

Abstract

Techniques for predicting the efficiency of multi-core processing associated with a set of tasks with varied CPU and main memory requirements are introduced. Prediction of CPU and memory availability is important in the context of making process assignment, load balancing, and scheduling decisions in distributed systems. Given a set of tasks each with varied CPU and main memory requirements, and a multi-core system (which generally has fewer cores than the number of tasks), we provide upper- and lower-bound models (formulas) for the efficiency with which the tasks are executed. In addition, a model for average CPU availability is introduced from the empirical study for applications that require a single predicted value instead of bounds. To facilitate scientific and controlled empirical evaluation, real-world benchmark programs with dynamic behaviour (CPU and memory requirements change in a short interval of time) are employed on UNIX systems that are parameterised by their CPU usage factor and memory requirement.