

Parallel Accelerators

Přemysl Šůcha

“Parallel algorithms”, 2017/2018
CTU/FEL


Topic Overview

- Graphical Processing Units (GPU) and CUDA
- Vector addition on CUDA
- Intel Xeon Scalable Processors

Graphical Processing Units



GPU – Nvidia - Roadmap



GPU - Use

- GPU is especially well-suited to address problems that can be expressed as **data-parallel computations**.
- The same program is executed on many data elements in parallel - with high **arithmetic intensity**.
- Applications that process **large data sets** can use a data-parallel programming model to speed up the computations (3D rendering, image processing, video encoding, ...)
- Many algorithms **outside the field of image rendering and processing** are accelerated by data-parallel processing too (machine learning, general signal processing, physics simulation, finance, ...).

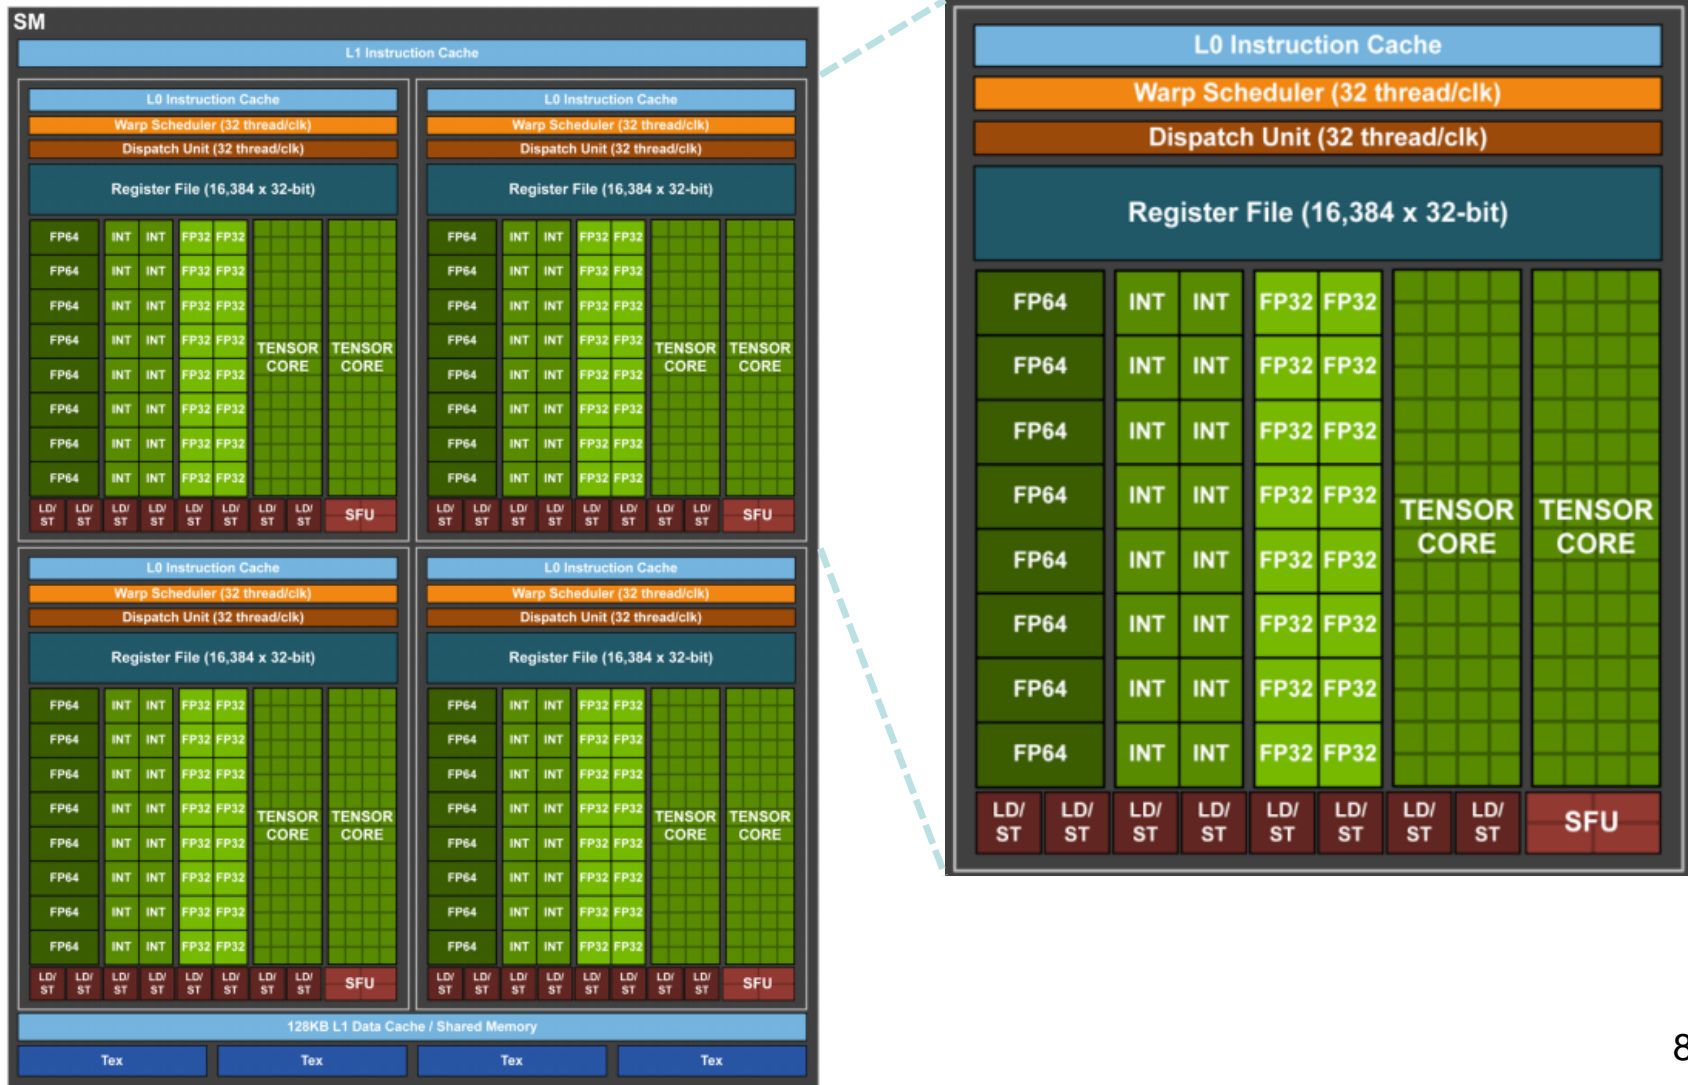
GPU - Overview

- CPU code runs on the **host**, GPU code runs on the **device**.
- A **kernel** consists of multiple threads.
- Threads execute in 32-thread groups called **warps**.
- Threads are grouped into **blocks**.
- A collection of blocks is called a **grid**.

GPU - Hardware Organization Overview

- GPU chip consists of one or more streaming **multiprocessors (SM)**.
- A multiprocessor consists of 1 (CC 1.x), 2 (CC 2.x), or 4 (CC 3.x, 5.x, 6.x, 7.x) **warp schedulers**. (CC = CUDA Capability)
- Each warp scheduler can issue to 2 (CC 5 and 6) or 1 (CC 7) **dispatch units**.
- A multiprocessor consists of **functional units** of several types.

Streaming Multiprocessor (SM) - Volta



GPU - Functional Units

- **INT, FP32 (CUDA Core)** - functional units that executes most types of instructions, including most integer and single precision floating point instructions.
- **FP64 (Double Precision)** - executes double-precision floating point instructions.
- **SFU (Special Functional Unit)** - executes reciprocal and transcendental instructions such as sine, cosine, and reciprocal square root.
- **LD/ST (Load/Store Unit)** – handles load and store instructions.
- **TENSOR CORE** – for deep learning matrix arithmetic.

GPU - Tensor Core

- V100 GPU contains 640 Tensor Cores: eight (8) per SM.
- Tensor Core performs 64 floating point FMA (fused multiply-add) operations per clock.
- Matrix-Matrix multiplication (GEMM) operations are at the core of neural network training and inferencing.
- Each Tensor Core operates on a 4x4 matrices.

$$\mathbf{D} = \begin{pmatrix} A_{0,0} & A_{0,1} & A_{0,2} & A_{0,3} \\ A_{1,0} & A_{1,1} & A_{1,2} & A_{1,3} \\ A_{2,0} & A_{2,1} & A_{2,2} & A_{2,3} \\ A_{3,0} & A_{3,1} & A_{3,2} & A_{3,3} \end{pmatrix} \begin{pmatrix} B_{0,0} & B_{0,1} & B_{0,2} & B_{0,3} \\ B_{1,0} & B_{1,1} & B_{1,2} & B_{1,3} \\ B_{2,0} & B_{2,1} & B_{2,2} & B_{2,3} \\ B_{3,0} & B_{3,1} & B_{3,2} & B_{3,3} \end{pmatrix} + \begin{pmatrix} C_{0,0} & C_{0,1} & C_{0,2} & C_{0,3} \\ C_{1,0} & C_{1,1} & C_{1,2} & C_{1,3} \\ C_{2,0} & C_{2,1} & C_{2,2} & C_{2,3} \\ C_{3,0} & C_{3,1} & C_{3,2} & C_{3,3} \end{pmatrix}$$

FP16 or FP32 FP16 FP16 FP16 or FP32

Streaming Multiprocessor (SM)

- Each SM has a set of temporary **registers** split amongst threads.
- Instructions can access high-speed **shared memory**.
- Instructions can access a cache-backed **constant space**.
- Instructions can access **local memory**.
- Instructions can access **global space**. (very slow in general)

GPU Architecture - Volta



GPU Architectures

Tesla Product	Tesla K40	Tesla M40	Tesla P100	Tesla V100
Manufacturing Process	28nm	28nm	16nm	12nm
Transistors	7.1 Billion	8.0 Billion	15.3 Billion	21.1 Billion
SMs / GPU	15	24	28	40
F32 Cores / SM	192	128	64	64
F32 Cores / GPU	2880	3072	3584	5120
Peak FP32 TFLOPS	5	6.8	10.6	15.7
Peak FP64 TFLOPS	1.7	0,21	5.3	7.8
GPU Boost Clock	810/870 MHz	1114 MHz	1480 MHz	1530 MHz
TDP	235 W	250 W	300 W	300 W
Memory Interface	384-bit GDDR5	384-bit GDDR5	4096-bit HBM2	4096-bit HBM2
Maximum TDP	244W	250W	250W	300W


Single-Instruction, Multiple-Thread

- SIMT is an execution model where single instruction, multiple data (**SIMD**) **is combined with multithreading**.
- The SM creates, manages, schedules, and executes threads in groups of 32 parallel threads called **warps**.
- A warp start together at the same program address, but they have their **own instruction address counter** and **register state** and are therefore **free to branch** and **execute independently**.

CUDA


- The NVIDIA GPU architecture is built around a **scalable** array of multithreaded Streaming Multiprocessors (SMs).
- CUDA (Compute Unified Device Architecture) provides a way how a CUDA **program can be executed on any number of SMs**.
- A multithreaded program is partitioned into **blocks** of threads that execute independently from each other.
- A GPU with more multiprocessors will automatically execute the program in less time than a GPU with fewer multiprocessors.

CUDA



Grid/Block/Thread

- threads can be identified using a 1-D, 2-D, or 3-D thread index, forming a 1-D, 2-D, or 3-D block of threads, called a **thread block**.
- Blocks are organized into a 1-D, 2-D, or 3-D **grid** of **thread blocks**.



2-D grid with 2-D thread blocks

Kernel

- CUDA C extends C by allowing the programmer to define C functions, called **kernels**.

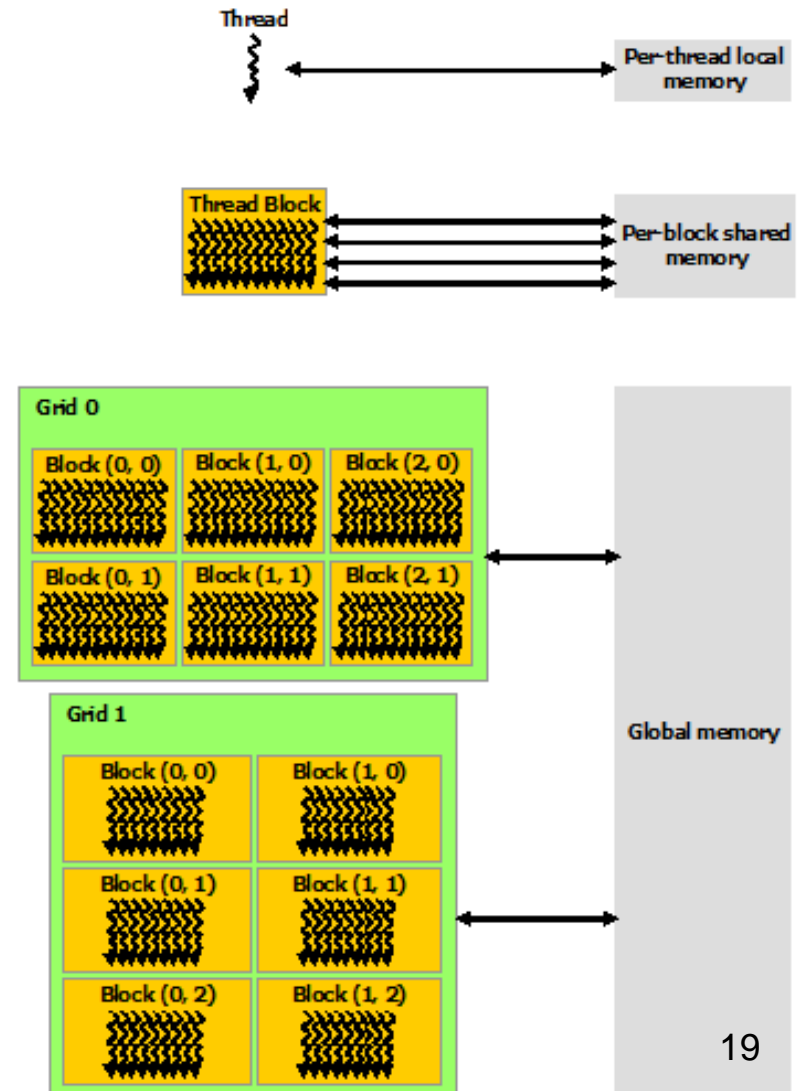
```
// Kernel definition
__global__ void VecAdd(float* A, float* B, float* C)
{
    int i = threadIdx.x;
    C[i] = A[i] + B[i];
}

int main()
{ ...
    // Kernel invocation with N threads inside 1 thread block
    VecAdd<<<1, N>>>(A, B, C);
}
```

- *threadIdx* is a 3-component vector, so that threads can be identified using a 1-D, 2-D, or 3-D **thread index**.

Memory Hierarchy

- Each thread has private set of **registers** and **local memory**.
- Each thread block has **shared memory** visible to all threads of the block.
- All threads have access to the same **global memory**.
- There are also two additional read-only memory spaces accessible by all threads (**constant** and **texture memory**).



GPU Programming - Example

- Element by element vector addition

[1] NVIDIA Corporation, *CUDA Toolkit Documentation* v9.0.176, 2017.

Element by element vector addition

```
/* Host main routine */
int main(void)
{
    int numElements = 50000;
    size_t size = numElements * sizeof(float);

    // Allocate the host input vectors A and B and output vector C
    float *h_A = (float *)malloc(size);
    float *h_B = (float *)malloc(size);
    float *h_C = (float *)malloc(size);

    // Initialize the host input vectors
    for (int i = 0; i < numElements; ++i)
    {
        h_A[i] = rand() / (float)RAND_MAX;
        h_B[i] = rand() / (float)RAND_MAX;
    }
}
```

Element by element vector addition

```
// Allocate the device input vectors A and B and output vector C
float *d_A = NULL;
cudaMalloc((void **)&d_A, size);
float *d_B = NULL;
cudaMalloc((void **)&d_B, size);
float *d_C = NULL;
cudaMalloc((void **)&d_C, size);
```

```
// Copy the host input vectors A and B in host memory to the device
// input vectors in device memory
cudaMemcpy(d_A, h_A, size, cudaMemcpyHostToDevice);
cudaMemcpy(d_B, h_B, size, cudaMemcpyHostToDevice);
```

Element by element vector addition

```
// Launch the Vector Add CUDA Kernel
int threadsPerBlock = 256;
int blocksPerGrid = (numElements + threadsPerBlock - 1) / threadsPerBlock;

vectorAdd<<<blocksPerGrid, threadsPerBlock>>>>(d_A, d_B, d_C, numElements);

// Copy the device result vector in device memory to the host result vector
// in host memory.
cudaMemcpy(h_C, d_C, size, cudaMemcpyDeviceToHost);
```

Element by element vector addition

```
// Free device global memory
err = cudaFree(d_A);
err = cudaFree(d_B);
err = cudaFree(d_C);

// Free host memory
free(h_A);
free(h_B);
free(h_C);

return 0;
}
```


Element by element vector addition

```
/**
 * CUDA Kernel Device code
 *
 * Computes the vector addition of A and B into C. The 3 vectors have the same
 * number of elements numElements.
 */
__global__ void vectorAdd(float *A, float *B, float *C, int numElements)
{
    int i = blockDim.x * blockIdx.x + threadIdx.x;


    if (i < numElements)
    {
        C[i] = A[i] + B[i];
    }
}
```

Intel Xeon Scalable Processors



Intel Xeon Scalable Processors

- **Intel Xeon** is a high-performance version of Intel desktop processors
- Intel Xeon Scalable Processors is a family of processors designed to provide scalable and reliable performance
- Scalable family are scalable from a two-socket configuration to an eight-socket configuration.

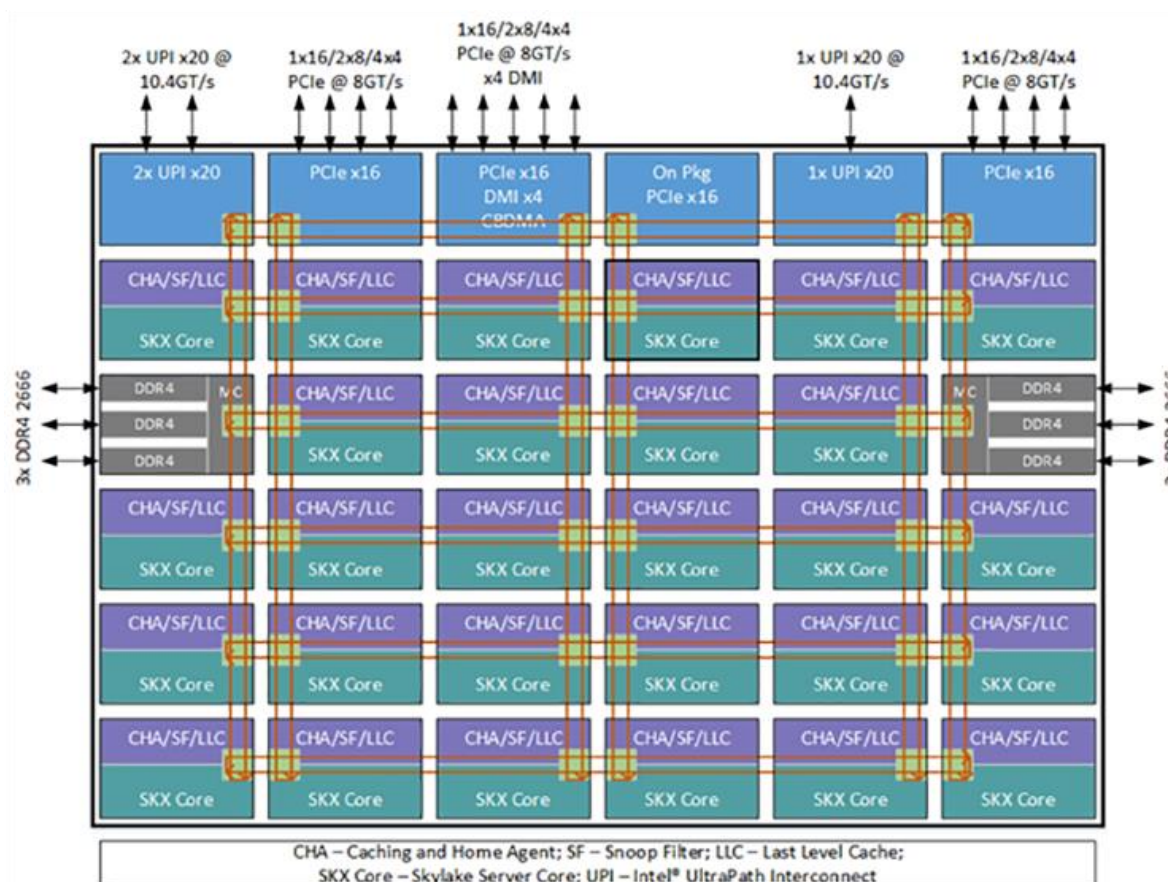


Intel Xeon Scalable Processors

- On the previous generations of Intel Xeon processor families the on-chip interconnect was realized using a **ring architecture**.
- With additional cores per processor the interconnect could become a performance limiter with the ring-based architecture.
- Intel Xeon processor Scalable family introduces a **mesh architecture** to mitigate the increased latencies and bandwidth constraints.
- The mesh architecture encompasses an array of vertical and horizontal communication paths allowing traversal from one core to another.

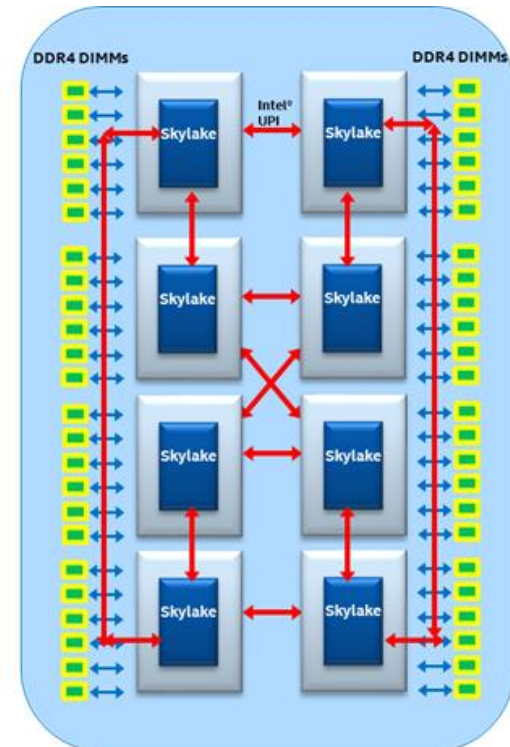
Intel Xeon Scalable Processors architecture

- The Intel Xeon processor Scalable family provides up to 28 cores per processor.



Intel Xeon Scalable Processors

- Intel uses **Ultra Path Interconnect (UPI)** as an interconnect for scalable systems containing multiple processors in a single shared address space.
- The operational is speed up to 10.4 GT/s



Intel Xeon Scalable Processors

- The Intel Xeon processor Scalable family introduces new Intel **AVX-512 instruction groups** (512-bit SIMD instructions)
- One of the extensions is Vector Neural Network Instruction (VNNI) that increases deep learning inference.

Intel Xeon Scalable Processors

Features	Intel® Xeon® processor E5 2600 product family	Intel® Xeon® processor E7 4600 product family	Intel® Xeon® Processor Scalable Family
Platform	Grantley	Brickland	Purley
CPU TDP (with IVR)	55-145W, 160W WS only	115-165W	45-205W
Socket	Socket R3	Socket R1	Socket P
Scalability	2S	2S, 4S, 8S	2S, 4S, 8S
Cores	Up to 22C with Intel® HT Technology	Up to 24C with Intel® HT Technology	Up to 28C with Intel® HT Technology
Mid-Level Cache	256kb private	256kb private	1 MB private
Last Level Cache	Up to 2.5MB/core (Inclusive)	Up to 2.5MB/core (Inclusive)	Up to 1.375MB/core (non-inclusive)
Memory	4 channels DDR4 per CPU RDIMM, LRDIMM	8 channels DDR4 per CPU with memory buffer RDIMM, LRDIMM	6 channels DDR4 per CPU RDIMM, LRDIMM
	1DPC=up to 2400, 2DPC= up to 2400, 3DPC=up to 1866	DDR3/4 Performance Mode 1333, 1600 DDR3/4 Lockstep mode 1333, 1600, 1866	2133, 2400, 2666 2DPC No 3 DPC support
Point to Point Link	Intel® QPI: 2 v1.1 channels per CPU 9.6 GT/s max	Intel® QPI: 3 v1.1 channels per CPU 9.6 GT/s max	UPI: 2-3 channels per CPU (9.6, 10.4 GT/s)
PCIe*	PCIe* 3.0 (2.5, 5.0, 8.0 GT/s)	PCIe* 3.0 (2.5, 5.0, 8.0 GT/s)	PCIe* 3.0 (2.5, 5.0, 8.0 GT/s)
	40 lanes per CPU	32 lanes per CPU	48 lanes per CPU Bifurcation support: x16, x8, x4
PCH	Wellsburg: DMI2 – 4 lanes; Up to 6xUSB3, 8x USB2 ports, 10xSATA3 ports; GbE MAC (+ External PHY)	Patsburg: 14x USB2 ports, 4x SATA2 ports, 2x SATA3 ports	Lewisburg: DMI3 – 4 lanes; 14xUSB2 ports Up to: 10xUSB3; 14xSATA3, 20xPCIe*3 New: Innovation Engine, Integrated Intel® Ethernet 4x10GbE ports, Intel® QuickAssist Technology
External Node Controller Support	None	3rd Party Node Controller	3rd Party Node Controller supported on select skus

References

[6] David Mulnix: Intel® Xeon® Processor Scalable Family Technical Overview, <https://software.intel.com/en-us/articles/intel-xeon-processor-scalable-family-technical-overview>, 2019.