ARM Cortex-A9 performance in HPEC applications

M. Boyd [1], C. Della Silva [2], K. Keville [3]
[1] Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA
mboyd@mit.edu
[2] Department of Mechanical Engineering, MIT, Cambridge, MA
clarkds@mit.edu
[3] Institute for Soldier Nanotechnologies, MIT, Cambridge, MA
kkeville@mit.edu

Abstract
The ARM architecture has been garnering a considerable amount of interest as a potential competitor in the HPC space that has been traditional dominated by servers built around Intel EM64T or AMD x86_64 based processors. These 2 processor families currently represent 90% of the Top500 list. Ongoing investigation into improving the performance of this processor indicate that this approach will be a legitimate contender in the energy-efficient HPC market in coming years and will score favorably on metrics used by the Green Grid and the Green 500 list.

Introduction
The rapid introduction of the ARM architecture into the consumer electronics market, mostly through the wide adoption of smartphones, has lead a number of researchers to investigate their value as a processor in HPC or HPEC applications. The Apple iPad 2 has been looked at for it's future utility in cluster computing by the researchers at the Innovative Computing Lab at the University of Tennessee[1]. While there are some promising aspects to it, it should clearly be looked at as a reference platform for a future implementation that has the embedded processor less the other hardware that is substantially placed there for applications important to smartphone and tablet users; location-based apps or improved user interface, for instance, things that are not of much importance to HPC users. Fortunately, there is considerable innovation in the systems integration ecosystem surrounding these offerings with many manufacturers freely revealing their reference designs, usually in the form of Gerber files and schematics.

ARM improvements
The shortcomings of ARM from an HPC perspective are well-known. They are designed primarily for integer or single precision floating point work. Clearly, the path to achieve better performance on the new classes of ARM SoCs is to leverage the graphics chipsets that many ship with these days. Popular additions are the SGX Series 5 (PowerVR), and the Mali 600 family. By August, in advance of the final paper publication deadline, we will have performance numbers on the current representative crop of ARM Cortex-A9 processors and will report their viability as HPC processors. We may also have initial data on ARM Cortex-A15 processors, which have been estimated to deliver 210 GFLOPS[2].

Prior Art
There has been a considerable amount of study on earlier generation ARM technologies to see what their utility as an embedded HPC processor might be. Prior models that have seen entrance into the consumer electronics world have generally been built around the ARM Cortex-A8 core. Performance numbers have been quite weak here. The MNM Team at LMU in Germany has demonstrated LINPACK scores of 57.5 MFLOPS in SP[3].

Performance under Android
There has been a number of attempts to harness the compute power of smartphones and tablets running the Android OS. These attempts have been stifled by Android's current inability to unlock the power of multicore processors. That having been said, there is a robust community using smartphones for peripherally HPC related applications. There is a port of LINPACK which uses the Dalvik Java VM to report scores on a variety of Android based phones. Scores of 60 MFLOPS are not unusual in the list published by the firm who has released the application (Green Computing). Additionally, there are ways to get incrementally better performance out of a processor by getting more operations per cycle that, in the aggregate, will reduce overall cycle use for a given process. These approaches generally fall under the rubric of “iterative refinement” and will be used where a IEEE 754 standard precision implementation is required for certification of the benchmarks.

Practical Application
Research into this area is primarily facilitated by the knowledge that many applications in HPC today need not be performed at native double precision. In fact, many float or fixed precision algorithms are used as sieves to rapidly determine the viability of certain approaches. If the research trends or converges to an unwanted result, then the simulation is terminated and another is initiated with the understanding that only promising approaches will be rerun at higher precision, possibly being migrated over to a traditional HPC cluster.

MIT ARM assessment effort
Published numbers on embedded ARM and GPU performance will be verified and validated. Currently there is wide agreement on the order of magnitude for performance but little agreement after that. We intend to perform a proper apples-to-apples comparison using well-known benchmark packages. Mini-clusters of 16 – 32 cores will be constructed for study of MPI-enabled applications and benchmarks.

Unverified performance numbers under review this

summer;

Tegra 2 - 4.8 GFLOPS (8, 1-way ALUs @ ~300MHz) PowerVR SGX543MP2 - 19.2 GFLOPS (8, 4-way ALUs @ ~300MHz)

TI Pandaboard EA1 – 610 MFLOPS (single core)

ZTSystems R1801e - TBD

Calxeda - TBD

Samsung Exynos - TBD

Nufront NuSmart2816 - TBD

References

[1] http://bits.blogs.nytimes.com/2011/05/09/the-ipad-in-your-hand-as-fast-as-a-supercomputer-of-yore

[2] http://www.stericsson.com/press_releases/NovaThor.jsp

[3] http://www.nm.ifi.lmu.de/projects/ATV2CLUSTER/