

The MVP

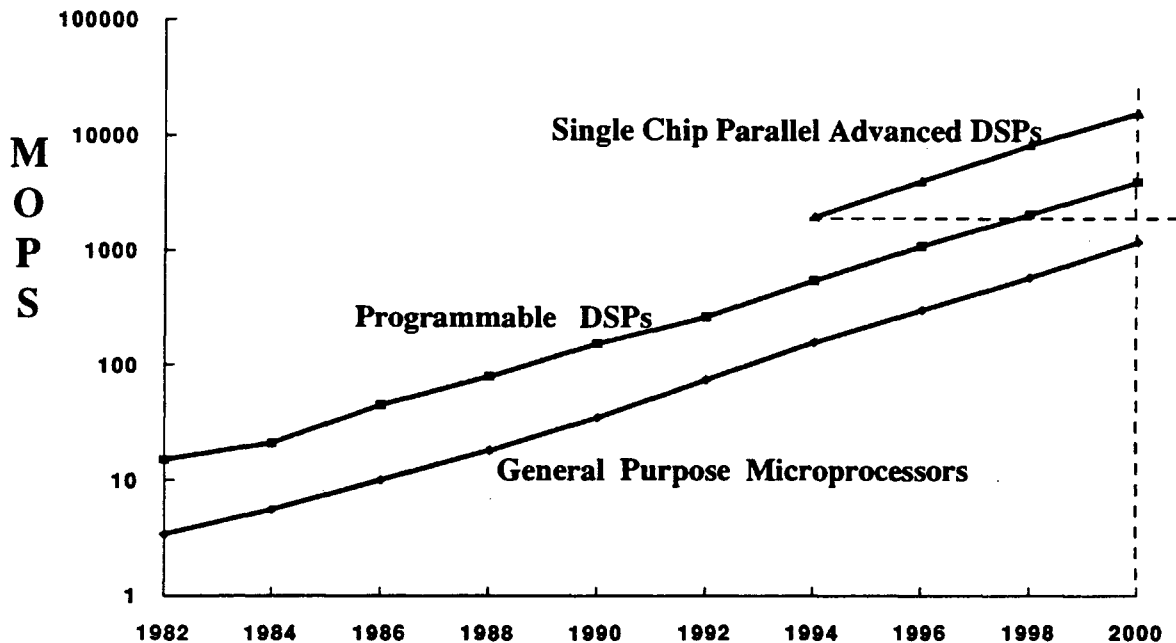
*(Or A Single Chip Crossbar Shared Memory
Heterogeneous MIMD Multi-Processor
That Can, Among Many Other Things, Perform
Video Compression/Decompression)*

By Karl M. Guttag
TI Fellow
Texas Instruments

KMG - MVP_PR13

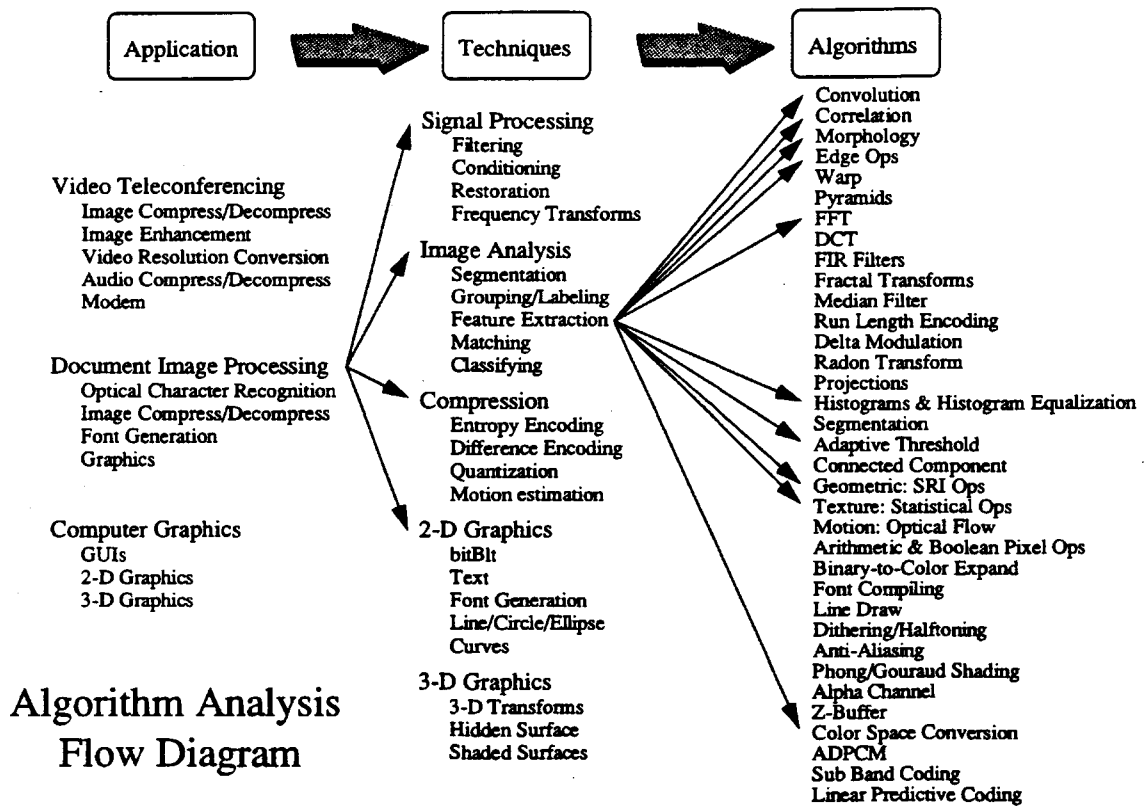


Trends In Programmable Processor Performance




KMG - MVP_PR13





What Is Done Poorly Can Dominate Performance



Programmable Multimedia Processor Requirements

Massive amounts of imaging, graphics, and signal processing

- Requires parallel processing (of some form)

Mixture of processing needs

- For high resolution still images, full motion video, 3-D graphics, and Audio
- Requires massive integer and fast floating point math
- Heterogeneous processing requirement

Flexible enough to handle a range of algorithms

- More than just compression standards
- Algorithms are in a state of flux (new ones constantly appearing)
- Algorithms vary dramatically even for the same task (ex Video Compression)

Very high bandwidth

- Wide and fast on-chip memory
- Large amount of on-chip memory to reduce off-chip bandwidth
- Low loss of bandwidth due to inter-processor communications

KMG - MVP_PR13



MIMD Inter-processor Connection Models

Pipeline (ex. Graphics)

- Bucket brigade processing

Multiple Channel Mesh/Array/Hypercube (Transputor, 320C40)

Mixed Private and Bus Shared Memory (Workstations use today)

- Single shared memory with only one processor access at a time
- Most processing out of private memory
- Typically multiple processors to 1 shared memory


Fully Shared Memory

- Multiple processors to multiple shared memories
- Cycle by cycle crossbar connection
- Requires huge busses
- Becomes practical only if the CPUs and memory are on the same chip


KMG - MVP_PR13




Parallel Processing Alternatives



Mixed dedicated & programmable processing




SIMD




MIMD


KMG - MVP_PR13



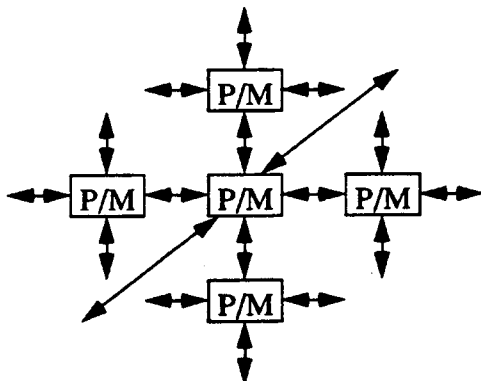
MIMD Processing Models




Pipeline



Bus Shared Memory



Comm. Port (Mesh/Array/Hypercube)



Crossbar Fully Shared Memory

Processing Models and Their Limitations

Mixed hardwired/programmable supports only a few Algorithms

SIMD only works well for massive uniform processing

- Very poor at decision making

Pipeline model is restrictive to algorithms that can be pipelined

Communication port systems are expandable but:

- Require that the communication of data is limited (ie coarse grain parallelism)
- The type of processing must be mappable onto the communication structure

Bus Shared Memory requires coarse grain parallelism

- Thrashes on access to shared memory otherwise.

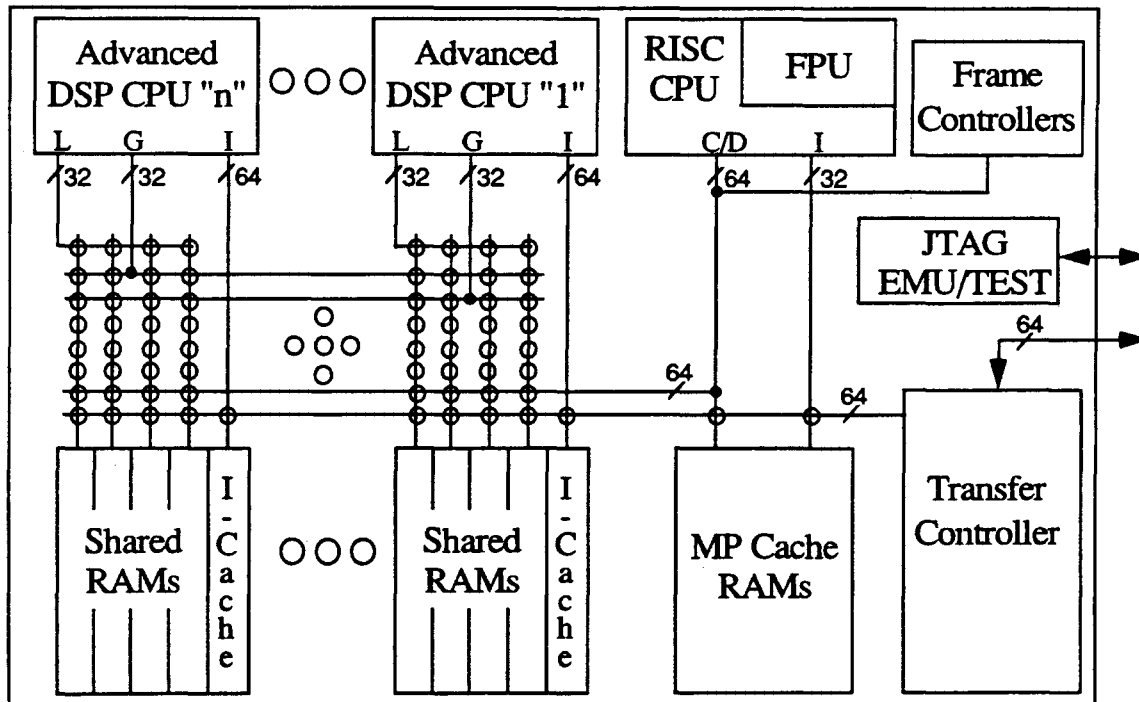
Fully shared memory requires a large crossbar bus

- Can be done economically only on a single chip


KMG - MVP_PR13

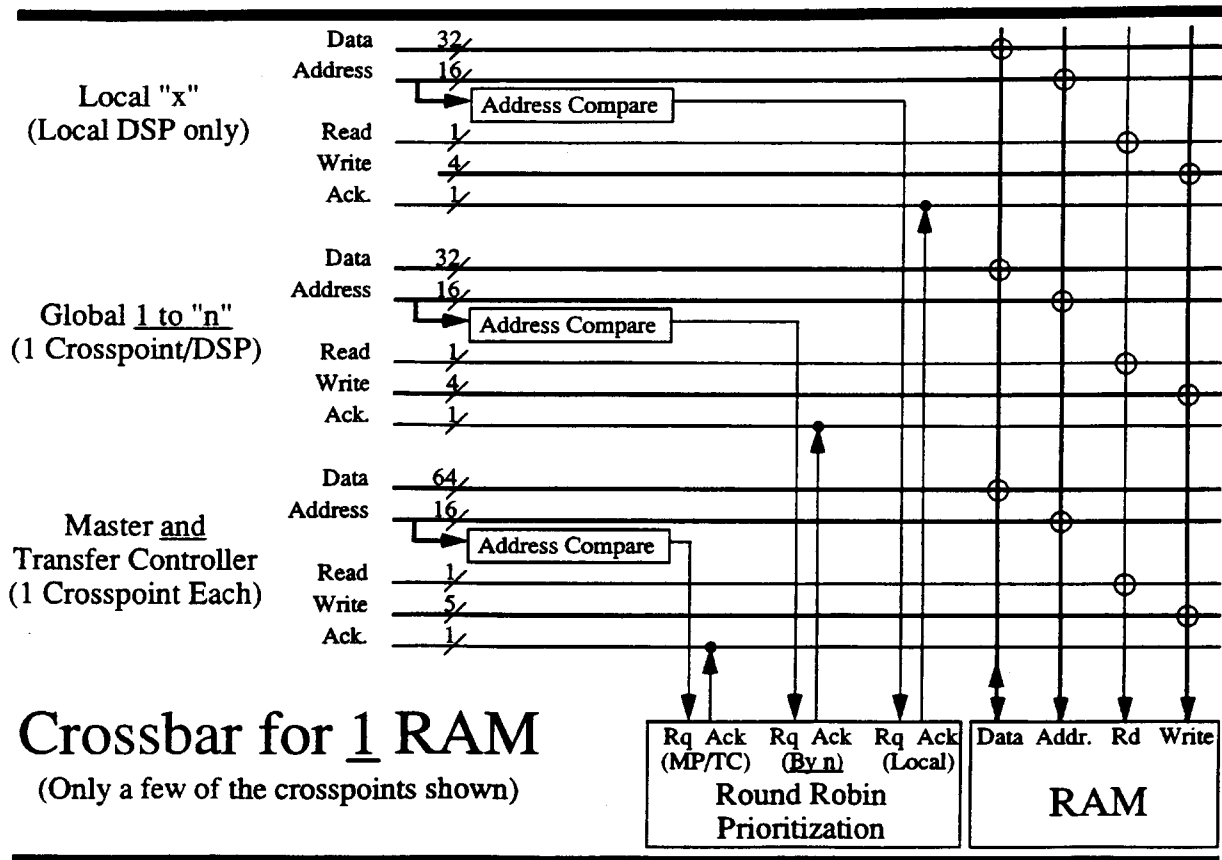


MVP Data and Instruction Routing



KMG - MVP_PR13





KMG - MVP_PR13



MVP Use Of Caches and Software Managed Memory

Master Processor has instruction and data caches

- For ease of programming in high level languages

Parallel DSPs have instruction caches

Shared memory directly mapped and not cached

- Common Reference by all processors
- No need for coherency detection/protection

Shared memory is software managed by any/all processor(s)

- Anticipatory loading of data and background saving of results by Transfer Controller
- Usually transferred in packets (ex. blocks) of data

Transfer Controller manages all memory traffic

- Autonomous transfers of multi-dimensional data packets
- Performs cache miss and data packet transfers
- Prioritizes requests for transfers
- Controls round robin access to each shared RAM

Processing Models and Their Limitations

Mixed hardwired/programmable supports only a few Algorithms

SIMD only works well for massive uniform processing

- Very poor at decision making

Pipeline model is restrictive to algorithms that can be pipelined

Communication port systems are expandable but:

- Require that the communication of data is limited (ie coarse grain parallelism)
- The type of processing must be mappable onto the communication structure

Bus Shared Memory requires coarse grain parallelism

- Thrashes on access to shared memory otherwise.

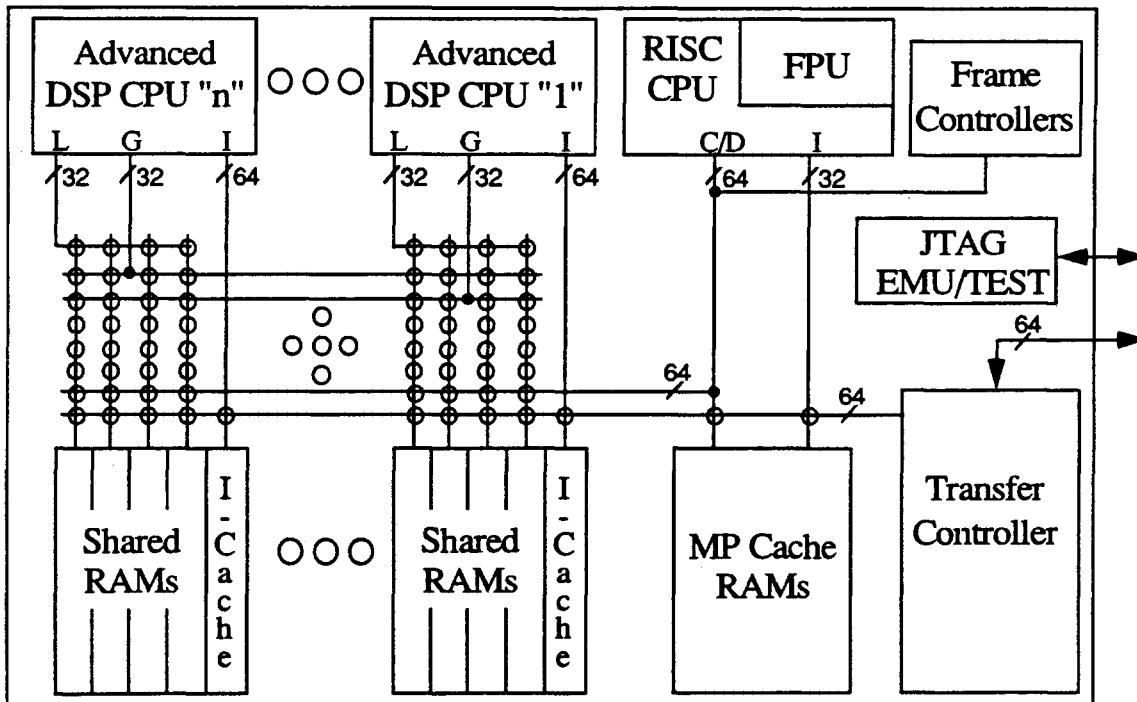
Fully shared memory requires a large crossbar bus

- Can be done economically only on a single chip

KMG - MVP_PR13

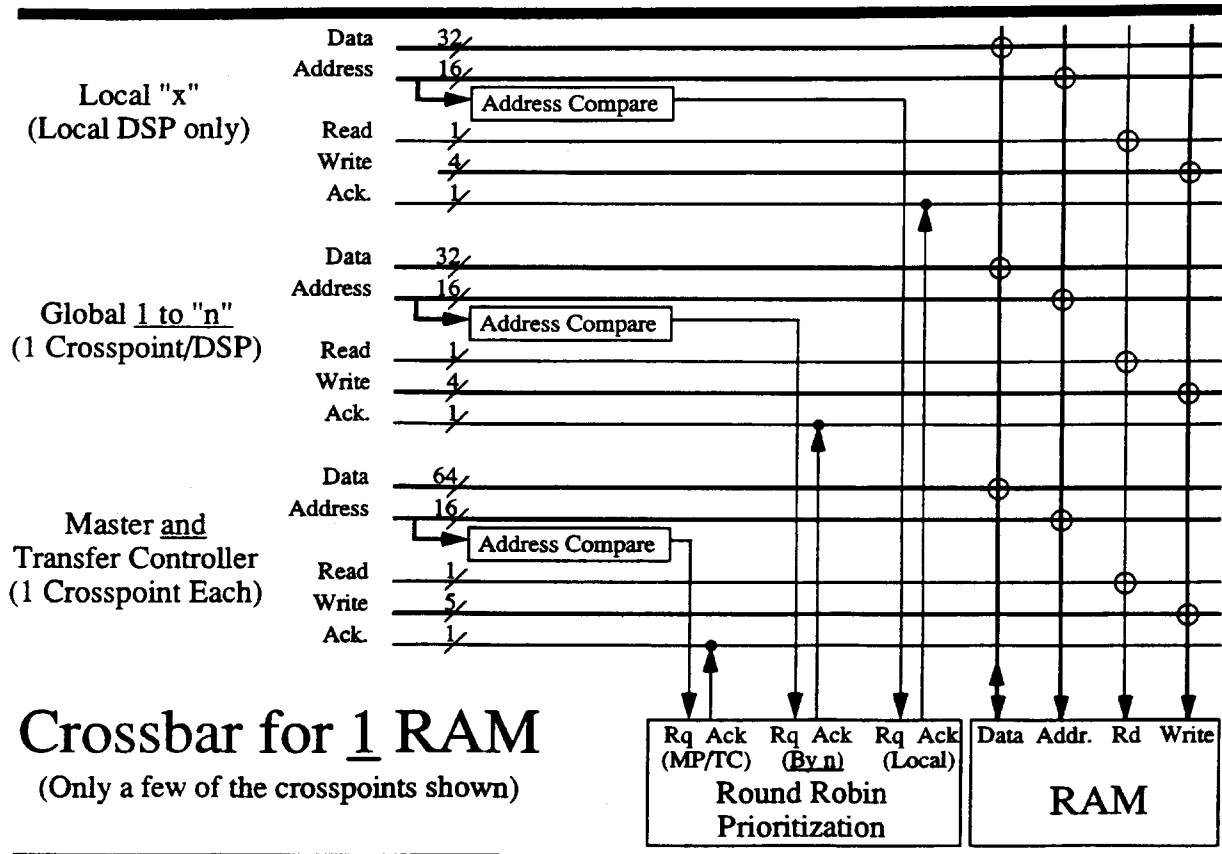


MVP Data and Instruction Routing



KMG - MVP_PR13





KMG - MVP_PR13



MVP Use Of Caches and Software Managed Memory

Master Processor has instruction and data caches

- For ease of programming in high level languages

Parallel DSPs have instruction caches

Shared memory directly mapped and not cached

- Common Reference by all processors
- No need for coherency detection/protection


Shared memory is software managed by any/all processor(s)

- Anticipatory loading of data and background saving of results by Transfer Controller
- Usually transferred in packets (ex. blocks) of data


Transfer Controller manages all memory traffic

- Autonomous transfers of multi-dimensional data packets
- Performs cache miss and data packet transfers
- Prioritizes requests for transfers
- Controls round robin access to each shared RAM


Flexible Parallel Programming Models With The MVP




Fully Shared Parallel



Pipelined



Master with Parallel Servers

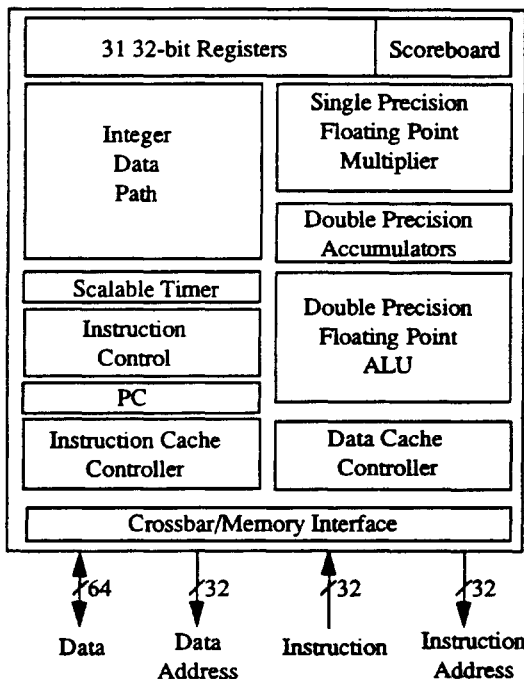


Mixed Separate Tasks

KMG - MVP_PR13



Master Processor



32-Bit RISC Processor
Load/Store Architecture

31 32-Bit General Purpose Registers
Plus R0=0
Scoreboarded for Loads and F.P. operations

IEEE-754 Floating Point Unit
Single Cycle Add/Sub (single or double)
Single Cycle Multiply (single)
Microcoded Multiply (double), Divide, Sq. root

Parallel Floating Point Instructions
Multiply || Add || 64-bit LD/ST(p++) in 1 cycle

Delayed Branches With 1 Delay Slot

15-Bit or 32-Bit Immediate Constants

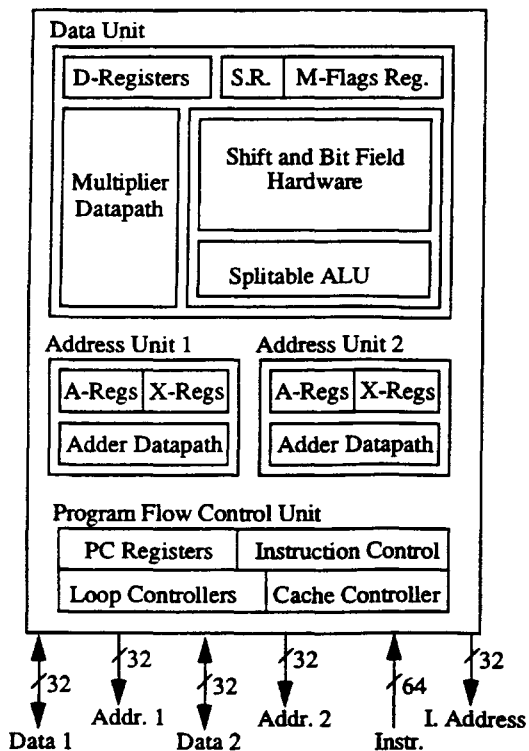
Left-most and Right-most one logic

Instruction and Data Caches

KMG - MVP_PR13



Advanced DSP CPUs



Parallel Execution Units

- Multiplier
- ALU (Splitable at 8- or 16-bit boundaries)
 - SIMD operation within an instruction
- 2 Independent Address Units
- Programmable loop controllers

Shift and bit field hardware

- Barrel Rotator
- Mask Generator
- 1-to-n Bit Expander
- LMO, RMO, LMBC, RMBC Bit Detectors

Two data and one instruction ports

Large Instruction Words (64-bit)

- Supports parallel independent operations

KMG - MVP_PR13



Texas
Instruments

Achieving The MVP's Performance

Aimed at tasks that run well on parallel processors

- I.E. not aimed a running existing binary object codes

Multiple on-chip RAMs to provide very high on-chip bandwidth

Crossbar to preserve bandwidth

Advanced DSP CPUs

- Parallel execution units
- Split ALU and bit expander hardware
- Barrel rotation, bit field masking, and bit detection hardware
- Hardware loop controllers

RISC Master Processor with FPU

- Prevents DSPs from being burdened with other system functions
- Parallel Floating Point operations (DSP like)

Transfer Controller (Very Intelligent DMA Controller)

- Autonomously controls transfers of data between memory spaces

