

Microcode enables programmability of (micro) architectural structures to enhance functionality and to apply patches to an existing design. As more features get added to a CPU core, the area and power costs associated with microcode increase. A recent Intel internal design targeted at low power and small footprint has estimated the costs of the microcode ROM to approach 20% of the total die area (and associated power consumption). Therefore, it is desirable to apply compression techniques to microcode.

Microcode poses unique challenges for compression due to the long instruction format, the hand-coded nature of the programs and the stringent performance requirements that require fast decompression. This paper describes techniques for microcode compression that achieve significant area and power savings, while presenting a streamlined architecture that enables high throughput within the constraints of a high performance CPU. The paper presents results for microcode compression on several commercial CPU designs which demonstrates compression ratios ranging from 50% to 62%.