EZ Encoding: A Class of Irredundant Low Power Codes for Data Address and Multiplexed Address Buses

Yazdan Aghaghiri
University of Southern California

Farzan Fallah
Fujitsu Labs of America, Inc.

Massoud Pedram
University of Southern California

Abstract *- In this paper, we introduce a class of irredundant low power encoding techniques for memory address buses. For a data address bus, the proposed encoding techniques make use of two working zones in the memory address space, whereas for a multiplexed data and instruction address bus, up to four working zones can be supported. The zones are dynamically updated to increase the saving in switching activity. Our techniques decrease the switching activity of data address and multiplexed address buses by an average of 55% and 77%, respectively, up from 25% and 64% achieved by previous methods.*

1. EZ Encoding

In our proposed techniques each address is encoded based on the zone in which it is accessed. When we say addresses reside in the same zone we mean that they are pretty close to each other, for instance, all instructions that access the memory for stack manipulation are in the same zone.

Our first encoding technique, named EZ for Enhanced Zone encoding, uses two registers, one associated with each of two working zones. To encode an N-bit address, its offset is computed with respect to both zone registers. The address encoding is done using the zone that yields the smaller offset. This offset is translated to N-1 bits and then concatenated with a single bit which identifies the zone register, that is, the encoding is done by using N bits and there is no need to provide an additional line for identifying the zone register that is used. The receiver detects the working zone with a simple computation. Values stored in the zone registers are changed dynamically to update the zones. More precisely, the last accessed address is written to the zone register that was used for the encoding. Zone registers can reside anywhere in the address space. The contribution of this technique is to introduce a one-to-one mapping between binary representation and the new representation that consists of an (N-1)-bit offset showing the distance to one of two arbitrary numbers and a single bit identifying which number was used.

Our second method, named EZ2, uses less complex encoding and decoding hardware. In EZ2, the two zone registers' values cannot be in the same half of the memory address space, i.e., one of the values should be in the lower half, whereas the other one should be in the upper half. The MSB of the address determines the zone register to be used for encoding the address. To further reduce the complexity of the encoder and the decoder, instead of using a subtractor to compute offsets, XOR gates are used.

Our final encoding technique, named EZ3, is a combination of the previous methods. It relies on four zone registers, two of which are in the upper half, whereas the other two are in the lower half of the address space. For addresses in each half of the memory, we use the EZ technique for encoding.

2. Results

Table 1 compares switching activity saving of our methods and working zone method [1]. For working zone method, we have used two registers while encoding data addresses and four registers while encoding multiplexed addresses to match the number of zones with the number of zones of our best encoding technique for that particular address bus.

	WZ	EZ	EZ2	EZ3
Data Address	24.9%	50.7%	**55.3%**	50.3%
Multiplexed Address	64.3%	43.4%	54.2%	**77.3%**

As one can see our methods perform better than working zone method. Despite this improvement, the hardware overhead of our methods is less than working zone method. Also, our methods are irredundant and do not require additional lines.

We used a 1.5v, 0.18u CMOS library to design our encoders and decoders. We used SPEC2000 address traces to compute the power consumption of an external bus with 10pF capacitance per line for 3.3v I/O voltage and 50MHz clock frequency. The power consumption in absence of encoding was 13.7 mW. The total power dissipation of the data address bus (including the power overhead of the encoder/decoder logic) when using EZ, EZ2 and EZ3 encodings were 8.15, 6.6 and 8.3 mW, respectively, which shows a considerable saving. The respective power dissipations of encoders were 0.67, 0.24, 0.76 mW.

3. Conclusion

We proposed three new low power bus-encoding techniques. Our techniques are irredundant meaning that they do not need any extra bus line. The proposed techniques are more effective than previous methods for reducing activity on a bus because they require simple encoders and decoders. The overhead of using these methods is negligible, that is, the power consumption of our encoders and decoders is 6% to 20% of the power saved on the bus.

4. References

1. E. Musoll, T. Lang, and J. Cortadella, "Exploiting the locality of memory references to reduce the address bus energy," *Proc. of Int'l Symp. on Low Power Electronics and Design*, pp. 202-207, Monterey CA, August 1997.