

Enhancing Linux Process Scheduling with Machine Learning Techniques

P. Srinivasa Rao¹, G. Snehalatha², S. Gayani³, G. Shivani⁴, V. Varnika⁵

Professor, Computer Science Engineering¹

Student, Computer Science Engineering²⁻⁵

ACE Engineering College, Ghatkesar, India

Abstract: This study explores the use of machine learning (ML) to improve Linux process scheduling, focusing on predicting CPU burst times by analyzing process attributes. The objective is to reduce Turn-around-Time (TaT) by accurately forecasting burst times and adjusting time slices accordingly. The Linux Kernel scheduler (v2.4.20-8) is modified to implement this predictive scheduling. Using the Waikato Environment for Knowledge Analysis (Weka), an open-source ML tool, we evaluate various algorithms to determine the most effective method for this task, with the C4.5 Decision Tree algorithm yielding the best results. The modified scheduler reduces TaT by 1.4% to 5.8% due to fewer context switches, demonstrating the potential of predictive scheduling in enhancing operating system performance.

Keywords: CPU Scheduling, Linux Kernel, Machine Learning, Special Time Slice (STS), Process Classification, Ensemble Learning, XGBoost, Random Forest, Stacking Classifier, SMOTE, Resthisce Allocation, Throughput, Latency, Cloud Computing, Edge Computing, High-Performance Computing

