UKCMG Annual Conference 2006

Hardware Architecture & Performance

Colin Butcher

Technical Director, XDelta Limited

Overview

Overall system performance depends on the underlying hardware architecture, the operating system behaviour and the system configuration. These all affect the design of application software and how well it performs, especially in a distributed environment.

Software developers need to understand the underlying behaviour of the systems in order to write code that is reliable, which has minimal performance impact and which scales well in a production environment.

Topics (1)

- **Basic principles**
- **The functional elements in a computer system**
- **History and development**
- **What do we mean by “hardware”?**
- **CPUs (Central Processing Units) & Multiprocessors**
- **Memory subsystem**
- **IO subsystem**
- **Partitioning**

Topics (2)

- **Interconnects**
- **Bus structures**
- **Performance characteristics**
- **Environmental characteristics**
- **Software implications:**
  - **Operating systems**
  - **Applications**

Basic principles of performance

- **Bandwidth – determines throughput**
- **Latency – determines response time**
- **Jitter – variation of latency with time**

- **Parallelism and reduction in "wait states"**
- **Synchronisation of access to data structures**
- **Serialisation of access to data structures**

- **Scalability**

Basic limits to performance

- Speed of light
- Feature size in “chip”
- Signal degradation
- Power consumption
- Heat dissipation
- Chemical impurities
- Mechanical contacts
- Manufacturing processes

Components of a computer system

- CPU (Central Processing Unit)
- Memory subsystem
- IO subsystem
- Storage subsystem
- Network connectivity

- Power
- Cooling
- EMI and RFI protection
- Mechanical protection

CPUs (1)

- **Processes instructions**
- **Manipulates data**
- **Instruction Set**
- **Protection mechanisms (processor modes)**
- **Clocking**
- **“bit width” (16bit, 32bit, 64bit …)**
- **Threads**
- **Multi-core “chips”**

CPUs (2)

- **Performance:**
  - **Registers**
  - **Cache**
  - **Pipelines**
  - **Parallel execution**
- **CISC (Complex Instruction Set Computer)**
- **RISC (Reduced Instruction Set Computer)**
- **EPIC (Explicitly Parallel Instruction Computing)**
- **Compilers are the key to performance**

Memory subsystem

- **Used to be small, slow and expensive, so tried hard to minimise memory usage**
- **Now plentiful, relatively fast and relatively cheap, so use memory to offset other constraints and gain performance**
- **Caches and synchronisation**
- **Error detection and correction**
- **Memory latency is one of the major constraints**
- **“Locality” to CPUs (NUMA effects)**
- **Memory interconnects to CPUs**

IO (Input / Output) subsystem

- **Provides interface with outside world (physical, electrical, logical …)**
- **Connects to storage subsystems**
- **Access times (memory, cache etc.)**
- **“Locality” to CPUs and IO processing**
- **IO devices operation (interrupts, registers, DMA)**
- **IO interconnects to CPU and memory subsystems**
- **Bus structures (eg: PCI-X)**
- **Bandwidth to storage devices (eg: fibrechannel)**

Multiprocessor systems

- Asymmetric Multiprocessing (attached processor)
- Symmetric Multiprocessing (SMP)
- Interconnections (switch, mesh, toroid etc.)
- Latency and Bandwidth
- Synchronisation and Serialisation
- IO processing and interrupt handling
- Partitioning
- “Locality of resources” and operating system constructs

Network interconnects

- All networks are multi-protocol networks
- Different protocols have different requirements in terms of bandwidth and latency
- Different protocols have different characteristics, especially with how they handle path load balancing and path failure
- Beware of complex network configurations
- Bandwidth to core switches
- Bandwidth between core switches
- Latency issues with WAN links – packet encapsulation and ‘managed services’ can create problems

Storage subsystems

- **Bandwidth to disc arrays over fibrechannel is now the primary constraint**
- **Disc arrays manage the allocation of data to disc blocks to maximise performance**
- **Storage interconnects to IO subsystem**
- **The “backup window” is a major issue in many production systems, especially 24x7 systems**
- **Backups need to be file structured, not block structured**

Configuring systems

- **Cost – need to understand trade-offs and benefits**
- **Performance**
- **Availability**
- **Maintainability**
- **Safety**
- **Simplicity**
- **Business continuity and Disaster recovery**
- **Representative test environment – both scale and function**

Software considerations

- Operating system design
- System configuration
- Performance tuning
- Application design (scalability, data integrity, availability)
- Compilers generate code better than you can
- Portability
- Reliability
- Supportability
- Testing and documentation

Thank you for you attention. Any questions?

Colin Butcher, XDelta Limited

Office: +44 117 904 8209
Cellphone: +44 7768 857615
E-mail: colin.butcher@xdelta.co.uk
Web: www.xdelta.co.uk