

# Code Crumpling: A Straight Technique to Improve Loop Performance on Cache Based Systems

S. Seidl

Bologna, September 1999

## Abstract

While parallelization has made important progress during the last few years, the sequential performance relatively stagnates, although we have taken the 1 GHz hurdle. Moreover, we have to put tens of processors into action to be competitive with vector machines. Once an application is analyzed by means of performance tools like Vampir, the user is typically faced with some loops which consume most of the CPU time.

To improve loop performance, the present paper assumes that it is just enough in many cases to apply a generalized blocking technique with respect to the size of one primary cache line, or two of them occasionally. Its primary goal is to maximize the number of user operations divided by the expected value for the number of drawn primary cache lines. Another experience is that the inner loop body has to be well-stuffed and has to have as little memory write accesses as possible.

More than 50 Fortran/C example loops have been investigated so far, and essential results are presented. The performance in speed correlates with results which we have seen for routines in the BLAS libraries on different platforms. Finally, because of its clear goals, this procedure can also easily be taught.