

Felix Wolf



Jülich Supercomputing Centre
Forschungszentrum Jülich
52425 Jülich, Germany
Phone: +49 2461 61 1583
E-mail: f.wolf@fz-juelich.de

Biography

Since 2006	Assistant Professor, Computer Science Department, RWTH Aachen University, Germany
Since 2005	Head of the Research Group <i>Performance Analysis of Parallel Programs</i> , Jülich Supercomputing Centre, Germany
2003-2008	Adjunct Assistant Professor, Computer Science Department, University of Tennessee
2003-2005	Sr. Research Associate (later Research Scientist), Innovative Computing Laboratory, University of Tennessee
2003	Ph.D. Computer Science, RWTH Aachen University
2001	Coop Pre-Professional Engineer, IBM Research, Yorktown Heights, NY
1999-2003	Doctoral Researcher, Jülich Supercomputing Centre
1998	M.Sc. Computer Science, RWTH Aachen University

Research Topic: Performance Analysis of Large-Scale Simulations

The objective of my research group is to make the optimization of parallel applications both more effective and more efficient. Driven by the widespread adoption of multi-core architectures, supercomputer applications are required to harness much higher degrees of parallelism in order to satisfy their growing demand for computing power. However, with an exponentially rising number of cores the often substantial gap between peak performance and that actually sustained by production codes is expected to widen even further. As a consequence, the productive use of large-scale parallel systems increasingly depends on the availability of robust and powerful performance-analysis tools. But growing parallelism places stronger scalability demands not only on applications but also on the parallel programming tools themselves. To improve the efficiency of large scale applications and, thus, to expand their potential, my group develops scalable tools that collect relevant data on code performance and identify the causes of performance problems. At the center of our activities lies the development of Scalasca [1], a performance analysis tool that has been specifically designed for large-scale systems and that allows the automatic identification of harmful wait states in applications running on thousands of processors [2].

[1] www.scalasca.org

[2] B. J.N. Wylie, M. Geimer, F. Wolf: Performance measurement and analysis of large-scale parallel applications on leadership computing systems. *Scientific Programming* 16(2-3):167-181, 2008.