

Commercial Computing and High Performance Computing How Different are They

Don DeSota

Introduction

- Commercial and HPC computing are viewed as being very different
- I will compare and contrast the requirements and characteristics of commercial and HPC computing

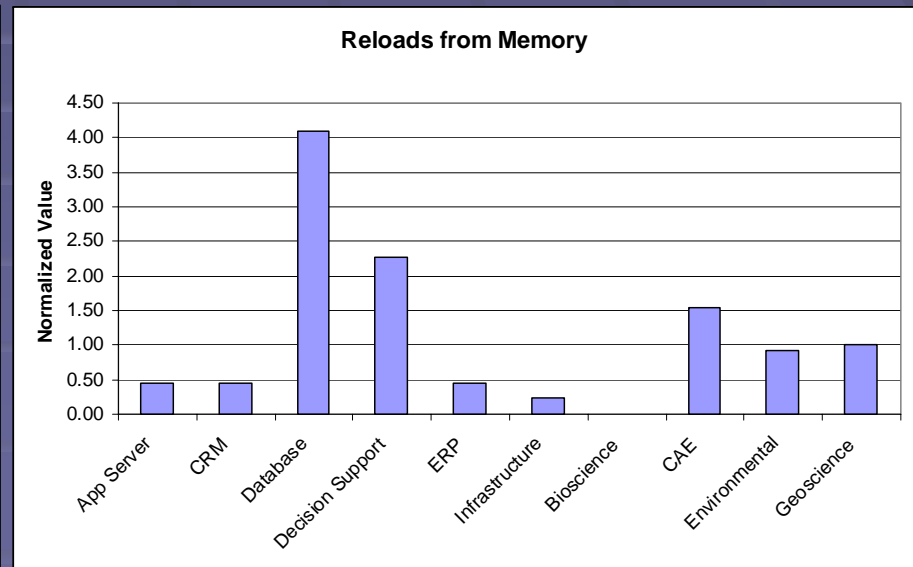
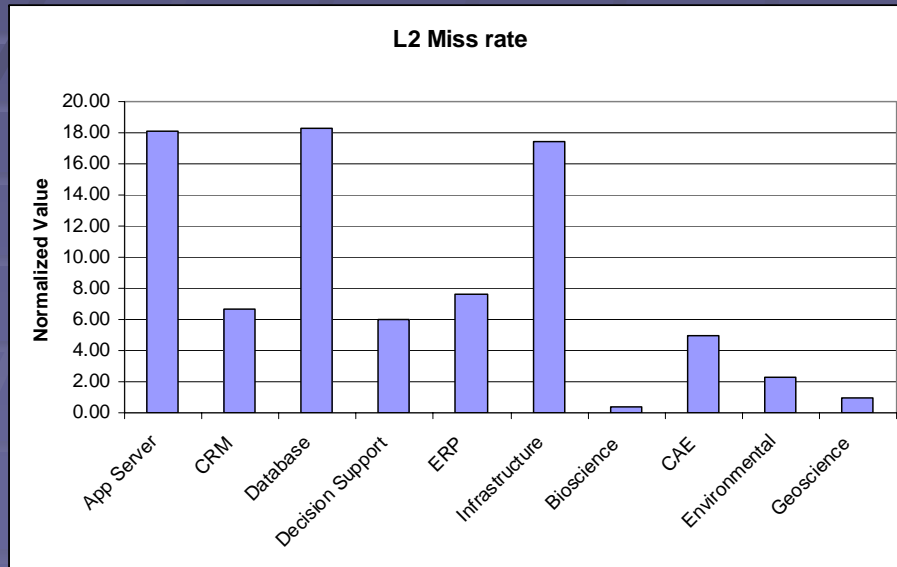

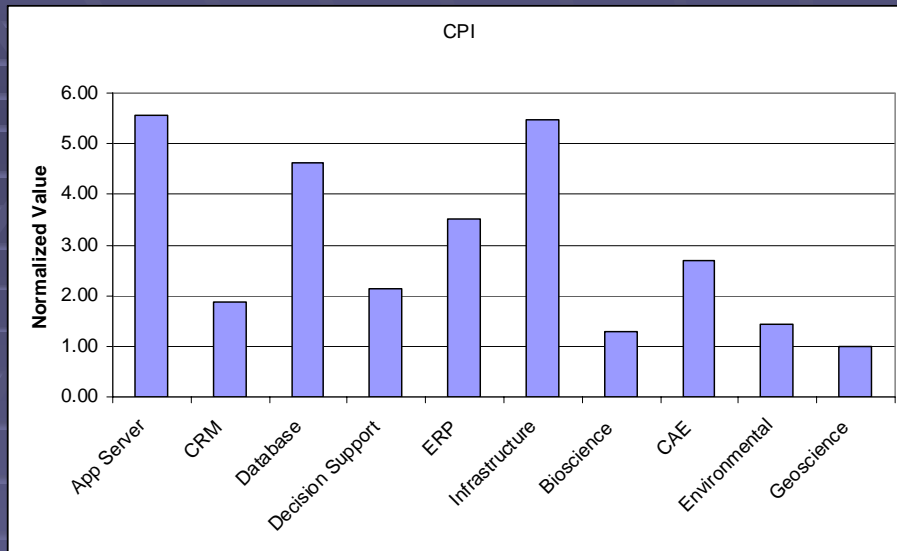
Outline

- Overview of the market segments
- Performance and system characteristics of the solution stack
- Future directions and challenges
- Conclusions

	Commercial	HPC
Market	<p>Multiple Segments</p> <ul style="list-style-type: none"> ▪ OLTP ▪ DSS ▪ CRM ▪ XML ▪ Infrastructure ▪ Web serving ▪ ERP ▪ Application serving 	<p>Multiple Segments</p> <ul style="list-style-type: none"> ▪ Life Sciences ▪ Environmental ▪ EDA ▪ CAE ▪ Geosciences ▪ Environmental ▪ Financial ▪ DCC ▪ Security ▪ Defense ▪ Imaging ▪ Government labs
	<p>Purchase Criteria</p> <p>Performance</p> <p>Cost</p> <p>Power</p> <p>Low Risk</p> <p>Current Investment</p> <p>Manageability</p>	<p>Purchase Criteria</p> <p>Performance</p> <p>Cost</p> <p>Power</p> <p>Compute Density</p> <p>Current Investment</p>
	<p>Servers from \$5K to \$10M</p>	<p>Servers from \$5K to \$100M</p> <p>Not just large government lab systems</p>

	Commercial	HPC
Requirements	Throughput Oriented Many threads per core Memory capacity	Capability Bandwidth Intensive Communication Single thread per core Throughput Light communications Single thread per core
	Scale-up or Scale-out Tending toward scale-out	Scale-out
	Mission Critical/RAS I/O Intensive	Cost/performance and power performance sensitive
Trends	Pervasive Web oriented SOA Data Explosion	Mission Critical/RAS Hybrid programming model Specialized processors Autoparallelization Grid computing

Performance Characteristics



Performance Characteristics

- Commercial workloads tend to have higher CPIs
 - Decision support looks similar to HPC applications particularly CAE
- Miss rates in commercial tend to be higher than HPC
- Branch misprediction between commercial and HPC is similar
- Memory bandwidth in HPC is generally higher for HPC
- Decision support has similar characteristics to some of the HPC segments and has similar machine requirements to CAE
- Certain subsegments of some HPC segments such as some biosciences applications have many attributes of commercial applications

Solution Stack

Layer	Commercial	HPC
Processor	High Branch Rate High cache miss rate <ul style="list-style-type: none">▪ Needed for both instructions and data	Floating Point Pipes High ILP Prefetch <ul style="list-style-type: none">▪ Uniform Strides Data Cache Friendly <ul style="list-style-type: none">▪ Algorithmic changes to make more effective
Memory	Capacity Limited Latency Sensitive	Bandwidth Limited A few are latency sensitive Low Bandwidth Capacity not an issue

Solution Stack

Layer	Commercial	HPC
Storage	<p>Large Volumes of Data</p> <p>Simultaneous IOs critical</p> <p>DSS Bandwidth critical</p> <p>Cause the need for many threads</p>	<p>Large amounts of data in a few segments</p> <p>Bandwidth critical</p> <p>Latency hidden</p>
Interconnect	Minimal	Heavy
RAS	<ul style="list-style-type: none"> ▪ Critical ▪ Provisioning 	<ul style="list-style-type: none"> ▪ Achieved more through the application and scale out ▪ Emerging ▪ Provisioning

Solution Stack

Layer	Commercial	HPC
Virtualization	Critical	Not Needed
Management	Provisioning	Provisioning
OS	RAS	Slim
Middleware	RDBMS TP SOA Web Server .NET Samba CGI	MPI MPI Runtime Scheduler
		Compilers and Libraries

Solution Stack

Layer	Commercial	HPC
Application		Algorithm
	SQL	Fortran
	Java	C or C++
	Ruby	
	Small Basic Block	Large Basic Block

Future Directions and Challenges

- Systems are judged on more than performance
 - The next frontier for modeling is to characterize and project other critical parameters such as power, reliability, usability, etc.
 - Performance and power trade-offs will become critical
- Multicore will impact applications
 - Many cores per node will need to be exploited by the software
 - Commercial applications will start to see more bandwidth bottlenecks
 - HPC applications will move toward a hybrid solution
 - This is also a good solution for DSS
- The line between HPC and commercial will blur
 - HPC applications will combine HPC and database
 - Commercial applications will include more computation
- Specialized processors will provide performance gains
 - They may be used for certain applications, to offload main processors or in hybrid systems
- More businesses are adopting HPC as being more cost effective

Conclusions

- Both areas have a wide range of price points and segments with different characteristics that drive different machine characteristics within the area
- DSS has similarities to HPC but the rest of traditional commercial computing has much different requirements than HPC
- Search and analysis is more similar to HPC but still does not have the interconnect requirements of HPC
- Techniques used to improve systems in HPC or commercial should be applied to the other segment to get improvements as well
- There is overlap between commercial and HPC characteristics but some differences exist
- There are plenty of opportunities to move characterization and modeling to aspects of the system other than performance