Contribution ID: **41** Type: **Poster**

Designing a Gated Recurrent Unit in the Versal AI Engines

Tuesday, 31 March 2026 12:00 (2 hours)

This poster presents the design and implementation of a Gated Recurrent Unit (GRU) on Xilinx Versal AI Engines. We outline the mapping of GRU computations to the AI Engine architecture, discuss dataflow and parallelization strategies, and highlight performance considerations for efficient recurrent neural network inference. The design supports unquantized models by leveraging 32-bit floating-point datatypes on the AI Engines.

Student

Yes

Primary author: SAPKAS, Michail (UniPD - INFN Padova)

Presenter: SAPKAS, Michail (UniPD - INFN Padova)

Session Classification: Poster session