General-Purpose vs. GPU: Comparison of Many-Cores on Irregular Workloads

University of Maryland, College Park

- **..we will not bring these** [100 core] **products to market until we have good solutions to the programming problem**
  J. Rattner, Intel CTO 3/2006
- Proposed solution **XMT (eXplicit Multi Threading)**
  - General-Purpose Many-Core platform
  - Issue: ease of parallel programming
  - Abstraction: (any) next instruction(s) execute immediately
  - Means: PRAM theory, programmer's workflow, HW+SW
  - Unmatched on: abstraction, teachability, and support by algorithms/theory, foundation of CS
- How much performance does one need to sacrifice for ease of programming?

Surprise. Performance **bonus** when using similar chip area:

- 6.05x average speedup over CUDA GPU on irregular applications
- 2.07x slowdown on regular applications


G.C. Caragea


F. Keceli

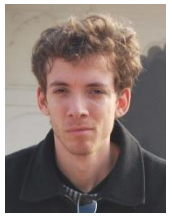
A. Tzannes


U.Vishkin