

QuantLib on Intel Xeon Phi

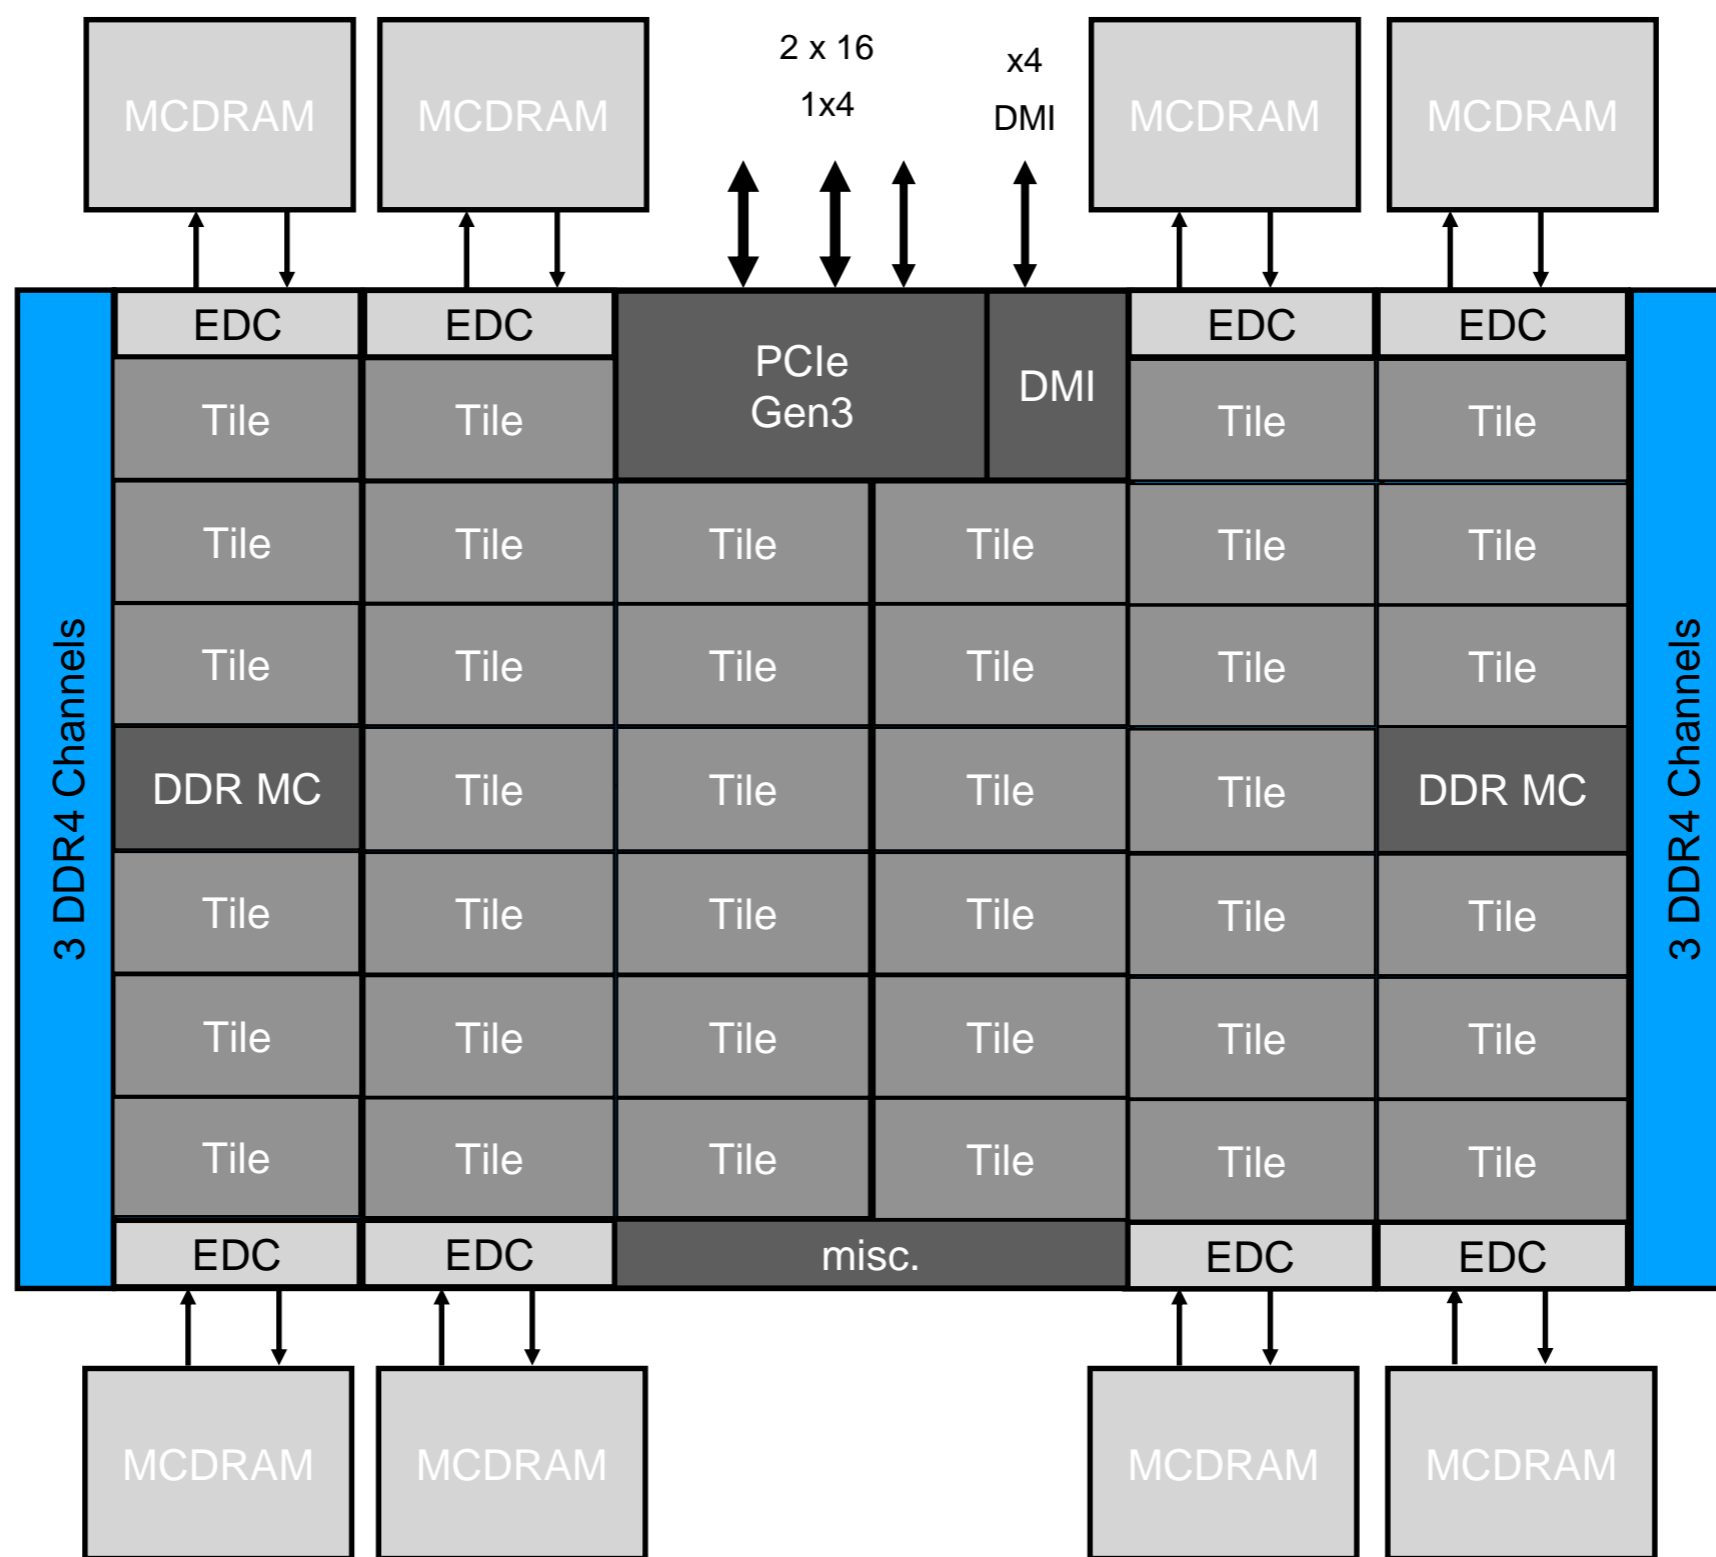
Thomas Rodgers
FCA Design
rodgert@twroddgers.com

My background

- Disclaimer - I am not a quant
- C++ developer with > 20 years of experience in Finance
- Member of the C++ Standards Committee
 - SG1 - Concurrency and Parallelism
 - SG14 - Games, HFT, Low-Latency systems
- My interest in this project is how to make C++ a better tool for heterogenous multi-core development

Project Background

- Work in support of research by Prof. Matthew Dixon, Stuart School of Business, Illinois Institute of Technology
- Help Quants target highly parallel, or vector-parallel architectures, without needing a deep background in vector/parallel software engineering
- Xeon Phi is one such target architecture
- Need a source of publicly releasable financial model codes
 - QuantLib is one possibility



Knights Landing

Die Layout

Knights Landing Die Features

- Chip organized around the concept of a Tile, each consists of -
 - 2 cores, 32kb L1 Instruction and 32kb L1 Data cache per core
 - 2 vector processing units per core, total 4 per Tile
 - 1M of L2 cache shared between the two cores
- All Tiles are connected to a 2D mesh with $> 700\text{GB/s}$ bandwidth
 - Organized into rows and columns of “half” rings that fold upon themselves at the endpoints
 - Enforces a “YX routing” rule, transactions travel vertically to target row, then travel horizontally to destination.

Knights Landing Die Features (cont)

- 2 Memory controllers with 3 DDR4 channels each
 - Maximum of 384GB of DDR4
 - Aggregate DDR4 bandwidth is ~90 GB/s
- 8 MCDRAM (High Bandwidth Memory) devices, each 2GB capacity
 - Each MCDRAM is connected to it's own memory controller (EDC)
 - Aggregate MCDRAM bandwidth is > 450 GB/s
 - Can be used in cache mode, flat mode (shares address space with DDR), or a mix of the two

Knights Landing

Core details

- Each Knights Landing die contains 36 tiles of 2 cores each for a total of 72 cores per die
- Each KNL core is based on a significantly modified Atom processor (Silvermont)
- Features -
 - ISA compatible with Haswell (except TSX support)
 - Support for up to 4 simultaneous threads (hyperthreading)
 - Out of order execution

Knights Landing

Core details

(cont)

- Each Knights Landing tile includes two Vector Processing Units per core for a total of 144 VPUs per KNL die
- Each VPU can perform 16 double precision, or 32 single precision floating point operations
- Each VPU supports an extended 512 bit instruction set, aka AVX-512
 - Includes “foundation” support for floating point and scalar vector operations, as well as dedicated support for exponentiation, prefetch, and conflict detection
 - AVX-512 ISA support is available via compiler intrinsics in ICC, GCC, and Clang.
 - Support for AVX-512 compiler intrinsics are also available in ICC, GCC, and Clang
 - see <https://software.intel.com/sites/landingpage/IntrinsicsGuide> for more information

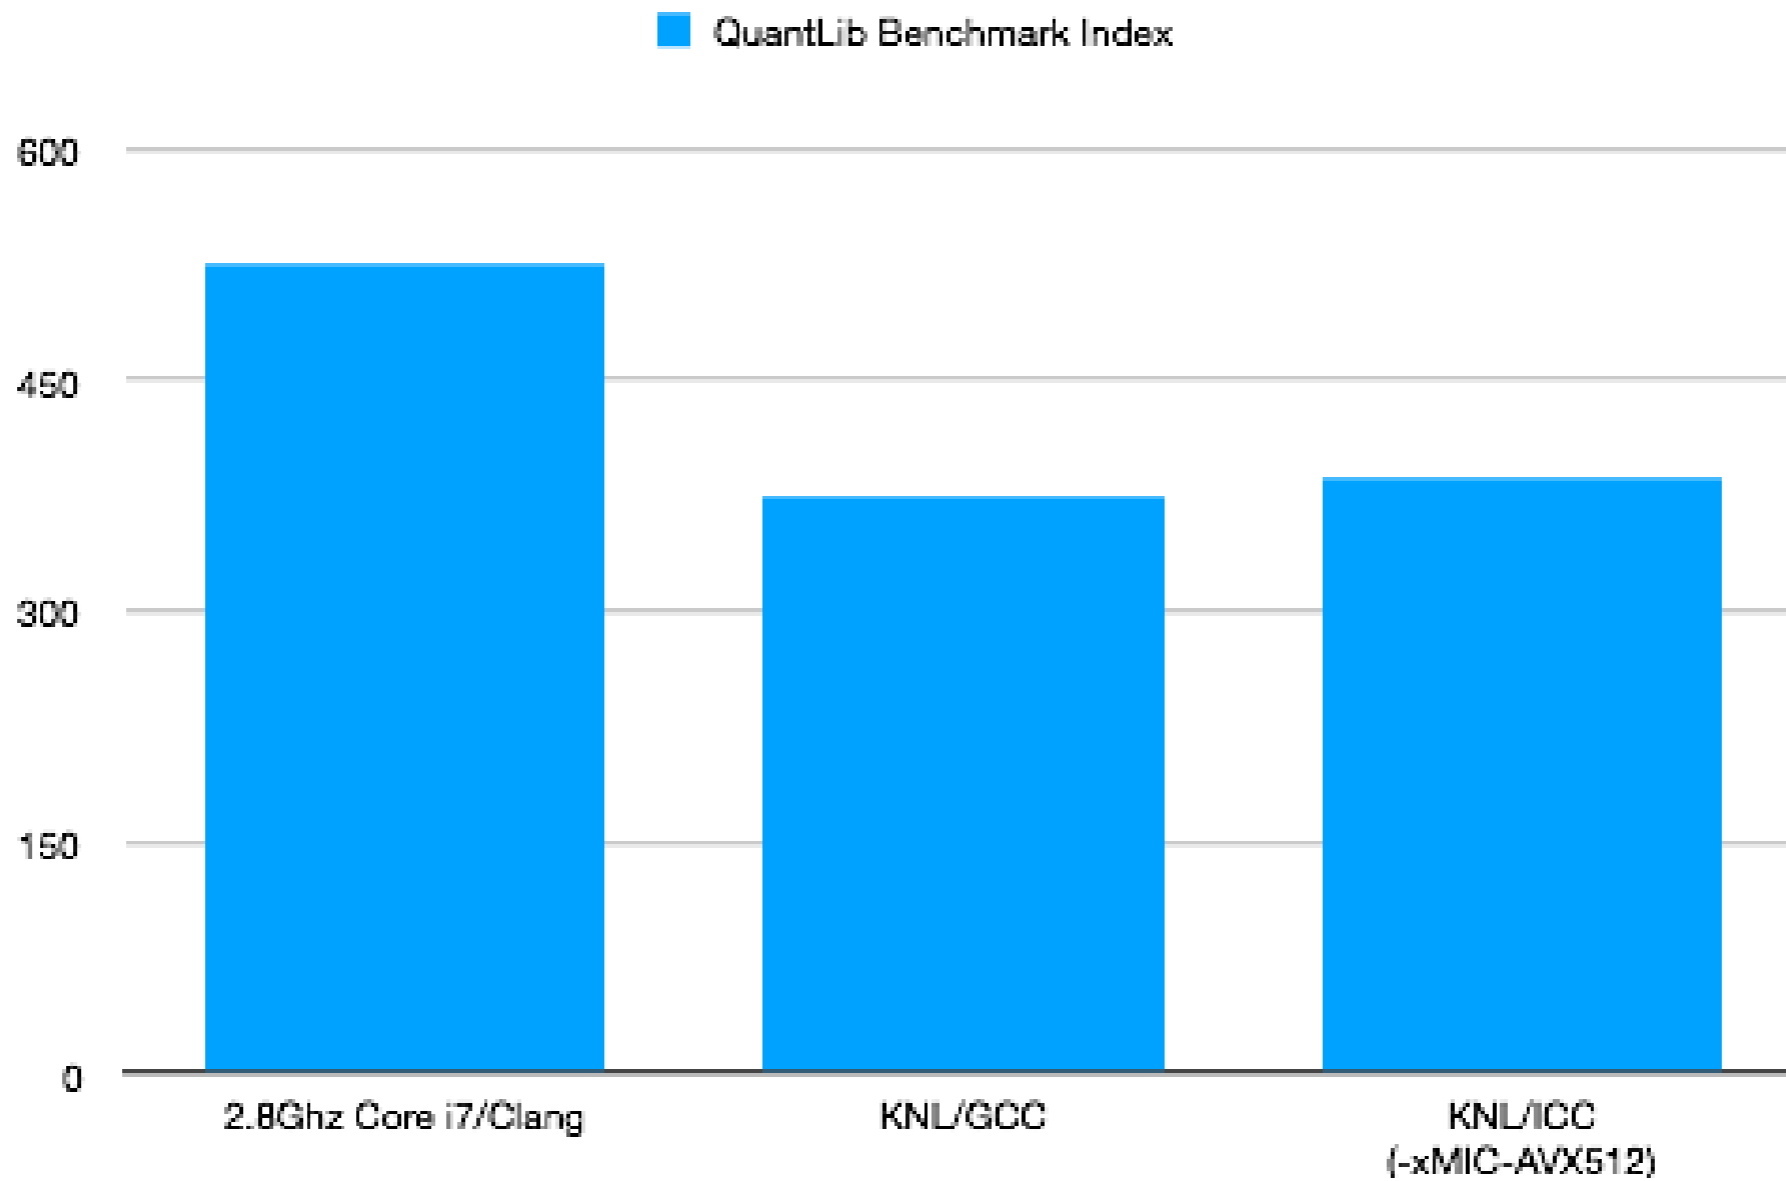
Knights Landing Configurations

- The Knights Landing processor, differs from the previous generation Knights Corner processor -
 - Standard Intel Architecture device
 - Fully capable of booting stock operating systems
 - Currently using CentOS 7.4

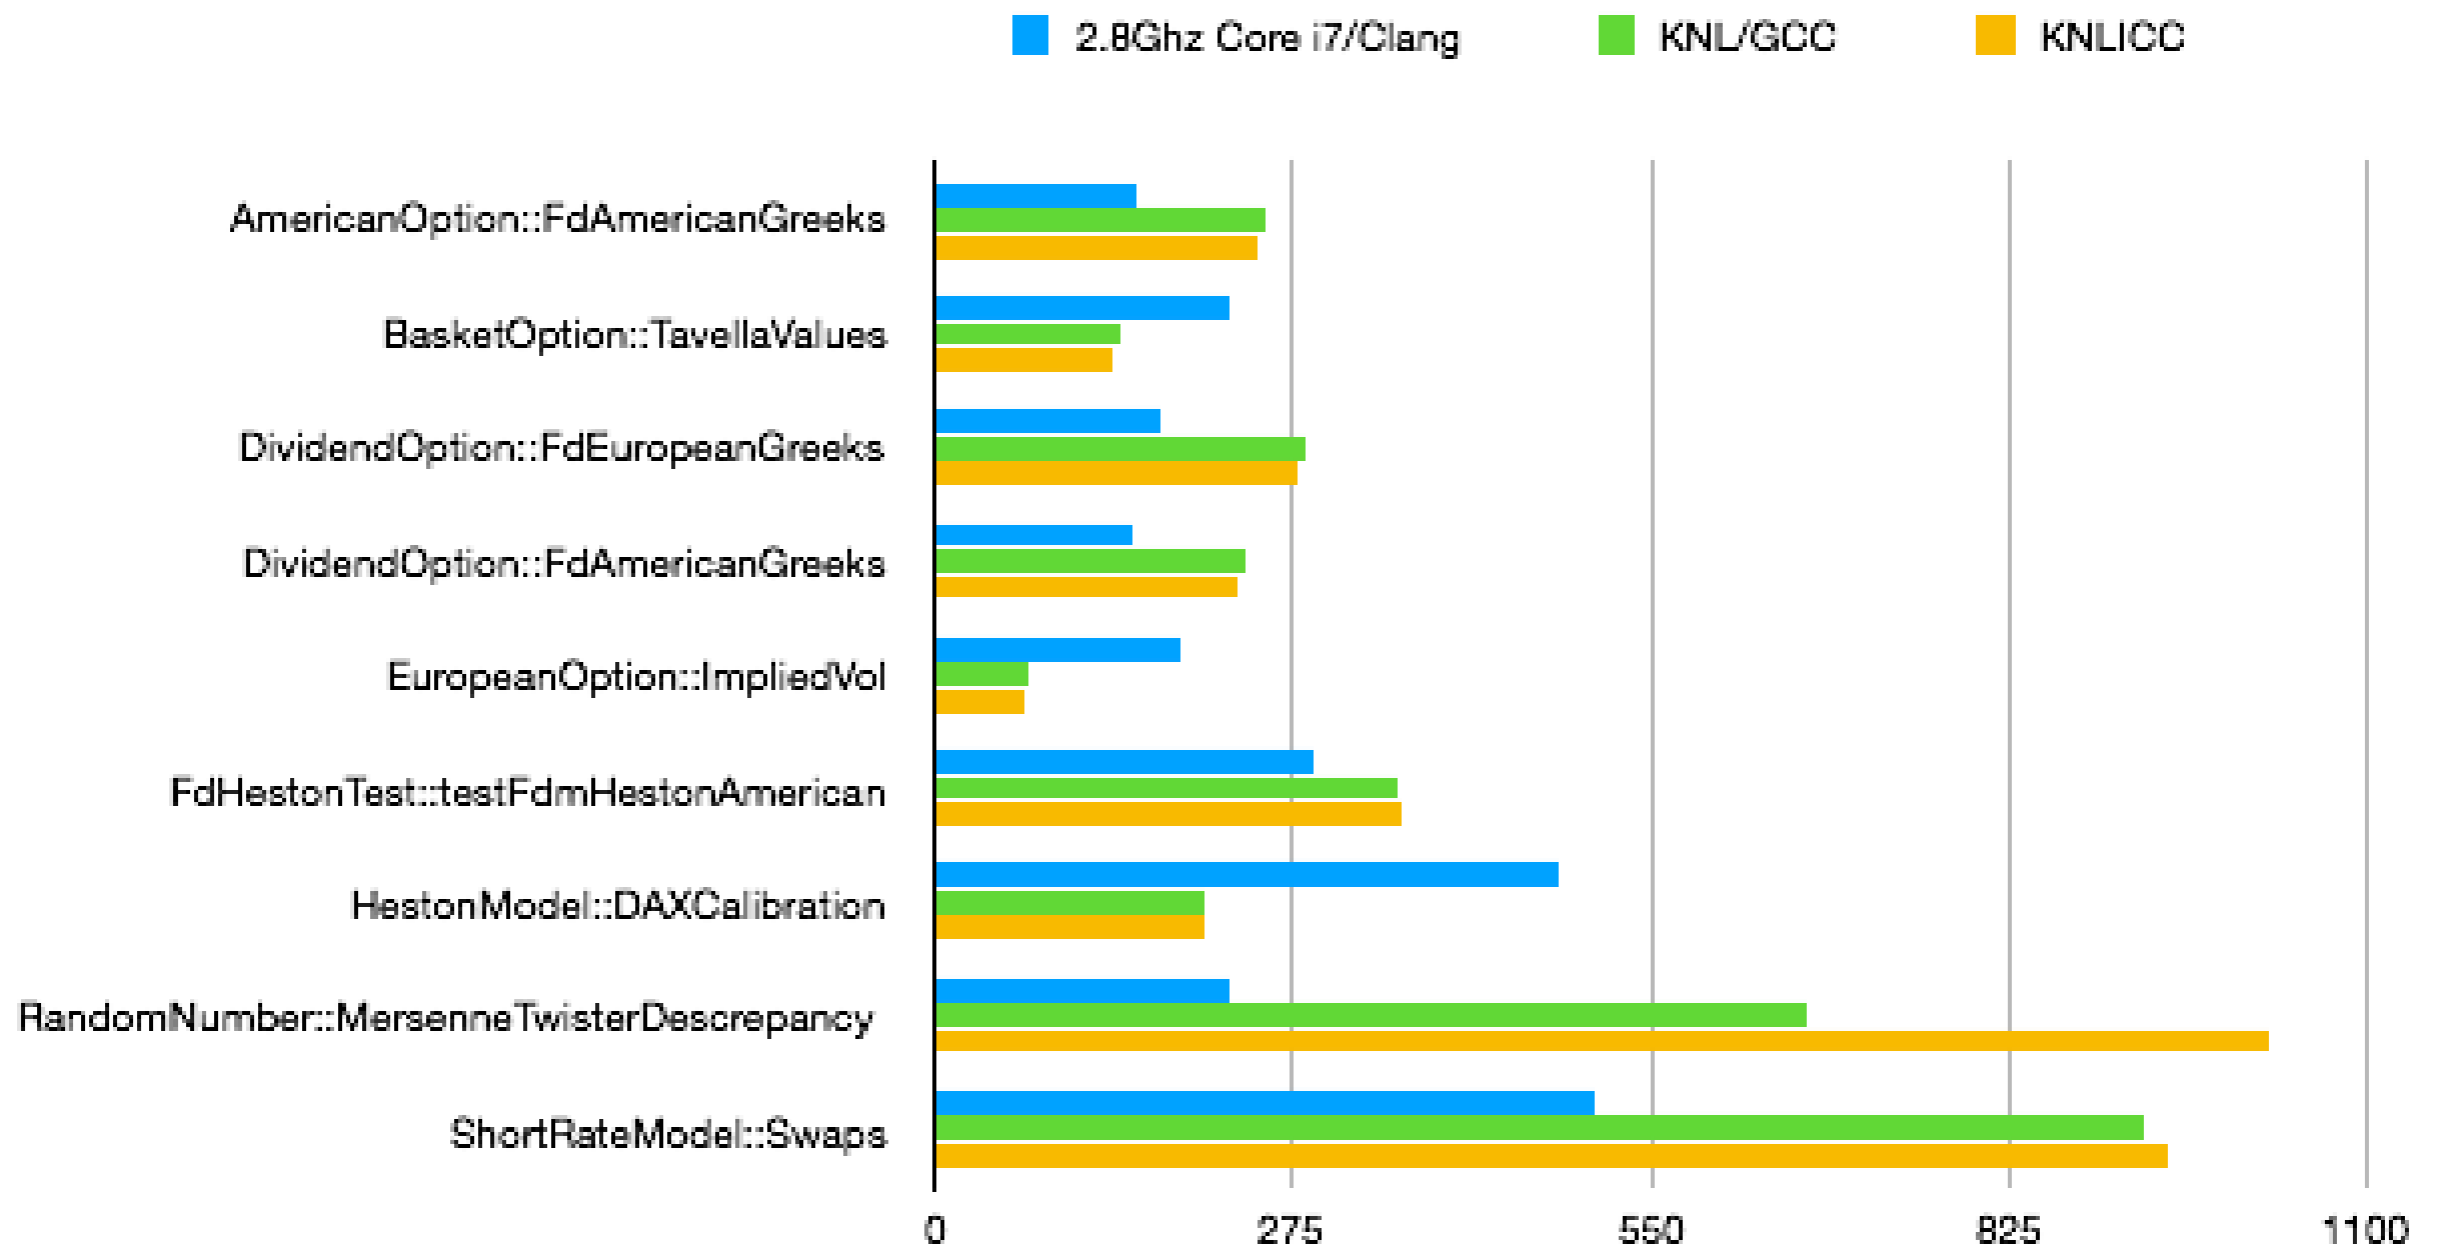
Knights Landing

Boot-time Configuration

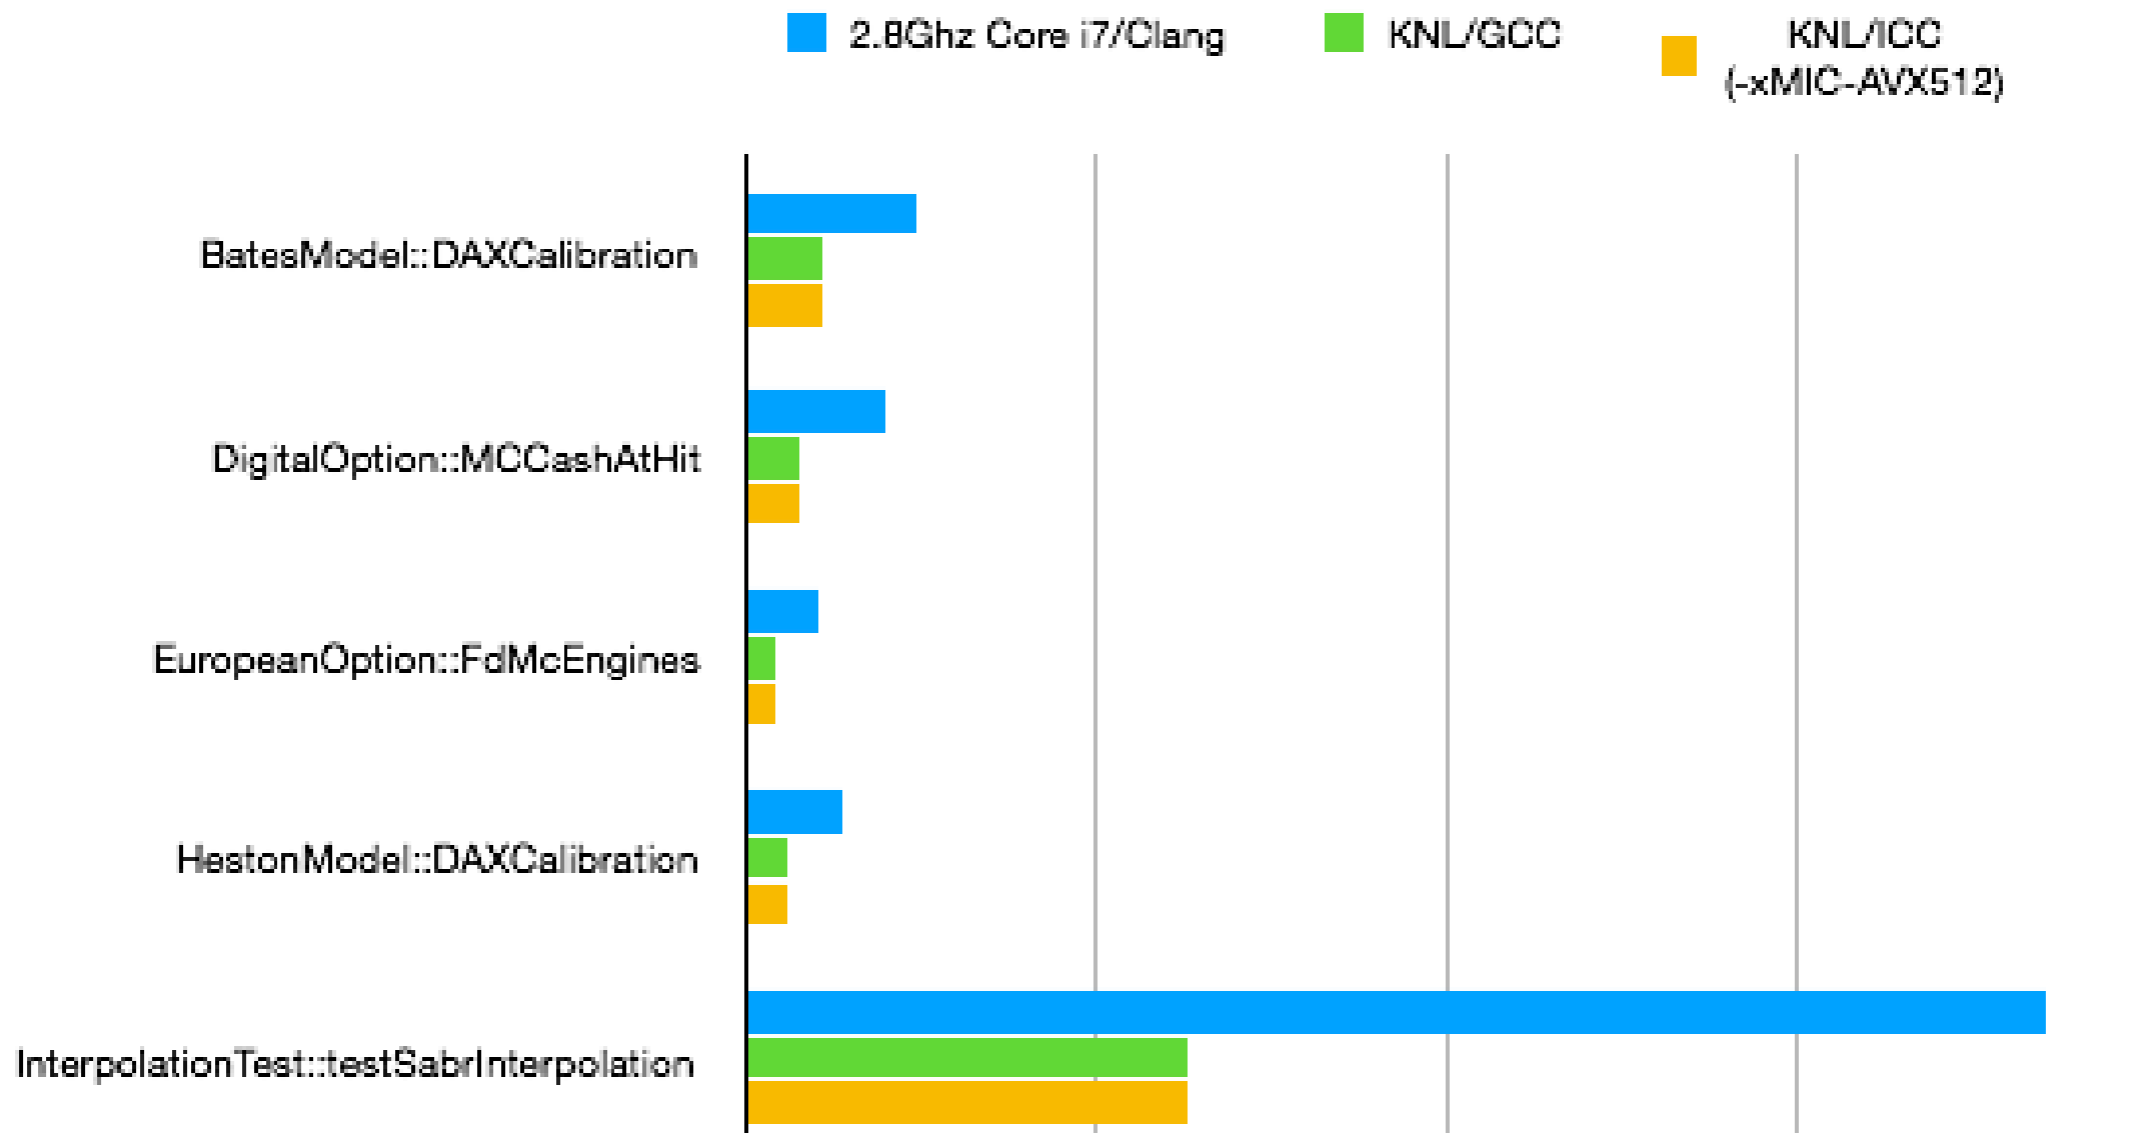
- Unlike standard Intel Xeon CPUs the Knights Landing CPU supports several boot-time configuration parameters
 - Cluster Mode - Configures the on chip mesh in one of
 - all-to-all mode (used with mixed DRAM configurations)
 - quadrant mode (default)
 - Sub-NUMA clustering (SNC)
 - Memory Mode -
 - cache, flat, hybrid



Benchmark Index (MIPS)



Benchmark Sample



Benchmark Sample

Observations

- The benchmark suite is single threaded
- Aggregate single core performance for Phi vs. a typical Intel Core/Xeon is lower, this result is generally expected, but...
 - Some individual models $>2x$ slower
- This appears at least in part due to how the QuantLib benchmark suite is structured
 - Significant amount of scalar set-up code inside deeply nested loops

It's Full of Cores...

```
1 [|||||] 98.4% 69 [|||||] 0.0% 137 [|||||] 100.0% 265 [|||||] 100.0%
2 [|||||] 2.9% 70 [|||||] 100.0% 138 [|||||] 100.0% 266 [|||||] 100.0%
3 [|||||] 21.9% 71 [|||||] 60.3% 139 [|||||] 100.0% 267 [|||||] 100.0%
4 [|||||] 82.3% 72 [|||||] 100.0% 140 [|||||] 100.0% 268 [|||||] 2.8%
5 [|||||] 7.1% 73 [|||||] 96.1% 141 [|||||] 96.9% 269 [|||||] 100.0%
6 [|||||] 93.7% 74 [|||||] 3.9% 142 [|||||] 100.0% 210 [|||||] 100.0%
7 [|||||] 46.4% 75 [|||||] 95.3% 143 [|||||] 61.5% 211 [|||||] 100.0%
8 [|||||] 58.3% 76 [|||||] 33.1% 144 [|||||] 100.0% 212 [|||||] 100.0%
9 [|||||] 9.4% 77 [|||||] 88.3% 145 [|||||] 97.6% 213 [|||||] 100.0%
10 [|||||] 96.9% 78 [|||||] 10.5% 146 [|||||] 100.0% 214 [|||||] 89.0%
11 [|||||] 4.3% 79 [|||||] 100.0% 147 [|||||] 100.0% 215 [|||||] 92.6%
12 [|||||] 97.6% 80 [|||||] 100.0% 148 [|||||] 100.0% 216 [|||||] 0.0%
13 [|||||] 27.2% 81 [|||||] 100.0% 149 [|||||] 100.0% 217 [|||||] 71.8%
14 [|||||] 0.4% 82 [|||||] 100.0% 150 [|||||] 100.0% 218 [|||||] 100.0%
15 [|||||] 39.9% 83 [|||||] 60.4% 151 [|||||] 100.0% 219 [|||||] 100.0%
16 [|||||] 58.1% 84 [|||||] 100.0% 152 [|||||] 100.0% 220 [|||||] 38.2%
17 [|||||] 98.8% 85 [|||||] 94.1% 153 [|||||] 100.0% 221 [|||||] 3.5%
18 [|||||] 8.2% 86 [|||||] 100.0% 154 [|||||] 89.8% 222 [|||||] 100.0%
19 [|||||] 33.1% 87 [|||||] 68.8% 155 [|||||] 95.3% 223 [|||||] 100.0%
20 [|||||] 84.6% 88 [|||||] 90.2% 156 [|||||] 15.9% 224 [|||||] 100.0%
21 [|||||] 1.6% 89 [|||||] 100.0% 157 [|||||] 100.0% 225 [|||||] 100.0%
22 [|||||] 98.4% 90 [|||||] 0.0% 158 [|||||] 100.0% 226 [|||||] 100.0%
23 [|||||] 16.5% 91 [|||||] 83.5% 159 [|||||] 100.0% 227 [|||||] 100.0%
24 [|||||] 97.3% 92 [|||||] 3.1% 160 [|||||] 100.0% 228 [|||||] 100.0%
25 [|||||] 1.6% 93 [|||||] 95.7% 161 [|||||] 100.0% 229 [|||||] 100.0%
26 [|||||] 98.4% 94 [|||||] 100.0% 162 [|||||] 100.0% 230 [|||||] 0.4%
27 [|||||] 0.0% 95 [|||||] 100.0% 163 [|||||] 100.0% 231 [|||||] 100.0%
28 [|||||] 99.2% 96 [|||||] 0.0% 164 [|||||] 100.0% 232 [|||||] 100.0%
29 [|||||] 100.0% 97 [|||||] 100.0% 165 [|||||] 13.2% 233 [|||||] 100.0%
30 [|||||] 100.0% 98 [|||||] 100.0% 166 [|||||] 0.0% 234 [|||||] 100.0%
31 [|||||] 98.0% 99 [|||||] 0.0% 167 [|||||] 100.0% 235 [|||||] 100.0%
32 [|||||] 2.0% 100 [|||||] 100.0% 168 [|||||] 100.0% 236 [|||||] 100.0%
33 [|||||] 74.0% 101 [|||||] 11.9% 169 [|||||] 93.7% 237 [|||||] 100.0%
34 [|||||] 49.0% 102 [|||||] 95.7% 170 [|||||] 100.0% 238 [|||||] 24.6%
35 [|||||] 9.4% 103 [|||||] 89.4% 171 [|||||] 100.0% 239 [|||||] 100.0%
36 [|||||] 98.8% 104 [|||||] 98.8% 172 [|||||] 91.1% 240 [|||||] 11.7%
37 [|||||] 8.2% 105 [|||||] 93.3% 173 [|||||] 100.0% 241 [|||||] 100.0%
38 [|||||] 95.7% 106 [|||||] 2.7% 174 [|||||] 100.0% 242 [|||||] 100.0%
39 [|||||] 4.7% 107 [|||||] 100.0% 175 [|||||] 100.0% 243 [|||||] 96.9%
40 [|||||] 94.9% 108 [|||||] 3.1% 176 [|||||] 100.0% 244 [|||||] 100.0%
41 [|||||] 93.0% 109 [|||||] 7.4% 177 [|||||] 100.0% 245 [|||||] 100.0%
42 [|||||] 7.1% 110 [|||||] 92.6% 178 [|||||] 100.0% 246 [|||||] 100.0%
43 [|||||] 84.8% 111 [|||||] 13.7% 179 [|||||] 100.0% 247 [|||||] 100.0%
44 [|||||] 13.4% 112 [|||||] 100.0% 180 [|||||] 100.0% 248 [|||||] 84.4%
45 [|||||] 42.3% 113 [|||||] 71.3% 181 [|||||] 100.0% 249 [|||||] 64.2%
46 [|||||] 66.0% 114 [|||||] 31.6% 182 [|||||] 100.0% 250 [|||||] 80.3%
47 [|||||] 0.0% 115 [|||||] 100.0% 183 [|||||] 100.0% 251 [|||||] 100.0%
48 [|||||] 100.0% 116 [|||||] 0.0% 184 [|||||] 100.0% 252 [|||||] 100.0%
49 [|||||] 41.5% 117 [|||||] 56.3% 185 [|||||] 91.8% 253 [|||||] 100.0%
50 [|||||] 61.5% 118 [|||||] 58.3% 186 [|||||] 100.0% 254 [|||||] 61.8%
51 [|||||] 98.8% 119 [|||||] 0.0% 187 [|||||] 100.0% 255 [|||||] 100.0%
52 [|||||] 1.2% 120 [|||||] 100.0% 188 [|||||] 100.0% 256 [|||||] 100.0%
53 [|||||] 51.8% 121 [|||||] 100.0% 189 [|||||] 100.0% 257 [|||||] 21.1%
54 [|||||] 59.9% 122 [|||||] 100.0% 190 [|||||] 100.0% 258 [|||||] 15.1%
55 [|||||] 96.9% 123 [|||||] 100.0% 191 [|||||] 100.0% 259 [|||||] 0.4%
56 [|||||] 4.0% 124 [|||||] 100.0% 192 [|||||] 93.3% 260 [|||||] 100.0%
57 [|||||] 2.0% 125 [|||||] 100.0% 193 [|||||] 100.0% 261 [|||||] 100.0%
58 [|||||] 98.0% 126 [|||||] 0.0% 194 [|||||] 100.0% 262 [|||||] 100.0%
59 [|||||] 42.0% 127 [|||||] 100.0% 195 [|||||] 50.4% 263 [|||||] 100.0%
60 [|||||] 63.9% 128 [|||||] 27.8% 196 [|||||] 100.0% 264 [|||||] 100.0%
61 [|||||] 3.5% 129 [|||||] 100.0% 197 [|||||] 100.0% 265 [|||||] 100.0%
62 [|||||] 96.5% 130 [|||||] 0.0% 198 [|||||] 100.0% 266 [|||||] 100.0%
63 [|||||] 98.8% 131 [|||||] 0.0% 199 [|||||] 100.0% 267 [|||||] 100.0%
64 [|||||] 1.2% 132 [|||||] 100.0% 200 [|||||] 100.0% 268 [|||||] 100.0%
65 [|||||] 53.0% 133 [|||||] 31.9% 201 [|||||] 100.0% 269 [|||||] 100.0%
66 [|||||] 54.0% 134 [|||||] 96.1% 202 [|||||] 100.0% 270 [|||||] 33.7%
67 [|||||] 0.0% 135 [|||||] 100.0% 203 [|||||] 100.0% 271 [|||||] 100.0%
68 [|||||] 98.8% 136 [|||||] 100.0% 204 [|||||] 100.0% 272 [|||||] 0.0%
Mem[|||||] 37.9G/110G Tasks: 743, 593 thr; 198 running
Swp[|||||] 0K/64.0G Load average: 166.23 92.59 43.19
Uptime: 1 day, 10:57:00
PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
93990 thomas 20 0 192M 134M 9084 R 106.0 0.1 0:11.55 /opt/intel/compilers_and_libraries_2018.0.128/linux/bin/intel64/mcpcom @/tmp/icpcargxNpQ5g
94996 thomas 20 0 168M 115M 8976 R 106.0 0.1 0:09.09 /opt/intel/compilers_and_libraries_2018.0.128/linux/bin/intel64/mcpcom @/tmp/icpcargaLuv1C
1?help F2Setup F3Search F4Filter F5Tree F6SortByF7Nice F8Nice F9Gll F10Quit
15:47:04 up 1 day, 10:56
```


Scaling up

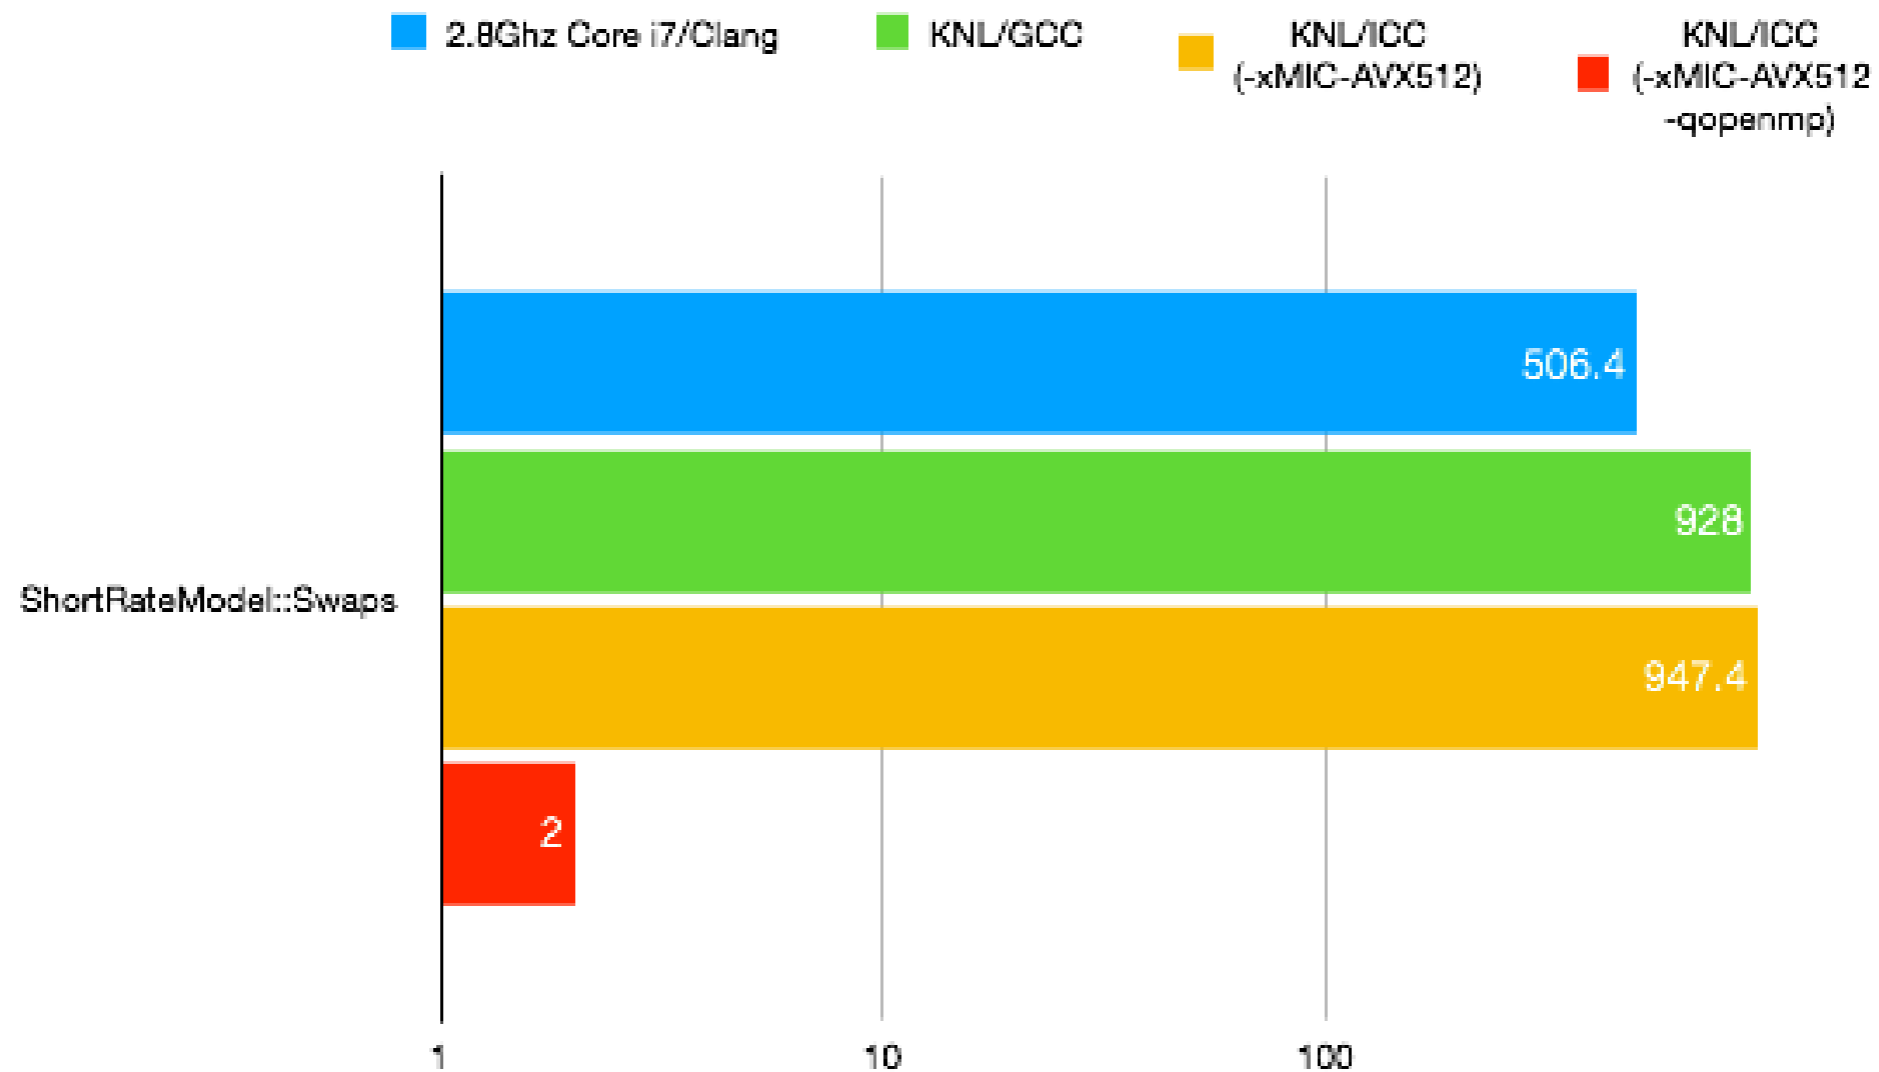
- Parallelization
 - 72 physical scalar cores
 - High bandwidth on-chip mesh
 - High bandwidth CPU adjacent memory (MCDRAM)
- Vectorization
 - Up to 16 double or 32 single precision floating point operations per vector unit, per clock
 - Two vector units per physical core

Intel Compiler Collection

Auto Parallelization

- Using `#pragma` directives to guide code generation
- Can parallelize candidate loops with either OpenMP or TBB
 - OpenMP for homogenous workloads
 - TBB for variable workloads

Auto parallelization is not always a win



What's going on here?

- Initial guess is that there's not enough work per thread to offset the communication overhead

Intel Compiler Collection

- Compiler vectorization reports
 - -qopt-report, -qopt-report-phase
- Vectorization Advisor
 - Extensive tools for analyzing and recommending vectorization opportunities
 - Not explicitly tied to code generated by the Intel Compiler
- VTune profiler
 - Sample based and “uncore” counter profiling

What's going on here?

- VTune Says -
 - ~21 of 272 logical cores utilized, ~7.8% utilization
 - ~57% time spent in serial code
 - ~50% scalar/50% SIMD instruction mix
 - Significant fraction of parallel time is load-imbalanced
 - 83% of pipeline slots remain empty

What's going on here?

- Top 3 serial hotspots -
 - `QuantLib::TrinomialTree::TrinomialTree`
 - `QuantLib::TreeLattice<QuantLib::OneFactorModel::ShortRateTree>::computeStatePrices`
- `malloc`

QuantLib is a non-trivial codebase

- Approximately 2280 source files, ~300k lines of code
- Makes heavy use of allocations, and shared pointers
 - Difficult to optimize without significant restructuring
- Detailed optimization of QuantLib is outside the scope of the work for this project

Narrowing the scope

- Currently evaluating optimization of
 - Heston model
 - SVD model w/parallel RNG
- Stripped down implementations, not part of QuantLib
- Explicit use of Intel's Math Kernel Libraries where appropriate

Early results

Scaling individual models

- Heston Model
 - Peak core utilization ~14%,
 - 75% scalar to 25% SIMD instruction mix
- SVD & Parallel RNG
 - Peak core utilization ~19%
 - 100% Packed SIMD instruction mix

Next steps

- Long term goal is not focused on individual model optimization
- Enable Quants to describe compute intensive problems in terms of high level composition of numerical algorithms
- Focus evaluation on scaling up compute utilization via explicitly parallel work partitioning

Next steps (cont)

- Custom C++ Allocators
 - Using C++ vector types backed by MCDRAM for frequently accessed data
 - Align vector types on cache-line boundary
- Non-Temporal Store
 - Avoid perturbing cache with data that is only ever written (e.g. result sets)