



Contribution ID: 40

Type: **Contributed Talk**

Practical Microsecond Real-Time Reinforcement Learning

Wednesday, February 7, 2024 1:00 PM (30 minutes)

Reinforcement Learning (RL) has been successfully applied to a wide range of problems. When the environment to control does not exhibit stringent real-time constraints, currently available techniques and computational infrastructures are sufficient. At particle accelerators, however, it is often possible to encounter stringent requirements on the time necessary for an action to be chosen, that in some extreme cases can fall in the microsecond scale.

These challenging conditions also present some benefits. For instance, the data throughput of the real-world environment can be orders of magnitude greater compared to a simulation. This opens the possibility of online training without the issues linked to transferring a simulation-trained agent to the real world.

In this contribution, real-time constraints and how they affect RL algorithms will be introduced, followed by a description of FPGAs and heterogeneous hardware platforms. This is then used to motivate the architecture of the state-of-the-art KINGFISHER RL system. Finally, an in-depth discussion of the use-cases where this approach can be beneficial will be provided, together with basic guidelines for structuring RL problems in a more hardware-friendly way.

Possible contributed talk

Yes

Are you a student?

Yes

Author: SCOMPARIN, Luca

Co-authors: SANTAMARIA GARCIA, Andrea (KIT); CASELLE, Michele (KIT); XU, Chenran (IBPT)

Presenter: SCOMPARIN, Luca

Session Classification: Contributed Talks