Applications – miniFE (Study Description)

- Parametric study focused on assessing miniFE weak scaling performance for ”small” and “big” dataset in quad/flat and quad/cache modes.
  - The “small” dataset (i.e., 400^3) was sized to fit within KNL MCDRAM.
- Comparisons are made between Mutrino, clogin80, ft-cdl, and current Trinity 8X Acceptance.
  - miniFE version 2.0.1-openmp-opt was built with Intel 16.0.3.

miniFE parameters

Mem. Footprint	Nodes	nx	ny	nz
big-76.2GB/node	2	718	718	718
big-76.2GB/node	8	1,138	1,138	1,138
big-76.2GB/node	32	1,805	1,805	1,805
big-76.2GB/node	128	2,865	2,865	2,865
big-76.2GB/node	512	4,547	4,547	4,547
small-13.3GB/node	2	400	400	400
small-13.3GB/node	8	800	800	400
small-13.3GB/node	32	1,600	800	800
small-13.3GB/node	128	1,600	1,600	1,600
small-13.3GB/node	512	3,200	3,200	1,600

Applications – miniFE (Build/Run Params.)

miniFE build environment & run parameters (below for quad/flat)

```
#PBS -l nodes=$PV_NNODES:ppn=8

module swap PrgEnv-cray PrgEnv-intel
module load craype-mic-knl

nodes=$PV_NNODES
cores=$PV_PESPERNODE
depth=16
export OMP_NUM_THREADS=${depth}
export MPICH_USE_DMAPP_COLL=1
export MPICH_DMAPP_HW_CE=1

aprun -n $(($cores*$nodes)) -N $cores -d $depth -j 2 -r 2 -cc depth numactl --membind=1 \
./miniFE.x_amagela_ft-cdl $PV_TEST_ARGS | tee "${PV_WS}/miniFE_output.log"
```

miniFE `Makefile` parameters

```
MINIFE_TYPES = -DMINIFE_SCALAR=double -DMINIFE_LOCAL_ORDINAL=int -DMINIFE_GLOBAL_ORDINAL=long

MINIFE_MATRIX_TYPE = -DMINIFE_CSR_MATRIX

CFLAGS = -O3 -qopenmp -restrict -qopt-report=3 -fast -fp-model fast=2 -no-prec-div -ansi-alias

CPPFLAGS = -I. -I../utils -I../fem $(MINIFE_TYPES) $(MINIFE_MATRIX_TYPE) \
-DMINIFE_RESTRICT=restrict -DDEBUG_DO_DETAILED_TIMING -DMINIFE_INFO=1 -DMINIFE_KERNELS=0 \
-DHAVE_MPI
```

Applications – miniFE (Calibration)

- The current 8X Acceptance SSP miniFE benchmark ratio between KNL and HSW is 2.57, which is in line with our own findings comparing clogin80 with Mutrino.
  - We expected similar, not exact, performance.
  - KNL utilized quad/flat mode forcing MCDRAM usage.
  - The figure of merit (FOM) used is the conjugate gradient solver's floating point operations per second per node.

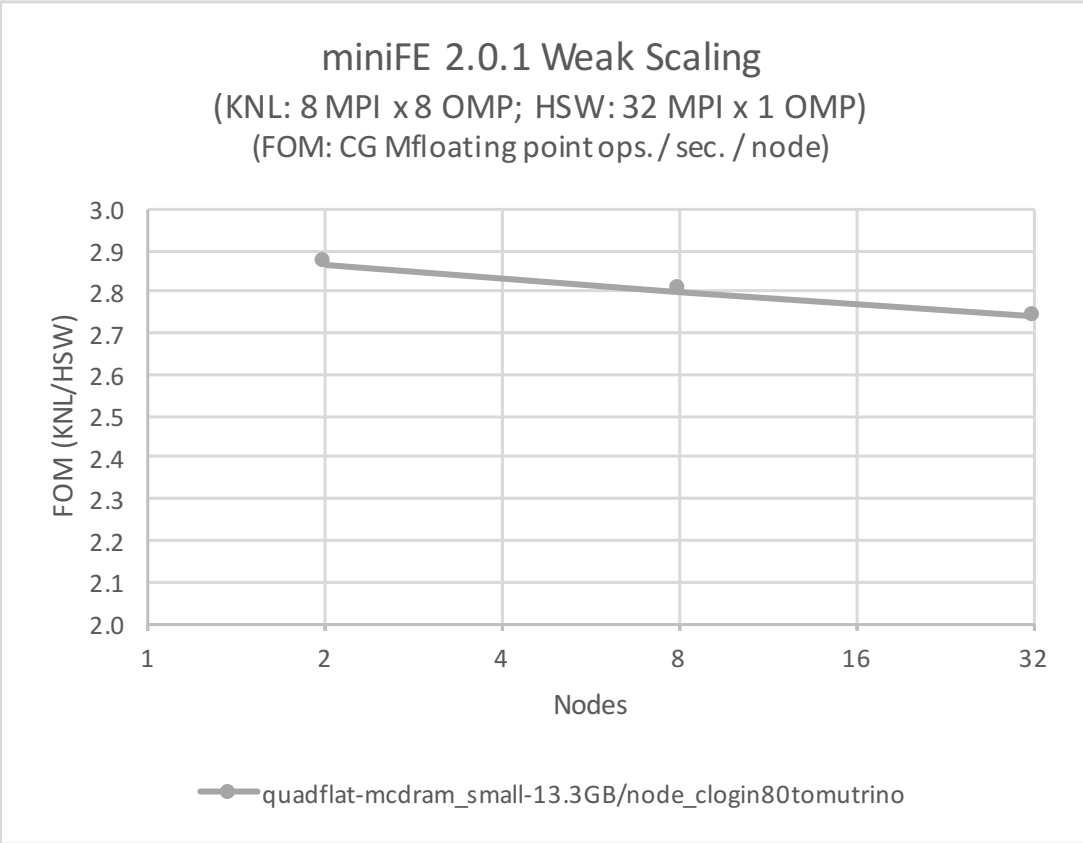

miniFE comparison between clogin80 (quad/flat, `numactl --membind=1`) and Mutrino

Applications – miniFE (quad/flat vs. quad/cache)

- *Points shown on plots are the **average** and that point's **min**, **max**, and **number of points in distribution** are also indicated.*
- Large variability is observed at high node counts for small problem.
  - The max case at 512 nodes carries the expected performance and outpaces its 8-, 32-, and 128-node counterparts.
- Overall, variability increases with node count.

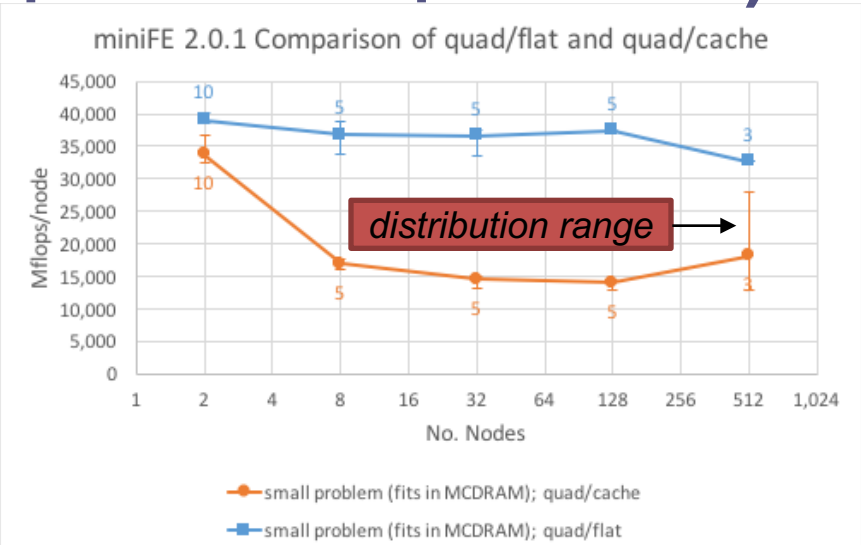

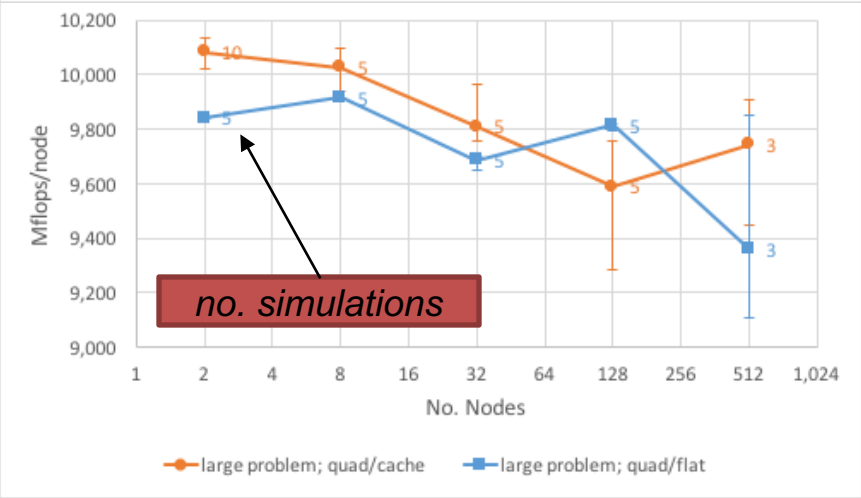

Applications – miniFE (2-Node Memory Sweep)

- *This plot compares the performance between quad/cache and quad/flat with varying miniFE dimension size (which changes the size of the problem).*
- Variability is present once the mean quad/cache and quad/flat performance is equal.
- The quad/cache performance degrades as the size of a problem approaches the MCDRAM capacity.

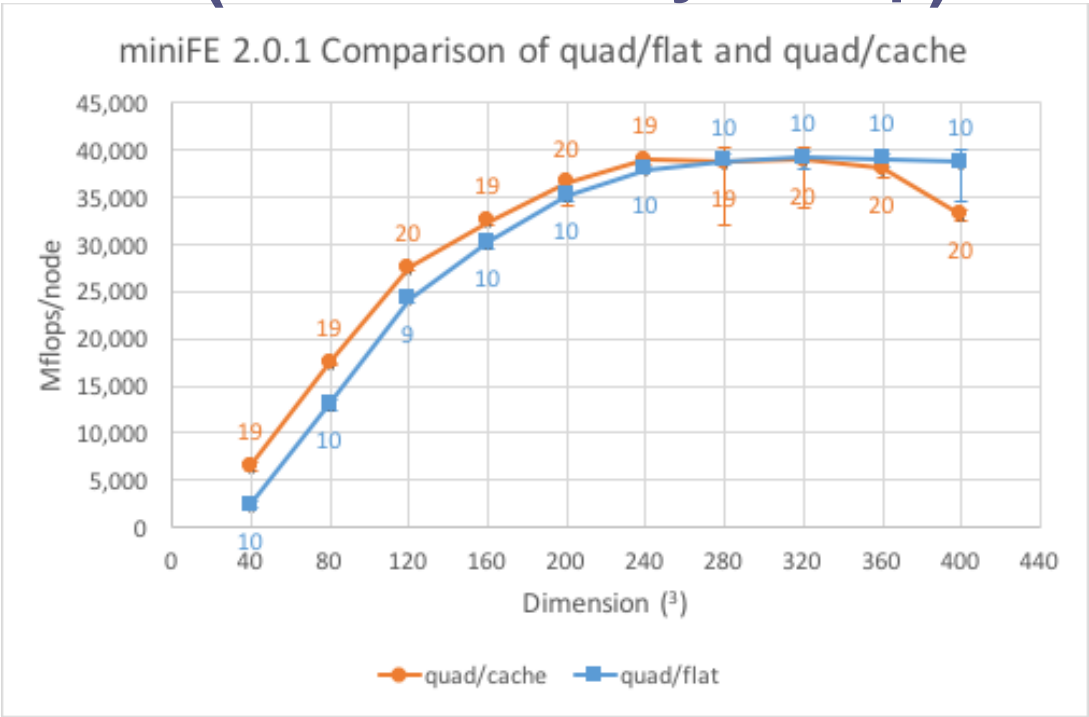

Applications – miniFE (Pstate Comparison)

- *The plots compare quad/cache (with and without pstate activated) and quad/flat (with and without pstate activated) for the **small** problem.*
- Activating pstate for quad/cache typically increases variability, ***however*** these cases' maximum is closer to expected performance.
  - The mean of the cases with pstate enabled had greater performance than with it disabled.
- The quad/flat performance and variability characteristics with and without pstate fall in line with expectations.

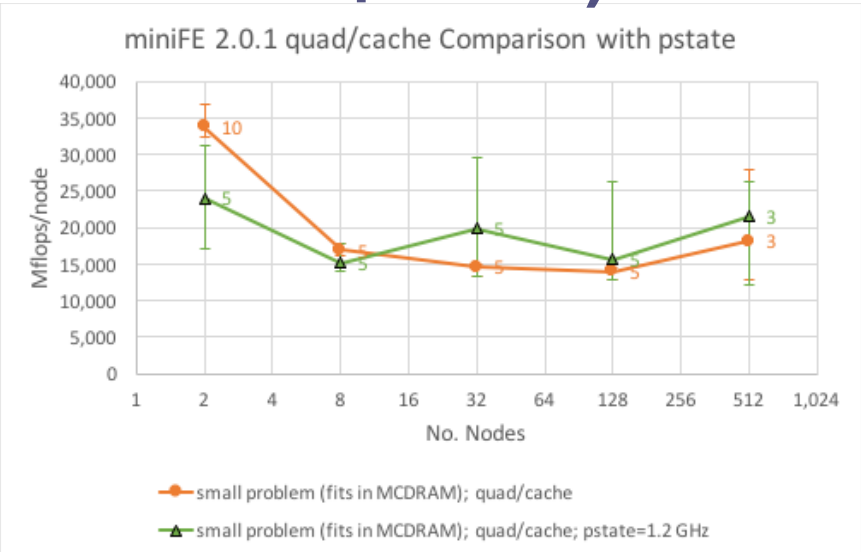

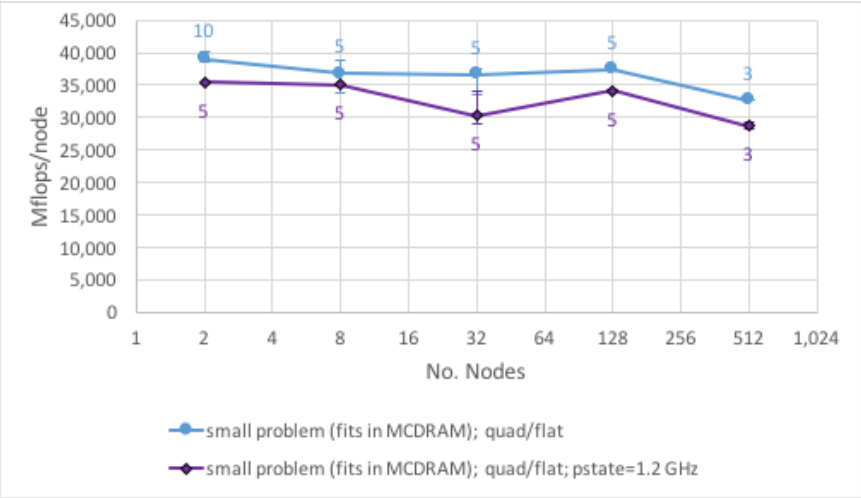

Applications – miniFE (Pstate Comparison)

- *The plots compare quad/cache (with and without pstate activated) and quad/flat (with and without pstate activated) for the **large** problem.*
- Little-to-no variability was observed for node counts less than 512 with and without pstate activated.
- Large variability was observed with pstate activated at 512 nodes for both quad/cache and quad/flat.
  - This is not due to direct-map cache thrashing.

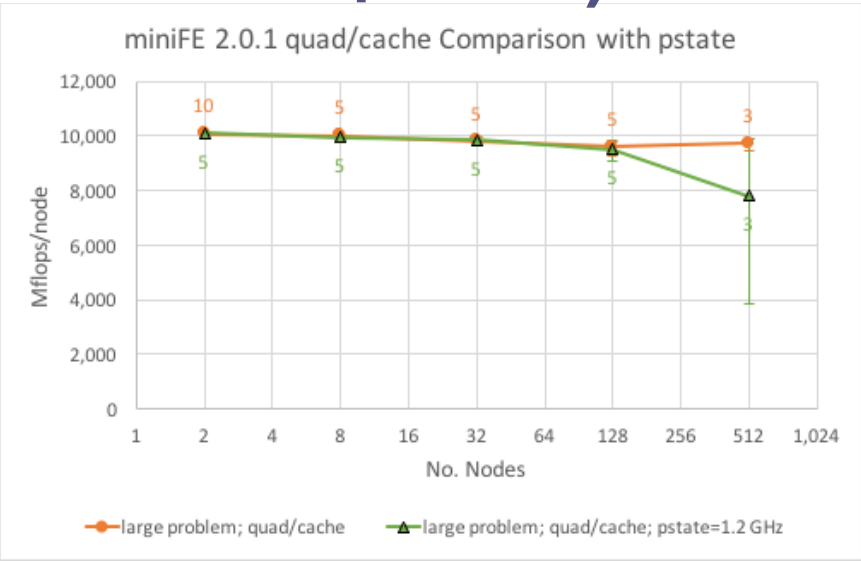

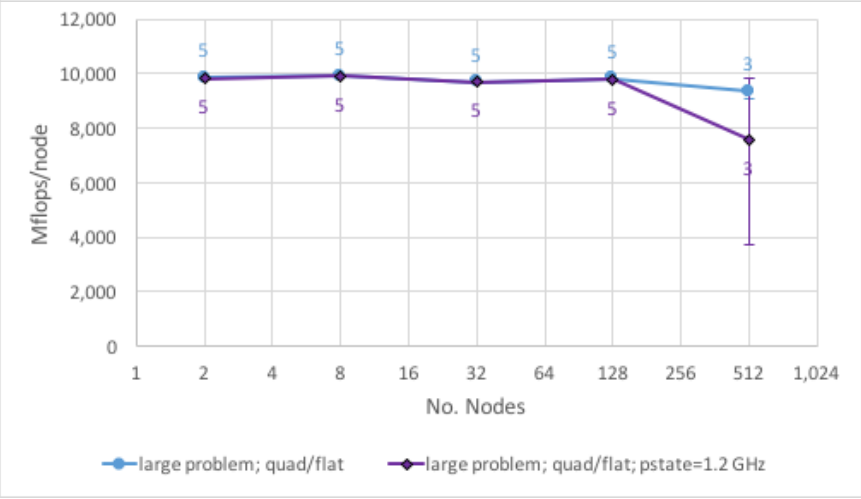

Applications – miniFE (conclusions)

- Performance of quad/flat meets expectations.
  - Variability is low.
- Performance of quad/cache running large problems meets expectations.
- Performance of quad/cache running problems designed to fit in MCDRAM while utilizing a large percentage of it has a lot of variability and, sometimes, meets expectations.
  - The cause of this is still being investigated and its behavior has also been observed on Avalon.
- I would like to acknowledge the assistance from the whole Factory Test Team. Mike Davis has impacted this testing with his debugging, suggestions, and double-checking.