Benchmark Testing of Embedded GPGPU Parallel Processing Systems

Furkan Kucukdemir, Ali Burak Ozakinci, Adnan Ozsoy

{furkan.kucukdemir, ali.ozakinci, adnan.ozsoy}@hacettepe.edu.tr

HACETTEPE UNIVERSITY

Introduction

Jetson TX1 is a credit-card sized module running under 10 watts with a high processing capabilities. It has 4 GB of LPDDR4 memory, 16 GB eMMC storage supported with 256 cores Maxwell GPU along with Quad ARM processor. With Jetson TX1, developers can cope with problems that has high computational needs under certain space and cost limitations. It is also suitable for applications that have limited access to an energy resource.

During our research, we tested Jetson TX1 against NVIDIA GeForce Titan X card on two benchmarks.

	Jetson TX1	GeForce Titan X
Architecture	Maxwell	Maxwell
CUDA Cores	256	3072
Memory	4 GB	12 GB
Max. Temp. (Cº)	80	91
Power	6.5-15W	Recommended:250W System Power:600W

PARALUTION

PARALUTION is a library for sparse iterative methods with CUDA support. It contains 25 applications which include parallel solvers and preconditioners. It also supports many sparse matrix representations like COO, CSR, DIA. Moreover, PARALUTION includes a benchmark application that applies some matrix operations on the matrices given as input and gives total execution time as output.

SHOC

The Scalable Heterogeneous Computing Benchmark Suite is a group of testing programs to analyze characteristics and behavior of GPUs and multicore processors. It includes 3 levels of programs :

Level 0 has bus speed, memory latency and bandwidth and peak flops tests for single and double precision.

Level 1 has Fast Fourier Transform, reduction, matrix multiplications, scan, sort tests.

Level 2 has high level programs like Quality Threshold Clustering algorithm and chemical applications.

Results

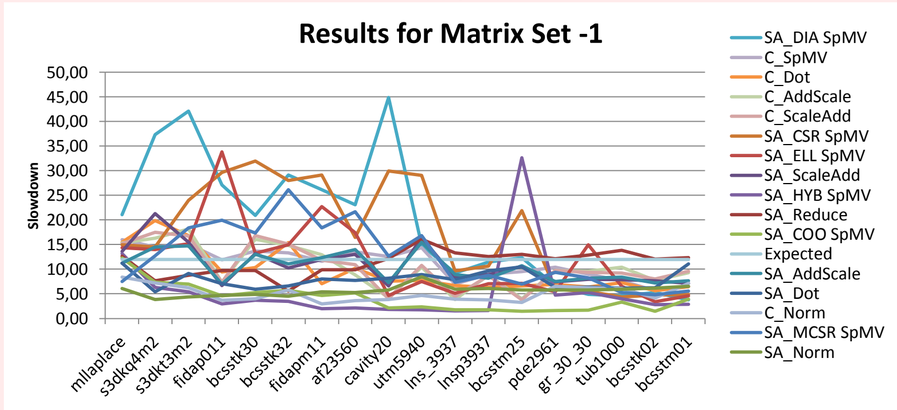

$$Slowdown = \frac{(Execution\ Time\ on\ Jetson\ TX1)}{(Execution\ Time\ on\ GeForce\ Titan\ X)}$$

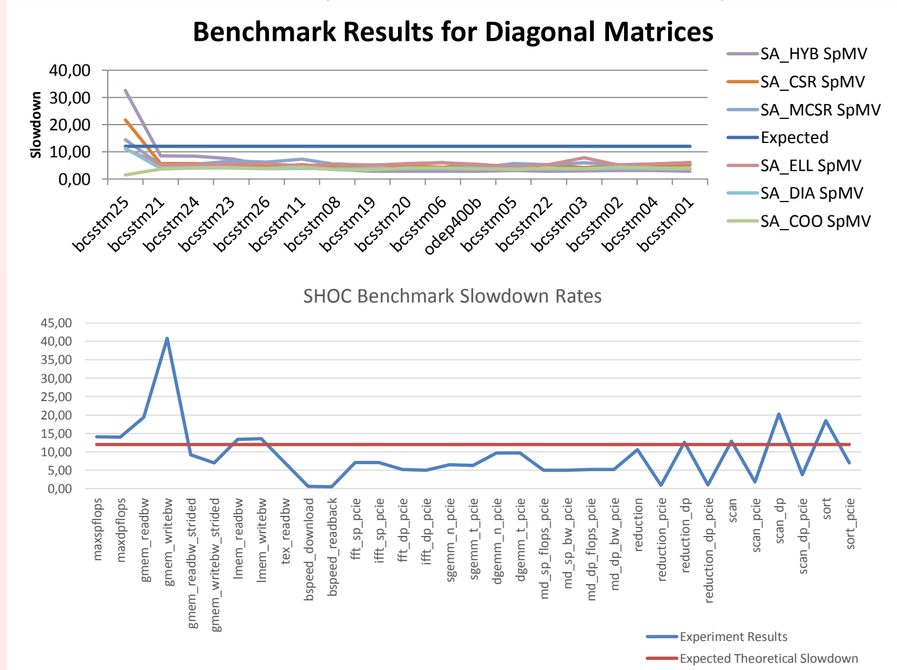

Results obtained from PARALUTION and SHOC benchmarks are shown in figures above

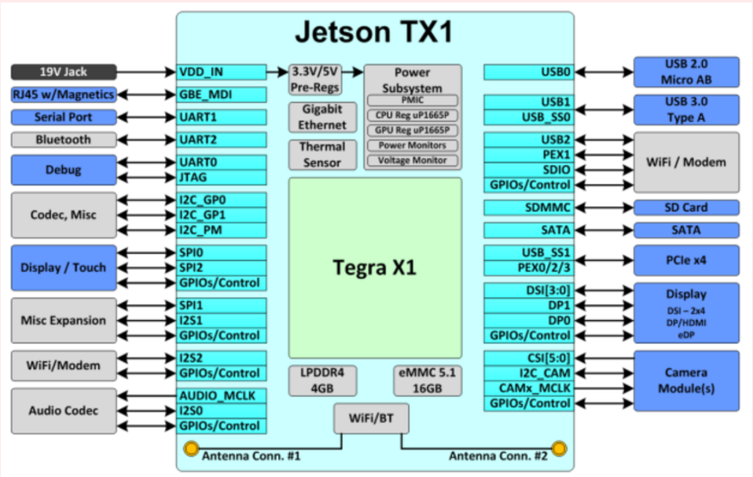

Conclusion

We experiment with a set of applications from different domains to fully understand the capabilities of the Jetson series embedded solutions. We compared Jetson TX1 results with a high-end desktop GPU, Titan X, and analyze the results. We conclude that although the Jetson TX1 is more than an order of magnitude less equipped device than Titan X in terms of the number of cores, memory, and speed, Jetson TX1 can achieve comparable performance to Titan X for certain applications and program behavior. Our main contribution in this research is identifying the strong side of Jetson TX1 device and provide a guide for researchers. Our tests showed us that Jetson series cards are not only devices for limited resource scenarios, but also capable to deliver good performance in certain applications.

Acknowledgement

This work is developed by undergraduate students under the Undergraduate Research Experience Program (DREAM) of Hacettepe University. Additionally, we would like to thank to NVIDIA GPU Grant Program for the Jetson TX1 hardware donation award.

DREAM

NVIDIA®