Exploiting Synchronization-aware Transformer for Aligning Large-Scale MPI Traces

Zhibo Xuan, Xin You, Hailong Yang, Haoran Kong, Jingqi Chen, Tianyu Feng, Zhongzhi Luan, Yi Liu, Depei Qian

Frontiers of Computer Science, DOI: 10.1007/s11704-026-51562-6

Problems & Ideas

- Problems of Large-Scale MPI Trace Timeline Alignment:
  - Time Skew Introduces Logical Errors in Performance Analysis.
  - Existing Methods Fail in Scalability or Accuracy.
  - Neglect of MPI Semantic Context.
- Ideas: Propose STAR, a large-scale Synced Trace Timeline Aligner tool designed to tackle the time skew problem for large-scale parallel programs.

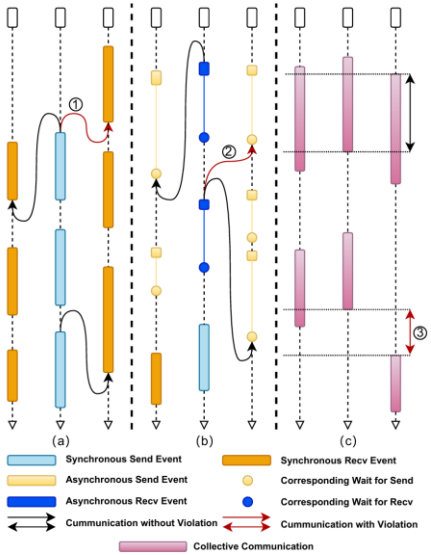

Time skew

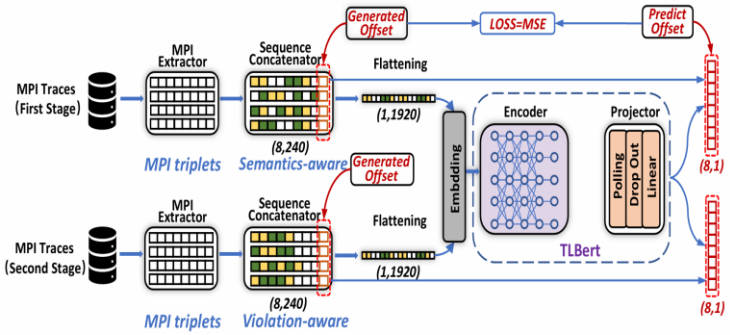

Model training

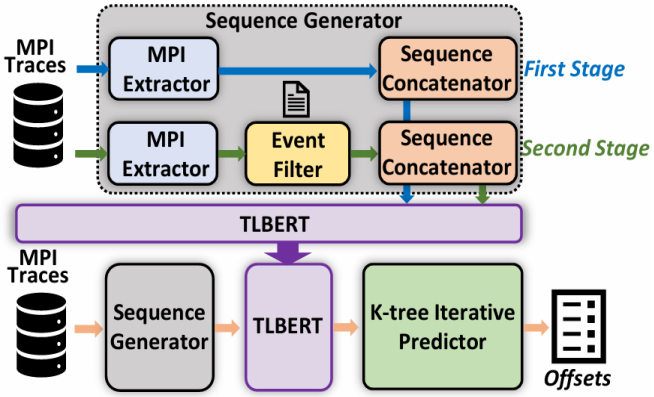

STAR overview

Main Contributions

- Contributions:
  - We propose the first Transformer-based model for MPI trace time offset prediction and design a divide-and-conquer predictor that scales TLBERT to large-scale scenarios, solving the high computational complexity of existing methods like interpolation-based approaches.
  - We implement STAR and evaluate STAR on up to 4,096 processes, demonstrating correctness, scalability, and low overhead, and show superiority over methods like Scalasca.

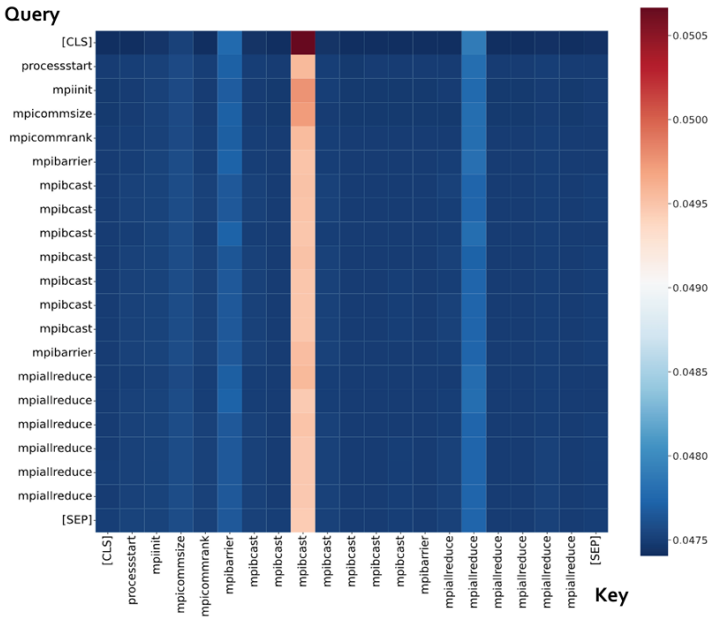

Visualization of Attention Map of TLBERT.

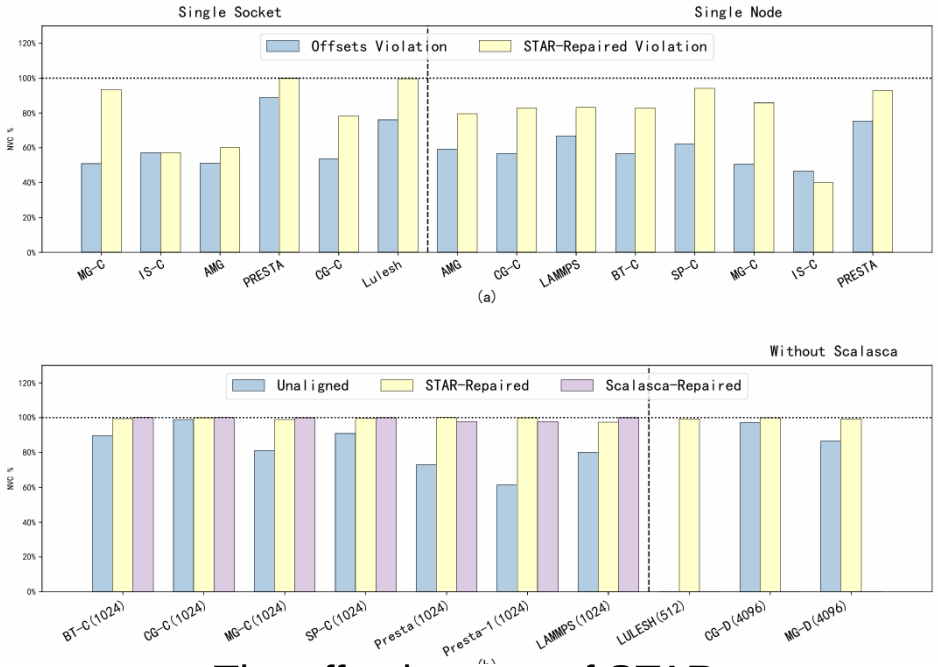

The effectiveness of STAR