

# Chip-Package Co-Implementation of a Triple DES Processor

Toby Schaffer, *Member, IEEE*, Alan Glaser, *Member, IEEE*, and Paul D. Franzon, *Senior Member, IEEE*

**Abstract**—This paper describes the design and implementation of a dedicated data encryption standard (DES) processor. The processor consists of three 0.6  $\mu\text{m}$  complementary metal oxide semiconductor (CMOS) integrated circuits (ICs) mounted on a single MCM-D thin-film substrate. Each chip can operate on an individual data stream, or the three can be cascaded to implement the so-called “triple-DES” (3DES) function for increased security. Measurements show 3DES operation at 110 MHz, which translates to a throughput of over 7 Gb/s, the highest reported 3DES throughput to date. System features which contribute to this throughput are the use of area-array (flip-chip) input/output (I/O) and global IC power/ground/clock distribution in the MCM package. In this case, package-level distribution reduced clock skew by 150 ps, and reduced the chip area required for power distribution by 20%. This paper also includes measurements of switching noise of the MCM’s  $V_{\text{dd}}$  plane and how it correlates with a simple model of the system power distribution.

**Index Terms**—CMOS, DES processor, IC, I/O, MCM package.

## I. INTRODUCTION

THE NEED for secure communications has existed since ancient times, but only recently has there been a need for secure communications over high-speed channels. This paper describes a high-bandwidth custom IC which uses a multichip module (MCM) substrate and novel power, ground and clock distribution strategies to yield the highest throughput single-package publicly reported 3DES implementation.

The module also demonstrates the practicality of package-level power, ground and clock distribution. Such distribution can be potentially useful in reducing the need for on-chip interconnect, by supplementing it with off-chip interconnect [1]. In this case, the global clock is distributed in the low-resistivity package wiring so as to reduce the clock skew and clock power. By distributing global power in the package, the on-chip power and ground grids are largely replaced by low-resistance planes in the package, leading to a reduction in the area required for power distribution. It also enables the chip to supplement on-chip decoupling with package-level decoupling.

Manuscript received April 30, 2001; revised December 9, 2003. This work was supported by Defense Advanced Research Projects Agency under Contract DASH04-94-G-003-P2 and National Science Foundation under Grant EIA-9703090.

T. Schaffer and A. Glaser were with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC 27695 USA. They are now with Integrated Device Technology, Inc., Duluth, GA 30097 USA (e-mail: toby.schaffer@idt.com; alan.glaser@idt.com).

P. D. Franzon is with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC 27695 USA (e-mail: paul\_franzon@ncsu.edu).

Digital Object Identifier 10.1109/TADVP.2004.824944

This paper first gives a brief overview of encryption and DES in particular. This is followed by a description of both the pipelined architecture of the IC—which enables single-cycle throughput—and the physical implementation of both the IC and the MCM. It continues with pre- and post-fabrication verification procedures and measurements of both data throughput, clock distribution and power supply switching noise before presenting the conclusions. The switching noise is compared with a predictive model.

## II. DES ALGORITHM

In general, the goal of an encryption system is to allow two parties to communicate in a manner such that transmitted information appears in an unintelligible form to an eavesdropping third party. Incoming *plaintext* is encrypted via some algorithm  $E$  under the control of a *key*. This encrypted data, referred to as *ciphertext*, is transmitted to the receiver, where it is decrypted ( $D$ ) under the control of a second key to restore the original plaintext, as shown below. In the general case the sender’s key  $K1$  and the receiver’s key  $K2$  are different

$$D_{K2}(E_{K1}(\text{data})) = \text{data}.$$

The data encryption standard (DES) specifies an algorithm for encrypting and decrypting digital data. In widespread use, it is both a U.S. government [2] and a private sector standard [3].

DES belongs to a class of algorithms known as *symmetric* or *secret-key*, as opposed to *asymmetric* or *public-key* algorithms such as RSA. The name arises because the same key is used to both encrypt and decrypt data; obviously, the key must remain secret as anyone who knows the key can decrypt any data encrypted with that key. DES also is an example of a *block cipher*, since it operates on blocks of data (64-b wide, in this case) at a time. The key size is 56-b. Finally, since DES is a *Feistel cipher*, encrypted data can be decrypted by simply running the algorithm in reverse, meaning the same hardware performs both functions.

A block diagram of the algorithm is shown in Fig. 1. An encryption begins by shuffling (a simple reordering of) the plaintext via the *initial permutation* (IP).

The IP is followed by sending each block of data through 16 identical rounds. A round, illustrated in Fig. 2, consists of a set of nonlinear substitutions and permutations applied under the control of a 48-b subset of the master key. The particular subset used is a function of the round number and whether the current operation is an encryption or a decryption. The operations performed are various bit-reorderings (E, P, and PC2), a bitwise rotate ( $\ll$ ), XORs, and a table lookup (S-box).




Fig. 1. Block diagram of DES algorithm.

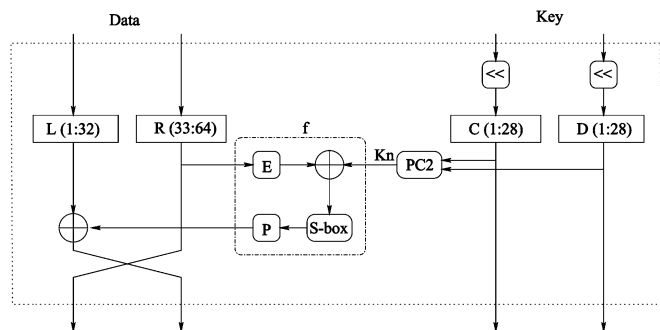


Fig. 2. DES round (encryption).

Finally, after the 16 rounds have been completed, the data undergoes an inverse initial permutation ( $IP^{-1}$ ). The output of  $IP^{-1}$  is the ciphertext.

Decryption, as mentioned before, is performed by essentially running the same algorithm in reverse.

### A. Security of DES

The security of any encryption algorithm against a “brute-force” attack (i.e., an exhaustive search of the keyspace) is directly related to the key length. DES’s key length of 56 b is small by today’s standards, and custom “cracking machines” have been built [4] which can break DES by brute-force in days.

For increased security, the DES operation can be performed three consecutive times. This operation, known as *triple-DES*




Fig. 3. Block diagram of DES IC.

(3DES), is U.S. government standard and expands the effective keyspace to 112 b [5]. A 3DES encrypted message is well out of reach of a brute-force attack for the foreseeable future [6].

## III. IC IMPLEMENTATION

This section first describes the architecture of the DES IC then follows with a discussion of the physical implementation.

### A. Architecture

Since the hardware requirements of each round are identical, previous implementations have instantiated one round and iterated the data through it 16 times [7]–[9]. One author described a gate-level design in which the rounds were unrolled into a 16-stage pipeline, increasing the throughput accordingly [10]. This design was geared for cipher-breaking, however, and as such did not include either high-bandwidth I/O capabilities or the ability to simultaneously execute encryptions and decryptions as does the chip described in this paper.

A block diagram of the IC is shown in Fig. 3.

The datapath, which consists of 16 pipeline stages, is shown as a single block. A block diagram of a pipeline stage is shown in Fig. 4. The eight “cells” in the middle contain the S-boxes, assorted logic and pipeline registers, and on either side is a 28-b register that holds half of the key. It is worth noting that each half of the key need only be routed through four of the eight cells, as this saves significant amounts of area over having to route the whole 56-b bus through the entire stage. Additionally, each stage contains an *opmode* bit, which is discussed below.

Fig. 5 shows a block diagram of an individual cell. The previous stage’s L, f, and R outputs (see Fig. 1) come in at the top, and the new L, f, and R data are sent out at the bottom. Inter-stage wiring implements the P and E functions.

Feeding the pipeline data inputs is a multiplexer, which selects as the data source either an on-chip pseudo-random number generator (PRNG) or the pads—actually, the solder bumps—which bring in data from off-chip. Key generation is done on chip via another PRNG (split into two parts). Every




Fig. 4. Pipeline stage (block diagram).




Fig. 5. Pipeline cell (block diagram).

clock cycle, the pipeline reads in new data and key vectors. The *pads* signal controls the multiplexor which routes either the PRNG or solder bumps to the pipeline. Both the initial seed and the polynomial of all PRNGs are fully programmable via scan chains.

Every cycle the pipeline reads in an *opmode* bit from off-chip. This bit, which is pipelined along with both the key and the data as shown in Fig. 4, directs an individual stage of the datapath to perform either an encryption or a decryption by appropriately shifting the key bits between stages, as illustrated in Fig. 6. When decrypting, the unshaded tri-states are active, shifting the key right as it travels down the pipeline; encryption activates the shaded tri-states to shift the key left. The net effect is that not only does each stage of the pipeline operate on a key/data pair independently of the other stages, but its mode (encryption or decryption) is independent of the other stages as well.

After traversing the pipeline, the output is sent to both the signature analyzer register (SAR), to be hashed with the previous results for functional verification, and the output drivers, to be sent off-chip. Like the PRNG's used to generate data vectors,

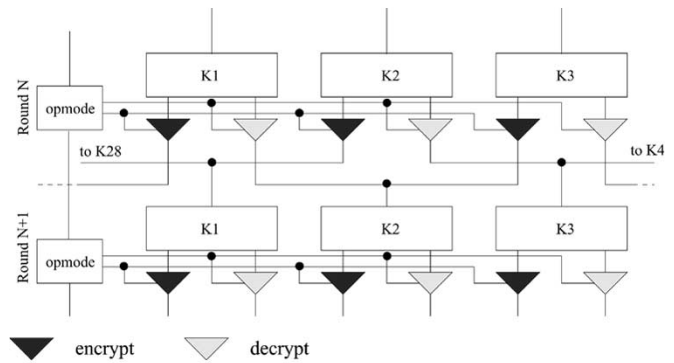


Fig. 6. Inter-stage key shifting.

TABLE I  
DESIGN SUMMARY OF DES IC IMPLEMENTATION

Process	3.3V, 3-metal, 0.6 $\mu$ m (drawn) N-well CMOS
Solder Bump Pads	210
Diameter (glasscut)	45 $\mu$ m
Pitch (glasscut)	250 $\mu$ m
I/O	131
Control	15
Vdd/Gnd/Clk	64
Die Size	5.78mm $\times$ 3.67mm
Transistor Count	123,104

the SAR is implemented as a fully-programmable linear-feedback shift register.

### B. Physical Implementation

The IC was designed in a 0.6  $\mu$ m (drawn) triple-metal CMOS process [11] using a full-custom methodology with EDA software from Cadence [12], [13]. All off-chip communication, as well as power and ground connections, is done through flip-chip solder bumps. Given the die size, 114 standard wirebond pads could have fit along the perimeter. However, this would not only have failed to meet the I/O requirements (128 signal I/O, in addition to control and test) but would have increased the die area 20%. Area-array I/O, which gives I/O capacity proportional to the IC's area rather than its perimeter, enables the IC to achieve its high bandwidth, since every cycle a 64-b data block can be read while another is being written.

With the exception of the S-box ROMs, the chip was implemented in conventional static CMOS. Clocking was done with a single phase clock distributed over the entire chip. Table I presents some information regarding the IC implementation.

Fig. 7 is an annotated photomicrograph of the IC. Data inputs are routed from the lower half of the chip to the "ESD protection" section (nwell resistors and  $V_{dd}/V_{ss}$  clamp diodes) before entering the pipeline. The ESD structures are congregated together rather than placed under the respective input pads because foundry rules dictated a 50  $\mu$ m minimum spacing between ESD structures and active devices. The exceptions are the clock drivers, each of which has its own local ESD protection structure (which does obey the spacing requirement). Data travels through the pipeline in a counter-clockwise direction as indicated in Fig. 7. At the end of the pipeline, the output drivers

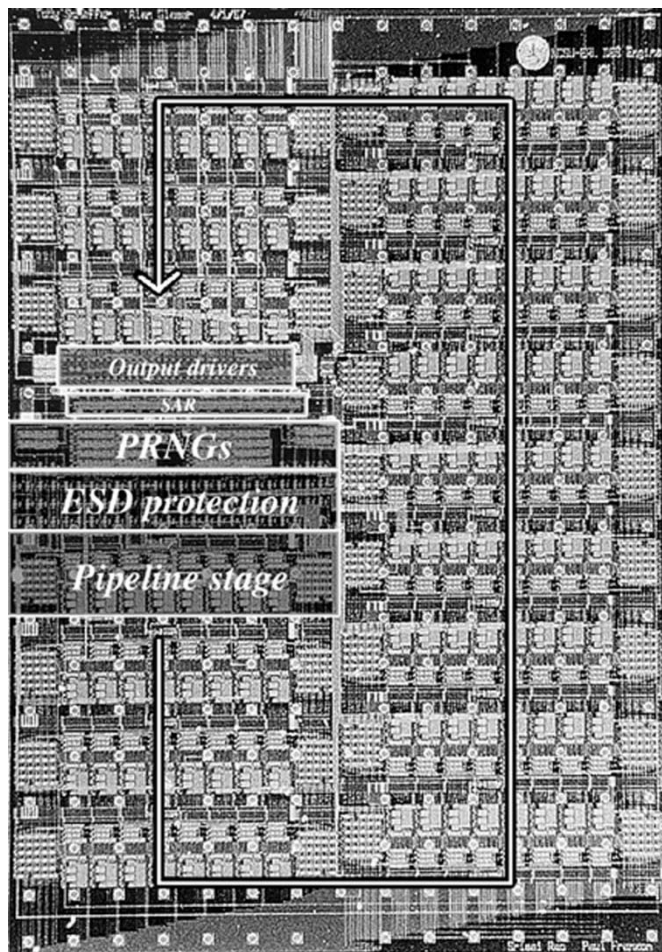


Fig. 7. Annotated photomicrograph of DES IC.




Fig. 8. Pipeline stage.

send the data both off-chip and into the SAR to be hashed with the previous results.

Each pipeline stage implements the operations shown in Fig. 2. A composite photomicrograph of a single stage is shown in Fig. 8. The three circled and labeled solder bump pads are the dedicated  $V_{dd}$ , clock, and ground connections for that particular stage; this is discussed further below.

Fig. 9 is a close-up of a cell (the pipeline registers and additional XOR gates are at the top of the cell and not shown), which is shown shaded in Fig. 8. At the top are the pipeline registers. Half of the key (28-b) comes on the horizontal bus (of which only a few wires are visible) above the XORs; the appropriate 6 b (which six depends on the cell number) are tapped off and xor'ed with 6 b from the registers. The sum is split into two parts: a 4-b number, which is decoded to index into the ROM, and a 2-b number, which multiplexes the correct 4-b ROM output. The eight cells form a 32-b output, which is permuted and expanded (operations P and E in Fig. 2) and sent to




Fig. 9. Pipeline cell (close-up).




Fig. 10. Clock distribution via solder bumps.

the next pipeline stage. The S-box ROM is constructed using a static pseudo-NMOS PLA structure, which accounts for the small S-box area ( $131 \mu\text{m} \times 116 \mu\text{m}$ , including associated decoders and multiplexors).

The chips' global clock, power and ground are distributed on the MCM—there is no global on-chip distribution for these. Each pipeline stage, as well as the data-generating PRNG's and the SAR, has its own clock input pad and driver. The clock distribution is conceptually illustrated in cross-section in Fig. 10, and a composite detail photomicrograph of a single clock input structure (including local ESD protection) is shown in Fig. 11. This arrangement—multiple clock entry points with local clock drivers—reduces clock skew by eliminating a single global driver which would then redistribute the clock to the local drivers. Circuit simulations show that across process variations a single driver can contribute up to 150 ps of skew; this potential skew can be avoided due to the low resistance





Fig. 11. Clock pad and local ESD/driver detail.

Fig. 12.  $V_{dd}/V_{ss}$  distribution via solder bumps.TABLE II  
PROCESS SUMMARY OF MCM-D

Size	substrate	25mm × 25mm
Size	package	47mm × 47mm
Substrate		Al
Interconnect		Cu (16 $\mu$ m wide, 10 mil pitch)
Signal Planes		2 (4 $\mu$ m thick Cu)
Vdd Plane		1 (2 $\mu$ m thick Cu)
Vss Plane		1 (50 mil Al substrate)
Dielectric		polyimide
Thickness	sig2-sig1	6 $\mu$ m
	sig1-Vdd	12 $\mu$ m
	Vdd-Vss	3.5 $\mu$ m
I/O Count	GND	40
	GND2	36
	VDD1	36
	VDD2	36
	Signal	144

of the MCM copper interconnect (4.5 m $\Omega$ /sq.) versus that of on-chip aluminum interconnect ( $\approx$ 40 m $\Omega$ /sq.).

In addition to dedicated clock inputs, each pipeline stage, PRNG and SAR has its own dedicated power and ground bumps which contact the MCM  $V_{dd}$  and  $V_{ss}$  plane, respectively, as conceptually illustrated in Fig. 12. There is no on-chip global mesh for power and ground distribution; however, there are thin metal straps connecting the power rails for each stage  $n$  and  $n+1$  (and similarly for the ground rails) to provide a low-inductance signal return path.

The advantage of using the MCM for global chip power and ground distribution is that it allows the die area to be smaller. In this case, the chip is 30% smaller (for the same peak IR drop) than one using a conventional peripheral-ring power

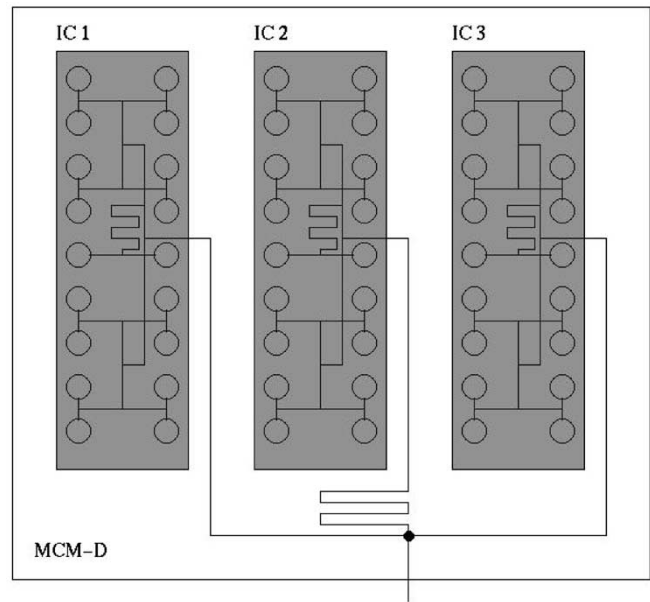


Fig. 13. On-MCM clock distribution H-tree.




Fig. 14. DES MCM (in package).

distribution system [1]. This savings results from the fact that low-resistance (i.e., wide) peripheral  $V_{dd}$ /ground metal rings are not needed in the area-array power-distribution scheme. Again, the low resistance of the MCM copper planes improves the performance/cost ratio. A side benefit of using the MCM for power distribution is the inherent parasitic decoupling capacitance (0.89 nF/cm<sup>2</sup>) between the MCM's power and ground planes.

As mentioned earlier, the I/O solder bumps on the lower half of the chip (looking at Fig. 7) are inputs while those on the upper half are outputs. This facilitates the inter-chip wiring necessary to implement 3DES when the ICs are laid end-to-end, as shown in Fig. 14.

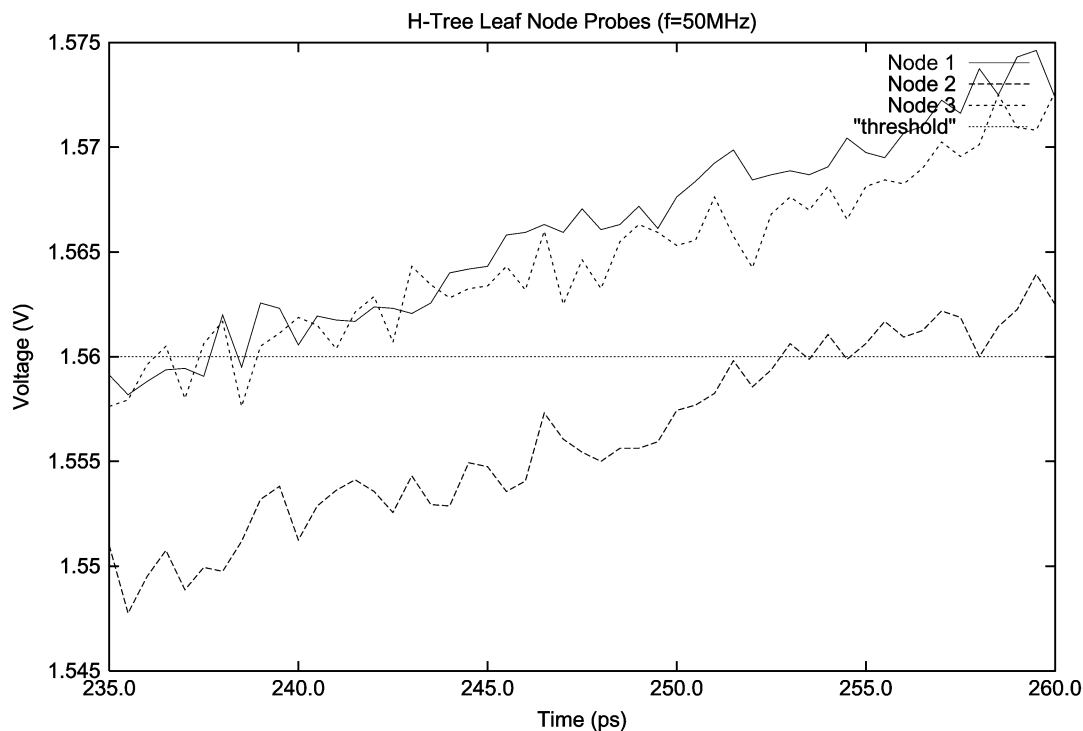


Fig. 15. Clock skew measurement.

#### IV. MCM-D

The ICs are mounted on an MCM-D substrate manufactured by MicroModule Systems, Inc. Process characteristics of the substrate are given in Table II.

A single-ended clock signal is distributed on a simple H-tree built on the MCM's signal layers, and each leaf of the tree is the input to an on-chip driver, as discussed in Section III.B. This is conceptually illustrated in Fig. 13.

Fig. 14 is a photograph of the populated MCM mounted in a 192-pin PGA package. Identifying letters have been superimposed on the chips. The three dark rectangles in the middle labeled A, B, and C are the ICs, face down. To the left of each die are mounting pads for decoupling capacitors, and various probe points are also visible.

To realize the 3DES function, the outputs of A are fed to the inputs of B, and the outputs of B to the inputs of C.

#### V. VERIFICATION

The entire chip was modeled in Verilog with a combination of structural and behavioral code. A predefined set of key-data vectors designed to test DES implementations [14] was successfully run against this model. Additionally, one million random vectors were run and the outputs successfully compared to that of a popular software DES implementation [15]. Standard LVS and DRC checks were also successfully run. As a final pre-fabrication check, a switch-level simulation of the entire chip was run using the Spectre circuit simulator.

With  $2^{64}$  possible data vectors, exhaustive testing is clearly impossible. Therefore, PRNGs and a signature analyzer (hash register), as described in Section III-B, were used to do functional testing.

Before beginning to feed data and keys into the pipeline, all flip-flops are initialized via scan chains. As mentioned above and shown in Fig. 3, every cycle of normal operation the pipeline output is hashed by the SAR. At the end of the test sequence, the contents of the SAR are scanned out and compared to the value predicted by the Verilog simulation. Due to area constraints, the SAR is only 16 b long. To verify that each bit of the output is correct, the 64-b output is divided into four 16-b groups which are multiplexed into the SAR (the "select" signal in Fig. 3). Thus, depending on which group is currently selected, all output bits can be hashed, albeit not simultaneously.

#### VI. MEASUREMENT RESULTS

A four-layer PCB was designed to accommodate the MCM package and provide interfaces for injecting stimuli and measuring outputs. Stimuli were created with a Tektronix HFS 9009, and a Tektronix TLA704 Logic Analyzer was used to examine the outputs. A ZIF socket was used to seat the MCM package rather than soldering it directly to the PCB to facilitate testing the various samples.

The basic method used to test for correct functionality was

- choose initial seeds and polynomials for PRNG's and SAR (polynomials were chosen to avoid cycles in the PRNGs and to minimize the probability of aliasing in the SAR [16]);
- run Verilog simulation and note final SAR contents;
- initialize chips identically and run same number of cycles;
- scan out SAR(s) and compare to Verilog simulation.

Four tests of 25 000 vectors each were run on each chip, one test for each 16-b group of the output. In each case the SAR contents matched those predicted by the Verilog simulations. This procedure was run for three different *opmode* settings: encrypt only,





Fig. 16. Noise model of pipeline stage.

Fig. 17. MCM  $V_{dd}$  plane voltage (simulated and measured).

decrypt only, and alternating encrypt and decrypt every cycle (i.e., eight encryptions and eight decryptions in the pipeline simultaneously). Again, in each case the SAR contents matched Verilog simulations.

3DES operation was also measured for maximum speed (i.e., throughput) using the same methodology as above. The modules functioned correctly at clock speeds up to 110 MHz, which corresponds to 7.04 Gb/s throughput.

Clock skew between the three ICs (between the H-tree leaf node at the northeast corner of the chips) was measured using a Tektronix 11801A oscilloscope and an HP 8133 pulse generator as a low-jitter clock source. Due to the symmetrical nature of the clock distribution H-tree, this is representative of the skew between any two clock entry points on the entire module.

A close-up view of the rising edge of the clock waveforms, along with a reference line representing the threshold of the on-chip clock driver, is shown in Fig. 15. The on-chip clock driver is not perfectly symmetrical, leading to a threshold voltage slightly less than  $V_{dd}/2$ . The jagged appearance of the traces is due to the oscilloscope's quantization error at the given level of magnification. At the threshold, skew was measured to be approximately 15 ps. Sources of error include manufacturing mismatches in the H-tree and asymmetrical capacitive loading due to crossover of signal interconnects.

It was important to verify that the flip-chip on-MCM power distribution scheme did not lead to increased noise levels. The voltage of the  $V_{dd}$  plane was measured and compared to that predicted by a simple model, in which the MCM planes were

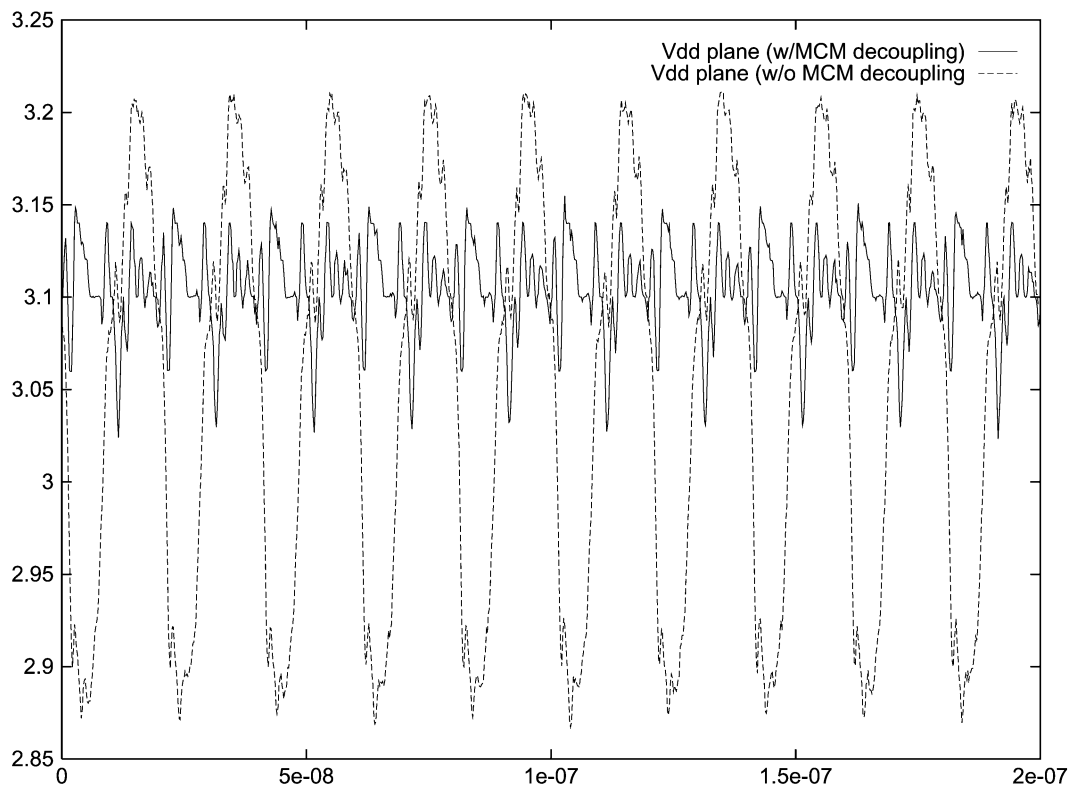


Fig. 18. Effect of decoupling capacitors on MCM  $V_{dd}$  plane voltage.

modeled as an RLC grid using vendor-provided parasitic values [17]. The vast majority of the on-chip ac current is drawn by the clock buffers driving each pipeline stage, while the S-box ROMs dissipate dc power. (Each ROM has a PMOS switch controlled by a “sleep” line, which, when asserted, reduces the dc current effectively to zero.) The output drivers also draw significant amounts of current. Each pipeline stage was modeled as a dc current source in conjunction with the transistor-level model of the clock driver driving an amount of gate capacitance equal to the load presented to the on-chip clock driver, as illustrated in Fig. 16. Each chip’s noise model was composed of 16 pipeline stages in addition to the 64 output drivers. To compare this model with the actual system, measurements of the  $V_{dd}$  plane during operation of the three ICs were taken with a GGB Model 34A high-impedance active probe.

Fig. 17 shows the results of a measurement compared to the predicted noise. The system clock is running at 50 MHz. Intra-cycle logic switching was not modeled in the simulation, which accounts for the smoothness of the simulated noise; however, the model predicts the mid- and low-frequency noise accurately. The peak noise is rather high, but on average the voltage fluctuations stay within 10% of the nominal  $V_{dd}$ . To see how much of the noise was due to the ZIF socket used in the test setup, the noise simulation was rerun with the socket inductance removed. This is labeled as “simulation (w/o ZIF).” Peak-peak noise predicted by this simulation is reduced to about 200 mV from 300 mV in the case with the socket, leading to the conclusion that about a third of the observed noise is due to parasitics of the test setup, not the MCM itself. It is also worth reiterating

that this noise is measured on the  $V_{dd}$  plane itself and therefore only a fraction of it is actually seen as simultaneous switching noise at the drivers’ outputs.

Adding three surface-mount  $3.3 \mu\text{F}$  capacitors to the MCM substrate reduced peak-peak noise magnitude to 100 mV. Fig. 18 shows before-and-after measurements illustrating this result. This shows that the native decoupling capacitance of the power-ground plane pair were insufficient.

## VII. CONCLUSION

This paper described a high-throughput 3DES IC implemented in a standard CMOS process. Utilizing area-array I/O and an MCM-D thin-film substrate for global power, ground and clock distribution, as well as inter-chip signal communication, 3DES operation with a sustained throughput of 7 Gb/s was obtained. Clock skew across the MCM was measured to be 15 ps, and the measured dynamic  $V_{dd}$  plane noise agreed well with the noise predicted by the model.

These results demonstrate the following:

- 1) custom DES design leads to very high throughputs, the highest demonstrated so far;
- 2) on-package global clock distribution is practical and leads to lower skews than comparable H-tree schemes with added buffer layers;
- 3) on-package power and ground distribution is practical and results in manageable noise levels, once sufficient on-package decoupling is employed.




## ACKNOWLEDGMENT

The authors wish to thank S. Rao for his work on the SAR layout and Dr. B. Beker, USC, for advice with the noise modeling.

## REFERENCES

- [1] T. Schaffer, A. Glaser, S. Rao, and P. Franzon, "A flip-chip implementation of the data encryption standard (DES)," in *Proc. 1997 IEEE Multi-Chip Module Conf.*, Feb. 1997.
- [2] "Data Encryption Standard (DES)," National Institute of Standards and Technology, FIPS 46-2 ed., Dec. 1993.
- [3] "American National Standard (X3.92-1981) Data Encryption Algorithm," ANSI, 1981.
- [4] "Cracking DES: Secrets of Encryption Research, Wiretap Politics, and Chip Design," Electronic Frontier Foundation, O'Reilly and Associates, Inc., Massachusetts, 1998.
- [5] R. Merkle and M. Hellman, "On the security of multiple encryption," *Commun. ACM*, vol. 24, no. 7, pp. 465-467, 1981.
- [6] M. Blaze *et al.*, Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security, Jan. 1996.
- [7] H. Eberle and C. Thacker, "A 1 Gbit/second GaAs DES chip," in *Proc. 1992 Custom Integrated Circuits Conf.*, 1992, pp. 19.7.1-19.7.4.
- [8] F. Hoornaert, J. Goubert, and Y. Desmedt, "Efficient hardware implementation of the DES," in *Proc. Adv. Cryptol. (CRYPTO'84)*, 1984, pp. 147-173.
- [9] (1997) VMS 110 16-bit DBS Coprocessor. VLSI Technology, Inc. [Online]. Available: <http://www.vlsi.com>
- [10] M. Wiener, "Efficient DES Key Search," School of Computer Science, Carleton University, Tech. Rep. TR-244, May 1994.
- [11] "CMOS14TB Design Reference Manual," Hewlett-Packard, 1994.
- [12] T. Schaffer, A. Glaser, and W. Ficken, "The N.C. State University Cadence Design Kit (CDK)," <http://www.ece.ncsu.edu/cadence/cdk.html>. WWW site.
- [13] T. Schaffer, A. Stanaski, A. Glaser, and P. Franzon, "The NCSU Design Kit for IC fabrication through MOSIS," in *Proc. Int. Cadence Users' Group Conf.*, 1998.
- [14] "Validating the Correctness of Hardware Implementations of the NBS Data Encryption Standard," National Bureau of Standards, NBS Special Publication 500-20 ed., Sept. 1980.
- [15] E. Young, Libdes (v. 3.23). anonymous FTP.
- [16] M. Abramovici, M. Breuer, and A. Friedman, *Digital Systems Testing and Testable Design*. New York: W.H. Freeman, 1990.
- [17] K. Lee and A. Barber, "Modeling and analysis of multichip module power supply planes," *IEEE Trans. Comp., Packag. Manufact. Technol. B*, vol. 18, pp. 628-639, Nov. 1995.




**Toby Schaffer** (M'01) received the B.S. degree in computer engineering, the M.S. degree in electrical engineering, and the Ph.D. degree in electrical engineering, from North Carolina State University (NCSU), Raleigh, in 1991, 1994, and 2000, respectively.

At NCSU, his research focused on power and clock distribution issues for multichip modules. Additionally, he was one of the original developers of the NCSU Cadence Design Kit. He is currently a Senior CAD Engineer with Integrated Device Technology, Inc., Duluth, GA.



**Alan Glaser** (S'95-M'01) received the B.E. degree in electrical engineering from Vanderbilt University, Nashville, TN, in 1991 and the M.S. degree in electrical engineering from North Carolina State University (NCSU), Raleigh, in 1995.

While at NCSU, he helped develop CAD software and frameworks that are currently in use in both industry and academia. He is presently a Senior CAD Engineer with Integrated Device Technology, Inc., Duluth, GA.



**Paul D. Franzon** (S'85-M'88-SM'99) received the Ph.D. degree from the University of Adelaide, Adelaide, Australia, in 1988.

He is currently an Alumni Distinguished Professor at North Carolina State University, Raleigh. He has also worked at AT&T Bell Laboratories, DSTO Australia, Australia Telecom and Communica Ltd. His current interests center on the technology and design of complex systems incorporating VLSI, MEMS, advanced packaging, and molecular computing. Application areas currently being explored include novel advanced packaging structures, network processors, SOI baseband radio circuit design for deep space, on-chip inductor and inductance issues, RF MEMS, and moleware circuits and characterization. He has led several major efforts and published over 120 papers in these areas.

Dr. Franzon received the NSF Young Investigators Award in 1993 and was selected to join the NCSU Academy of Outstanding Teachers in 2001.