Tal El-Nun, **Stateful Dataflow Multigraphs: A Data-Centric Approach for Performance Portability on Heterogeneous Architectures**

Abstract:

The ubiquity of accelerators in high-performance computing has driven programming complexity beyond the skill-set of the average domain scientist. To maintain performance portability in the future, it is imperative to decouple architecture-specific programming paradigms from the underlying scientific computations. The talk highlights a data-centric approach for heterogeneous parallel programming: the DaCe framework and the Stateful DataFlow multiGraph (SDFG), a data-centric intermediate representation that enables separating program definition from its optimization. By combining fine-grained data dependencies with high-level control-flow, SDFGs are both expressive and amenable to program transformations, such as tiling and double-buffering. These transformations are applied to the SDFG in an interactive process, using extensible pattern matching, graph rewriting, and an IDE-embedded graphical user interface. The DaCe framework is shown on CPUs, GPUs, and FPGAs over various motifs - from fundamental computational kernels to graph analytics. DaCe and SDFGs deliver competitive performance, allowing domain scientists to develop applications naturally and port them to approach peak hardware performance without modifying the original scientific code.