perform the autoranging: should a particular DSCU coefficient value be less than $1/4$, but greater than $1/16$, the coefficient word is shifted two bits to the left and the gain (exponent) is changed to $1/4$, maintaining the overall coefficient invariant, but errors are reduced by four at the output. A similar transition occurs should the coefficient value be less than $1/16$.

The extended-range multiplier design, much like the DSCU, provides large improvements in accuracy for small output signal levels. Since $(1/4)^2$ multipliers have an inherent percentage of full scale error, the error as a percentage of output becomes very large for small outputs. This multiplier design reduces this inherent drawback by automatically switching or rescaling the internal multiplier signals such as to approach a percent of output error characteristic. Window comparators inside the multiplier detect when either (X or Y) or (X and Y) decrease below $1/4$ of reference. When this happens, the X and Y signals are amplified back up to full scale and the multiplier output is attenuated by either $1/4$ or $1/16$, respectively, reducing the error by the same factor. See Figure 9a for a block diagram and Figure 9b for the operating equations of the extended range multiplier. Figure 10 shows the relative error reduction in the X, Y plane over conventional multipliers. Figure 11 shows the improved error characteristics vs. conventional multipliers and Figure 12 the improvement in the squaring mode.

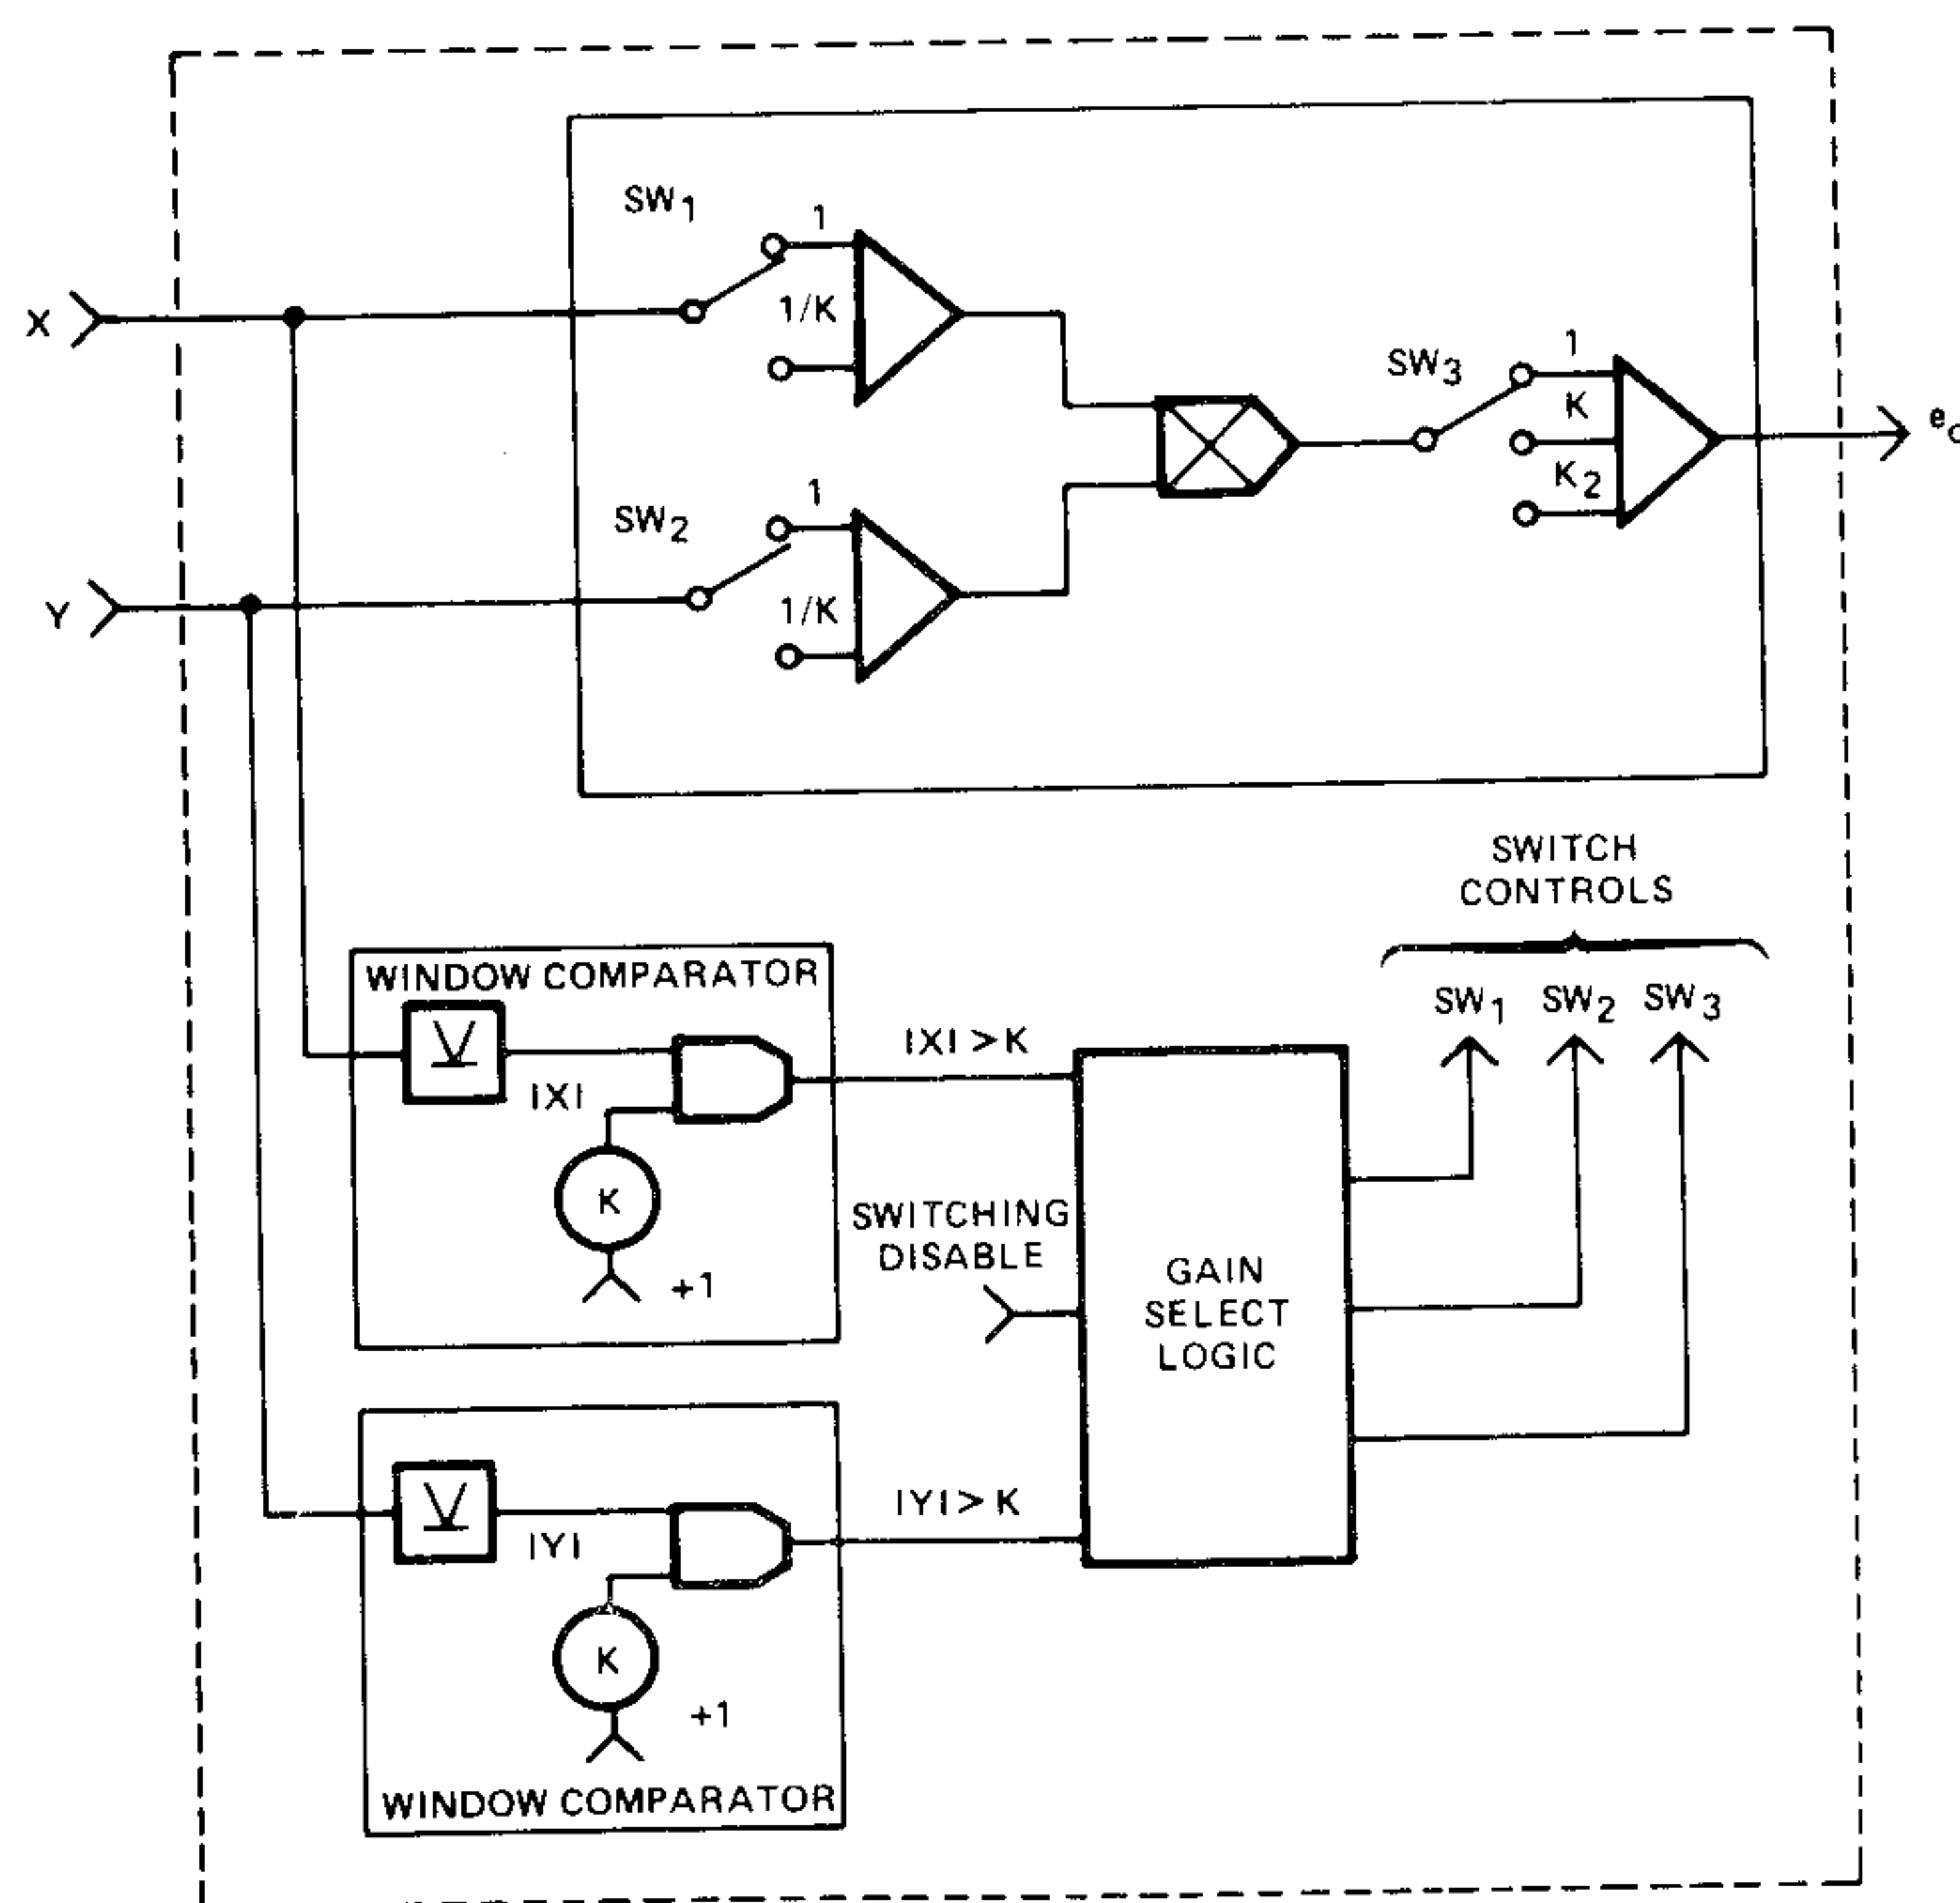


Figure 9a Block Diagram of the Extended Range Multiplier

for: $|X|$ AND $|Y| > K$ $e_0 = XY \pm \epsilon$
 $|X|$ OR $|Y| \leq K$ $e_0 = Y \left[\frac{X}{K} \right] K \pm K\epsilon$
OR
 $e_0 = X \left[\frac{Y}{K} \right] K \pm K\epsilon$
 $|X|$ AND $|Y| \leq K$ $e_0 = \left[\frac{X}{K} \right] \left[\frac{Y}{K} \right] K^2 \pm K^2\epsilon$

$0 < K \leq 1$, $K = 1/4$ in Simstar design

Figure 9b Extended Range Multiplier Operating Equations




Figure 10 Extended Range Multiplier Error Reduction

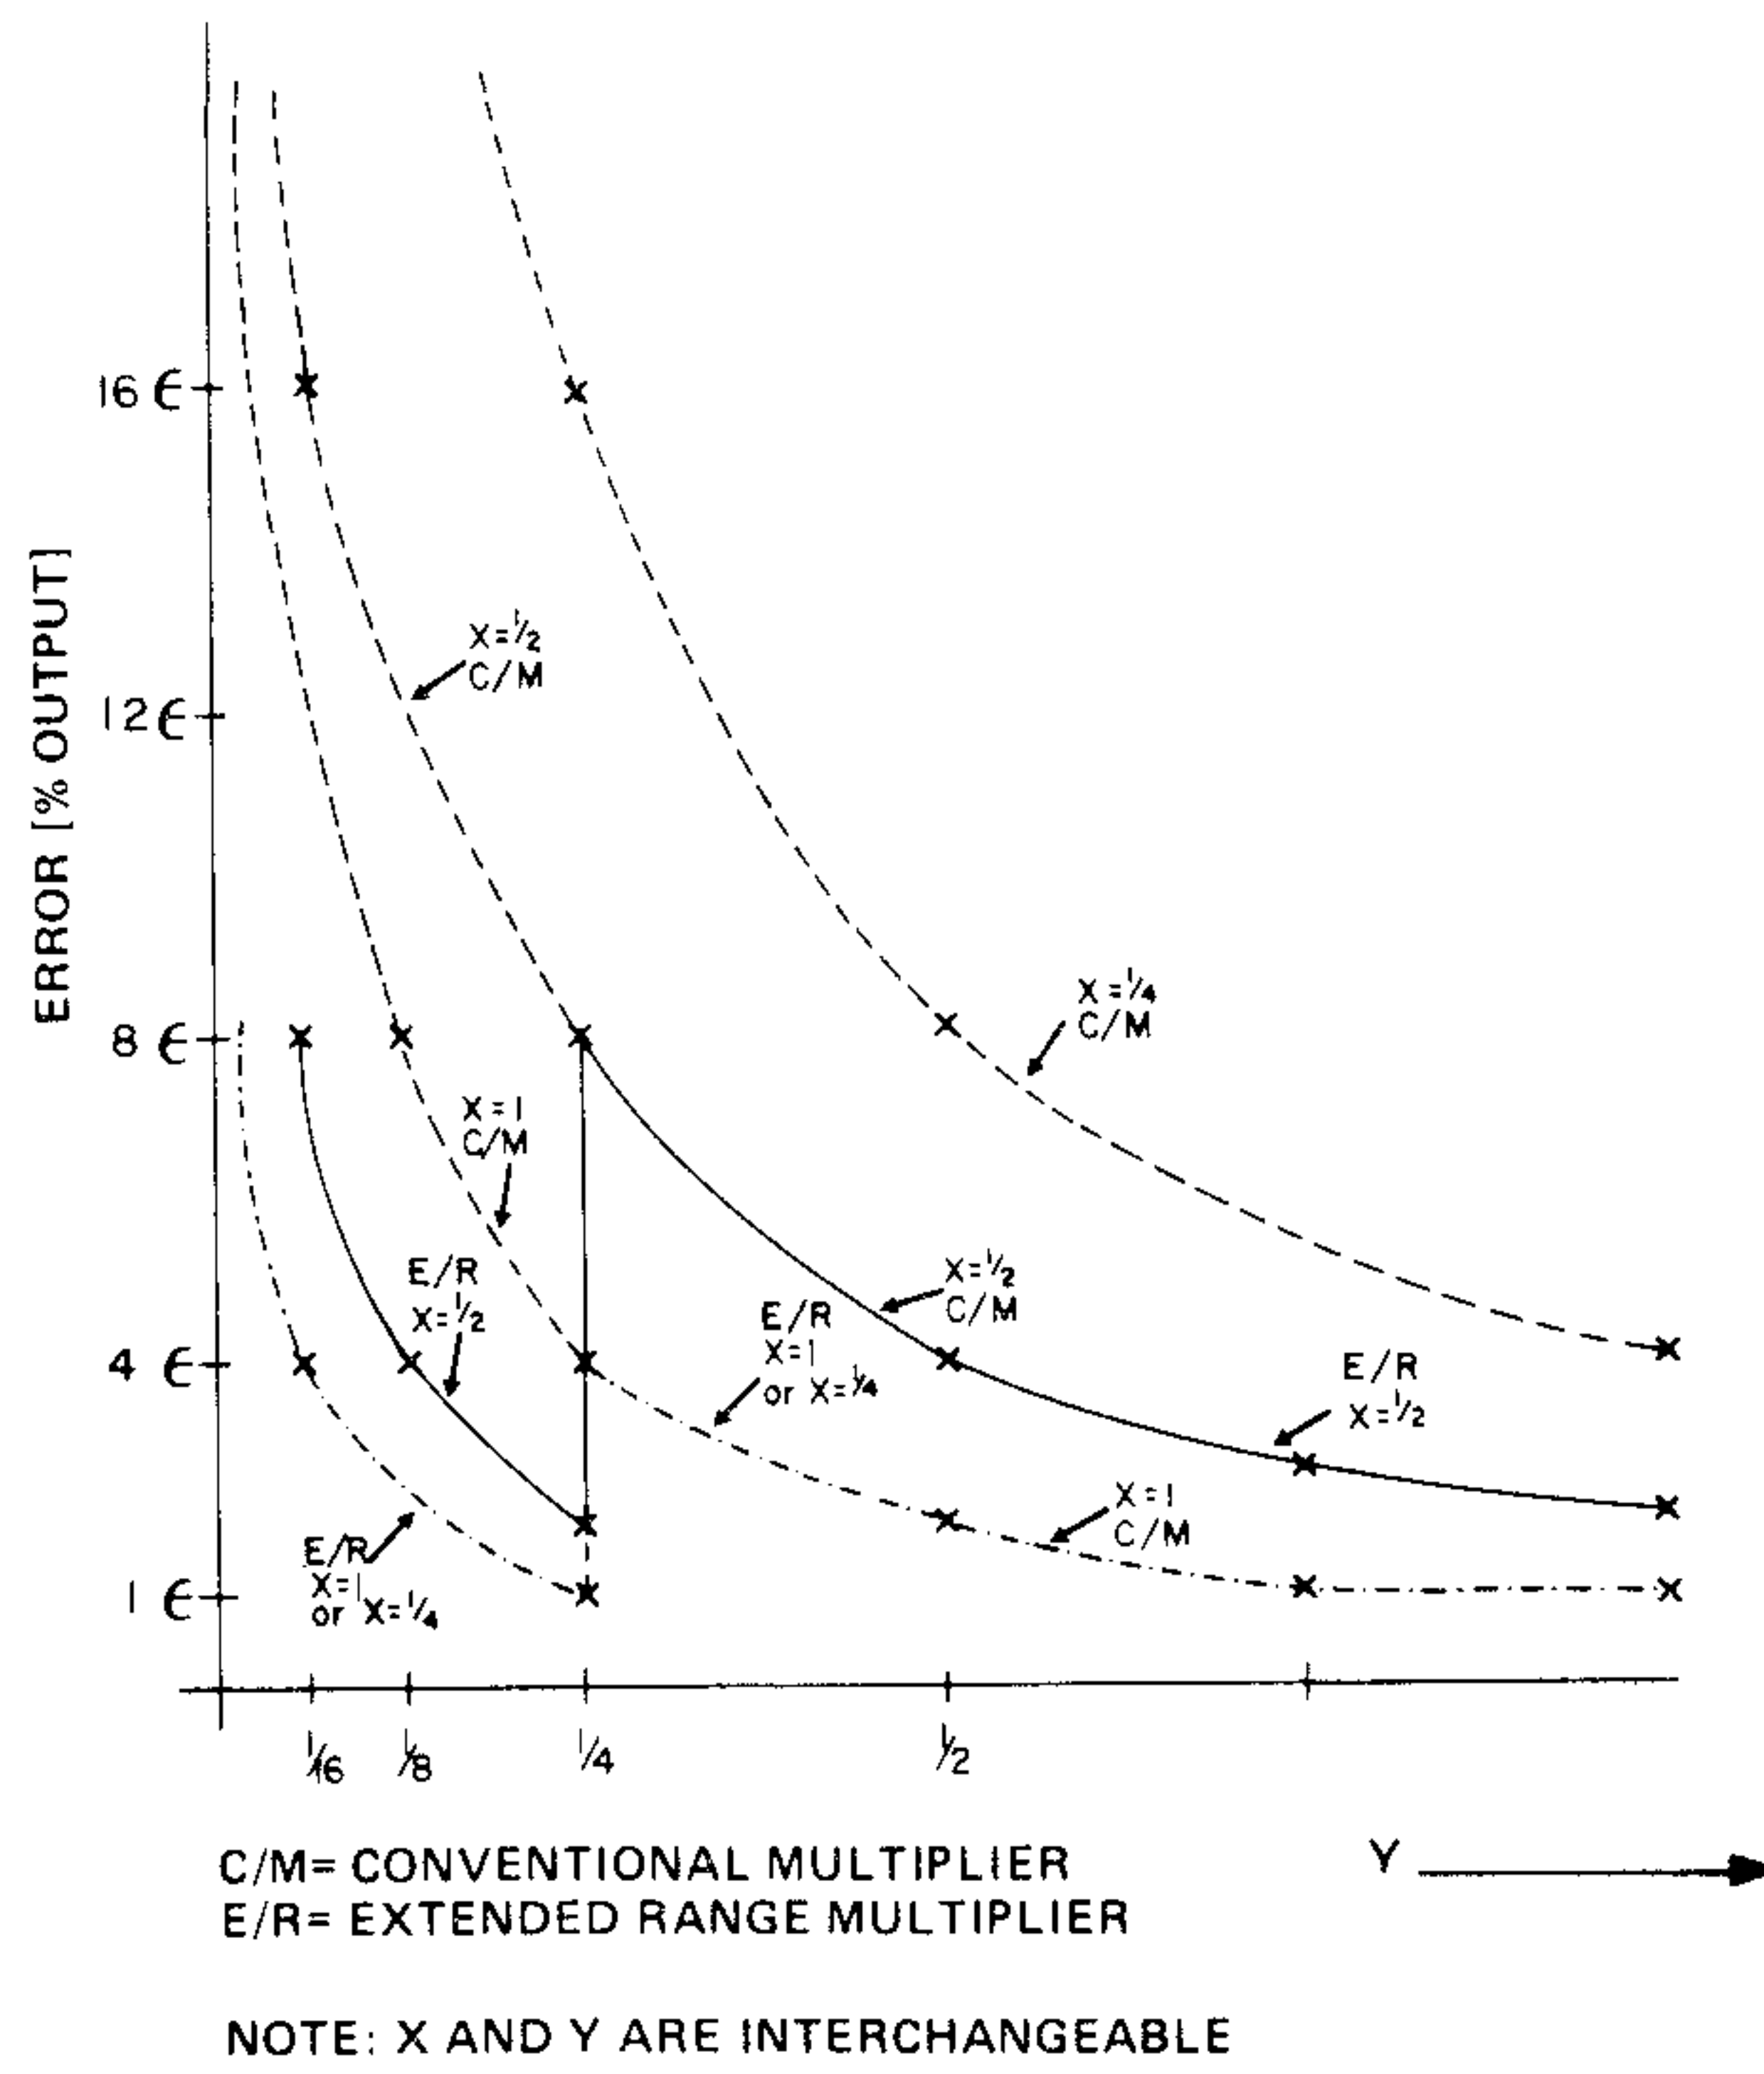


Figure 11 Extended Range Multiplier Error Characteristics

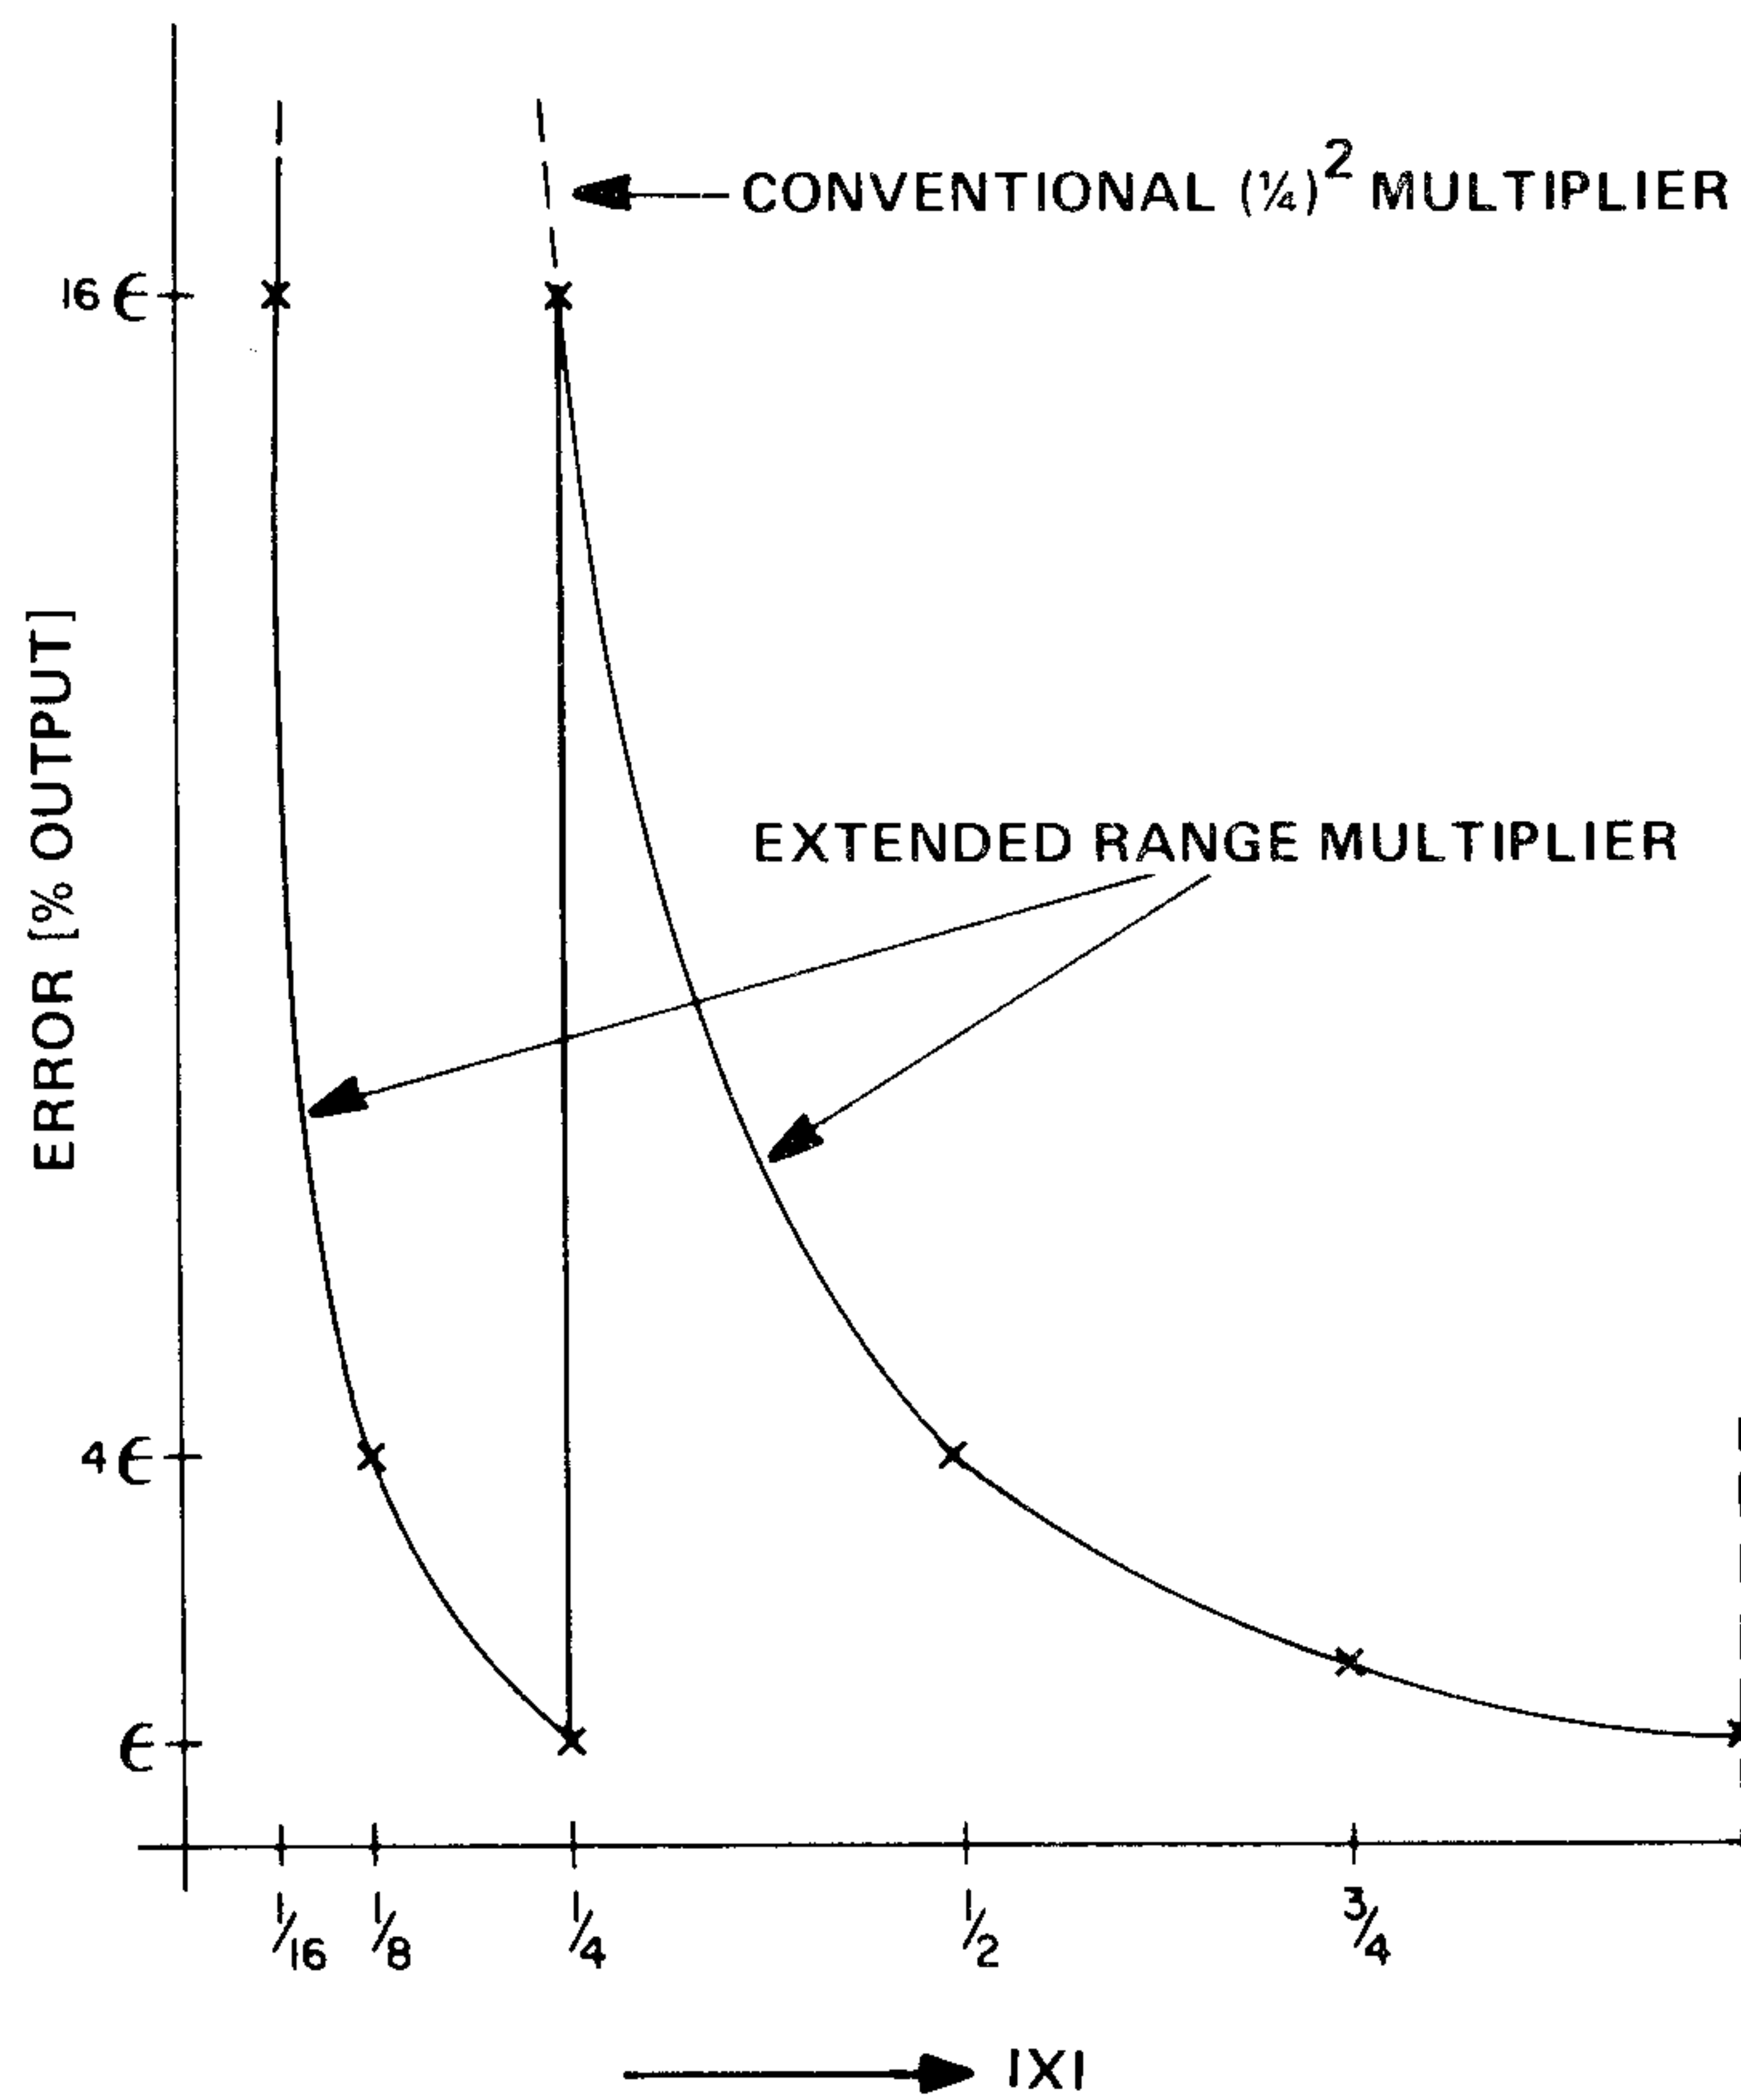


Figure 12 Extended Range Multiplier Error Characteristic (X^2 Mode)

SUMMARY

Eight of the key technological developments for EAI's new SIMSTAR Simulation Multiprocessor have been presented. Each is believed to be an engineering achievement in itself. Together, they and the other developments which were mentioned but not discussed, have laid the basis for a state-of-the-art simulation system to be used by all the major government and industrial based R&D organizations throughout the world.

REFERENCES

1. Benes, V.E., "Mathematical Theory of Connecting Networks and Telephone Traffic", Academic Press, New York 1965.
2. Hannauer, G., "Stored Program Concept for Analog Computers", Final Report, NASA Project #NAS8-21228 June, 1968.
3. Hannauer, G. and Asthana, A., "Recent Advances in Automatic Patching Software", Proceedings of Special Symposium on Advanced Hybrid Computing, San Francisco, July, 1975.
4. Gary, M.R. and Johnson, D.S., "Computers and Intractability", W.H. Freeman and Company, San Francisco 1979.
5. Wirth, N., "Algorithms + Data Structures = Programs", Prentiss-Hall, Englewood Cliffs, New Jersey, 1976.