

Exploiting Idle Register Classes for Fast Spill Destination

Fang Lu, Lei Wang, Xiaobing Feng, Zhiyuan Li, Zhaoqing Zhang

Abstract. On today's microprocessor, there often exist a few different types of registers, e.g. general purpose registers and floating point registers. A given program may use one type of registers much more than the other. This creates an opportunity to employ the less used registers as the spill destination for the more often used ones. In this paper, we present a code optimization method named *idle register exploitation* (IRE) to exploit such opportunities. We have developed a model, called the *IRE model, or IREM*, to guide the static analysis for the cost of using IRE versus spilling to the stack. On a microprocessor that has fast data paths between different types of registers, We find that IRE method speeds up the execution of the SPECint benchmark suite in the range of 1.7 to 10%. In contrast, on microprocessors whose data transfer is much less efficient, the speedup is greatly reduced, and performance degradation may even result. This result argues strongly for the adoption of fast data paths between different types of registers for the purpose of reducing register spills, which is important in view of the increased significance of memory bottlenecks on future microprocessors.

Keywords: spilling cost, register conversion