

Automated Reliability Classification of Queueing Models for Streaming Computation

Jonathan Beard, Cooper Epstein,
and Roger Chamberlain



Stream Based
Supercomputing Lab
<http://sbs.wustl.edu>



Washington
University in St. Louis


Work also supported by:




Stream Processing

➔ for $i \leftarrow 0$ through N do
 $a[i] \leftarrow (b[i] + c[i])$
 $i++$
end do

Traditional Control Flow




Streaming



Stream Processing

Kernel



Assigned To Any Compute Resource



RaftLib

C++ Streaming Template Library


Auto-parallelizes code

Manages resources, buffers, TCP links




Automatically Optimized Online

software download: <http://raftlib.io>


Modeling Streams as Queues




“Stream” is modeled as a Queue




Buffer Sizing



Traditional Service Rate




Isolated Compute Kernel



Counter - In


Counter - Out

Example Streaming App




Matrix Multiply

Example Streaming App




Rabin-Karp String Search

More Complex Example



“Knowledge” Transfer



“Knowledge”

Machine Learning

Support Vector Machine


Can we use a Support Vector Machine (SVM) to classify a model as use or don't use?

M/D/1

M/M/1

None

?




Features




Kernel

Dual Core AMD Opteron Processor 280
Dual Core AMD Opteron Processor 875
Level Three Cache Associativity
Quad-Core AMD Opteron Processor 2387
2.6.32-358.23.2.el6.x86_64
Level One ICache Associativity
2.6.32-431.5.1.el6.x86_64
Quad-Core AMD Opteron Processor 2374
traditional_with_pset_runqueue
Intel Xeon CPU X5472 @ 3.00GHz
Dual-Core AMD Opteron Processor 2218
Processor Frequency
Level One DCache Line Size
Level Two Cache Size
Fifteen Min Load
3.10.37+
Intel Xeon CPU E5-2640 v4 @ 2.40GHz
Total Swap
Darwin
armv6l
Linux
ppc64
Level One D Cache Associativity
AMD Opteron Processor 6200
Level Two Cache Associativity
Level Three Cache Line Size
Level Three Cache Size
Number of Processors
Free Ram
Level One DCache Size
Level Two Cache Line Size
Total Main Memory
PPC970
Arrival Process Service Time
ARM1176JZF-S
Number of Processes Running
SCHED_OTHER
3.2.14
Buffer Ram
Level One ICache Size
Free Swap
13.1.0
3.0.27

Monitor Arrangement



Online Service Rate




Kernel

Dual Core AMD Opteron Processor 280
Dual Core AMD Opteron Processor 875
Level Three Cache Associativity
Quad-Core AMD Opteron Processor 2387
2.6.32-358.23.2.el6.x86_64
Level One ICACHE Associativity
2.6.32-431.5.1.el6.x86_64
Quad-Core AMD Opteron Processor 2374
traditional_with_pset_runqueue
Intel Xeon CPU X5472 @ 3.00GHz
Dual-Core AMD Opteron Processor 2218
Processor Frequency
Level One DCACHE Line Size
Level Two Cache Size
Fifteen Min Load
3.10.37+
Intel Xeon CPU E5-2640 v4 @ 2.10GHz
Total Swap
Darwin
armv6l
Linux
ppc64
Level One D Cache Associativity
AMD Opteron Processor 6200
Level Two Cache Associativity
Level Three Cache Line Size
Level Three Cache Size
Number of Processors
Free Ram
Level One DCACHE Size
Level Two Cache Line Size
Total Main Memory
PPC
Level One ICACHE Line Size
Arrival Process
Intel Xeon CPU E5345 @ 2.33GHz
Dual-Core AMD Opteron Processor 2212 SE
3.13.0-24-PowerPc64-SMP
Dual-Core AMD Opteron Processor 2435
Number of Processes Running
ARM1176
SCHED_OTHER
3.2.14
Buffer Ram
Level One ICACHE Size
Dual-Core AMD Opteron Processor 8214
Dual-Core AMD Opteron Processor 8214
Free Swap
13.1.0
3.0.27

Kernel

SVM

Applying the SVM



Dual Core AMD Opteron Processor 2
 Dual Core AMD Opteron Processor 8
 Level Three Cache Associativit
 Quad-Core AMD Opteron Processor 2387
 2.6.32-358.23.2.el6.x86_64
 Level One ICache Associativity
 2.6.32-431.5.1.el6.x86_64
 Quad-Core AMD Opteron Processor 2374
 traditional_with_pset_runqueue
 Intel Xeon CPU X5472 @ 3.00GHz Intel(R)
 Dual-Core AMD Opteron Processor 2218
 Processor Frequency
 Level One DCache Line Size
 Level Two Cache Size
 Fifteen Min Load
 3.10.37+
 Intel Xeon CPU E3-1410 @ 3.40GHz
 Total Swap
 armv6l
 Darwin
 Linux
 ppc6

Training & Evaluation Methodology

> 45k micro-benchmark executions

> 1,000 full application executions (2 apps)


Multiple hardware types (ARM, x86, PowerPC)

Multiple operating systems (Linux, OS X, BSD)

Training / Testing


- 20% of both micro-benchmark and application data used for training
- 80% used for testing

Test App I




Matrix Multiply

Test App 2




Rabin-Karp String Search


Correct Classification by M/M/I




Correct Classification by M/D/I



Test Classification Rates M/M/I



Test Classification Rates M/D/I





RaftLib: <http://raftlib.io>

My Page: <http://cse.wustl.edu/~beardj>

Lab Page: <http://sbs.wustl.edu>

