Selecting the right Tesla/GTX GPU from a Drunken Baker's Dozen

GPU Computing Applications

Here's what Nvidia says its Tesla K20(X) card excels at doing - Seismic processing, CFD, CAE, Financial computing, Computational chemistry and Physics, Data analytics, Satellite imaging & Weather modeling.
Here's what Nvidia says its Tesla K10 card excels at doing - Seismic processing, signal and image processing & video analytics.

The Tesla cards have important features that GTX cards lack, such as ECC memory and an interconnect feature that allows direct communication between the Tesla card and other devices, such as storage devices, connected to PCIe. Here's some of the hardware features of some Tesla cards and some GTX cards. You should take note of the similarities and differences of the Tesla cards and their closest GTX card counterpart.

I. Number and Type of GPU

1) Tesla C2075 GPU Computing Module---1x GF100
2) GTX 480----------1x GF100
3) Tesla M2090 GPU Computing Module---1x GF110
4) GTX 580----------1x GF110-375-A1

5) Tesla K20X----------1x Kepler GK110
6) GTX Titan----------1x Kepler GK110-400-A1
7) Tesla K40 (Atlas) */----------1x GK180 (GK110 w/mods)
8) GTX 780 Ti----------1x Kepler GK110-425-B1
9) Tesla K20----------1x Kepler GK110
10) Tesla K10----------2x Kepler GK104s
11) GTX 690----------2× Kepler GK104-355-A2s

12) GTX 680 (4GB)----------1x Kepler GK104-400-A2
13) GT 640----------1x Kepler GK107

II, Peak double precision floating point performance

1) Tesla C2075 GPU Computing Module---.515 Tflops (DP - FMA)
2) GTX 480----------480 Fragment Pipelines;
----------60 Texture Units; 48 Raster Units
3) Tesla M2090 GPU Computing Module-------.666 Tflops (DP - FMA)
4) GTX 580----------512 Fragment Pipelines;
----------128 Texture Units; 48 Raster Units

5) Tesla K20X----------1.32 Tflops
6) GTX Titan----------1.30-1.50 Tflops
7) Tesla K40 (Atlas)---------->1.4 Tflops
8) GTX 780 Ti----------.210 Tflops
9) Tesla K20----------1.17 Tflops
10) Tesla K10----------.19 Tflops (.095 Tflops per GPU)
11) GTX 690----------8 GPCS / SMXS 16
----

12) GTX 680 (4GB)----------???
13) GT 640----------???

III. Peak single precision floating point performance

1) Tesla C2075 GPU Computing Module------1.29 Tflops Total(MUL+ADD +SF);
----------1.03 Tflops MAD(MUL+ADD)
2) GTX 480----------1.34 Tflops (FMA) (TE =.51)
3) Tesla M2090 GPU Computing Module-----1.66 Tflops Total(MUL+ADD +SF);
----------1.33 Tflops MAD(MUL+ADD)
4) GTX 580----------1.58 Tflops (FMA) (TE = .59)

5) Tesla K20X----------3.95 Tflops
6) GTX Titan----------4.50 Tflops

7) Tesla K40 (Atlas)----------->4.00 Tflops
8) GTX 780 Ti----------5.05 Tflops
9) Tesla K20----------3.52 Tflops
10) Tesla K10----------5.34 Tflops (2.67 Tflops per GPU)
11) GTX 690----------5.62 Tflops (2.81 Tflops per GPU) (FMA)
----------(TE = 1.1+)

12) GTX 680 (4GB)----------3.09 Tflops (FMA) (TE = .54)
13) GT 640----------.69 Tflops (FMA)

IV. Memory Clock (Effective Memory Clock) MHz
1) Tesla C2075 GPU Computing Module----?????
2) GTX 480----------1848 (3696); Memory Speed: 0.4ns
3) Tesla M2090 GPU Computing Module----?????
4) GTX 580----------2004 MHz

5) Tesla K20X----------?????
6) GTX Titan----------1502 (6008); Memory Speed: 0.33ns
7) Tesla K40 (Atlas)----------?????
8) GTX 780 Ti----------1752 (7008)
9) Tesla K20----------1502 (6008)
10) Tesla K10----------1502 (6008)
11) GTX 690----------1502 (6008); Memory Speed: 0.33ns

12) GTX 680 (4GB)----------1502 (6008); Memory Speed: 0.33ns
13) GT 640----------667 vs 891 (1334 vs 1782)

V. Memory Bus width
1) Tesla C2075 GPU Computing Module----384 bit
2) GTX 480----------384 bit

3) Tesla M2090 GPU Computing Module----384 bit
4) GTX 580----------384 bit

5) Tesla K20X----------384 bit
6) GTX Titan----------384 bit
7) Tesla K40 (Atlas)----------384 bit
8) GTX 780 Ti----------384 bit
9) Tesla K20----------320 bit
10) Tesla K10----------2x256 bit = 512
11) GTX 690----------2x256 bit = 512

12) GTX 680 (4GB)----------256 bit
13) GT 640----------128 bit

VI. Memory bandwidth (ECC off)
1) Tesla C2075 GPU Computing Module----144 GB/sec
2) GTX 480----------177 GB/sec
3) Tesla M2090 GPU Computing Module----177 GB/sec
4) GTX 580----------192 GB/sec

5) Tesla K20X----------250 GB/sec
6) GTX Titan----------288 GB/sec
7) Tesla K40 (Atlas)----------288 GB/sec
8) GTX 780 Ti----------336 GB/sec
9) Tesla K20----------208 GB/sec
10) Tesla K10----------320 GB/sec (160 GB/sec per GPU)
11) GTX 690----------384 GB/sec (2× 192)

12) GTX 680 (4GB)----------192 GB/s
13) GT 640----------21 vs 29 GB/sec

VII. Memory size (GDDR5)
1) Tesla C2075 GPU Computing Module-----6 GB GDDR5
2) GTX 480----------1.5 GB GDDR5
3) Tesla M2090 GPU Computing Module ----6 GB GDDR5

4) GTX 580---------- 1.5 / 3.0 GB GDDR5

5) Tesla K20X----------6 GB GDDR5
6) GTX Titan ----------6 GB GDDR5
7) Tesla K40 (Atlas)---------- 12 GB GDDR5
8) GTX 780 Ti---------- 3 GB GDDR5 **/
9) Tesla K20----------5 GB GDDR5
10) Tesla K10----------8 GB (4 GB per GPU) GDDR5
11) GTX 690----------4 GB (2 GB per GPU) GDDR5

12) GTX 680 (4GB)----------2 GB / 4 GB GDDR5
13) GT 640----------4 GB DDR3

VIII. CUDA Cores
1) Tesla C2075 GPU Computing Module----448
2) GTX 480----------480: (TMUS) 60: (ROPS) 48
3) Tesla M2090 GPU Computing Module----512
4) GTX 580----------512: (TMUS) 64: (ROPS) 48

5) Tesla K20X----------2688 FP32, 896 FP64
6) GTX Titan----------2688: (TMUS) 224: (ROPS) 48
7) Tesla K40 (Atlas)----------2880: (TMUS) 240: (ROPS) 48
8) GTX 780 Ti----------2880: (TMUS) 240: (ROPS) 48
9) Tesla K20----------2496 FP32, 832 FP64
10) Tesla K10----------3072 (1536 per GPU):
----------(TMUS) 2x128=256 :
----------(ROPS) 2x32=64)
11) GTX 690----------3072 (2×1536=3072 cores :
----------(TMUS) 2x128=256 :
----------(ROPS) 2x32=64)

12) GTX 680 (4GB)----------1536: (TMUS) 128: (ROPS) 32
13) GT 640----------384: (TMUS) 32: (ROPS) 16

IX. Core Clock
1) Tesla C2075 GPU Computing Module---575 MHz

2) GTX 480----------700 MHz
3) Tesla M2090 GPU Computing Module---650 MHz
4) GTX 580----------772 MHz

5) K20X----------732 MHz
6) GTX Titan SC----------876 / 928 / ??? MHz ***/
7) Tesla K40 (Atlas)----------??? MHz
8) GeForce GTX 780 Ti----------875 / 928 / 1020 MHz ***/
9) Tesla K20----------705 MHz
10) Tesla K10----------745 MHz
11) GTX 690----------915 / 1019 / 1058 MHz ***/

12) GTX 680 (4 GB)----------1019 / 1084 / ??? MHz ***/
13) GT 640----------901 MHz

X. Shaders Thread Processors Total or SM Count or SMX Count
1) Tesla C2075 GPU Computing Module---448
2) GTX 480----------15 SM
3) Tesla M2090 GPU Computing Module---512
4) GTX 580----------2x16 SM

5) K20X----------2688 FP32, 896 FP64
6) GTX Titan----------14 SMX
7) Tesla K40 (Atlas)----------???
8) GeForce GTX 780 Ti----------15 SMX
9) Tesla K20----------2496 FP32, 832 FP64
10) Tesla K10----------2 x 1536 (3072)
11) GTX 690----------2x8 SM

12) GTX 680 (4 GB)----------8 SM
13) GT 640----------2 SM

XI. Shaders Clock MHz
1) Tesla C2075 GPU Computing Module---1,150
2) GTX 480----------1,401
3) Tesla M2090 GPU Computing Module---1,300

4) GTX 580----------1,544

5) K20X----------732
6) GTX Titan----------???
7) Tesla K40 (Atlas)----------???
8) GeForce GTX 780 Ti----------???
9) Tesla K20----------705
10) Tesla K10----------745
11) GTX 690----------915

12) GTX 680 (4 GB)----------1005-6
13) GT 640----------797 or 900

XII. Open CL
1) Tesla C2075 GPU Computing Module---N/A
2) GTX 480----------1.1
3) Tesla M2090 GPU Computing Module---N/A
4) GTX 580----------1.1

5) K20X----------N/A
6) GTX Titan----------1.1
7) Tesla K40 (Atlas)----------N/A
8) GeForce GTX 780 Ti----------1.1
9) Tesla K20----------N/A
10) Tesla K10----------N/A
11) GTX 690----------1.1

12) GTX 680 (4 GB)----------1.1
13) GT 640----------1.1

XIII. Open GL
1) Tesla C2075 GPU Computing Module---N/A
2) GTX 480----------4.4
3) Tesla M2090 GPU Computing Module---N/A
4) GTX 580----------4.4

5) K20X----------N/A
6) GTX Titan----------4.4
7) Tesla K40 (Atlas)----------N/A
8) GeForce GTX 780 Ti----------4.4
9) Tesla K20----------N/A
10) Tesla K10----------N/A
11) GTX 690----------4.4

12) GTX 680 (4 GB)----------4.4
13) GT 640----------4.4

XIV. DirectX
1) Tesla C2075 GPU Computing Module---N/A
2) GTX 480----------11.06
3) Tesla M2090 GPU Computing Module---N/A
4) GTX 580----------11.0

5) K20X----------N/A
6) GTX Titan----------11.1
7) Tesla K40 (Atlas)----------N/A
8) GeForce GTX 780 Ti----------11.2
9) Tesla K20----------N/A
10) Tesla K10----------N/A
11) GTX 690----------11.0

12) GTX 680 (4 GB)----------11.0
13) GT 640----------11.0

XV. Bus Interface
1) Tesla C2075 GPU Computing Module---PCIe 2.0 x16
2) GTX 480----------PCIe 2.0 x16
3) Tesla M2090 GPU Computing Module---PCIe 2.0 x16
4) GTX 580----------PCIe 2.0 x16

5) K20X----------PCIe 3.0 x16
6) GTX Titan----------PCIe 3.0 x16

7) Tesla K40 (Atlas)----------PCIe 3.0 x16
8) GeForce GTX 780 Ti----------PCIe 3.0 x16
9) Tesla K20----------PCIe 3.0 x16
10) Tesla K10----------PCIe 3.0 x16
11) GTX 690----------PCIe 3.0 x16

12) GTX 680 (4 GB)----------PCIe 3.0 x16
13) GT 640----------PCIe 3.0 x16

XVI. TDP
1) Tesla C2075 GPU Computing Module---225W
2) GTX 480----------250W
3) Tesla M2090 GPU Computing Module---250W
4) GTX 580----------244W

5) K20X----------235W
6) GTX Titan----------250W
7) Tesla K40 (Atlas)----------235W
8) GeForce GTX 780 Ti----------250W
9) Tesla K20----------225W
10) Tesla K10----------250W
11) GTX 690----------300W

12) GTX 680 (4 GB)----------195W
13) GT 640----------65W

See more at: http://www.nvidia.com/object/tesla-servers.html#sthash.OkVHZxWN.dpuf

*/ http://videocardz.com/46388/nvidia-launch-tesla-k40-atlas-gk180-gpu
**/ EVGA is said to be making a 6 gig version
***/ Base MHz / Average Boost MHz / Max Boost MHz