



ORGANIZACIJA RAČUNALNIKOV

5 Paralelizem na nivoju procesorjev

Namen in cilji 5. poglavja:

- Paralelnost poskušamo izkoristiti na **več različnih nivojih**
- Prednosti in omejitve paralelizma na nivoju procesorjev
 - Kako **gradimo in povezujemo** ?
 - Kako **programiramo** ?
- **Multiprocesorji, multiračunalniki**
- **Superračunalniki, podatkovni centri:**
 - Z/brez cenovnega kompromisa
 - Razmerje cena/zmogljivost :
 - **Kako to dela Google** ?
- **Zanimivi pristopi:**
 - **GPU, podatkovno pretokovno računanje, „Big-brain chips“, Spinnaker**

Paralelnost:

- je problem (programiranje, učinkovitost)
- je edino upanje (omejitve tehnologije)

Primer:

- želimo: 1 CPE s $t_{CPE} = 0.001\text{ns}$
- dobimo: 1000 CPE s $t_{CPE} = 1\text{ns}$

- teoretično enaka zmogljivost, v čem je težava?
 - slabo: zaporedno razmišljanje, problemi, algoritmi, programi
 - dobro: imamo probleme, ki so „paralelni“
- razlika v porabi (v prid paralelnosti)

- Za $IPC > 10$ edina možnost !

Omejitev pohitritve v paralelizmu: Amdahlov zakon

Paralelnost :

- Problem ali rešitev ?

Pohitritev programa zaradi paralelnega izvrševanja v multiprocesorskem sistemu **je omejena** z deležem **programa, ki ga lahko paralelno izvršujemo.**

$$S(N) = \frac{1}{f + (1-f)/N} = \frac{N}{1 + (N-1)f}$$

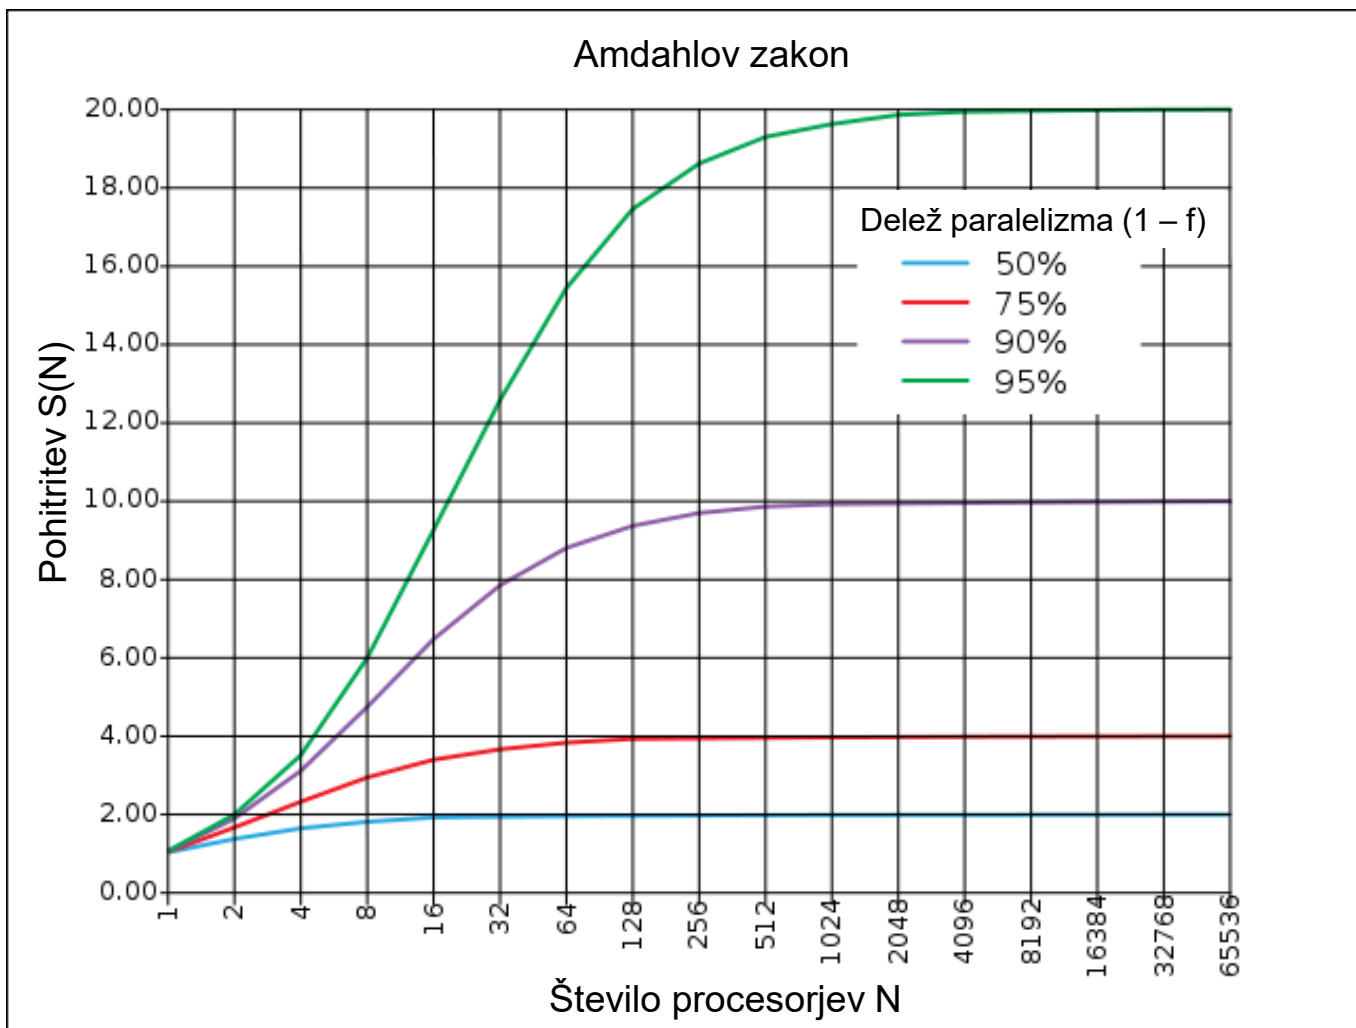
f = delež _operacij, ki _niso _pohitrene

$(1-f)$ = delež _operacij, ki _so _ N - krat _pohitrene

$$\lim_{N \rightarrow \infty} S(N) = \frac{1}{\frac{1}{N} + \frac{N}{N}f - \frac{1}{N}f} = \frac{1}{f}$$

f = delež _operacij, ki _niso _pohitrene

Pohitritev izvrševanja programa v odvisnosti od števila procesorjev in deleža izkoriščenega paralelizma v programu



Vrste paralelnih računalnikov glede na povezanost in lokacijo

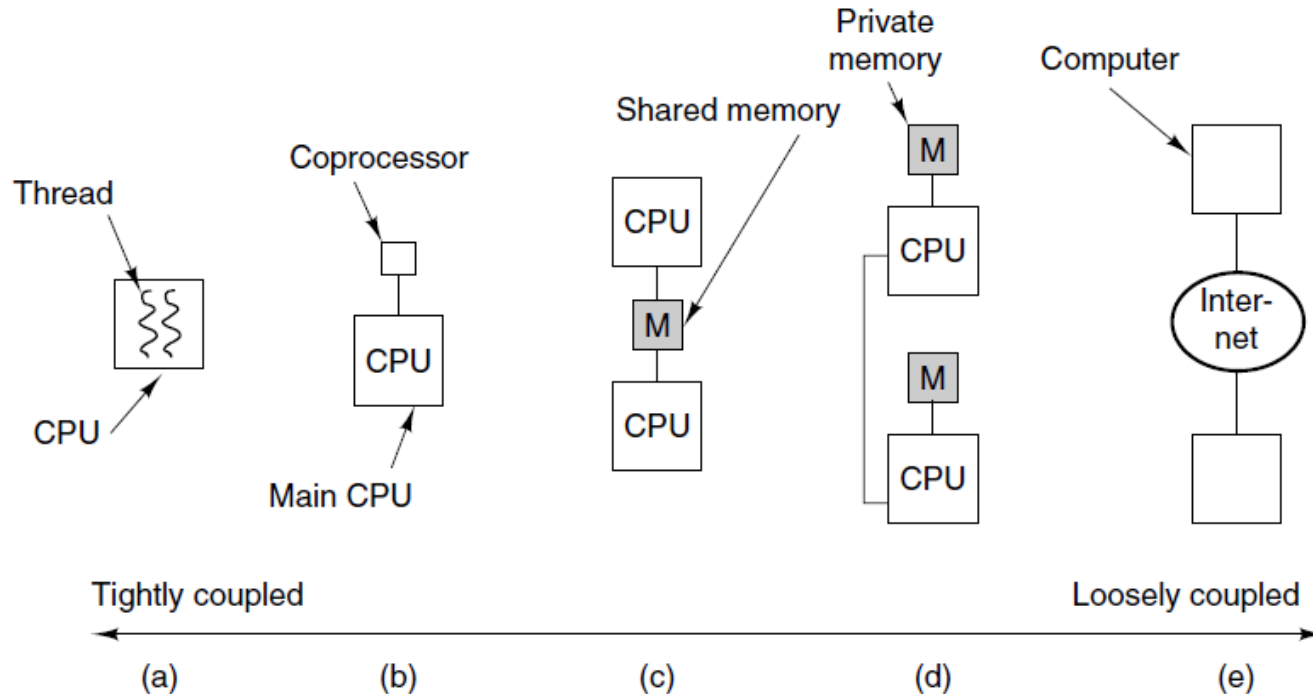


Figure 8-1. (a) On-chip parallelism. (b) A coprocessor. (c) A multiprocessor. (d) A multicomputer. (e) A grid.

- Na čipu (»On-Chip«) na nivoju ukazov (pogl. 4) večnitnost (pogl. 4.9) multiprocesorji na 1 čipu

- Koprocesorji (»Coprocessor«)
- MultiProcesorji (»MultiProcessor«)
- MultiRačunalniki (»MultiComputer«)
- Omrežje (»Grid«)


Različni nivoji uvajanja paralelizma z bolj tesno ali ohlapno povezanostjo:

■ **Na čipu (»On-Chip«)**

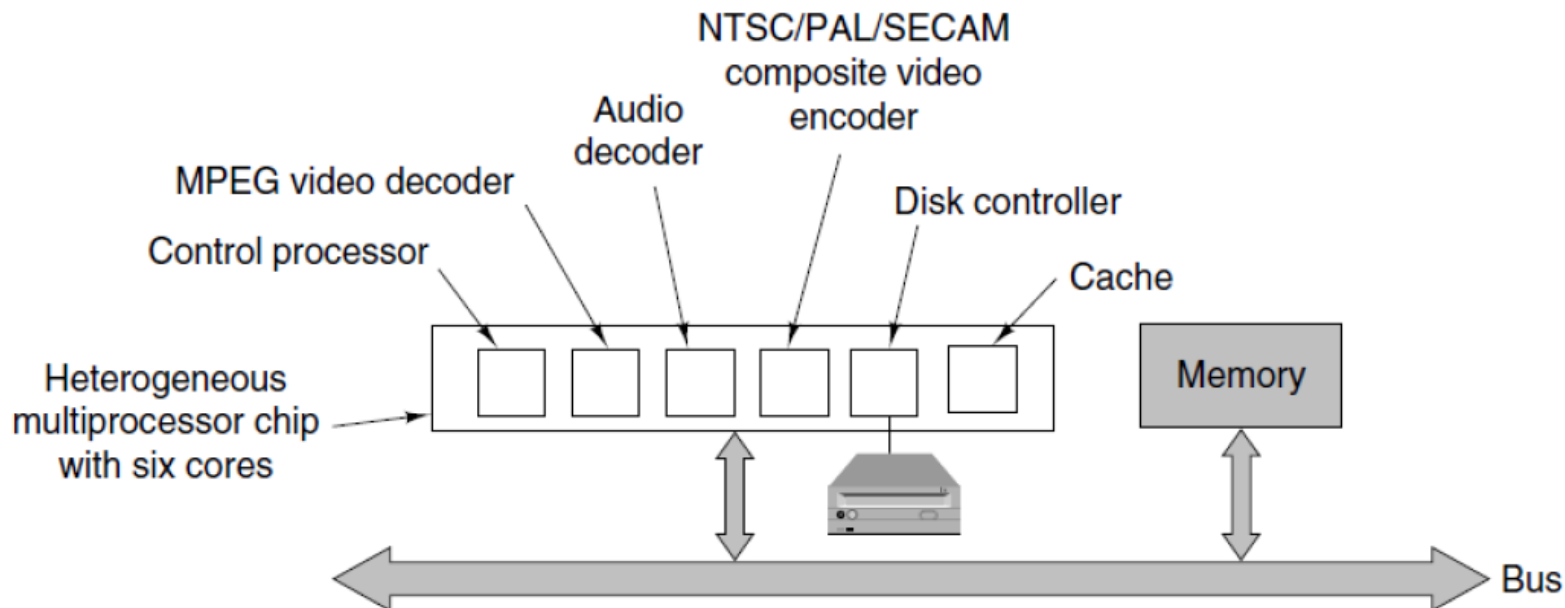
- na nivoju ukazov (**pogl. 4**)
- večnitnost (**pogl. 4.9**)
- multiprocesorji na 1 čipu**

5.1.1 Multiprocesorji na čipu („On Chip“) – „homogeni“

Core i7:



Primer heterogenega multiprocesorja : DVD predvajalnik.



Samsung Galaxy S5, S6, S7, S8 in S20

SoC		CPU				GPU	GPU Performance GFLOPS	Memory technology
Model number	fab	Instruction set	Microarchitecture	Cores	Frequency (GHz)			
Exynos 5 Octa ^[53] (Exynos 5422)	28 nm HKMG	ARMv7	Cortex-A15+ Cortex-A7 (big.LITTLE with GTS)	4+4	1.9 1.3	ARM Mali-T628 MP6 @ 533 MHz	102.4 (FP32)	32-bit dual-channel 933 MHz LPDDR3/DDR3 (14.9 GB/s)
Exynos 7 Octa 7420 ^{[64][65]}	14 nm LPE	ARMv8-A	Cortex-A57+ Cortex-A53 (big.LITTLE with GTS)	4+4	2.1 1.5	Mali-T760 MP8 @ 772 MHz	210 (FP32) ^[66]	32-bit dual-channel 1552 MHz LPDDR4 (24.88 GB/s) ^[67]
Exynos 8 Octa 8890 ^[92]	14 nm LPP		Exynos M1 "Mongoose"+ Cortex-A53 (GTS) ^[93]	4+4	M1: 2.6 (1-2 cores load) 2.3 (3-4 load) A53: 1.6 M1: 2.0 A53: 1.5 (Lite)	Mali-T880 MP12 Mali-T880 MP10 (Lite)	650 265.2 221 (Lite)	LPDDR4
Exynos 9 Octa 8895 ^{[95][96][97]}	10 nm LPE		Exynos M2 "Mongoose"+ Cortex-A53 (GTS)	4+4	M2: 2.314 / A53: 1.69	Mali-G71 MP20	546 375	LPDDR4X
Exynos 2100 (S5E9840) ^[157]			1 + 3 + 4 cores (2.91 GHz Cortex-X1 + 2.81 GHz Cortex-A78 + 2.2 GHz Cortex-A55)			Mali G78 MP14	854 1,530	LPDDR5

Različni nivoji uvajanja paralelizma z bolj tesno ali ohlapno povezanostjo (nadaljevanje):

■ **Koprocessori** (»Coprocessor«)

- Mrežni, medijski, kriptografski, GPU

■ **MultiProcesorji** (»MultiProcessor«)

- enostavni (skupen naslovni prostor) za programiranje
- tesneje povezani, slabše skalirajo


■ **MultiRačunalniki** (»MultiComputer«)

- ločeni naslovni prostori, težje programiranje („porazdeljen pomn.“)
- ohlapnejše povezave, boljše skaliranje

■ **Omrežje** (»Grid«)


- Heterogeni, še ohlapneje povezani , geografsko bolj razpršeni
- SETI@Home, „Smart Energy Grid“

I.1966 Flynn· poda· **klasifikacijo· paralelnih· računalnikov·** glede· na· število· ukazov· in· število· operandov·,ki·se·izvajajo·hkrati:



5.2.1 SISD (Single Instruction, Single Data)

- Von Neumannov model




Source: ARS Technica

5.2.2 SIMD (Single Instruction, Multiple Data)

Ideja: poenostaviti KE, izkoristiti paralelizem na nivoju podatkov

- SIMD razširitve:
 - Intel: MMX, SSE ukazi
 - ARM: NEON kot SIMD razširitve
- vektorski, „array“ procesorji,
 - CRAY (najbolj znani vektorski rač.)

Figure 4.5. NEON vector-by-scalar multiplication




5.2.2 SIMD (Single Instruction, Multiple Data):

Primer: matrično množenje: (ARM: NEON kot SIMD razširitev) :

Figure 4.4. Matrix multiplication showing one column of results

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \cdot \begin{bmatrix} b_{11} & b_{12} & b_{13} & b_{14} \\ b_{21} & b_{22} & b_{23} & b_{24} \\ b_{31} & b_{32} & b_{33} & b_{34} \\ b_{41} & b_{42} & b_{43} & b_{44} \end{bmatrix} = \begin{bmatrix} a_{11} \cdot b_{11} + a_{12} \cdot b_{21} + a_{13} \cdot b_{31} + a_{14} \cdot b_{41} & \dots & \dots & \dots \\ a_{21} \cdot b_{11} + a_{22} \cdot b_{21} + a_{23} \cdot b_{31} + a_{24} \cdot b_{41} & \dots & \dots & \dots \\ a_{31} \cdot b_{11} + a_{32} \cdot b_{21} + a_{33} \cdot b_{31} + a_{34} \cdot b_{41} & \dots & \dots & \dots \\ a_{41} \cdot b_{11} + a_{42} \cdot b_{21} + a_{43} \cdot b_{31} + a_{44} \cdot b_{41} & \dots & \dots & \dots \end{bmatrix}$$

Figure 4.5. NEON vector-by-scalar multiplication



5.2.3 MISD (Multiple Instruction, Single Data):

ne obstajajo, nekateri sem štejejo cevovod...



5.2.4 MIMD (Multiple Instruction, Multiple Data):

Ideja: izvajati hkrati več ukazov, vsakega nad svojimi operandi (podatki)

Značilnosti:

- Splošnejši od SIMD, izkoristijo več paralelizma
- V praksi pogosto SPMD („Single Program Multiple Data“)

Ločimo 2 skupini:

- Multiprocesorji (skupen naslovni prostor) – „Shared Memory Multiproc.“
 - CPE tesno povezane s skupnim pomnilnikom
 - + enostavno programiranje, - slabša skalabilnost
- Multiračunalniki (ločeni pomnilniki) – „Message Passing Multiproc.“
 - rahlo povezani, vsaka CPE svoj pomnilnik
 - + cenejši, + bolj skalabilni, - težje programiranje

Primerjava MIMD in SIMD:

■ z vidika MIMD:

- + za gradnjo se lahko uporabijo obstoječi procesorji (nižja cena)
- + skalabilnost; uporabnik število procesorjev prilagaja svojim potrebam in zmožnostim.
- - težko narediti **prevajalnik**, ki bi poleg podatkovne izkoriščal še ukazno paralelnost

■ Klasifikacija Tanenbaum:

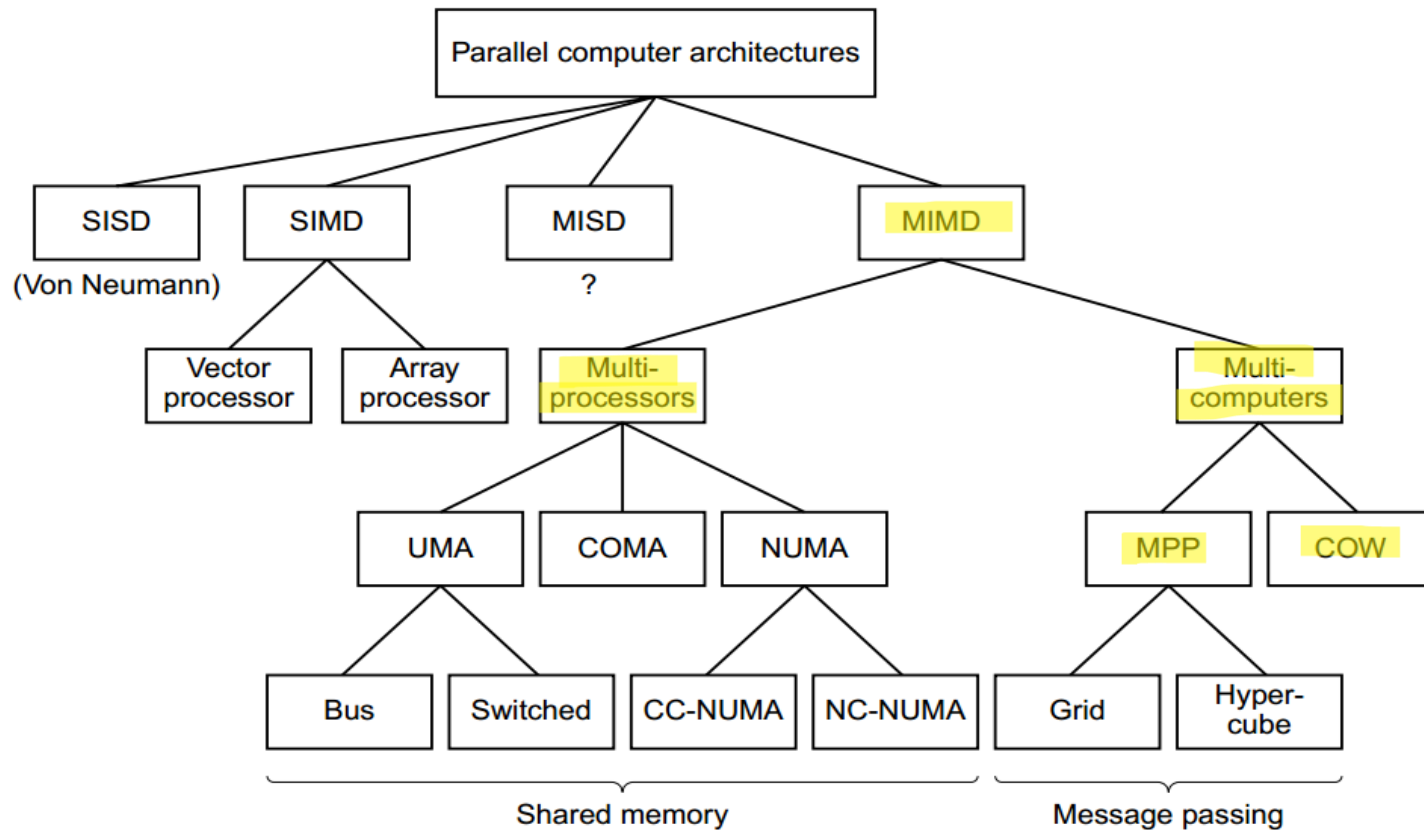



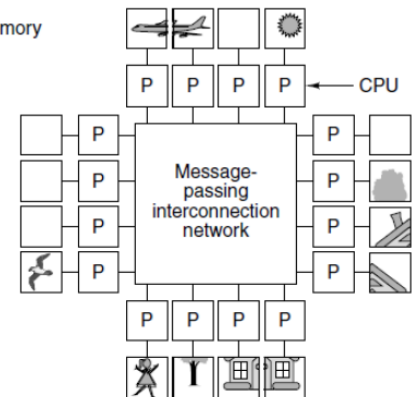
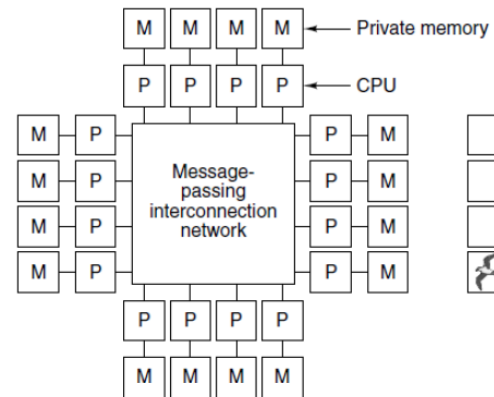


Figure 8-14. A taxonomy of parallel computers.

5.2.4 MIMD (Multiple Instruction, Multiple Data):



Multiprocessors

Multicomputers

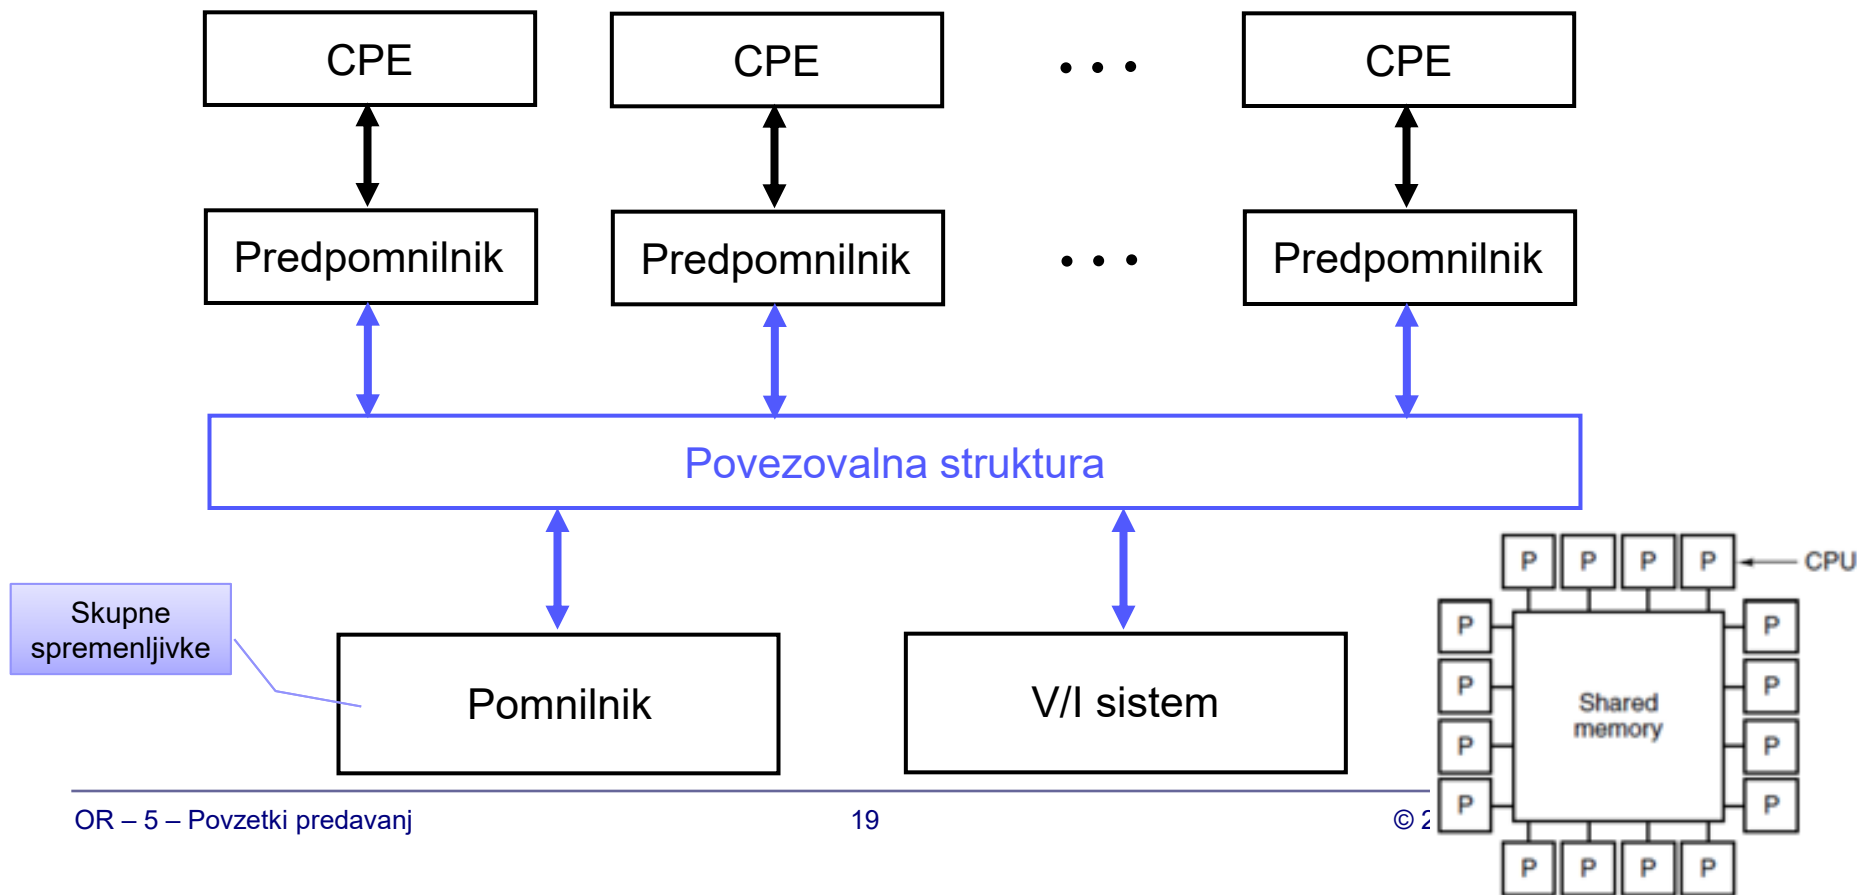


5.2.4.1 Multiprocesorji s skupnim pomnilnikom

Ideja:

- predelava starih programov za paralelne računalnike je zahtevna
- ena od možnih poenostavitvev je skupen pomnilniški prostor

Tipična organizacija multiprocesorja s skupnim pomnilnikom (tesna povezanost)



UMA - nekaj tipičnih multiprocesorjev s skupnim pomnilnikom in vodilom

UMA – Uniform Memory Access




Figure 8-22. Three bus-based multiprocessors. (a) Without caching. (b) With caching. (c) With caching and private memories.

Prednosti :

+ enostavno programiranje

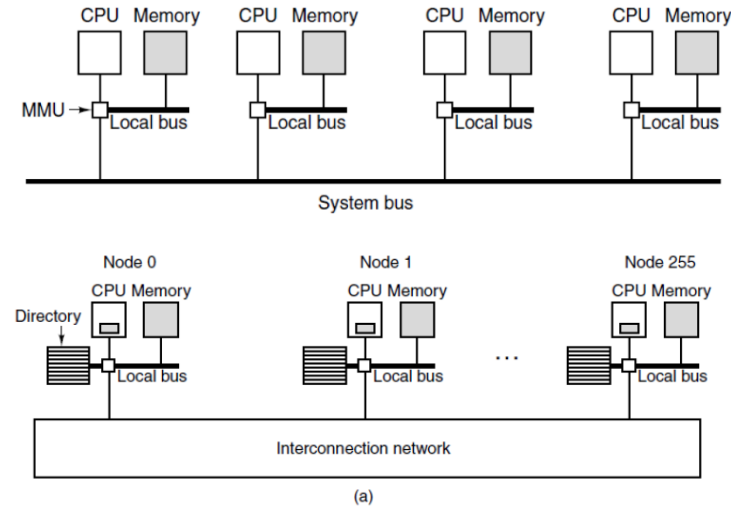
+ enotni čas dostopa do kateregakoli podatka

- je hkrati tudi največja omejitev sistema! (v katerem smislu ?)

Tipični multiprocesorji s skupnim pomnilnikom: NUMA, COMA

NUMA:

- Prednosti:
 - hitrejši lokalni (počasnejši oddaljeni) dostopi, poenostavitev sistema, skalabilnost, hitrost
- Slabost
 - težje programiranje
- Izvedbi:
 - NC-NUMA (Non Cached NUMA)
 - vodilo je ozko grlo
 - + ni problema skladnosti vsebine predpomnilnikov
 - CC-NUMA (Cache Coherent NUMA)
 - + hitrejši dostopi
 - problem skladnosti vsebine predpomnilnikov

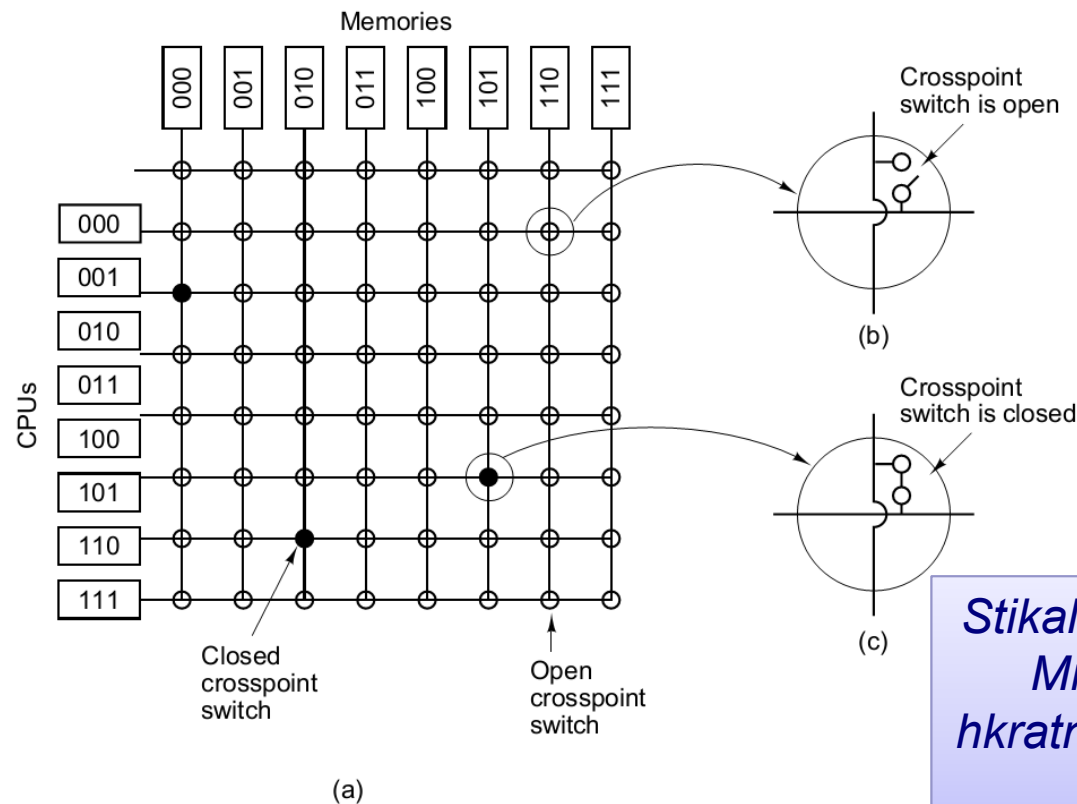


COMA: („Cache Only Memory Access“)

- naslovni prostor se razdeli na vrstice v predpomnilnikih
- lokalni pomnilnik = predpomnilnik
- ni veliko realizacij po tem pristopu

UMA: tipični multiprocesor s skupnim pomnilnikom

UMA in stikalno mrežo



*Stikalna mreža vs. vodilo ?
Mreža omogoča več
hkratnih povezav, je pa bolj
kompleksna.*

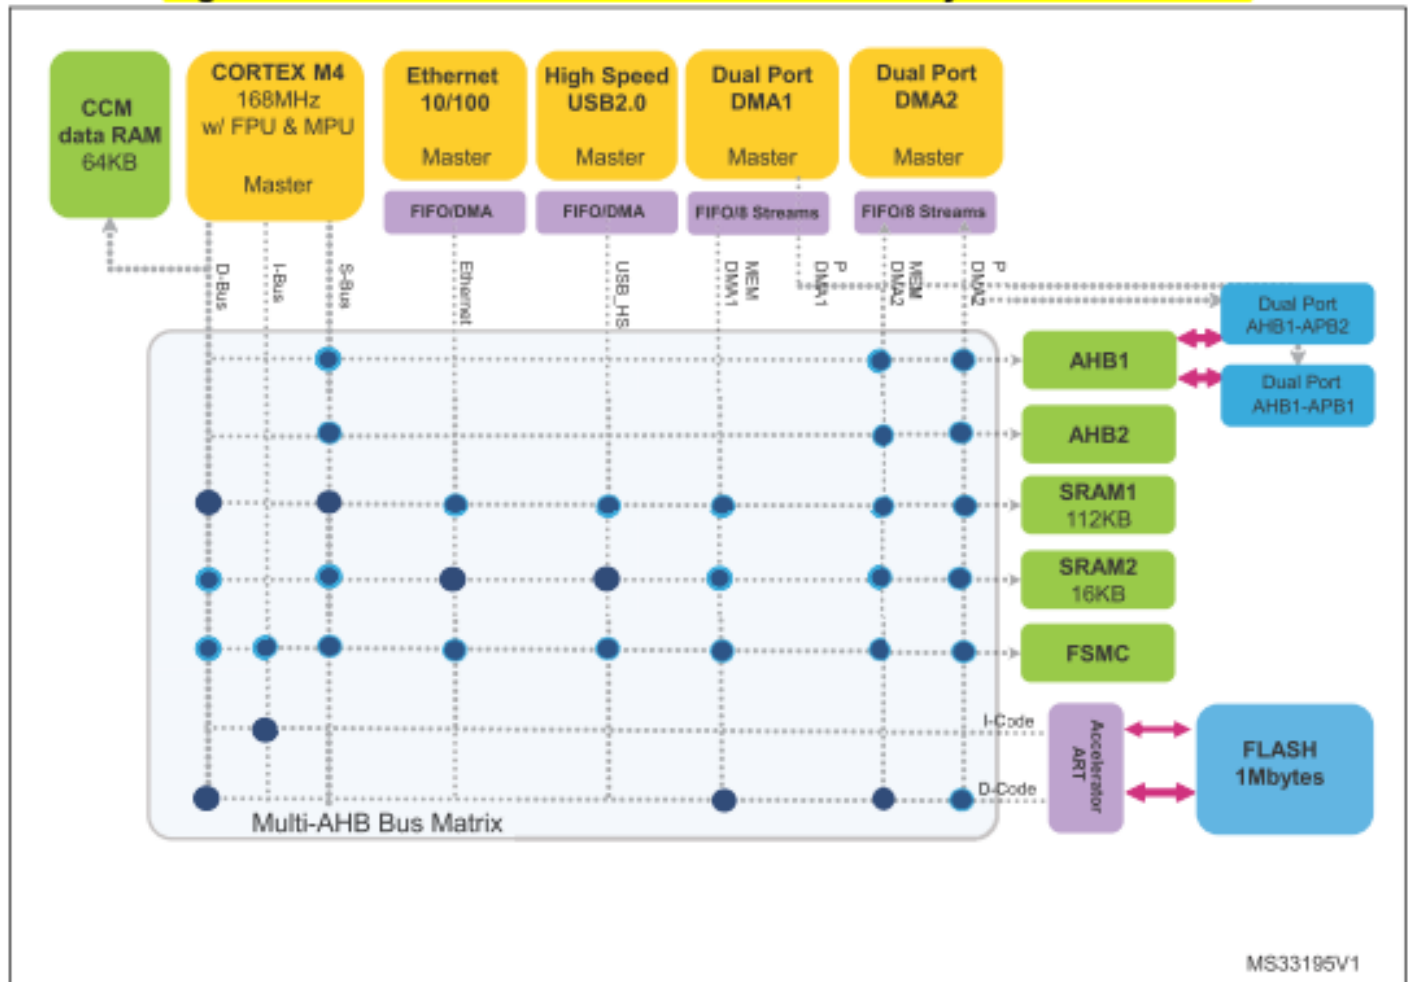
Figure 8-25. (a) An 8 × 8 crossbar switch. (b) An open crosspoint. (c) A closed crosspoint.

Primer stikalne matrike v mikrokrmilnikih ARM Cortex M4

AN4031

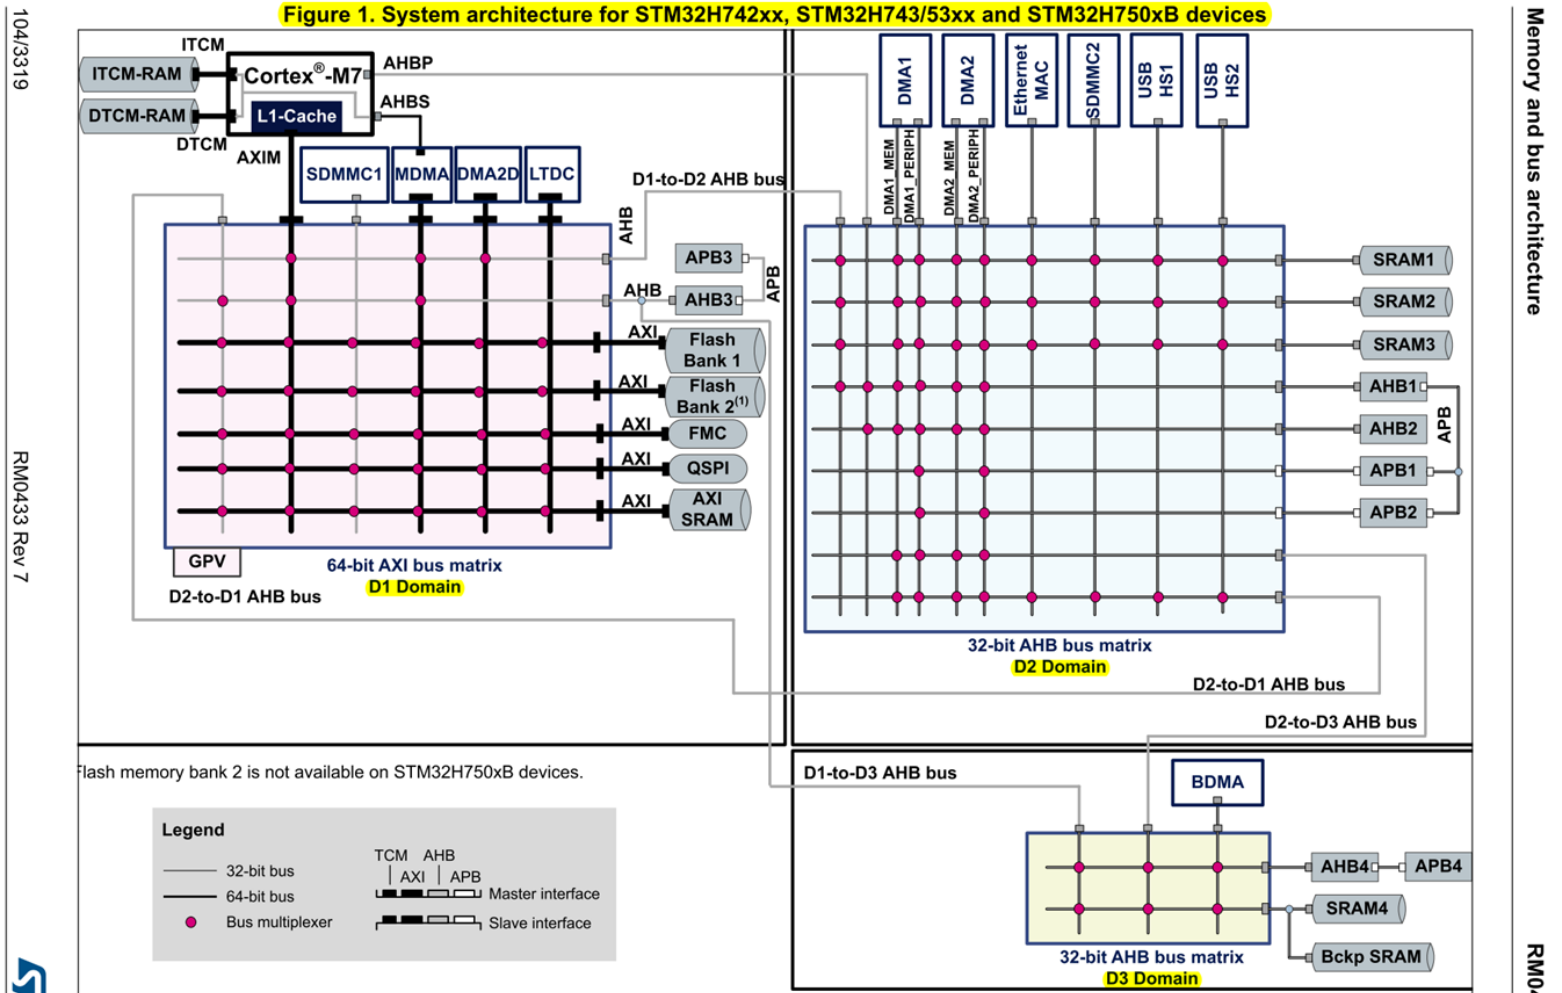
System performance considerations

Figure 7. STM32F405/415 and STM32F407/417 system architecture



OR

Primer stikalne matrike v mikrokrmilnikih ARM Cortex M7 (H7)



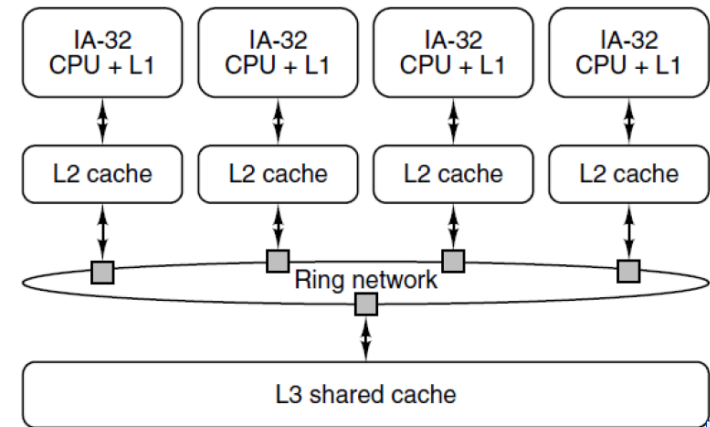
Multiprocesorji s skupnim pomnilnikom - skladnost vsebine predpomnilnikov I

Predpomnilniki :

- „snooping caches“ – spremljajo dogajanje na vodilu in ustrezno reagirajo

Skladnost vsebine predpomnilnikov :

- „Write-through“ – Cache coherence protocol
 - Enostaven,
 - Neučinkovit (vsakič pisanje (tudi) v pomnilnik)


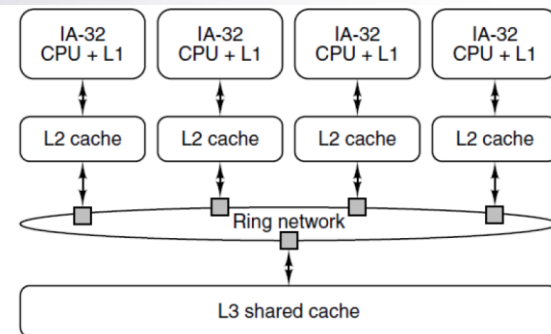


Action	Local request	Remote request
Read miss	Fetch data from memory	
Read hit	Use data from local cache	
Write miss	Update data in memory	
Write hit	Update cache and memory	Invalidate cache entry


Multiprocesorji s skupnim pomnilnikom skladnost vsebine predpomnilnikov II

- „Write-Back“ Cache Coherence Protocol :


- kompleksnejši, kasnejše pisanje v pomn. (uskladitev)
- Primer: „MESI“ – Cache coherence protocol (tudi Core i7):
 - **I**nvalid – ni vrstice
 - **S**hared – več predp. vsebuje aktualno vrstico (skladno s pomn.)
 - **E**xclusive – samo dol. predp. vsebuje aktualno vrstico (skladno s pomn.)
 - **M**odified – samo dol. predp vsebuje aktualno vrstico (neskladno s pomn.)




CPU 1 reads block A




CPU 3 reads block A




CPU 2 reads block A



CPU 2 writes block A



CPU 2 writes block A



CPU 1 writes block A

Multiprocesorji s skupnim pomnilnikom - programiranje

Sinhronizacija je postopek koordinacije delovanja dveh ali več procesov, ki se izvršujejo na različnih procesorjih.

Eden od načinov sinhronizacije je **zакlepanje (lock)** skupnih spremenljivk:

- dovoljuje dostop do podatka samo enemu procesorju istočasno.
- samo en procesor lahko v nekem trenutku doseže zакlepanje spremenljivke, vsi ostali procesorji morajo čakati, dokler ta procesor ne sprostí spremenljivke.

Vsota 64,000 števil na 64 CPE

Vsak procesor ima ID oz. $P_n : 0 \leq P_n \leq 63$

```
sum[Pn] = 0;  
for (i = 1000*Pn;  
     i < 1000*(Pn+1); i = i + 1)  
    sum[Pn] = sum[Pn] + A[i];
```

```
half = 64;  
repeat
```


```
    synch();
```

⇒ if (half%2 != 0 && Pn == 0) **sum[0] = sum[0] + sum[half-1];**

```
    half = half/2; /* dividing line on two sums */
```

⇒ if (Pn < half) **sum[Pn] = sum[Pn] + sum[Pn+half];**


```
until (half == 1);
```



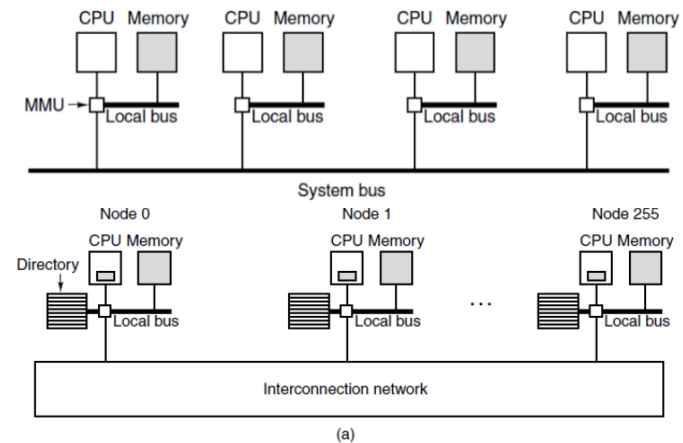
/* Samo, če je
število liho;
Procesor0 prišteje
podatek brez para*/

Multiprocesorji s skupnim pomnilnikom povzetek

- UMA (enoten čas dostopa)
 - + enostavno programiranje
 - *slaba skalabilnost*




- NUMA : (različen čas dostopa)
 - + hitrejši dostopi, *boljša skalabilnost*
 - težje programiranje
 - NC-NUMA (Non Cached NUMA)
 - *vodilo je ozko grlo*
 - + ni problema skladnosti vsebine predpomnilnikov
 - CC-NUMA (Cache Coherent NUMA)
 - + hitrejši dostopi
 - *problem skladnosti vsebine predpomnilnikov*



- COMA: („Cache Only Memory Access“)
 - *naslovni prostor se razbije na vrstice v predpomnilnikih*
 - *lokalni pomnilnik = predpomnilnik*
 - *ni veliko realizacij po tem pristopu*

Ideja: podatki nevezani, se selijo bližje tistemu, ki jih potrebuje



5.2.4.2 Multiračunalniki (MR) s sporočili

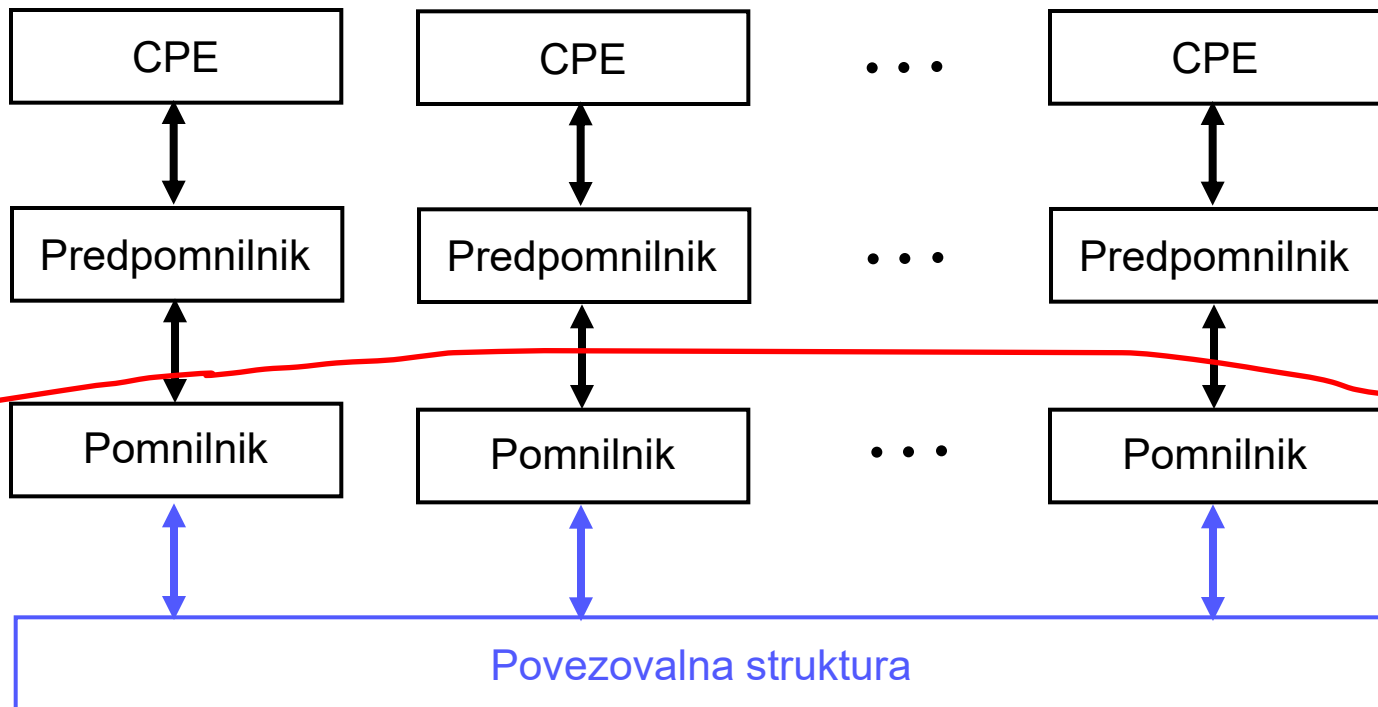
Ideja izhaja iz omejitev multiprocesorjev:

- slabo skaliranje: veliko dodatne logike za npr. 72 CPE (MR lahko enostavneje do 65536!)
- ozko grlo v deljenem (skupnem) pomnilniku: npr. 100 CPE hkrati piše v določeno lokacijo

Rešitev: vsaka CPE ima svoj fizični naslovni prostor:

- komunikacija preko sporočil (ni skupnega pomnilnika)

Tipična organizacija multiprocesorja z več lastnimi pomnilniki (rahla povezanost)



Multiračunalniki (MR) s sporočili

Značilnosti:

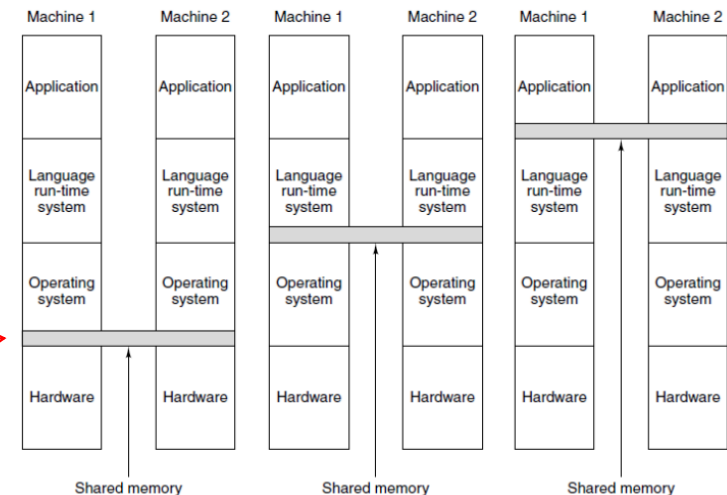
- vsaka CPE ima **svoj pomnilnik**
- komunikacija preko posebne infrastrukture – **povezovalne mreže**
 - zakomplicira zadevo, programi težje komunicirajo, prenosi trajajo
- **se lažje gradijo, bolje skalirajo, a težje programirajo**
- **zakaj** so sploh zanimivi?
 - so lahko bistveno **VEČJI** in **CENEJŠI** !!!

Sistem mora vsebovati:

- podprograme za pošiljanje in sprejemanje sporočil,
- prenos sporočil :
 - vgrajena tudi koordinacija
 - sistem potrjevanja prejetih sporočil

Ideja skupnega pomnilnika ni izključena

Multicomputers



Povezovalne topografije

Povezovalne strukture – parametri :

- »fanout« : št. povezav iz vozlišča
- »diameter« : št. povezav med najbolj oddaljenimi
- »bisection BW« : kapaciteta prereza
- »dimensionality« : št. različnih poti med točkama




Figure 8-4. Various topologies. The heavy dots represent switches. The CPUs and memories are not shown. (a) A star. (b) A complete interconnect. (c) A tree. (d) A ring. (e) A grid. (f) A double torus. (g) A cube. (h) A 4D hypercube.



Vsota 64,000 števil na 64 CPE

■ Operacije za pošiljanje: *send()* in *receive()*

```
limit = 64; half = 64; /* 64 processors */  
repeat
```

```
    half = (half+1)/2; /* send vs.receive  
                        dividing line */
```

```
    if (Pn >= half && Pn < limit)
```

```
        ⇒ send(Pn - half, sum);
```

```
    if (Pn < (limit/2))
```

```
        ⇒ sum = sum + receive();
```

```
    limit = half; /* upper limit of senders */  
until (half == 1); /* exit with final sum */
```


Multiračunalniki (MR) s sporočili

Visoko zmogljivi multiprocesorski sistemi zahtevajo tudi **zelo zmogljive povezovalne** strukture, kar pa pomeni tudi **visoko ceno**:

- obstaja **malo aplikacij**, ki bi opravičevale tako velike stroške,
- nižje cenovna alternativa: **gruče računalnikov**, pri katerih poteka komunikacija med procesorji preko standardnih V/I omrežij.

Vsled opisanih razlik ločimo multiračunalnike v 2 večji skupini glede na njihovo ceno (ki je pogojena z izbiro komponent), zgradba je podobna v obeh:

- MPP
 - Zmogljivost pred ceno!
- COW, WSC (gruče)
 - „Clusters Of Workstations“
 - gruče
 - „Warehouse Scale Computers“
 - spletne storitve



5.2.4.2.1 MPP („Massively Parallel Processors“)

Ideja:

- narediti superračunalnik s pomočjo:
 - velikega števila **zmogljivih vozlišč** in
 - zelo **zmogljivo povezovalno mrežo**....
- cena ni ovira (sicer gruča - COW)...

Značilnosti:

- **najprej** se uporabijo v **znanosti**, danes vse bolj že v **komercialne namene („Big Data“)**
- so večmilj. superračunalniki – nadomestili so vektorske SIMD, itd...
- večinoma se uporabijo **standardni CPUji** (Pentium, UltraSPARC, PowerPC)
- zelo **zmogljive in hitre povezave**
- **lasten SW**, knjižnice
- vedno prisoten HW/SW, ki **preverja napake** in jih odpravlja...

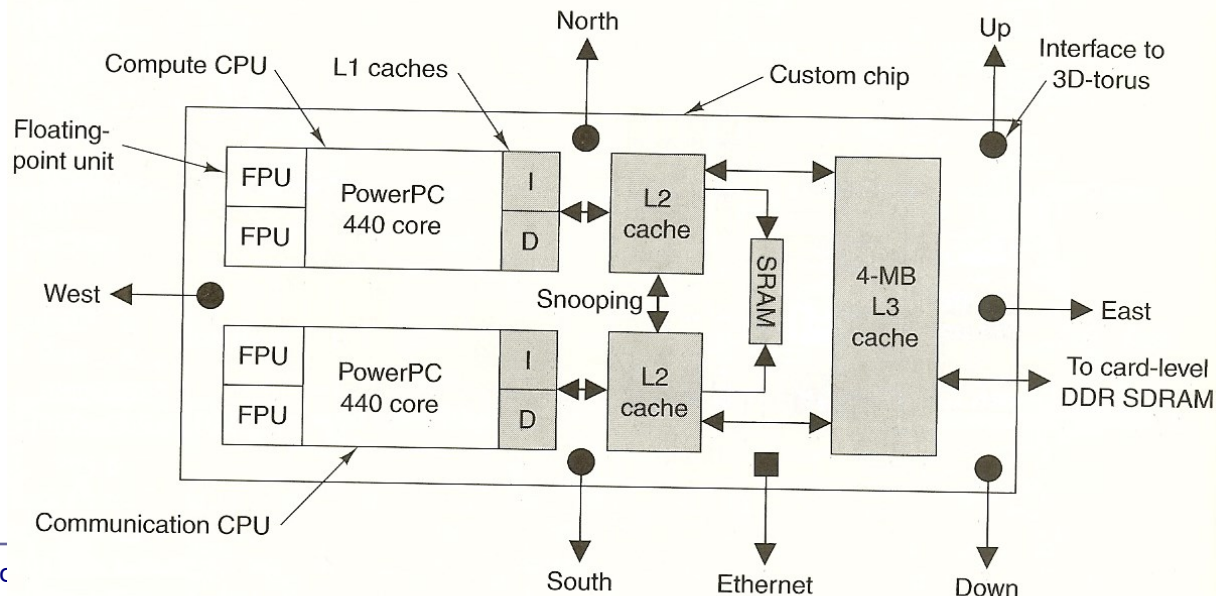
Primer:

- **IBM Blue Gene/L** je l.2005 zmagal za »Gordon Bell« nagrado
- teoretična zmoglj. **360 000 GFLOPS**
- zmagovalni program izkoristi **28%**
- pri nekaterih je izkoristek **1% !!!**
- **#1 na top500.org v letih 2004 do 2008**

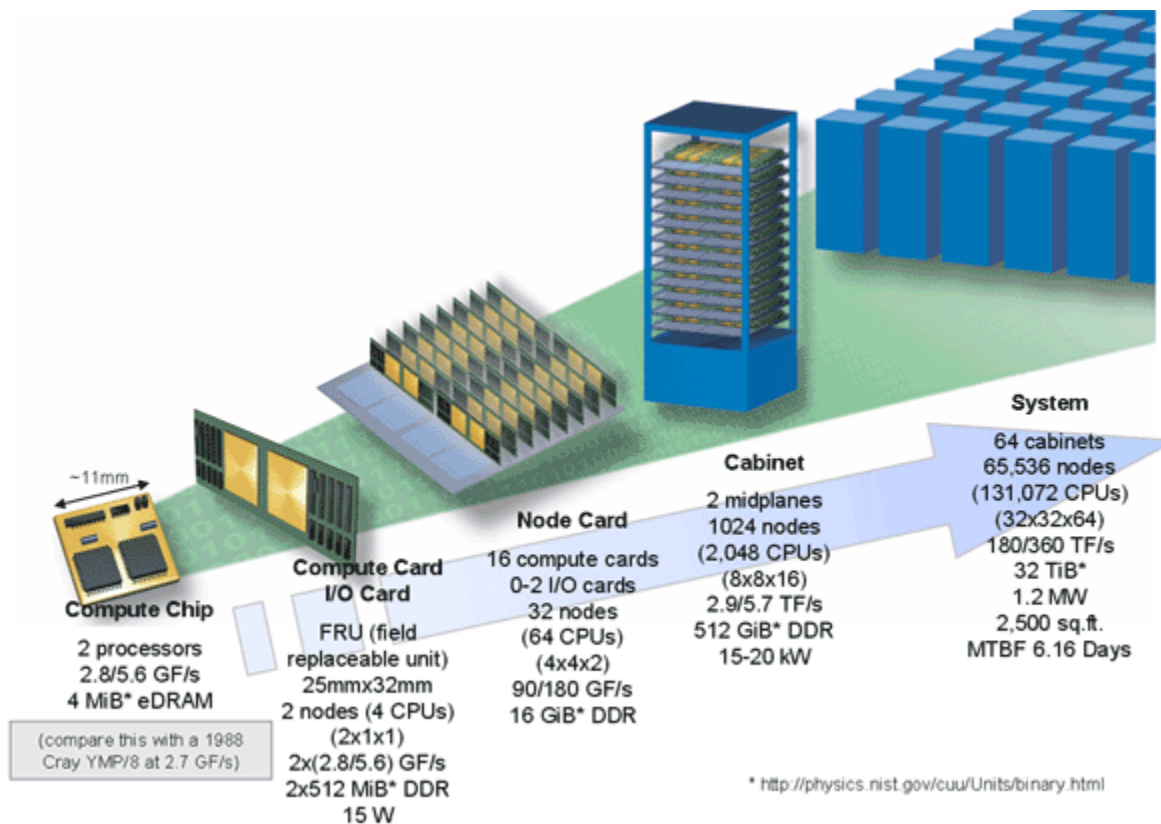
IBM Blue Gene/L

Osnova računalnika Blue Gene/L je **ASIC** («Application Specific Integrated Circuit») procesorski čip z dvema jedri PowerPC:

- **PowerPC** je cevovodni 2-kratni (dvo-izstavitveni) superskalarni procesor
- **FPU funkcijski enoti** sta dvovhodni - lahko izvršita štiri operacije v plavajoči vejici v 1 t_{CPE}
- Procesor ni med najzmogljivejšimi.
 - eno jedro je namenjeno računanju, drugo pa komunikaciji med 65.536 procesorji.
- **Povezave med CPUji** so:
 - 3D torus in še 4 druge povezovalne mreže različnih tipov
- Enormna **V/I zmogljivost** (1.4Gbps na P2P povezavo, oz. 275Tbit/s=900 000 knjig/s)
- Cilji: optimalna razmerja: TFLOPS/USD, TFLOPS/W, TFLOPS/m³



IBM Blue Gene



Procesorski
 ASIC čip


Kartica

Plošča


Omara

Sistem

BLUEGENE/L COMPUTE CARD



BLUE GENE/ L NODE BOARD

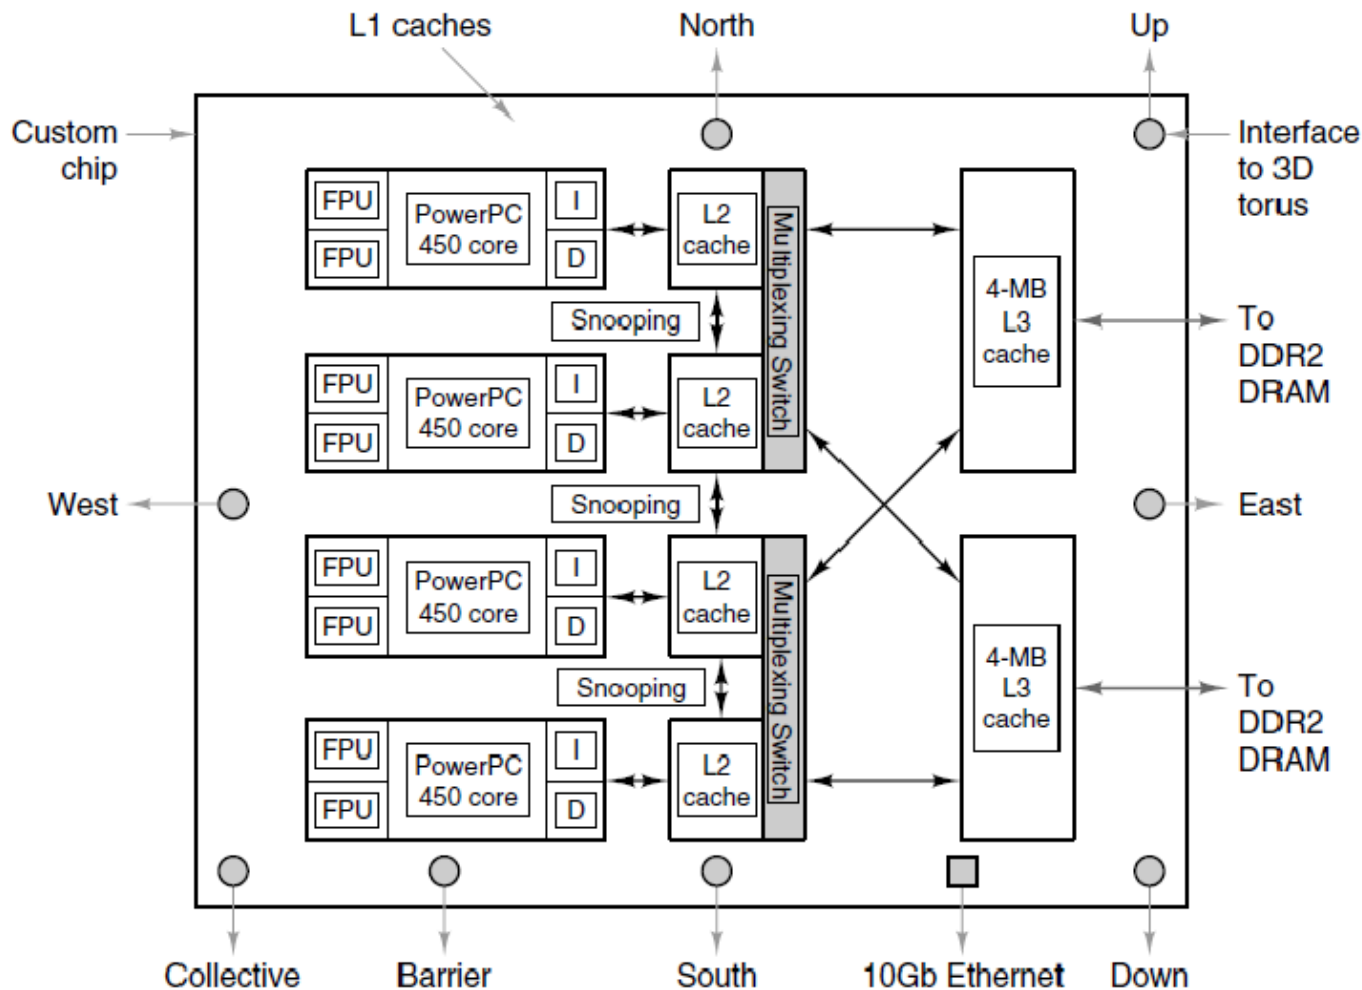


Paralelni računalnik IBM Blue Gene



Paralelni računalnik IBM Blue Gene/P

- Od I. 2007 naprej
- 2009: prebije mejo PFLOP/sec



Paralelni računalnik IBM Blue Gene/Q

<http://www-03.ibm.com/systems/technicalcomputing/solutions/bluegene/>




■ #3 na top500.org v letu 2014 (#3 v 2013, #1 v 2012):

- 96 omar x 1024 nodes x 16 jeder = 98304 nodes x 16 jeder =
1 572 864 jeder
- vozlišče: 16 jedrni PowerPC A2 + 16GB DDR3 RAM
- povezava: 5D torus
- ≈ 16 PetaFlops/s
- 7x boljši zmoglj./Watt od serije P, 3x boljši zmoglj./Watt od takrat #1 na top500 !!!
- zavzema 280m², porabi **7.9 MW** (HE Medvode ~ **25MW**)

Razširljiv do 512 omar.

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>



Rapid growth of supercomputer performance, based on data from the top500.org website. The logarithmic y-axis shows performance in GFLOPS.

Sum... Combined performance of 500 largest supercomputers

Top... Fastest supercomputer

#500... Supercomputer in 500th place

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov. 2012, Nov. 2013, Nov. 2014, Nov. 2015

Rank	Rmax Rpeak (Pflops)	Name	Computer design Processor type, interconnect	Vendor	Site Country, year	Operating system
1 ()	33.863 54.902	<u>Tianhe-2</u>	<u>NUDT</u> <u>Xeon E5-2692 +</u> <u>Xeon Phi 31S1P,</u> <u>TH Express-2</u>	<u>NUDT</u>	<u>National</u> <u>Supercomputing</u> <u>Center in Guangzhou</u> <u>China, 2013</u>	<u>Linux</u> (<u>Kylin</u>)
2 (1)	17.590 27.113	<u>Titan</u>	<u>Cray XK7</u> <u>16 core AMD</u> <u>Opteron CPU +</u> <u>Nvidia K20 GPU,</u> <u>Custom</u>	<u>Cray</u>	<u>Oak Ridge National</u> <u>Laboratory (ORNL) in</u> <u>Tennessee</u> <u>United States, 2012</u>	Cray Linux Env (<u>SuSE</u> based)
3 (2)	16.325 20.133	<u>Sequoia</u>	<u>Blue Gene/Q</u> <u>PowerPC A2,</u> <u>Custom</u>	<u>IBM</u>	<u>Lawrence Livermore</u> <u>National Laboratory</u> <u>United States, 2011</u>	<u>Linux</u> (<u>RHEL</u> and <u>CNK</u>)
4 (3)	10.510 11.280	<u>K computer</u>	<u>RIKEN</u> <u>SPARC64 VIIIfx,</u> <u>Tofu</u>	<u>Fujitsu</u>	<u>RIKEN</u> <u>Japan, 2011</u>	<u>Linux</u>
5 (4)	8.162 10.066	<u>Mira</u>	<u>Blue Gene/Q</u> <u>PowerPC A2,</u> <u>Custom</u>	<u>IBM</u>	<u>Argonne National</u> <u>Laboratory</u> <u>United States, 2012</u>	<u>Linux</u> (<u>RHEL</u> and <u>CNK</u>)

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov. 2015, **Nov. 2016**, Nov 2017

Rank	Rmax Rpeak (Pflops)	Name	Computer design Processor type, interconnect	Vendor	Site Country, year	Operating system
1 ()	93.015 125.436	<u>Sunway</u> <u>TaihuLight</u>	<u>NUDT</u> <u>Xeon E5-2692 +</u> <u>Xeon Phi 31S1P, TH</u> Express-2	<u>NRCPC</u>	<u>National</u> <u>Supercomputing</u> <u>Center in Wuxi</u> <u>China, 2016</u>	<u>Linux</u> (Raise)
2 (1)	33.863 54.902	<u>Tianhe-2</u>	<u>NUDT</u> <u>Xeon E5-2692 +</u> <u>Xeon Phi 31S1P, TH</u> Express-2	<u>NUDT</u>	<u>National</u> <u>Supercomputing</u> <u>Center in</u> <u>Guangzhou</u> <u>China, 2013</u>	<u>Linux</u> (Kylin)
3 (2)	17.590 27.113	<u>Titan</u>	<u>Cray XK7</u> <u>16 core AMD Opteron</u> <u>CPU + Nvidia K20</u> <u>GPU, Custom</u>	<u>Cray</u>	<u>Oak Ridge National</u> <u>Laboratory (ORNL) in</u> <u>Tennessee</u> <u>United States, 2012</u>	Cray Linux Env (<u>SuSE</u> based)
4 (3)	16.325 20.133	<u>Sequoia</u>	<u>Blue Gene/Q</u> <u>PowerPC A2, Custom</u>	<u>IBM</u>	<u>Lawrence Livermore</u> <u>National Laboratory</u> <u>United States, 2011</u>	<u>Linux</u> (<u>RHEL</u> and <u>CNK</u>)
5 (-)	14.015 27.881	<u>Cori</u>	<u>Cray XC40</u> <u>Xeon Phi 7250, Aries</u>	<u>Cray</u>	<u>National Energy</u> <u>Research Scientific</u> <u>Computing Center</u> <u>United States, 2016</u>	<u>Linux</u> (<u>CLE</u>)

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov. 2015, Nov. 2016, Nov 2017, Nov 2018 in 2019

Rank	Rmax Rpeak (PFL OPS)	Name	Model	Processor	Interconnect	Vendor
1	143.500 200.795	Summit	Power System AC922	POWER9, Tesla V100	Infiniband EDR	IBM
2	94.640 125.436	Sierra	Power System S922LC	POWER9, Tesla V100	Infiniband EDR	IBM
3	93.015 125.436	Sunway TaihuLight	Sunway MPP	SW26010	Sunway ^[20]	NRCPC
4	61.445 100.679	Tianhe-2A	TH-IVB-FEP	Xeon E5-2692 v2, Matri x-2000	TH Express-2	NUDT
5	21.230 27.154	Piz Daint	Cray XC50	Xeon E5-2690 v3, Tesla P100	Aries	Cray

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500> Nov 2017, Nov 2018 in 2019, Nov 2020

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
5	Selene - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63,460.0	79,215.0	2,646

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov 2017, Nov 2018 in 2019, Nov 2020, Nov 2021

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
5	Perlmutter - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70,870.0	93,750.0	2,589

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov 2021, **Nov 2022**

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112	1,102.00	1,685.65	21,100
2	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	Leonardo - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,463,616	174.70	255.75	5,610
5	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096
8	Perlmutter - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70.87	93.75	2,589

Lestvica Top 500

<http://en.wikipedia.org/wiki/TOP500>

Nov 2022, **Nov 2023**

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,699,904	1,194.00	1,679.82	22,703
2	Aurora - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, Intel DOE/SC/Argonne National Laboratory United States	4,742,808	585.34	1,059.33	24,687
3	Eagle - Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR, Microsoft Microsoft Azure United States	1,123,200	561.20	846.84	
4	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
5	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,752,704	379.70	531.51	7,107

Vega : SLO superračunalnik

List	Rank	System	Vendor	Total Cores	Rmax (TFlops)	Rpeak (TFlops)	Power (kW)
11/2021	157	BullSequana XH2000, AMD EPYC 7H12 64C 2.6GHz, NVIDIA A100, Infiniband HDR	Atos	33,600	3,096.0	4,680.0	
06/2021	134	BullSequana XH2000, AMD EPYC 7H12 64C 2.6GHz, NVIDIA A100, Infiniband HDR	Atos	33,600	3,096.0	4,680.0	

Uvrstitev : 2021

O Vegi

Superračunalniki EuroHPC JU:

1. MareNostrum5
Barcelona Supercomputing Centre
2. LUMI, 375 PFLOPS,
CSC, IT Center for Science (Finska)
3. Leonardo, 250 PFLOPS
CINECA (Italija)
4. EURO_IT4I, 15,2 PFLOPS
IT4Innovations (Češka)
5. MeluXina, 10 PFLOPS
LuxProvide (Luxemburg)
6. Deucalion, 5+5 (ARM) PFLOPS
Minho Advance Computing Centre (Portugalska)
7. PetaSC, 4,44 PFLOPS
Sofia Tech Park



8. HPC Vega: osnovne informacije

Cena 17,2 MEUR. Poraba 1 MW.
Delovna zmogljivost 6,8 PFLOPS.
Več kot 1000 vozlišč, preko 120.000 jeder.
Okoli 20 PB diskovnih kapacitet.

<https://eurohpc-ju.europa.eu/discover-eurohpc#ecl-inpage-211>

VEGA HPC GPU - BullSequana XH2000, AMD EPYC 7H12 64C 2.6GHz...

VEGA HPC GPU - BULLSEQUANA XH2000, AMD EPYC 7H12 64C 2.6GHZ, NVIDIA A100, INFINIBAND HDR

Site: [IZUM](#)

Manufacturer: Atos

Cores: 33,600

Memory: 30,720 GB

Processor: AMD EPYC 7H12 64C 2.6GHz

Interconnect: Infiniband HDR

Performance

Linpack Performance (Rmax) 3,096 TFlop/s

Theoretical Peak (Rpeak) 4,680 TFlop/s

Nmax 1,152,000

HPCG [TFlop/s] 77.547

Software

Operating System: Linux

Vega, Maister : SLO superračunalnika

HPC VEGA

Prilagodljiv, hibridni in večnamenski
produksijski HPC sistem

Vrednost investicije

21.653.444,85 EUR

Skupna računska zmogljivost: **6,8 PFLOP/s.**

Računska particija: **320 troprocesorskih računskih rezin s skupno 122.880 jedri.**

Vektorska (GPU) particija: **60 dvoprocorskih računskih vozlišč s po 4 dodanimi GPU enotami s skupno preko 1,6 milijona jedri.**

Diskovno polje s **4 PB hitrega (SSD) in 18 PB trajnega (HDD) shranjevalnega prostora.**

Vsa vozlišča in diskovni sistemi povezani na hitro lokalno omrežje **Ethernet** in nizko latenčno omrežje **Infiniband**.

Sistem bo povezan v **ARNES** in **GEANT** omrežje.

OPEN SOURCE

Odpri tokodna sistemska programska oprema

Operacijski sistem: **Linux CentOS**

Avtomatizacija postopkov: **Foreman, Puppet, Ansible**

Shranjevanje, dostop, prenos in obdelava podatkov:
CephFS, Ceph RDB, ObjectStore (Webdav, XrootD, gsiftp), dCache, Rucio, iRODS

Baze podatkov **InfluxDB, MariaDB, PostgreSQL**

Zabojniki/Vsebniki (container) – podpora aplikacij

HPC MAISTER

Prototipni HPC

Vrednost investicije

2.764.022,14 EUR

76 dvoprocorskih računskih vozlišč s skupno 4.256 jedri.

6 dvoprocorskih računskih vozlišč s po 4 dodanimi grafičnimi procesnimi enotami s skupno 122.952 jedri.

Skupno **40 TB delovnega pomnilnika in 158 TB hitrega shranjevalnega prostora na vozliščih.**

Hitro diskovno polje (SSD) kapacitete 138 TB.

Diskovno polje za trajno hranjenje podatkov (HDD) kapacitete 2,88 PB.

Vsa vozlišča in diskovna polja so povezani na hitro lokalno omrežje **Ethernet** in **Infiniband** omrežje nizko latenco.

Posodobljena vsa komunikacijska oprema za dostop do omrežja **ARNES**.

Tehnologija hlajenja s tekočino

PDU + Power Management Controller

Up to 6 x 15kW PSU shelves
DLC & hot-pluggable

Up to 32 compute blades
Technologies provided by:

- Intel
- Nvidia
- ARM
- AMD


DLC & hot-pluggable

Up to 3 Hydraulic chassis for Direct Liquid Cooling

- Inlet water temperature up to 40°C
- NH1 redundancy



42U cabinet front view



42U cabinet rear view

Up to 2 Management network switches
hot-pluggable

Up to 10 Interconnect Network Switches
Technologies:


- ESI
- Mellanox HDR
- High-speed Ethernet DLC & hot-pluggable

Flexible mid plane

Allows for:

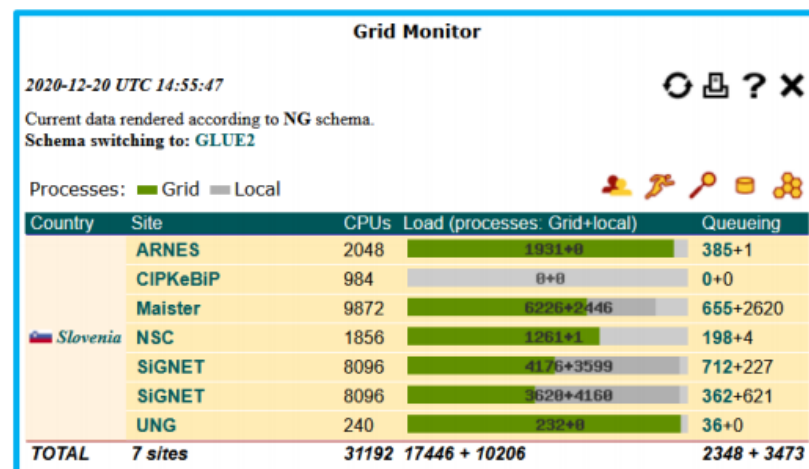
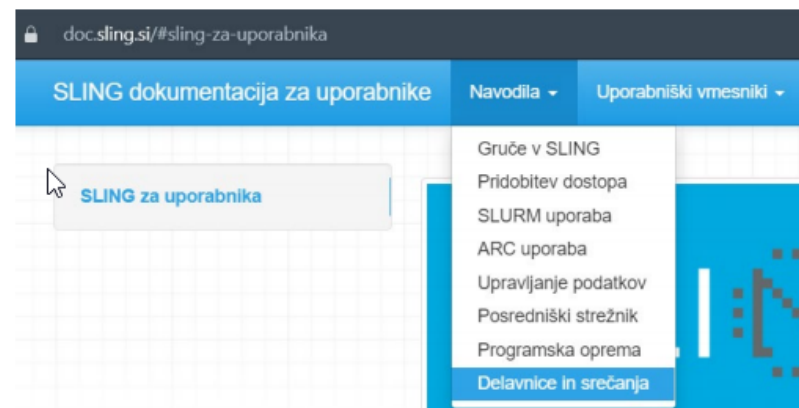
- routing customization
- mixed interconnect bandwidths
- two possible topologies (Full Fat Tree and DragonFly+)

SLING | Slovenska iniciativa za nacionalni grid



SLING za skupnost

- <https://www.sling.si> uradna stran
- <https://signet-ca.ijs.si> digitalna potrdila
- <https://doc.sling.si> navodila za uporabnike
- <https://fido.sling.si> Centos Identity Management
- <http://www.sling.si/gridmonitor/loadmon.php> ARC monitor
- <https://voms.sling.si:8443/voms>
Virtual Organization Membership Service
- <https://wiki.sling.si> za administratorje
- <https://repo.sling.si> GitLab za administratorje
- <https://mapa.sling.si> dokumenti na NextCloud
- Splošne informacije: info@sling.si
- Podpora za slovenske uporabnike: support@sling.si



SLING – Slovenska iniciativa za nacionalni grid

Delavnice

- Osnove superračunalništva
- Superračunalništvo bolj zares
- Računajmo na vsa jedra
- Programiranje grafičnih procesnih enot (OpenCL)

- [O delavnici](#)

- Uvod

- [Heterogeni sistemi](#)

- [Anatomija GPE](#)

- [Grafične procesne enote](#)

- [Nvidia Tesla K40](#)

- [Zvalami modeli](#)

- [Pomnilniška hierarhija](#)

- Uvod v OpenCL

- [OpenCL](#)

- [Ustvarjanje OpenCL okolja na gostitelju](#)

- Programiranje GPE

- [Seštevanje vektorjev](#)

- [Skalarni produkt vektorjev](#)

- [Množenje matrik](#)

- [Obdelava slik](#)

- Programiranje grafičnih procesnih enot (CUDA)

- [O delavnici](#)

- Uvod

- Grafične procesne enote

- [CUDA](#)

- Programiranje GPE

Dan slovenskega superračunalniškega omrežja

13 October 2022

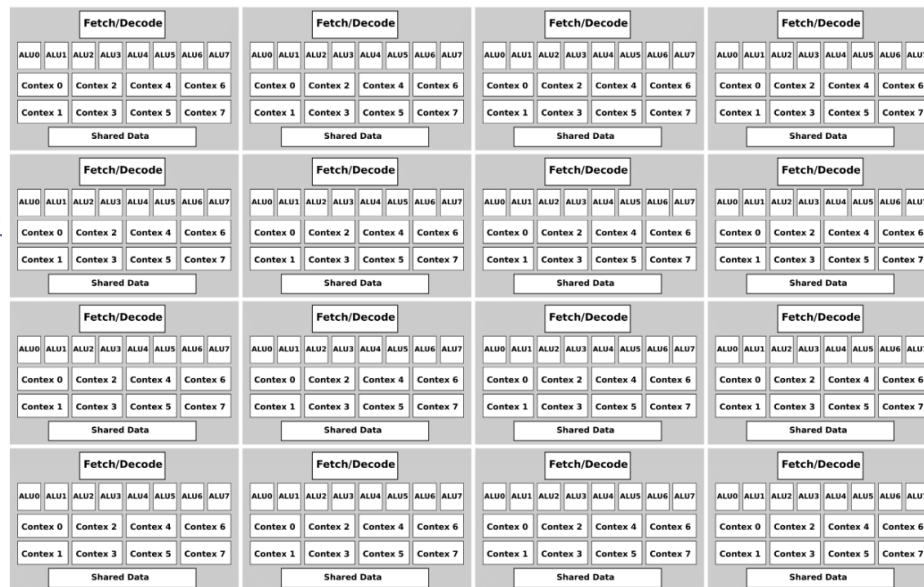
Grafična procesna enota

<https://indico.ijs.si/event/1533/>

Gremo še en korak naprej. Namesto ene računske enote v sistemu uporabimo kar 16 računskih enot, kot prikazuje spodnja slika.

```
add r3,r0,#tid
slli r3,r3,#2
lfp f1,r3(vecA)
lfp f2,r3(vecB)
addf f1,f1,f2
sfp f1,r3(vecC)
```

Each GPU core (CU) executes this code



Sedaj nam ni treba zanke ponoviti 16 krat temveč vsaki računski enoti dodelimo le eno iteracijo zanke. Zgornja slika prikazuje poenostavljeno zgradbo grafičnih procesnih enot.

<https://doc.sling.si/workshops/programming-gpu/intro/anatomija/>


SLING – Slovenska iniciativa za nacionalni grid

Gruče superračunalniškega omrežja SLING Prosto dostopne gruče

Nekatere od gruč v omrežju SLING podpirajo odprti dostop za upravičence SLING:

- Univerza v Mariboru - gruča Maister (projekt [HPC RIVR](#))
 - [opis gruče Maister](#)
 - vstopno vozlišče: **rmaister.hpc-rivr.um.si**
 - navodila za uporabo
- NSC - skupna gruča IJS, omogoča odprti dostop
 - opis gruče [NSC.ijs.si](#)
 - tehnični opis in navodila za dostop in uporabo
 - vstopno vozlišče: **nsc-login.ijs.si**
- ARNES - gruča Arnes
 - [opis gruče Arnes](#)
- Fakulteta za informacijske študije Novo mesto - gruči Rudolf in Trdina
 - [opis gruče Trdina](#)
 - vstopno vozlišče: **trdina-login.fis.unm.si**
- IZUM - Institut informacijskih znanosti - gruča Vega
 - [opis gruče Vega](#)
 - vstopno vozlišče: **login.vega.izum.si**

From: <https://doc.sling.si/navodila/clusters/>



Energy efficiency of top-ranked computers (gigaflops/watt)

The Green500 List


Listed below are the June 2015 The Green500's energy-efficient supercomputers ranked from 1 to 10.

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	7,031.58	RIKEN	Shoubu - ExaScaler-1.4 80Brick, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband FDR, PEZY-SC	50.32
2	6,842.31	High Energy Accelerator Research Organization /KEK	Suiren Blue - ExaScaler-1.4 16Brick, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband, PEZY-SC	28.25
3	6,217.04	High Energy Accelerator Research Organization /KEK	Suiren - ExaScaler 32U256SC Cluster, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, PEZY-SC	32.59
4	5,271.81	GSI Helmholtz Center	ASUS ESC4000 FDR/G2S, Intel Xeon E5-2690v2 10C 3GHz, Infiniband FDR, AMD FirePro S9150	57.15
5	4,257.88	GSIC Center, Tokyo Institute of Technology	TSUBAME-KFC - LX 1U-4GPU/104Re-1G Cluster, Intel Xeon E5-2620v2 6C 2.100GHz, Infiniband FDR, NVIDIA K20x	39.83
6	4,112.11	Stanford Research Computing Center	XStream - Cray CS-Storm, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR, Nvidia K80	190.00
7	3,962.73	Cray Inc.	Storm1 - Cray CS-Storm, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, Nvidia K40m	44.54
8	3,631.70	Cambridge University	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	52.62
9	3,614.71	TU Dresden, ZIH	Taurus GPUs - Bull bullx R400, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Nvidia K80	58.01
10	3,543.32	Financial Institution	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	54.60

* Performance data obtained from publicly available sources including TOP500

Lestvica Green Top 500 – I. 2016

<https://www.top500.org/lists/green500/>



Green500 Rank	TOP500 Rank	MFLUPS/W	Site	System	Total Power(kW)
1	28	9462.1	NVIDIA Corporation	NVIDIA DGX-1, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla P100	349.5
2	8	7453.5	Swiss National Supercomputing Centre (CSCS)	Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100	1312
3	116	6673.8	Advanced Center for Computing and Communication, RIKEN	ZettaScaler-1.6, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband FDR, PEZY-SCnp	150.0
4	1	6051.3	National Supercomputing Center in Wuxi	Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway	15371
5	375	5806.3	Fujitsu Technology Solutions GmbH	PRIMERGY CX1640 M1, Intel Xeon Phi 7210 64C 1.3GHz, Intel Omni-Path	77

Lestvica Green Top 500 – I. 2017

<https://www.top500.org/lists/green500/>



Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	259	Shoubu system B - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 , PEZY Computing / Exascaler Inc. Advanced Center for Computing and Communication, RIKEN Japan	794,400	842.0	50	17.009
2	307	Suiren2 - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 , PEZY Computing / Exascaler Inc. High Energy Accelerator Research Organization /KEK Japan	762,624	788.2	47	16.759
3	276	Sakura - ZettaScaler-2.2, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband EDR, PEZY-SC2 , PEZY Computing / Exascaler Inc. PEZY Computing K.K. Japan	794,400	824.7	50	16.657
4	149	DGX SaturnV Volta - NVIDIA DGX-1 Volta36, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla V100 , Nvidia NVIDIA Corporation United States	22,440	1,070.0	97	15.113
5	4	Gyokou - ZettaScaler-2.2 HPC system, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 700Mhz , ExaScaler Japan Agency for Marine-Earth Science and Technology Japan	19,860,000	19,135.8	1,350	14.173

Lestvica Green Top 500 – I. nov/2018

<https://www.top500.org/lists/green500/>



Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	375	Shoubu system B - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 , PEZY Computing / Exascaler Inc. Advanced Center for Computing and Communication, RIKEN Japan	953,280	1,063.3	60	17.604
2	374	DGX SaturnV Volta - NVIDIA DGX-1 Volta36, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla V100 , Nvidia NVIDIA Corporation United States	22,440	1,070.0	97	15.113
3	1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,397,824	143,500.0	9,783	14.668
4	7	AI Bridging Cloud Infrastructure (ABCI) - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , Fujitsu National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	1,649	14.423
5	22	TSUBAME3.0 - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2 , HPE GSIC Center, Tokyo Institute of Technology Japan	135,828	8,125.0	792	13.704

Lestvica Green Top 500 – I. nov/2019

<https://www.top500.org/lists/green500/>




TOP500				Rmax	Power	Power
Rank	Rank	System	Cores	(TFlop/s)	(kW)	Efficiency (GFlops/watts)
1	159	A64FX prototype - Fujitsu A64FX, Fujitsu A64FX 48C 2GHz, Tofu interconnect D , Fujitsu Lestvica Green Top 500 – I. 2016... Fujitsu Numazu Plant Japan	36,864	1,999.5	118	16.876
2	420	NA-1 - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 700Mhz , PEZY Computing / Exascaler Inc. PEZY Computing K.K. Japan	1,271,040	1,303.2	80	16.256
3	24	AiMOS - IBM Power System AC922, IBM POWER9 20C 3.45GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Volta GV100 , IBM Rensselaer Polytechnic Institute Center for Computational Innovations (CCI) United States	130,000	8,045.0	510	15.771
4	373	Satori - IBM Power System AC922, IBM POWER9 20C 2.4GHz, Infiniband EDR, NVIDIA Tesla V100 SXM2 , IBM MIT/MGHPCC Holyoke, MA United States	23,040	1,464.0	94	15.574
5	1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	10,096	14.719



Lestvica Green Top 500 – I. nov/2020

<https://www.top500.org/lists/green500/>

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	170	NVIDIA DGX SuperPOD - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	19,840	2,356.0	90	26.195
2	330	MN-3 - MN-Core Server, Xeon Platinum 8260M 24C 2.4GHz, Preferred Networks MN-Core, MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	1,664	1,652.9	65	26.039
3	7	JUWELS Booster Module - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite, Atos Forschungszentrum Juelich (FZJ) Germany	449,280	44,120.0	1,764	25.008
4	146	Spartan2 - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR Infiniband, Atos Atos France	23,040	2,566.0	106	24.262
5	5	Selene - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63,460.0	2,646	23.983



Lestvica Green Top 500 – I. nov/2021

<https://www.top500.org/lists/green500/>



Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	301	MN-3 - MN-Core Server, Xeon Platinum 8260M 24C 2.4GHz, Preferred Networks MN-Core, MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	1,664	2,181.2	55	39.379
2	291	SSC-21 Scalable Module - Apollo 6500 Gen10 plus, AMD EPYC 7543 32C 2.8GHz, NVIDIA A100 80GB, Infiniband HDR200, HPE Samsung Electronics South Korea	16,704	2,274.1	103	33.983
3	295	Tethys - NVIDIA DGX A100 Liquid Cooled Prototype, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 80GB, Infiniband HDR, Nvidia NVIDIA Corporation United States	19,840	2,255.0	72	31.538
4	280	Wilkes-3 - PowerEdge XE8545, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 80GB, Infiniband HDR200 dual rail, DELL EMC University of Cambridge United Kingdom	26,880	2,287.0	74	30.797
5	30	HiPerGator AI - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Infiniband HDR, Nvidia University of Florida United States	138,880	17,200.0	583	29.521



Lestvica Green Top 500 – I. nov/2022


<https://www.top500.org/lists/green500/>



Rank	TOP500 Rank	System	Cores	Rmax (PFlop/s)	Power (kW)	Energy Efficiency (GFlops/watts)
	405	Henri - Lenovo ThinkSystem SR670 V2, Intel Xeon Platinum 8362 2800Mhz (32C), NVIDIA H100 80GB PCIe, Infiniband HDR, Lenovo Flatiron Institute United States	5,920	2.04	31	65.091
2	32	Frontier TDS - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	120,832	19.20	309	62.684
3	11	Adastra - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE Grand Equipement National de Calcul Intensif - Centre Informatique National de l'Enseignement Suprieur (GENCI-CINES) France	319,072	46.10	921	58.021
4	15	Setonix - GPU - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE Pawsey Supercomputing Centre, Kensington, Western Australia Australia	181,248	27.16	477	56.983
5	68	Dardel GPU - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE KTH - Royal Institute of Technology Sweden	52,864	8.26	146	56.491

Lestvica Green Top 500 – I. nov/2023(nov/2022)

<https://www.top500.org/lists/green500/>



Rank	TOP500 Rank	System	Cores	Rmax (PFlop/s)	Power (kW)	Energy Efficiency (GFlops/watts)
1	293	Henri - ThinkSystem SR670 V2, Intel Xeon Platinum 8362 32C 2.8GHz, NVIDIA H100 80GB PCIe, Infiniband HDR, Lenovo Flatiron Institute United States	8,288	2.88	44	65.396
2	44	Frontier TDS - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	120,832	19.20	309	62.684
3	17	Adastra - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE Grand Equipement National de Calcul Intensif - Centre Informatique National de l'Enseignement Suprieur (GENCI-CINES) France	319,072	46.10	921	58.021
4	25	Setonix - GPU - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE Pawsey Supercomputing Centre, Kensington, Western Australia Australia	181,248	27.16	477	56.983
5	92	Dardel GPU - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE KTH - Royal Institute of Technology Sweden	52,864	8.26	146	56.491

5.2.4.2.2 COW („Clusters of Workstations“)

Ideja - V bistvu so po sami zgradbi podobni MPP sistemom, vendar imajo vse rešitve:

- običajne, množično uporabljane komponente
- in so zato bistveno cenejši

Prednosti :

- poceni, zanesljivi
- energetsko učinkoviti
- enostavno razširljivi

Pomanjkljivosti:

- strošek upravljanja z gručo n - računalnikov je zelo podoben kot strošek upravljanja z n - neodvisnimi računalniki
- manj zmogljive V/I povezave (še posebej v primerjavi z multiproc. sistemi).
- več kopij operacijskega sistema

2 tipa gruč:

- centraliziran
 - homogeni računalniki v omari
- decentraliziran
 - povezava bolj oddaljenih, samostojnih sistemov

Google „Data Center“ (gruča Googlovih računalnikov)



Warehouse-Scale Computers (WSC)

Razsežnost: pomnijo, indeksirajo, :


- l. 2008: **10¹² strani** (angleško trillion, slovensko bilijarda)
- l. 2013: 30 bilijard strani (30x več/5 let)

Kako ? :

- razpršeni podatkovni centri (»Data center - DC«)
- »www.google.com« se razrešuje na najbližjem DC

Zakaj zanimivo ?


- ?
- ?
- superračun./standardni ?
- cena,zmogljivost, oboje



(vir: <http://www.internetlivestats.com/>)

Google „Data Center“ (gruča Googlovih računalnikov)

- <http://www.google.com/about/datacenters/>
 - 2014: model izboljšanja energetske učinkovitosti s strojnim učenjem
 - www.google.com/about/datacenters/efficiency/internal/assets/machine-learning-applications-for-datacenter-optimization-finalv2.pdf



$$\text{PUE} = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

<https://www.google.com/about/datacenters/efficiency/internal/>

- [Google Data Center 360° Tour \(2016\)](#)
- [Google Data Center 360° Tour \(2020\)](#)

Gruča Googlovih računalnikov

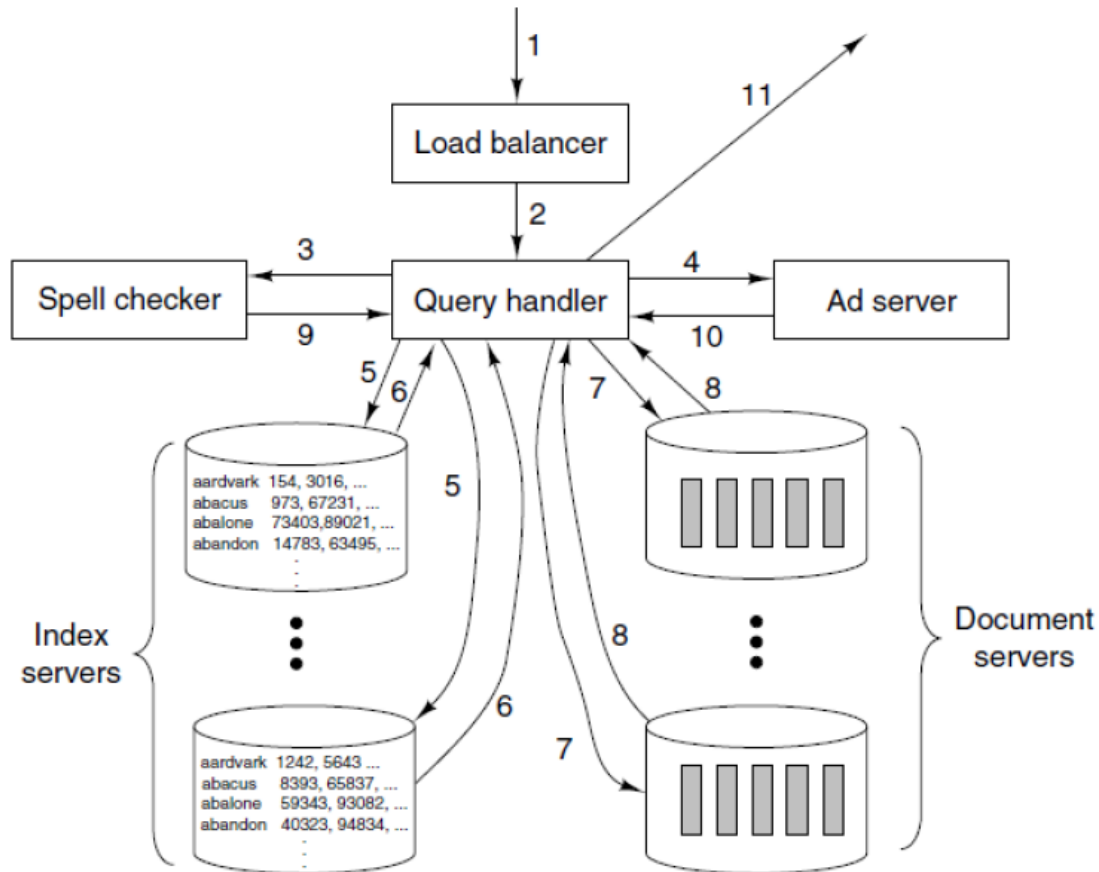
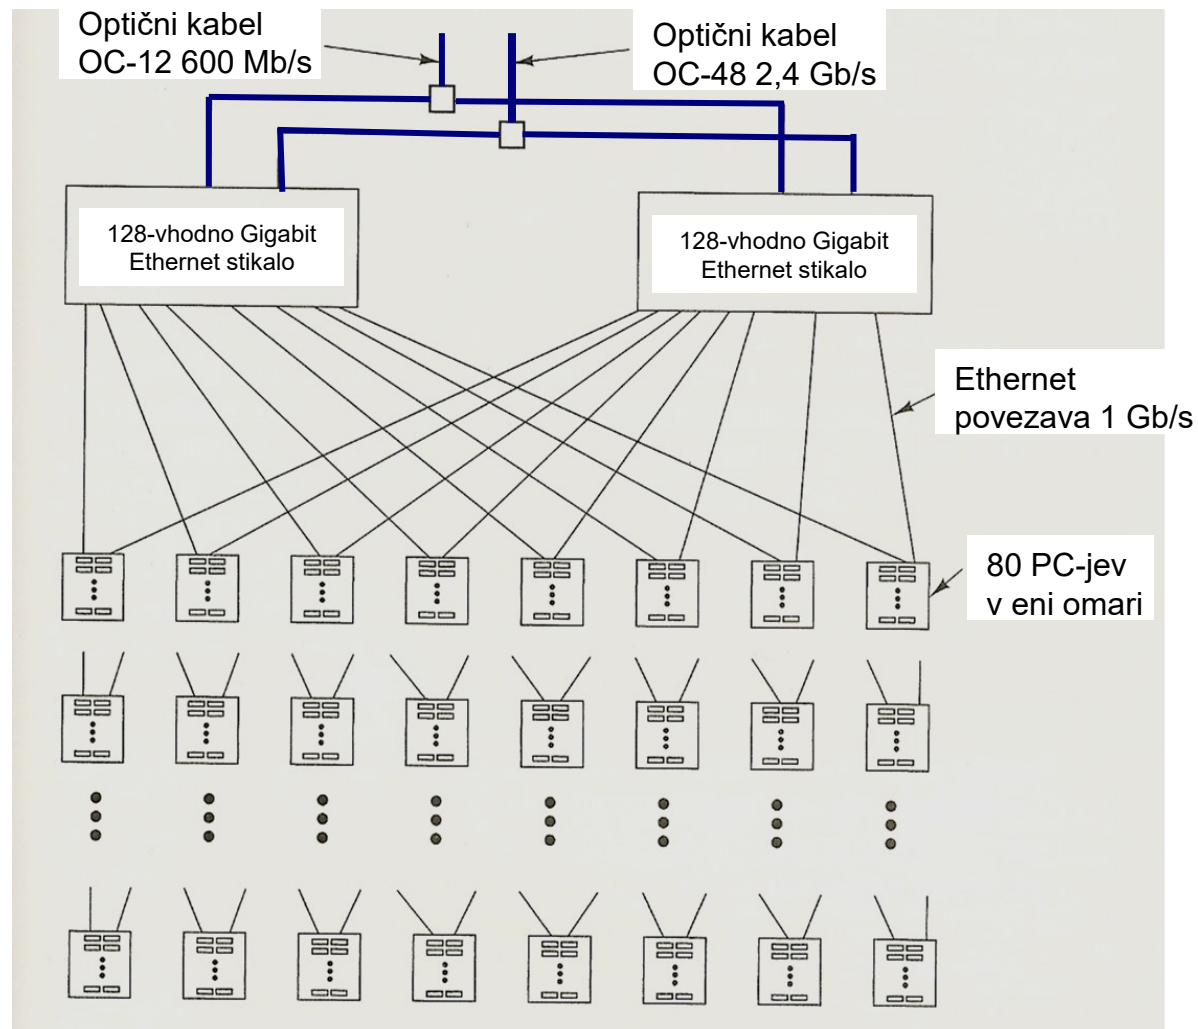


Figure 8-43. Processing of a Google query.

Gruča Googlovih računalnikov


Zgradba DC :

- PC-ji so v **ohišju 1U** (1U rack višina 5cm)
- po 80 jih je v omari (40 spredaj, 40 zadaj).
- PC-ji v omari so povezani preko Ethernet stikala v omari.
- omare so med seboj povezane z dvema redundantnima Gigabit Ethernet stikaloma s po 128 vhodi.
- maksimalno število omar je 64
 - (po dve povezavi iz vsake omare na vsako 128 vhodno stikalo), oziroma **64*80=5120** PC-jev.
- poraba je približno 10kW na omaro (80 * 120W)
- omara zasede približno 3m² (servisiranje, hlajenje), kar pomeni 3000 W/m² (posebno hlajenje)




Past Design

- In-house rack design
- PC-class motherboards
- Low-end storage and networking hardware
- Linux
- + in-house software



OR – 5 – Povzetki predavanj



Container Datacenter



Container Datacenter



Google „Data Center“

(gruča Googlovih računalnikov)

Praktične izkušnje :

- Fault tolerant SW
- Raje redundanca kot najbolj zmožljiva (in najdražja) tehnologija

Vsako leto se pokvari pribl. 2% računalnikov (našteto po pogostosti) :

- največji izvor težav : Programska oprema - SW
- 1/2 okvar diski
- napajalniki
- pomnilniki
- CPE skoraj nikoli

Gruča Googlovih računalnikov – danes: COW -> WSC



5.2.4.3 Izvajanje programov v multiračunalnikih (MR)

Medsebojna komunikacija s sporočili:

- sinhronska predaja sporočil (»synchronous message passing«)
 - pošiljatelj blokira, dokler sprejemnik ne prevzame sporočila

- posredna predaja sporočil (»buffered message passing«)
 - se sporočila hranijo v izravnalnikih

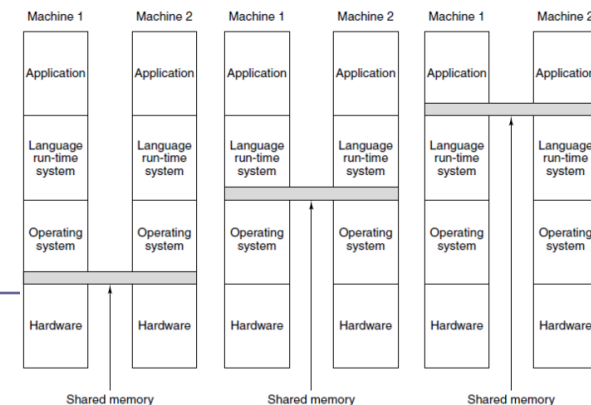
- zakasnjena predaja sporočil (»non-blocking message passing«)
 - se sporočila hranijo v izhodnih izravnalnikih
 - proces se nadaljuje takoj brez ustavljanja

5.2.4.3 Izvajanje programov v multiračunalnikih

Najpogostejša načina realizacija medsebojne komunikacije :

- MPI (Message Passing Interface): MPI-1, MPI-2
 - MPI_Send (buffer, count, data_type, destination, tag, proc_group)
 - $\text{buffer} = \text{count} * \text{data_type} \rightarrow \text{destination}$ z oznako „tag“ v „proc_group“
 - MPI_Recv (&buffer, count, data_type, destination, tag, communicator, &status)
- DSM (Distributed Shared Memory)
 - Ideja: „navideznega“ skupnega pomnilnika (Ikažje programiranje) :
 - skupina procesorjev si deli skupni navidezni pomnilnik
 - realizacija na nivoju programske opreme

Multicomputers (2)



Primer programa z uporabo MPI

```
// Find out rank (process ID), size
```

```
int world_rank;
```

```
MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);
```

```
int number;
```

```
if (world_rank == 0) {  
    number = -1;
```


```
MPI_Send(&number, 1, MPI_INT, 1, 0, MPI_COMM_WORLD);
```

```
} else if (world_rank == 1) {
```

```
MPI_Recv(&number, 1, MPI_INT, 0, 0, MPI_COMM_WORLD, MPI_STATUS_IGNORE);
```

```
printf("Process 1 received number %d from process 0\n", number);
```

```
}
```



5.2.4.3 Izvajanje programov v multiračunalnikih

Dodeljevanje procesorjev in opravil:

■ FIFO :

- zaporedno izvajanje iz vrste (proces pove potrebno št. procesorjev)

■ »no head-of-line blocking«:

- prednost tisti, ki ustrezajo po št. procesorjev naprej

■ »tiling«:

- zahteva podatke o št. procesorjev in času (boljša izkoriščenost)

Primer z 8 procesorji

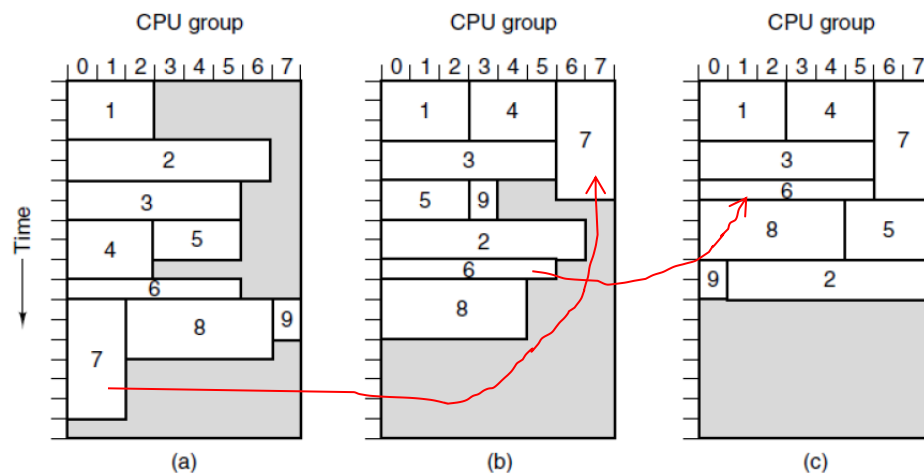


Figure 8-45. Scheduling a cluster. (a) FIFO. (b) Without head-of-line blocking. (c) Tiling. The shaded areas indicate idle CPUs.

5.2.4.4 GPU – Nvidia CUDA (Compute Unified Device Architecture)

Vzporedno z razvojem splošno namenskih procesorjev so se razvijali tudi sistemi na grafičnih karticah:

- **nekdaj:** toga, specializirana vezja (»VGA controller«) pa
- **do današnjih** sodobnih multiprocesorjev z ekstenzivnim paralelizmom.

Nosilec razvoja je bila industrija računalniških iger:

- povzroči **celo hitrejši razvoj od splošnih CPE:**
 - **kompatibilnost je vezana le na API nivo:**
 - več svobode eksperimentiranja, uvajanja sprememb in novih tehnologij
 - bistveno manjše breme kompatibilnosti za nazaj

- **Izpolnjujejo prostor podatkovne paralelnosti**

	Statično določanje (ob prevajanju)	Dinamično določanje (med delovanjem)
Paralelizem – ukazi	VLIW	Superskalarni rač.
Paralelizem - podatki	SIMD, vektorski rač.	GPU

2022: RTX4090 1600USD
• **73 TFLOPS**, 16384 jeder, **450W**

2020: RTX3090 1500USD
• **35 TFLOPS**, 10496 jeder, **350W**

2015: GTX960 cca. 250 EUR:
• **2.3 TFLOPS**, 1024 jeder, **120W**
2013: GTX660 cca. 200 EUR:
• **1.8 TFLOPS**, 960 jeder, **140W**
1990: Cray-2 cca 30Milj. USD:
• **0.002 TFLOPS**, **150kW**,
(najhitrejši takrat, sedaj 1000x počasnejši)

5.2.4.4 GPU – Nvidia CUDA (Compute Unified Device Architecture)

Glavne razlike v arhitekturi GPU vs. CPU:

- GPU: ni toliko poudarka na predpomnilnikih:
 - **Veliko št. niti, se zastoje „maskira“**
- GPU uporabljajo **ekstenzivni paralelizem**
 - Dinamični podatkovni paralelizem (veliko niti)
 - Mnogo paralelnih, enostavnejših „procesorjev“ (npr. FP enot)
- GPU pomnilnik :
 - **Pomembnejša širina poti** od same hitrosti (veliko niti „kompenzira“ hitrost)
- GPU: omejena podpora za double prec.(DP) FP
 - DP-FP: Vsebolj prisotna , a do 10x počasnejša
 - Npr. : **GeForce RTX 4090:**
 - **SP 73.1 TFLOPS**
 - **DP 1.14 TFLOPS**

GPU – primer GeForce 8800 GTX (Tesla)

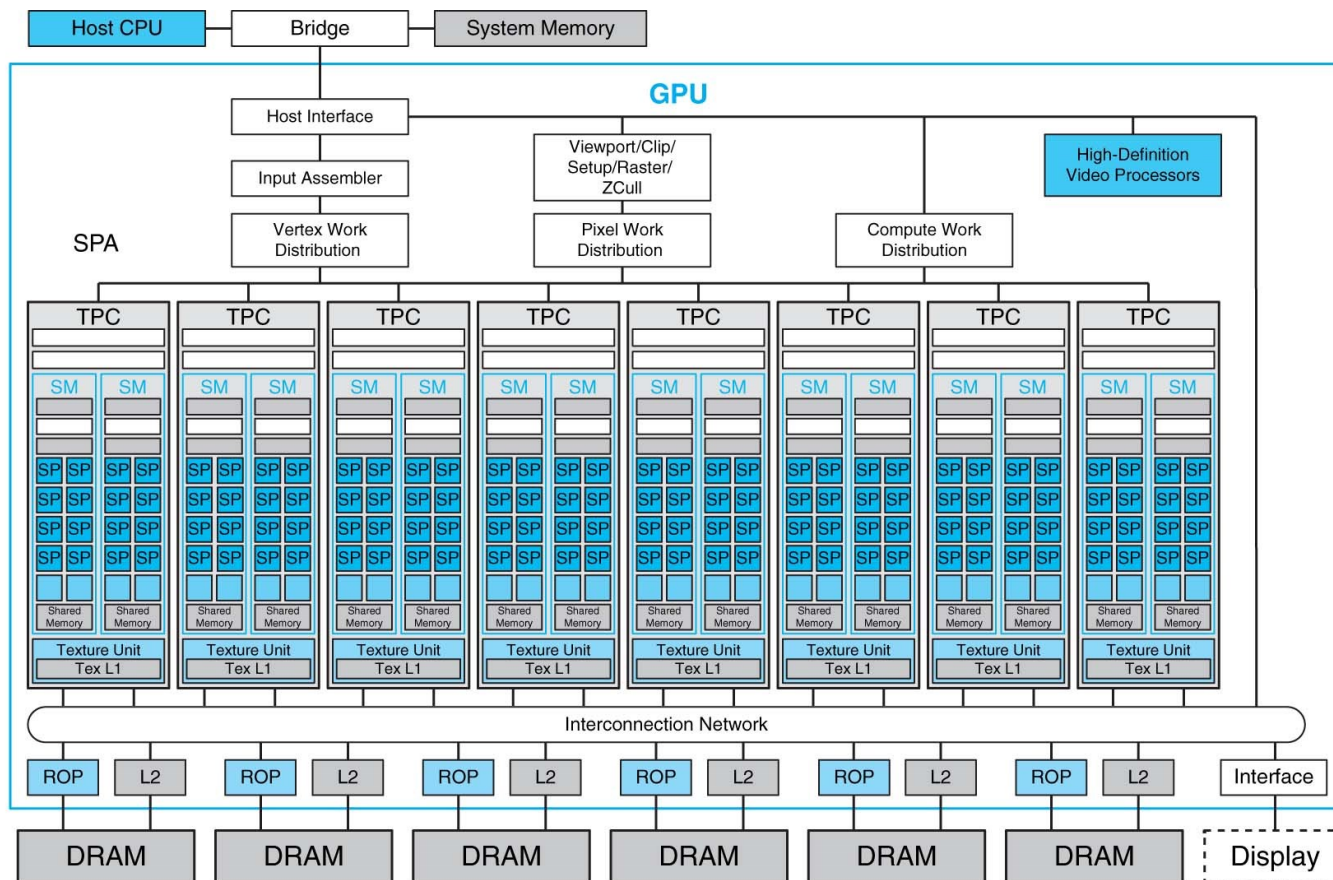


FIGURE A.7.1 NVIDIA Tesla unified graphics and computing GPU architecture. This GeForce 8800 has **128 streaming processor (SP) cores in 16 streaming multiprocessors (SM)**, arranged in eight **texture/processor clusters (TPC)**. The processors connect with six 64-bit-wide DRAM partitions via an interconnection network. Other GPUs implementing the Tesla architecture vary the number of SP cores, SMs, DRAM partitions, and other units. Copyright © 2009 Elsevier, Inc. All rights reserved.

GPU – primer GeForce 8800 GTX (Tesla)




FIGURE A.7.2 Texture/processor cluster (TPC) and a streaming multiprocessor (SM). Each SM has eight streaming processor (SP) cores, two SFUs, and a shared memory. Copyright © 2009 Elsevier, Inc. All rights reserved.

GPU – primer GeForce 8800 GTX (Tesla) - MP

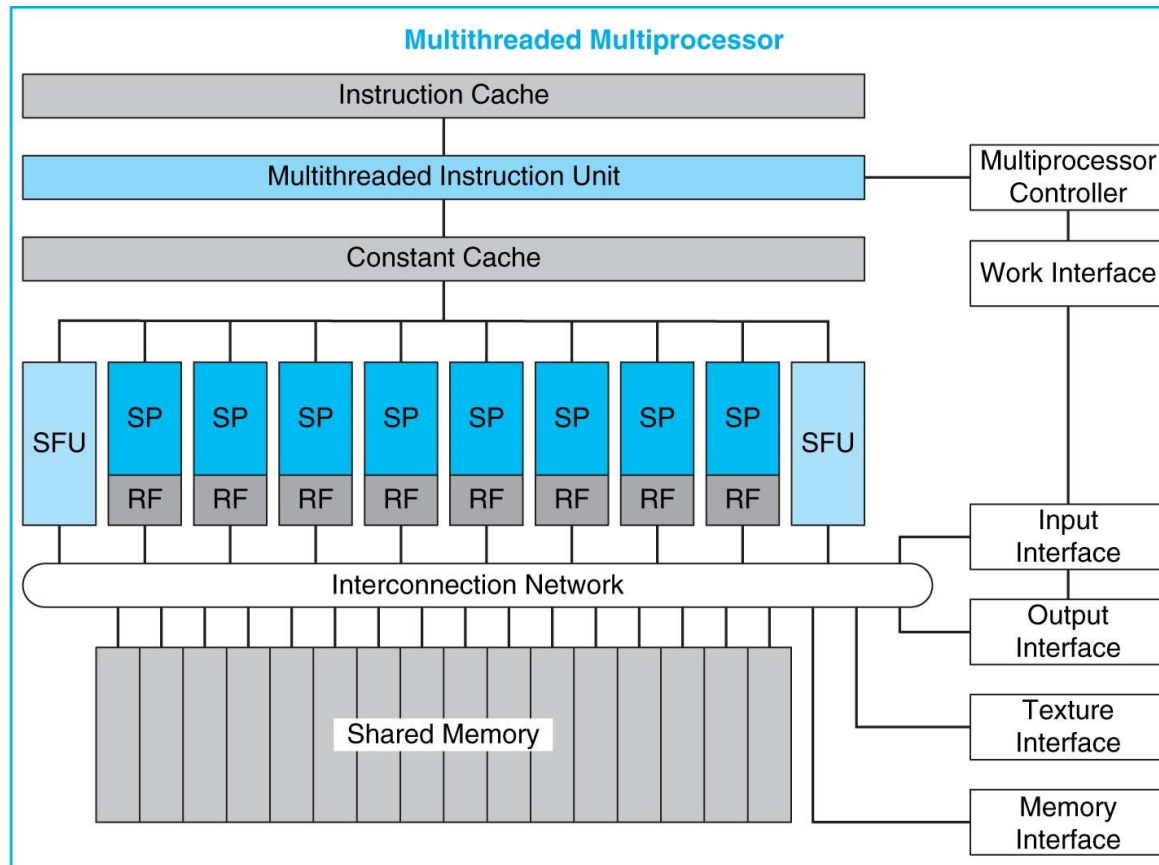


FIGURE A.4.1 Multithreaded multiprocessor with eight scalar processor (SP) cores. The eight SP cores each have a large multithreaded register file (RF) and share an instruction cache, multithreaded instruction issue unit, constant cache, two special function units (SFUs), interconnection network, and a multibank shared memory. Copyright © 2009 Elsevier, Inc. All rights reserved.

GPU – primer GeForce 8800 GTX (Tesla)


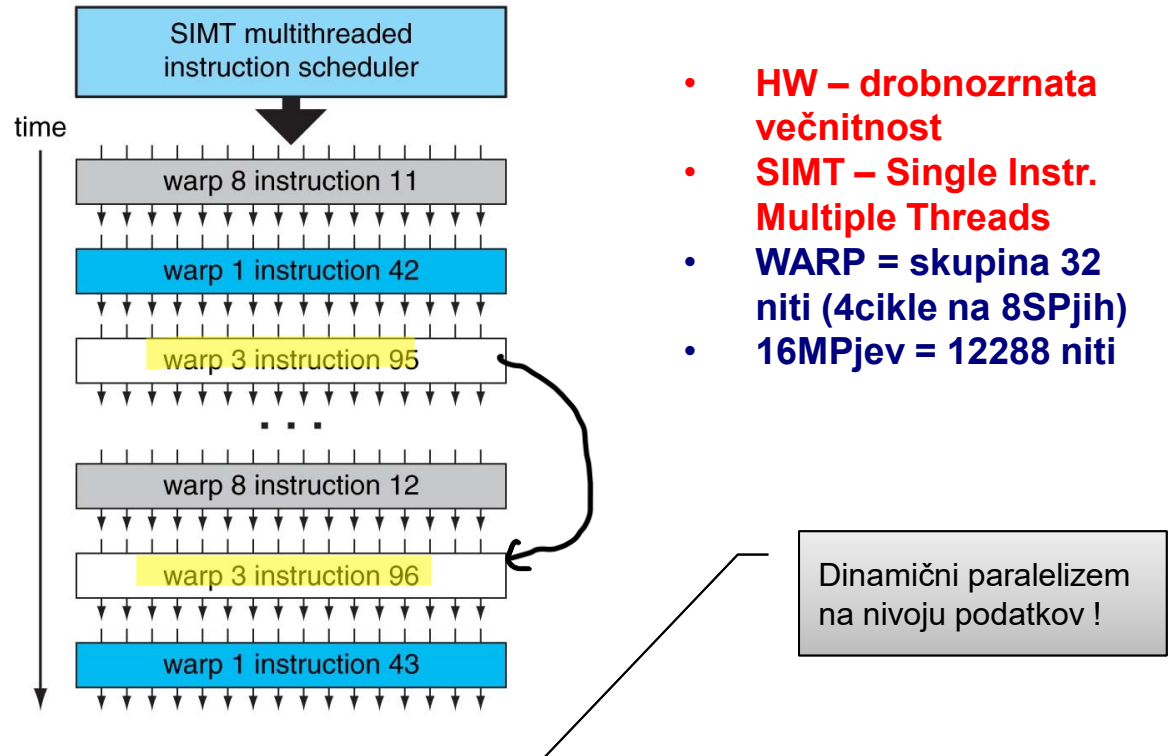


Photo: Judy Schoonmaker

Izvajanje niti (Warp-ov)



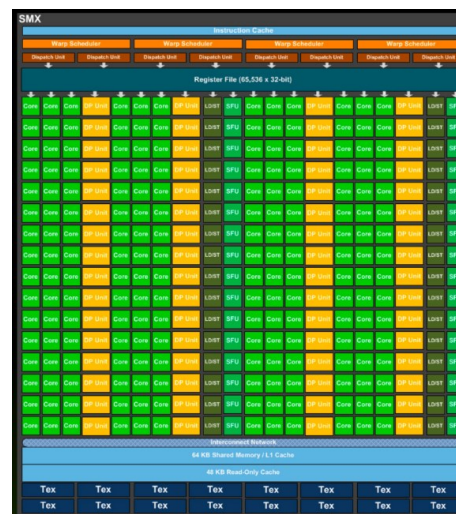
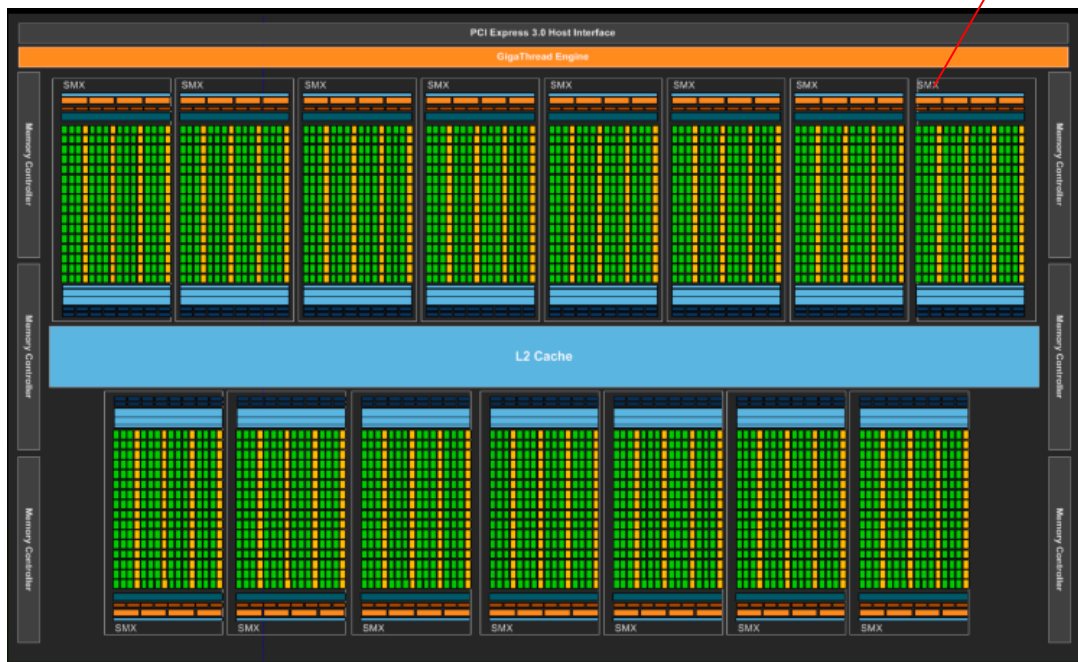
- **HW – drobnozrnata večnitnost**
- **SIMT – Single Instr. Multiple Threads**
- **WARP = skupina 32 niti (4cikle na 8SPjih)**
- **16MPjev = 12288 niti**

FIGURE A.4.2 SIMT multithreaded warp scheduling. The scheduler selects a ready warp and issues an instruction synchronously to the parallel threads composing the warp. Because warps are independent, the scheduler may select a different warp each time. Copyright © 2009 Elsevier, Inc. All rights reserved.

GPU – primer Tesla K40 (tudi v SLING)

Vsebuje 15 računskih enot (CU) :

- ime „NeXt Generation Streaming Multiprocessor“ (SMX)
 - Vsaka 128 procesnih enot (jedra ali „kernel“ ali „Stream Processor“
- Skupaj: $15 * 128 = 1920$ procesnih enot



- Programiranje grafičnih procesnih enot
- O delavnici
- Uvod
- Heterogeni sistemi
- Anatomija GPE
- Grafične procesne enote
- Nvidia Tesla K40
- Izvajalni model
- Pomnilniška hierarhija
- Uvod v OpenCL
- OpenCL
- Ustvarjanje OpenCL okolja na gostitelju
- Programiranje GPE
- Seštevanje vektorjev
- Skalarni produkt vektorjev
- Množenje matrik
- Obdelava slik

<https://doc.sling.si/workshops/programming-gpu/GPE/teslak40/>
<https://doc.sling.si/workshops/programming-gpu/intro/course/>

SISD – primer programa

```
void add( int *a, int *b, int *c ) {  
    int tid = 0;    // this is CPU zero, so we start at zero  
    while (tid < N) {  
        c[tid] = a[tid] + b[tid];  
        tid += 1;    // we have one CPU, so we increment by one  
    }  
}
```

CUDA – primer programa

BLOCK 1

```
__global__ void  
add( int *a, int *b, int *c ) {  
    int tid = 0;  
    if (tid < N)  
        c[tid] = a[tid] + b[tid];  
}
```

BLOCK 2

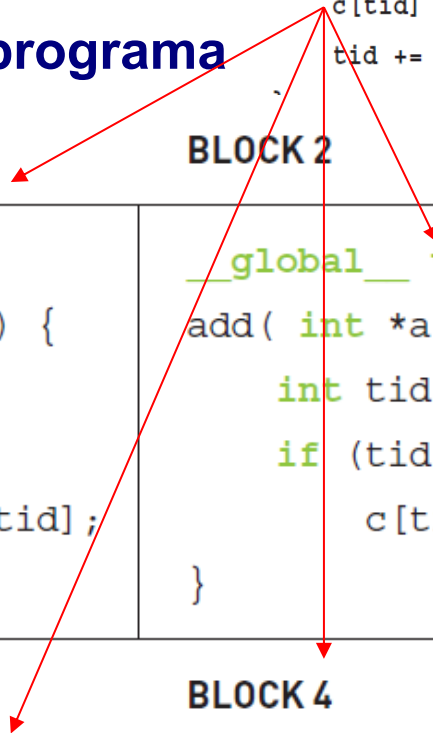
```
__global__ void  
add( int *a, int *b, int *c ) {  
    int tid = 1;  
    if (tid < N)  
        c[tid] = a[tid] + b[tid];  
}
```

BLOCK 3


```
__global__ void  
add( int *a, int *b, int *c ) {  
    int tid = 2;  
    if (tid < N)  
        c[tid] = a[tid] + b[tid];  
}
```

BLOCK 4

```
__global__ void  
add( int *a, int *b, int *c ) {  
    int tid = 3;  
    if (tid < N)  
        c[tid] = a[tid] + b[tid];  
}
```



CUDA – Primerjava izv. časov programa : $B[i] = A[i]*2$



Sling – Primerjava izv. časov programa za vektorski produkt

```
srun --reservation=fri --gpu=1 prog DotProd_reduction 8192 64 4096 1
```

Argumenti so naslednji.

Prvi določa dolžini vektorja,	8192
drugi določa število niti znotraj delovne skupine,	64
tretji skupno število niti,	4096
zadnji pa je 0 ali 1 in pove ali naj se pri računanju skalarnega produkta uporabi redukcija na GPE.	1

Dolžina vektorja	Niti - del.sk.	Skupno niti	Redukcija	Čas [ms]
8192	64	4096	1	0.018
32768	64	4096	1	0.014
262144	64	4096	1	0.052

Delavnice

- Osnove superračunalništva >
- Superračunalništvo bolj zares >
- Računajmo na vsa jedra >
- Programiranje grafičnih procesnih enot (OpenCL) ✓
- 0 delavnici**
- Uvod >
- Grafične procesne enote >
- Uvod v OpenCL >
- Programiranje GPE >
- Programiranje grafičnih procesnih enot (CUDA) >



<https://doc.sling.si/workshops/programming-gpu/intro/course/>

Programiranje grafičnih procesnih enot¹²

Patricio Bulić, Univerza v Ljubljani, Fakulteta za računalništvo in informatiko

Vsebina


Na delavnici se bomo seznanili z zgradbo grafičnih procesnih enot (GPE) in njihovim programiranjem. Skozi preproste primere se bomo naučili pisati in poganjati programe na GPE, prenašati podatke med gostiteljem in GPE ter učinkovito uporabljati pomnilnik na GPE. Na koncu bomo napisali paralelni program za obdelovanje slik, ki se bo izvajal na GPE. Programirali bomo v programskem jeziku C in uporabljali ogrodje OpenCL, ki podpira GPE različnih proizvajalcev (Nvidia, AMD, Intel). Enake pristope k uporabi GPE vključujejo tudi druga ogrodja, zato se boste s pridobljenim znanjem zlahka lotili programiranja grafičnih kartic Nvidia v ogrodju CUDA ali pa programski jezik C zamenjali s kakšnim drugim.

5.2.4.4 GPU

NEWS COMPUTING

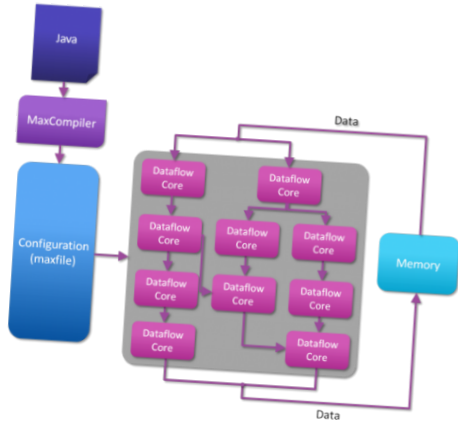

GPUs Can Now Analyze a Billion Complex Vectors in Record Time

> The advancement boosts the speed of GPU image analysis eight-fold

BY MICHELLE HAMPSON | 16 JUL 2021 | 2 MIN READ | 




5.2.5 Primeri drugačnega pristopa – veliko različnih pristopov






SpiNNaker




SpiNNaker chip



Multi-chip packaging by UNISem Europe



Use cases demand more embedded intelligence



5.2.5 Primeri drugačnega pristopa - Podatkovno pretokovni računalniki

Maxeler : Za t.i. “**Big Data algorithms**” in **enako ceno HW** sedaj dobimo :

- **pohitritev, 20-200x**
- **manjša poraba** (manjši račun za elektriko 20x)
- velikost **20x manjša**

<https://www.maxeler.com/>

Bistvo konkretnega („podatkovno pretokovnega“) pristopa:

- **prevajanje pod nivojem strojnega jezika** prinaša pohitritve, manjšo moč, velikost in ceno
- cena, ki jo zato plačamo:
 - bistveno **težje in zamudnejše programiranje**
 - cikel prevajanja traja daljši čas (tudi 4 ure). na srečo obstaja simulator...
- Posledica:
 - Idealno le za t.i. **WORM (“Write Once Run Many”)** aplikacije

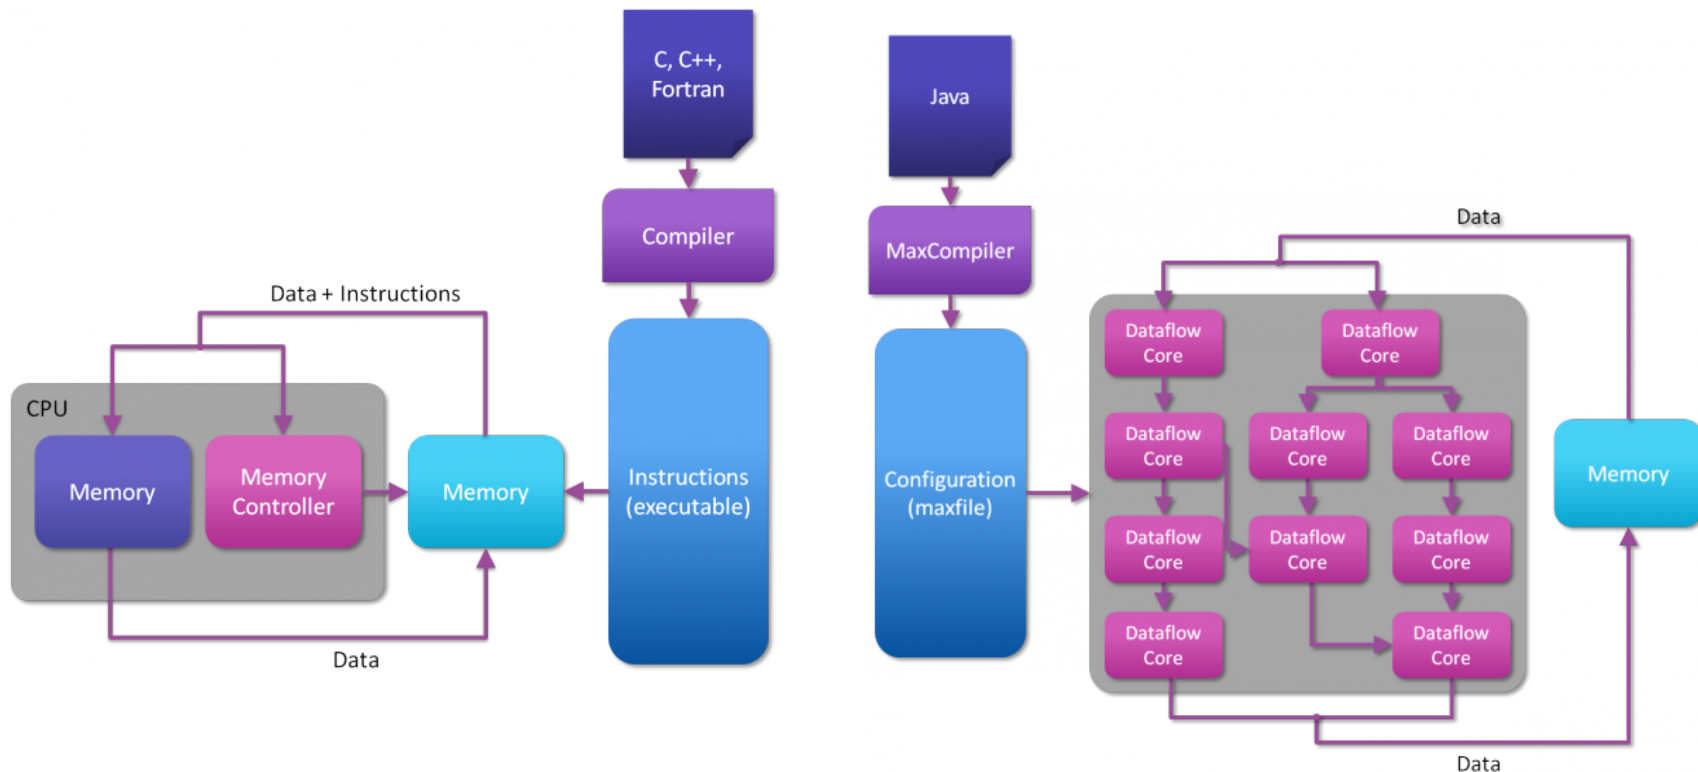
Maxeler Technologies Develops Real Time FPGA-
Based Processing for European XFEL
March 13, 2020

<https://www.maxeler.com/technology/dataflow-computing/>



Podatkovno pretokovni pristop „DataFlow Computing“ – by Maxeler

„ControlFlow“ vs. „DataFlow“



<https://www.maxeler.com/technology/dataflow-computing/>

5.2.5 Primer drugačnega pristopa: Podatkovno pretokovni računalniki

■ Osnovne ideje:

- Hitrost : programiranje na nižjem nivoju logičnih vrat (FPGA) !

- Varčnost: več počasnejših jeder!
 - optimalna prilagoditev porabe delovanju
 - nižja frekvenca delovanja
 - večja učinkovitost - manjša poraba !!!

- Manj prostora: ekstenzivni paralelizem!
 - manjši kontrolni del
 - večina logike v končni funkciji obdelave podatkov

MultiCore/ManyCore



Machine Level Code

Dataflow



Gate Transfer Level

MultiCore/ManyCore




Dataflow




$$P = kfU^2$$

MultiCore/ManyCore



DataFlow

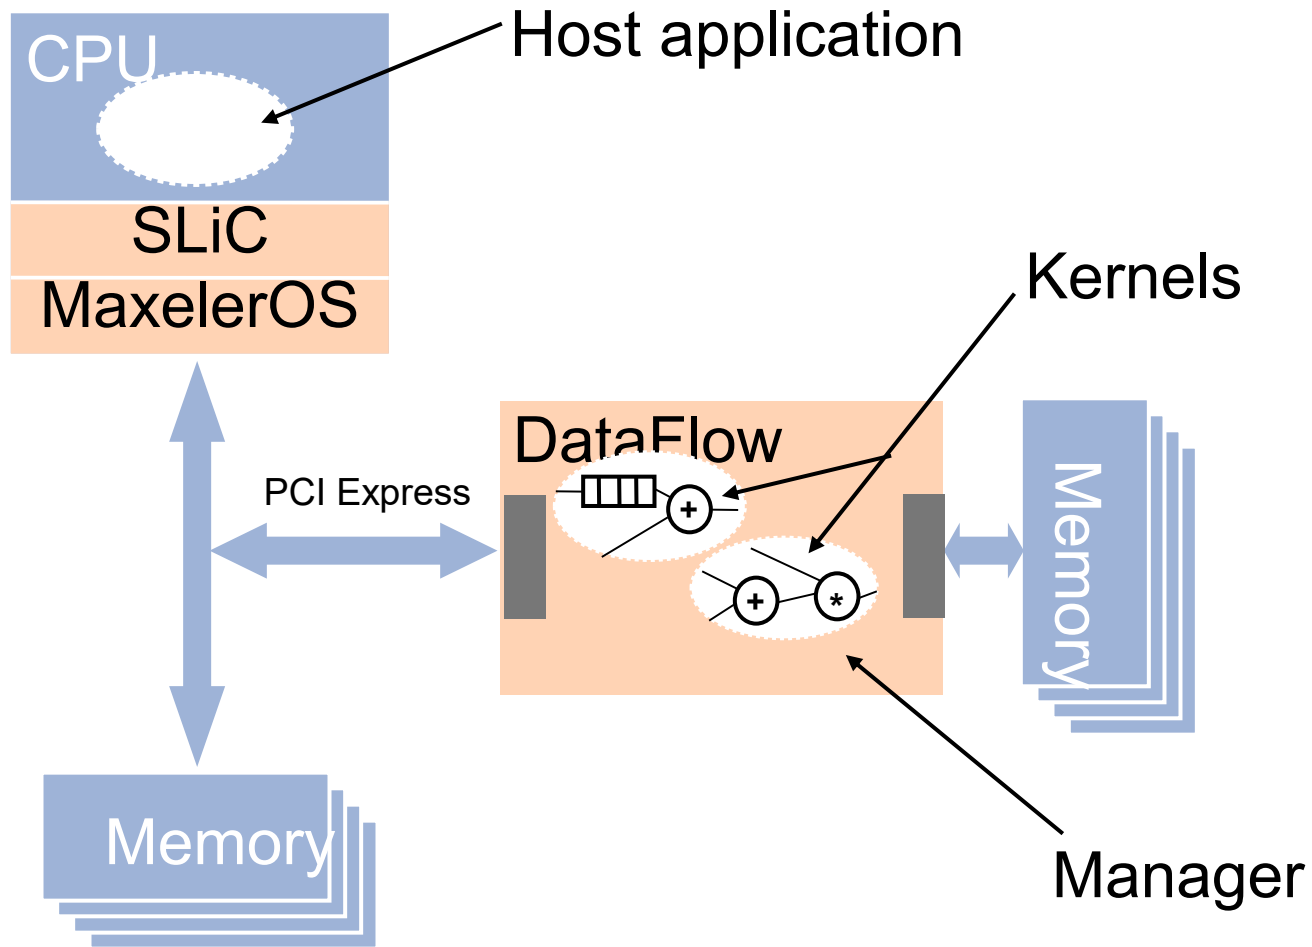


5.2.5 Primer drugačnega pristopa: Podatkovno pretokovni računalniki

- **Primeri uporabe** Maxeler tehnologije:
 - geofizika (pohitritev 20x-40x, analiza tal za vrtine,...),
 - bančništvo (pohitritev 200x-1000x, JP Morgan: plačilni promet s karticami)
 - podatkovno rudarjenje (“Datamining” - Google), ...

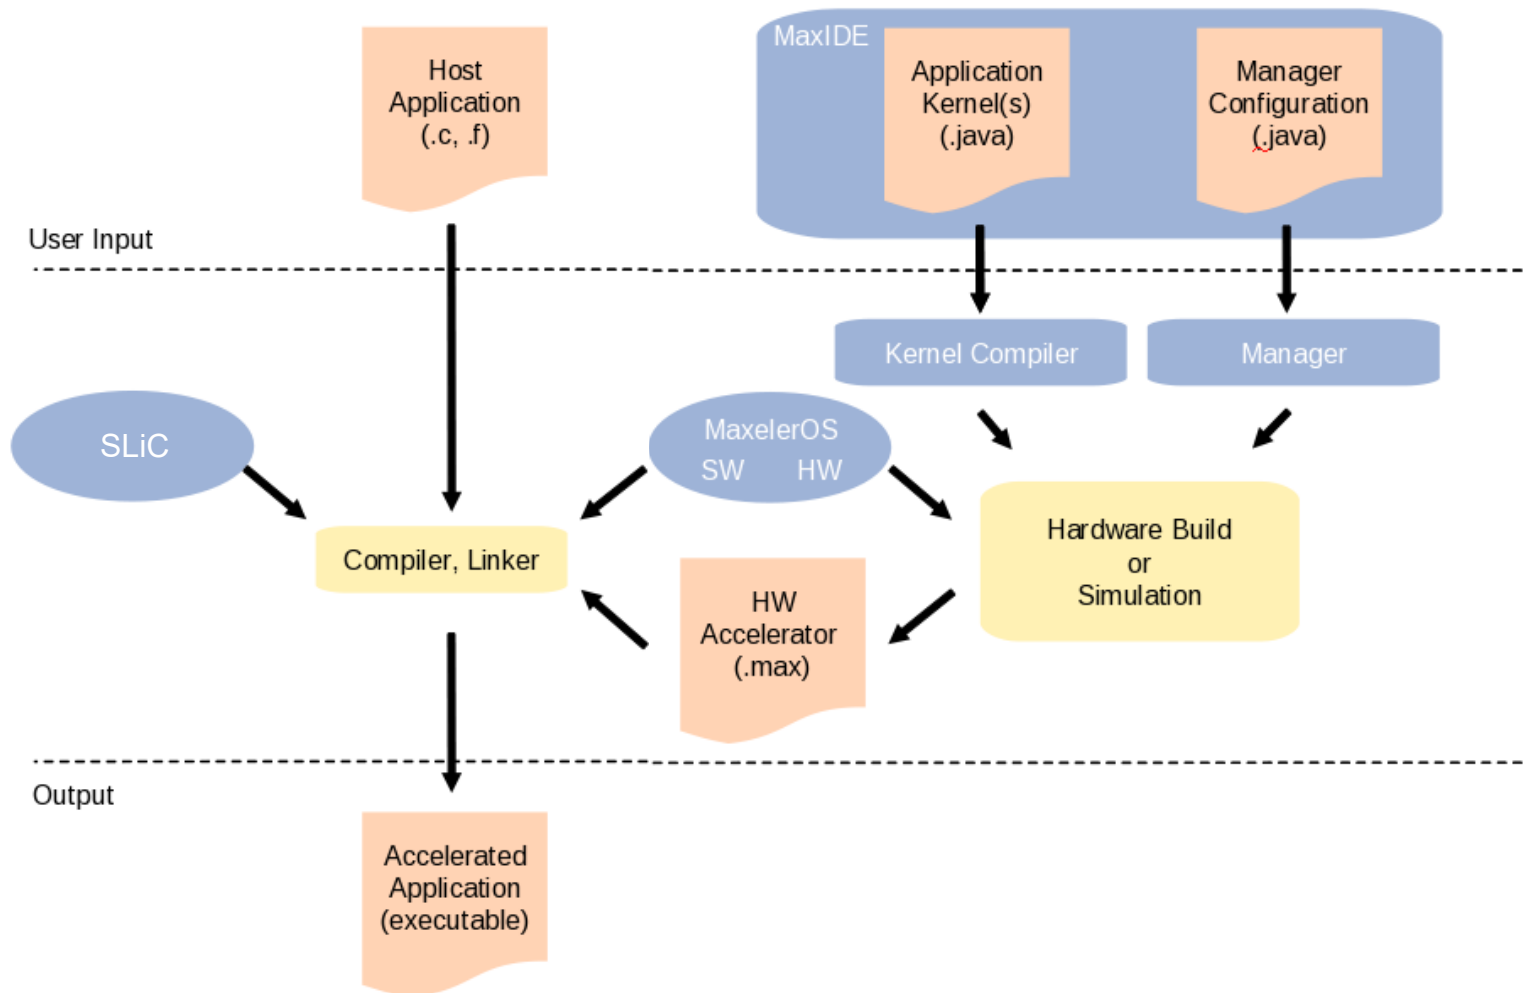
- **Veliko odvisno od programerja** in njegove spretnosti (izkušenj):
 - if-then-else z MUX-i
 - razvezava zank do zapolnitve prostora
 - odločanje, kaj gre v paralelno izvajanje (čimveč ponovitev, podatkov)
 - težko razhroščevanje („debugging“):
 - preventiva – programer hkrati napiše tudi testne programe

Maxeler Application Components



C / C++ / Fortran

Java




DataFlow Programming (C +Java)

program



aparturna izvedba

```
7 public class MovingAverageKernel extends Kernel {  
8  
9     public MovingAverageKernel(KernelParameters parameters, int N) {  
10         super(parameters);  
11  
12         // Input  
13         HWVar x = io.input("x", hwFloat(8, 24));  
14  
15         // Data  
16         HWVar prev = stream.offset(x, -1);  
17  
18         HWVar next = stream.offset(x, 1);  
19  
20         HWVar sum = prev+x+next;  
21  
22         HWVar result = sum/3;  
23  
24         // Output  
25         io.output("y", result, hwFloat(8, 24));  
26     }  
27 }
```



Dataflow Applications

HUXLEY MUSCLE MODEL

A skeletal muscle fiber generates tension when properly stimulated, for instance by the nervous system or by electrical impulses.

Authors: Ogrjen Andric, Ivan Milankovic, Milos Ivanovic

☆☆☆☆

$$r_{xy} = \frac{\sum x_i y_i - n \bar{x} \bar{y}}{\sqrt{n \sum x_i^2 - (\sum x_i)^2} \sqrt{n \sum y_i^2 - (\sum y_i)^2}}$$

Correlation

Correlation is a statistical measure that indicates the extent to which two or more variables fluctuate together. A positive

Author: Maxeler Analytics

☆☆☆☆

Poisson Solver

Application solves the 3D Poisson Equation for n input sets of N x N x N size, where N has to be power of 2 and at least 32. N

Author: Marko Stojanovic

☆☆☆☆

Line Rate

High Speed Packet Capture

Provides sustained line-rate packet capture in distributed write mode and at bursts of up to 24GB in size in local write mode. Custom

Author: Maxeler Networking

☆☆☆☆

Brain Network

Linear correlation analysis of brain images to detect brain activity.

Author: Maxeler London

☆☆☆☆

Linear Regression

In statistics, linear regression is an approach for modeling the relationship between a scalar dependent variable and one or more

Author: Maxeler Intern

☆☆☆☆

Classification

Cluster analysis or clustering is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are

Author: Maxeler Networking

☆☆☆☆

Packet Pusher

Takes a PCAP (packet capture) file and replays its contents through a network interface, either at the original timing or at a constant

Author: Maxeler Networking

☆☆☆☆

Breast Mammogram ROI Extraction

This App extracts the region of interest from breast mammogram images. Basically, this app removes pectoral muscle and background

Authors: Faculty of Engineering, University of Kragujevac

☆☆☆☆

Low-Latency HTTP Web-Server

This App implements an HTTP Web-Server in a DFE. The App serves static webpages directly from LMEM, routes uncached

Author: Maxeler Belgrade

☆☆☆☆

Lattice Boltzmann

This lattice Boltzmann app demonstrates some simple concepts behind doing a finite difference type simulation on Maxeler DFE.

Author: David Packwood

☆☆☆☆

Boosted Decision Tree

An ensemble of decision trees predicts the class of input data in the DFE with low latency.

Author: Sioni Summers

☆☆☆☆

Simplex

Simplex algorithm for the Maxeler data-flow computer architecture.

Authors: Uroš Čehaj, Jurij Mihelič

☆☆☆☆

Network Sorting

Network sorting algorithms are parallel comparison-based sorting algorithms with a fixed structure, that

Authors: Vukasin Baric

☆☆☆☆

<https://appgallery.maxeler.com/#/>

5.2.5 Primeri drugega pristopa – Stohastično računanje



By the numbers: Conventional binary numbers, just like the decimal numbers in everyday use, rely on the concept of place value [left]. Stochastic bitstreams don't use place value; the value they represent is determined by how often 1s appear [right].



Many times better: Using stochastic bitstreams, multiplication can be carried out with a single AND gate. Here two bitstreams, representing $1/2$ and $3/4$, provide the inputs. The output has 1s in three of eight positions, meaning that it represents a value of $3/8$ —the product of the two inputs.

5.2.5 Primeri drugačnega pristopa – Stohastično računanje

Conventional binary



Stochastic computing



COMPUTING WITH RANDOMNESS



Stochastic computing,
a 50-year-old idea,
is set for
a comeback

Always on edge: Edge detection is commonly used in image processing. Here, an edge-detection algorithm that uses conventional binary numbers [top row] is compared with one that uses stochastic bitstreams [bottom row]. The stochastic results hold up much better as the bit-error rate is increased from 0.1 percent [far left] to 0.5 percent [middle left] to 1.0 percent [middle right] and finally to 2.0 percent [far right].

LOBO: A HYBRID APPROXIMATE MULTIPLIER FOR ENERGY-EFFICIENT COMPUTING


RATKO PILIPOVIĆ AND PATRICIO BULIĆ

UNIVERSITY OF LJUBLJANA, FACULTY OF COMPUTER AND INFORMATION SCIENCE,
VEČNA POT 113, 1000 LJUBLJANA, SLOVENIA



ENERGY AND AREA SAVINGS

Bar diagram below presents achieved reductions in area and energy compared to exact radix-4 Booth multiplier in TSMC 180nm technology. Reduction goes up to 70 % in energy consumption and 50 % in area utilization.



CONCLUSIONS


- Novel hybrid approximate multiplier LOBO is presented
- Datapath pruning is introduced in order to reduce delay and overall hardware complexity
- LOBO delivers great reductions in area and energy consumption with small computational error
- The feasibility of LOBO multiplier is shown in image sharpening application

IMAGE SHARPENING SHOWCASE


In Table below we present peak-signal-to-noise ratio (PSNR) between an image filtered with an exact multiplier and an image filtered with an approximate multiplier. For test image we selected *Mandrill.tiff*

Multiplier	PSNR [dB]
$L(d = 10, T_H = 0, T_S = 0)$	61.54
$L(d = 12, T_H = 0, T_S = 0)$	46.57
$L(d = 10, T_H = 8, T_S = 6)$	55.14
$L(d = 12, T_H = 8, T_S = 8)$	40.01
$L(d = 10, T_H = 12, T_S = 6)$	55.14
$L(d = 12, T_H = 12, T_S = 8)$	40.01

Filtered with exact multiplier



Filtered with $L(d = 12, T_H = 12, T_S = 8)$




5.2.5 Primeri drugačnega pristopa - Neuro-morfni čipi («neurosynaptic, brain-like architecture«)

Ideja:

- posnemati **nevronska mrežno zgradbo** človeških možganov
- specifičen računski model – oblika nevronske mreže - ni namenjen splošnemu računanju


Primer 1 : IBM TrueNorth

- 1 milijon programabilnih nevronov
- 256 milijonov programabilnih sinaps (povezav med nevroni)
- 4096 nevrosinaptičnih jeder (pomnilnik, procesor in komunikacija)
- 5.4 milijarde tranzistorjev
- Izredno nizka poraba: cca **20mW/cm² ali 70mW/čip**. Kako ?
 - ni ure – asinhrono delovanje – »event driven computing«
 - zmogljive medsebojne povezave jeder
 - tehnologija izdelave vezij za mobilne naprave
 - digitalni način dela namesto analognega !



<http://spectrum.ieee.org/computing/hardware/how-ibm-got-brainlike-efficiency-from-the-truenorth-chip>

5.2.6 Neuro-morfni čipi («brain-like architecture«)



Neuro-morfni čipi («brain-like architecture»)

Primerjava : TrueNorth (komercialna), SpiNNaker (raziskovalna)

TrueNorth IBM, DARPA SyNAPSE

Project features

Low-power neuromorphic chip designed for applications in mobile sensors, cloud computing, and so on.

Chip specs

1 million neurons
256 million synapses

Largest current configuration



16 million neurons
4 billion synapses

Power density



Final configuration 10 billion neurons; 100 trillion synapses

SpiNNaker University of Manchester, Human Brain Project


Project features

Enables low-power, large-scale digital model of brain; helps improve models of brain diseases.

Chip specs

Up to 16,000 neurons
16 million synapses

Largest current configuration



Up to 20 million neurons
20 billion synapses

Power density



Final configuration Up to 1 billion neurons; 1 trillion synapses

Neuro-morfni čipi («brain-like architecture«)

Nanowire Synapses 30,000x Faster Than Nature's > This way to artificial, ultraefficient, optoelectronic neurons?

BY CHARLES Q. CHOI | 28 OCT 2022 | 4 MIN READ |

New AI Chip Twice as Energy Efficient as Alternatives > RRAM-based neuromorphic chips are more versatile and accurate as well

BY PAYAL DHAR | 29 AUG 2022 | 3 MIN READ |

Toward Optoelectronic Chips That Mimic the Human Brain > An interview with a NIST researcher keen to improve spiking neuromorphic networks


BY DAVID SCHNEIDER | 18 APR 2022 | 5 MIN READ |

Brain-Inspired Chips Good for More than AI, Study Says > Neuromorphic tech from IBM and Intel may prove useful for analyzing X-rays, stock markets, and more

BY CHARLES Q. CHOI | 15 FEB 2022 | 4 MIN READ |

IEEE Spectrum FOR THE TECHNOLOGY INSIDER

NEWS | ARTIFICIAL INTELLIGENCE



Intel's Neuromorphic Chip Gets A Major Upgrade > Loihi 2 packs 1 million neurons in a chip half the size of its predecessor

BY SAMUEL K. MOORE | 05 OCT 2021 | 4 MIN READ |

The *SpiNNaker* Project


Research Groups: APT - Advanced Processor Technologies (School of Computer Science - The University of Manchester)

Steve Furber

ICL Professor of Computer Engineering

The University of Manchester

<https://apt.cs.manchester.ac.uk/projects/SpiNNaker/>



European Research Council
Established by the European Commission



Human Brain Project



The University of Manchester

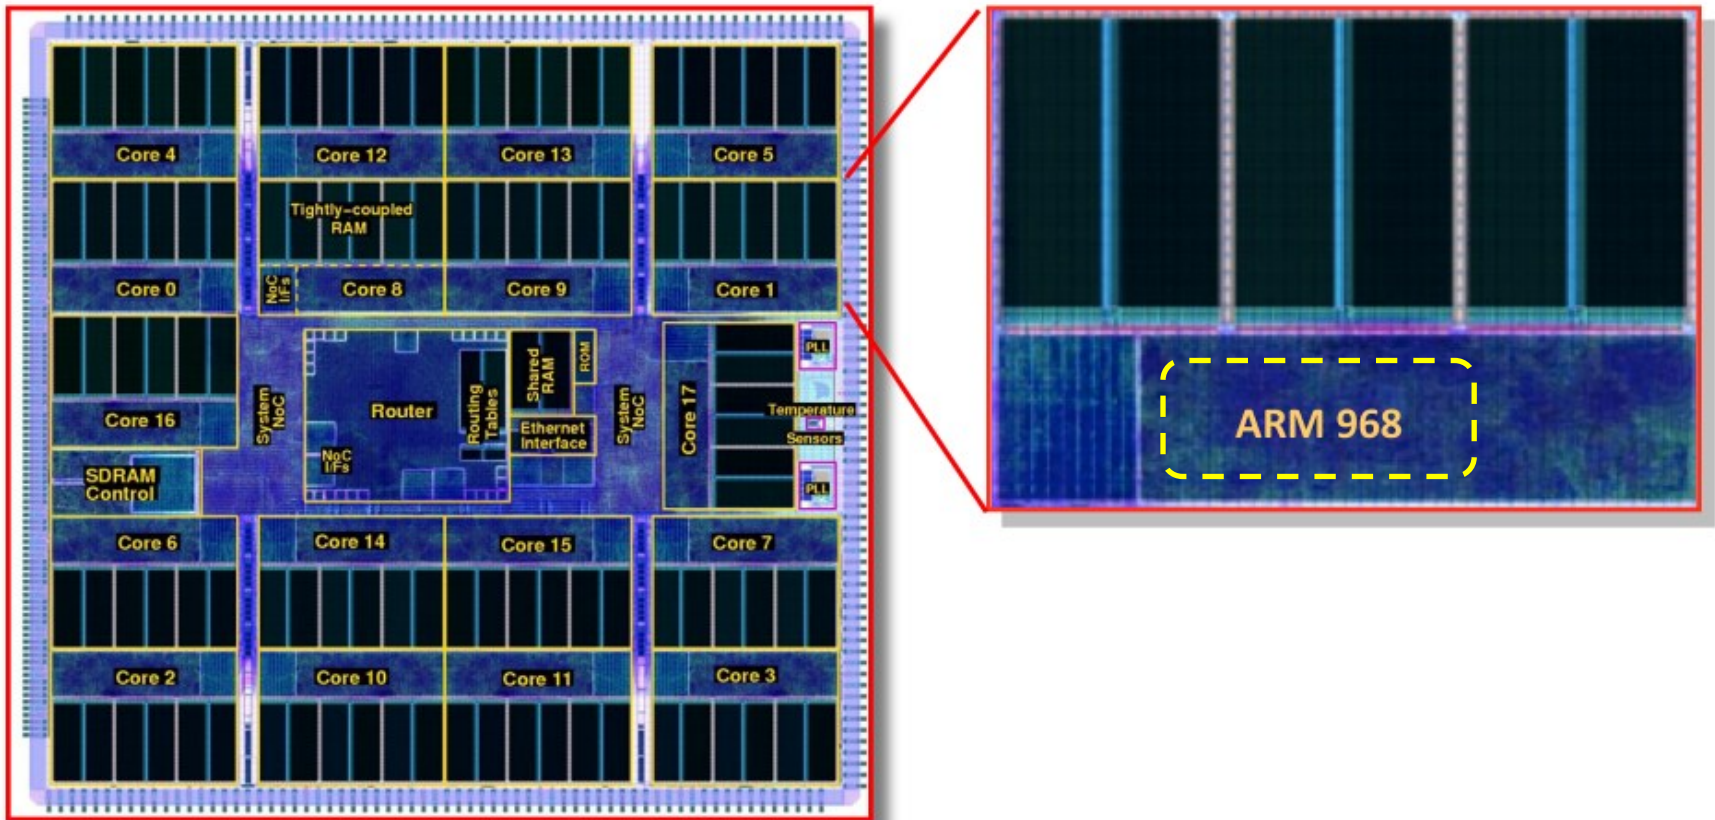


The University of Sheffield.

UNIVERSITY OF
Southampton



SpiNNaker CPU (2011)




Building brains

- Brains demonstrate
 - massive parallelism (10^{11} neurons)
 - massive connectivity (10^{15} synapses)
 - excellent power-efficiency
 - much better than today's microchips
 - low-performance components (~ 100 Hz)
 - low-speed communication (\sim metres/sec)
 - adaptivity – tolerant of component failure
 - autonomous learning



The **world's largest neuromorphic supercomputer** designed and built to work in the same way a human brain



*„SpiNNaker completely re-thinks the way conventional computers work. We've essentially created **a machine that works more like a brain than a traditional computer**, which is extremely exciting.“*

SpiNNaker


Biologically Inspired
Massively Parallel
Architecture

SpiNNaker project


- A million mobile phone processors in one computer
- Able to model about 1% of the human brain...
- ...or 10 mice!




EPSRC



SpiNNaker chip



Multi-chip
packaging by
UNISEM Europe



SpiNNaker machines

Spinnaker

102



72 cores
- pond snail scale



103



864 cores
- drosophila scale




104



20,000 cores
- frog scale



105



100,000 cores
- mouse scale

106

- HBP platform
 - 500,000 cores
 - 6 cabinets (including server)
- Launch
 - 22 March 2016



72 cores
- pond snail scale



1 000,000 cores
10 cabinets
~1% of human brain
~10 mice brains

Launch: 2018



Spinnaker – Trenutno 1

SpiNNaker (Spiking Neural Network Architecture) is a massively parallel, manycore supercomputer architecture designed by the Advanced Processor Technologies Research Group (APT) at the Department of Computer Science, University of Manchester.

- It is composed of 57,600 ARM9 processors (specifically **ARM968**),
 - each with 18 cores and 128 MB of mobile DDR SDRAM,
 - totalling 1,036,800 cores and over 7 TB of RAM.^[3]
 - The computing platform is based on spiking neural networks, useful in simulating the human brain (see Human Brain Project).^{[4][5][6][7][8][9][10][11][12]}
- The completed design is housed
 - in 10 19-inch racks,
 - with each rack holding over 100,000 cores.^[13]
 - The cards holding the chips are held in 5 Blade enclosures, and
 - each core emulates 1000 Neurons.^[13]
 - In total, the goal is to simulate the behavior of aggregates of up to a billion neurons in real time.^[14] This machine requires **about 100 kW from a 240 V supply** and an air-conditioned environment.^[15]
- On October 14, 2018 the HBP announced that the million core milestone had been achieved.^{[18][19]}
- On September 24, 2019 HBP announced that a **8 million euro grant**, that will fund construction of the **second generation machine**, (called spincloud) has been given to TU Dresden.^[20]



Spinnaker – 2


- On September 24, 2019 HBP announced that a **8 million euro grant**, that will fund construction of the **second generation machine**, (called spincloud) has been given to TU Dresden.^[20]

The full-scale SpiNNaker 2 will consist of **10 Mio ARM cores** distributed across 70.000 Chips in 10 server racks.

SpiNNaker is unique because, unlike traditional computers, it doesn't communicate by sending large amounts of information from point A to B via a standard network. Instead it **mimics the massively parallel communication architecture of the brain, sending billions of small amounts of information simultaneously to thousands of different destinations.**

Saxon Science Ministry delivers 8 Mio Euro to TU Dresden for second generation SpiNNaker machine

Press Release, For Immediate Release



Pictured From left, Professor Steve Furber (University of Manchester), Professor Christian Mayr (TU Dresden) and Professor Katrin Amunts, Scientific Director of the Human Brain Project.

<https://spinncloud.com/>


Our present

SpiNNaker2 is its new generation, architected by Professor Dr. Christian Mayr, who has disrupted the SpiNNaker architecture by creating a hybrid system that combines statistical AI and Neuromorphic computing


https://tu-dresden.de/ing/elektrotechnik/iee/hpsn/forschung/forschungsprojekte/spinncloud?set_language=en

5.2.5 Primeri drugačnega pristopa – DSA (Domain Specific Architectures)

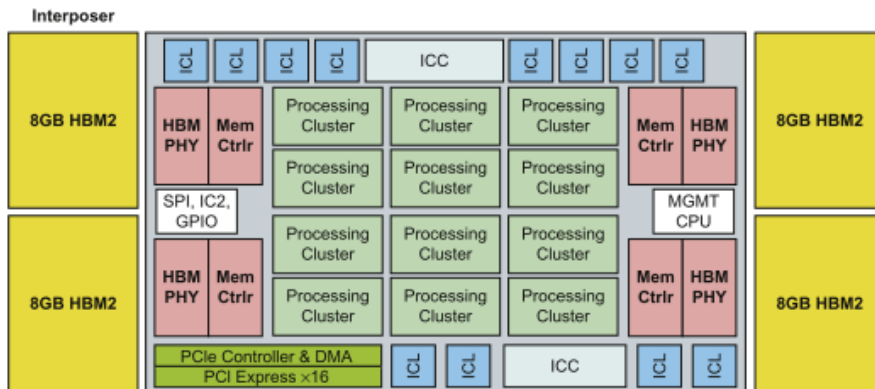
Google's Tensor Processing Unit, an Inference Data Center Accelerator




Microsoft Catapult, a Flexible Data Center Accelerator



Intel Crest, a Data Center Accelerator for Training




Pixel Visual Core, a Personal Mobile Device Image Processing Unit



5.2.5 Primeri drugačnih pristopov/tehnologij

These Super-Efficient, Artificial Neurons Do Not Use Electrons > So could the *brain's* super-efficiency have to do with ions?


BY PAYAL BHAM | 03 SEP 2021 | 3 MIN READ | □



Designing **electronic systems that mimic the human brain**, both in terms of energy use and ability to carry information, is a holy grail of scientific research.


New AI Chip Twice as Energy Efficient as Alternatives > RRAM-based neuromorphic chips are more versatile and accurate as well

BY PAYAL BHAM | 29 AUG 2022 | 3 MIN READ | □



Nanowire Synapses 30,000x Faster Than Nature's > This way to artificial, ultraefficient, optoelectronic neurons?


BY CHARLES Q. CHOI | 28 OCT 2022 | 4 MIN READ | □



The NIST superconducting circuit pictured above behaves like an artificial version of a synapse, a connection between nerve cells, or neurons, in the brain. ©. SHAM AND S. PRINAVERA/NIST


Intel's Neuromorphic Chip Gets A Major Upgrade > Loihi 2 packs 1 million neurons in a chip half the size of its predecessor

BY SAMUEL K. MOORE | 06 OCT 2021 | 4 MIN READ | □



Computing With Chemicals Makes Faster, Leaner AI > Battery-inspired artificial synapses are gaining ground

BY DINA GEMINA | 02 DEC 2022 | 6 MIN READ | □



5.2.6 „Deep Learning“ – HW


NEWS SEMICONDUCTORS

Cerebras' New Monster AI Chip Adds 1.4 Trillion Transistors > Shift to 7-nanometer process boosts the second-generation chip's transistor count to a mind boggling 2.6-trillion

BY SAMUEL K. MOORE | 20 APR 2021 | 4 MIN READ | □

Photonic Chip Performs Image Recognition at the Speed of Light > New photonic deep neural network could also analyze audio, video, and other data

BY CHARLES Q. CHOI | 06 JUN 2022 | 3 MIN READ | □




The chip uses a deep neural network of optical waveguides smaller than a square centimeter. The network can detect and classify an image in less than a nanosecond, without the need for a separate processor or memory unit. ELLA MARU STUDIO



5.2.6 „Deep Learning“ – HW

NVIDIA Turing Tensor Cores Second Generation

NVIDIA Turing™ Tensor Core technology features multi-precision computing for efficient AI inference. Turing Tensor Cores provide a range of precisions for deep learning training and inference, from FP32 to FP16 to INT8, as well as INT4, to provide giant leaps in performance over NVIDIA Pascal™ GPUs.



Sanja Fidler

Sanja Fidler is the senior director of AI at NVIDIA, leading a research lab in Toronto focused on 3D Deep learning for Omniverse. She is also an associate professor at the Department of Computer Science, University of Toronto, Ontario, which she joined in 2014.




<https://www.nvidia.com/en-us/data-center/tensor-cores/>

NVIDIA Turing

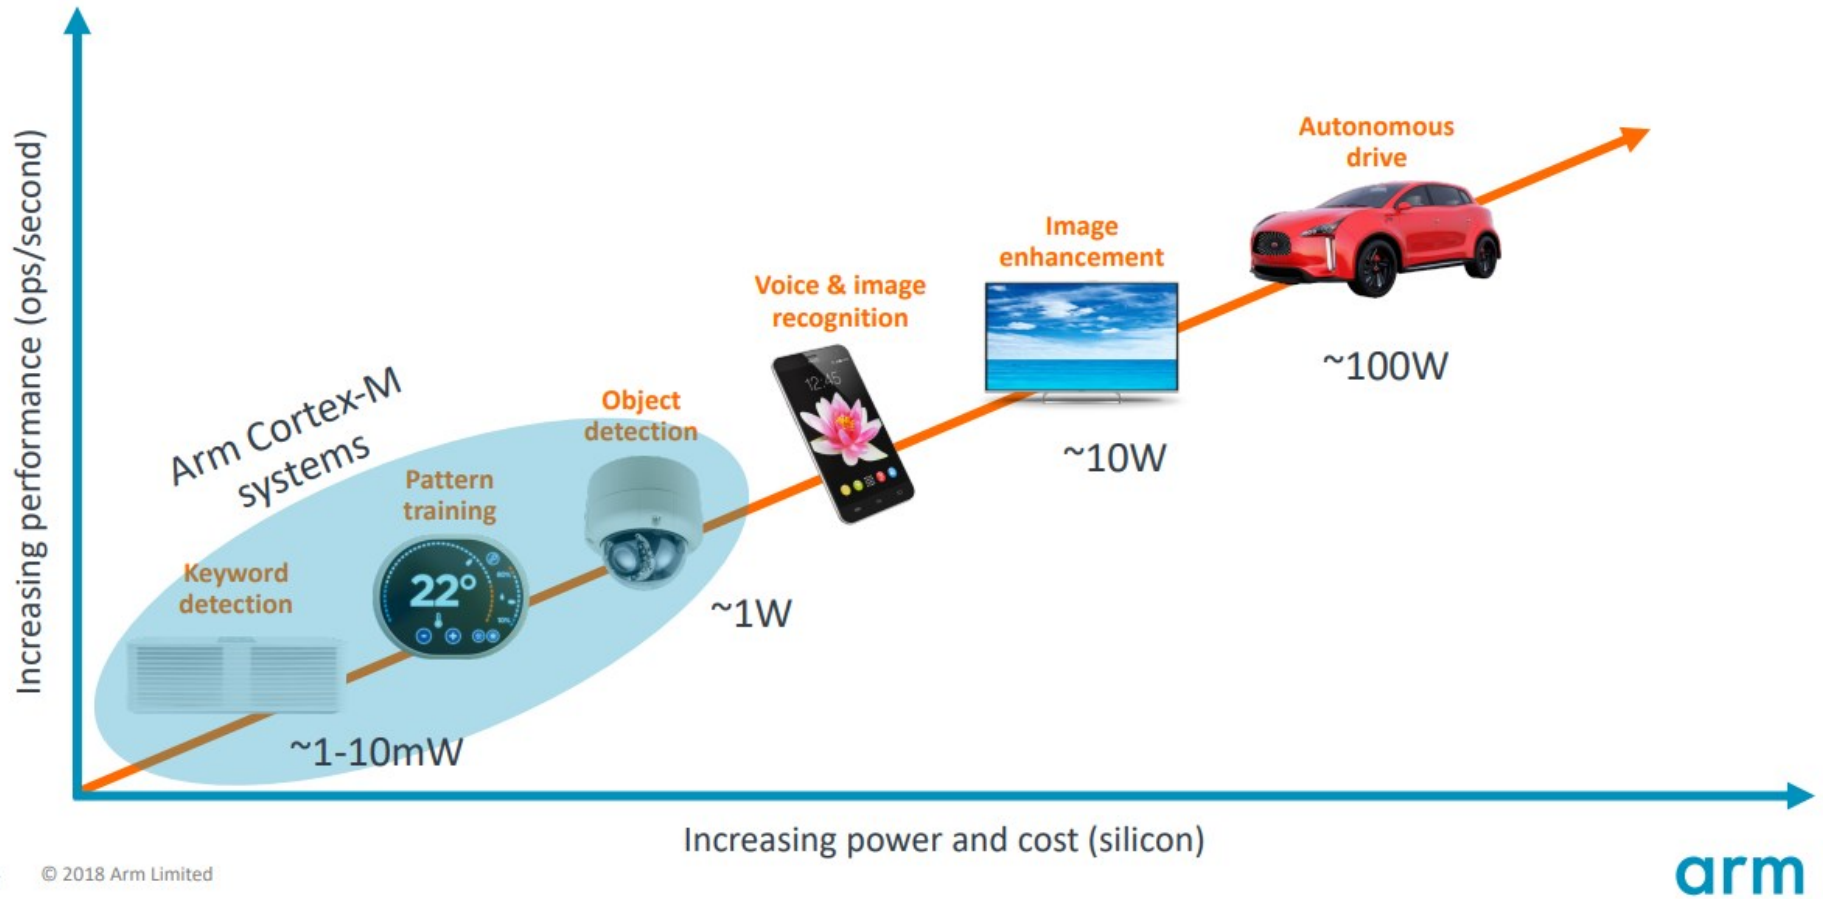
Graphics Reinvented

The greatest leap since the invention of the NVIDIA® CUDA® GPU in 2006, the NVIDIA Turing™ architecture fuses real-time ray tracing, AI, simulation, and rasterization to fundamentally change computer graphics.




5.2.8 „Embedded AI“ – ARM Cortex-M

ML Edge Use cases



5.2.8 „Embedded AI“ – ARM Cortex-M

Use cases demand more embedded intelligence



More data


More processing

More algorithms


More compute

5.2.8 „Embedded AI“ – ARM Cortex-M

Use cases demand more embedded intelligence



Edge computing



Edge computing

Smart system challenges Moving to edge computing

CLOUD COMPUTING

EDGE COMPUTING

Collect and send data

Protocol translation and device management

Big Data and heavy computation

Time-sensitive applications should be locally processed


Time-sensitive applications are limited by remote cloud

Collect, Process And Send Data

Local Processing of Data


Optimized computation and Advanced Analysis

Opportunity: move computation to sensor nodes with local processing for real-time elaboration and best power efficiency




AI is moving to the edge

Capabilities included in a project



*AI and GPU were added in 2019.

Advanced technology in a project



68% of EMEA users are considering using Machine Learning

- 15% of embedded projects already include AI in 2019
- Pervasion of Machine Learning and other AI capabilities

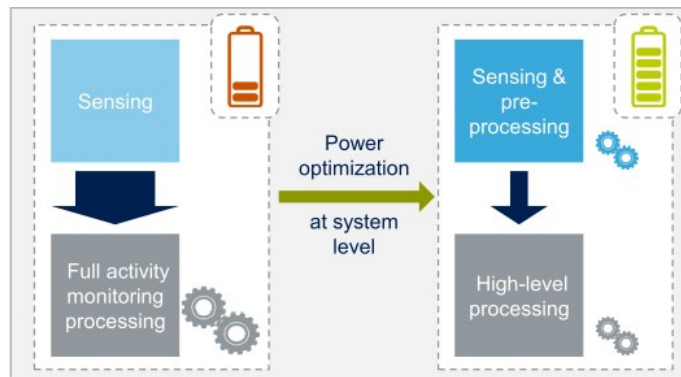
Edge computing – moduli, tipala

BHI260AP

Ultra-low power, high performance, **self-learning AI** smart sensor with integrated accelerometer and gyroscope

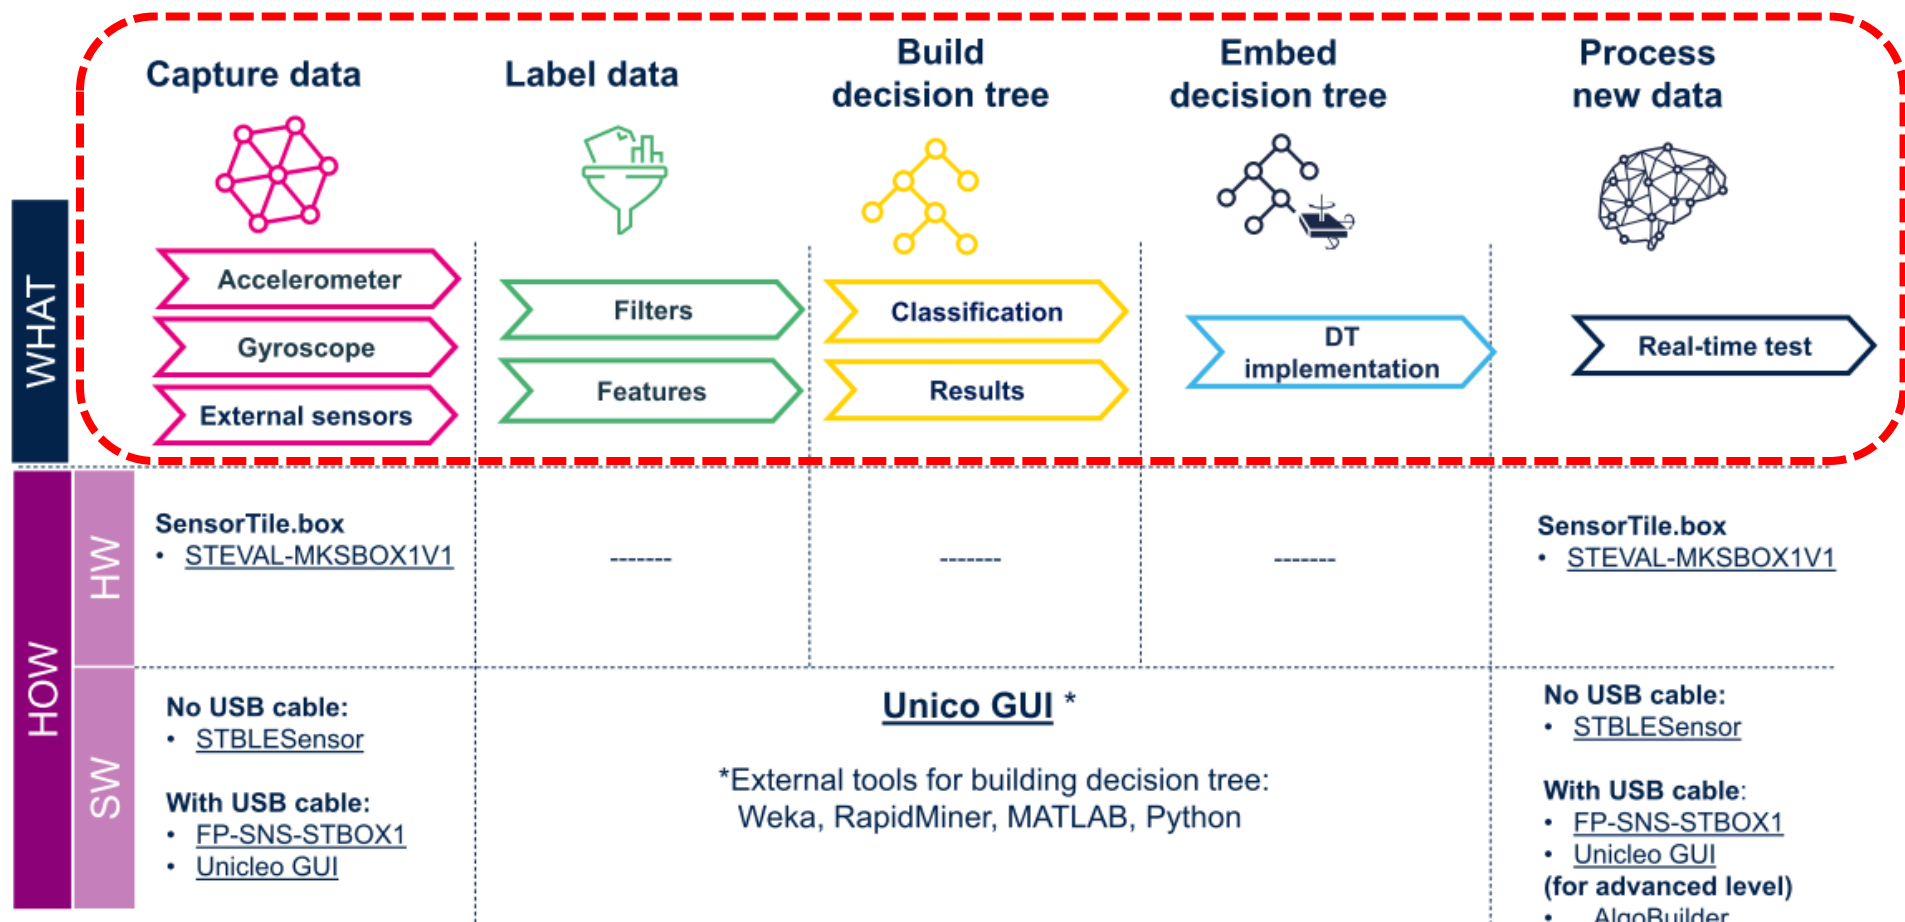


LSM6DSOX Unique Performance



Edge computing – moduli, tipala

LSM6DSOX – SensorTile.box




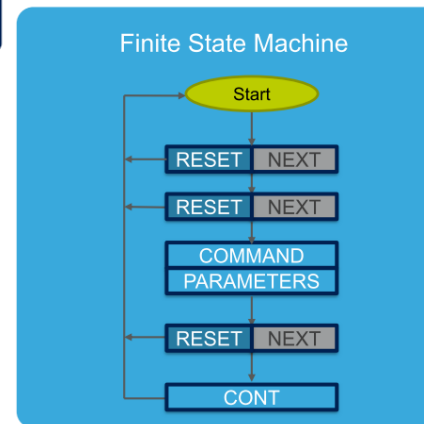
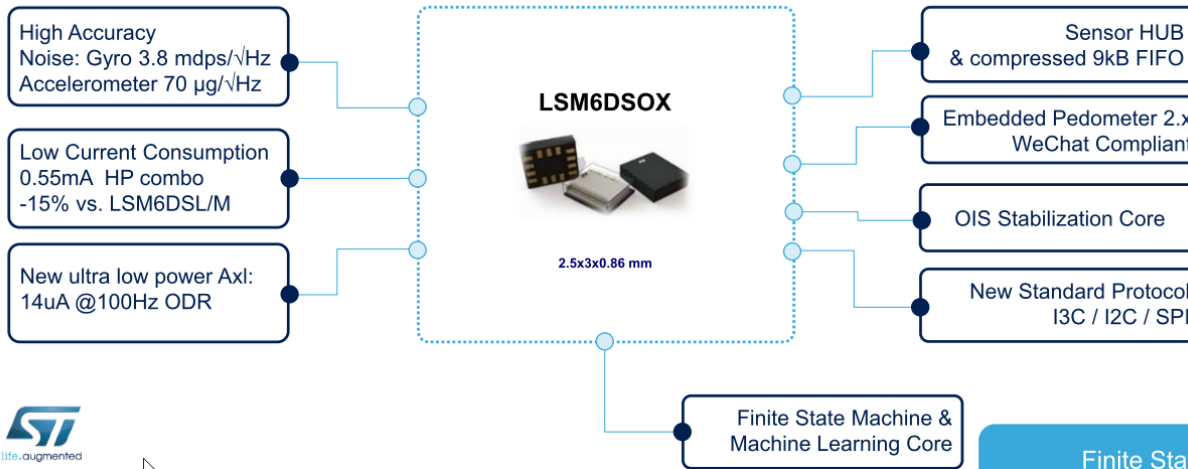
Edge computing – moduli, tipala



LSM6DSOX Unique Performance



Improved Accuracy, Optimized System Power



FSM & MLC allows sensors to process data with reduced help of a host MCU

5.2.8 „Embedded AI“ – ARM Cortex-M

Overview




Nicla Vision allows you to build your next smart project. Ever wanted an automated house? Or a smart garden? Well, now it's easy with the Arduino IoT Cloud compatible boards. It means: you can connect devices, visualize data, control and share your projects from anywhere in the world. Whether you're a beginner or a pro, we have a wide range of [plans](#) to make sure you get the features you need.

Nicla Vision combines a powerful **STM32H747AI16 Dual ARM® Cortex® M7/M4 IC** processor with a **2MP color camera** that supports TinyML, as well as a **smart 6-axis motion sensor, integrated microphone and distance sensor**.

You can easily include it into any project because it's designed to be compatible with all Arduino Portenta and MKR products, fully integrates with OpenMV, supports MicroPython and also offers both WiFi and Bluetooth® Low Energy connectivity.

It's so compact – with its 22.86 x 22.86 mm form factor – it can physically fit into most scenarios, and requires so little energy it can be powered by battery for standalone applications.

<https://store.arduino.cc/products/nicla-vision>



Tech specs

Microcontroller	<p>STM32H747AI16 Dual Arm® Cortex® M7/M4 IC:</p> <ul style="list-style-type: none"> • 1x Arm® Cortex® M7 core up to 480 MHz • 1x Arm® Cortex® M4 core up to 240 MHz
Sensors	<ul style="list-style-type: none"> • 2 MP Color Camera • 6-Axis IMU (LSM6DSOX) • Distance / Time Of Flight sensor (VL53L1CBV0FY/1) • Microphone (MP34DT05)