

montblanc-project.eu | @MontBlanc_EU

HPC Requirements for SVE. A Mont-Blanc 2020 Perspective

Adrià Armejach, Ruben Langarita, Filippo Mantovani,
Víctor Soria, Darío Suárez and Miquel Moreto
Barcelona Supercomputing Center (BSC)





**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

arm



→ Arm – BSC Centre of Excellence

- “In recognition of the leadership and hard work of BSC in pioneering Arm in HPC and the success of EU-funded MontBlanc projects”
- Multi-year agreement starting in 2019.
- Broaden the current scope of interaction and collaboration between BSC and Arm beyond HPC

→ Research activities

- Acceleration of precision medicine workloads
- Advanced memory systems for sparse and irregular memory access patterns
- Autonomous driving

→ Outreach and education activities

- Arm SVE Hackathon: "Arm Scalable Performance for HPC and ML" (Feb 2019)
- Future events organized together with PRACE
- Participation in the Student Cluster Competition at ISC (Jun 2019)

The Student Cluster Competition

Competition rules

- Six undergraduate students
- One cluster managed by the team
- 3 kW power limit

Competition awards

- Best Linpack performance
- 1st, 2nd and 3rd overall
- Fan favorite

The challenge

- HPC Benchmarks
 - HPL, HPCG
- Real scientific applications
 - Quantum Espresso, TensorFlow
- Judge interviews
- Secret challenges
 - Secret application, Power outage



The educational package

Effort


- ~160h of work
Amounts to six ECTS
European Credit Transfer System
- ~530h of cluster usage
- 23 students (and counting!)

Cluster and other hardware resources

- Supported by sponsors
- Absorbed by university

Daily cost when traveling

- ~100 EUR/day per student during competition



■ High Education - 8 (34.78%) ■ Research - 6 (26.09%)
■ Industry - 4 (17.39%) ■ Unrelated to HPC - 5 (21.74%)

**~80% of students actively engaged in HPC
after going through the educational package**

F. Mantovani, and F. Banchelli, "Filling the gap between education and industry: evidence-based methods for introducing undergraduate students to HPC", EduHPC18 workshop held in conjunction of Supercomputing 2018, Dallas.

Mont-Blanc roadmap

Today

Vision: leverage the fast growing market of mobile technology for scientific computation

2012

2013

2014

2015

2016

2017

2018

2019

2020



- Design SoC
- Develop IPs

Prepare an industrial solution
Test market acceptance



Mont-Blanc 2020

Extend the concept
and explore new
possibilities



Mont-Blanc 3

Proof of concept : HPC
computing based
on mobile embedded
technology



Mont-Blanc 2



Mont-Blanc

MONT-BLANC

- **Define a low-power System-on-Chip (SoC) architecture targeting Exascale using Arm ISA**
 - Leverage Arm's strong technological relevance covering embedded to high end solutions
 - Dynamic ecosystem, needed to deliver the system software and applications mandatory for successful market acceptance.

- **To design, implement and leverage new technologies to improve the performance, power efficiency, reliability and security of the processor:**
 - Exploit SVE compute units
 - Specific Network-on-Chip (NoC) with high bandwidth and low power characteristics
 - Specific SoC Power Management controller
 - Innovative packaging technologies
 - State of the art silicon process

- **To improve on the economic sustainability of processor development through a modular design that allows to retarget the SoC for different markets**

Requirements for SVE: Design-space exploration

→ Explore hardware trade-offs

- Different vector lengths (128 to 2048 bits)
 - Off-chip memory bandwidth demands
 - Pin-based with DDR interfaces
 - Silicon interposer with HBM interfaces
- Sizing of key core structures – e.g., load/store queue
- Number of vector units – i.e., 2 x 256bit vs. 1 x 512bit
- Recommended memory bandwidth per core

→ Characterize performance of applications

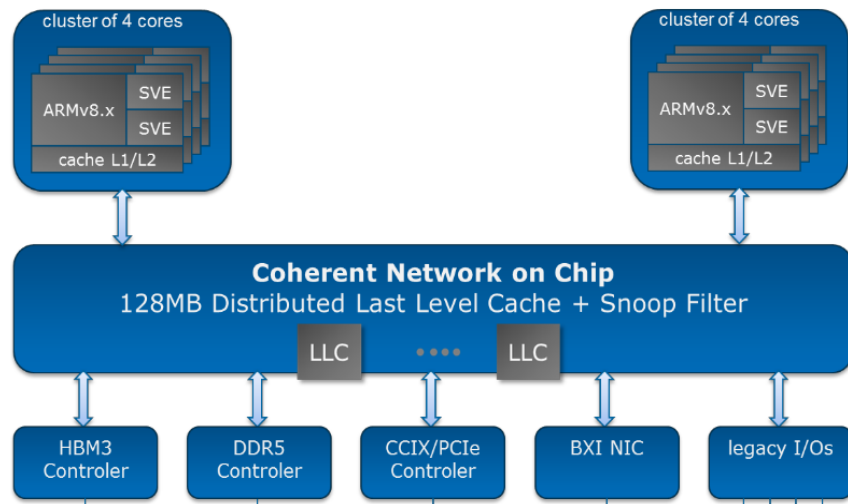
- Floating point operations per demanded bytes – operational intensity (OI)
 - End goal: Understand how applications perform under different vector lengths and memory bandwidth constraints

→ Main challenge at the beginning of the project

- Tools maturity for SVE (compiler, emulator, simulator, etc.)...
- Has improved a lot in the last few years!

→ Simulation-based studies using gem5

- Result using simple but relevant loops
 - Simulate a representative subset of the RAJAPerf loops
- Resemble the architecture envisioned in MB2020
 - Clusters of 4 cores with private L1/L2 caches
 - Shared multi-bank L3 cache
 - Evaluate DDR and HBM interfaces



Requirements for SVE – Benchmarks

- **Simulate a representative subset of the RAJAPerf loops**
 - Good auto-vectorization coverage with gcc and Arm HPC compiler

Benchmark	Description
MULADDSUB	Simple synthetic loop for quick testing
HYDRO_1D	Main computational loop of the HYDRO application
EOS	Calculates the equation of state
INT_PREDICTOR	Integral predictor
ENERGY	Kernel to calculate energy states, extracted from an LLNL application
PRESSURE	Kernel to calculate pressure in a system, extracted from an LLNL application
FIR	Finite impulse filter, extracted from an LLNL application
LTIMES	Kernel extracted from an LLNL application
Stream-DOT	Stream benchmark, dot-product version
VOL3D	Kernel to calculate a 3D volume, extracted from an LLNL application

Requirements for SVE – Simulation parameters

- Resemble architecture envisioned in MB2020
- Simulation parameters

	Description
Processor size	8 cores - 2 clusters of 4 cores each
Cores	3-wide issue/retire, 64-entry instruction queue, 192-entry ROB, 48LDQ + 48STQ, 2 Vector processing units (VPU), 2GHz
L1I	64KB, 4-way, 2 cycle, 8MSHR
L1D	64KB, 4-way, 2 cycle, 24MSHR
L2 (private)	256KB, 2-way, 7 cycle, 24MSHR, stride prefetcher
Last-level Cache	16MB in 8 banks, 16-way, 20 cycles, 64MSHR
NoC	Coherent crossbar, 128-bit wide, 2 cycles
Main memory	4 DDR4-2400, 2 ranks/channel, 16 banks/rank, 8KB row-buffer 76.8 GB/s peak bandwidth
	1 HBM stack, 8 channels/stack 128-bit wide, 8 banks/channel 128 GB/s peak bandwidth

Requirements for SVE – Vector length results (DDR)

→ Peak performance (compute ceilings)

- Peak SVE128 = 2GHz (cycles/s) × 8 (#cores) × 2 (flop/instr) × 2 (instr/cycle) = **64GFlops/s**

→ Memory ceilings:


- Peak 76.8 GB/s
- Measured ~72GB/s**

→ Operational intensity

- Most benchmarks <1
- Common values

→ Bandwidth quickly becomes a bottleneck

- Most benchmarks below scalar peak performance
- VOL3D benefits from SVE 256 bits
- DDR-based memory not an option to scale
- SVE 512 bit performance only achievable with OI > 3.5



Requirements for SVE – Vector length results (HBM)

→ Peak performance (compute ceilings)

- Peak SVE128 = 2GHz (cycles/s) × 8 (#cores) × 2 (flop/instr) × 2 (instr/cycle) = **64GFlops/s**

→ Memory ceilings:

- Peak 128 GB/s
- 16GB/s per core

→ Operational intensity


- Most benchmarks <1
- Common values

→ HBM improves SVE scalability

- Memory bound benchmarks benefit greatly
- EOS shows significantly better performance with SVE256

→ SVE 256 likely to be the best candidate

- Need at least OI == 2 to reach peak SVE 512 bit performance



Roofline Models for ThunderX2 (Dibona)

→ Peak performance (compute ceilings)

- Peak NEON = 2GHz × 32 (#cores) × 2 (elem/vector) × 2 (flop/instr) × 2 (instr/cycle) = **512GFlops/s**

→ Memory ceilings:


- Peak 170 GB/s

→ Operational intensity

- Most benchmarks <1
- Common values

→ All benchmarks scale well until memory bandwidth is exhausted

- Most benchmarks below scalar peak performance
- JACOBI_2D, VOL3D, HACCK and FIR benefit from NEON 128 bits SIMD
- Compilers can squeeze the memory bandwidth in FIR and HACCK



→ Summary

- Use of HBM technology mandatory to scale to more than 8 cores per socket
- Recommend the use of SVE 256 or SVE 512 bits
 - SVE 512bits starts to be beneficial with $OI > 1.5$
 - SVE 512bits offers marginal benefits for $OI < 1$
- Contemporary aggressive OoO HW structures sufficient for SVE 256 & 512 bits
 - 48 LQ, 48SQ, 92 IQ, 192 ROB
 - 33% larger HW structures offer modest improvements – 17% max. 5% avg.
- We recommend a minimum of 16GB/s of off-chip memory bw per core

→ Next steps

- Look at a wider range of benchmarks
 - HACCKernels, HPCG, MiniAMR, SWFFT, XSBench, ...
- Larger SoC configurations
 - 16 cluster of 4 cores (64 cores in total)
- Explore different compiler optimizations



montblanc-project.eu | @MontBlanc_EU

HPC Requirements for SVE. A Mont-Blanc 2020 Perspective

Adrià Armejach, Ruben Langarita, Filippo Mantovani,
Víctor Soria, Darío Suárez and Miquel Moreto
Barcelona Supercomputing Center (BSC)

