TWINS: Server Access Coordination in the I/O Forwarding Layer

Jean Luca Bez
Informatics Institute
Federal University of Rio Grande do Sul
jean.bez@inf.ufrgs.br

Francieli Zanon Boito
Department of Informatics and Statistics
Federal University of Santa Catarina
francieli.boito@posgrad.ufsc.br

Lucas M. Schnorr
Informatics Institute
Federal University of Rio Grande do Sul
schnorr@inf.ufrgs.br

Philippe O. A. Navaux
Informatics Institute
Federal University of Rio Grande do Sul
navaux@inf.ufrgs.br

Jean-Francois Mehaut
Universite de Grenoble Alpes
INRIA
CEA LIG Laboratory
Grenoble, France
jean-francois.mehaut@imag.fr

Abstract

This paper presents a study of I/O scheduling techniques applied to the I/O forwarding layer. In high-performance computing environments, applications rely on parallel file systems (PFS) to obtain good I/O performance even when handling large amounts of data. To alleviate the concurrency caused by thousands of nodes accessing a significantly smaller number of PFS servers, intermediate I/O nodes are typically applied between processing nodes and the file system. Each intermediate node forwards requests from multiple clients to the system, a setup which gives this component the opportunity to perform optimizations like I/O scheduling. We evaluate scheduling techniques that improve spatiality and request size of the access patterns. We show they are only partially effective because the access pattern is not the main factor for read performance in the I/O forwarding layer. A new scheduling algorithm, TWINS, is presented to coordinate the access of intermediate I/O nodes to the data servers. Our proposal decreases concurrency at the data servers, a factor previously proven to negatively affect performance. The proposed algorithm is able to improve read performance from shared files by up to 28% over other scheduling algorithms and by up to 50% over not forwarding I/O.