



Survey on Efficient Scheduling Policies for Microsecond-Scale Tasks

Presented by Yuxuan Liu

Abstract - Modern datacenter applications often involve many short Remote Procedure Calls (RPCs) to other servers. These RPCs allow applications with large memory footprints to access memory on other servers, enable applications to leverage large amounts of compute over short timescales. The service times of these tasks grow ever smaller, and today are often a single microsecond or less.

Tasks with short service times are particularly vulnerable to latency inflation; even small overheads can increase the latency of a $1 \mu s$ task by an order of magnitude. This is problematic for today's applications, which depend on low latency both at the median and at the tail of the distribution (e.g., 99% latency). As a result, many techniques have been proposed to reduce the overheads of handling these short tasks. However, performance with respect to latency is not the only measurement. As Moore's Law slows, datacenter operators are increasingly concerned not just with providing low latency but also with achieving high CPU efficiency. Here comes with the optimization target: what is the best OS scheduling policy that balances both performance in microsecond-level latency and CPU efficiency? In this survey, a large number of state-of-the-art techniques will be discussed and analyzed to answer this question.