INTRODUCTION

Accelerators & Coprocessors (A&C)

- Dedicated logic for specific workloads
  - In HPC this means: Flop/s, mem BW, high parallelism
- Tradeoffs
  - Limitations in general-purpose compute capabilities
  - Programs must typically be adapted to some extent
- Different families of technologies
  - GPGPU (Nvidia Tesla, AMD)
  - Manycores (**Intel MIC**, Adapteva)
  - FPGA (Convey etc.)

Accelerated systems in Top500

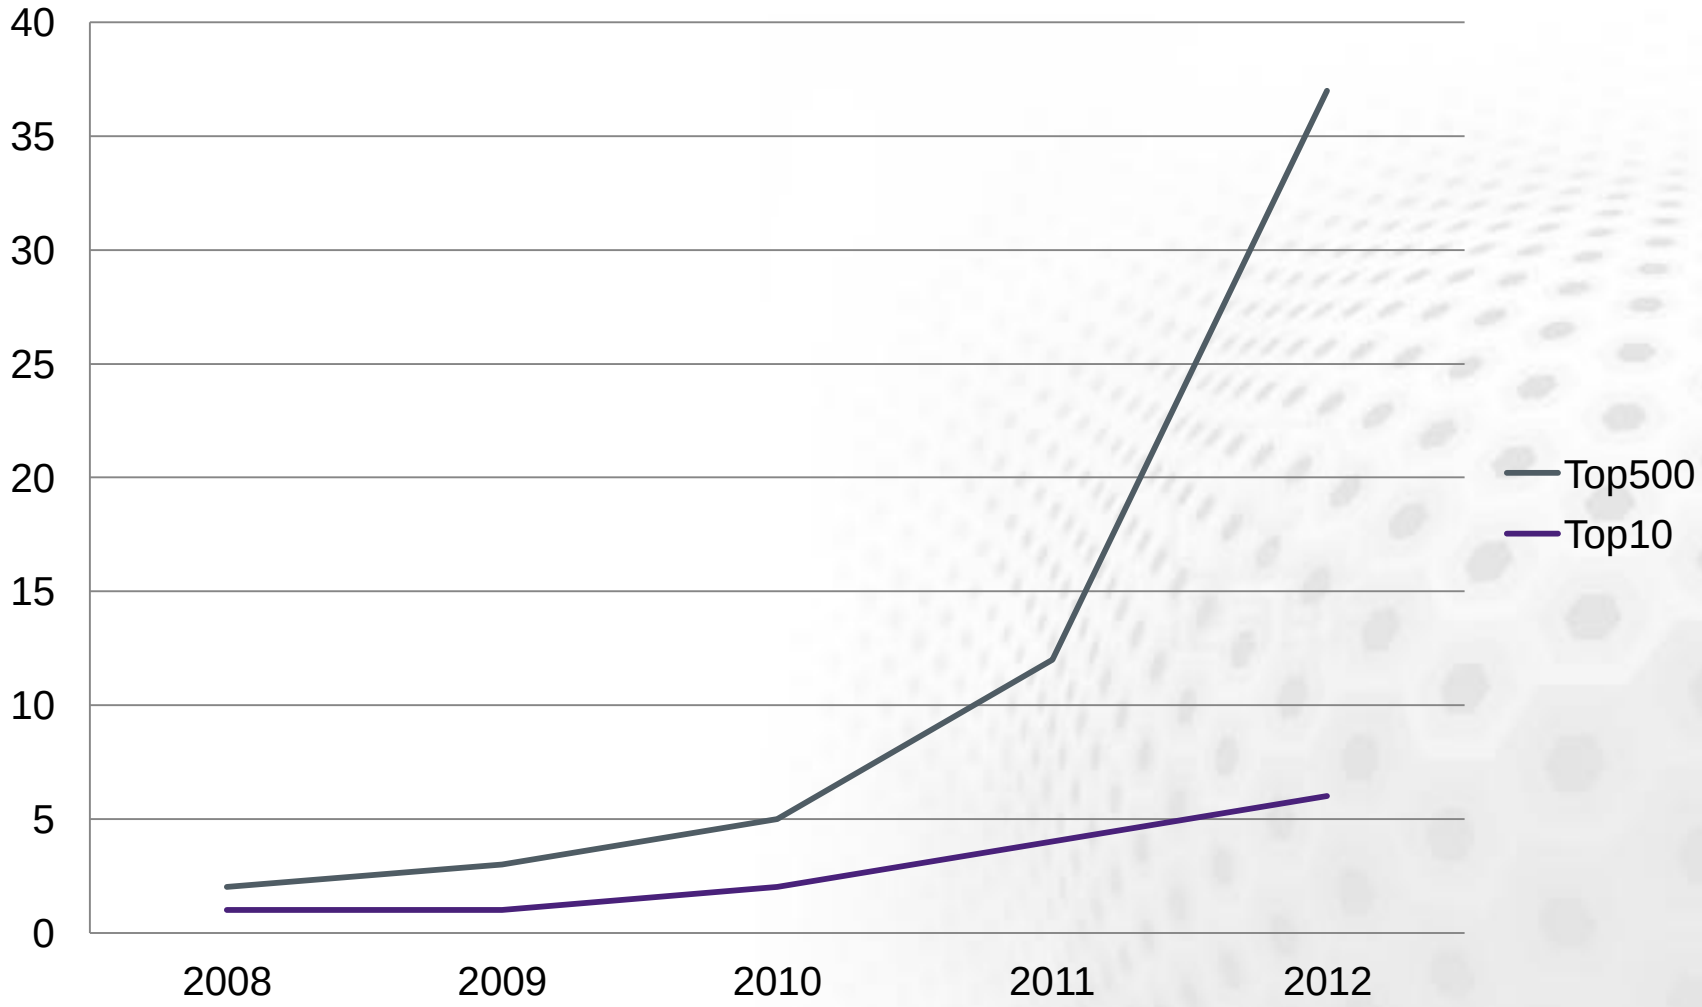

Evolution of Performance

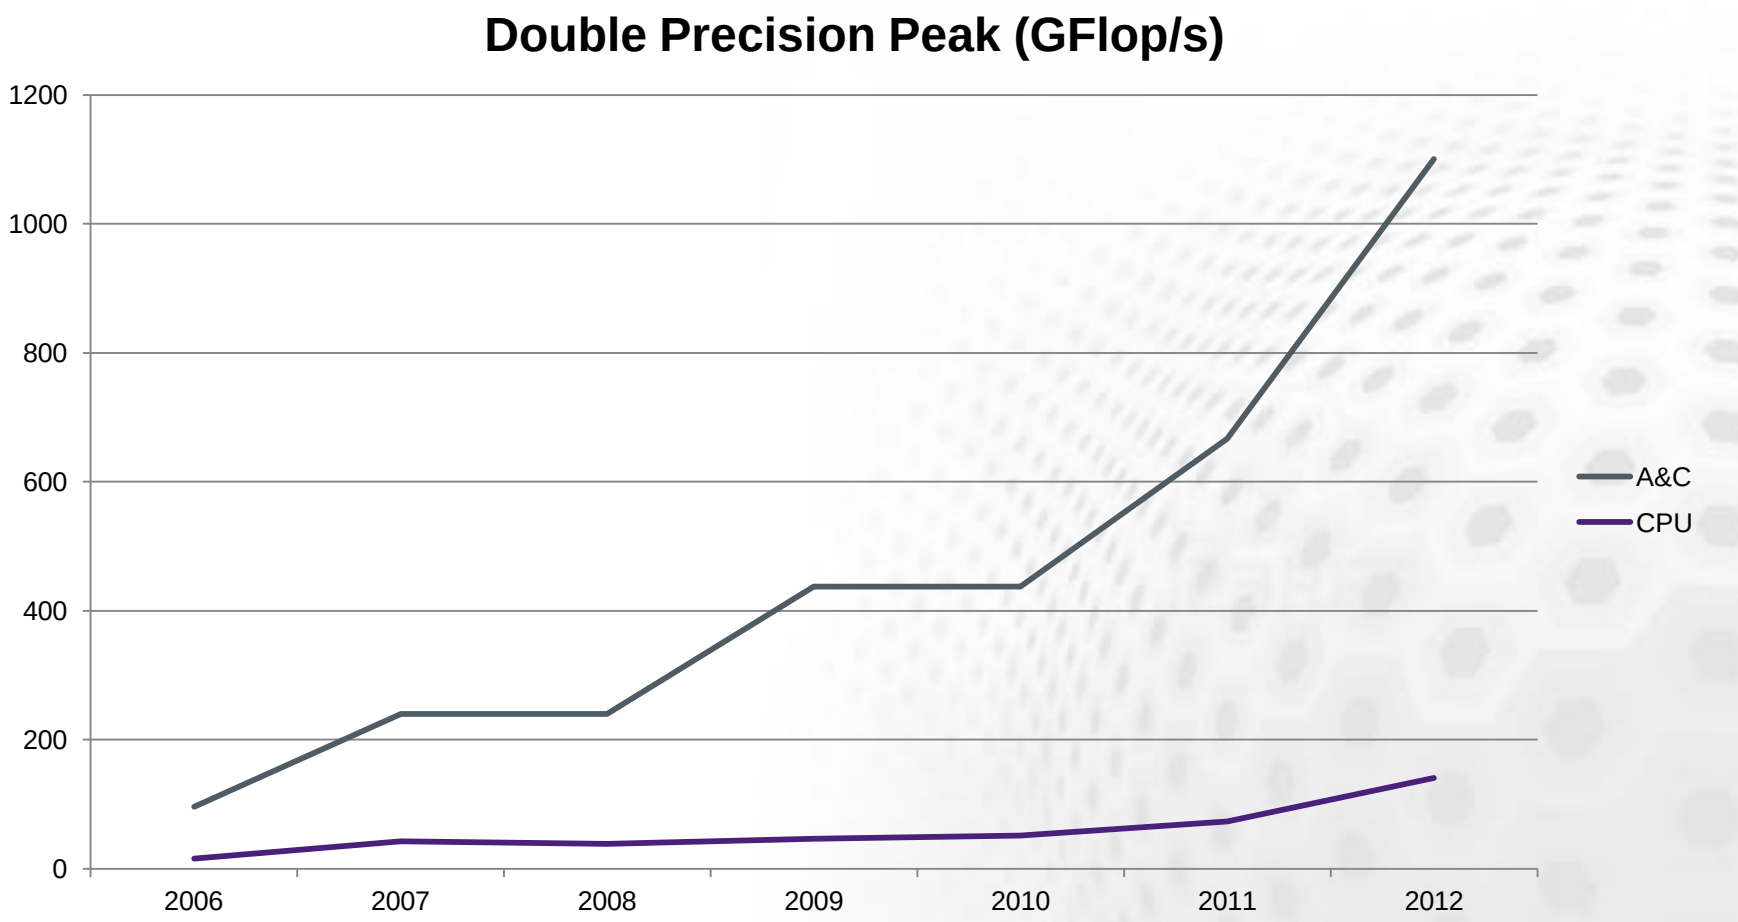

Evolution of Performance

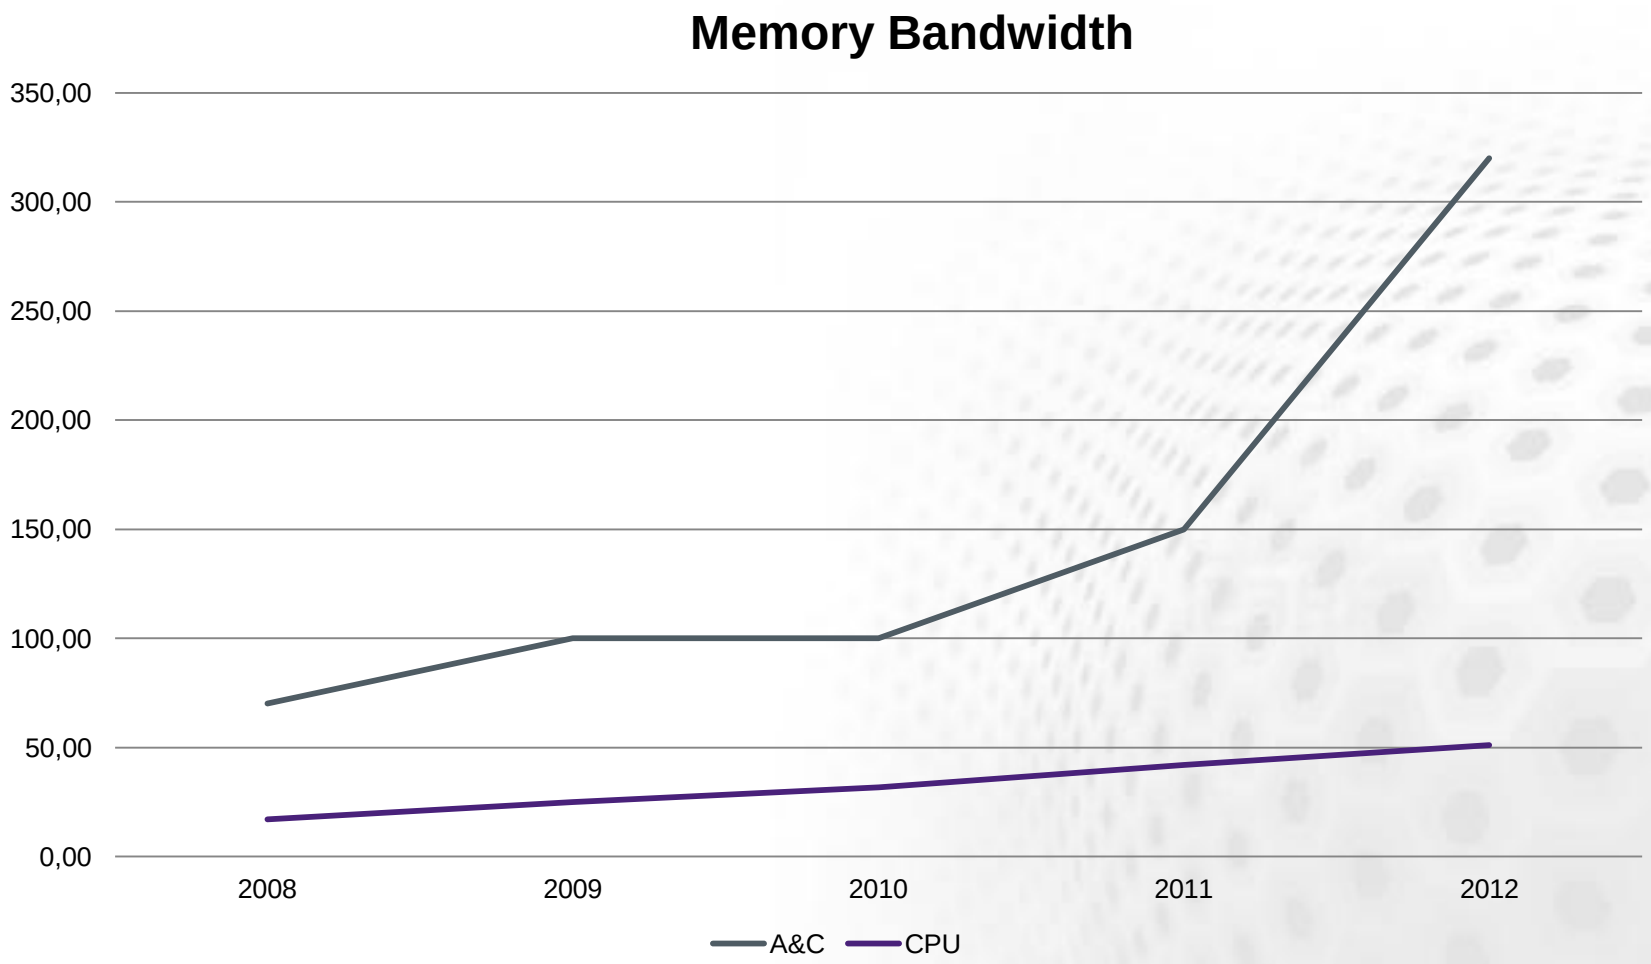

Recent Developments

- Focus shifting to productivity
  - Xeon Phi
  - Improvements in compilers, profilers, debuggers
  - Directive-based offloading (OpenACC, OpenMP 4.x)
  - Improved hardware features
    - Automatic caches, more registers etc.
  - Increasing library support and application ecosystem
- Major deployments in US and EU
  - Oak Ridge Titan, CEA Curie, TACC Ranger

Things evolve at a very rapid pace!

Conventional wisdom may be misleading!