

Speculative Monte-Carlo Tree Search

Scott Cheng^{1,2}, Mahmut Taylan Kandemir¹, Ding-Yong Hong²

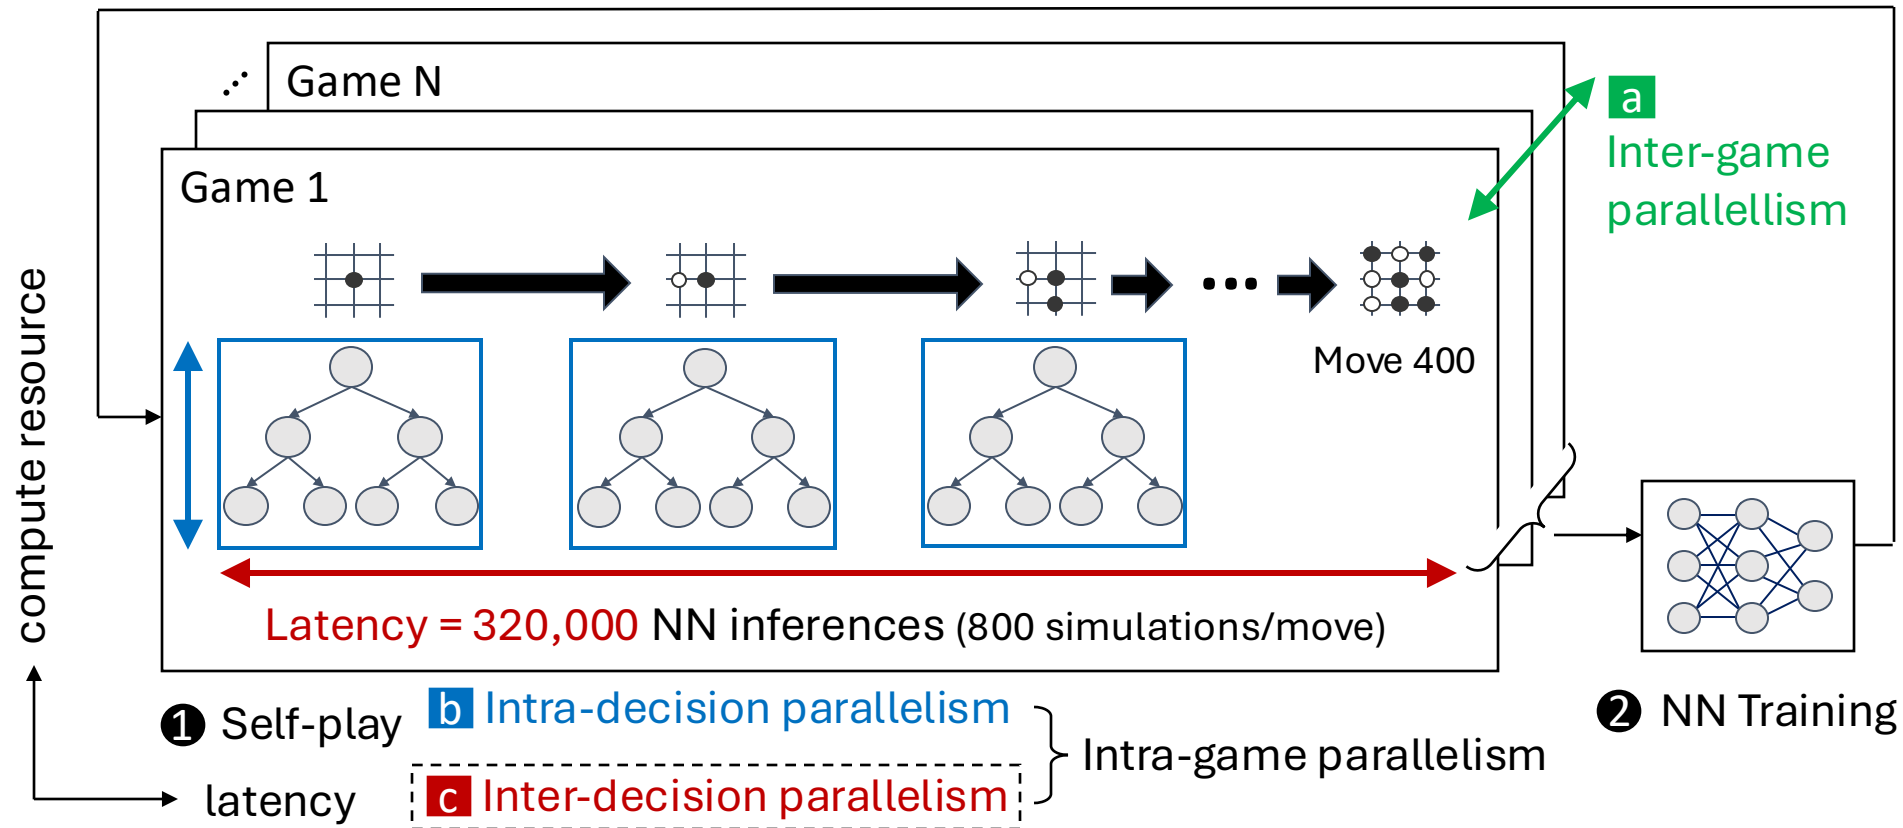
¹ The Pennsylvania State University, USA

² Academia Sinica, Taiwan





Sequential MCTS limits AlphaZero training latency

AlphaZero has achieved breakthroughs in chemistry, medicine, and many other fields, but the training is considerably time-consuming.



Speculative Parallelism in MCTS

Motivation: partial MCTS can reasonably predict the full MCTS results.



Speculative MCTS Analysis


Based on the speculative pipeline design, we can estimate the latency by solving Markov chain:

$$\begin{cases} \mathbf{q}\mathbf{M} = \mathbf{q} \\ \|\mathbf{q}\| = 1 \end{cases} \iff$$


\mathbf{q} : steady-state distribution

\mathbf{M} : transition probability

(a) Original (0-step look-ahead)

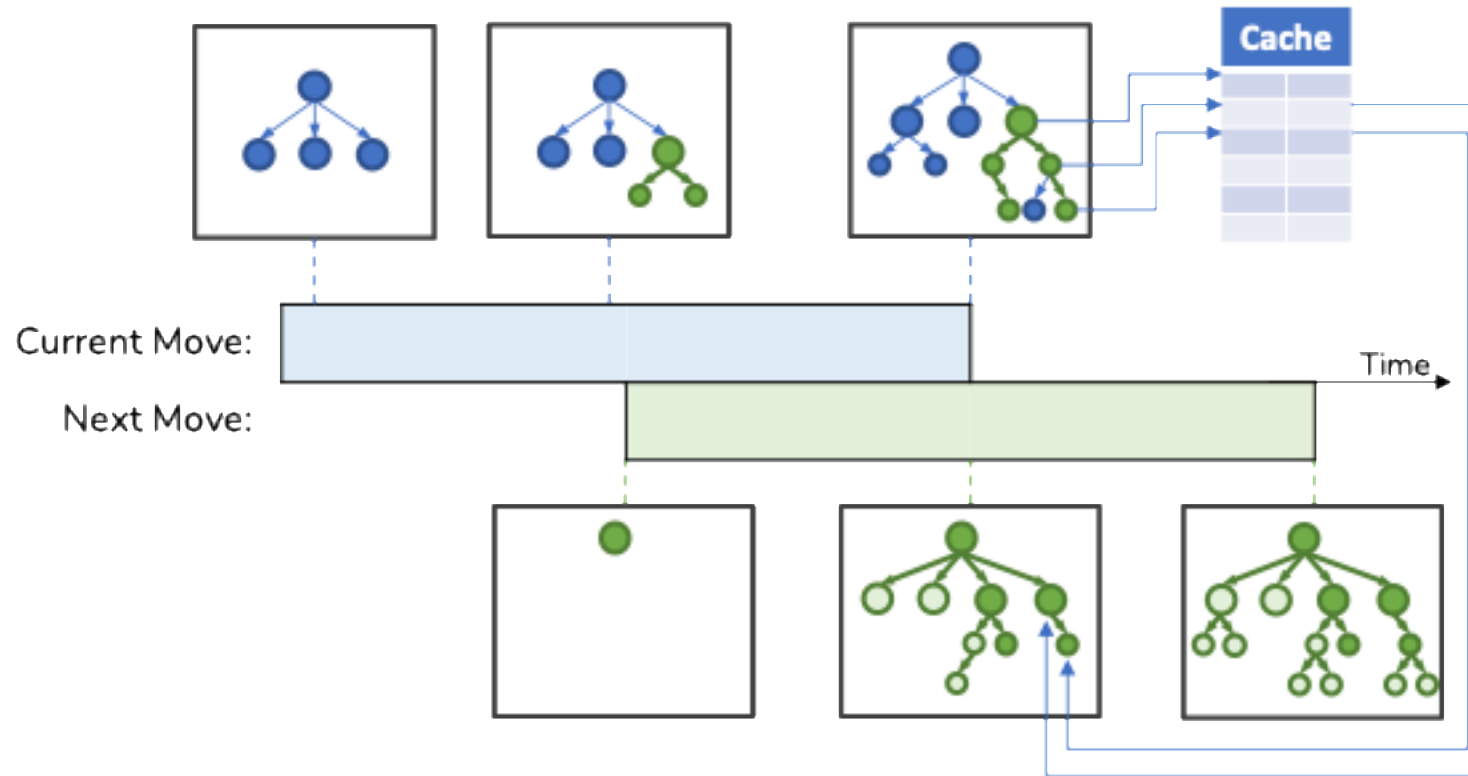


(b) Speculation (1-step look-ahead)



Synergizing Speculation and Caching

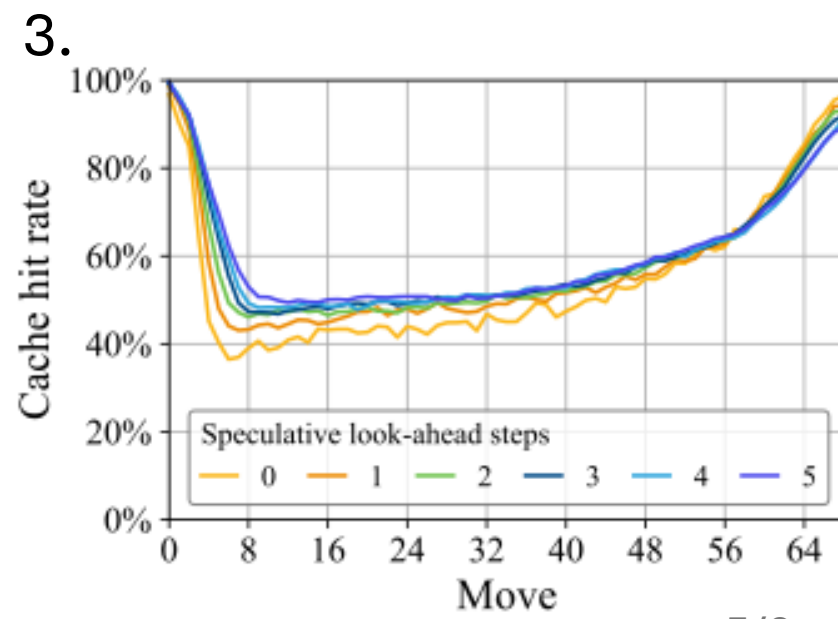
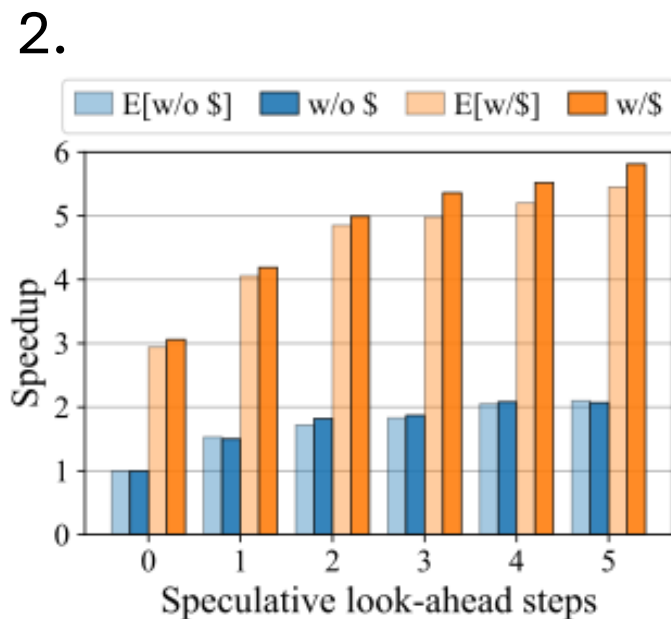
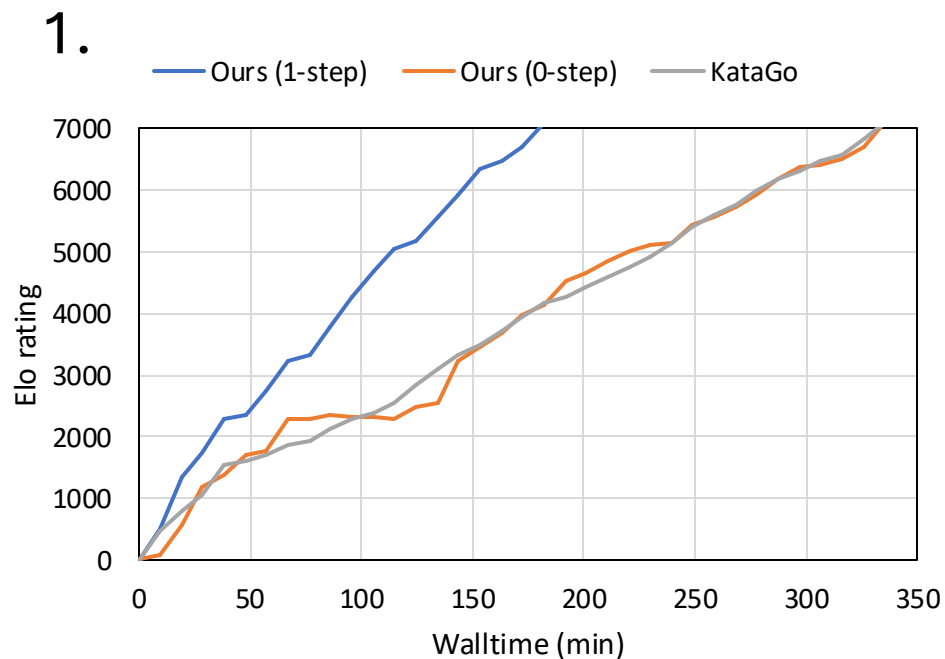
When speculation fails to predict the next move, the inference results remain valuable in the neural network cache.



Source of cache hits: (1) intra-game and (2) inter-game

Evaluations

1. End-to-end training: 1.91x speedup compared to the state-of-the-art KataGo program
2. Training latency estimation with an average RMSE of 22.2
3. Insights into speculation with caching and more ...



For More Details ... Please check our paper!

Why is inter-decision parallelism necessary beyond inter-game and intra-decision parallelism?

Why do you choose speculation to achieve inter-decision parallelism?

Can you show a design alternative for Speculative MCTS?

What are the design choices for Speculative MCTS?

What are the limitations of Speculative MCTS?



Thank you

Acknowledgements to:

- National Science and Technology Council, Taiwan
- National Center for High-performance Computing, Taiwan
- Mr. Meng-Yu Tsai