

Fabrication and Performance of Mesa Interconnect

L. Richard Carley, David F. Guillou, and Suresh Santhanam

Department of Electrical and Computer Engineering
Carnegie Mellon University, Pittsburgh, PA 15213 USA

Abstract

This paper explores the decrease in interconnect capacitance that can be achieved by the use of mesa-shaped wires. A methodology for creating mesa interconnect is described that uses simple postprocessing steps following standard CMOS IC foundry processes. By removing the oxide between metal lines and replacing it with air, the fringing capacitance between metal lines and the substrate is reduced as is the capacitance between adjacent metal lines. The decrease in capacitance resulting from the removal of the oxide between metal conductors is predicted using a 2-D electrostatic field solver and verified by measurements on test structures fabricated in an 0.5 μ HP CMOS process. Capacitance of mesa interconnect structures was reduced by roughly 40% for these test structures. For circuits dominated by interconnect capacitance, cutting the parasitic capacitance by 40% cuts the power dissipation by 40% and cuts the gate delay by 40%. Therefore the power-delay product is reduced by roughly 65% and the energy-delay product is reduced by roughly 80%. Although many issues must still be addressed to make this form of interconnect practical, the results presented in this paper suggest that the rewards in terms of decreasing capacitance could well be worth investing the effort to solve these issues.

1. Introduction

Two of the most important goals in designing digital ICs are achieving the highest possible clock speed and the lowest possible power dissipation. Both of these goals are hampered by the capacitance that interconnect wiring adds to logic circuits. To a first order, both power dissipation and gate delay are linearly proportional to interconnect capacitive load. And, for typical digital ICs, the power used to drive on-chip interconnect capacitances ranges from 50% to 80% of the total intrachip (excluding off-chip I/Os) power consumption. As process feature sizes continue to shrink, this percentage will grow even higher because the capacitance of a metal line is dominated by its fringing fields and not its planar area.

In this paper a post-processing fabrication methodology is demonstrated that can reduce on-chip interconnect capacitance by roughly **40%** [1] depending on the density of interconnect in the wiring channels. For circuit nodes whose capacitance is dominated by interconnect, this translates into a **40%** saving in power and a **40%** decrease in delay. Note, in applications requiring a fixed delay we can convert the **40%** decrease in delay into additional power savings by decreasing the power supply voltage. We achieve this decrease in interconnect capacitance by replacing much of the silicon dioxide and silicon nitride (which have dielectric constants of 3.9 and 7.5 respectively) that surround metal interconnect with air (which has a dielectric constant of 1). If we were to remove all of the oxide, we could theoretically reduce the capacitance by a factor of at least **3.9x**. However, in order to maintain mechanical rigidity, we leave the oxide underneath metal interconnect while removing the oxide from its sides. And, in order to insulate the exposed metal lines against particulate contamination and also to provide some mechanical support at the top and sidewalls of aluminum interconnect, after stripping off all of the oxide and nitride between wires, we isotropically recoat the interconnect with a thin layer of nitride. This results in interconnect wiring that consists of a series of mesas – hence the name mesa interconnect.

First we describe the process by which we created mesa interconnect as a series of post-processing steps after conventional CMOS fabrication. Next we present the degree of capacitance reduction that can be expected for various interconnect structures derived via electrostatic 2D field analysis. This is followed by experimental verification of two of these structures using a prototype fabricated through MOSIS in a 0.5 μ HP CMOS process. Finally, some limitations and problems with this process will be discussed.

2. Fabrication of Mesa Interconnect

The basic method we use for fabricating mesa interconnect structures was originally developed by us as part of the process we use for creating movable microstructures integrated with conventional CMOS fabrication [2] [3]. Specifically, we use reactive ion etch (RIE) to anisotropically remove the oxide and nitride from silicon ICs. Note, using RIE to remove oxide anisotropically is not a new idea. For example, this

technique has been used to prepare samples of interconnect structures for imaging in scanning electron microscopes, e.g. [4]. Isotropic etching of oxide from silicon ICs while preserving the electronic devices has also been demonstrated before [5]. Fig. 1a illustrates a cross-section of a normal CMOS circuit. We then apply an RIE anisotropic oxide etch which removes oxide and nitride not covered by metal as seen in Fig. 1b. The mesa interconnect structure can be clearly seen in the middle of Fig. 1b. In order not to damage underlying device structures, the RIE etch is terminated before the field oxide is completely removed. Finally, a thin layer of silicon nitride is isotropically sputter deposited on the entire die as shown in Fig. 1c.

Although the example shown in Fig. 1 shows a three metal process, this approach can be applied to any number of layers of metallization. However, as will be discussed below, as more layers of metal are used, there is more etching of the topmost metal layer. In addition, the idea of only cutting part way down through the oxide and nitride could be used to create mesa interconnect for only the upper metal layers in a process. This would be particularly valuable for fabrication processes that use some form of silicided first level interconnect which could be sensitive to the RIE. Note, not reducing capacitance of the first level interconnect in these cases is acceptable because it is typically used for local wiring only, where gate input capacitance dominates over wiring capacitance.

The metal layers act as etch-resistant masks during the dry etching step that anisotropically removes the oxide and nitride layers. This dry etch step was performed in a Plasma-Therm 790 reactor. Oxide areas not covered by metal were anisotropically etched in a CHF_3/O_2 reactive ion etch (RIE) for 160 min, resulting in the cross-section shown in Fig. 1b. The thickness of the etched oxide is $5.6\mu\text{m}$. We have specifically tailored the etch process to remove as little of the metal as possible. Since ion milling during RIE is the main cause of metal loss, this is achieved by adjusting the pressure and power to keep a low DC bias. Metal etching is important for two reasons. First, if too much of the top metal were removed then vias to the next lower metal layer might fail. And second, any metal removed might be redeposited on the surface of the die which could possibly result in a short circuit between two metal lines. Approximately $0.2\mu\text{m}$ out of $1.7\mu\text{m}$ of the third-level aluminum is etched away in this step. First-level metal and second-level metal are not significantly eroded because they have more oxide covering them.

At the left of Fig. 1, MOS devices are shown with third-level metal acting as an etch shield. To date, we have always covered all active devices with these metal shields to prevent etching the oxide above them. The reason for this approach is that we initially envisioned applying this technique to stan-




Figure 1: Process flow cross-sections. (a) after CMOS foundry processing. (b) after anisotropic oxide etch. (c) after deposition of silicon nitride.

dard cell designs that have distinct cell and routing channel regions. We would create mesa interconnect through the channel regions, and we would shape the metal power and ground lines in the standard cells in order to cover all device regions. However, if we can determine that the process of forming mesa interconnect above devices does not affect them, then we plan to apply this methodology to entire designs. Oxide etching is normally stopped just before the field oxide is completely removed. This leaves an oxide layer covering the silicon everywhere helping to maintain protection of the active devices. Note, in all cases all metal lines are firmly anchored to the substrate by pillars of oxide underneath them. The process presented in this paper results in extremely good alignment of the oxide pedestals with the metal lines because they are inherently self-aligned. However, if resputtering of the top metal layer is a problem, a patterned resist layer could be used to block the oxide etch. But, its alignment to the metal lines would have to be accurate.

A scanning electron micrograph (SEM) of an oxide etched interconnect structure is shown in Fig. 2. The nearly vertical walls of the oxide structures are clear in this picture. In addition, regions of metal-1, metal-2, metal-3, and even vias between metal-1 and metal-2 and between metal-2 and metal-3 are clearly visible. It is clear that the metal-3 to metal-2 vias, which are the thinnest part of metal-3 due to the step coverage, are not damaged by the etching. In addition, we have seen no evidence of resputtered metal bridges forming between metal lines in any of our test structures.




Figure 2: Side view of a mesa interconnect structure after etching.

3. Capacitance Prediction

Using a 2D electrostatic field solver we can predict the capacitance of any collection of interconnect wires both before and after forming mesa interconnect. Not only can we predict the change in capacitance to the substrate, but we can also predict the reduction in wire-to-wire capacitance; i.e., the reduction in cross-talk. Table 1 summarizes the reduction in capacitance caused by removing the oxide for isolated metal-1 conductors, metal-2 conductors, and metal-3 conductors. In addition, the reduction in metal-1 to metal-1 capacitance is computed together with metal-1 to metal-2 capacitance for lines spaced 1.2μ apart in both cases. This data assumes complete oxide etch, down to the substrate.

TABLE 1. Summary of Predicted Capacitance Decrease

Layer combination	% decrease
Metal-1 alone over substrate	53%
Metal-2 alone over substrate	59%
Metal-3 alone over substrate	53%
Metal-1 with Metal-1 adjacent over substrate	59%
Metal-2 with Metal-1 adjacent over substrate	56%

4. Experimental Verification

A 0.5μ CMOS test chip was fabricated to verify the predicted change in capacitance. In order to accurately measure the capacitance, the interconnect was attached to a ring oscillator

designed so that the interconnect capacitance would be dominant in determining its oscillation frequency. Two different configurations of metal lines were considered in the experimental design. In the first case (SEM shown in Fig. 3a) the interconnect was a set of 0.9μ wide metal-1 lines running with a 1.2μ spacing over the substrate. In the second case (shown in Fig. 3b) the interconnect was a set of 0.9μ wide metal-2 lines running over the substrate with a set of 0.9μ wide metal-1 lines 1.2μ away on either side of the metal-2 lines.

We then measured the change in oscillation frequency as the oxide was removed. The measured data are summarized in Table 2. We can see that the experimentally measured reduc-

TABLE 2. Summary of Experimental Capacitance Decrease

Layer combination	Field oxide thickness remaining	% decrease
Metal-1 with Metal-1 adjacent over substrate	1000nm	25%
Metal-2 with Metal-1 adjacent over substrate	1000nm	29%
Metal-1 with Metal-1 adjacent over substrate	200nm	46%
Metal-2 with Metal-1 adjacent over substrate	200nm	48%

tion in capacitance was lower than that predicted by the electrostatic field analysis. There are three primary reasons for these inaccuracies. First, our simulations assumed complete etching of oxide and nitride, with no field oxide remaining. Second, the exact shape of the oxide walls was not taken into account in the field analysis. Third, the relative thickness of the oxide and nitride layers is not known precisely and they have significantly different dielectric constants.

Deposition of 100nm of silicon nitride after mesa interconnect formation to prevent contamination problems results in about 5% increase in capacitance over the data presented in table 2, which yields better than 40% capacitance reduction with 200nm of field oxide still remaining.

5. Issues In Using Mesa Interconnect

There are a number of “complications” that must be addressed in order to use mesa interconnect as described in this paper. The primary ones are: (1) MOS devices stability, (2) packaging issues, and (3) contamination issues. We will consider the problem posed in each of these areas.

To date all MOS devices have been fully covered with a metal mask in order to act as an etch shield. This mask covers the devices and extends several microns beyond each device. However, in a modern digital design, all metal layers are normally used to distribute signals, clocks, and power. Unless the power lines can be cleverly used to cover devices, the need to shield MOS transistors will decrease the metal rout-





Figure 3: SEM of the two interconnect structures on the experimental IC. Oxide has been removed to form a mesa interconnect in both cases. (a) Top: metal-1 lines adjacent to each other. (b) Bottom: metal-1 and metal-2 lines running adjacent to each other. The horizontal line is a metal-3 power line running across the interconnect.

ing space available resulting in some increase in overall die size. One alternative is to not shield the devices. We are currently studying the effects of the oxide etch on unshielded MOS devices.

The etched dice must be packaged carefully. Simply putting them in a standard injection molded plastic package may result in an increase in capacitance as plastic may fill in the air spaces around the interconnect. One possible solution is to package these parts in a ceramic hermetically sealed package in which there is a natural air space above the die. However, the high cost of a ceramic package would significantly limit the applications to which mesa interconnect can be applied. We are currently exploring protective layers that would protect the surface of the mesa interconnect, thus making it compatible with standard plastic packaging.

Sputtering a thin (roughly 100nm) layer of silicon nitride back onto the etched metal lines addresses two problems. The first problem is that of conductive contaminants landing on the surface of the die and causing shorts between metal lines. The second concern is that without the mechanical support of adjacent insulators, the electromigration limit of the Al lines might be significantly reduced. Coating the lines with silicon nitride again fully encases the Al lines within an insulating layer. However, because the nitride layer is thin, we still need to characterize the resulting electromigration limits.

6. Conclusions

We have demonstrated a simple post-processing fabrication methodology that can modify the interconnect in a standard foundry CMOS IC to form a mesa interconnect with roughly 40% less capacitance. In this paper we have demonstrated that the capacitance decreases resulting from oxide etching to form mesa interconnect can be predicted with reasonable accuracy by 2D electrostatic field analysis. The primary source of error appears to be inaccurate modeling of the exact geometry of the conductors and insulators. Finally, we have performed an initial fabrication demonstration with a few simple interconnect structures. From these measurements, we conclude that mesa interconnect does have the potential to significantly reduce capacitance, decreasing both the power-delay product and the energy-delay product of CMOS ICs. Further, we believe that these results should encourage additional work to overcome the remaining obstacles to the commercial application of mesa interconnect.

7. Acknowledgments

We would like to acknowledge the invaluable help of both Dr. Herman Schmit and David Collier in testing out our ideas. In addition, the authors would like to thank the SLPE committee for numerous valuable comments on our original abstract. This work was supported in part by the Defense Ad-

vanced Research Projects Agency under Order Number: A564, Amendments #6 & #8. The United States government has certain rights to this material.

8. References

- [1] L.R. Carley, "Methodology for reducing capacitance in integrated circuit interconnect structures", U.S. patent pending
- [2] G. K. Fedder, S. Santhanam, M. L. Reed, S. C. Eagle, D. F. Guillou, M. S.-C. Lu, and L. R. Carley, "Laminated high-aspect-ratio microstructures in a conventional CMOS process," *Proceedings of the IEEE 9th International Workshop on Micro Electro Mechanical Systems*, San Diego, CA USA, Feb. 11-15, 1996
- [3] L.R. Carley, M.L. Reed, G.K. Fedder, and S. Santhanam, "Microelectromechanical data storage apparatus and method for fabricating same", U.S. patent pending
- [4] L. Wissel, D. Stout, and N. Buck, "Gate-array library design using local interconnect", *Proc. Custom Integrated Circuits Conf.*, pp. 509-512, San Diego, CA USA, May 1996
- [5] M. Parameswaran, H. P. Baltes, L. Ristic, A. C. Dhaded, and A. M. Robinson, "A new approach for the fabrication of micromechanical structures," *Sensors and Actuators*, vol. 19, no. 3, pp. 289-307, Sept. 1989.