# Towards Fault-Tolerant Energy-Efficient High Performance Computing in the Cloud

Kurt Keville, Rohan Garg, David J. Yates, Kapil Arya, and Gene Cooperman

MIT and Bentley U. and Northeastern U., Boston, USA

## Abstract

We examine the use of ARM-based clusters for low-power, high performance computing. This work examines two likely use-modes: (i) a standard dedicated cluster; and (ii) a cluster of pre-configured virtual machines in the cloud. A 40-node department-level cluster based on an ARM Cortex-A9 is compared against a similar cluster based on an Intel Core 2 Duo, in contrast to a recent similar study on just a 4-node cluster. For the NAS benchmarks on 32-node clusters, ARM was found to have a power efficiency ranging from 1.3 to 6.2 times greater than that of Intel. This is despite Intel's approximately five times greater performance. The particular efficiency ratio depends primarily on the size of the working set relative to L2 cache. In addition to energy-efficient computing, this study also emphasizes fault tolerance: an important ingredient in high performance computing. It relies on two recent extensions to the DMTCP checkpoint-restart package. DMTCP was extended (i) to support ARM CPUs, and (ii) to support checkpointing of the Qemu virtual machine in user-mode. DMTCP is used both to checkpoint native distributed applications, and to checkpoint a network of virtual machines. This latter case demonstrates the ability to deploy pre-configured software in virtual machines hosted in the cloud, and further to migrate cluster computation between hosts in the cloud.

## ARM-based cluster with 40 nodes, with chair

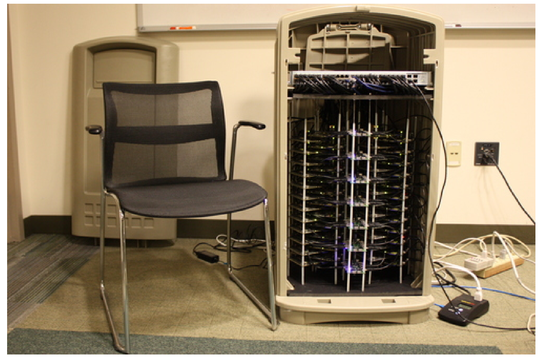

## DMTCP Checkpoint-Restart Architecture

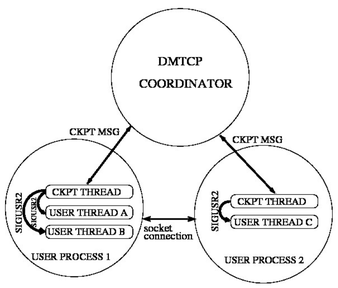


Figure: Architecture of DMTCP

The three core commands of DMTCP are:

```
dmtcp_command a.out
dmtcp_command --checkpoint
dmtcp_restart ckpt_a.out_*.dmtcp
```

## NAS Parallel Benchmarks: ARM vs Intel

| Benchmark | ARM | | Intel | | Ratio |
|---|---|---|---|---|---|
| | Time (s) | Energy (KJ) | Time (s) | Energy (KJ) | (ARM/Intel) |
| NAS/CG.A | 7.79 | 1.18 | 3.27 | 7.35 | 6.2 |
| NAS/EP.A | 13.00 | 1.97 | 1.12 | 2.52 | 1.3 |
| NAS/FT.A | 13.25 | 2.01 | 3.61 | 8.12 | 4.0 |
| NAS/IS.A | 3.37 | 0.51 | 1.16 | 2.61 | 5.1 |
| NAS/LU.A | 138.98 | 21.08 | 12.03 | 27.04 | 1.3 |
| NAS/MG.A | 4.50 | 0.68 | 0.68 | 1.53 | 2.3 |

Table: Native applications (no VM). Clusters for ARM and Intel both contain 32 nodes. The ARM/Intel ratio compares the performance-to-energy ratios for the two architectures. The energy is the reading from a power meter, integrated over the time of the benchmark for the 32 nodes. An ARM/Intel ratio greater than 1.0 favors ARM.

Table 1, above, presents results for a variety of NAS-MPI benchmarks. The tests show a performance to power advantage for ARM over Intel that ranges from a ratio of 1.3 to 6.2. The ratios correlate well with whether the working set fits in L2 cache. The Intel CPU has 6 MB of cache, while ARM has only 1 MB. CG (conjugate gradient: eigenvalue of a single matrix) and IS (integer sort using bucket sort) have relatively small working sets, while LU matrix factorization (**64** $\times$ **64** $\times$ **64** for LU.A), and MG (MultiGrid; 3-D discrete Poisson equation) have larger working sets due to the problems being of higher dimension. The EP benchmark (embarassingly parallel) seems to be a special case due to its extensive use of square roots and logarithms. This takes advantage of Intel's CPU hardware support.

## Scalability: From 1 to 40 nodes

Table 2, below, demonstrates the scalability of the ARM and Intel clusters running a native application as the number of nodes in the cluster varies. The NAS/LU.A benchmark used in Table 2 was chosen for scalability tests. NAS/LU.A was conservatively chosen because it favored the Intel CPU the most (aside from the EP special case). The energy efficiency ratios in Table 2 immediately highlight the two regimes of one to four nodes, and eight or more nodes. One to four nodes favor Intel due to the heavy pressure on the CPU cache, which is only 1 MB in the case of ARM. For eight nodes or more, network communication stress the computation to a greater extent, thus bringing the performance of ARM and Intel closer together.. We stress again that LU was chosen to be conservative. The choice of NAS/CG or NAS/IS would very strongly favor ARM, as already seen by the large ratios in Table 1.

| Num. nodes | ARM | | Intel | | Ratio |
|---|---|---|---|---|---|
| | Time (s) | Energy (KJ) | Time (s) | Energy (KJ) | (ARM/Intel) |
| NAS-MPI/LU.A | | | | | |
| 1 | 949.11 | 10.7 | 125.7 | 8.8 | 0.86 |
| 2 | 502.72 | 10.4 | 63.21 | 8.9 | 0.85 |
| 4 | 380.5 | 12.30 | 39.93 | 11.2 | 0.91 |
| 8 | 341.5 | 15.25 | 31.79 | 17.9 | 1.17 |
| 16 | 183.8 | 16.41 | 16.11 | 18.1 | 1.10 |
| 32 | 139.2 | 21.39 | 12.03 | 27.0 | 1.26 |
| 40 | 120.0 | 22.60 | 10.12 | 28.4 | 1.25 |

Table: Native application (no VM). Scalability is demonstrated as the number of nodes in a cluster varies. The NAS/LU.A benchmark is used. The ratio in the last column is defined as in Table 1.

## Scalability for Checkpointing of NAS/LU.A benchmark

| Num. nodes | ARM | Intel |
|---|---|---|
| | Ckpt (s) | Ckpt (s) |
| NAS-MPI/LU.A | | |
| 1 | 8.2 | 2.65 |
| 2 | 8.4 | 2.95 |
| 4 | 8.2 | 3.13 |
| 8 | 9.6 | 3.28 |
| 16 | 12.0 | 3.9 |
| 32 | 20.1 | 4.68 |
| 40 | 26.7 | 9.36 |

Table: Native application (no VM). Scalability for checkpointing is demonstrated as the number of nodes in a cluster varies. Checkpoints are taken to /tmp, residing on a local disk (Intel) or a local SD card (ARM). For the 40-node Intel case, it is conjectured that one node was particularly slow.

Although writing to local disk (Intel) or an SD card (ARM) can contribute to the total time, operating systems hide this time by buffering disk writes. Further, DMTCP dynamically invokes gzip for compression. So, we believe the times to be limited largely by speed of RAM. The ARM clock speed is three times slower than that of Intel, which is consistent with checkpoint times for ARM being up to three times slower than Intel. Nevertheless the ARM CPU consumes less than a tenth the energy of the Intel CPU. This gives ARM an overall energy efficiency advantage.

## Scalability for Checkpointing of Qemu virtual machine

| Num. nodes | ARM | | Intel | | Ratio |
|---|---|---|---|---|---|
| | Time (s) | Energy (KJ) | Time (s) | Energy (KJ) | (ARM/Intel) |
| NAS-MPI/LU.A | | | | | |
| 1 | 19.3 | 0.22 | 13.3 | 0.93 | 4.2 |
| 2 | 22.3 | 0.46 | 16.6 | 2.32 | 5.0 |
| 4 | 22.2 | 0.65 | 20.8 | 5.84 | 9.0 |
| 8 | 24.1 | 1.08 | 21.3 | 11.97 | 11.1 |
| 16 | 29.9 | 2.67 | 21.8 | 24.50 | 9.2 |
| 32 | 35.2 | 5.34 | 25.2 | 56.64 | 10.6 |
| 40 | 60.2 | 11.34 | 40.8 | 113.65 | 10.0 |

Table: Checkpointing efficiency on ARM versus Intel using the Qemu VM. Qemu is run in user-mode. Clusters on ARM and Intel both contain up to 40 nodes. The ARM/Intel ratio is defined as in Table 1.

In Table 4, the times for checkpointing a virtual machine are now much closer between ARM and Intel. We believe this is because the memory footprint of a virtual machine is too large for the operating system to buffer most writes. Since checkpointing in both clusters is limited by disk-write speed, Intel loses much of its speed advantage, thereby yielding energy efficiency ratios heavily in favor of ARM.

## Conclusions

The 1 GHz dual-core ARM Cortex-A9 demonstrated an energy efficiency ratio 1.3 to 6.2 times greater than a 3 GHz Intel Core 2 Duo in 32-node cluster configurations. ARM was more energy efficient despite Intel having approximately five times greater performance. For those benchmarks with the largest working sets and smallest cluster sizes (just one, two, or four nodes), the ratio fell as low as 0.9, since this stressed primarily the CPU cache and not the network. The benchmarks favored ARM especially when the benchmark had a small working set, since the ARM CPU had only a 1 MB L2 cache, as opposed to Intel's 6 MB cache. A large cluster size also favored ARM, since the problem data could be distributed across the network. This stresses the network, but not the CPU cache. As the current ARM Cortex-A9 generation is replaced by by the 32-bit ARM Cortex-A15 and then the future 64-bit ARMv8, this may further bias the ratios toward ARM. In particular, AppliedMicro projects in their roadmap a 3 GHz ARMv8 CPU consuming 2 Watts per core. The comparison here using the Cortex-A9 provides designers of energy-efficient systems with initial estimates of energy consumption on realistic benchmarks for department-level clusters, or in the cloud.