Experience with Beignet OpenCL on Intel Core M

Felice Pantaleo, Vincenzo Innocente

CERN – EP Department

Outline

- Low Power Architectures
- Intel Skylake Core M
  - Test and results
- Beignet OpenCL
  - Test and results
- Conclusion

Low Power Architectures

- Power consumption is becoming a hot-spot in the total bill
  - Especially true in Europe
- At CERN, this will be even more important with the HL-LHC upgrade
  - LHC experiments will have to cope with 2-3x the amount of data
- Increasing interest in alternative low power architectures based on ARM
- Increasing interest in complementary highly efficient accelerators like GPUs

Intel Skylake Core M

Credits: Intel

- Turbo Boost, Hyperthreading
- AVX2.0
- 4MB of L3 Cache
- Intel HD 515 GPU
- 4.5 W TDP, passively cooled…

Who needs Lugano Lake?

Hardware test configuration

Product Name	Intel® Core™ m3-6Y30 Processor (4M Cache, up to 2.20 GHz)	Intel® Core™ i7-6700K Processor (8M Cache, up to 4.20 GHz)
Code Name	Skylake	Skylake
Essentials		
Performance		
# of Cores	2	4
# of Threads	4	8
Processor Base Frequency	900 MHz	4 GHz
Max Turbo Frequency	2.2 GHz	4.2 GHz
TDP	4.5 W	91 W
Configurable TDP-up	7 W	
Configurable TDP-down	3.8 W	
Memory Specifications		
Graphics Specifications		
Processor Graphics ‡	Intel® HD Graphics 515	Intel® HD Graphics 530
Graphics Base Frequency	300 MHz	350 MHz
Graphics Max Dynamic Frequency	850 MHz	1.15 GHz
Graphics Video Max Memory	1.7 GB	1.7 GB

Credits: Intel

FKDTree

- Parallel Heapified KDTree
- Branchless search in TBB, OpenCL, CUDA
- Mainly used for track seeding or clustering (nearest neighbor search)

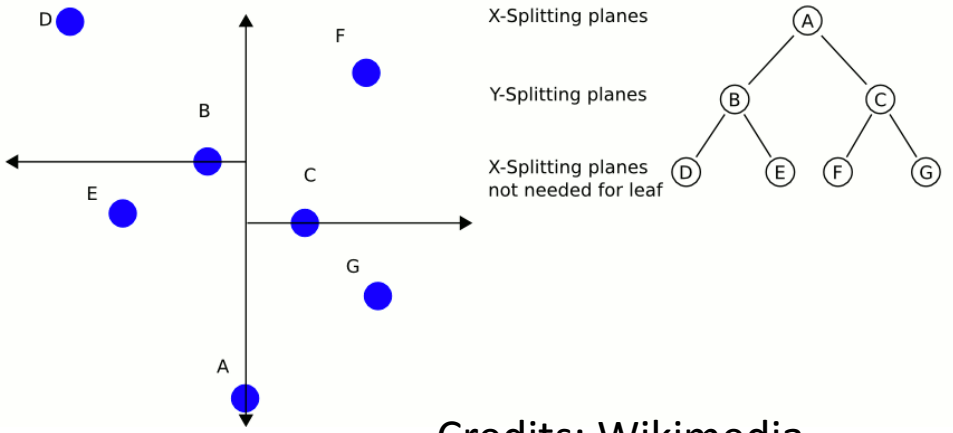

Credits: Wikimedia

- Test configuration:
  - 3D Cloud of 500k points
  - Searching for points inside a box around each point
  - Density of points ~32 avg points inside each search box

FKDTree results

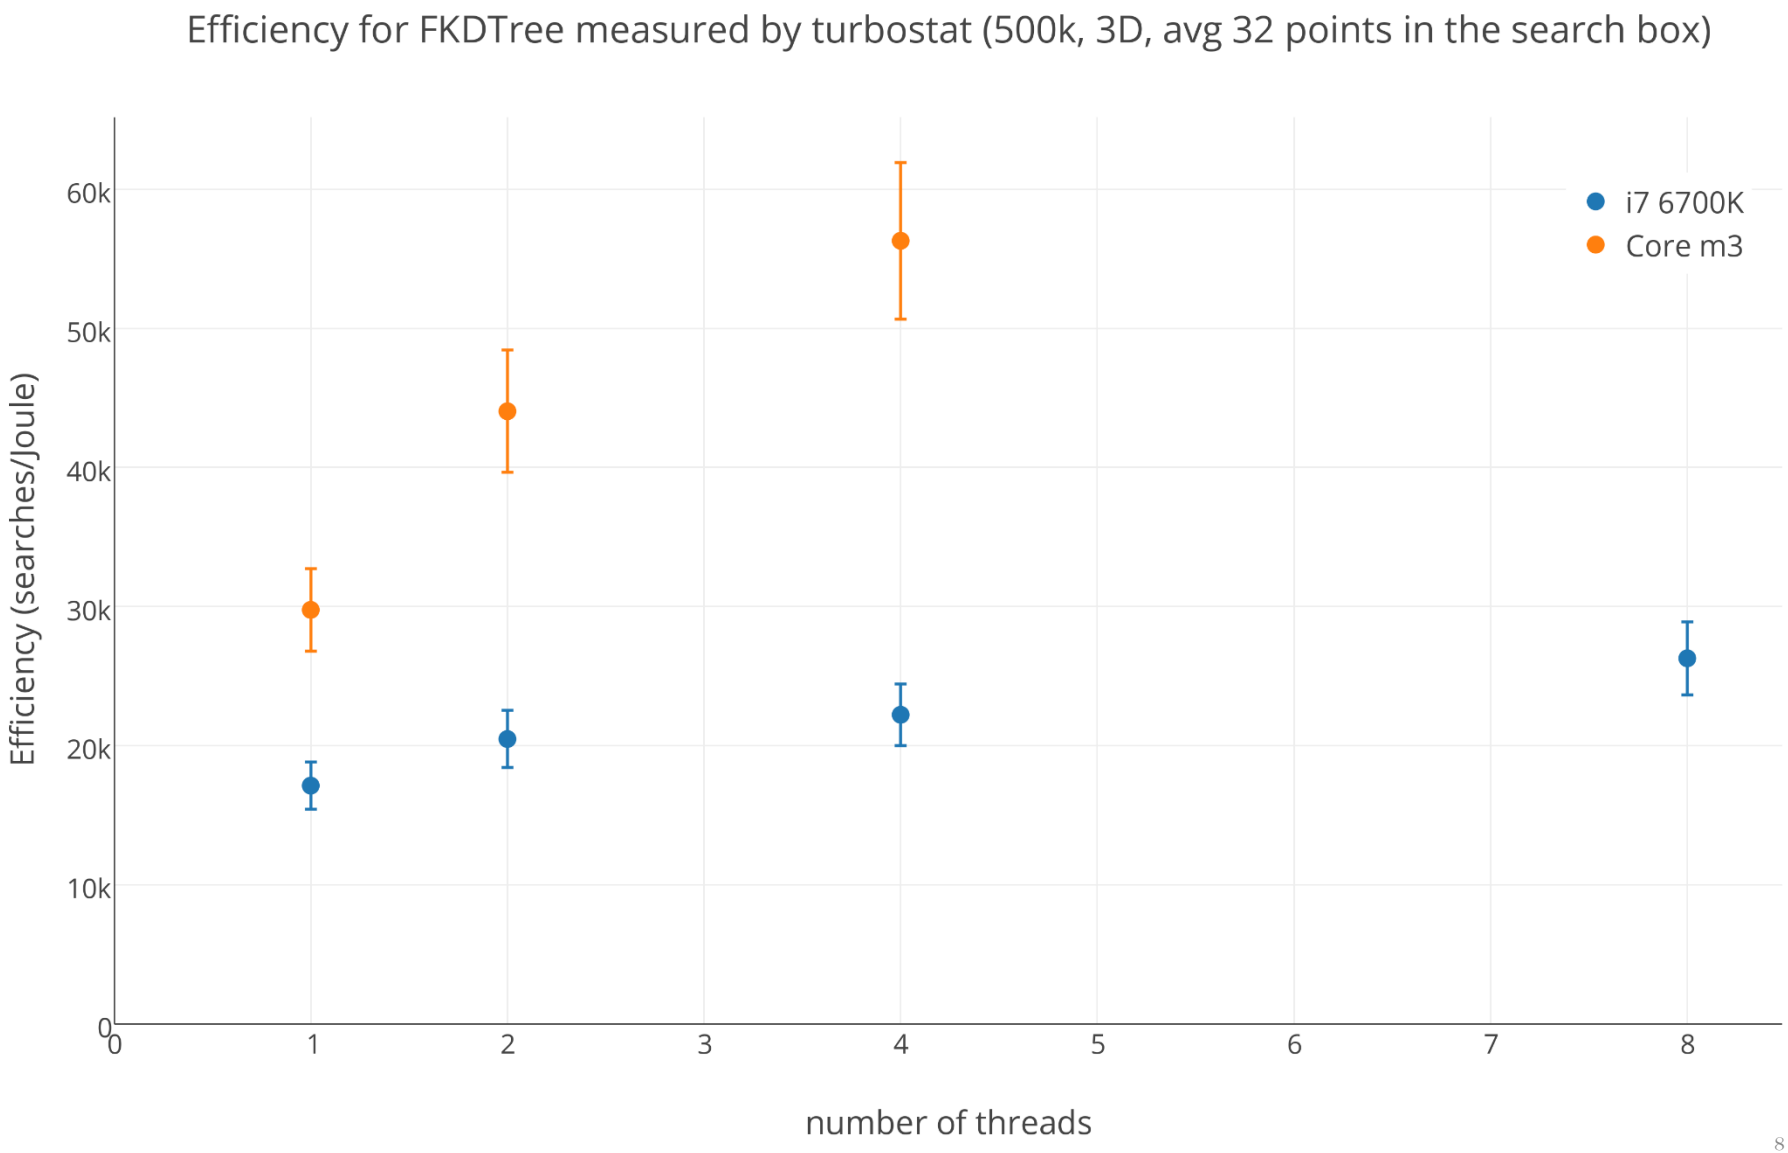

FKDTree results (ctd.)

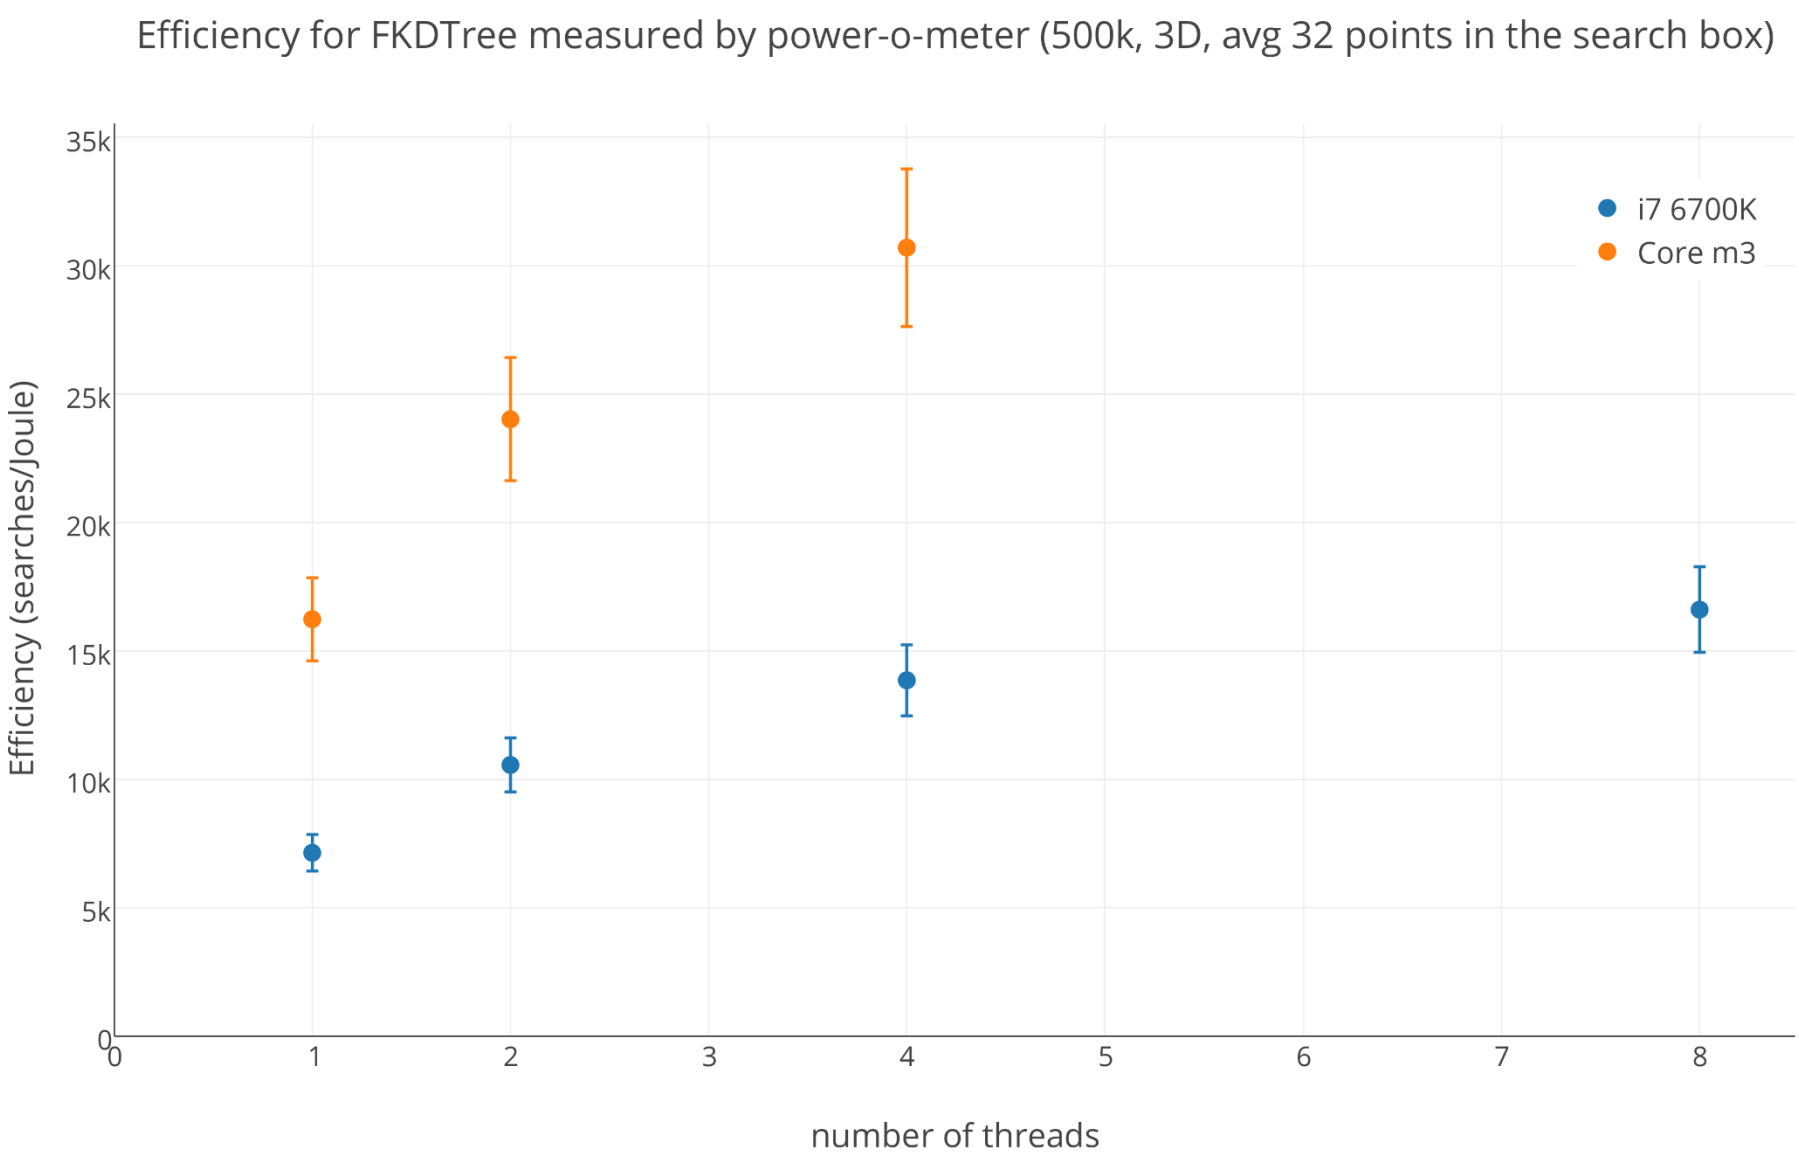

ParfullCMS

- Standalone CMS simulation using Geant4 (v10.1) with representative geometry
- Simplified Physics
- Compiled with GCC 4.9.x
  - static binaries and multithreading support.

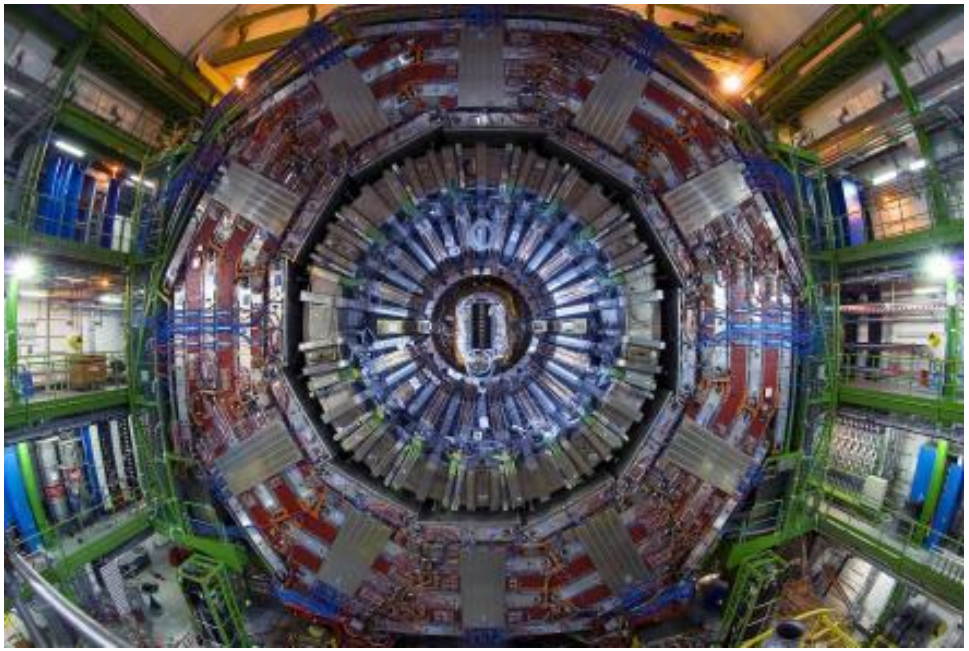

Credits: CDS CERN

ParfullCMS results

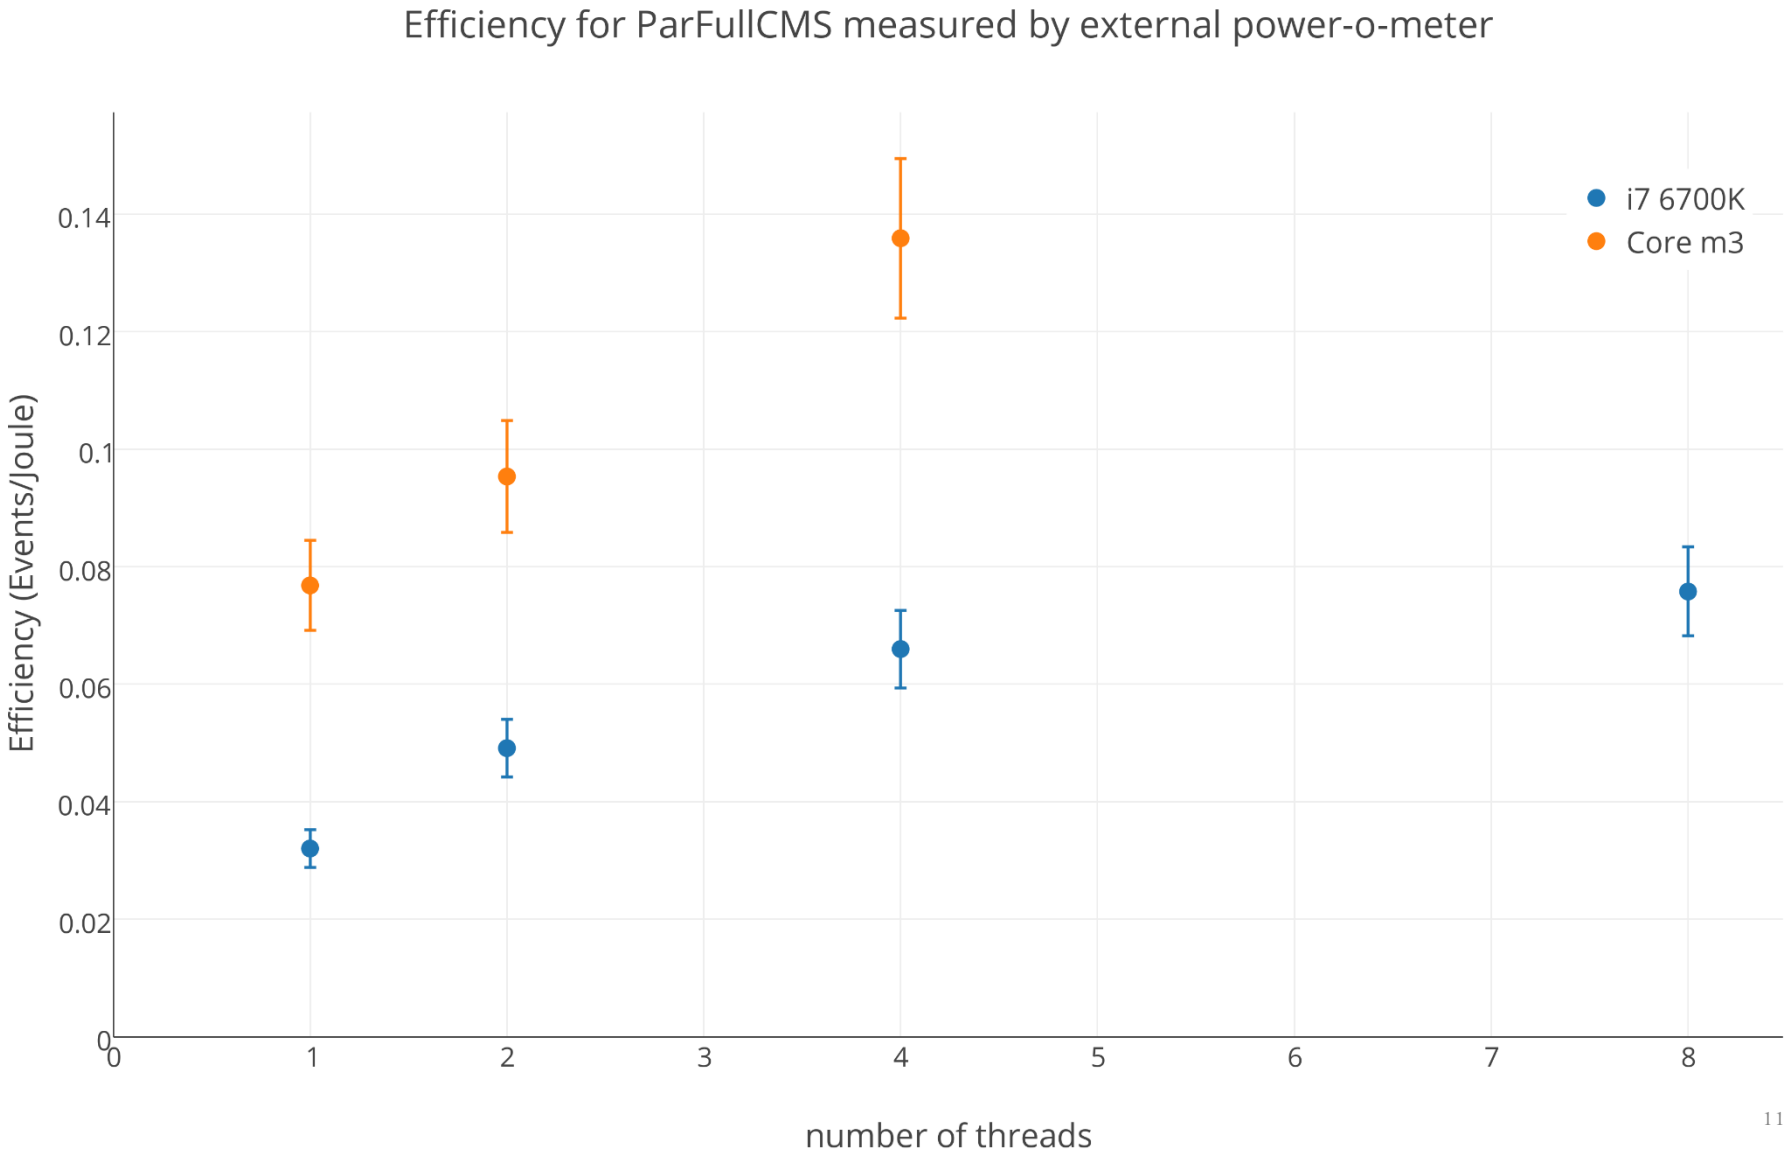

ParfullCMS results

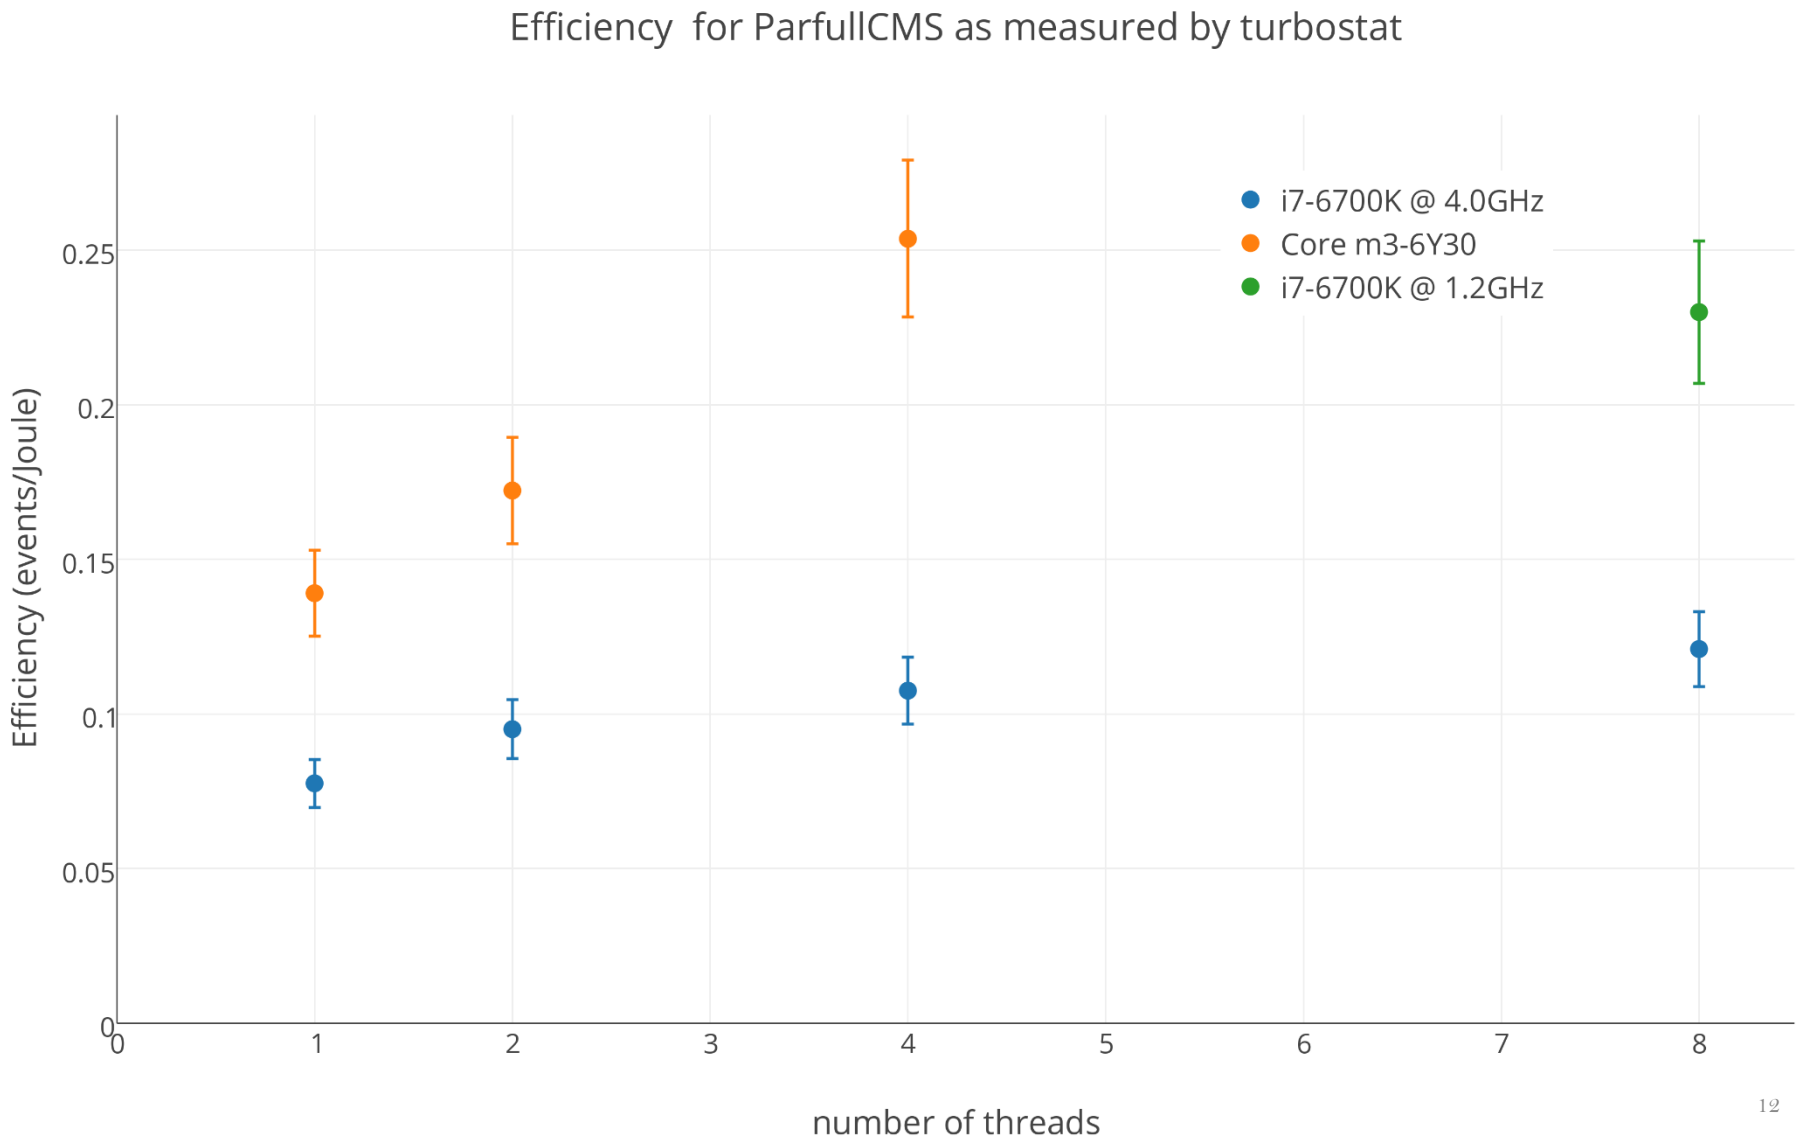

ParfullCMS results

- ParfullCMS scales with the frequency
- Running the i7 6700K @ 1.2 GHz with 8 threads
  - 2.35 events/second
  - 0.23 events/Joule
- Running the Core m3 with 4 threads
  - 1.2 events/second
  - 0.26 events/Joule

- It looks like they cut the i7 in two, downclocked and downvolted

Credits: Fantasia, Walt Disney

But wait...

- Are we actually squeezing all the processing power from the SoC?
- There is still an integrated GPU in the package…

Beignet OpenCL

- Open-source implementation
- 70000+ lines of C and C++
- Distributed under the LGPLv2.1
- Supported GPUs: Intel HD, Iris, Iris Pro
- Supported CPUs: Intel Core, Atom

First installation on Skylake Core m3 (end of 2015) not straightforward:

- Required manual kernel patch now included in 4.3.3+
- libdrm version 2.4.66+
- llvm 3.5

Memory management

- Applications can inform the driver of their memory usage scenarios
  - during allocation
  - memory transfer API
- Driver implementations create internal copies of memory buffers
  - beneficial for improving caching behavior
  - dramatic impact on performance
  - device-specific knowledge to avoid these copies

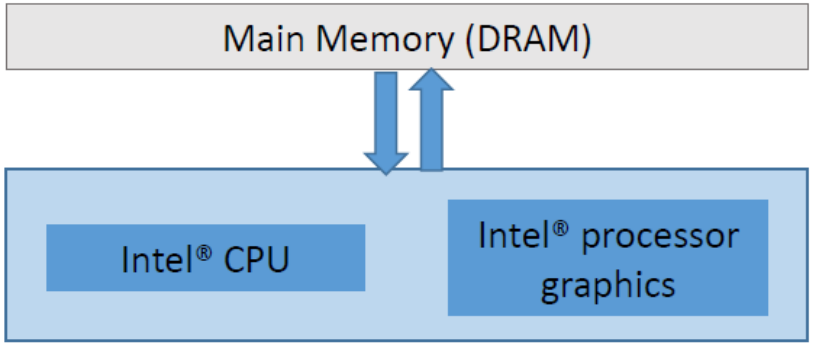

Memory management (ctd.)

- Integrated graphics
  - best performance when using zero copy
  - no need to create a host and a device version of data
  - No NUMA effects: memory shared between the CPU and GPU can be efficiently accessed by both devices.
- Adding OpenCL in existing codebase:
  - must create a buffer that is aligned to a 4096 byte boundary and have a size that is a multiple of 64 bytes

```
int *pbuf = (int *)_aligned_malloc(sizeof(int) * 256, 4096);
cl_mem myZeroCopyCLMemObj =
clCreateBuffer(ctx,…CL_MEM_USE_HOST_PTR…);
```

- OpenCL-managed host allocation

```
buf = clCreateBuffer(ctx, …CL_MEM_ALLOC_HOST_PTR, …)
```

FKDTree performance

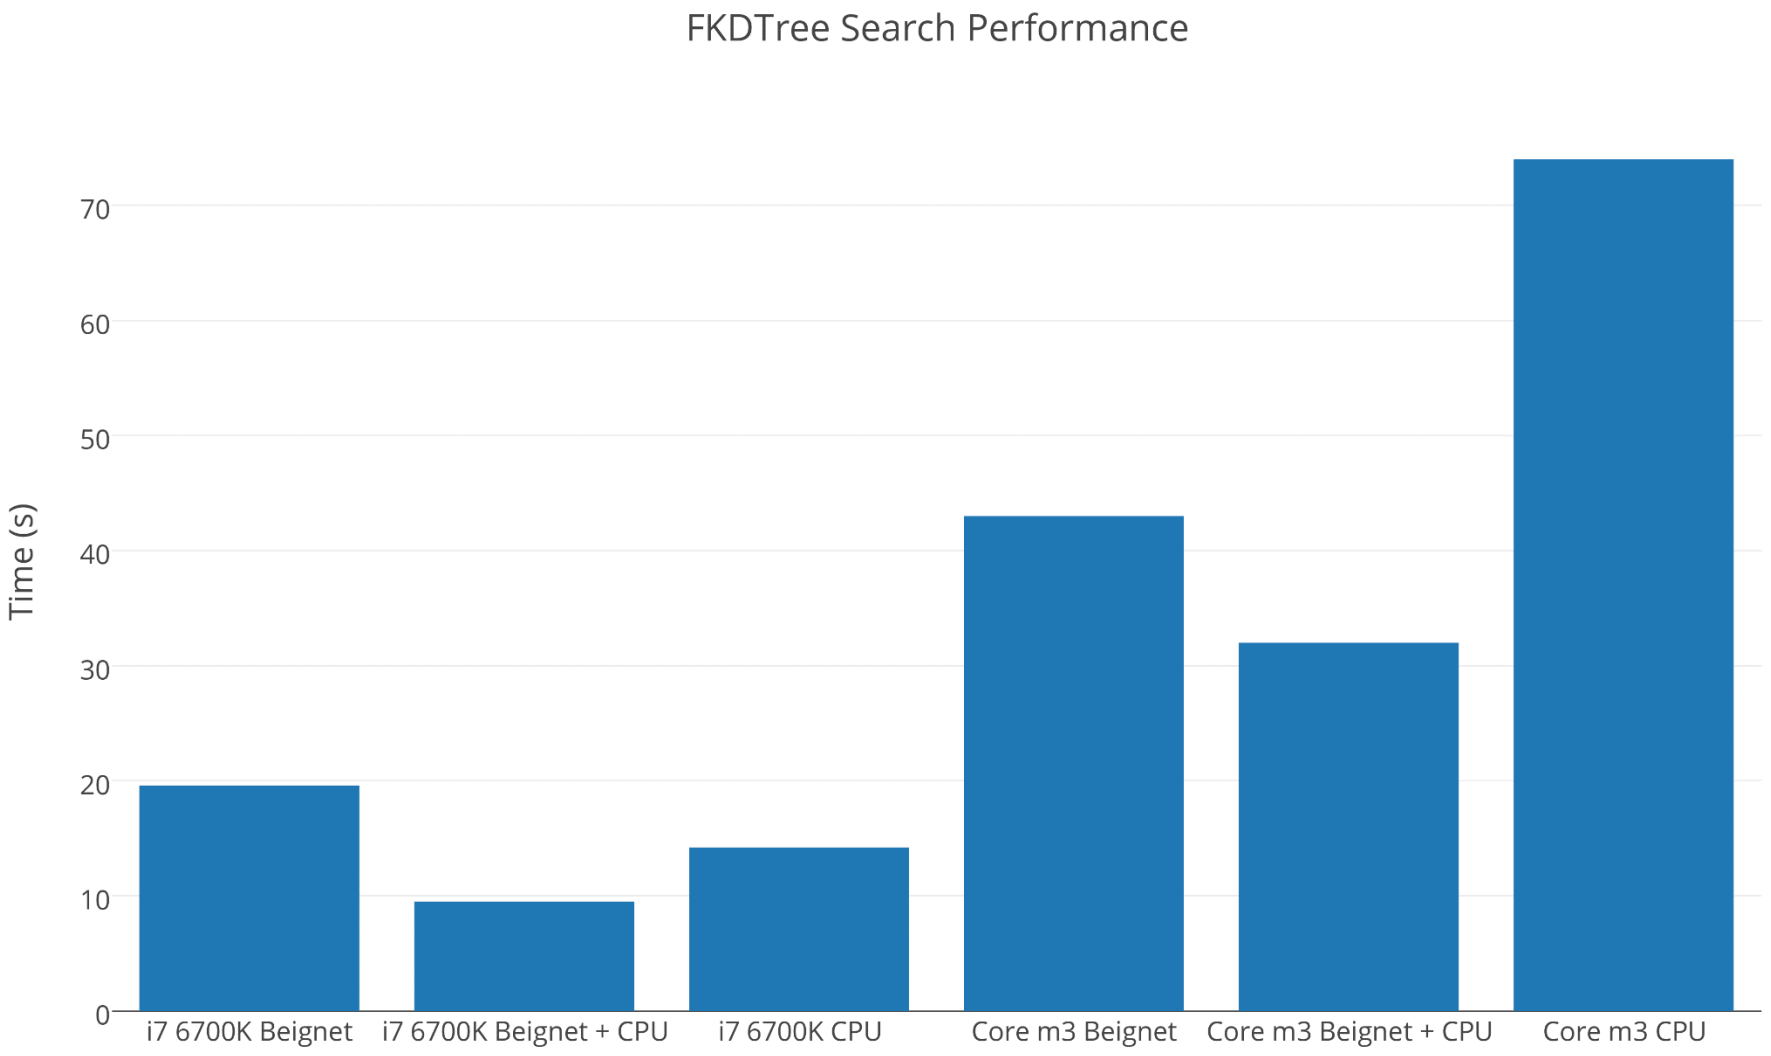

FKDTree efficiency

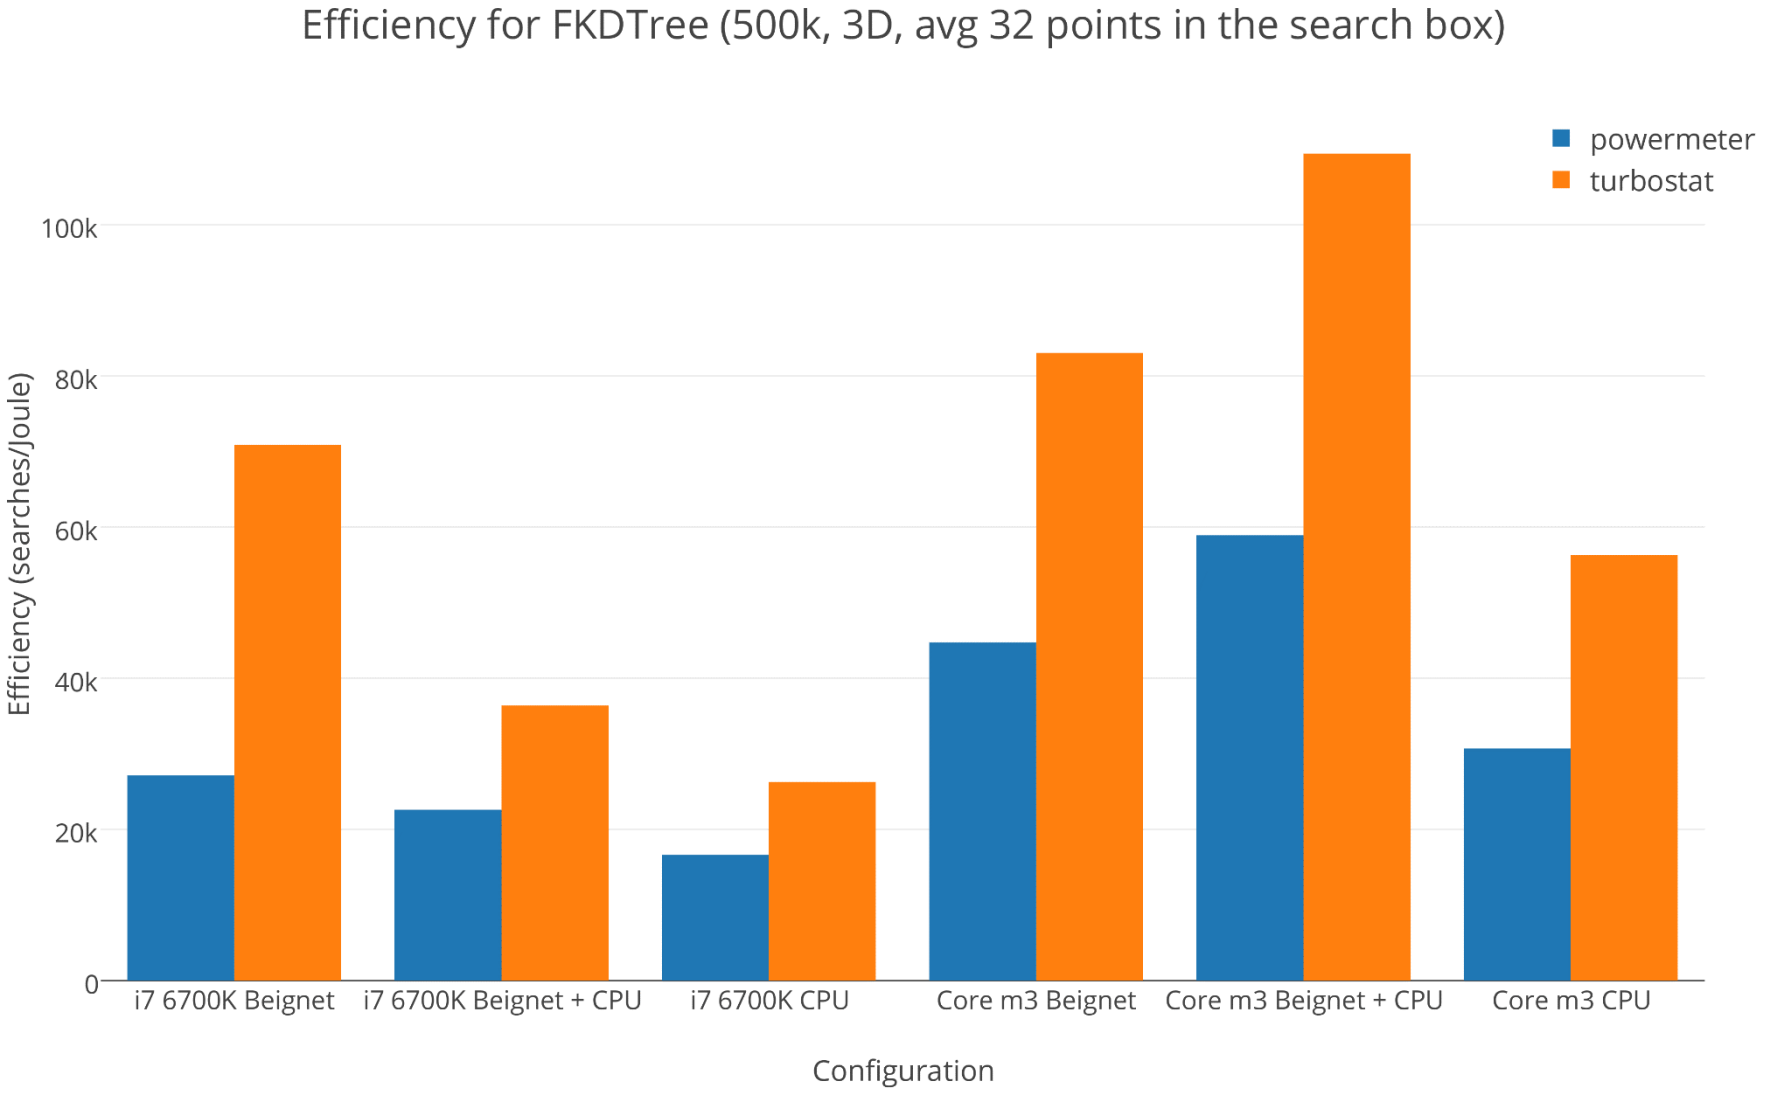

Conclusion

- Intel is working hard to fill the gap with ARM and NVIDIA
- Frequency is one of the dominating factors in high throughput computing
  - Higher frequency need improved cooling systems and decrease density
- Compare energy density efficiency? events/(Joule liter)
- Exploiting unused SoC resources now possible and “easy”
- Allows to achieve higher energy efficiency, throughput and latency
- Useful to offload parallel friendly C-kernels
- Still problems getting everything to run on CentOS 7