

Russian Supercomputing Days September 25-26, 2017, Moscow

# Как приблизиться к пиковому значению Flops/s? Сравнение архитектур x86 и ARMv8



UNIVERSIT

#### V. Nikolskiy, V. Stegailov

International Laboratory for Supercomputer Atomistic Modelling and Multi-scale Analysis, Higher School of Economics National Research University Higher School of Economics Intl Lab



Supercomputer Atomistic Modelling and Multi-scale Analysis

Web: samma.hse.ru Email: {vnikolskiy, vstegailov}@hse.ru



# Mobile and embedded market leader





# Architecture evolution

Свойство	ARMv6	ARMv7	ARMv8
64-bit	-	-	+
Multicore	-	+	+
Cache hierarchy	-	+	+
SP Scalar Floating-point	Optional	Optional	+
DP Scalar Floating-point	Optional	Optional	+
SP Floating-point SIMD	-	Optional	+
DP Floating-point SIMD	-	-	+



#### Mont-Blanc 2011-2018:

European approach towards energy efficient high performance computation

- To produce an architecture that will provide Exascale performance using 15 to 30 times less energy
- To develop the Mont-Blanc software ecosystem, with emphasis on programmer tools and system resiliency
- To design a well-balanced architecture for an ARM based SoC or SoP capable of providing preexascale performance when implemented in the 2019-2020





# ARMv8 server-level processors

**Cavium** ThunderX2

- 54 cores @ 3.0 GHz (14 nm FinFET)
- 6 DDR4 controllers
- Integrated SATAv3
- Integrated PCIe Gen3

#### MACOM (AppliedMicro) X-Gene

- 32 cores @ 3.0+ GHz (16 nm FinFET)
- 8 DDR4-2667 memory channels
- 42 PCIe Gen3 lanes with eight controllers

#### Qualcomm

- 48 cores (10-nm FinFET)
- Multiple DDR4 memory controllers
- PCIe Gen3 x16
- Integrated IO and SATAv3 ports

#### AMD Opteron A1100

- 8 Cortex-A57 cores @ 2.0 GHz
- 2 DDR4-1866 memory channels
- 8 PCIe Gen3 lanes
- 14 SATA 3 (6 GB/s)



# The Scalable Vector Extension (SVE)

Aarch64 extension which expand vector up to 2048 bits

- Expand fine-grain parallelism for HPC scientific workloads
- Reduce software deployment effort
- Enables vectorization of complex data structures with non-linear access patterns
- Permits vectorization of uncounted loops with data-dependent exits



# Fujitsu Post-K

- Based on 64-bit ARMv8-A + SVE
- Exaflops supercomputer
- 6D mesh/torus interconnect "Tofu"
- 10nm FinFET
- Be completed by 2022
- Budget \$910m
- Power consumption 30-40 MW





- Quad core processor Amlogic S805 Cortex-A5
  - 1.5GHz
  - VFPv4 modules on each core
  - graphics accelerator Mali-450 MP2 (was not used)
- 1GB DDR3 SDRAM
- OS:
  - Linux Ubuntu 14.05.1
  - or Android 4.4.2 OS





#### LAMMPS performance for different CPUs





#### Nvidia Jetson

### Nvidia Jetson TK1 (2014)





- 32-bit Tegra K1 SoC
- 4 GB memory
- 4 Cortex-A15 cores @ 2.3 GHz
- 1 Kepler SM @ 852 МГц, 128 CUDA cores
- 64-bit Tegra X1 SoC
- 4 GB 64-bit memory
- 4 Cortex-A57 cores @ 2.1 GHz
- 2 Maxwell SM @ 998 МГц, 256 CUDA cores



# MD algorithms performance



Nikolskiy, V.P., Stegailov, V.V., Vecher, V.S.: Efficiency of the Tegra K1 and X1 systems-on-chip for classical molecular dynamics. In: 2016 International Conference on High Performance Computing Simulation (HPCS), pp. 682–689 (2016). DOI 10.1109/HPCSim.2016.7568401



### Nvidia Jetson

# Nvidia Jetson TX2 (2017)



- 2 Denver2 cores + 4 Cortex-A57 cores
  @ 2.0 GHz
- 8 GB 128-bit memory
- Pascal GPU



# Julich Prototype cluster

- 32 Cortex-A57 cores @ 2.1 GHz (64-bit)
- 1 MB L2 cache, 32 MB L3 cache
- 128 GB memory

- Ubuntu 16.04 LTS
- gcc v5.3.1



## ARMv8 Cortex-A pipeline



 $R(op) = CPU_{freq} \cdot k_{OP}(op) \cdot k_{SIMD}(op) \cdot throughput(op)$ 



# Microbenchmark development

#### Based on github.com/Mysticial/Flops

😒 🖱 💿 bench_f64v1_aarch64.h — ~/workspace/Flops/version2/source — Atom						
bench_f	54v1_aarch64.h main.cpp arch_Aa	irch64.h		arch_2013_Haswell.h		
444	float64x2_t rA = $\{2.0f, 2.0f\}$	01f};				
445	$float64x2_t rB = \{2.1f, 2.1\}$	11f};				
446						
447	<pre>for (size_t i = 0; i &lt; iter</pre>	rations; i++)	{			
448	r0 = vfmaq_f64(r0, r	mul0, add0);	r1 = vaddq_f64(r	1, add0);		
449	r2 = vfmaq_f64(r2, r	mul0, add0);	r3 = vaddq_f64(r	3, add0);		
	r4 = vfmaq_f64(r4, r	mul0, add0);	r5 = vaddq_f64(r	5, add0);		
	r6 = vfmaq_f64(r6, i	nul0, add0);	r7 = vaddq_f64(r	7, add0);		
	r8 = vfmaq_f64(r8, r	nul0, add0);	r9 = vaddq_f64(r	9, add0);		
	rA = vfmaq_f64(rA, r	nul0, add0);	rB = vaddq_f64(rl	3, add0);		
	r0 = vfmsq_f64(r0, r	nul1, add0);	r1 = vsubq_f64(r	1, sub0);		
	r2 = vfmsq_f64(r2, r	nul1, add0);	r3 = vsubq_f64(r3	3, sub0);		
	r4 = vfmsq_f64(r4, r	nul1, add0);	r5 = vsubq_f64(r	5, sub0);		
	r6 = vfmsq_f64(r6, r	nul1, add0);	r7 = vsubq_f64(r	7, sub0);		
	r8 = vfmsq_f64(r8, r	nul1, add0);	r9 = vsubq_f64(r	9, sub0);		
	rA = vfmsq_f64(rA, r	nul1, add0);	rB = vsubq_f64(rl	3, sub0);		
462						
	<pre>flops_reduce_chains12(</pre>					
	vaddq_f64,					
	r0, r1, r2, r3, r4, i	r5, r6, r7, r	8, r9, rA, rB			
	);					
467	<pre>result = vaddvq_f64(r0);</pre>					
bench_128/be	nch_F64v1_aarch64.h 467:32			C++ 🖗 master 🛨+479		

bench_f64v1_aa	arch64.h — ~/workspace/Flops,	/version2/source — Atom		
ench_f64v1_aarch64.h				
	vaddq_f64,	vmulq_f64,		
	add0, mul1	, mul0,		
	r10, r11,	r12, r13, r14, r15	ō, r16, r17, r18,	r19, r1A, r1B
	);			
	flops_muladd_c	hains12_ops12(		
	vsubq_f64,	vmulq_f64,		
	sub0, mul0	), mul1,		
	r0, r1, r2	2, r3, r4, r5, r6,	r7, r8, r9, rA, i	-B
	);			
	flops_muladd_c	chains12_ops12(		
	vsubq_f64,	vmulq_f64,		
	sub0, mul1	, mul0,		
	r10, r11,	r12, r13, r14, r15	5, r16, r17, r18,	r19, r1A, r1B
	);			
	}			
	flops_reduce_chair	ns12(		
	vaddq_f64,			
	r0, r1, r2, r3	3, r4, r5, r6, r7,	r8, r9, rA, rB	
	);			
	flops_reduce_chair	ıs12(		
	vaddq_f64,			
	r10, r11, r12,	. r13, r14, r15, r1	16, r17, r18, r19	, r1A, r1B
	);			
	result = vaddvq_f6	54(r10)+vaddvq_f64	(r0);	
	return iterations	* 2 * 12 * 4.		
28/bench_f64v1_aarch6	4.h 467:32			8 C++ 🖗 master 🛨+4



# Microbenchmark results



Haswell



#### Jülich Prototype Cluster

Jetson TX2 Denver2



Jetson TX2 Cortex-A57





# Benchmark comparison

Intel Haswell



$$a = a \cdot c + b$$
  
 $a = b \cdot a + c$   
 $a = b \cdot c + a$ 



 $a = b \cdot c + a$ 

ARMv8



## Conclusions

- The efficiency of ARM Cortex-A was compared with other CPUs (Intel x86\_64 and legacy) for the classical MD algorithm example (LJ liquid in LAMMPS).
- The theoretical peak performance of the processors ARM Cortex and Nvidia Denver2 was calculated
- The highly optimized test code was created, that achieves the highest proportion of the calculated peak performance
- We considered the case of incorrect results of popular test Empirical Roofline Toolkit with ARMv8 processors
- We have shown that LINPACK gives correct and predictable results