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 space whereas for a multiplexed data and instruction address bus, up to four working zones can be supported. The zones are dynamically updated to improve the power efficiency of the proposed encoding techniques. These 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 truncated in a special manner to N-1 bits and then concatenated with a single bit for 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 representaion that consists of an 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. However, it is not as general as EZ. The restriction is that the two zone registers cannot be both in the same half of the memory space, i.e., one of these numbers should be in the lower half whereas the other one should be in upper half of the memory. The MSB of the address determines the zone register to be used for encoding the address. Instead of offset calculation, we send the exclusive-or of the address with the corresponding zone register.

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

2. Results

We compared EZ encoding and its variants with different configurations of the working zone method proposed by Musoll et al [1]. Table below shows the transition savings of the various encoding methods for data address and multiplexed data and instruction address buses. We used a two register working zone for data addresses and a four register one for multiplexed addresses.

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

All three encoders and decoders were designed and their netlists were generated in BLIF format. The netlists were optimized using the SIS *script.rugged* and mapped to a 1.5-volt 0.18µ CMOS library using the SIS technology mapper. I/O voltage was assumed to be 3.3v. SPEC2000 address traces were fed into a gate-level simulation program called, *sim-power*, to calculate the power consumption of the encoders and decoders. The system clock frequency was set at 50 MHZ. We also calculated the power consumed on the data address bus when the encoded addresses are sent. In this calculation, we assumed that the bus capacitance is 10pF per line. The total power consumption of the data address bus in absence of any encoding was 13.7 mW for SPEC2000 benchmark programs. The encoder power dissipation overhead of EZ, EZ2 and EZ3 were 0.67, 0.24, 0.76 mW, respectively. 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. These results demonstrate the considerable power savings of the EZ code and its variants.

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 for our different techniques.

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