CPU + GPU HETEROGENEOUS COMPUTING (THE LAB)

Ana Lucia Varbanescu, University of Amserdam, NL
a.l.varbanescu@uva.nl

With significant contributions by:
Jie Shen @ Delft Unviersity of Technology
Stijn Heldens @ U Twente

UvA Universiteit van Amsterdam

CPU + GPU

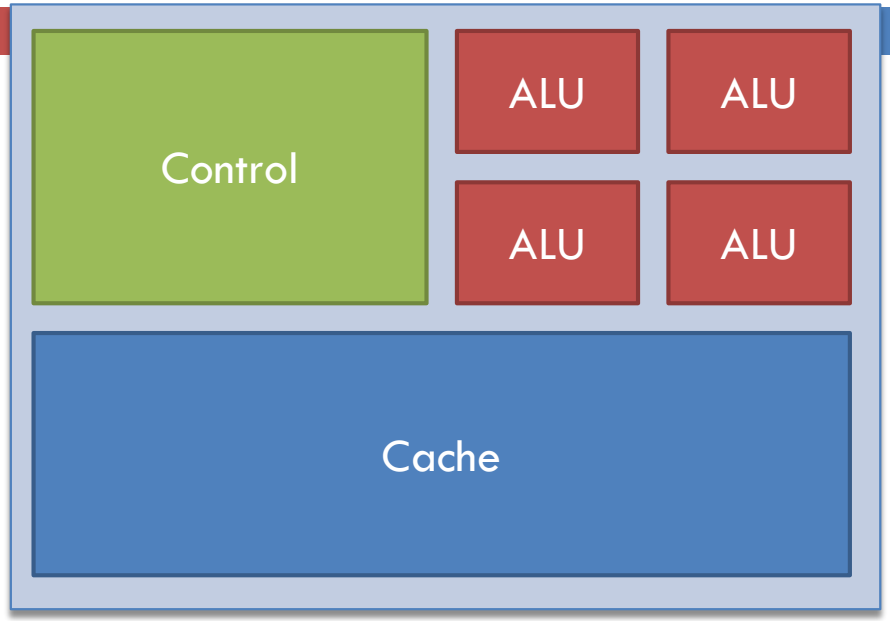

CPU

Low latency, high flexibility.
Excellent for irregular codes with limited parallelism.

GPU

High throughput.
Excellent for massively parallel workloads.

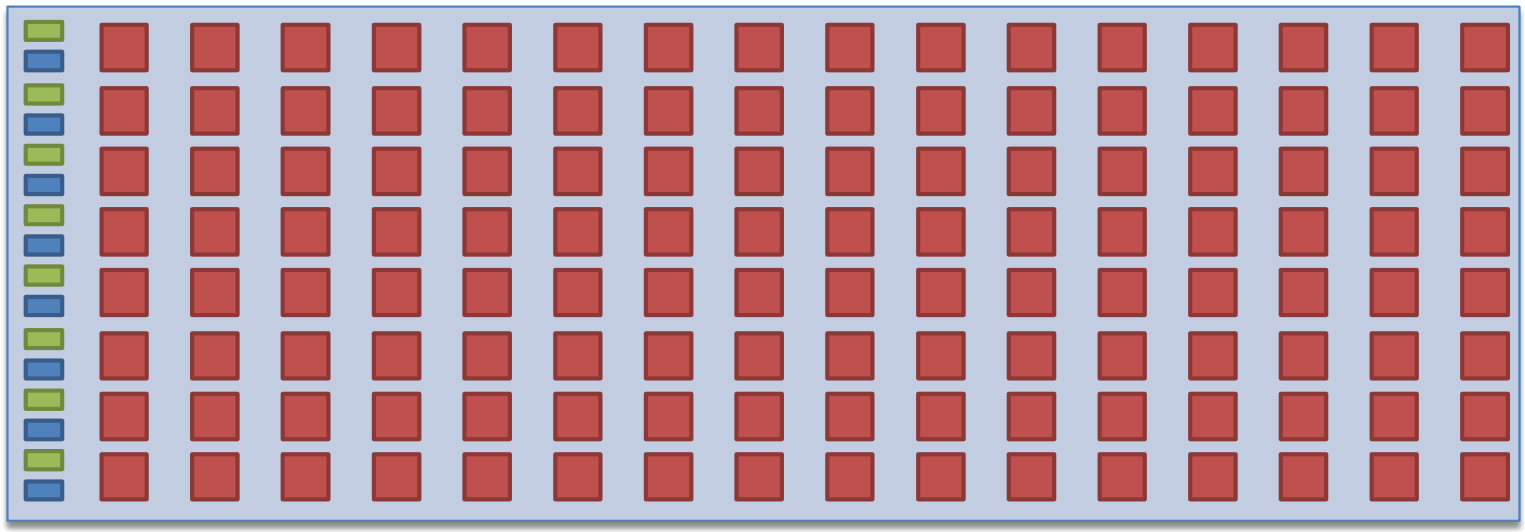

Workload

- DAG (directed acyclic graph) of "kernels"

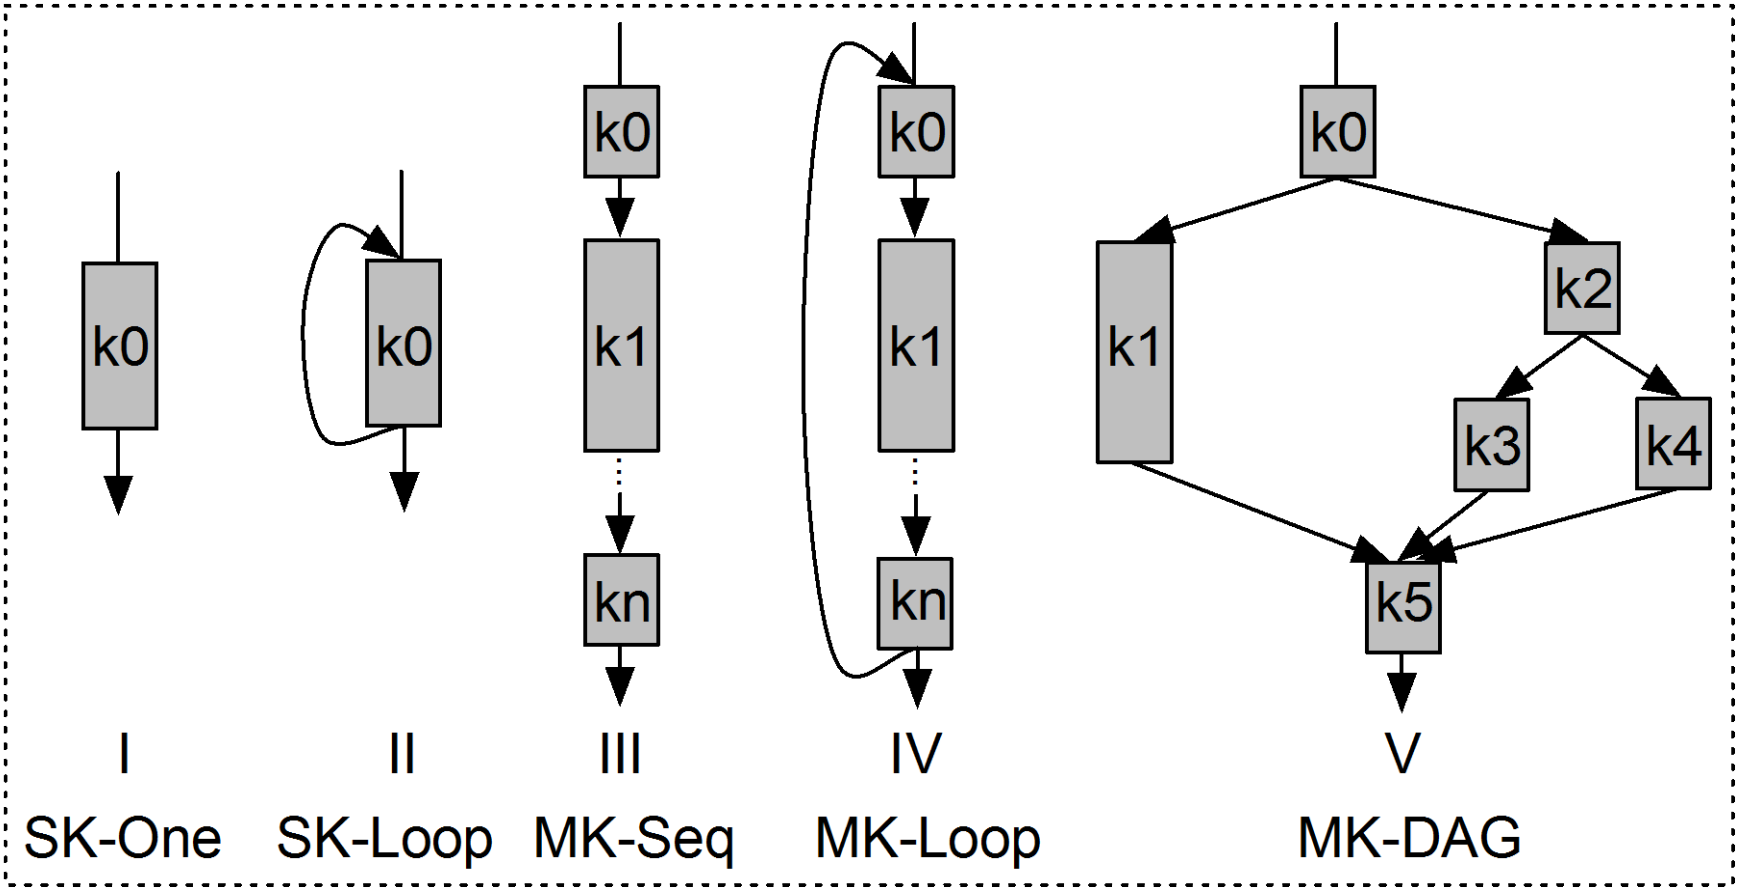

Two main challenges

- How to **program** the workload in a “user-friendly” and “heterogeneity-friendly” way
  - Unified programming models ?
  - Low- or high-level ?
  - …
- How to **split** the work between the different processors
  - Static
  - Dynamic
  - Hybrid … ?

(1) Heterogeneous Computing PMs

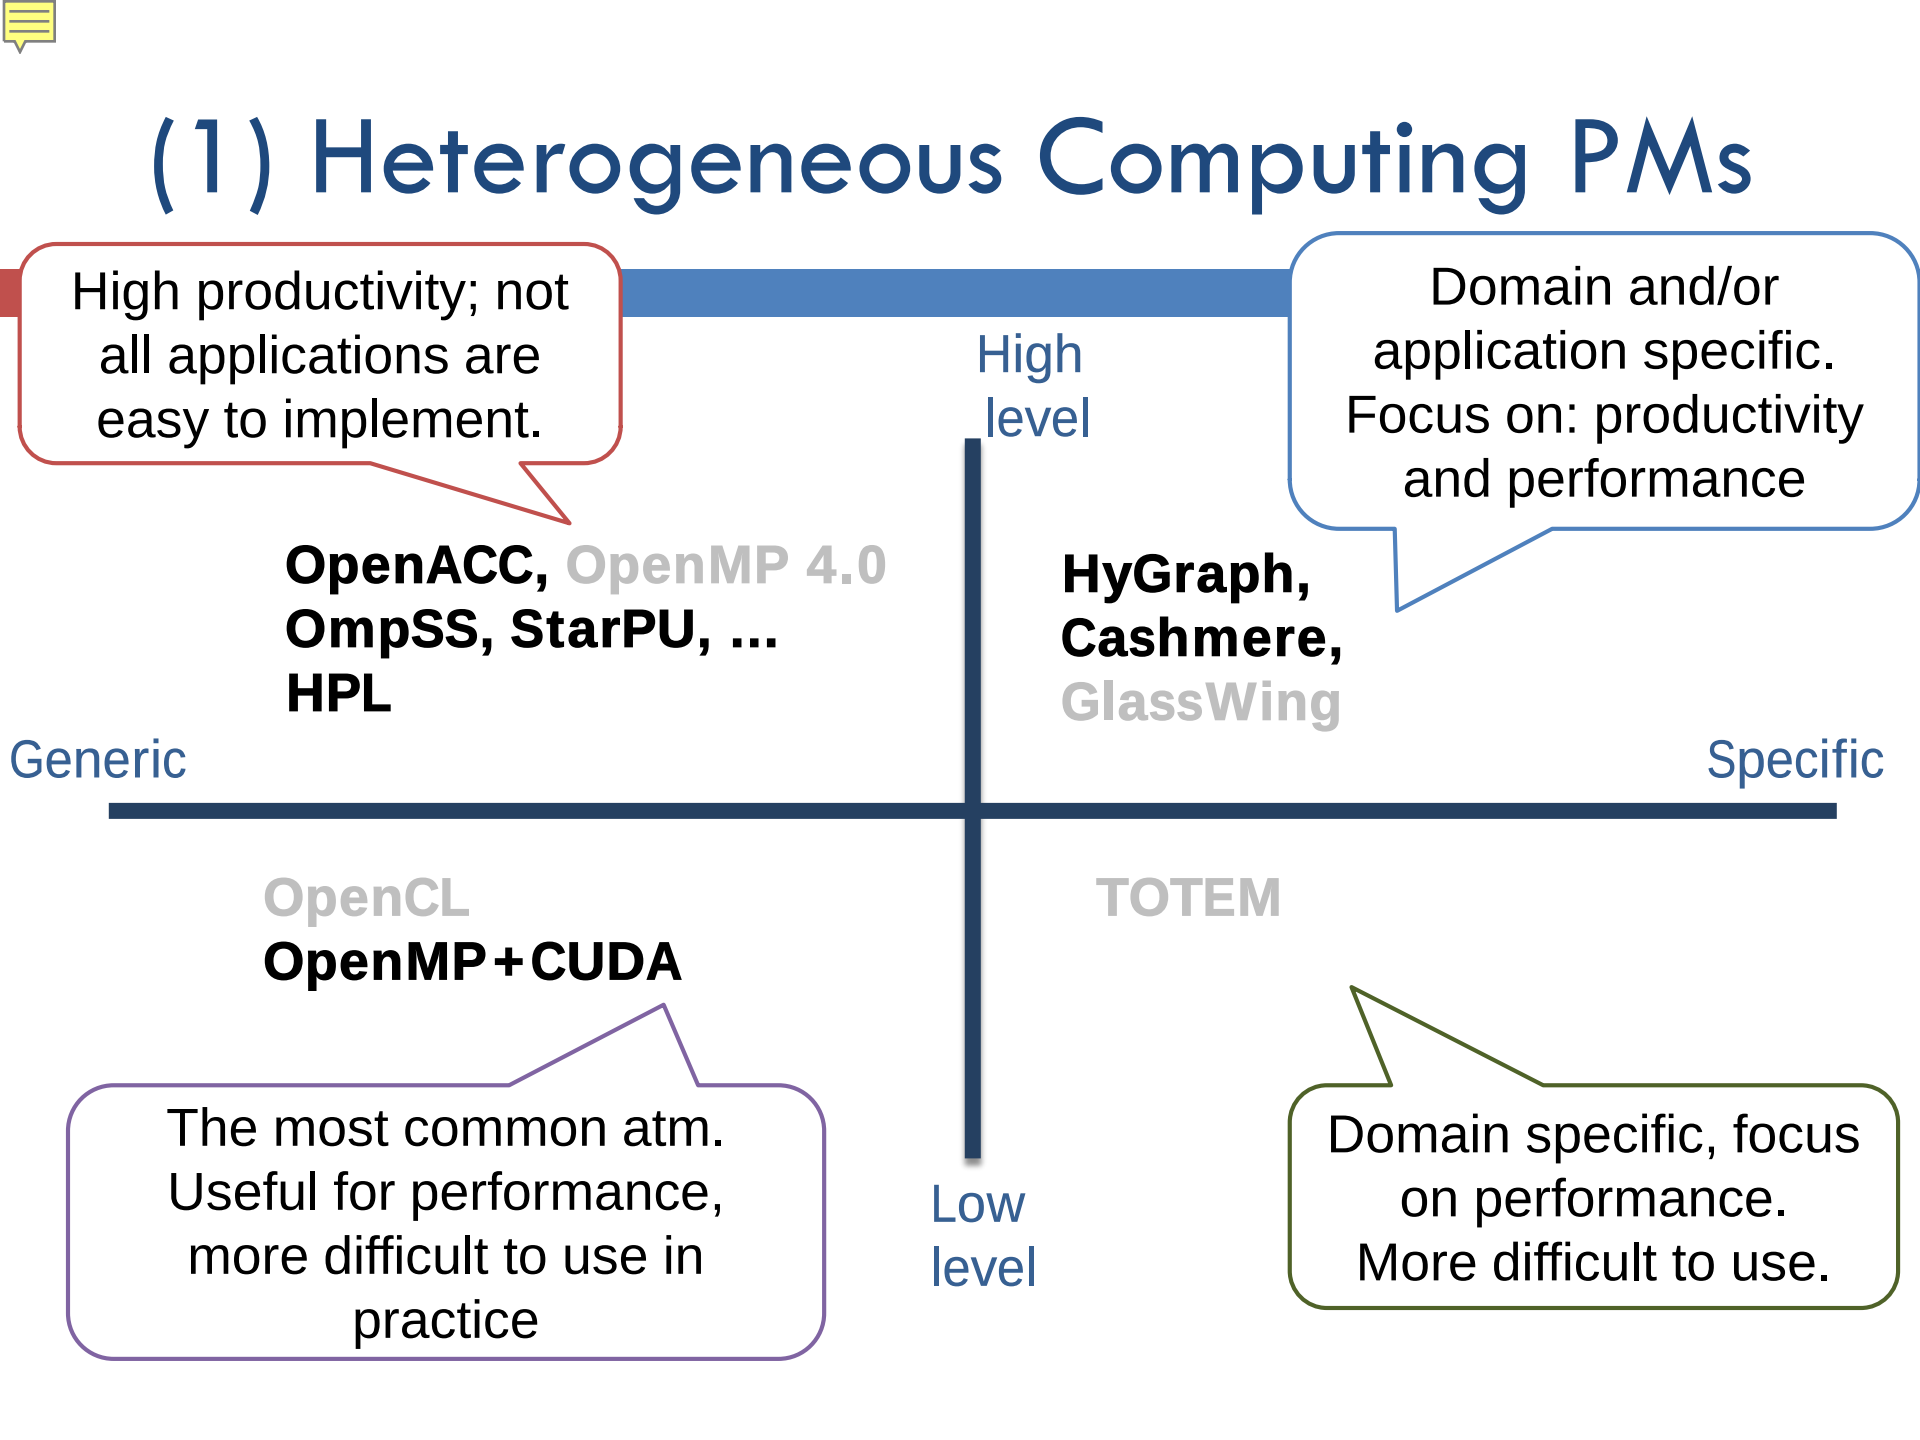

Heterogeneous Computing today

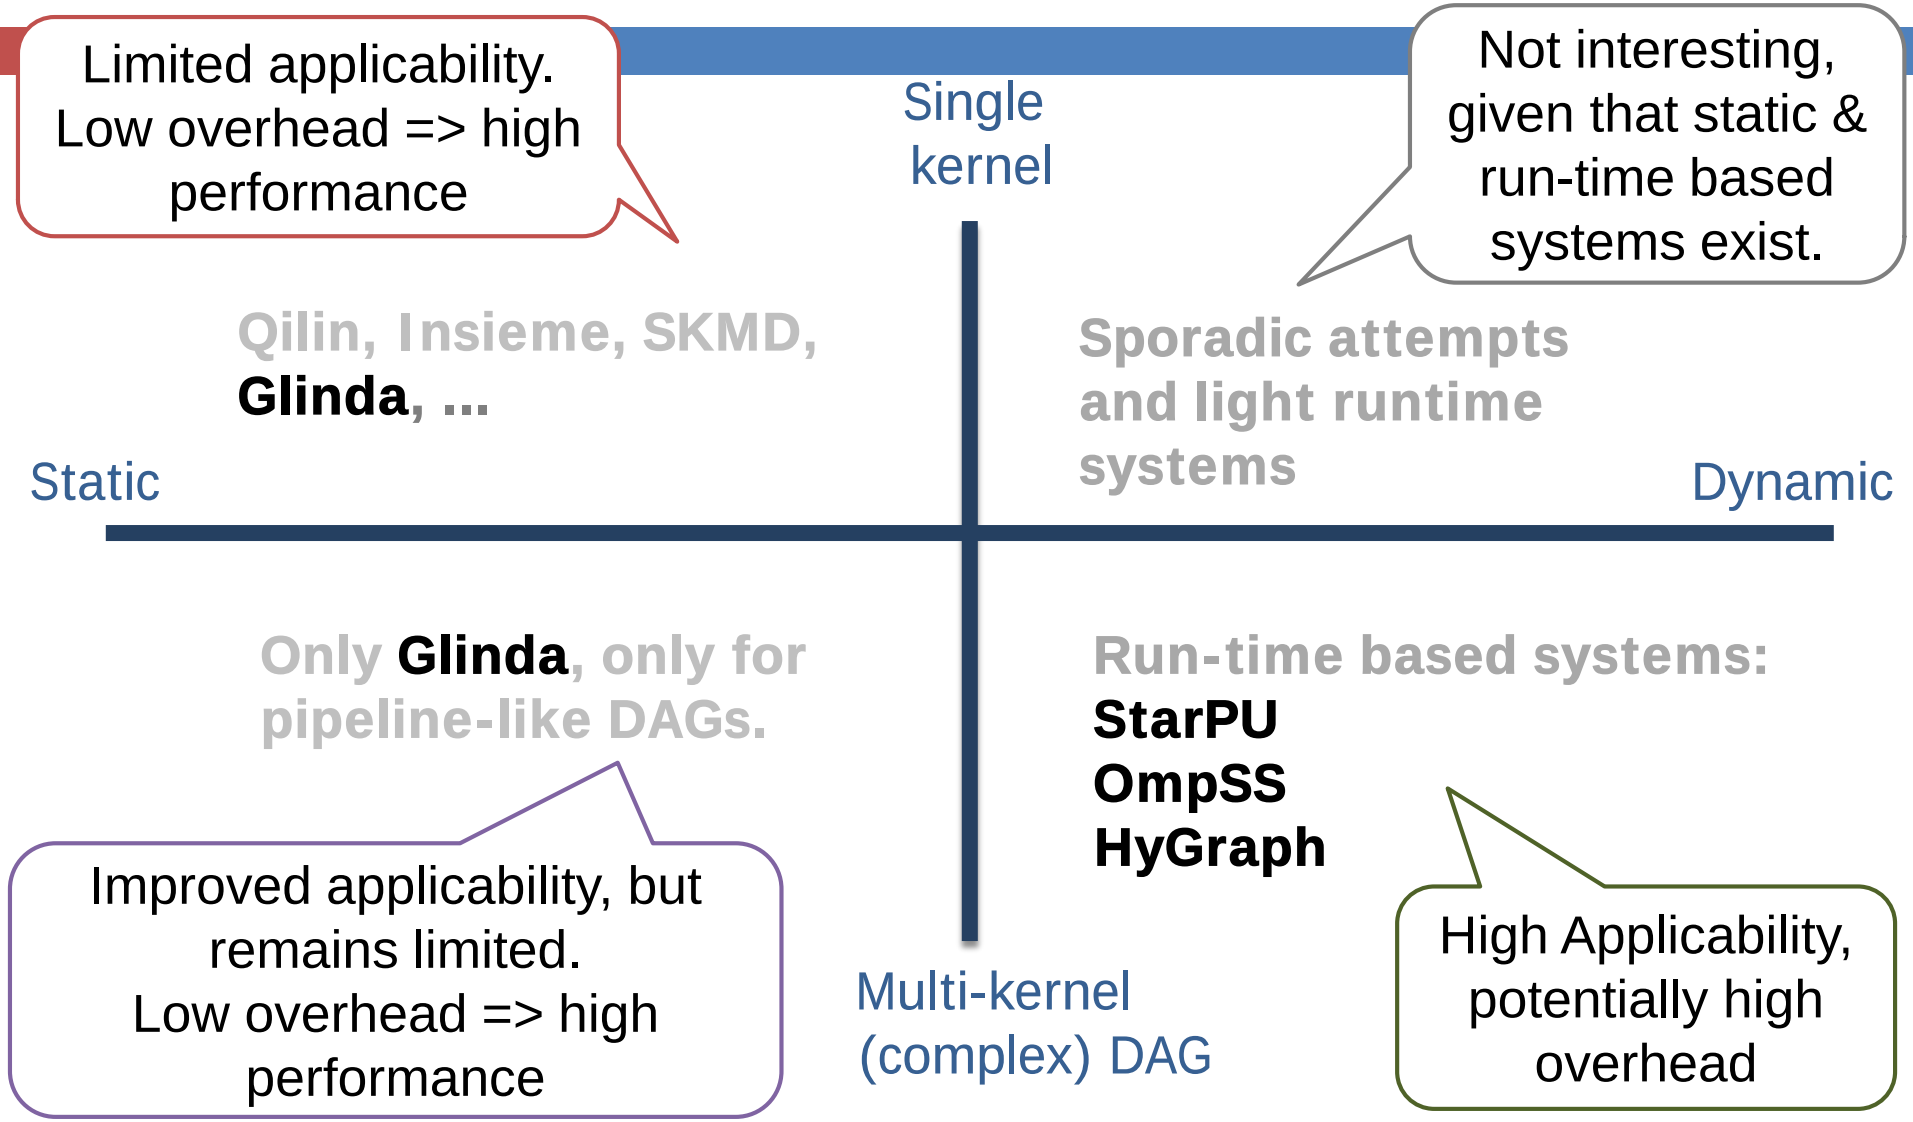

The lab

- Three application (classes)
  - Single-kernel: matrix multiplication
  - Multi-kernel: image processing pipeline
  - Domain-specific: graph processing

- Given code templates
  - Fill in the blanks
  - Compare & improve
  - Reflect

Practicalities

My email:
A.L.Varbanescu@uva.nl

- Where ?
  - DAS5 (http://www.cs.vu.nl/das5/)
    - Ask me for *temporary* accounts
- What ?
  - Code + basic report
- When ?
  - Deadlines given earlier
- What if …
  - "I want my own project?" – OK on a case-by-case basis
  - "I don't know programming?" – it will be difficult