

On Programming Models for Multi-Core Computing


Frank Feinbube

HPI Research School
Operating Systems and Middleware
Prof. Dr. Andreas Polze

The free lunch is over...

2

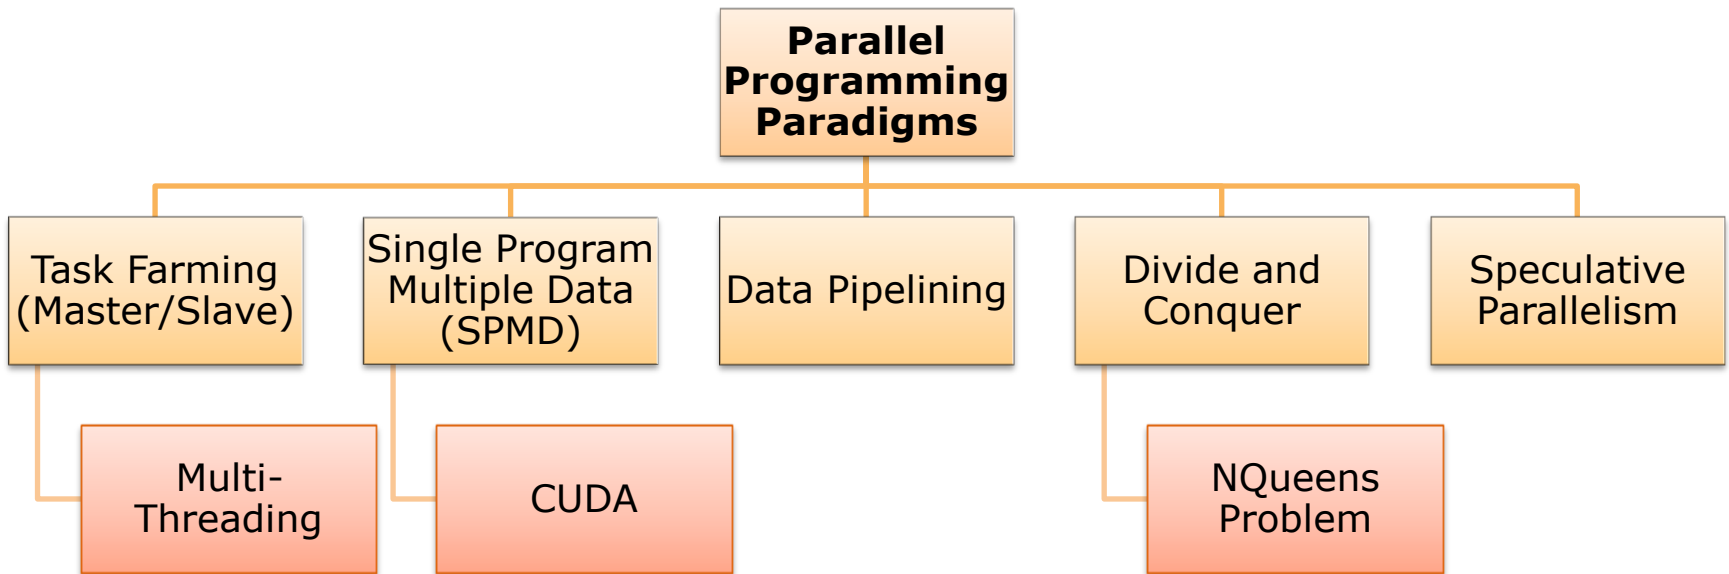
- ... the end of Moore's law:
 - The number of transistors continues to climb, at least for now
 - Clock speed, however, is a different story



based on http://jonasboner.com/talks/state_youre_doing_it_wrong/html/all.html

Overview: Parallel Programming Paradigms

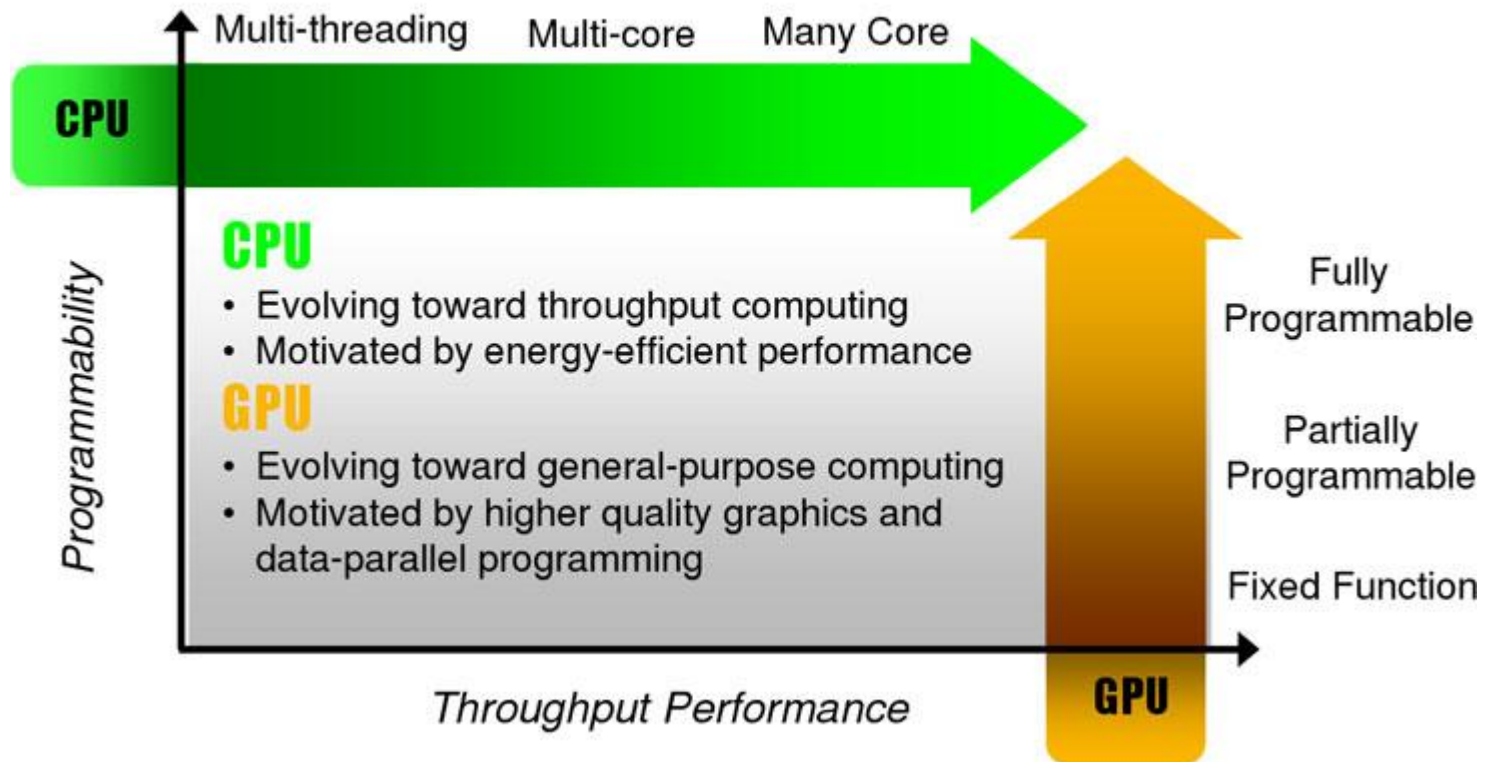
3



based on Rajkumar Buyya, High Performance Cluster Computing: Programming and Applications

CPU and GPU convergence

4




http://images.bit-tech.net/content_images/2009/08/does-nvidia-have-a-future/convergence.jpg

Die Damen können sich nicht mehr gegenseitig schlagen.

22,317,699,616,364,044

Evaluation of NQueens optimizations

6



- Time needed for calculation of the 16 queens problem on GeForce GTX 275
- Time needed for calculation of the 15 queens problem on GeForce 8600M GS

Conclusion

7

- There is still a big gap between General purpose graphic cards and CPUs
 - Predicting the performance of a CUDA program is hard
 - ◇ Need for a model to derive performance from register count, shared memory size and thread count for a known architecture
 - Long running functions -> driver will crash
 - Double precision ... work in progress
 - No recursion

- Memory access patterns and data structures sizes are the challenges of the multi-core future

- Preparing a paper for „Facing the Multicore-Challenge“, 03/2010