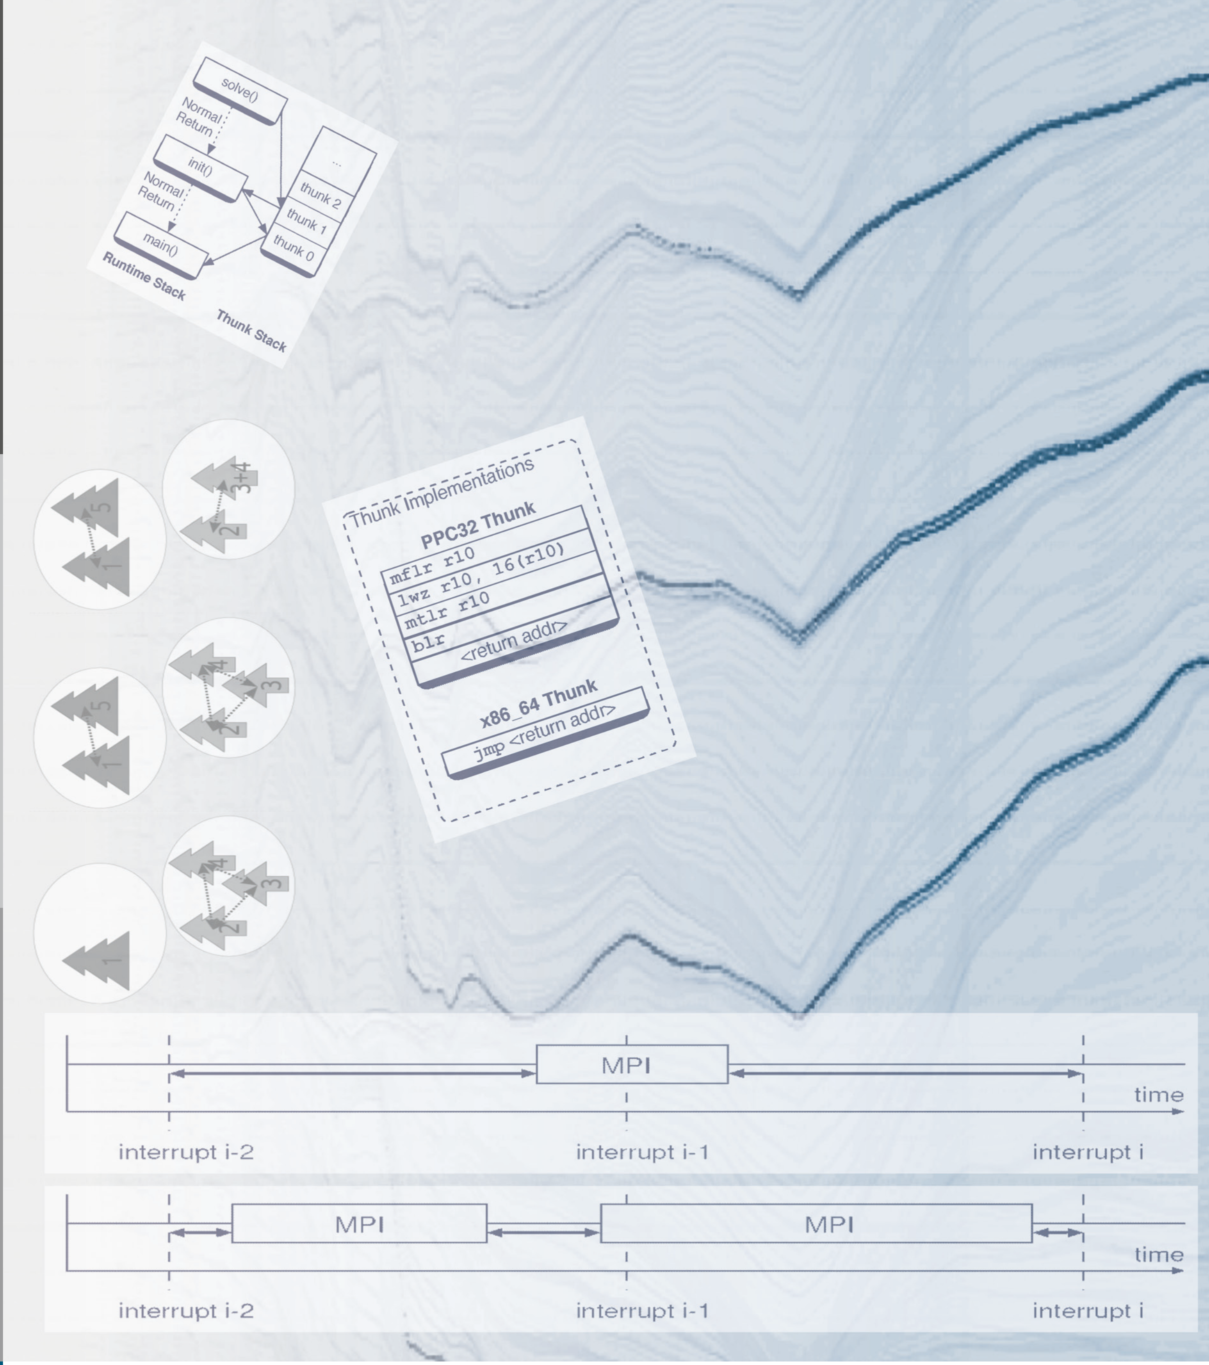

Capturing Parallel Performance Dynamics

Zoltán Péter Szebenyi

AI ces

JÜLICH
FORSCHUNGSZENTRUM

Forschungszentrum Jülich GmbH
Institute for Advanced Simulation (IAS)
Jülich Supercomputing Centre (JSC)

Capturing Parallel Performance Dynamics

Zoltán Péter Szebenyi

Schriften des Forschungszentrums Jülich
IAS Series Volume 12

ISSN 1868-8489 ISBN 978-3-89336-798-6

Contents

Runtime call-path profiling is a conventional, well-known method used for collecting summary statistics of a parallel application's execution such as the time spent in different call paths of the code. However, these kinds of measurements give the user only a summary overview of the entire execution, without regard to changes in performance behavior over time. As present day scientific applications tend to be run for extended periods of time, understanding the patterns and trends in the performance data along the time axis becomes crucial.

As shown by our analysis of a representative set of scientific codes, profiling every iteration separately provides a wealth of new data that often leads to invaluable new insights. However, with the introduction of the time dimension, memory usage and file sizes grow considerably. To counter this problem, a low-overhead on-line compression algorithm was developed that exploits similarities between different iterations.

While standard, direct instrumentation, which is assumed by the initial version of the compression algorithm, results in fairly low overhead with many scientific codes, in some cases the high frequency of events makes such measurements impractical. Therefore, a hybrid solution was developed that seamlessly integrates sampling and direct instrumentation techniques in a single unified measurement, using direct instrumentation for message passing constructs, while sampling the rest of the code. Finally, the compression algorithm was adapted to the hybrid profiling approach, avoiding the overhead of pure direct instrumentation.

Evaluation of the above methodologies shows that our similarity-based compression algorithm provides a very good approximation of the original data with very little measurement dilation, while the hybrid combination of sampling and direct instrumentation fulfills its purpose by showing the expected reduction of measurement dilation in cases unsuitable for direct instrumentation.

This publication was written at the Jülich Supercomputing Centre (JSC) which is an integral part of the Institute for Advanced Simulation (IAS). The IAS combines the Jülich simulation sciences and the supercomputer facility in one organizational unit. It includes those parts of the scientific institutes at Forschungszentrum Jülich which use simulation on supercomputers as their main research methodology.

IAS Series
Volume 12
ISBN 978-3-89336-798-6

JÜLICH
FORSCHUNGSZENTRUM