1 Description of the Omnilab Setup

This was very straightforward, and followed the documentation handed out in class very closely. Attached files will show which of the chip inputs/outputs were attached to which of the Omnilab stimulus outputs and analyzer inputs.

2 Testing

The testing of the chip was relatively unsophisticated. There was no aspect of our design which led us to suspect that a particular input or class of inputs would provide any more liklihood of chip failure than any other input or class of inputs, so we essentially selected random data and random keys and tested those. The command files used to do the testing are attached.

Two of the chips were completely functional and two of the chips were completely non-functional. All test vectors that were used to stimulate the chips which we labelled as functional produced the output that we expected. Again, while it is true that we made no attempt to test all possible inputs exhaustively (and indeed, this would not be feasible given that the data is 16 bits and the key is 14 bits, resulting in a huge number of possible inputs), there is nothing about our design that would suggest that any input or class of inputs would behave differently from any other input or class of inputs. Therefore, while we are away of the fallacy of proof by example, we are quite willing to declare success for the two functional chips.

The two non-functional chips are a different story. In both of them, seven of the eight outputs are stuck either at Vdd or GND, regardless of what the inputs are doing. Examination of the dies themselves shows no obvious fabrication errors, and we do not know exactly what went wrong. However, it is the case that we neglected to size up our buffers when they were driving lots of inputs. It seems possible to us that the stuck-at faults are the result of an under-sized buffer attempting to drive an unacceptably high number of inputs, resulting in self-destruction and shorting to either Vdd or GND. If this scenario is what actually happened, then it was just dumb luck that we got any functional chips at all. In addition, we should be able to destroy the functional ones with further use, although we did not attempt to do this.

Using Omnilab, the maximum clock frequency that we observed on the functional chips without output errors was 17 MHz. However, since we were using a two-phase non-overlapping clock, this means that the actual maximum clock frequency is essentially 4.25 MHz. We had projected a maximum clock frequency of about 25 MHz, so results this low are a bit dissappointing. However, since you warned only to expect a correct clock frequency within an order of magnitude, we do not view this as being unacceptably low.

The Rowboat Encryption Chip

Frank Alejano, John Cusey, and Chris Pickett
ELEC423
May 1, 1996