GPU Bench

GPU Performance Details: Tesla K40c

Contents:

System Configuration

⚠ **Note that this is previously stored data and does not reflect your system configuration.**

MATLAB Release: R2016a

Host

Name	Intel(R) Xeon(R) CPU E5-1620 0 @ 3.60GHz
Clock	3601 MHz
Cache	1024 KB
NumProcessors	4
OSType	Windows
OSVersion	Microsoft Windows 7 Enterprise

GPU

Name	Tesla K40c
Clock	745 MHz
NumProcessors	15
ComputeCapability	3.5
TotalMemory	11.25 GB
CUDAVersion	7.5
DriverVersion	8.17.13.5456 (354.56)

Results for MTimes (double)

These results show the performance of the GPU or host PC when calculating a matrix multiplication of two NxN real matrices. The number of operation assumed to be `2*N^3 - N^2`.

This calculation is usually compute-bound, i.e. the performance depends mainly on how fast the GPU or host PC can perform floating-point operations.

Raw data for Tesla K40c - MTimes (double)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	64,512	0.05	1.38
4,096	520,192	0.05	11.16
16,384	4,177,920	0.05	81.11
65,536	33,488,896	0.09	363.21
262,144	268,173,312	0.47	572.05
1,048,576	2,146,435,072	2.25	954.23
4,194,304	17,175,674,880	15.69	1094.71
16,777,216	137,422,176,256	119.08	1154.03
67,108,864	1,099,444,518,912	944.18	1164.44
268,435,456	8,795,824,586,752	7528.32	1168.36

(N gigaflops = $Nx10^9$ operations per second)

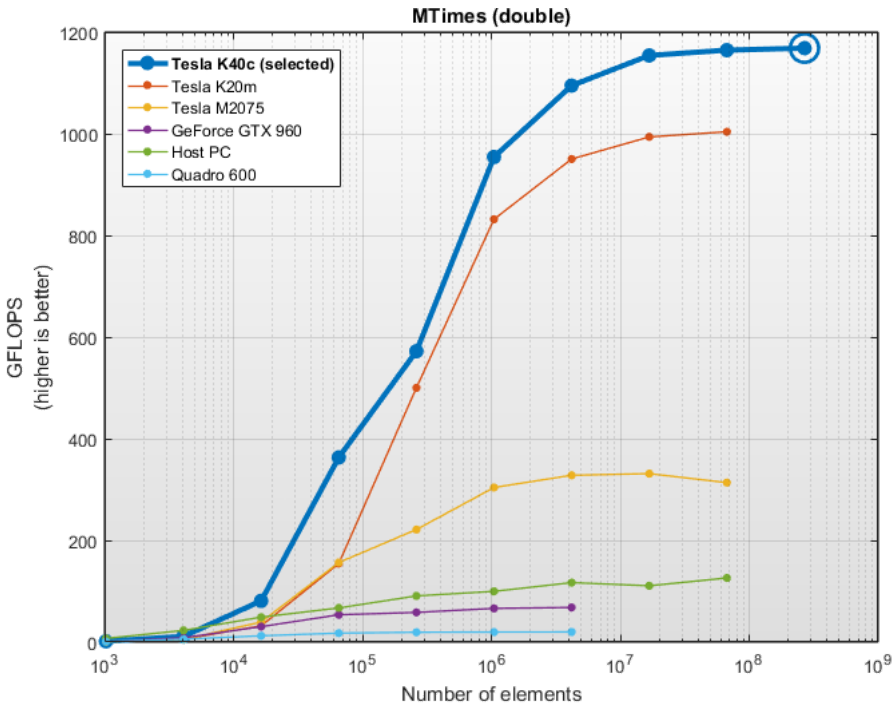

Results for Backslash (double)

These results show the performance of the GPU or host PC when calculating the matrix left division of an `NxN` matrix with an `Nx1` vector. The number o operations is assumed to be `2/3*N^3 + 3/2*N^2`.

This calculation is usually compute-bound, i.e. the performance depends mainly on how fast the GPU or host PC can perform floating-point operations.

Raw data for Tesla K40c - Backslash (double)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	23,381	0.19	0.13
4,096	180,907	0.24	0.76
16,384	1,422,677	0.43	3.31
65,536	11,283,115	1.68	6.72
262,144	89,871,701	3.42	26.26
1,048,576	717,400,747	9.16	78.29
4,194,304	5,732,914,517	30.42	188.46
16,777,216	45,838,150,315	139.56	328.45
67,108,864	366,604,539,221	792.23	462.75

(N gigaflops = $Nx10^9$ operations per second)

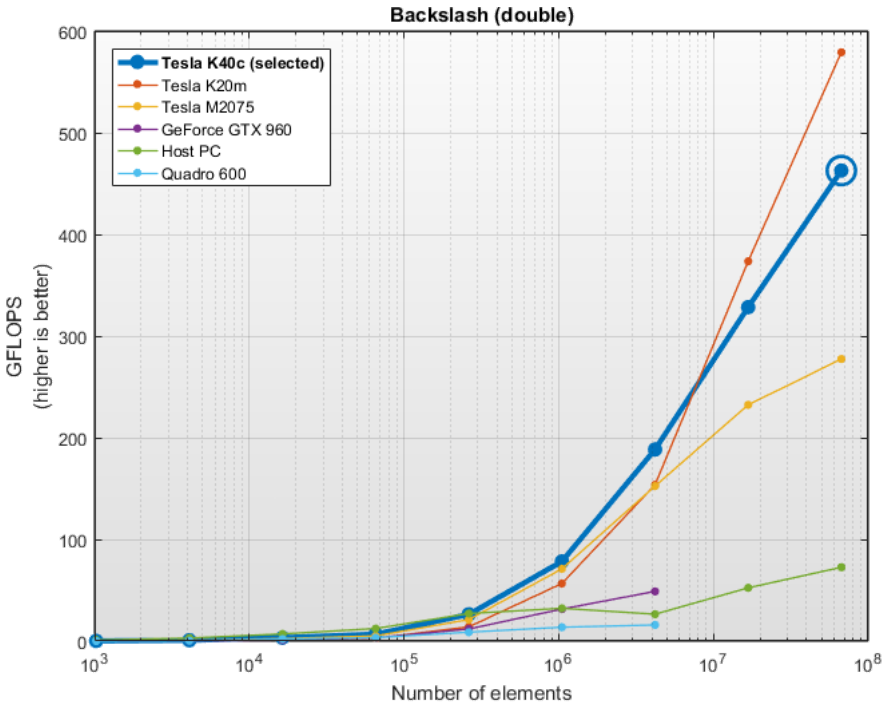

Results for FFT (double)

These results show the performance of the GPU or host PC when calculating the Fast-Fourier-Transform of a vector of complex numbers. The number operations for a vector of length `N` is assumed to be `5*N*log2(N)`.

This calculation is usually memory-bound, i.e. the performance depends mainly on how fast the GPU or host PC can read and write data.

Raw data for Tesla K40c - FFT (double)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	51,200	0.06	0.79
4,096	245,760	0.08	3.03
16,384	1,146,880	0.08	13.79
65,536	5,242,880	0.10	51.34
262,144	23,592,960	0.24	99.39
1,048,576	104,857,600	0.82	127.52
4,194,304	461,373,440	3.38	136.51
16,777,216	2,013,265,920	15.75	127.81
67,108,864	8,724,152,320	64.88	134.47

(N gigaflops = $Nx10^9$ operations per second)

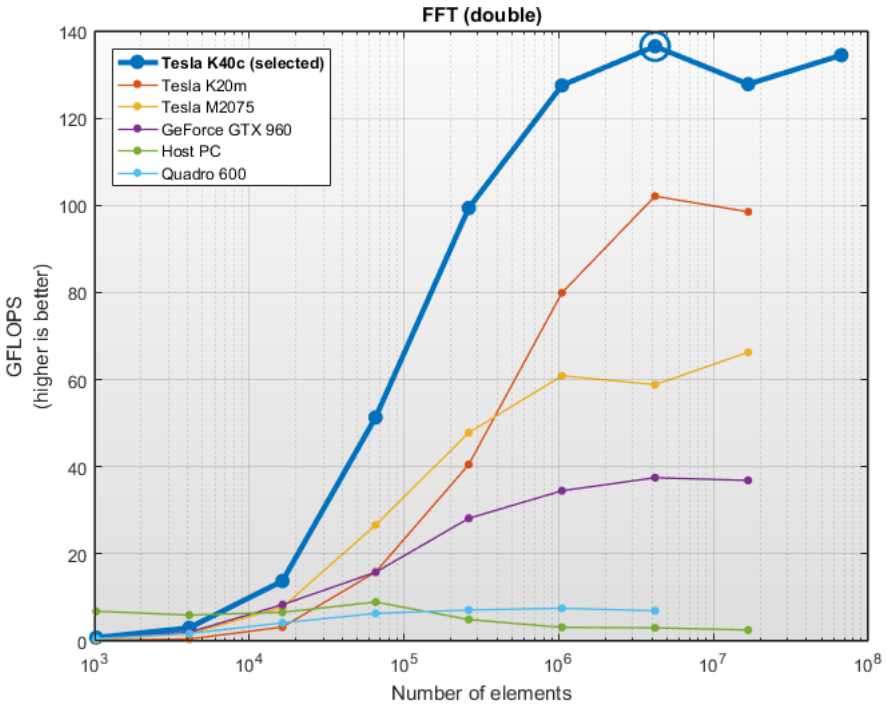

Results for MTimes (single)

These results show the performance of the GPU or host PC when calculating a matrix multiplication of two NxN real matrices. The number of operation assumed to be `2*N^3 - N^2`.

This calculation is usually compute-bound, i.e. the performance depends mainly on how fast the GPU or host PC can perform floating-point operations.

Raw data for Tesla K40c - MTimes (single)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	64,512	0.05	1.39
4,096	520,192	0.05	11.32
16,384	4,177,920	0.05	89.78

65,536	33,488,896	0.06	598.96
262,144	268,173,312	0.22	1219.43
1,048,576	2,146,435,072	1.09	1962.62
4,194,304	17,175,674,880	6.09	2819.34
16,777,216	137,422,176,256	45.54	3017.66
67,108,864	1,099,444,518,912	355.83	3089.83
268,435,456	8,795,824,586,752	2867.51	3067.41

(N gigaflops = Nx10^9 operations per second)

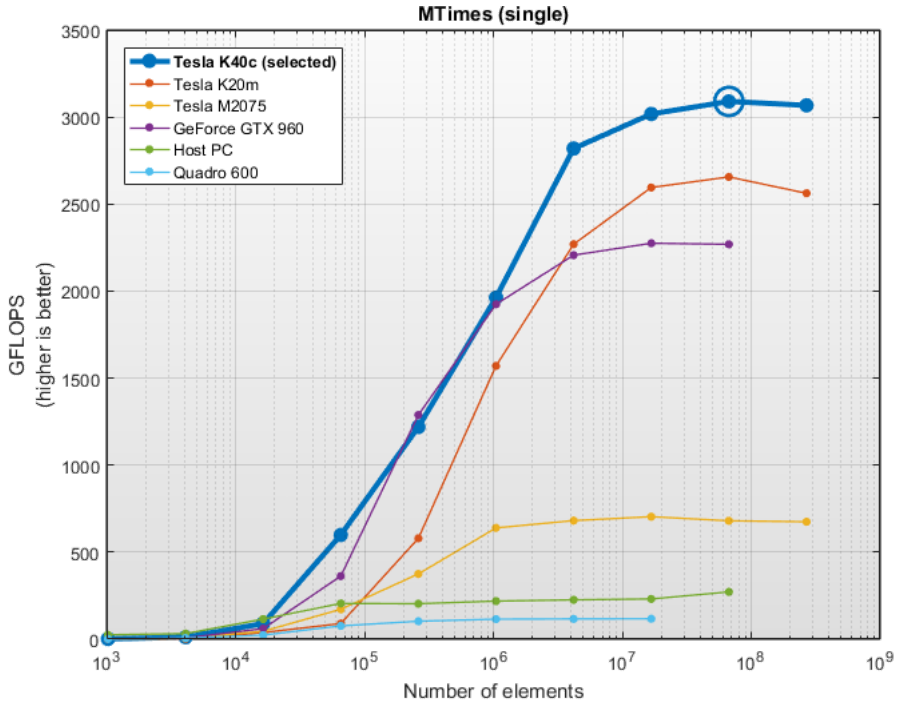

Results for Backslash (single)

These results show the performance of the GPU or host PC when calculating the matrix left division of an NxN matrix with an Nx1 vector. The number o operations is assumed to be `2/3*N^3 + 3/2*N^2`.

This calculation is usually compute-bound, i.e. the performance depends mainly on how fast the GPU or host PC can perform floating-point operations.

Raw data for Tesla K40c - Backslash (single)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	23,381	0.20	0.12
4,096	180,907	0.22	0.83
16,384	1,422,677	0.37	3.86
65,536	11,283,115	0.83	13.67
262,144	89,871,701	2.67	33.62
1,048,576	717,400,747	8.49	84.47
4,194,304	5,732,914,517	24.02	238.70
16,777,216	45,838,150,315	92.21	497.11
67,108,864	366,604,539,221	360.54	1016.83

(N gigaflops = Nx10^9 operations per second)

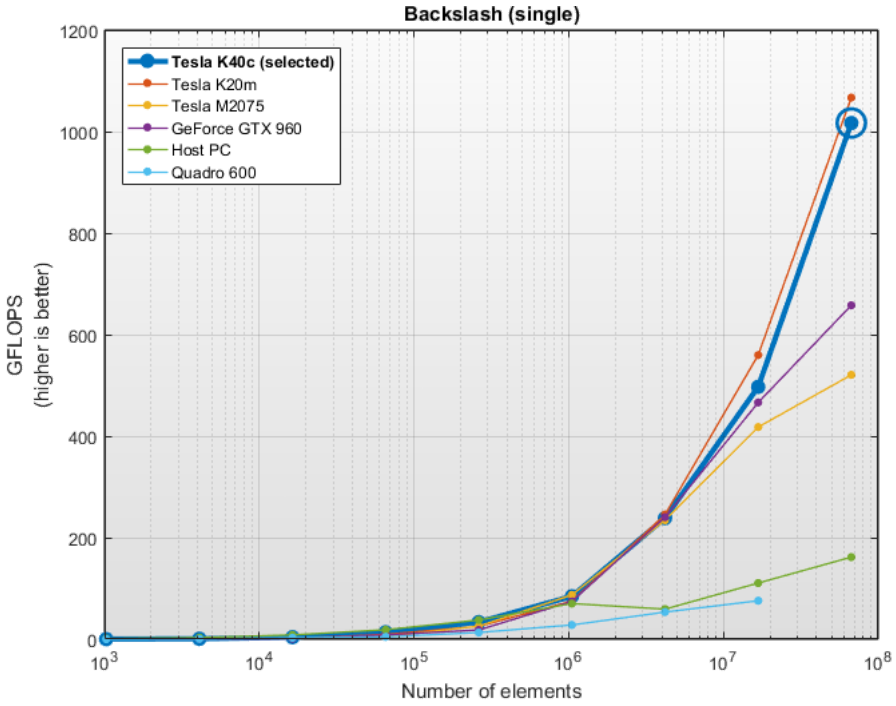

Results for FFT (single)

These results show the performance of the GPU or host PC when calculating the Fast-Fourier-Transform of a vector of complex numbers. The number operations for a vector of length N is assumed to be `5*N*log2(N)`.

This calculation is usually memory-bound, i.e. the performance depends mainly on how fast the GPU or host PC can read and write data.

Raw data for Tesla K40c - FFT (single)

Array size (elements)	Num Operations	Time (ms)	GigaFLOPS
1,024	51,200	0.10	0.53
4,096	245,760	0.06	3.87
16,384	1,146,880	0.08	14.08
65,536	5,242,880	0.08	64.07
262,144	23,592,960	0.12	193.54

1,048,576	104,857,600	0.43	244.82
4,194,304	461,373,440	1.64	281.67
16,777,216	2,013,265,920	6.69	300.87
67,108,864	8,724,152,320	31.56	276.44
268,435,456	37,580,963,840	142.55	263.64

(N gigaflops = $Nx10^9$ operations per second)

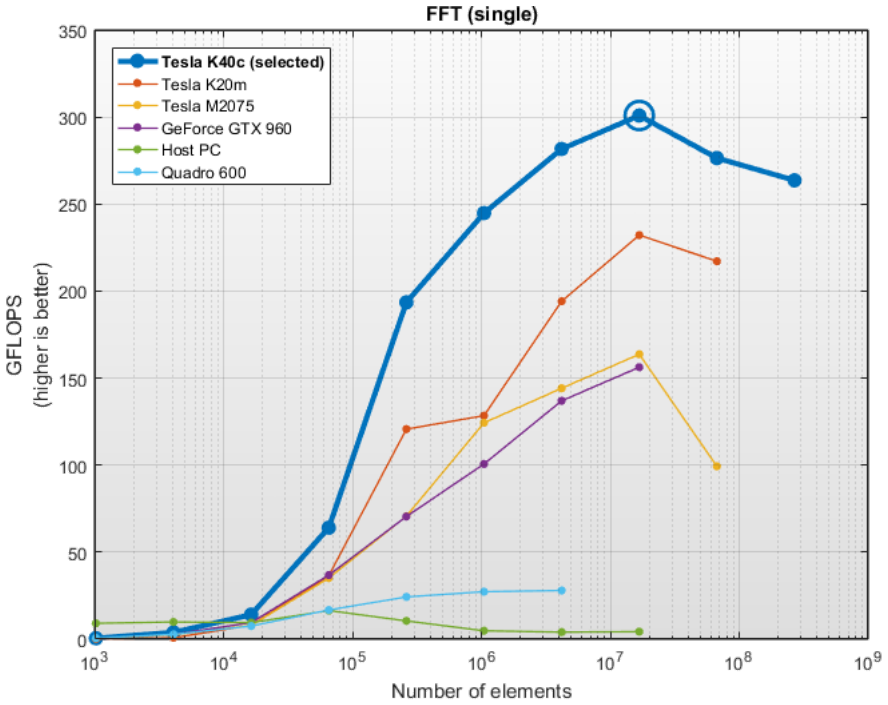

Generated by gpuBench v1.7: 2017-02-19 15:02:52

Back to su