

Code compression, initially conceived as an effective technique to reduce code size in embedded systems, today also brings advantages in terms of performance and energy consumption, due to an increase in the cache hit ratio. This paper proposes the design of a code decompressor engine for our dictionary-based method, embedding it into the Leon (SPARC V8) processor. Our design guarantees that the processor cycle-time is maintained and the decompression is performed on-the-fly. We have achieved a functional implementation on a FPGA, with compression ratios ranging from 72% to 88%, performance improvement as high as 45% and a reduction on energy consumption reaching 35%, validated through two real-world benchmarks suites: MediaBench and MiBench. We also explore some trade-offs between compression ratio and performance