

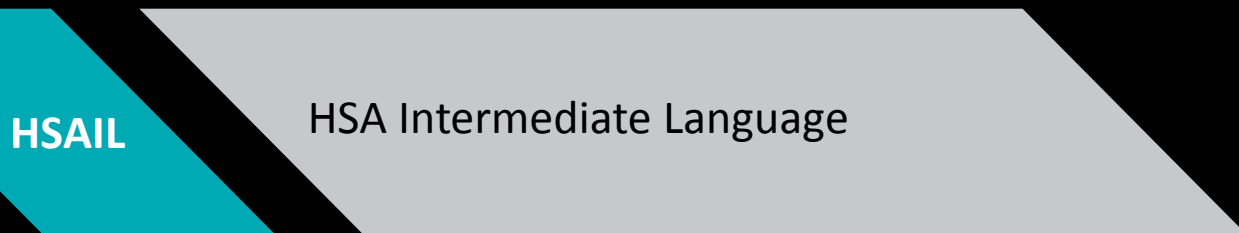


A Progress Report on COHX: Chapel on HSA+XTQ

Mauricio Breternitz, Bibek Ghimire, Mike Chu, Steve Reinhardt


AMD Research

CHI UW, June 13, 2015

HSA & XTQ




EQUAL ACCESS TO ENTIRE MEMORY

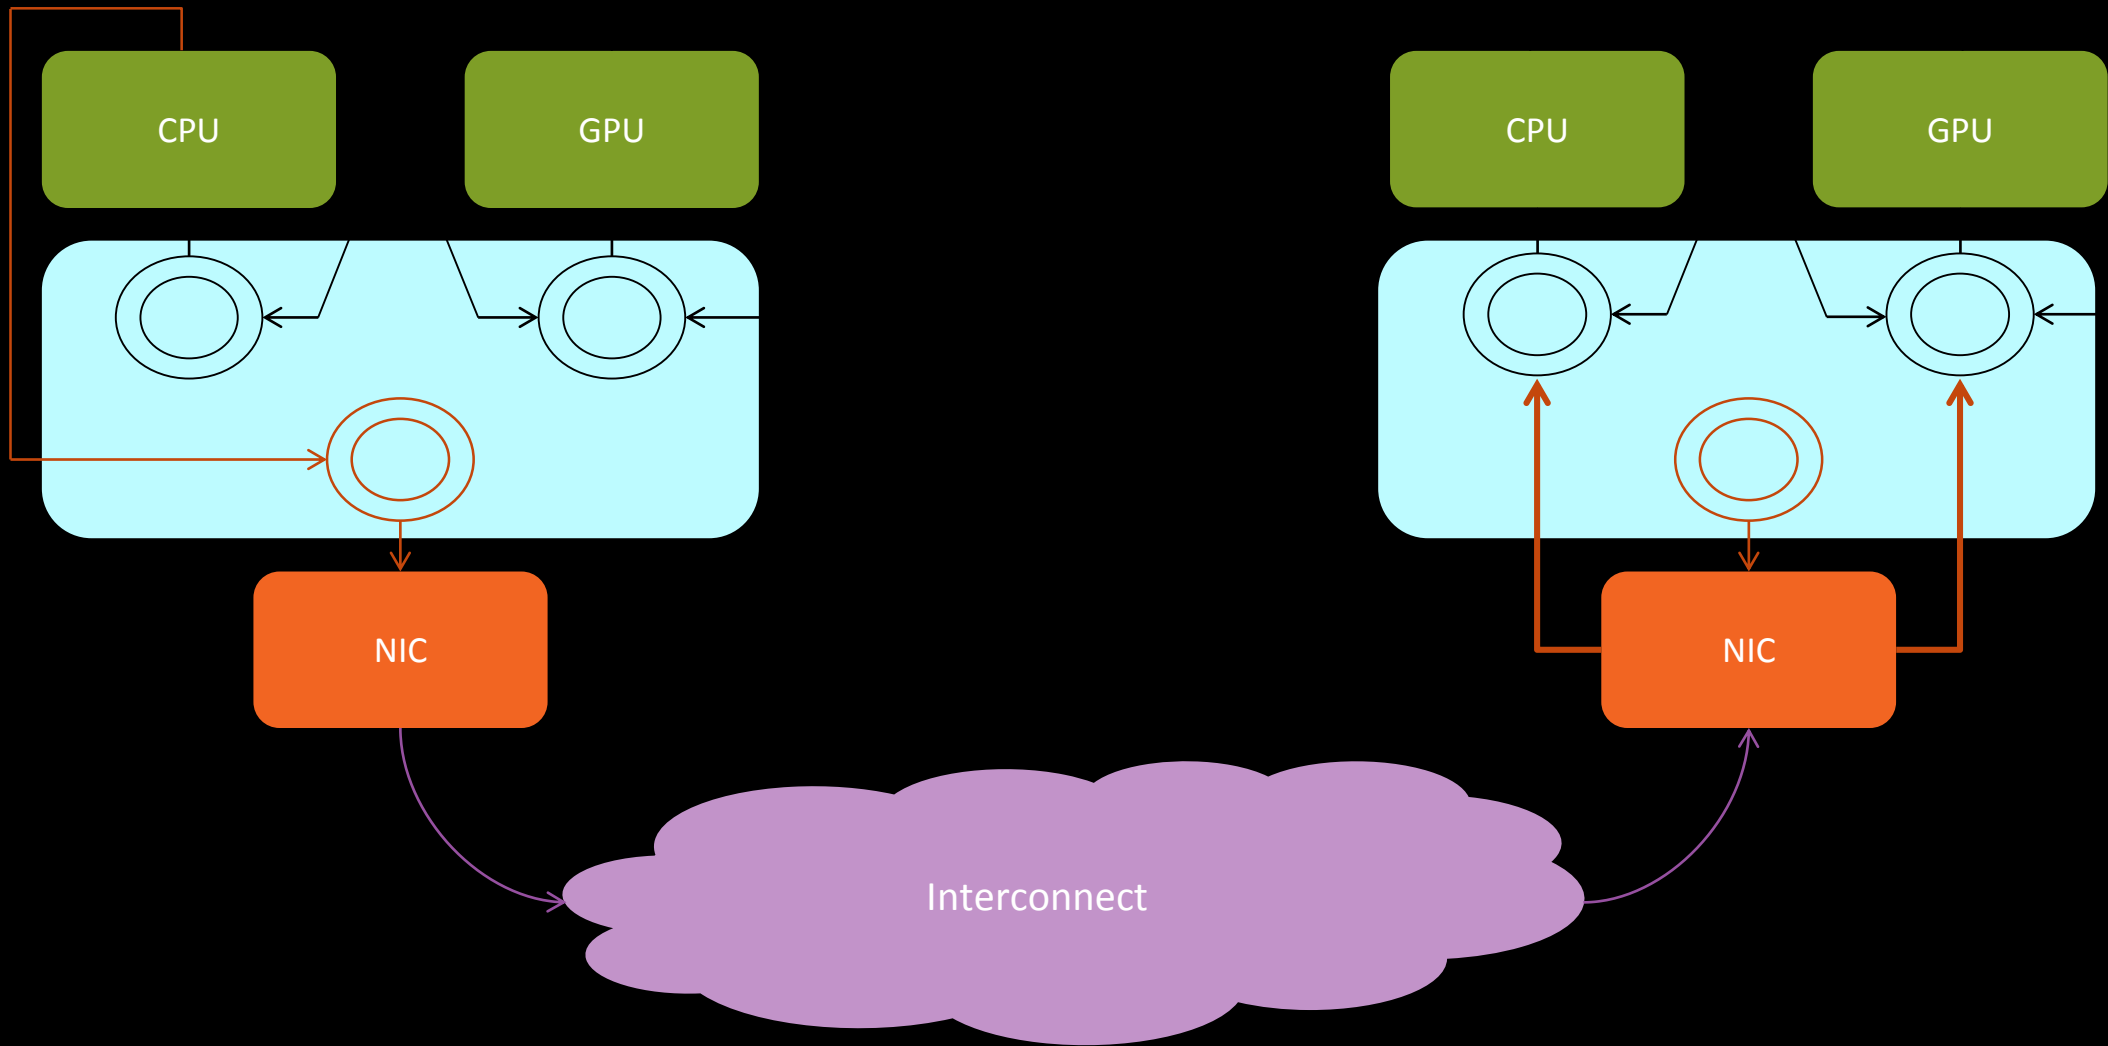


GPU and CPU have uniform visibility into entire memory space

ALL-PROCESSORS-EQUAL



GPU and CPU have equal flexibility to be used to create and dispatch work items



A locale is a unit of storage and processing

Specify at run-time the number of locales

Can specify at which locale objects are stored,
computations performed

Can specify locale by a variable

- `on TaskLocalA do taskA()`

Can specify locale by data location

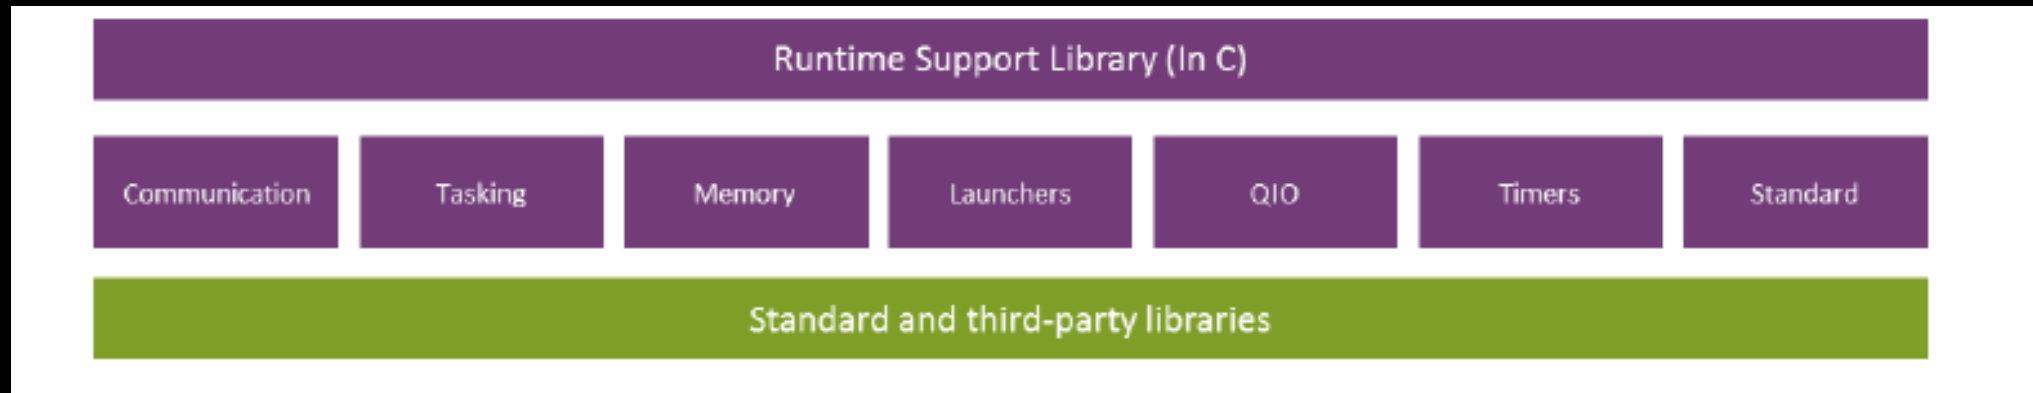
- `forall i in D on(a(i)) do ...`

Locality Control

```
writeln("on locale 0");  
  
on Locales[1] do  
  writeln("now on locale ", here.id);  
  
writeln("on locale 0 again");
```

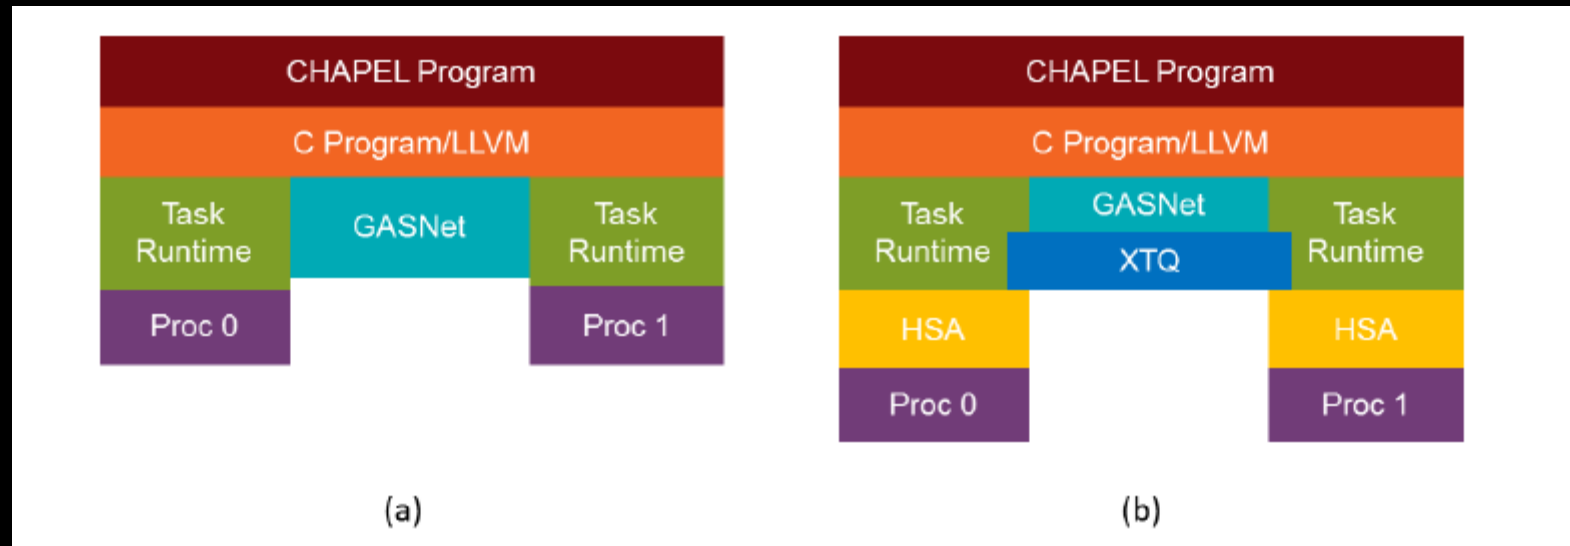
```
on locale 0  
now on locale 1  
on locale 0 again
```

CHAPEL RUNTIME ORGANIZATION



Focus of changes for XTQ

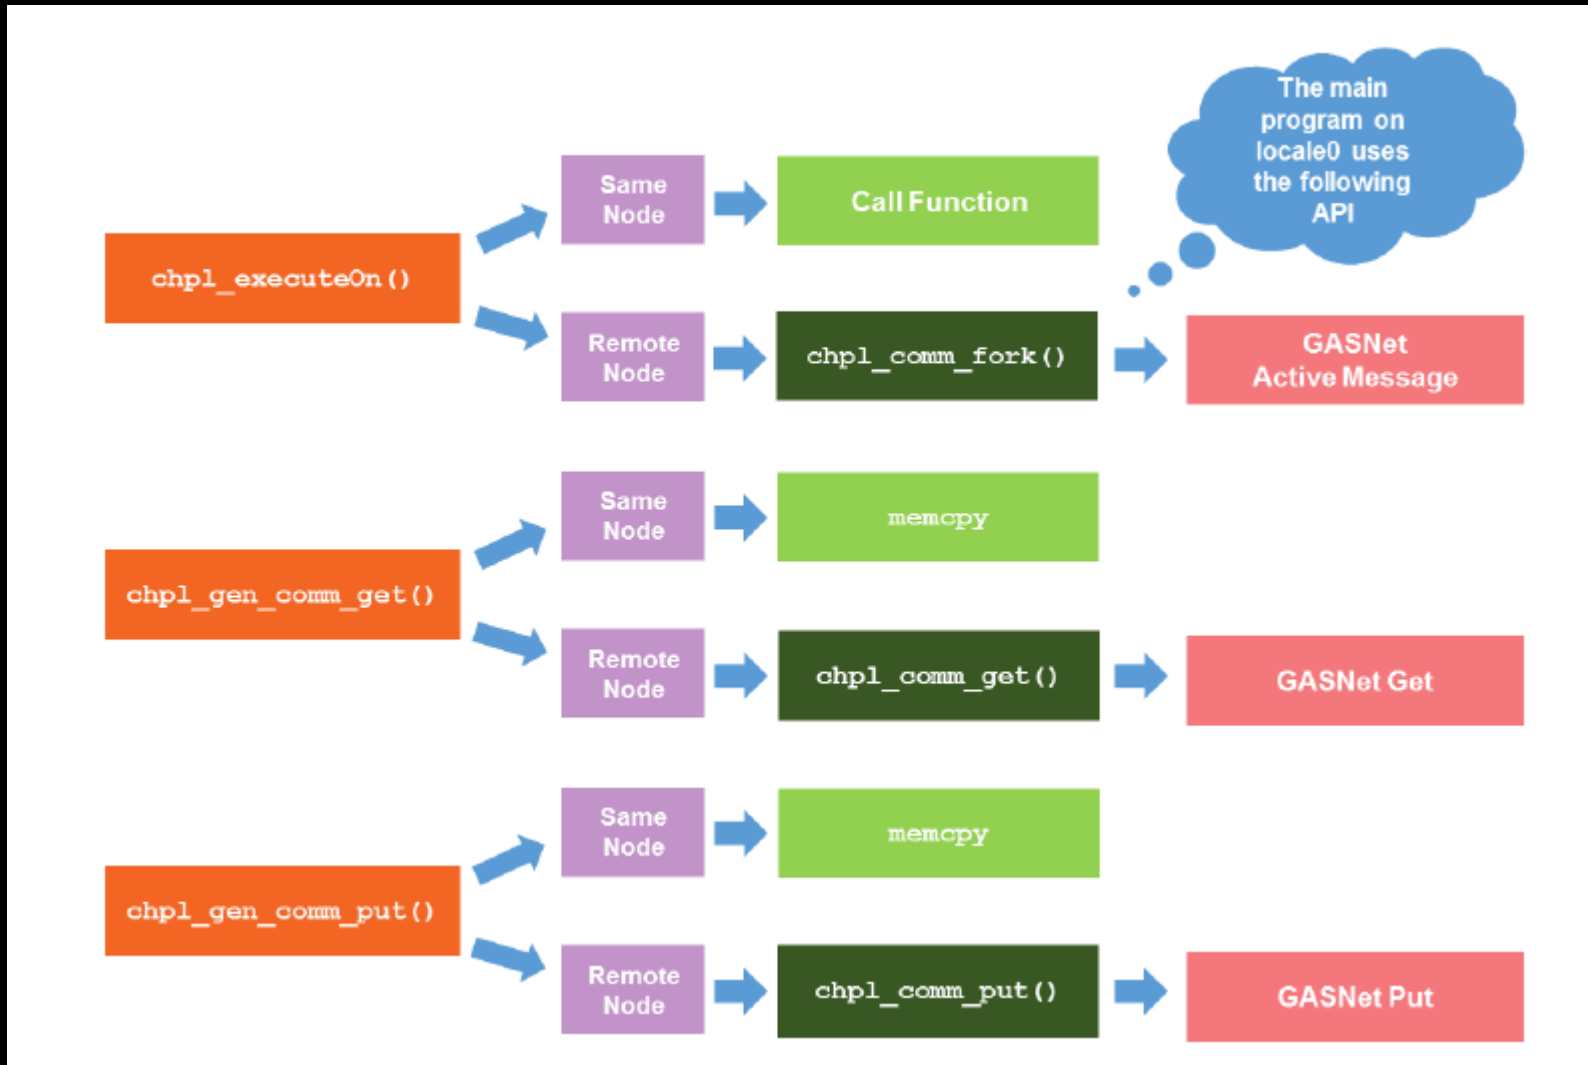
EXAMPLE: TWO-NODE CHAPEL EXECUTION



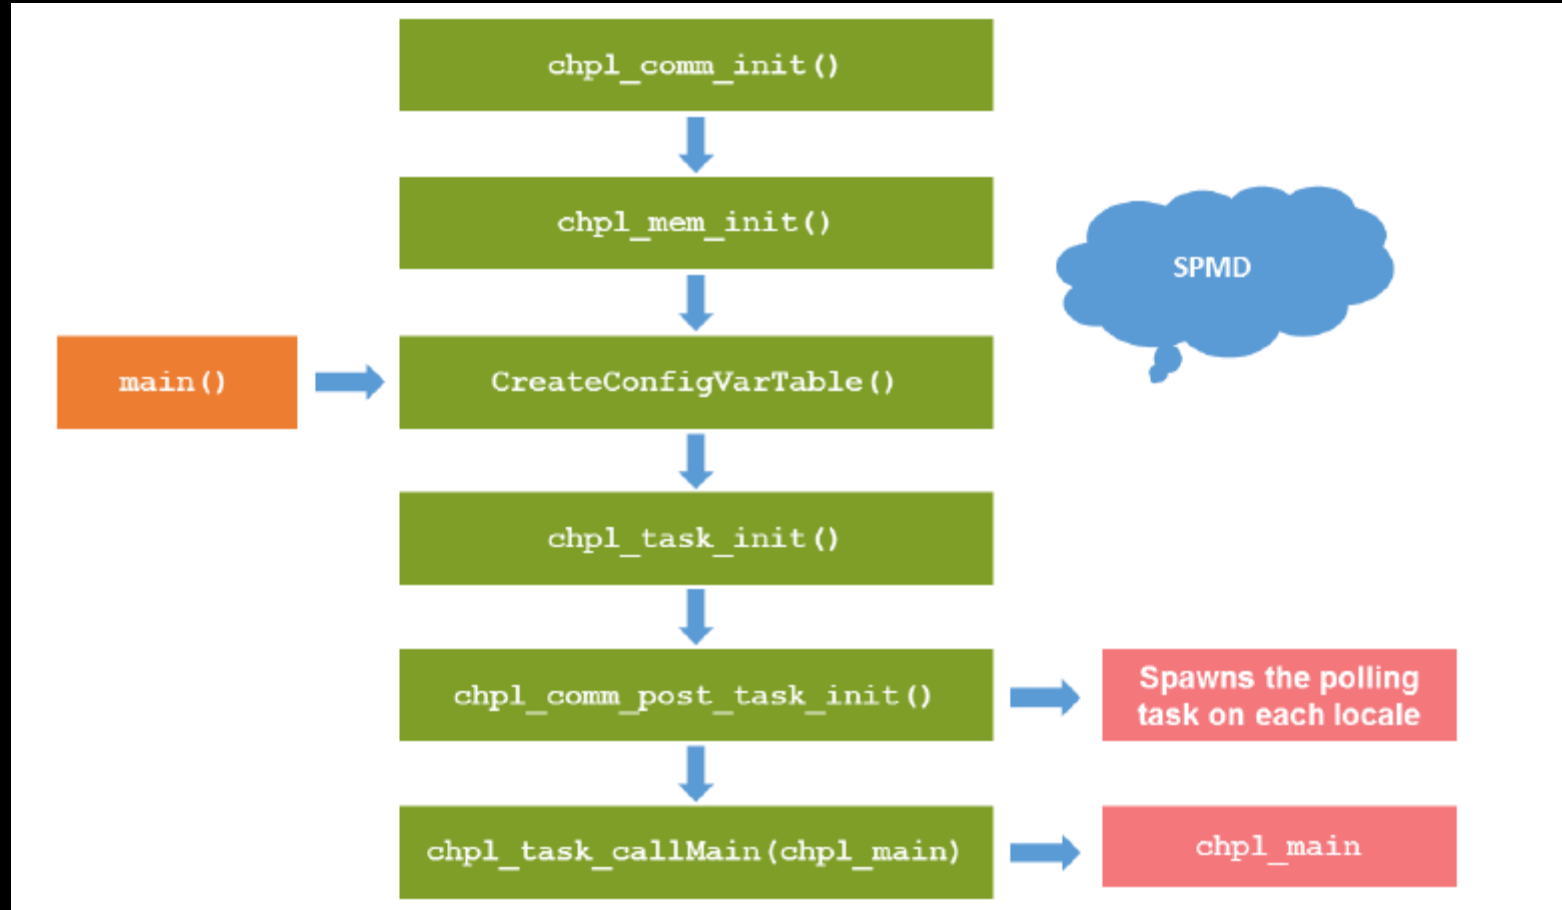
(a) GasNET


(b) XTQ+HSA

CHAPEL: 'ON' INVOCATION AND REMOTE DATA ACCESS

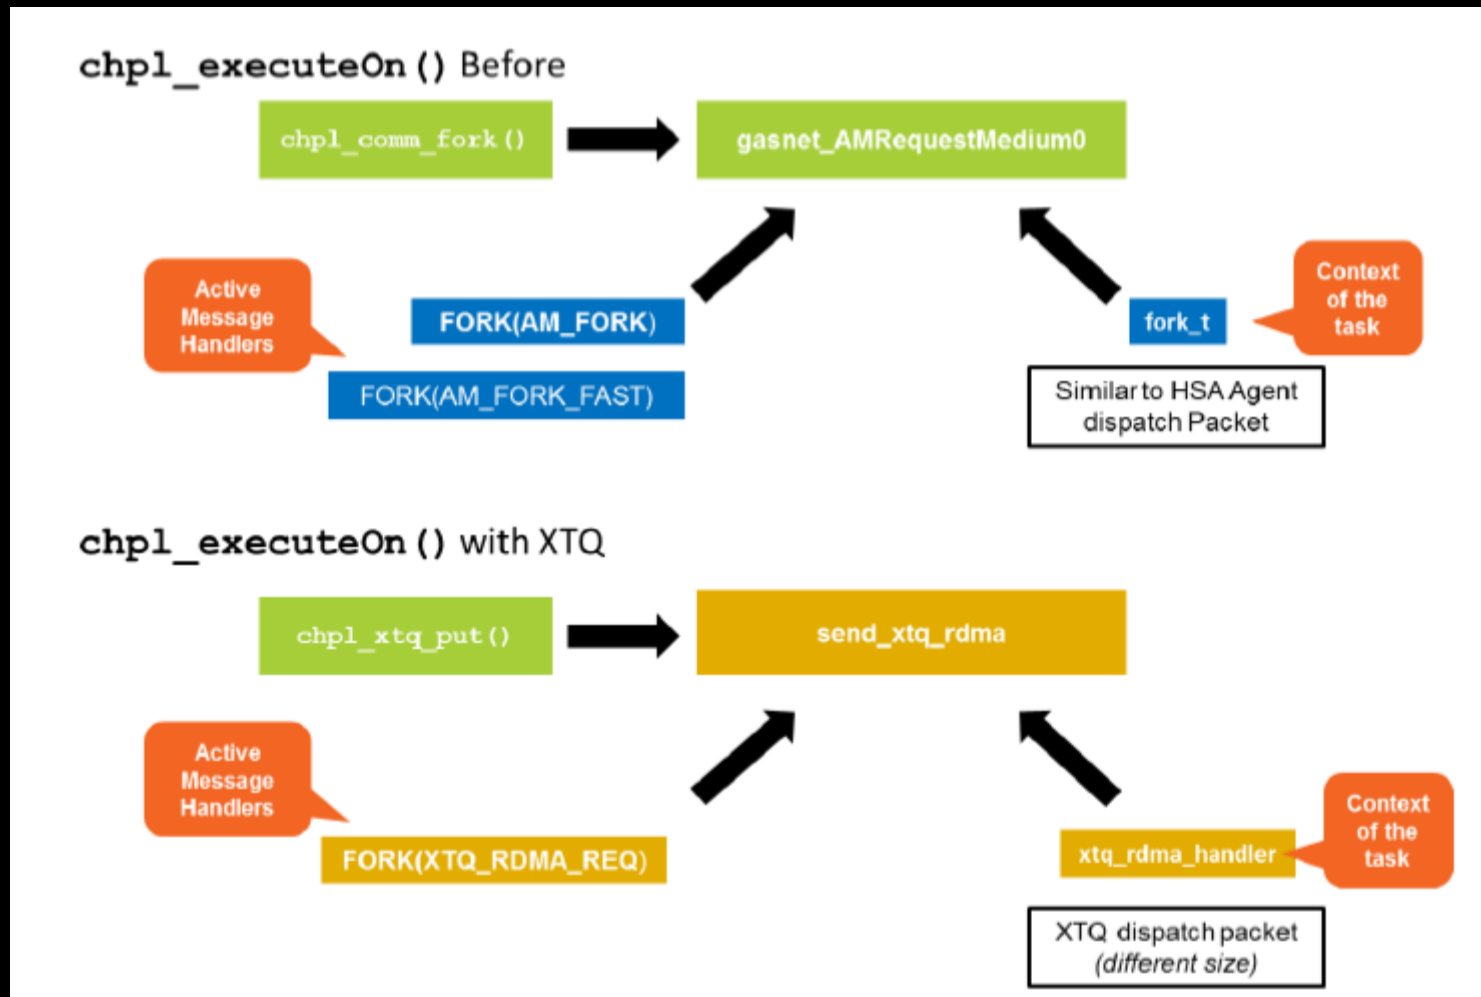


CHAPEL PROGRAM INITIALIZATION






CHAPEL ACTIVE MESSAGES OVER XTQ



CHAPEL REMOTE LAUNCH LATENCY MICROBENCHMARK



```
if (here.id == 0 )
{
    tstart();
    for i in 1..1000 {
        on Locales[(1-here.id) ]
        {
            global += 1;
        };
    }
    tfinish();
}
```

Chapel Microbenchmark

▲ Simple active message invokes remote method

```
(gdb) b XtqPt1Put
Breakpoint 1 at 0x404a90
(gdb) b xtq_rdma_handler
Breakpoint 2 at 0x4f5594: file to.Chapel.c, line 208.
(gdb) run -nl 2
Starting program: /home-nis/mbretern/two/nlchu.chapel-1.9.0_rdma_xtq/examples/tempu2/tempu2 -nl 2
[Thread debugging using libthread_db enabled]
Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
[New Thread 0x7ffff47bd700 (LWP 1265)]
[New Thread 0x7ffff3a48700 (LWP 1268)]
[New Thread 0x7ffff2f49700 (LWP 1269)]
ENV parameter: GASNET_FREEZE = NO (default)
ENV parameter: GASNET_RM_BUFFER_SIZE = 1048576 (default)
ENV parameter: GASNET_RM_NUM_ENTRIES = 16 (default)
ENV parameter: GASNET_RM_EVENT_QUEUE_LENGTH = 8192 (default)
ENV parameter: GASNET_EXTENDED_EVENT_QUEUE_LENGTH = 8192 (default)
ENV parameter: GASNET_NODEMAP_EXACT = NO (default)
ENV parameter: GASNET_SUPERNODE_MAXSIZE = 0 (default)
ENV parameter: GASNET_MAX_SEGSIZE = 13747867KB (13 GB) (default)
ENV parameter: GASNET_MALLOC_INIT = NO (default)
ENV parameter: GASNET_MALLOC_INITVAL = NAN (default)
ENV parameter: GASNET_MALLOC_CLOBBER = NO (default)
ENV parameter: GASNET_MALLOC_CLOBBERVAL = NAN (default)
ENV parameter: GASNET_MALLOC_LEAKKILL = NO (default)
ENV parameter: GASNET_MALLOC_SCANFREED = NO (default)
ENV parameter: GASNET_MALLOC_EXTRACHECK = NO (default)
ENV parameter: GASNET_BARRIER = DISSEM (default)
ENV parameter: GASNET_COLL_MIN_SCRATCH_SIZE = 1 KB (default)
ENV parameter: GASNET_COLL_SCRATCH_SIZE = 2 MB (default)
ENV parameter: GASNET_DISABLE_ARGDECODE = NO (default)
ENV parameter: GASNET_DISABLE_ENVDECODE = NO (default)
ENV parameter: GASNET_BACKTRACE = YES (default)
ENV parameter: GASNET_BACKTRACE_NODES = *not set* (default)
ENV parameter: GASNET_TMPDIR = *not set* (default)
ENV parameter: TMPDIR = *not set* (default)
ENV parameter: GASNET_BACKTRACE_TYPE = GDB_EXECINFO (default)
ENV parameter: GASNET_FREEZE_ON_ERROR = NO (default)
ENV parameter: GASNET_FREEZE_SIGNAL = *not set* (default)
ENV parameter: GASNET_BACKTRACE_SIGNAL = *not set* (default)
ENV parameter: GASNET_TRACEFILE = *empty* (default)
ENV parameter: GASNET_STATSFILE = *empty* (default)
ENV parameter: GASNET_TRACEFLUSH = NO (default)
ENV parameter: GASNET_TRACELOCAL = YES (default)
ENV parameter: GASNET_MALLOCFILE = *empty* (default)
ENV parameter: GASNET_DISABLE_MUNMAP = NO (default)
ENV parameter: GASNET_MAX_THREADS = 256 (default)
ENV parameter: GASNET_VIS_RMPPIPE = NO (default)
ENV parameter: GASNET_VIS_MAXCHUNK = 2032 B (default)
ENV parameter: GASNET_VIS_REMOTECONTIG = NO (default)
[New Thread 0x7ffff3df4700 (LWP 1270)]
[New Thread 0x7ffff35fc700 (LWP 1271)]
[New Thread 0x7ffffd8df9700 (LWP 1275)]
[New Thread 0x7ffffd85f8700 (LWP 1276)]
TLS=16
[Thread 0x7ffffd85f8700 (LWP 1276) exited]
[New Thread 0x7ffffd85f8700 (LWP 1278)]
[New Thread 0x7ffffd7df7700 (LWP 1280)]
[Switching to Thread 0x7ffffd7df7700 (LWP 1280)]

Breakpoint 1, XtqPt1Put (nd_handle=83888276,
local_offset=140737018599664, length=248,
ack_req=1 '\001', target_id=...,
pt_index=31, match_bits=170,
remote_offset=0, user_ptr=0x7ffff40019f0,
hdr_data=0) at ptl_move.c:96
96 ptl_size_t remote_offset, void *user_ptr, ptl_hdr_data_t hdr_data) {
(gdb) c
Continuing.
```

1 XtqPut

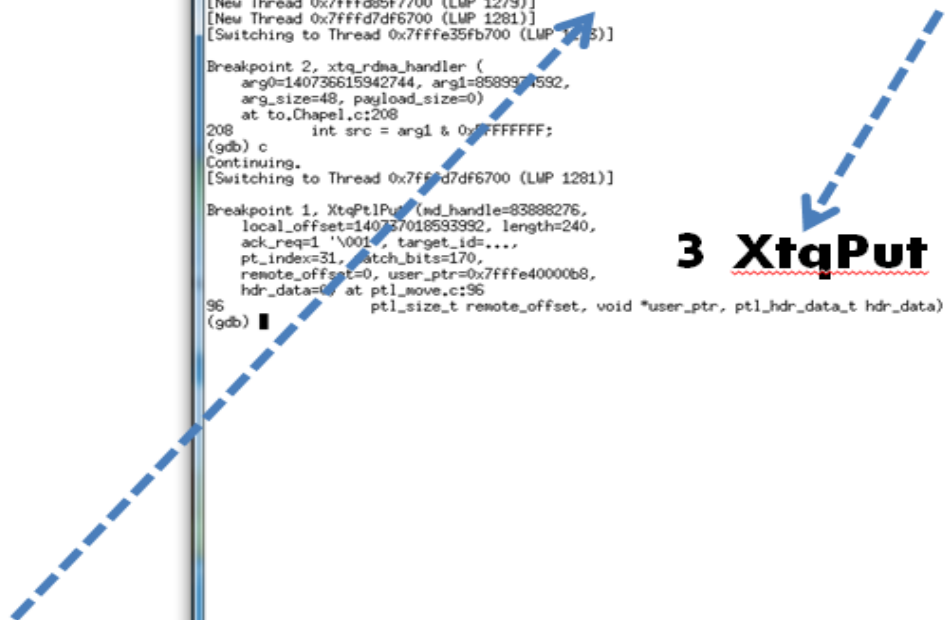
```
(gdb) b XtqPt1Put
Breakpoint 1 at 0x404a90
(gdb) b xtq_rdma_handler
Breakpoint 2 at 0x4f5594: file to.Chapel.c, line 208.
(gdb) run -nl 2
Starting program: /home-nis/mbretern/two/nlchu.chapel-1.9.0_rdma_xtq/examples/tempu2/tempu2 -nl 2
[Thread debugging using libthread_db enabled]
Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
[New Thread 0x7ffff47bd700 (LWP 1256)]
[New Thread 0x7ffff3a48700 (LWP 1259)]
[New Thread 0x7ffff2f49700 (LWP 1260)]
[New Thread 0x7ffff3dfc700 (LWP 1272)]
[New Thread 0x7ffff35fb700 (LWP 1273)]
[New Thread 0x7ffffd8df8700 (LWP 1274)]
[New Thread 0x7ffffd85f7700 (LWP 1277)]
TLS=16
[Thread 0x7ffffd85f7700 (LWP 1277) exited]
[New Thread 0x7ffffd85f7700 (LWP 1279)]
[New Thread 0x7ffffd7df6700 (LWP 1281)]
[Switching to Thread 0x7ffff35fb700 (LWP 1259)]

Breakpoint 2, xtq_rdma_handler (
arg0=140736615942744, arg1=858997592,
arg_size=48, payload_size=0)
at to.Chapel.c:208
208 int src = arg1 & 0xffffffff;
(gdb) c
Continuing.
[Switching to Thread 0x7ffffd7df6700 (LWP 1281)]

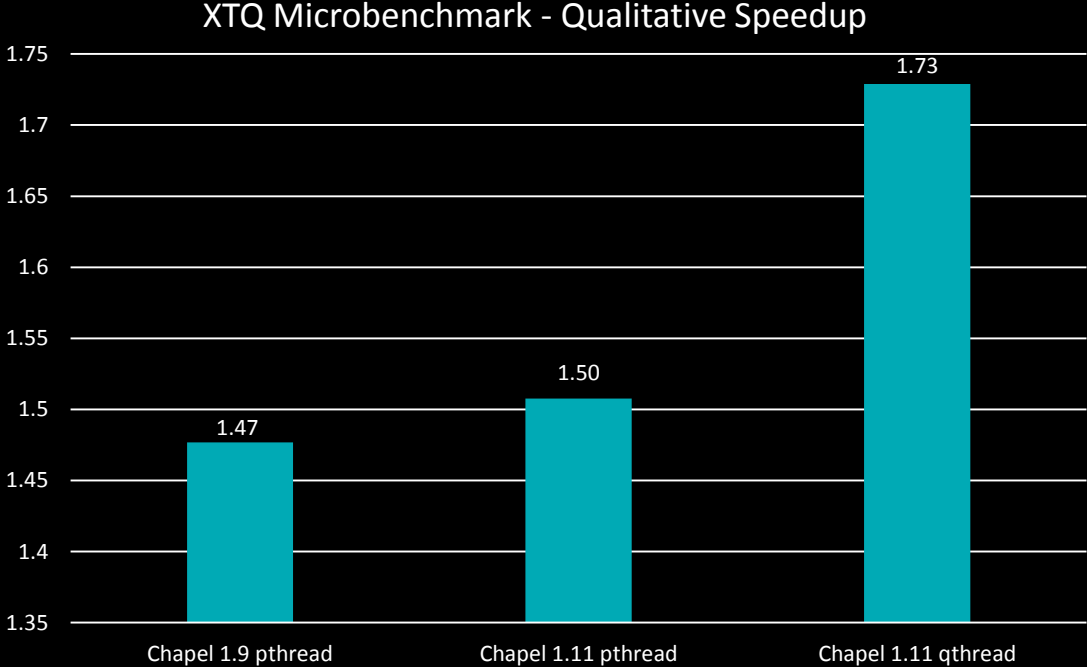
Breakpoint 1, XtqPt1Put (nd_handle=83888276,
local_offset=140737018593992, length=240,
ack_req=1 '\001', target_id=...,
pt_index=31, match_bits=170,
remote_offset=0, user_ptr=0x7ffff40000b8,
hdr_data=0) at ptl_move.c:96
96 ptl_size_t remote_offset, void *user_ptr, ptl_hdr_data_t hdr_data) {
(gdb) |
```

2 rdma_handler

3 XtqPut



XTQ MICROBENCHMARK – QUALITATIVE SPEEDUP



- ▲ XTQ provides Chapel Runtime Speedup
- ▲ ExecuteOn utilized by Runtime → speedup ?
- ▲ LocaleInference + Aggregation → XTQ: data + remote method
- ▲ HSA generation
- ▲ HSA signals

DISCLAIMER & ATTRIBUTION



The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

© 2015 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. SPEC is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). Other names are for informational purposes only and may be trademarks of their respective owners.