



# **EINLADUNG**

## **zum**

# **ZIH - KOLLOQUIUM**

**Title:**           **ScalaTrace: scalable compression and replay of communication traces**

**Referent:**   **Frank Müller**  
**North Carolina State University**

**Zeit:**           **Dienstag, den 16. Oktober 2007, 13:00 Uhr**  
**Ort:**            **Willers-Bau C 207**

**Abstract:**

Characterizing the communication behavior of large-scale applications is a difficult and costly task due to code/system complexity and their long execution times. An alternative to running actual codes is to gather their communication traces and then replay them, which facilitates application tuning and future procurements. While past approaches lacked lossless scalable trace collection, we contribute an approach that provides orders of magnitude smaller, if not near constant-size, communication traces regardless of the number of nodes while preserving structural information. We introduce intra- and inter-node compression techniques of MPI events and present results of our implementation for BlueGene/L. Given this novel capability, we discuss its impact on communication tuning and beyond. We further develop scalable time stamp encoding schemes for communication traces, show that timing behavior can be preserved during replay of communication events and discuss the inherent benefits for efficient timeline-based searches over our compressed trace representation. To the best of our knowledge, such a concise representation of MPI traces in a scalable manner combined with deterministic MPI call replay in a time-preserving manner is without any precedence.

**gez. Prof. Dr. Wolfgang E. Nagel**