On-Chip Network-Enabled Multi-Core Platforms Targeting Maximum Likelihood Phylogeny Reconstruction

Turbo Majumder, PhD Candidate, School of EECS (tmajumde@eecs.wsu.edu)

WASHINGTON STATE UNIVERSITY
World Class. Face to Face.

Objective

- Finding a ***phylogenetic tree*** that best explains the evolutionary relationship among a given set of species is ***computationally complex*** because of the ***greater than exponential search-space*** (real, multi-dimensional) and ***floating point arithmetic*** computation.

Why Network-on-Chip?

- Phylogenetic tree reconstruction is ***data/computation intensive***.
- ***Use parallelism***: Divide into a large number of smaller ***semi-independent*** sub-problems that can be computed ***concurrently***.

NoC
- A large number of cores to deal with sub-problems
- Low-latency inter-core communication

Multi-core System Design

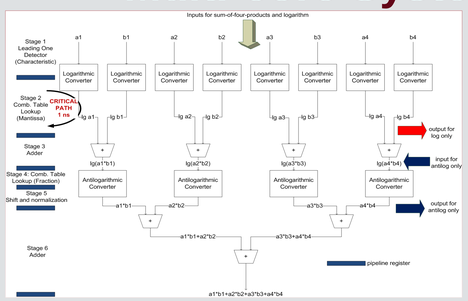

Computation core

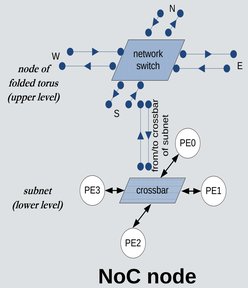

NoC node

Network architectures:
2-D folded torus (*2D_serial, 2D_parallel*)

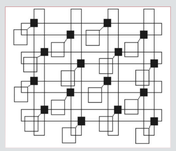

3-D folded torus *(3D_torus)*

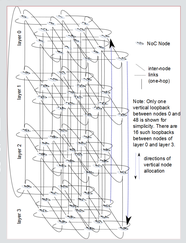

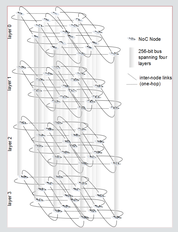

3-D stacked torus (*3D_sttorus*)

Dynamic Node Allocation

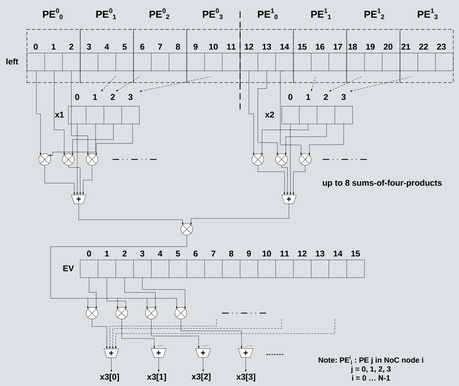

Schematic representation of computation tree of *newviewGTRCAT* (*f2*)

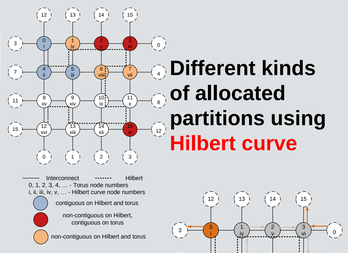

Different kinds of allocated partitions using Hilbert curve

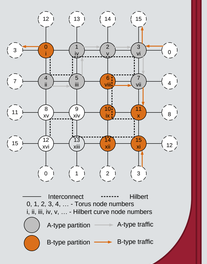

Traffic types and routing

Experimental Setup

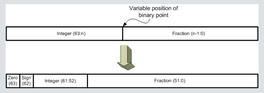

❑**Computation core**
- Datapath: 64 bit
- Number representation accuracy of 2^{-52} using Fixed-Point Hybrid Number System
- All components designed with Verilog HDL and synthesized with 65 nm standard libraries

❑**Multi-core System**
- Interconnects laid out, parasitics (resistance, capacitance) extracted to determine physical parameters (power dissipation, delay)
- *N*=16 and *N*=64 system sizes simulated using TreeSim
- 32-lane PCIe 2.0 interface (5 Gbps)

❑**Software**
- RAxML-VI-HPC (version 7.0.4) on three inputs sourced from 2,177-taxon 68-gene mammalian dataset
- Pentium IV 3.2 GHz dual-core CPU; GNU *gprof* utility for profiling
- Best software runtime used as the baseline
- Functions *coreGTRCAT* (*f3*) (48%), *newviewGTRGAMMA* (*f6*) (21%) and *newviewGTRCAT* (*f2*) (17%) collectively account for more than 85% of the total runtime

Results

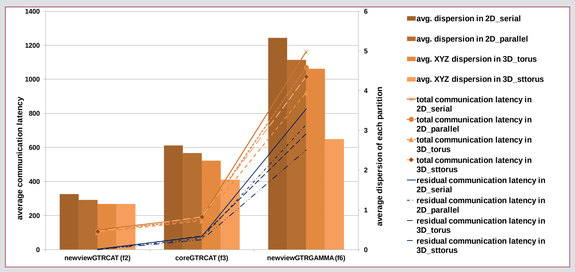

Variation of partition dispersion and function communication latency across different NoC architectures

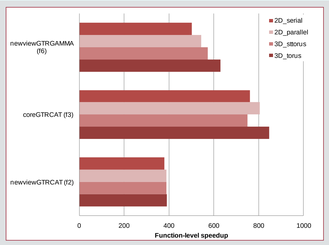

Function-level speedup across different NoC architectures

Results (cont'd.)

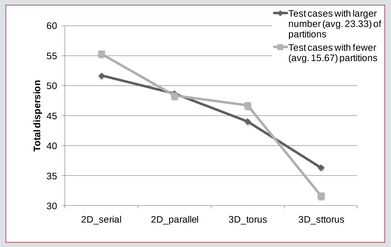

Total dispersion across different NoC architectures

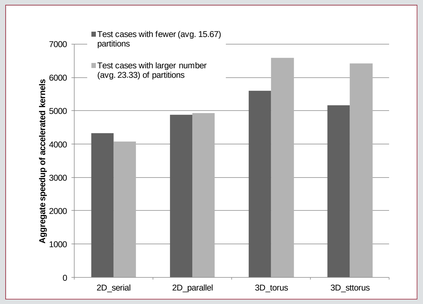

Average aggregate speedup of the accelerated kernels across different NoC architectures

Total run-times for different inputs using different NoC-based platforms vis-à-vis only software

Input data (DNA)		Unaccelerated software run-time (s)	Time spent in accelerated kernels (s)	Allocation time (s)	PCIe interface time (s)	Total run-time using NoC platform as hardware accelerator (s)	Total 4T software run-time (s)
50_5000	2D_serial	292.000444	0.515478	0.130387	0.145065	292.791374	924.052039
	2D_parallel	292.000444	0.481303	0.104805	0.145065	292.731617	924.052039
	3D_torus	292.000444	0.433625	0.050889	0.145065	292.630024	924.052039
	3D_sttorus	292.000444	0.474657	0.050889	0.145065	292.671056	924.052039
500_5000	2D_serial	7038.847538	19.1142	8.467062	8.273363	7074.702162	37124.7233
	2D_parallel	7038.847538	18.04733	6.805803	8.273363	7071.974034	37124.7233
	3D_torus	7038.847538	16.766102	3.304655	8.273363	7067.191658	37124.7233
	3D_sttorus	7038.847538	18.102936	3.304655	8.273363	7068.528491	37124.7233

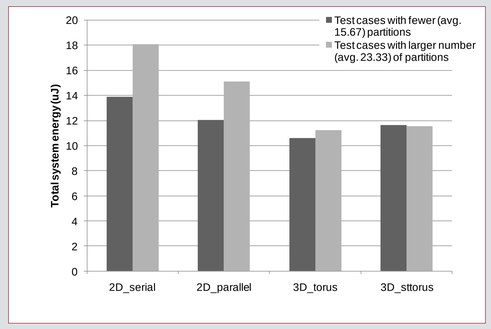

Total system energy consumption across different NoC architectures

Summary

•Network-on-Chip (NoC) based multi-core platform for accelerating Maximum Likelihood (ML) based phylogeny reconstruction
•Chief contributions
→design of a fine-grained parallel PE architecture
→novel algorithm to dynamically allocate nodes to tasks based on Hilbert space-filling curves
→design and extensive evaluation of different 2-D and 3-D NoC architectures
•Function-level speedup of ~**847x**, aggregate speedup of the accelerated portion up to ~**6,500x**, and overall run-time reduction of more than **5x** over multithreaded software; exceeds the performance of all state-of-the-art hardware accelerators for this application class.