Breaking Master-Slave Model between Host and FPGAs

Jaume Bosch, Miquel Vidal, Antonio Filgueras, Carlos Álvarez, Daniel Jiménez-González, Xavier Martorell, Eduard Ayguadé

Barcelona Supercomputing Center (BSC), Universitat Politècnica de Catalunya (UPC)

Summary

- Support for task spawn and synchronization inside FPGA devices
  - FPGA tasks for other task accelerators in the FPGA bitstream
  - SMP or other arch tasks through reverse offload to the host runtime (Nanos++).
    This enables the ability to do I/O and perform syscalls from FPGA devices

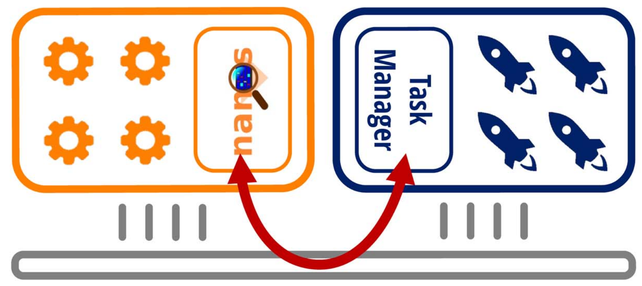

OmpSs@FPGA Ecosystem

- OmpSs (OpenMP forerunner) extends standard OpenMP syntax
- Designed for high productivity on heterogeneous systems
- Automatic data movements between host and FPGA address-spaces
- Support for Xilinx FPGAs (SoC and PCIe boards)

```
#pragma omp target device(fpga) copy_inout([bs*bs]A)
#pragma omp task
void potrf(float *A);

#pragma omp target device(fpga) copy_in([bs*bs]A)\
  copy_inout([bs*bs]B)
#pragma omp task
void trsm(const float *A, float *B);

#pragma omp target device(fpga) copy_inout([nb*nb*bs*bs]A)
#pragma omp task
void cholesky_blocked(const int nb, float *A) {
  for(int k = 0; k < nb; k++) {
    potrf( A + (k*nb + k)*bs*bs );
    #pragma omp taskwait

    for(int i = k + 1; i < nt; i++) {
      trsm( A + (k*nb + k)*bs*bs,
            A + (k*nb + i)*bs*bs );
    }
    #pragma omp taskwait
    //...
} }
```

Proposed Design

- Cooperation of OmpSs runtime (Nanos++, in the host) with the FPGA hardware runtime (Task Manager)
  - Communication based on queues in shared memory regions
  - Modular and reconfigurable hardware runtime design
- Keep task management near to task creation to minimize latency and maximize application performance

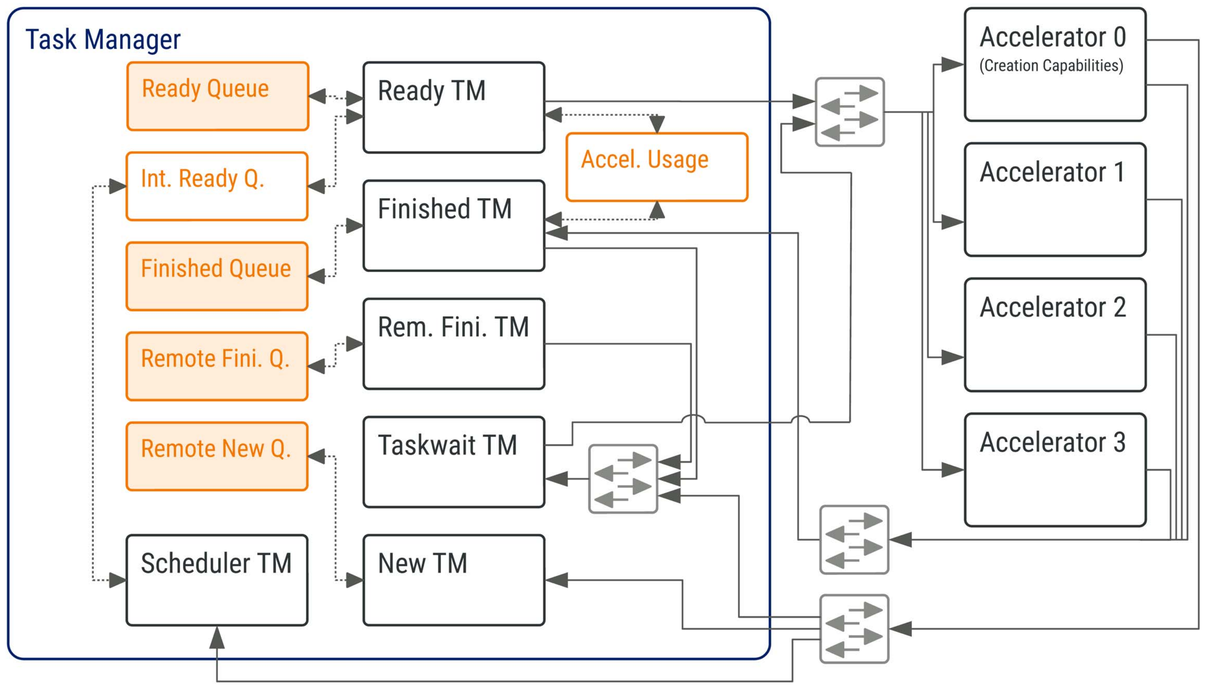

Evaluation

- Real implementation tested on Xilinx Zynq UltraScale+ MPSoC ZCU102
  - 4x A53 ARM cores, 4GB DDR4 memory, 1x ZU9EG Xilinx FPGA
  - 912 RAM36K, 2.520 DSP, 548.160 FF, 247.080 LUT
- Matrix Multiply: 7, 3 or 3 FPGA task accelerators at 300Mhz
- N-Body simulation
  - 4, 3 or 1 calculate_forces FPGA task accelerators at 250Mhz
  - 1 update_positions FPGA task accelerator at 250Mhz
- Cholesky factorization
  - trsm (3,1,1), gemm (6,4,2), syrk (1) FPGA task accelerators at 250Mhz
  - potrf at host threads with OpenBLAS
- FPGA task spanners consume 30mW (vs 800mW of ARM core)
  - 3 RAM36K (0.3%), 2.200 FF (0.4%), 8.800 LUT (3.6%)
- Peak performance moves to smaller task granularity (more parallelism)

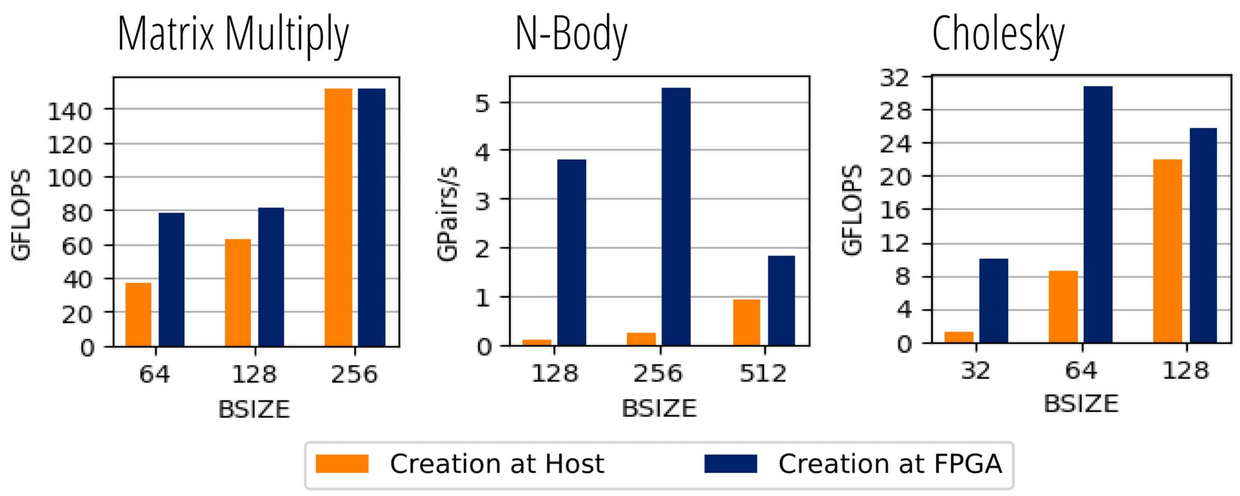

For more information visit pm.bsc.es/ompss-at-fpga or email me to jbosch@bsc.es

UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH
Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EPEEC PRODUCTIVITY AT EXASCALE
GOBIERNO DE ESPAÑA
Generalitat de Catalunya