

Software-Configured Compute Environments

Ulrich Drepper <drepper@redhat.com>

November 2018

Driving Factors

































Abstraction Problem

```
int cnt1;
int cnt2;
int main() {
   auto f1 = [](){ while (! f1done) { ++cnt1; f1work(); } };
   auto f2 = [](){ while (! f2done) { ++cnt2; f2work(); } };
   std::thread t1(f1), t2(f2);
   t1.join();
   t2.join();
}
```



Abstraction Problem

```
int cnt1;
int cnt2; (Possibly) False Sharing
int main() {
    auto f1 = [](){ while (! f1done) { ++cnt1; f1work(); } };
    auto f2 = [](){ while (! f2done) { ++cnt2; f2work(); } };
    std::thread t1(f1), t2(f2);
    t1.join();
    t2.join();
}
```



Breaking the Abstraction

Matrix Multiplication





Breaking the Abstraction

Matrix Multiplication





While We Are at Matrix Math...

```
Eigen::Matrix<M,N> m1;
Eigen::Matrix<N,P> m2;
auto m3 = m1 * m2;
```

Quality of Implementation



While We Are at Matrix Math...





One-Size-Fits-All Architecture





Cache Number Explosion

Many different access times:

- 1+2 Level 1 Cache
- 3 Level 2 Cache
- 4 Last Level Cache
- 5 Caches on neighboring core
- 6 Caches on distant core
- 7 Local DRAM
- 8 Local NVRAM
- 9 Cache on remote cores
- 10 On var remote core
- 11 RAM attached to remote socket
- 12 NVRAM attached to remote socket

	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		
	SKX Core	SKX Core	SKX Core	SKX Core	SKX Core	SKX Core		
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		SF/LLC
	SKX Core	5KX Core	SKX Core	5KX Core	SKX Core	SKX Care		
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	c 🕤 📖	CHA/SF/LLC		L1D
	SKX Core	SKX Core	SKX Core	5KX Core	SKX Core	5KX Core	L2	111
	DDR4 MC	CHA/SF/LLC	CHA/SF/LLC		CHA/SF/LLC			
	DDR4	SKX Core	SKX Core	5KX Core	SKX Core	DDR4	3	2
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		
	SKX Core	SKX Core	SKX Core	5KX Core	SKX Core	5KX Core		
RAM DRAM	2x UPI x20	PCIe x16	PCIe x16 DMI x4	On Pkg PCle x16	1x UPI x20	PCle x16		
						=		
	2x U? x20	PCle x16	PCIe x16 DMI x4	On Pkg PCle x16	1x UPI x20	PCle x16		
RAM DRAM		H				8		
	ci 🧐 ilic	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		
	SKX Core	SKX Core	SKX Core	SKX Core	SKX Core	SKX Core		
	DDR4 MC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	MC DDR4		
	DDR4	5KX Core	SKX Core	5KX Core	SKX Core	DDR4 DDR4		
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		
	SKX Core	SKX Core	SKX Core	SK) (cr)	SKX Core	5KX Core		
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		
<mark>2) (11</mark>)	SKX Core	5KX Core	SKX Core	5KX Core	SKX Core	5KX Core		
	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC	CHA/SF/LLC		



Hot Issue in NNs

Data types

- FP8, FP16, FP32, FP64
- Int4, Int8, Int16, Int32
- Coarse granularity

Why not FP12?



Kernel Steering

- Execution
 - Processes
 - Kernel Threads
- Memory
- Interrupts



Predictability Matters!









Normal Networking





Kernel Bypass





Real-Time

One Approach





What If ...?





Adjust Compute Environment Dynamically for each Process



Bump-In-The-Wire





Bump-In-The-Wire

Instead of kernel bypass:

- Implement function on FPGA
- Implement decoding on FPGA
 - Reduced communication with host
 - Complex operations on host





Cache Allocation

Restrict shared resource use

- Some hardware support available
- Not process property
- Static configuration



PID2TID2

PID3TID1

PID2TID1

PID1TID1



Cache Allocation

Reconfigure hardware

- Generate softcore processor for problem
- For VMs/containers, not processes
- Common hypervisor?

Core	Core	Core	Core	
L1I L1D	L1I L1D	L1I L1D	L1I L1D	
L2	L2	L2	L2	
L3	L	3	L3	

PID2TID1

PID2TID2

PID3TID1



Deploy Unikernels

• Normal OS on bare metal





Deploy Unikernels

- Normal OS on bare metal
- Direct access to devices
 - Virtual functions
 - SR-IOV
- Consecutive memory range
- Dedicated CPU sockets or cores





(FROM AHMED'S SLIDES)

Architecture



Summary



Tasks

- CPU research
 - Softcore implementation
 - Reconfigurable with IP blocks addressed through custom instructions
 - Accelerator IP blocks
 - Cache isolation
- Integrated development platform
 - Open source HDL toolchain
 - Integration of softcore as accelerator in OpenMP/...
- Research compiler techniques
 - C/C++ vectorization
 - DSL, translate to C/C++
 - Unikernel binary generation

- OS research
 - Run unikernels as executables
 - Dedicated resources
 - For efficiency, latency, RT
 - (automatic) configure cache isolation
- FPGA "OS" research
 - API, basic services, security
 - Toolchain to create services
- Runtime research
 - Automatic parallelization control
 - Automatic resource allocation and deallocation
 - Automatic selection of accelerator and workload split





THANK YOU



plus.google.com/+RedHat



linkedin.com/company/red-hat



youtube.com/user/RedHatVideos



facebook.com/redhatinc



twitter.com/RedHatNews