SC10:

A Video Processor And Pixel-Shading GPU for Handheld Devices

Edward Hutchins

What is SC10?

- **Companion chip for handheld devices**
  (Cell-phones, PDAs, dedicated game devices)
- **Assumes external host processor + memory**
- **Accelerates image, video, 2D and 3D processing**
- **Designed for battery-powered applications**
- **Small package**
- **Host interface scales from 8-bit to full 32-bit I/O**

SC10 Block Diagram

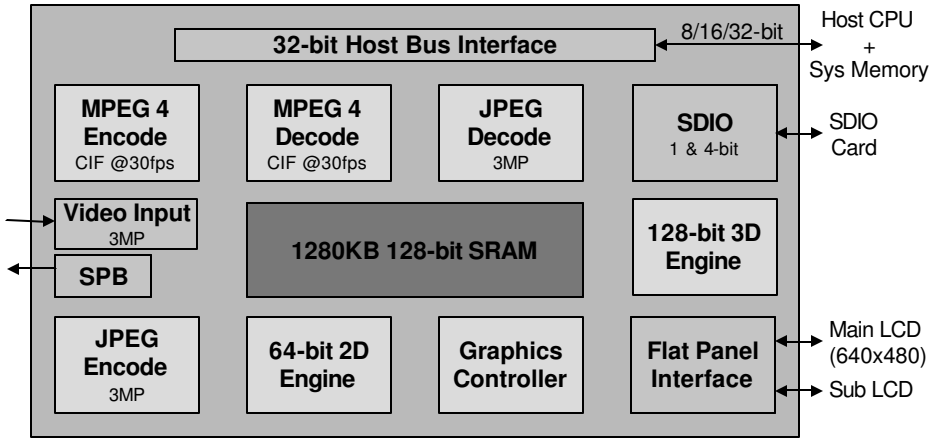

Multimedia Features

- **Full-duplex hardware MPEG4 codec**
- **CIF MPEG4 encode @ 30fps**
- **CIF MPEG4 decode @ 30fps**
- **MPEG4 simple profile level 0, 1, 2, 3**
- **3MP camera support**
- **1.3MP preview @ 24fps**
- **640x480 display support**
- **72MHz camera clock**
- **Serial Peripheral Bus for camera control**

3D Design Goals

- Efficient use of power
- Efficient use of area
- Broad API support
  - OpenGL-ES
  - D3D Mobile
  - non-3D uses (signal/image processing)
- Modular, expandable architecture
  - quick response as the market grows
  - simple, consistent interfaces and data-types
  - fully programmable to support evolving APIs

3D Pipeline Overview

- Setup – prepares triangles for rasterization
- Raster – interpolation of parameters
- Gatekeeper – scoreboarding and data flow control
- Data Fetch (DF) – color, depth, texture data reads
- Arithmetic Units (ALUs) – blending/combiner ops
- Data Write (DW) – color & depth writeback

3D Block Diagram

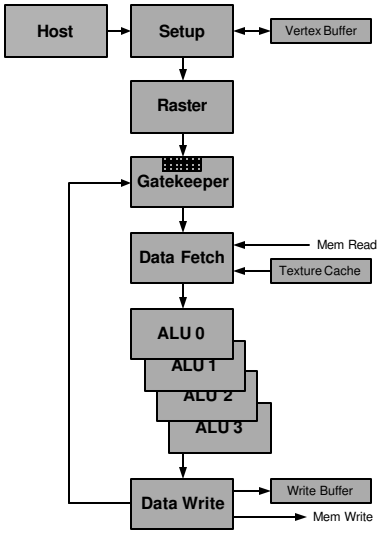

Setup Unit

- Simple packet-based host interface
- 32-bit IEEE float, S15.16 fixed and packed .8 inputs
- Up to 24 parameters + x,y,z,w per triangle
  - meaning of parameters left up to software
- Large vertex cache (software controlled)
- Performs simple transform, clip & viewport (1/w)
- Culls back-facing triangles
- Calculates interpolated LOD quantity per vertex

Raster Unit

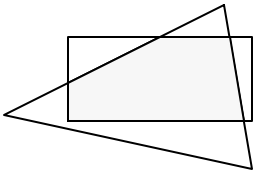

- Iterates intersection between triangle and scissor rectangle
- Supports “free” guard-band clipping
  - reduces complex clipping cases by 10x
- Follows OGL-ES/D3DM rasterization rules
- Generates pixel packets used by rest of pipe
  - 4 high-precision and 4 low-precision perspective-correct iterated values per row (plus Z for “free”)
  - span location (X,Y) sent via SPAN_START packet
  - even/odd pixels interleaved to hide ALU latency
- Reduces precisions as early in pipeline as possible

Pixel Packet

Packet contains iterated data and sideband information about even/odd, kill and instruction sequence

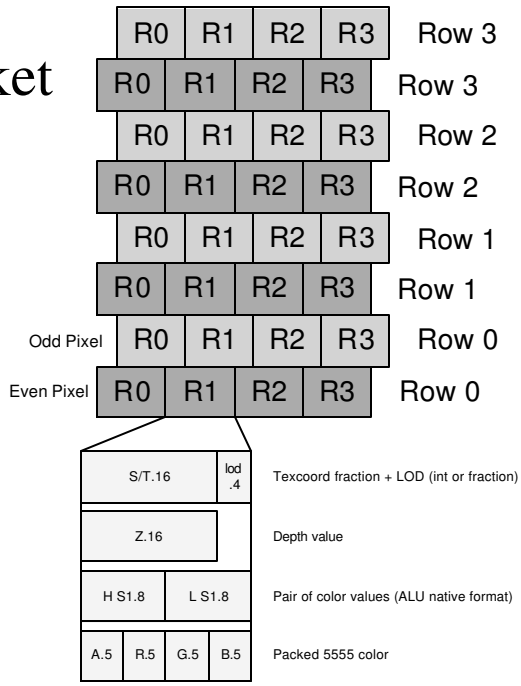

Gatekeeper (GK) Unit

- Controls pixel packet flow via counting
  - keeps pipeline as full as possible during recirculation
  - tracks recirculated X,Y positions
- Prevents coincident pixels from entering shader
  - scoreboarding via bitmap marks X,Y pairs busy
- Detects pipeline idle condition
- Collects debug readback info from shader blocks
- Maintains cache coherency

Data Fetch (DF) Unit

- 8 instructions (e.g. 4 rows w/ 2 recirculations)
- 8 surfaces (color, depth and up to 6 textures)
- One depth fetch per row (typically one per pixel), AND
- One color or filtered texture fetch per row
  - supports RGBA, palettized and compressed formats
  - one clock nearest/bilinear, two clocks per trilinear
  - replaces texcoords with fetched 40-bit color (S1.8)
- Manages color/depth and texture caches separately
- Performs depth test, marking pixel killed on failure
  - Z-write deferred till semantically correct place

Arithmetic Logic (ALU) Units

- 8 instructions
- Performs a*b + c*d on one S1.8 scalar per row
- Operates on 4 20-bit R variables from row
- Can use 2 20-bit constants per unit (read-only)
- Can also use 2 20-bit temporary values per unit

  consts/temps can be 4 values in S1.8 format,
  8 values if 5555 packing is used
  both are software loadable
- Can complement/negate operands, clamp result
- 4 ALU units in SC10 (typically one per channel)

ALU Block Diagram

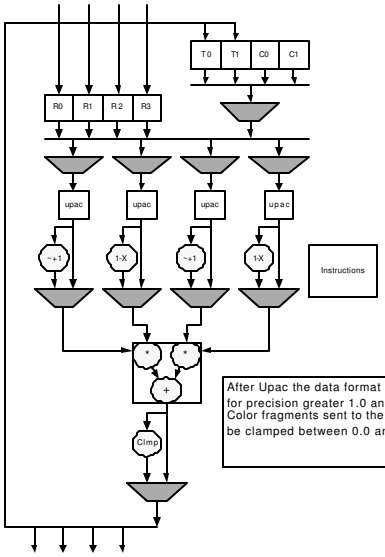


```
Instructions

MAD  r = a*b + c*d
MBA  r = a*b & c*d
MBO  r = a*b | c*d
MBX  r = a*b ^ c*d
MUL  r(lo) = a*b  r(hi) = c*d
MIN  r = min(a*b, c*d)
MAX  r = max(a*b, c*d)
SNE  r = (a*b) != (c*d) ? 1 : 0
SEQ  r = (a*b) == (c*d) ? 1 : 0
SLT  r = (a*b) <  (c*d) ? 1 : 0
SLE  r = (a*b) <= (c*d) ? 1 : 0
KNE  kill if (a*b) != (c*d)
KEQ  kill if (a*b) == (c*d)
KLT  kill if (a*b) <  (c*d)
KLE  kill if (a*b) <= (c*d)
```

After Upac the data format is S1.8 allowing for precision greater 1.0 and less than 0.0 . Color fragments sent to the frame buffer will be clamped between 0.0 and 1.0

Data Write (DW) Unit

- 8 instructions
- Write coalescing buffers for color and depth
  - can write to all 8 surfaces at lower performance
- Optionally dithers to 565 color
- Suppresses killed pixel writes
- Indicates retired writes to GK for scoreboarding

SC10 Chip Plot

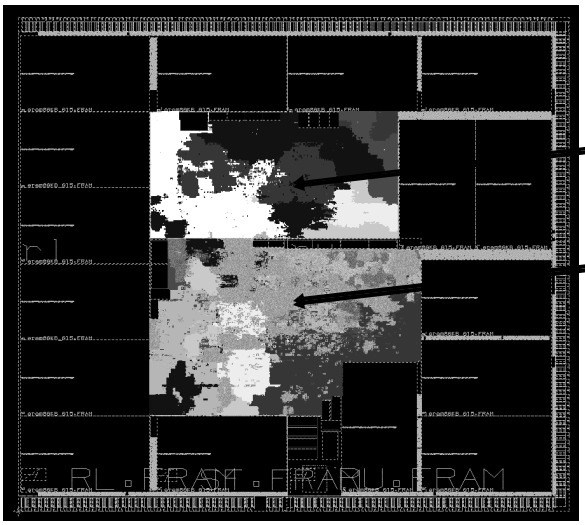

Power Reduction Techniques

- Hardware Co-Processor Engines
  - 2D/3D Graphics Engines, Graphics Controller Video Input Port, Flat Panel Interface
- Low Leakage Process
- Embedded Memory
- Relaxation Oscillator
- Fully Asynchronous Clocking
- Dynamic Clock Switching
- Sub-Module Clock Gating
- Automatic Pipeline Shutdown

Multimedia Statistics

- 6.8M transistors in UMC 0.15µ LL

Use Case	Conditions	Core Power (mW)
1.3MP JPEG Decode	15fps @ 1.3MP Continuous Decode	10.6 mW
1.3MP JPEG Encode	15fps @ 1.3MP Continuous Encode	19.9 mW
MPEG4 Decode	CIF @ 30fps	13.1 mW
MPEG4 Encode	CIF @ 30fps	28.1 mW

CVDD = 1.5V FVDD = 3.3V	QVGA TFT LCD Display 1.3MP CMOS Camera	BVDD = 3.3V VVDD = 2.8V

3D Statistics

- 3.7M transistors in UMC 0.15µ LL
- Capable of:
  - useful pixels at 1 clock/pixel (color + texture + depth)
  - complex at 2 clocks/pixel (e.g. + blending or 2nd texture)
  - supporting 4 textures with all OGL-ES features enabled
- Measured pixel fill-rates of 96% theoretical peak on XSCALE/ Accelent development system
- Given 1 vertex/triangle, can draw ~1M tris/sec at 72MHz (988k triangles/sec measured on XSCALE/ Accelent system)
- Preliminary core power measurements at 1.5V/72MHz show apps (furry demo, Quake II) consume ~50-75mW at 30FPS
  - Drivers not yet tuned for power

3D Examples