Programming Texas Instruments Keystone SoCs using OpenMP

OpenMP BoF, SC'14

Eric Stotzer

High Performance Embedded Computing

DVR / NVR & smart camera

Networking

Mission critical systems

Medical imaging

Video and audio infrastructure

High-performance and cloud computing

Portable mobile radio

Industrial imaging

Home AVR and automotive audio

Analytics

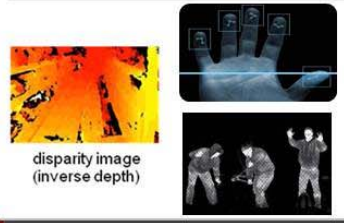

Wireless testers

Industrial control

media processing | *computing* | *radar & communications* | *industrial electronics*

TI Keystone Architecture

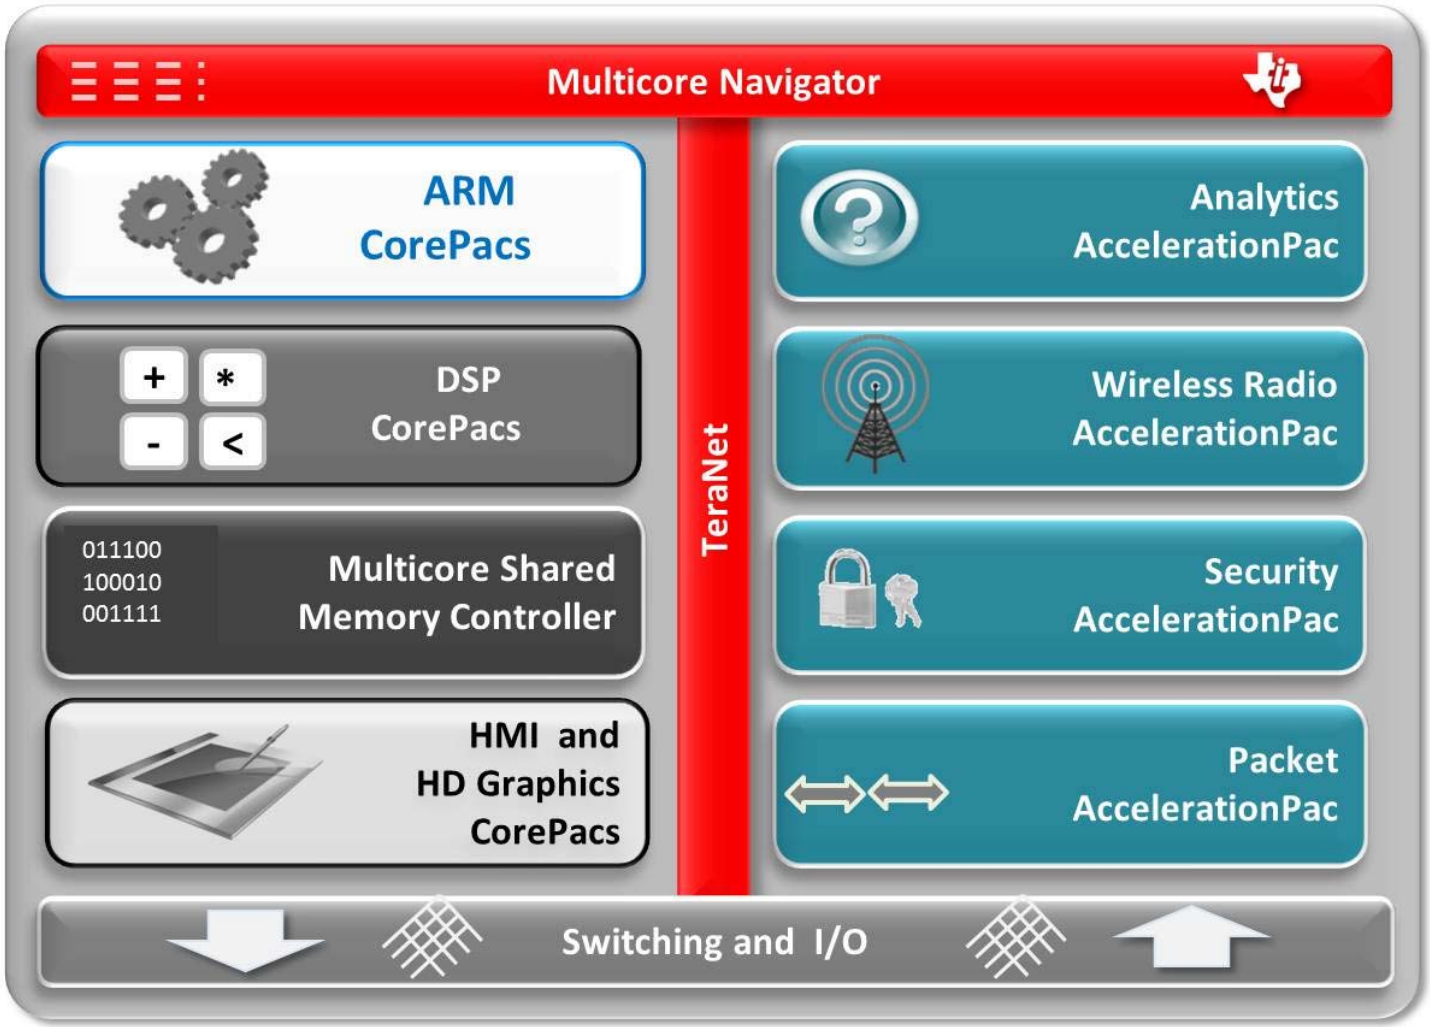

Keystone I: C6678 SoC

- Eight 8 C66x cores
- Each with 32k L1P, 32k L1D, 512k L2
- 1 to 1.25 GHz
- 320 GMACS
- 160 SP GFLOPS
- 512 KB/Core of local L2
- 4MB Multicore Shared Memory (MSMC)
- Multicore Navigator (8k HW queues) and TeraNet
- Serial-RapidIO, PCIe-II, Ethernet, 1xHyperlink

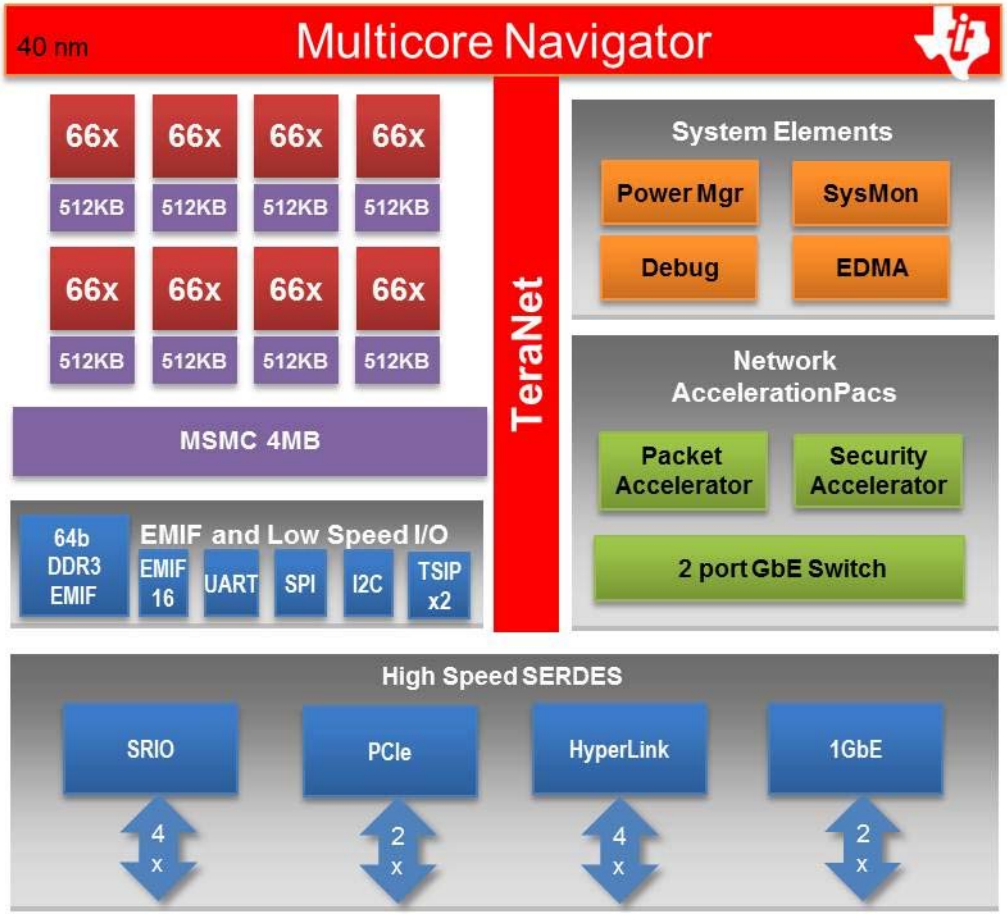

24mm x 24mm package

Energy Efficiency

LINPACK running on C6678 achieves 25.6 Gflops, ~2.1 Gflops/W

PRACE First Implementation Project, Grant RI-261557, Final Report on Prototypes Evaluation. Lennart Johnsson, Gilbert Netzer, SNIC/KTH, 3/29/2013.

Linpack Power Profile

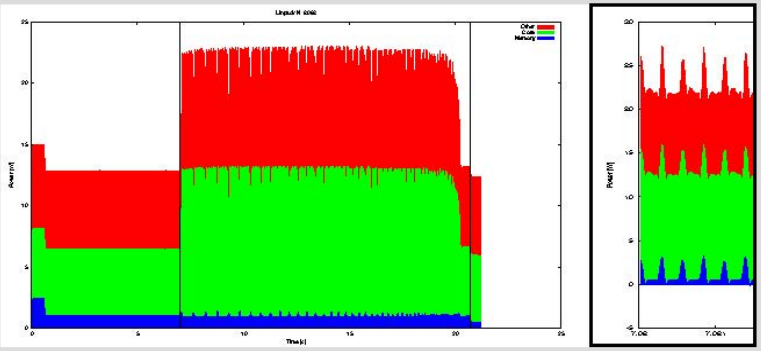

The plot shows the power consumption over time during a single execution of the Linpack benchmark code. Blue shows memory power, green is added power fed to the DSP and red other module consumers stacked atop. The vertical lines denote the timed section of the code. Distinct phases of execution can be seen, for instance the serial back-substitution at the end of the run. A zoom in also reveals the power peaks caused by DMA block copies to and from the main memory.

DSP Linpack Energy Efficiency

Size	Perf.	Eff.	Core	Mem.	Other	Total	Core + Mem	Tot
	GF/s	*%*	*W*	*W*	*W*	*W*	*MF/J*	*MF/J*
127	1.3	4	5.95	1.26	6.87	14.08	176	90
255	2.8	9	4.78	0.99	5.17	10.95	493	260
511	6.0	19	6.40	1.12	6.58	14.09	796	425
1023	11.3	35	8.02	1.19	7.65	16.86	1230	672
2047	16.9	53	9.16	1.10	8.13	18.40	1649	920
4095	22.0	69	10.30	1.03	8.70	20.03	1939	1097
8063	25.6	80	11.20	0.99	9.20	21.39	2097	1195

The table shows the power and energy consumption of the major components of the C6678 DSP EVM. Core refers to the C6678 DSP SoC excluding I/O power. Mem is the DDR3 memory subsytem. The 5-9 watts of "Other" power is to a large part, except for about 1.5 W DC converter losses, consumed by debugging and unused hardware features that would not be present in an HPC server node. Therefore the "Core + Mem" power is a good estimate for the energy efficiency of an HPC server node. The values for small problem sizes show deviations due to various parts of the benchmark not being executed. The outmost loop step size is 128 columns, another breakpoint occurs at 1024.

High Density COTS boards

DSPC-8681 ½ length PCIe card - 54Watts

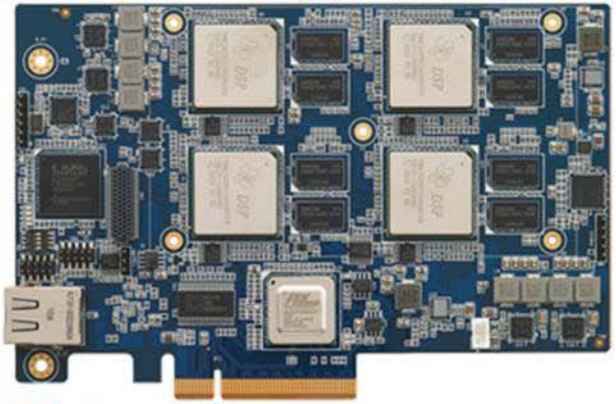

DSPC-8682 PCIe Full-Length Card - 110Watts

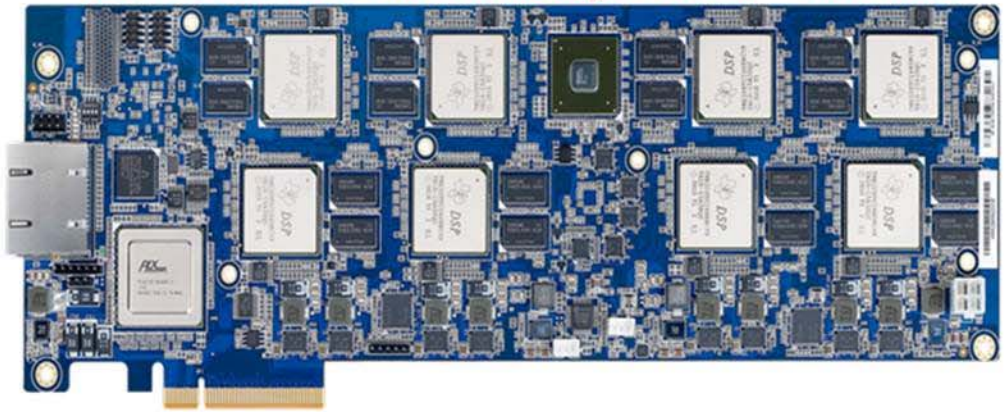

DSPC-8682 ATCA blade 350Watts

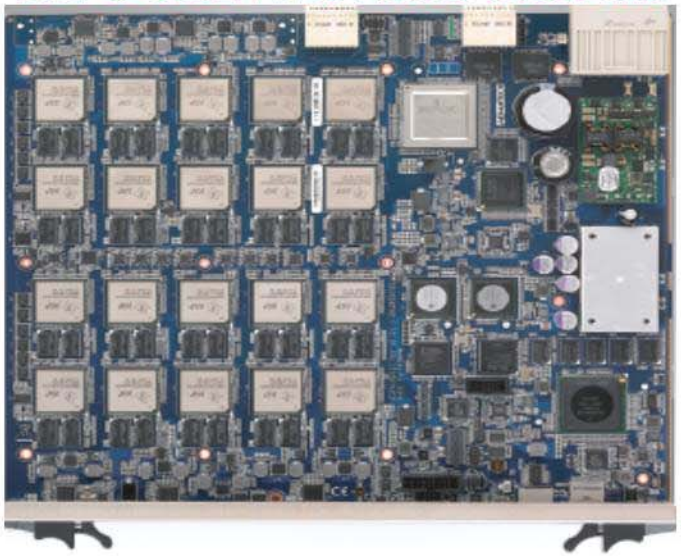

Keystone II: 66AK2H12/06 SoC

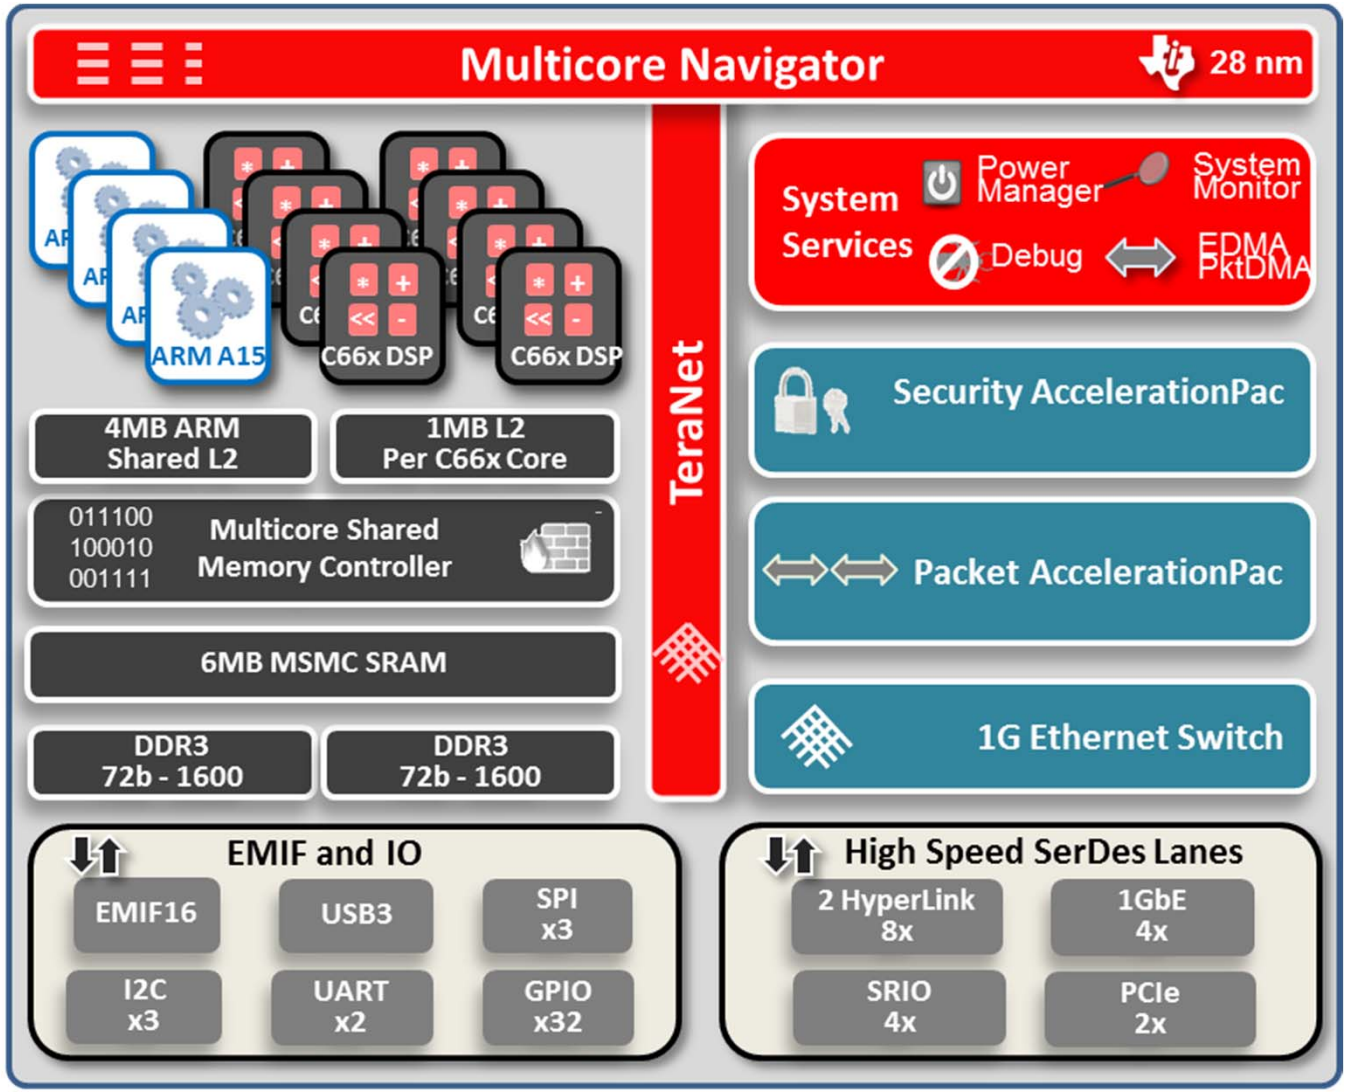

Available HPC Platforms

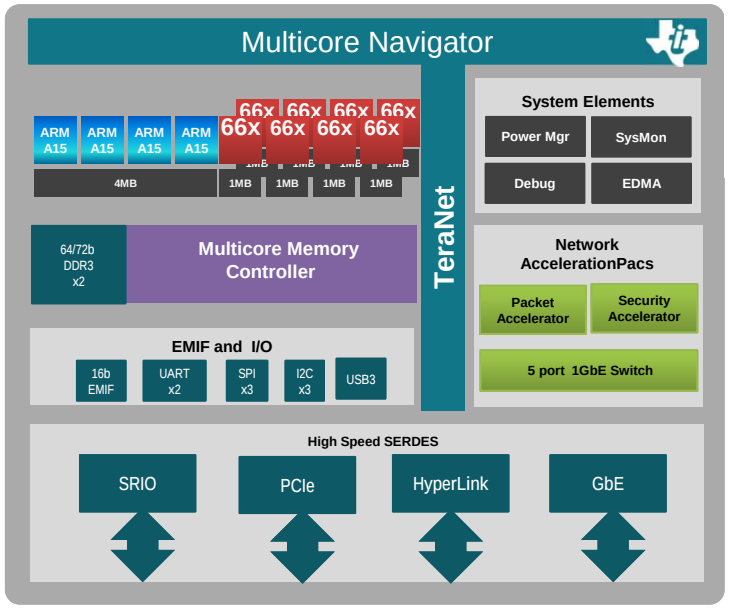

Programming Heterogeneous Multicore SoCs

HP Moonshot
ProLiant m800

nCore
Brown Dwarf

ProDrive PDAK2H

66AK2H
Evaluation Module

Heterogeneous Multicore Programming

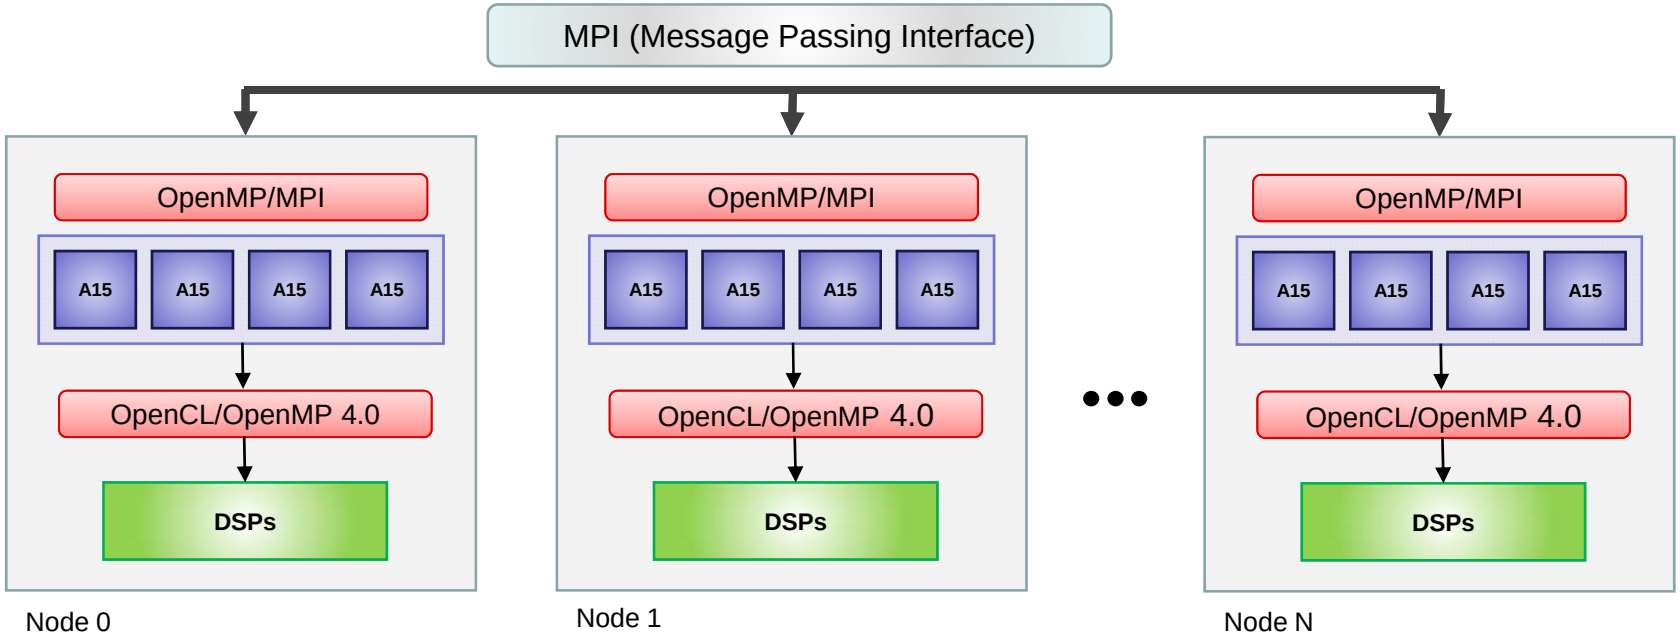

- Within a node, OpenCL™ or OpenMP® 4.0 can be used to program heterogeneous compute cores
- Across nodes, MPI is used to partition the application and manage program execution, data transfer and synchronization

ARM + OpenMP 4.0

```
// OpenMP Accelerator vector add
// OpenMP for loop parallelization
void ompVectorAdd(int   N,
                  float *a,
                  float *b,
                  float *c)
{
  #pragma omp target           \
  map(to:   N, a[0:N], b[0:N]) \
  map(from: c[0:N])
  {
    int i;
    #pragma omp parallel for
    for (i = 0; i < N; i++)
      c[i] = a[i] + b[i];
  }
}
```

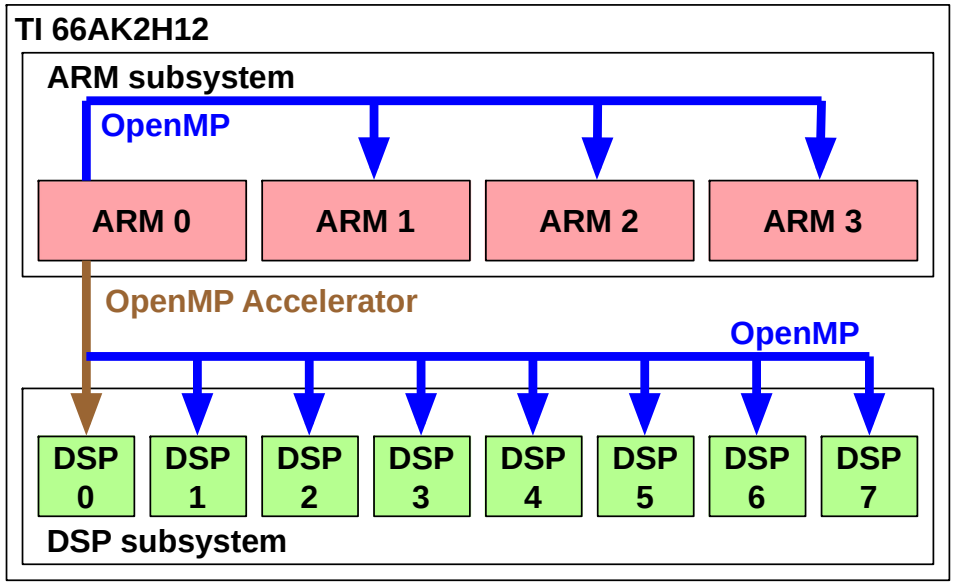

Data movement

- **`to`** copies variables from the ARM memory to the DSP memory
- **`from`** copies variables from the DSP memory to the ARM memory
- TI provides special **`alloc`** and **`free`** functions to allocate DSP memory such that copies are not needed

Calling existing DSP code from the ARM

- Wrapping existing DSP functions with OpenMP Accelerator code is straightforward

OpenMP Compiler support at TI

- OpenMP 3.0 on C667x (DSP only)
  - BIOS MCSDK 2.1.x
- OpenMP 3.0 + OpenMP 4.0 device constructs
  - HPC MCSDK 3.0.x
- Ongoing development to get to 4.0
- Come by the Texas Instruments Booth 3745 at SC