
Gallium Documentation

Release 0.4

VMware, X.org, Nouveau

Oct 09, 2017

Contents

1	Introduction	3
1.1	What is Gallium?	3
2	Debugging	5
2.1	Debug Variables	5
2.2	Remote Debugger	6
3	TGSI	7
3.1	Basics	7
3.2	Instruction Set	8
3.3	Explanation of symbols used	45
3.4	Other tokens	46
3.5	Texture Sampling and Texture Formats	56
4	Screen	59
4.1	Flags and enumerations	59
4.2	Methods	70
4.3	Thread safety	72
5	Resources and derived objects	73
5.1	Transfers	73
5.2	Resource targets	73
5.3	Surfaces	76
5.4	Sampler views	76
6	Formats in gallium	77
6.1	References	78
7	Context	79
7.1	Methods	79
7.2	Using several contexts	89
8	CSO	91
8.1	Blend	91
8.2	Depth, Stencil, & Alpha	93
8.3	Rasterizer	94
8.4	Sampler	99

8.5	Shader	100
8.6	Vertex Elements	100
9	Distribution	103
9.1	Drivers	103
9.2	State Trackers	104
9.3	Auxiliary	105
10	Drivers	107
10.1	Freedreno	107
10.2	OpenSWR	115
11	Glossary	123
12	Indices and tables	125

Contents:

What is Gallium?

Gallium is essentially an API for writing graphics drivers in a largely device-agnostic fashion. It provides several objects which encapsulate the core services of graphics hardware in a straightforward manner.

Debugging utilities in gallium.

Debug Variables

All drivers respond to a set of common debug environment variables, as well as some driver-specific variables. Set them as normal environment variables for the platform or operating system you are running. For example, for Linux this can be done by typing “export var=value” into a console and then running the program from that console.

Common

GALLIUM_PRINT_OPTIONS Type: bool **Default: false**

This option controls if the debug variables should be printed to stderr. This is probably the most useful variable, since it allows you to find which variables a driver uses.

GALLIUM_RBUG Type: bool **Default: false**

Controls if the *Remote Debugger* should be used.

GALLIUM_TRACE Type: string **Default: “”**

If set, this variable will cause the *Trace* output to be written to the specified file. Paths may be relative or absolute; relative paths are relative to the working directory. For example, setting it to “trace.xml” will cause the trace to be written to a file of the same name in the working directory.

GALLIUM_DUMP_CPU Type: bool **Default: false**

Dump information about the current CPU that the driver is running on.

TGSI_PRINT_SANITY Type: bool **Default: false**

Gallium has a built-in shader sanity checker. This option controls whether the shader sanity checker prints its warnings and errors to stderr.

DRAW_USE_LLVM Type: bool **Default: false**

Whether the *Draw* module will attempt to use LLVM for vertex and geometry shaders.

State tracker-specific

ST_DEBUG Type: flags **Default: 0x0**

Debug *Flags* for the GL state tracker.

Driver-specific

I915_DEBUG Type: flags **Default: 0x0**

Debug *Flags* for the i915 driver.

I915_NO_HW Type: bool **Default: false**

Stop the i915 driver from submitting commands to the hardware.

I915_DUMP_CMD Type: bool **Default: false**

Dump all commands going to the hardware.

LP_DEBUG Type: flags **Default: 0x0**

Debug *Flags* for the llvmpipe driver.

LP_NUM_THREADS Type: int **Default: number of CPUs**

Number of threads that the llvmpipe driver should use.

FD_MESA_DEBUG Type: flags **Default: 0x0**

Debug *Flags* for the freedreno driver.

Flags

The variables of type “flags” all take a string with comma-separated flags to enable different debugging for different parts of the drivers or state tracker. If set to “help”, the driver will print a list of flags which the variable accepts. Order does not matter.

Remote Debugger

The remote debugger, commonly known as rbug, allows for runtime inspections of *Context*, *Screen*, *Resources and derived objects* and *Shader* objects; and pausing and stepping of *Draw* calls. Is used with rbug-gui which is hosted outside of the main mesa repository. rbug is can be used over a network connection, so the debugger does not need to be on the same machine.

TGSI, Tungsten Graphics Shader Infrastructure, is an intermediate language for describing shaders. Since Gallium is inherently shaderful, shaders are an important part of the API. TGSI is the only intermediate representation used by all drivers.

Basics

All TGSI instructions, known as *opcodes*, operate on arbitrary-precision floating-point four-component vectors. An opcode may have up to one destination register, known as *dst*, and between zero and three source registers, called *src0* through *src2*, or simply *src* if there is only one.

Some instructions, like *I2F*, permit re-interpretation of vector components as integers. Other instructions permit using registers as two-component vectors with double precision; see *Double ISA*.

When an instruction has a scalar result, the result is usually copied into each of the components of *dst*. When this happens, the result is said to be *replicated* to *dst*. *RCP* is one such instruction.

Modifiers

TGSI supports modifiers on inputs (as well as saturate and precise modifier on instructions).

For arithmetic instruction having a precise modifier certain optimizations which may alter the result are disallowed. Example: *add(mul(a,b),c)* can't be optimized to TGSI_OPCODE_MAD, because some hardware only supports the fused MAD instruction.

For inputs which have a floating point type, both absolute value and negation modifiers are supported (with absolute value being applied first). The only source of TGSI_OPCODE_MOV and the second and third sources of TGSI_OPCODE_UCMP are considered to have float type for applying modifiers.

For inputs which have signed or unsigned type only the negate modifier is supported.

Instruction Set

Core ISA

These opcodes are guaranteed to be available regardless of the driver being used.

ARL (Address Register Load)

$$\begin{aligned}dst.x &= (int)[src.x] \\dst.y &= (int)[src.y] \\dst.z &= (int)[src.z] \\dst.w &= (int)[src.w]\end{aligned}$$

MOV (Move)

$$\begin{aligned}dst.x &= src.x \\dst.y &= src.y \\dst.z &= src.z \\dst.w &= src.w\end{aligned}$$

LIT (Light Coefficients)

$$\begin{aligned}dst.x &= 1 \\dst.y &= \max(src.x, 0) \\dst.z &= (src.x > 0) ? \max(src.y, 0)^{\text{clamp}(src.w, -128, 128)} : 0 \\dst.w &= 1\end{aligned}$$

RCP (Reciprocal)

This instruction replicates its result.

$$dst = \frac{1}{src.x}$$

RSQ (Reciprocal Square Root)

This instruction replicates its result. The results are undefined for $src \leq 0$.

$$dst = \frac{1}{\sqrt{src.x}}$$

SQRT (Square Root)

This instruction replicates its result. The results are undefined for $src < 0$.

$$dst = \sqrt{src.x}$$

EXP (Approximate Exponential Base 2)

$$\begin{aligned}
 dst.x &= 2^{\lfloor src.x \rfloor} \\
 dst.y &= src.x - \lfloor src.x \rfloor \\
 dst.z &= 2^{src.x} \\
 dst.w &= 1
 \end{aligned}$$

LOG (Approximate Logarithm Base 2)

$$\begin{aligned}
 dst.x &= \lfloor \log_2 |src.x| \rfloor \\
 dst.y &= \frac{|src.x|}{2^{\lfloor \log_2 |src.x| \rfloor}} \\
 dst.z &= \log_2 |src.x| \\
 dst.w &= 1
 \end{aligned}$$

MUL (Multiply)

$$\begin{aligned}
 dst.x &= src0.x \times src1.x \\
 dst.y &= src0.y \times src1.y \\
 dst.z &= src0.z \times src1.z \\
 dst.w &= src0.w \times src1.w
 \end{aligned}$$

ADD (Add)

$$\begin{aligned}
 dst.x &= src0.x + src1.x \\
 dst.y &= src0.y + src1.y \\
 dst.z &= src0.z + src1.z \\
 dst.w &= src0.w + src1.w
 \end{aligned}$$

DP3 (3-component Dot Product)

This instruction replicates its result.

$$dst = src0.x \times src1.x + src0.y \times src1.y + src0.z \times src1.z$$

DP4 (4-component Dot Product)

This instruction replicates its result.

$$dst = src0.x \times src1.x + src0.y \times src1.y + src0.z \times src1.z + src0.w \times src1.w$$

DST (Distance Vector)

$$\begin{aligned}dst.x &= 1 \\dst.y &= src0.y \times src1.y \\dst.z &= src0.z \\dst.w &= src1.w\end{aligned}$$
MIN (Minimum)
$$\begin{aligned}dst.x &= \min(src0.x, src1.x) \\dst.y &= \min(src0.y, src1.y) \\dst.z &= \min(src0.z, src1.z) \\dst.w &= \min(src0.w, src1.w)\end{aligned}$$
MAX (Maximum)
$$\begin{aligned}dst.x &= \max(src0.x, src1.x) \\dst.y &= \max(src0.y, src1.y) \\dst.z &= \max(src0.z, src1.z) \\dst.w &= \max(src0.w, src1.w)\end{aligned}$$
SLT (Set On Less Than)
$$\begin{aligned}dst.x &= (src0.x < src1.x)?1.0F : 0.0F \\dst.y &= (src0.y < src1.y)?1.0F : 0.0F \\dst.z &= (src0.z < src1.z)?1.0F : 0.0F \\dst.w &= (src0.w < src1.w)?1.0F : 0.0F\end{aligned}$$
SGE (Set On Greater Equal Than)
$$\begin{aligned}dst.x &= (src0.x \geq src1.x)?1.0F : 0.0F \\dst.y &= (src0.y \geq src1.y)?1.0F : 0.0F \\dst.z &= (src0.z \geq src1.z)?1.0F : 0.0F \\dst.w &= (src0.w \geq src1.w)?1.0F : 0.0F\end{aligned}$$
MAD (Multiply And Add)

Perform $a * b + c$. The implementation is free to decide whether there is an intermediate rounding step or not.

$$\begin{aligned}dst.x &= src0.x \times src1.x + src2.x \\dst.y &= src0.y \times src1.y + src2.y \\dst.z &= src0.z \times src1.z + src2.z \\dst.w &= src0.w \times src1.w + src2.w\end{aligned}$$

LRP (Linear Interpolate)

$$\begin{aligned}
 dst.x &= src0.x \times src1.x + (1 - src0.x) \times src2.x \\
 dst.y &= src0.y \times src1.y + (1 - src0.y) \times src2.y \\
 dst.z &= src0.z \times src1.z + (1 - src0.z) \times src2.z \\
 dst.w &= src0.w \times src1.w + (1 - src0.w) \times src2.w
 \end{aligned}$$

FMA (Fused Multiply-Add)

Perform $a * b + c$ with no intermediate rounding step.

$$\begin{aligned}
 dst.x &= src0.x \times src1.x + src2.x \\
 dst.y &= src0.y \times src1.y + src2.y \\
 dst.z &= src0.z \times src1.z + src2.z \\
 dst.w &= src0.w \times src1.w + src2.w
 \end{aligned}$$

FRC (Fraction)

$$\begin{aligned}
 dst.x &= src.x - \lfloor src.x \rfloor \\
 dst.y &= src.y - \lfloor src.y \rfloor \\
 dst.z &= src.z - \lfloor src.z \rfloor \\
 dst.w &= src.w - \lfloor src.w \rfloor
 \end{aligned}$$

FLR (Floor)

$$\begin{aligned}
 dst.x &= \lfloor src.x \rfloor \\
 dst.y &= \lfloor src.y \rfloor \\
 dst.z &= \lfloor src.z \rfloor \\
 dst.w &= \lfloor src.w \rfloor
 \end{aligned}$$

ROUND (Round)

$$\begin{aligned}
 dst.x &= round(src.x) \\
 dst.y &= round(src.y) \\
 dst.z &= round(src.z) \\
 dst.w &= round(src.w)
 \end{aligned}$$

EX2 (Exponential Base 2)

This instruction replicates its result.

$$dst = 2^{src.x}$$

LG2 (Logarithm Base 2)

This instruction replicates its result.

$$dst = \log_2 src.x$$

POW (Power)

This instruction replicates its result.

$$dst = src0.x^{src1.x}$$

LDEXP (Multiply Number by Integral Power of 2)

src1 is an integer.

$$dst.x = src0.x * 2^{src1.x} \quad dst.y = src0.y * 2^{src1.y} \quad dst.z = src0.z * 2^{src1.z} \quad dst.w = src0.w * 2^{src1.w}$$

COS (Cosine)

This instruction replicates its result.

$$dst = \cos src.x$$

DDX, DDX_FINE (Derivative Relative To X)

The fine variant is only used when `PIPE_CAP_TGSI_FS_FINE_DERIVATIVE` is advertised. When it is, the fine version guarantees one derivative per row while DDX is allowed to be the same for the entire 2x2 quad.

$$\begin{aligned} dst.x &= \text{partial}x(src.x) \\ dst.y &= \text{partial}x(src.y) \\ dst.z &= \text{partial}x(src.z) \\ dst.w &= \text{partial}x(src.w) \end{aligned}$$

DDY, DDY_FINE (Derivative Relative To Y)

The fine variant is only used when `PIPE_CAP_TGSI_FS_FINE_DERIVATIVE` is advertised. When it is, the fine version guarantees one derivative per column while DDY is allowed to be the same for the entire 2x2 quad.

$$\begin{aligned} dst.x &= \text{partial}y(src.x) \\ dst.y &= \text{partial}y(src.y) \\ dst.z &= \text{partial}y(src.z) \\ dst.w &= \text{partial}y(src.w) \end{aligned}$$

PK2H (Pack Two 16-bit Floats)

This instruction replicates its result.

$$dst = f32_to_f16(src.x) | f32_to_f16(src.y) \ll 16$$

PK2US (Pack Two Unsigned 16-bit Scalars)

This instruction replicates its result.

$$dst = f32_to_unorm16(src.x) | f32_to_unorm16(src.y) \ll 16$$

PK4B (Pack Four Signed 8-bit Scalars)

This instruction replicates its result.

$$dst = f32_to_snorm8(src.x) | (f32_to_snorm8(src.y) \ll 8) | (f32_to_snorm8(src.z) \ll 16) | (f32_to_snorm8(src.w) \ll 24)$$

PK4UB (Pack Four Unsigned 8-bit Scalars)

This instruction replicates its result.

$$dst = f32_to_unorm8(src.x) | (f32_to_unorm8(src.y) \ll 8) | (f32_to_unorm8(src.z) \ll 16) | (f32_to_unorm8(src.w) \ll 24)$$

SEQ (Set On Equal)

$$\begin{aligned} dst.x &= (src0.x == src1.x)?1.0F : 0.0F \\ dst.y &= (src0.y == src1.y)?1.0F : 0.0F \\ dst.z &= (src0.z == src1.z)?1.0F : 0.0F \\ dst.w &= (src0.w == src1.w)?1.0F : 0.0F \end{aligned}$$

SGT (Set On Greater Than)

$$\begin{aligned} dst.x &= (src0.x > src1.x)?1.0F : 0.0F \\ dst.y &= (src0.y > src1.y)?1.0F : 0.0F \\ dst.z &= (src0.z > src1.z)?1.0F : 0.0F \\ dst.w &= (src0.w > src1.w)?1.0F : 0.0F \end{aligned}$$

SIN (Sine)

This instruction replicates its result.

$$dst = \sin src.x$$

SLE (Set On Less Equal Than)

$$\begin{aligned} dst.x &= (src0.x \leq src1.x)?1.0F : 0.0F \\ dst.y &= (src0.y \leq src1.y)?1.0F : 0.0F \\ dst.z &= (src0.z \leq src1.z)?1.0F : 0.0F \\ dst.w &= (src0.w \leq src1.w)?1.0F : 0.0F \end{aligned}$$

SNE (Set On Not Equal)
$$\begin{aligned}dst.x &= (src0.x! = src1.x)?1.0F : 0.0F \\dst.y &= (src0.y! = src1.y)?1.0F : 0.0F \\dst.z &= (src0.z! = src1.z)?1.0F : 0.0F \\dst.w &= (src0.w! = src1.w)?1.0F : 0.0F\end{aligned}$$
TEX (Texture Lookup)

for array textures src0.y contains the slice for 1D, and src0.z contain the slice for 2D.

for shadow textures with no arrays (and not cube map), src0.z contains the reference value.

for shadow textures with arrays, src0.z contains the reference value for 1D arrays, and src0.w contains the reference value for 2D arrays and cube maps.

for cube map array shadow textures, the reference value cannot be passed in src0.w, and TEX2 must be used instead.

$$\begin{aligned}coord &= src0 \\shadow_{ref} &= src0.z \text{ or } src0.w \text{ (optional)} \\unit &= src1 \\dst &= texture_sample(unit, coord, shadow_{ref})\end{aligned}$$
TEX2 (Texture Lookup (for shadow cube map arrays only))

this is the same as TEX, but uses another reg to encode the reference value.

$$\begin{aligned}coord &= src0 \\shadow_{ref} &= src1.x \\unit &= src2 \\dst &= texture_sample(unit, coord, shadow_{ref})\end{aligned}$$
TXD (Texture Lookup with Derivatives)
$$\begin{aligned}coord &= src0 \\ddx &= src1 \\ddy &= src2 \\unit &= src3 \\dst &= texture_sample_deriv(unit, coord, ddx, ddy)\end{aligned}$$
TXP (Projective Texture Lookup)

$$\begin{aligned}coord.x &= src0.x/src0.w \\ coord.y &= src0.y/src0.w \\ coord.z &= src0.z/src0.w \\ coord.w &= src0.w \\ unit &= src1 \\ dst &= texture_sample(unit, coord)\end{aligned}$$
UP2H (Unpack Two 16-Bit Floats)
$$\begin{aligned}dst.x &= f16_to_f32(src0.x \& 0xffff) \\ dst.y &= f16_to_f32(src0.x \gg 16) \\ dst.z &= f16_to_f32(src0.x \& 0xffff) \\ dst.w &= f16_to_f32(src0.x \gg 16)\end{aligned}$$

Note: Considered for removal.

UP2US (Unpack Two Unsigned 16-Bit Scalars)
TBD

Note: Considered for removal.

UP4B (Unpack Four Signed 8-Bit Values)
TBD

Note: Considered for removal.

UP4UB (Unpack Four Unsigned 8-Bit Scalars)
TBD

Note: Considered for removal.

ARR (Address Register Load With Round)
$$\begin{aligned}dst.x &= (int)round(src.x) \\ dst.y &= (int)round(src.y) \\ dst.z &= (int)round(src.z) \\ dst.w &= (int)round(src.w)\end{aligned}$$
SSG (Set Sign)

$$\begin{aligned}dst.x &= (src.x > 0)?1 : (src.x < 0)? -1 : 0 \\dst.y &= (src.y > 0)?1 : (src.y < 0)? -1 : 0 \\dst.z &= (src.z > 0)?1 : (src.z < 0)? -1 : 0 \\dst.w &= (src.w > 0)?1 : (src.w < 0)? -1 : 0\end{aligned}$$
CMP (Compare)
$$\begin{aligned}dst.x &= (src0.x < 0)?src1.x : src2.x \\dst.y &= (src0.y < 0)?src1.y : src2.y \\dst.z &= (src0.z < 0)?src1.z : src2.z \\dst.w &= (src0.w < 0)?src1.w : src2.w\end{aligned}$$
KILL_IF (Conditional Discard)

Conditional discard. Allowed in fragment shaders only.

$$if(src.x < 0 || src.y < 0 || src.z < 0 || src.w < 0) discard endif$$
KILL (Discard)

Unconditional discard. Allowed in fragment shaders only.

TXB (Texture Lookup With Bias)

for cube map array textures and shadow cube maps, the bias value cannot be passed in src0.w, and TXB2 must be used instead.

if the target is a shadow texture, the reference value is always in src.z (this prevents shadow 3d and shadow 2d arrays from using this instruction, but this is not needed).

$$\begin{aligned}coord.x &= src0.x \\coord.y &= src0.y \\coord.z &= src0.z \\coord.w &= none \\bias &= src0.w \\unit &= src1 \\dst &= texture_sample(unit, coord, bias)\end{aligned}$$
TXB2 (Texture Lookup With Bias (some cube maps only))

this is the same as TXB, but uses another reg to encode the lod bias value for cube map arrays and shadow cube maps. Presumably shadow 2d arrays and shadow 3d targets could use this encoding too, but this is not legal.

shadow cube map arrays are neither possible nor required.

$$\begin{aligned}coord &= src0 \\bias &= src1.x \\unit &= src2 \\dst &= texture_sample(unit, coord, bias)\end{aligned}$$

DIV (Divide)

$$\begin{aligned}
 dst.x &= \frac{src0.x}{src1.x} \\
 dst.y &= \frac{src0.y}{src1.y} \\
 dst.z &= \frac{src0.z}{src1.z} \\
 dst.w &= \frac{src0.w}{src1.w}
 \end{aligned}$$

DP2 (2-component Dot Product)

This instruction replicates its result.

$$dst = src0.x \times src1.x + src0.y \times src1.y$$

TEX_LZ (Texture Lookup With LOD = 0)

This is the same as TXL with LOD = 0. Like every texture opcode, it obeys `pipe_sampler_view::u.tex.first_level` and `pipe_sampler_state::min_lod`. There is no way to override those two in shaders.

$$\begin{aligned}
 coord.x &= src0.x \\
 coord.y &= src0.y \\
 coord.z &= src0.z \\
 coord.w &= none \\
 lod &= 0 \\
 unit &= src1 \\
 dst &= texture_sample(unit, coord, lod)
 \end{aligned}$$

TXL (Texture Lookup With explicit LOD)

for cube map array textures, the explicit lod value cannot be passed in `src0.w`, and TXL2 must be used instead.

if the target is a shadow texture, the reference value is always in `src.z` (this prevents shadow 3d / 2d array / cube targets from using this instruction, but this is not needed).

$$\begin{aligned}
 coord.x &= src0.x \\
 coord.y &= src0.y \\
 coord.z &= src0.z \\
 coord.w &= none \\
 lod &= src0.w \\
 unit &= src1 \\
 dst &= texture_sample(unit, coord, lod)
 \end{aligned}$$

TXL2 (Texture Lookup With explicit LOD (for cube map arrays only))

this is the same as TXL, but uses another reg to encode the explicit lod value. Presumably shadow 3d / 2d array / cube targets could use this encoding too, but this is not legal.

shadow cube map arrays are neither possible nor required.

$$\begin{aligned}coord &= src0 \\ lod &= src1.x \\ unit &= src2 \\ dst &= texture_sample(unit, coord, lod)\end{aligned}$$

Compute ISA

These opcodes are primarily provided for special-use computational shaders. Support for these opcodes indicated by a special pipe capability bit (TBD).

XXX doesn't look like most of the opcodes really belong here.

CEIL (Ceiling)

$$\begin{aligned}dst.x &= \lceil src.x \rceil \\ dst.y &= \lceil src.y \rceil \\ dst.z &= \lceil src.z \rceil \\ dst.w &= \lceil src.w \rceil\end{aligned}$$

TRUNC (Truncate)

$$\begin{aligned}dst.x &= trunc(src.x) \\ dst.y &= trunc(src.y) \\ dst.z &= trunc(src.z) \\ dst.w &= trunc(src.w)\end{aligned}$$

MOD (Modulus)

$$\begin{aligned}dst.x &= src0.x \bmod src1.x \\ dst.y &= src0.y \bmod src1.y \\ dst.z &= src0.z \bmod src1.z \\ dst.w &= src0.w \bmod src1.w\end{aligned}$$

UARL (Integer Address Register Load)

Moves the contents of the source register, assumed to be an integer, into the destination register, which is assumed to be an address (ADDR) register.

TXF (Texel Fetch)

As per NV_gpu_shader4, extract a single texel from a specified texture image or PIPE_BUFFER resource. The source sampler may not be a CUBE or SHADOW. src 0 is a four-component signed integer vector used to identify the single texel accessed. 3 components + level. If the texture is multisampled, then the fourth component indicates the sample, not the mipmap level. Just like texture instructions, an optional offset vector is provided, which is subject to various driver restrictions (regarding range, source of offsets). This instruction ignores the sampler state.

TXF(uint_vec coord, int_vec offset).

TXQ (Texture Size Query)

As per NV_gpu_program4, retrieve the dimensions of the texture depending on the target. For 1D (width), 2D/RECT/CUBE (width, height), 3D (width, height, depth), 1D array (width, layers), 2D array (width, height, layers). Also return the number of accessible levels (last_level - first_level + 1) in W.

For components which don't return a resource dimension, their value is undefined.

$$\begin{aligned} lod &= src0.x \\ dst.x &= texture_width(unit, lod) \\ dst.y &= texture_height(unit, lod) \\ dst.z &= texture_depth(unit, lod) \\ dst.w &= texture_levels(unit) \end{aligned}$$

TXQS (Texture Samples Query)

This retrieves the number of samples in the texture, and stores it into the x component as an unsigned integer. The other components are undefined. If the texture is not multisampled, this function returns (1, undef, undef, undef).

$$dst.x = texture_samples(unit)$$

TG4 (Texture Gather)

As per ARB_texture_gather, gathers the four texels to be used in a bi-linear filtering operation and packs them into a single register. Only works with 2D, 2D array, cubemaps, and cubemaps arrays. For 2D textures, only the addressing modes of the sampler and the top level of any mip pyramid are used. Set W to zero. It behaves like the TEX instruction, but a filtered sample is not generated. The four samples that contribute to filtering are placed into xyzw in clockwise order, starting with the (u,v) texture coordinate delta at the following locations (-, +), (+, +), (+, -), (-, -), where the magnitude of the deltas are half a texel.

PIPE_CAP_TEXTURE_SM5 enhances this instruction to support shadow per-sample depth compares, single component selection, and a non-constant offset. It doesn't allow support for the GL independent offset to get i0,j0. This would require another CAP is hw can do it natively. For now we lower that before TGSI.

$$\begin{aligned} coord &= src0 \\ component &= src1 \\ dst &= texture_gather4(unit, coord, component) \end{aligned}$$

(with SM5 - cube array shadow)

$$\begin{aligned} coord &= src0 \\ compare &= src1 \\ dst &= texture_gather(uint, coord, compare) \end{aligned}$$

LODQ (level of detail query)

Compute the LOD information that the texture pipe would use to access the texture. The Y component contains the computed LOD lambda_prime. The X component contains the LOD that will be accessed, based on min/max lod's and mipmap filters.

$$\begin{aligned} coord &= src0 \\ dst.xy &= lodq(uint, coord); \end{aligned}$$
CLOCK (retrieve the current shader time)

Invoking this instruction multiple times in the same shader should cause monotonically increasing values to be returned. The values are implicitly 64-bit, so if fewer than 64 bits of precision are available, to provide expected wraparound semantics, the value should be shifted up so that the most significant bit of the time is the most significant bit of the 64-bit value.

$$dst.xy = clock()$$

Integer ISA

These opcodes are used for integer operations. Support for these opcodes indicated by PIPE_SHADER_CAP_INTEGERS (all of them?)

I2F (Signed Integer To Float)

Rounding is unspecified (round to nearest even suggested).

$$\begin{aligned} dst.x &= (float)src.x \\ dst.y &= (float)src.y \\ dst.z &= (float)src.z \\ dst.w &= (float)src.w \end{aligned}$$
U2F (Unsigned Integer To Float)

Rounding is unspecified (round to nearest even suggested).

$$\begin{aligned} dst.x &= (float)src.x \\ dst.y &= (float)src.y \\ dst.z &= (float)src.z \\ dst.w &= (float)src.w \end{aligned}$$
F2I (Float to Signed Integer)

Rounding is towards zero (truncate). Values outside signed range (including NaNs) produce undefined results.

$$\begin{aligned} dst.x &= (int)src.x \\ dst.y &= (int)src.y \\ dst.z &= (int)src.z \\ dst.w &= (int)src.w \end{aligned}$$
F2U (Float to Unsigned Integer)

Rounding is towards zero (truncate). Values outside unsigned range (including NaNs) produce undefined results.

$$\begin{aligned}dst.x &= (\text{unsigned})src.x \\dst.y &= (\text{unsigned})src.y \\dst.z &= (\text{unsigned})src.z \\dst.w &= (\text{unsigned})src.w\end{aligned}$$

UADD (Integer Add)

This instruction works the same for signed and unsigned integers. The low 32bit of the result is returned.

$$\begin{aligned}dst.x &= src0.x + src1.x \\dst.y &= src0.y + src1.y \\dst.z &= src0.z + src1.z \\dst.w &= src0.w + src1.w\end{aligned}$$

UMAD (Integer Multiply And Add)

This instruction works the same for signed and unsigned integers. The multiplication returns the low 32bit (as does the result itself).

$$\begin{aligned}dst.x &= src0.x \times src1.x + src2.x \\dst.y &= src0.y \times src1.y + src2.y \\dst.z &= src0.z \times src1.z + src2.z \\dst.w &= src0.w \times src1.w + src2.w\end{aligned}$$

UMUL (Integer Multiply)

This instruction works the same for signed and unsigned integers. The low 32bit of the result is returned.

$$\begin{aligned}dst.x &= src0.x \times src1.x \\dst.y &= src0.y \times src1.y \\dst.z &= src0.z \times src1.z \\dst.w &= src0.w \times src1.w\end{aligned}$$

IMUL_HI (Signed Integer Multiply High Bits)

The high 32bits of the multiplication of 2 signed integers are returned.

$$\begin{aligned}dst.x &= (src0.x \times src1.x) \gg 32 \\dst.y &= (src0.y \times src1.y) \gg 32 \\dst.z &= (src0.z \times src1.z) \gg 32 \\dst.w &= (src0.w \times src1.w) \gg 32\end{aligned}$$

UMUL_HI (Unsigned Integer Multiply High Bits)

The high 32bits of the multiplication of 2 unsigned integers are returned.

$$\begin{aligned}dst.x &= (src0.x \times src1.x) \gg 32 \\dst.y &= (src0.y \times src1.y) \gg 32 \\dst.z &= (src0.z \times src1.z) \gg 32 \\dst.w &= (src0.w \times src1.w) \gg 32\end{aligned}$$

IDIV (Signed Integer Division)

TBD: behavior for division by zero.

$$\begin{aligned}dst.x &= \frac{src0.x}{src1.x} \\dst.y &= \frac{src0.y}{src1.y} \\dst.z &= \frac{src0.z}{src1.z} \\dst.w &= \frac{src0.w}{src1.w}\end{aligned}$$

UDIV (Unsigned Integer Division)

For division by zero, 0xffffffff is returned.

$$\begin{aligned}dst.x &= \frac{src0.x}{src1.x} \\dst.y &= \frac{src0.y}{src1.y} \\dst.z &= \frac{src0.z}{src1.z} \\dst.w &= \frac{src0.w}{src1.w}\end{aligned}$$

UMOD (Unsigned Integer Remainder)

If second arg is zero, 0xffffffff is returned.

$$\begin{aligned}dst.x &= src0.x \bmod src1.x \\dst.y &= src0.y \bmod src1.y \\dst.z &= src0.z \bmod src1.z \\dst.w &= src0.w \bmod src1.w\end{aligned}$$

NOT (Bitwise Not)

$$\begin{aligned}dst.x &= \sim src.x \\dst.y &= \sim src.y \\dst.z &= \sim src.z \\dst.w &= \sim src.w\end{aligned}$$

AND (Bitwise And)
$$\begin{aligned}dst.x &= src0.x \& src1.x \\dst.y &= src0.y \& src1.y \\dst.z &= src0.z \& src1.z \\dst.w &= src0.w \& src1.w\end{aligned}$$
OR (Bitwise Or)
$$\begin{aligned}dst.x &= src0.x | src1.x \\dst.y &= src0.y | src1.y \\dst.z &= src0.z | src1.z \\dst.w &= src0.w | src1.w\end{aligned}$$
XOR (Bitwise Xor)
$$\begin{aligned}dst.x &= src0.x \oplus src1.x \\dst.y &= src0.y \oplus src1.y \\dst.z &= src0.z \oplus src1.z \\dst.w &= src0.w \oplus src1.w\end{aligned}$$
IMAX (Maximum of Signed Integers)
$$\begin{aligned}dst.x &= \max(src0.x, src1.x) \\dst.y &= \max(src0.y, src1.y) \\dst.z &= \max(src0.z, src1.z) \\dst.w &= \max(src0.w, src1.w)\end{aligned}$$
UMAX (Maximum of Unsigned Integers)
$$\begin{aligned}dst.x &= \max(src0.x, src1.x) \\dst.y &= \max(src0.y, src1.y) \\dst.z &= \max(src0.z, src1.z) \\dst.w &= \max(src0.w, src1.w)\end{aligned}$$
IMIN (Minimum of Signed Integers)
$$\begin{aligned}dst.x &= \min(src0.x, src1.x) \\dst.y &= \min(src0.y, src1.y) \\dst.z &= \min(src0.z, src1.z) \\dst.w &= \min(src0.w, src1.w)\end{aligned}$$

UMIN (Minimum of Unsigned Integers)
$$\begin{aligned}dst.x &= \min(src0.x, src1.x) \\dst.y &= \min(src0.y, src1.y) \\dst.z &= \min(src0.z, src1.z) \\dst.w &= \min(src0.w, src1.w)\end{aligned}$$
SHL (Shift Left)

The shift count is masked with 0x1f before the shift is applied.

$$\begin{aligned}dst.x &= src0.x \ll (0x1f \& src1.x) \\dst.y &= src0.y \ll (0x1f \& src1.y) \\dst.z &= src0.z \ll (0x1f \& src1.z) \\dst.w &= src0.w \ll (0x1f \& src1.w)\end{aligned}$$
ISHR (Arithmetic Shift Right (of Signed Integer))

The shift count is masked with 0x1f before the shift is applied.

$$\begin{aligned}dst.x &= src0.x \gg (0x1f \& src1.x) \\dst.y &= src0.y \gg (0x1f \& src1.y) \\dst.z &= src0.z \gg (0x1f \& src1.z) \\dst.w &= src0.w \gg (0x1f \& src1.w)\end{aligned}$$
USHR (Logical Shift Right)

The shift count is masked with 0x1f before the shift is applied.

$$\begin{aligned}dst.x &= src0.x \gg (unsigned)(0x1f \& src1.x) \\dst.y &= src0.y \gg (unsigned)(0x1f \& src1.y) \\dst.z &= src0.z \gg (unsigned)(0x1f \& src1.z) \\dst.w &= src0.w \gg (unsigned)(0x1f \& src1.w)\end{aligned}$$
UCMP (Integer Conditional Move)
$$\begin{aligned}dst.x &= src0.x ? src1.x : src2.x \\dst.y &= src0.y ? src1.y : src2.y \\dst.z &= src0.z ? src1.z : src2.z \\dst.w &= src0.w ? src1.w : src2.w\end{aligned}$$
ISSG (Integer Set Sign)

$$\begin{aligned}
 dst.x &= (src0.x < 0)? -1 : (src0.x > 0)?1 : 0 \\
 dst.y &= (src0.y < 0)? -1 : (src0.y > 0)?1 : 0 \\
 dst.z &= (src0.z < 0)? -1 : (src0.z > 0)?1 : 0 \\
 dst.w &= (src0.w < 0)? -1 : (src0.w > 0)?1 : 0
 \end{aligned}$$
FSLT (Float Set On Less Than (ordered))

Same comparison as SLT but returns integer instead of 1.0/0.0 float

$$\begin{aligned}
 dst.x &= (src0.x < src1.x)? \sim 0 : 0 \\
 dst.y &= (src0.y < src1.y)? \sim 0 : 0 \\
 dst.z &= (src0.z < src1.z)? \sim 0 : 0 \\
 dst.w &= (src0.w < src1.w)? \sim 0 : 0
 \end{aligned}$$
ISLT (Signed Integer Set On Less Than)

$$\begin{aligned}
 dst.x &= (src0.x < src1.x)? \sim 0 : 0 \\
 dst.y &= (src0.y < src1.y)? \sim 0 : 0 \\
 dst.z &= (src0.z < src1.z)? \sim 0 : 0 \\
 dst.w &= (src0.w < src1.w)? \sim 0 : 0
 \end{aligned}$$
USLT (Unsigned Integer Set On Less Than)

$$\begin{aligned}
 dst.x &= (src0.x < src1.x)? \sim 0 : 0 \\
 dst.y &= (src0.y < src1.y)? \sim 0 : 0 \\
 dst.z &= (src0.z < src1.z)? \sim 0 : 0 \\
 dst.w &= (src0.w < src1.w)? \sim 0 : 0
 \end{aligned}$$
FSGE (Float Set On Greater Equal Than (ordered))

Same comparison as SGE but returns integer instead of 1.0/0.0 float

$$\begin{aligned}
 dst.x &= (src0.x \geq src1.x)? \sim 0 : 0 \\
 dst.y &= (src0.y \geq src1.y)? \sim 0 : 0 \\
 dst.z &= (src0.z \geq src1.z)? \sim 0 : 0 \\
 dst.w &= (src0.w \geq src1.w)? \sim 0 : 0
 \end{aligned}$$
ISGE (Signed Integer Set On Greater Equal Than)

$$\begin{aligned}
 dst.x &= (src0.x \geq src1.x)? \sim 0 : 0 \\
 dst.y &= (src0.y \geq src1.y)? \sim 0 : 0 \\
 dst.z &= (src0.z \geq src1.z)? \sim 0 : 0 \\
 dst.w &= (src0.w \geq src1.w)? \sim 0 : 0
 \end{aligned}$$

USGE (Unsigned Integer Set On Greater Equal Than)
$$\begin{aligned}dst.x &= (src0.x \geq src1.x)? \sim 0 : 0 \\dst.y &= (src0.y \geq src1.y)? \sim 0 : 0 \\dst.z &= (src0.z \geq src1.z)? \sim 0 : 0 \\dst.w &= (src0.w \geq src1.w)? \sim 0 : 0\end{aligned}$$
FSEQ (Float Set On Equal (ordered))

Same comparison as SEQ but returns integer instead of 1.0/0.0 float

$$\begin{aligned}dst.x &= (src0.x == src1.x)? \sim 0 : 0 \\dst.y &= (src0.y == src1.y)? \sim 0 : 0 \\dst.z &= (src0.z == src1.z)? \sim 0 : 0 \\dst.w &= (src0.w == src1.w)? \sim 0 : 0\end{aligned}$$
USEQ (Integer Set On Equal)
$$\begin{aligned}dst.x &= (src0.x == src1.x)? \sim 0 : 0 \\dst.y &= (src0.y == src1.y)? \sim 0 : 0 \\dst.z &= (src0.z == src1.z)? \sim 0 : 0 \\dst.w &= (src0.w == src1.w)? \sim 0 : 0\end{aligned}$$
FSNE (Float Set On Not Equal (unordered))

Same comparison as SNE but returns integer instead of 1.0/0.0 float

$$\begin{aligned}dst.x &= (src0.x! = src1.x)? \sim 0 : 0 \\dst.y &= (src0.y! = src1.y)? \sim 0 : 0 \\dst.z &= (src0.z! = src1.z)? \sim 0 : 0 \\dst.w &= (src0.w! = src1.w)? \sim 0 : 0\end{aligned}$$
USNE (Integer Set On Not Equal)
$$\begin{aligned}dst.x &= (src0.x! = src1.x)? \sim 0 : 0 \\dst.y &= (src0.y! = src1.y)? \sim 0 : 0 \\dst.z &= (src0.z! = src1.z)? \sim 0 : 0 \\dst.w &= (src0.w! = src1.w)? \sim 0 : 0\end{aligned}$$
INEG (Integer Negate)

Two's complement.

```

dst.x = -src.x
dst.y = -src.y
dst.z = -src.z
dst.w = -src.w

```

IABS (Integer Absolute Value)

```

dst.x = |src.x|
dst.y = |src.y|
dst.z = |src.z|
dst.w = |src.w|

```

Bitwise ISA

These opcodes are used for bit-level manipulation of integers.

IBFE (Signed Bitfield Extract)

Like GLSL bitfieldExtract. Extracts a set of bits from the input, and sign-extends them if the high bit of the extracted window is set.

Pseudocode:

```

def ibfe(value, offset, bits):
    if offset < 0 or bits < 0 or offset + bits > 32:
        return undefined
    if bits == 0: return 0
    # Note: >> sign-extends
    return (value << (32 - offset - bits)) >> (32 - bits)

```

UBFE (Unsigned Bitfield Extract)

Like GLSL bitfieldExtract. Extracts a set of bits from the input, without any sign-extension.

Pseudocode:

```

def ubfe(value, offset, bits):
    if offset < 0 or bits < 0 or offset + bits > 32:
        return undefined
    if bits == 0: return 0
    # Note: >> does not sign-extend
    return (value << (32 - offset - bits)) >> (32 - bits)

```

BFI (Bitfield Insert)

Like GLSL bitfieldInsert. Replaces a bit region of 'base' with the low bits of 'insert'.

Pseudocode:

```

def bfi(base, insert, offset, bits):
    if offset < 0 or bits < 0 or offset + bits > 32:
        return undefined
    # << defined such that mask == ~0 when bits == 32, offset == 0
    mask = ((1 << bits) - 1) << offset
    return ((insert << offset) & mask) | (base & ~mask)

```

BREV (Bitfield Reverse)

See SM5 instruction BFREV. Reverses the bits of the argument.

POPC (Population Count)

See SM5 instruction COUNTBITS. Counts the number of set bits in the argument.

LSB (Index of lowest set bit)

See SM5 instruction FIRSTBIT_LO. Computes the 0-based index of the first set bit of the argument. Returns -1 if none are set.

IMSB (Index of highest non-sign bit)

See SM5 instruction FIRSTBIT_SHI. Computes the 0-based index of the highest non-sign bit of the argument (i.e. highest 0 bit for negative numbers, highest 1 bit for positive numbers). Returns -1 if all bits are the same (i.e. for inputs 0 and -1).

UMSB (Index of highest set bit)

See SM5 instruction FIRSTBIT_HI. Computes the 0-based index of the highest set bit of the argument. Returns -1 if none are set.

Geometry ISA

These opcodes are only supported in geometry shaders; they have no meaning in any other type of shader.

EMIT (Emit)

Generate a new vertex for the current primitive into the specified vertex stream using the values in the output registers.

ENDPRIM (End Primitive)

Complete the current primitive in the specified vertex stream (consisting of the emitted vertices), and start a new one.

GLSL ISA

These opcodes are part of *GLSL*'s opcode set. Support for these opcodes is determined by a special capability bit, GLSL. Some require glsl version 1.30 (UIF/SWITCH/CASE/DEFAULT/ENDSWITCH).

CAL (Subroutine Call)

push(pc) pc = target

RET (Subroutine Call Return)

pc = pop()

CONT (Continue)

Unconditionally moves the point of execution to the instruction after the last bgnloop. The instruction must appear within a bgnloop/endloop.

Note: Support for CONT is determined by a special capability bit, TGSI_CONT_SUPPORTED. See *Screen* for more information.

BGNLOOP (Begin a Loop)

Start a loop. Must have a matching endloop.

BGNSUB (Begin Subroutine)

Starts definition of a subroutine. Must have a matching endsub.

ENDLOOP (End a Loop)

End a loop started with bgnloop.

ENDSUB (End Subroutine)

Ends definition of a subroutine.

NOP (No Operation)

Do nothing.

BRK (Break)

Unconditionally moves the point of execution to the instruction after the next endloop or endswitch. The instruction must appear within a loop/endloop or switch/endswitch.

IF (Float If)

Start an IF ... ELSE .. ENDIF block. Condition evaluates to true if

```
src0.x != 0.0
```

where src0.x is interpreted as a floating point register.

UIF (Bitwise If)

Start an UIF ... ELSE .. ENDIF block. Condition evaluates to true if

```
src0.x != 0
```

where src0.x is interpreted as an integer register.

ELSE (Else)

Starts an else block, after an IF or UIF statement.

ENDIF (End If)

Ends an IF or UIF block.

SWITCH (Switch)

Starts a C-style switch expression. The switch consists of one or multiple CASE statements, and at most one DEFAULT statement. Execution of a statement ends when a BRK is hit, but just like in C falling through to other cases without a break is allowed. Similarly, DEFAULT label is allowed anywhere not just as last statement, and fallthrough is allowed into/from it. CASE src arguments are evaluated at bit level against the SWITCH src argument.

Example:

```
SWITCH src[0].x
CASE src[0].x
(some instructions here)
(optional BRK here)
DEFAULT
(some instructions here)
(optional BRK here)
CASE src[0].x
(some instructions here)
(optional BRK here)
ENDSWITCH
```

CASE (Switch case)

This represents a switch case label. The src arg must be an integer immediate.

DEFAULT (Switch default)

This represents the default case in the switch, which is taken if no other case matches.

ENDSWITCH (End of switch)

Ends a switch expression.

Interpolation ISA

The interpolation instructions allow an input to be interpolated in a different way than its declaration. This corresponds to the GLSL 4.00 `interpolateAt*` functions. The first argument of each of these must come from `TGSI_FILE_INPUT`.

INTERP_CENTROID (Interpolate at the centroid)

Interpolates the varying specified by `src0` at the centroid

INTERP_SAMPLE (Interpolate at the specified sample)

Interpolates the varying specified by `src0` at the sample id specified by `src1.x` (interpreted as an integer)

INTERP_OFFSET (Interpolate at the specified offset)

Interpolates the varying specified by `src0` at the offset `src1.xy` from the pixel center (interpreted as floats)

Double ISA

The double-precision opcodes reinterpret four-component vectors into two-component vectors with doubled precision in each component.

DABS (Absolute)

$$\begin{aligned}dst.xy &= |src0.xy| \\dst.zw &= |src0.zw|\end{aligned}$$

DADD (Add)

$$\begin{aligned}dst.xy &= src0.xy + src1.xy \\dst.zw &= src0.zw + src1.zw\end{aligned}$$

DSEQ (Set on Equal)

$$\begin{aligned}dst.x &= src0.xy == src1.xy? \sim 0 : 0 \\dst.z &= src0.zw == src1.zw? \sim 0 : 0\end{aligned}$$

DSNE (Set on Not Equal)

$$\begin{aligned}dst.x &= src0.xy != src1.xy? \sim 0 : 0 \\dst.z &= src0.zw != src1.zw? \sim 0 : 0\end{aligned}$$

DSLTT (Set on Less than)

$$\begin{aligned}dst.x &= src0.xy < src1.xy? \sim 0 : 0 \\dst.z &= src0.zw < src1.zw? \sim 0 : 0\end{aligned}$$

DSGE (Set on Greater equal)

$$dst.x = src0.xy \geq src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw \geq src1.zw? \sim 0 : 0$$

DFRAC (Fraction)

$$dst.xy = src.xy - \lfloor src.xy \rfloor$$

$$dst.zw = src.zw - \lfloor src.zw \rfloor$$

DTRUNC (Truncate)

$$dst.xy = trunc(src.xy)$$

$$dst.zw = trunc(src.zw)$$

DCEIL (Ceiling)

$$dst.xy = \lceil src.xy \rceil$$

$$dst.zw = \lceil src.zw \rceil$$

DFLR (Floor)

$$dst.xy = \lfloor src.xy \rfloor$$

$$dst.zw = \lfloor src.zw \rfloor$$

DROUND (Fraction)

$$dst.xy = round(src.xy)$$

$$dst.zw = round(src.zw)$$

DSSG (Set Sign)

$$dst.xy = (src.xy > 0)?1.0 : (src.xy < 0)? -1.0 : 0.0$$

$$dst.zw = (src.zw > 0)?1.0 : (src.zw < 0)? -1.0 : 0.0$$

DFRACEXP (Convert Number to Fractional and Integral Components)

Like the `frexp()` routine in many math libraries, this opcode stores the exponent of its source to `dst0`, and the significand to `dst1`, such that $dst1 \times 2^{dst0} = src$. The results are replicated across channels.

$$dst0.xy = dst.zw = frac(src.xy)$$

$$dst1 = frac(src.xy)$$

DLDEXP (Multiply Number by Integral Power of 2)

This opcode is the inverse of *DFRACEXP*. The second source is an integer.

$$\begin{aligned}dst.xy &= src0.xy \times 2^{src1.x} \\dst.zw &= src0.zw \times 2^{src1.z}\end{aligned}$$

DMIN (Minimum)

$$\begin{aligned}dst.xy &= \min(src0.xy, src1.xy) \\dst.zw &= \min(src0.zw, src1.zw)\end{aligned}$$

DMAX (Maximum)

$$\begin{aligned}dst.xy &= \max(src0.xy, src1.xy) \\dst.zw &= \max(src0.zw, src1.zw)\end{aligned}$$

DMUL (Multiply)

$$\begin{aligned}dst.xy &= src0.xy \times src1.xy \\dst.zw &= src0.zw \times src1.zw\end{aligned}$$

DMAD (Multiply And Add)

$$\begin{aligned}dst.xy &= src0.xy \times src1.xy + src2.xy \\dst.zw &= src0.zw \times src1.zw + src2.zw\end{aligned}$$

DFMA (Fused Multiply-Add)

Perform $a * b + c$ with no intermediate rounding step.

$$\begin{aligned}dst.xy &= src0.xy \times src1.xy + src2.xy \\dst.zw &= src0.zw \times src1.zw + src2.zw\end{aligned}$$

DDIV (Divide)

$$\begin{aligned}dst.xy &= \frac{src0.xy}{src1.xy} \\dst.zw &= \frac{src0.zw}{src1.zw}\end{aligned}$$

DRCP (Reciprocal)

$$dst.xy = \frac{1}{src.xy}$$

$$dst.zw = \frac{1}{src.zw}$$

DSQRT (Square Root)

$$dst.xy = \sqrt{src.xy}$$

$$dst.zw = \sqrt{src.zw}$$

DRSQ (Reciprocal Square Root)

$$dst.xy = \frac{1}{\sqrt{src.xy}}$$

$$dst.zw = \frac{1}{\sqrt{src.zw}}$$

F2D (Float to Double)

$$dst.xy = double(src0.x)$$

$$dst.zw = double(src0.y)$$

D2F (Double to Float)

$$dst.x = float(src0.xy)$$

$$dst.y = float(src0.zw)$$

I2D (Int to Double)

$$dst.xy = double(src0.x)$$

$$dst.zw = double(src0.y)$$

D2I (Double to Int)

$$dst.x = int(src0.xy)$$

$$dst.y = int(src0.zw)$$

U2D (Unsigned Int to Double)

$$\begin{aligned}dst.xy &= double(src0.x) \\dst.zw &= double(src0.y)\end{aligned}$$

D2U (Double to Unsigned Int)

$$\begin{aligned}dst.x &= unsigned(src0.xy) \\dst.y &= unsigned(src0.zw)\end{aligned}$$

64-bit Integer ISA

The 64-bit integer opcodes reinterpret four-component vectors into two-component vectors with 64-bits in each component.

I64ABS (64-bit Integer Absolute Value)

$$\begin{aligned}dst.xy &= |src0.xy| \\dst.zw &= |src0.zw|\end{aligned}$$

I64NEG (64-bit Integer Negate)

Two's complement.

$$\begin{aligned}dst.xy &= -src.xy \\dst.zw &= -src.zw\end{aligned}$$

I64SSG (64-bit Integer Set Sign)

$$\begin{aligned}dst.xy &= (src0.xy < 0)? -1 : (src0.xy > 0)? 1 : 0 \\dst.zw &= (src0.zw < 0)? -1 : (src0.zw > 0)? 1 : 0\end{aligned}$$

U64ADD (64-bit Integer Add)

$$\begin{aligned}dst.xy &= src0.xy + src1.xy \\dst.zw &= src0.zw + src1.zw\end{aligned}$$

U64MUL (64-bit Integer Multiply)

$$\begin{aligned}dst.xy &= src0.xy * src1.xy \\dst.zw &= src0.zw * src1.zw\end{aligned}$$

U64SEQ (64-bit Integer Set on Equal)

$$dst.x = src0.xy == src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw == src1.zw? \sim 0 : 0$$
U64SNE (64-bit Integer Set on Not Equal)

$$dst.x = src0.xy! = src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw! = src1.zw? \sim 0 : 0$$
U64SLT (64-bit Unsigned Integer Set on Less Than)

$$dst.x = src0.xy < src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw < src1.zw? \sim 0 : 0$$
U64SGE (64-bit Unsigned Integer Set on Greater Equal)

$$dst.x = src0.xy \geq src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw \geq src1.zw? \sim 0 : 0$$
I64SLT (64-bit Signed Integer Set on Less Than)

$$dst.x = src0.xy < src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw < src1.zw? \sim 0 : 0$$
I64SGE (64-bit Signed Integer Set on Greater Equal)

$$dst.x = src0.xy \geq src1.xy? \sim 0 : 0$$

$$dst.z = src0.zw \geq src1.zw? \sim 0 : 0$$
I64MIN (Minimum of 64-bit Signed Integers)

$$dst.xy = \min(src0.xy, src1.xy)$$

$$dst.zw = \min(src0.zw, src1.zw)$$
U64MIN (Minimum of 64-bit Unsigned Integers)

$$dst.xy = \min(src0.xy, src1.xy)$$

$$dst.zw = \min(src0.zw, src1.zw)$$

I64MAX (Maximum of 64-bit Signed Integers)

$$\begin{aligned}dst.xy &= \max(src0.xy, src1.xy) \\dst.zw &= \max(src0.zw, src1.zw)\end{aligned}$$

U64MAX (Maximum of 64-bit Unsigned Integers)

$$\begin{aligned}dst.xy &= \max(src0.xy, src1.xy) \\dst.zw &= \max(src0.zw, src1.zw)\end{aligned}$$

U64SHL (Shift Left 64-bit Unsigned Integer)

The shift count is masked with 0x3f before the shift is applied.

$$\begin{aligned}dst.xy &= src0.xy \ll (0x3f \& src1.x) \\dst.zw &= src0.zw \ll (0x3f \& src1.y)\end{aligned}$$

I64SHR (Arithmetic Shift Right (of 64-bit Signed Integer))

The shift count is masked with 0x3f before the shift is applied.

$$\begin{aligned}dst.xy &= src0.xy \gg (0x3f \& src1.x) \\dst.zw &= src0.zw \gg (0x3f \& src1.y)\end{aligned}$$

U64SHR (Logical Shift Right (of 64-bit Unsigned Integer))

The shift count is masked with 0x3f before the shift is applied.

$$\begin{aligned}dst.xy &= src0.xy \gg (\text{unsigned})(0x3f \& src1.x) \\dst.zw &= src0.zw \gg (\text{unsigned})(0x3f \& src1.y)\end{aligned}$$

I64DIV (64-bit Signed Integer Division)

$$\begin{aligned}dst.xy &= \frac{src0.xy}{src1.xy} \\dst.zw &= \frac{src0.zw}{src1.zw}\end{aligned}$$

U64DIV (64-bit Unsigned Integer Division)

$$\begin{aligned}dst.xy &= \frac{src0.xy}{src1.xy} \\dst.zw &= \frac{src0.zw}{src1.zw}\end{aligned}$$

U64MOD (64-bit Unsigned Integer Remainder)

$$dst.xy = src0.xy \bmod src1.xy$$

$$dst.zw = src0.zw \bmod src1.zw$$

I64MOD (64-bit Signed Integer Remainder)

$$dst.xy = src0.xy \bmod src1.xy$$

$$dst.zw = src0.zw \bmod src1.zw$$

F2U64 (Float to 64-bit Unsigned Int)

$$dst.xy = (uint64_t)src0.x$$

$$dst.zw = (uint64_t)src0.y$$

F2I64 (Float to 64-bit Int)

$$dst.xy = (int64_t)src0.x$$

$$dst.zw = (int64_t)src0.y$$

U2I64 (Unsigned Integer to 64-bit Integer)

This is a zero extension.

$$dst.xy = (int64_t)src0.x$$

$$dst.zw = (int64_t)src0.y$$

I2I64 (Signed Integer to 64-bit Integer)

This is a sign extension.

$$dst.xy = (int64_t)src0.x$$

$$dst.zw = (int64_t)src0.y$$

D2U64 (Double to 64-bit Unsigned Int)

$$dst.xy = (uint64_t)src0.xy$$

$$dst.zw = (uint64_t)src0.zw$$

D2I64 (Double to 64-bit Int)

$$dst.xy = (int64_t)src0.xy$$
$$dst.zw = (int64_t)src0.zw$$
U642F (64-bit unsigned integer to float)
$$dst.x = (float)src0.xy$$
$$dst.y = (float)src0.zw$$
I642F (64-bit Int to Float)
$$dst.x = (float)src0.xy$$
$$dst.y = (float)src0.zw$$
U642D (64-bit unsigned integer to double)
$$dst.xy = (double)src0.xy$$
$$dst.zw = (double)src0.zw$$
I642D (64-bit Int to double)
$$dst.xy = (double)src0.xy$$
$$dst.zw = (double)src0.zw$$

Resource Sampling Opcodes

Those opcodes follow very closely semantics of the respective Direct3D instructions. If in doubt double check Direct3D documentation. Note that the swizzle on SVIEW (src1) determines texel swizzling after lookup.

SAMPLE

Using provided address, sample data from the specified texture using the filtering mode identified by the given sampler. The source data may come from any resource type other than buffers.

Syntax: `SAMPLE dst, address, sampler_view, sampler`

Example: `SAMPLE TEMP[0], TEMP[1], SVIEW[0], SAMP[0]`

SAMPLE_I

Simplified alternative to the `SAMPLE` instruction. Using the provided integer address, `SAMPLE_I` fetches data from the specified sampler view without any filtering. The source data may come from any resource type other than `CUBE`.

Syntax: `SAMPLE_I dst, address, sampler_view`

Example: `SAMPLE_I TEMP[0], TEMP[1], SVIEW[0]`

The 'address' is specified as unsigned integers. If the 'address' is out of range `[0...(# texels - 1)]` the result of the fetch is always 0 in all components. As such the instruction doesn't honor address wrap modes, in cases

where that behavior is desirable ‘SAMPLE’ instruction should be used. address.w always provides an unsigned integer mipmap level. If the value is out of the range then the instruction always returns 0 in all components. address.yz are ignored for buffers and 1d textures. address.z is ignored for 1d texture arrays and 2d textures.

For 1D texture arrays address.y provides the array index (also as unsigned integer). If the value is out of the range of available array indices [0.. (array size - 1)] then the opcode always returns 0 in all components. For 2D texture arrays address.z provides the array index, otherwise it exhibits the same behavior as in the case for 1D texture arrays. The exact semantics of the source address are presented in the table below:

resource type	X	Y	Z	W
PIPE_BUFFER	x			ignored
PIPE_TEXTURE_1D	x			mpl
PIPE_TEXTURE_2D	x	y		mpl
PIPE_TEXTURE_3D	x	y	z	mpl
PIPE_TEXTURE_RECT	x	y		mpl
PIPE_TEXTURE_CUBE	not allowed as source			
PIPE_TEXTURE_1D_ARRAY	x	idx		mpl
PIPE_TEXTURE_2D_ARRAY	x	y	idx	mpl

Where ‘mpl’ is a mipmap level and ‘idx’ is the array index.

SAMPLE_I_MS

Just like SAMPLE_I but allows fetch data from multi-sampled surfaces.

Syntax: `SAMPLE_I_MS dst, address, sampler_view, sample`

SAMPLE_B

Just like the SAMPLE instruction with the exception that an additional bias is applied to the level of detail computed as part of the instruction execution.

Syntax: `SAMPLE_B dst, address, sampler_view, sampler, lod_bias`

Example: `SAMPLE_B TEMP[0], TEMP[1], SVIEW[0], SAMP[0], TEMP[2].x`

SAMPLE_C

Similar to the SAMPLE instruction but it performs a comparison filter. The operands to SAMPLE_C are identical to SAMPLE, except that there is an additional float32 operand, reference value, which must be a register with single-component, or a scalar literal. SAMPLE_C makes the hardware use the current samplers compare_func (in pipe_sampler_state) to compare reference value against the red component value for the source resource at each texel that the currently configured texture filter covers based on the provided coordinates.

Syntax: `SAMPLE_C dst, address, sampler_view.r, sampler, ref_value`

Example: `SAMPLE_C TEMP[0], TEMP[1], SVIEW[0].r, SAMP[0], TEMP[2].x`

SAMPLE_C_LZ

Same as SAMPLE_C, but LOD is 0 and derivatives are ignored. The LZ stands for level-zero.

Syntax: `SAMPLE_C_LZ dst, address, sampler_view.r, sampler, ref_value`

Example: `SAMPLE_C_LZ TEMP[0], TEMP[1], SVIEW[0].r, SAMP[0], TEMP[2].x`

SAMPLE_D

SAMPLE_D is identical to the SAMPLE opcode except that the derivatives for the source address in the x direction and the y direction are provided by extra parameters.

Syntax: `SAMPLE_D dst, address, sampler_view, sampler, der_x, der_y`

Example: `SAMPLE_D TEMP[0], TEMP[1], SVIEW[0], SAMP[0], TEMP[2], TEMP[3]`

SAMPLE_L

SAMPLE_L is identical to the SAMPLE opcode except that the LOD is provided directly as a scalar value, representing no anisotropy.

Syntax: `SAMPLE_L dst, address, sampler_view, sampler, explicit_lod`

Example: `SAMPLE_L TEMP[0], TEMP[1], SVIEW[0], SAMP[0], TEMP[2].x`

GATHER4

Gathers the four texels to be used in a bi-linear filtering operation and packs them into a single register. Only works with 2D, 2D array, cubemaps, and cubemaps arrays. For 2D textures, only the addressing modes of the sampler and the top level of any mip pyramid are used. Set W to zero. It behaves like the `SAMPLE` instruction, but a filtered sample is not generated. The four samples that contribute to filtering are placed into `xyzw` in counter-clockwise order, starting with the (u,v) texture coordinate delta at the following locations (-, +), (+, +), (+, -), (-, -), where the magnitude of the deltas are half a texel.

SVIEWINFO

Query the dimensions of a given sampler view. `dst` receives width, height, depth or array size and number of mipmap levels as `int4`. The `dst` can have a writemask which will specify what info is the caller interested in.

Syntax: `SVIEWINFO dst, src_mip_level, sampler_view`

Example: `SVIEWINFO TEMP[0], TEMP[1].x, SVIEW[0]`

`src_mip_level` is an unsigned integer scalar. If it's out of range then returns 0 for width, height and depth/array size but the total number of mipmap is still returned correctly for the given sampler view. The returned width, height and depth values are for the mipmap level selected by the `src_mip_level` and are in the number of texels. For 1d texture array width is in `dst.x`, array size is in `dst.y` and `dst.z` is 0. The number of mipmaps is still in `dst.w`. In contrast to `d3d10 resinfo`, there's no way in the `tgsi` instruction encoding to specify the return type (float/rcpfloat/uint), hence always using `uint`. Also, unlike the `SAMPLE` instructions, the swizzle on `src1 resinfo` allowing swizzling `dst` values is ignored (due to the interaction with `rcpfloat` modifier which requires some swizzle handling in the state tracker anyway).

SAMPLE_POS

Query the position of a sample in the given resource or render target when per-sample fragment shading is in effect.

Syntax: `SAMPLE_POS dst, source, sample_index`

`dst` receives `float4 (x, y, undef, undef)` indicated where the sample is located. Sample locations are in the range `[0, 1]` where 0.5 is the center of the fragment.

`source` is either a sampler view (to indicate a shader resource) or temp register (to indicate the render target). The source register may have an optional swizzle to apply to the returned result

`sample_index` is an integer scalar indicating which sample position is to be queried.

If per-sample shading is not in effect or the source resource or render target is not multisampled, the result is `(0.5, 0.5, undef, undef)`.

NOTE: no driver has implemented this opcode yet (and no state tracker emits it). This information is subject to change.

SAMPLE_INFO

Query the number of samples in a multisampled resource or render target.

Syntax: `SAMPLE_INFO dst, source`

`dst` receives `int4 (n, 0, 0, 0)` where `n` is the number of samples in a resource or the render target.

`source` is either a sampler view (to indicate a shader resource) or temp register (to indicate the render target). The source register may have an optional swizzle to apply to the returned result

If per-sample shading is not in effect or the source resource or render target is not multisampled, the result is `(1, 0, 0, 0)`.

NOTE: no driver has implemented this opcode yet (and no state tracker emits it). This information is subject to change.

LOD (level of detail)

Same syntax as the SAMPLE opcode but instead of performing an actual texture lookup/filter, return the computed LOD information that the texture pipe would use to access the texture. The Y component contains the computed LOD lambda_prime. The X component contains the LOD that will be accessed, based on min/max lod's and mipmap filters. The Z and W components are set to 0.

Syntax: LOD dst, address, sampler_view, sampler

Resource Access Opcodes

For these opcodes, the resource can be a BUFFER, IMAGE, or MEMORY.

LOAD (Fetch data from a shader buffer or image)

Syntax: LOAD dst, resource, address

Example: LOAD TEMP[0], BUFFER[0], TEMP[1]

Using the provided integer address, LOAD fetches data from the specified buffer or texture without any filtering.

The 'address' is specified as a vector of unsigned integers. If the 'address' is out of range the result is unspecified.

Only the first mipmap level of a resource can be read from using this instruction.

For 1D or 2D texture arrays, the array index is provided as an unsigned integer in address.y or address.z, respectively. address.yz are ignored for buffers and 1D textures. address.z is ignored for 1D texture arrays and 2D textures. address.w is always ignored.

A swizzle suffix may be added to the resource argument this will cause the resource data to be swizzled accordingly.

STORE (Write data to a shader resource)

Syntax: STORE resource, address, src

Example: STORE BUFFER[0], TEMP[0], TEMP[1]

Using the provided integer address, STORE writes data to the specified buffer or texture.

The 'address' is specified as a vector of unsigned integers. If the 'address' is out of range the result is unspecified.

Only the first mipmap level of a resource can be written to using this instruction.

For 1D or 2D texture arrays, the array index is provided as an unsigned integer in address.y or address.z, respectively. address.yz are ignored for buffers and 1D textures. address.z is ignored for 1D texture arrays and 2D textures. address.w is always ignored.

RESQ (Query information about a resource)

Syntax: RESQ dst, resource

Example: RESQ TEMP[0], BUFFER[0]

Returns information about the buffer or image resource. For buffer resources, the size (in bytes) is returned in the x component. For image resources, .xyz will contain the width/height/layers of the image, while .w will contain the number of samples for multi-sampled images.

FBFETCH (Load data from framebuffer)

Syntax: FBFETCH dst, output

Example: FBFETCH TEMP[0], OUT[0]

This is only valid on COLOR semantic outputs. Returns the color of the current position in the framebuffer from before this fragment shader invocation. May return the same value from multiple calls for a particular output

within a single invocation. Note that result may be undefined if a fragment is drawn multiple times without a blend barrier in between.

Inter-thread synchronization opcodes

These opcodes are intended for communication between threads running within the same compute grid. For now they're only valid in compute programs.

BARRIER (Thread group barrier)

BARRIER

This opcode suspends the execution of the current thread until all the remaining threads in the working group reach the same point of the program. Results are unspecified if any of the remaining threads terminates or never reaches an executed BARRIER instruction.

MEMBAR (Memory barrier)

MEMBAR *type*

This opcode waits for the completion of all memory accesses based on the type passed in. The type is an immediate bitfield with the following meaning:

Bit 0: Shader storage buffers Bit 1: Atomic buffers Bit 2: Images Bit 3: Shared memory Bit 4: Thread group

These may be passed in in any combination. An implementation is free to not distinguish between these as it sees fit. However these map to all the possibilities made available by GLSL.

Atomic opcodes

These opcodes provide atomic variants of some common arithmetic and logical operations. In this context atomicity means that another concurrent memory access operation that affects the same memory location is guaranteed to be performed strictly before or after the entire execution of the atomic operation. The resource may be a BUFFER, IMAGE, or MEMORY. In the case of an image, the offset works the same as for LOAD and STORE, specified above. These atomic operations may only be used with 32-bit integer image formats.

ATOMUADD (Atomic integer addition)

Syntax: ATOMUADD *dst*, *resource*, *offset*, *src*

Example: ATOMUADD TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned}dst_x &= resource[offset] \\ resource[offset] &= dst_x + src_x\end{aligned}$$

ATOMXCHG (Atomic exchange)

Syntax: ATOMXCHG *dst*, *resource*, *offset*, *src*

Example: ATOMXCHG TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned}dst_x &= resource[offset] \\ resource[offset] &= src_x\end{aligned}$$

ATOMCAS (Atomic compare-and-exchange)

Syntax: ATOMCAS *dst*, *resource*, *offset*, *cmp*, *src*

Example: ATOMCAS TEMP[0], BUFFER[0], TEMP[1], TEMP[2], TEMP[3]

The following operation is performed atomically:

$$\begin{aligned} dst_x &= resource[offset] \\ resource[offset] &= (dst_x == cmp_x ? src_x : dst_x) \end{aligned}$$

ATOMAND (Atomic bitwise And)

Syntax: ATOMAND *dst*, *resource*, *offset*, *src*

Example: ATOMAND TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned} dst_x &= resource[offset] \\ resource[offset] &= dst_x \& src_x \end{aligned}$$

ATOMOR (Atomic bitwise Or)

Syntax: ATOMOR *dst*, *resource*, *offset*, *src*

Example: ATOMOR TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned} dst_x &= resource[offset] \\ resource[offset] &= dst_x | src_x \end{aligned}$$

ATOMXOR (Atomic bitwise Xor)

Syntax: ATOMXOR *dst*, *resource*, *offset*, *src*

Example: ATOMXOR TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned} dst_x &= resource[offset] \\ resource[offset] &= dst_x \oplus src_x \end{aligned}$$

ATOMUMIN (Atomic unsigned minimum)

Syntax: ATOMUMIN *dst*, *resource*, *offset*, *src*

Example: ATOMUMIN TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned} dst_x &= resource[offset] \\ resource[offset] &= (dst_x < src_x ? dst_x : src_x) \end{aligned}$$

ATOMUMAX (Atomic unsigned maximum)

Syntax: ATOMUMAX *dst*, *resource*, *offset*, *src*

Example: ATOMUMAX TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned}dst_x &= resource[offset] \\ resource[offset] &= (dst_x > src_x ? dst_x : src_x)\end{aligned}$$

ATOMIMIN (Atomic signed minimum)

Syntax: ATOMIMIN *dst*, *resource*, *offset*, *src*

Example: ATOMIMIN TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned}dst_x &= resource[offset] \\ resource[offset] &= (dst_x < src_x ? dst_x : src_x)\end{aligned}$$

ATOMIMAX (Atomic signed maximum)

Syntax: ATOMIMAX *dst*, *resource*, *offset*, *src*

Example: ATOMIMAX TEMP[0], BUFFER[0], TEMP[1], TEMP[2]

The following operation is performed atomically:

$$\begin{aligned}dst_x &= resource[offset] \\ resource[offset] &= (dst_x > src_x ? dst_x : src_x)\end{aligned}$$

Inter-lane opcodes

These opcodes reduce the given value across the shader invocations running in the current SIMD group. Every thread in the subgroup will receive the same result. The BALLOT operations accept a single-channel argument that is treated as a boolean and produce a 64-bit value.

VOTE_ANY (Value is set in any of the active invocations)

Syntax: VOTE_ANY *dst*, *value*

Example: VOTE_ANY TEMP[0].x, TEMP[1].x

VOTE_ALL (Value is set in all of the active invocations)

Syntax: VOTE_ALL *dst*, *value*

Example: VOTE_ALL TEMP[0].x, TEMP[1].x

VOTE_EQ (Value is the same in all of the active invocations)

Syntax: VOTE_EQ *dst*, *value*

Example: VOTE_EQ TEMP[0].x, TEMP[1].x

BALLOT (Lanemask of whether the value is set in each active)

invocation

Syntax: `BALLOT dst, value`

Example: `BALLOT TEMP[0].xy, TEMP[1].x`

When the argument is a constant true, this produces a bitmask of active invocations. In fragment shaders, this can include helper invocations (invocations whose outputs and writes to memory are discarded, but which are used to compute derivatives).

READ_FIRST (Broadcast the value from the first active)**invocation to all active lanes**

Syntax: `READ_FIRST dst, value`

Example: `READ_FIRST TEMP[0], TEMP[1]`

READ_INVOC (Retrieve the value from the given invocation)**(need not be uniform)**

Syntax: `READ_INVOC dst, value, invocation`

Example: `READ_INVOC TEMP[0].xy, TEMP[1].xy, TEMP[2].x`

`invocation.x` controls the invocation number to read from for all channels. The invocation number must be the same across all active invocations in a sub-group; otherwise, the results are undefined.

Explanation of symbols used

Functions

$|x|$ Absolute value of x .

$\lceil x \rceil$ Ceiling of x .

clamp(x,y,z) Clamp x between y and z . $(x < y) ? y : (x > z) ? z : x$

$\lfloor x \rfloor$ Floor of x .

$\log_2 x$ Logarithm of x , base 2.

max(x,y) Maximum of x and y . $(x > y) ? x : y$

min(x,y) Minimum of x and y . $(x < y) ? x : y$

`partialx(x)` Derivative of x relative to fragment's X .

`partially(x)` Derivative of x relative to fragment's Y .

`pop()` Pop from stack.

x^y x to the power y .

`push(x)` Push x on stack.

`round(x)` Round x .

`trunc(x)` Truncate x , i.e. drop the fraction bits.

Keywords

`discard` Discard fragment.

`pc` Program counter.

target Label of target instruction.

Other tokens

Declaration

Declares a register that is will be referenced as an operand in Instruction tokens.

File field contains register file that is being declared and is one of TGSI_FILE.

UsageMask field specifies which of the register components can be accessed and is one of TGSI_WRITEMASK.

The Local flag specifies that a given value isn't intended for subroutine parameter passing and, as a result, the implementation isn't required to give any guarantees of it being preserved across subroutine boundaries. As it's merely a compiler hint, the implementation is free to ignore it.

If Dimension flag is set to 1, a Declaration Dimension token follows.

If Semantic flag is set to 1, a Declaration Semantic token follows.

If Interpolate flag is set to 1, a Declaration Interpolate token follows.

If file is TGSI_FILE_RESOURCE, a Declaration Resource token follows.

If Array flag is set to 1, a Declaration Array token follows.

Array Declaration

Declarations can optional have an ArrayID attribute which can be referred by indirect addressing operands. An ArrayID of zero is reserved and treated as if no ArrayID is specified.

If an indirect addressing operand refers to a specific declaration by using an ArrayID only the registers in this declaration are guaranteed to be accessed, accessing any register outside this declaration results in undefined behavior. Note that for compatibility the effective index is zero-based and not relative to the specified declaration

If no ArrayID is specified with an indirect addressing operand the whole register file might be accessed by this operand. This is strongly discouraged and will prevent packing of scalar/vec2 arrays and effective alias analysis. This is only legal for TEMP and CONST register files.

Declaration Semantic

Vertex and fragment shader input and output registers may be labeled with semantic information consisting of a name and index.

Follows Declaration token if Semantic bit is set.

Since its purpose is to link a shader with other stages of the pipeline, it is valid to follow only those Declaration tokens that declare a register either in INPUT or OUTPUT file.

SemanticName field contains the semantic name of the register being declared. There is no default value.

SemanticIndex is an optional subscript that can be used to distinguish different register declarations with the same semantic name. The default value is 0.

The meanings of the individual semantic names are explained in the following sections.

TGSI_SEMANTIC_POSITION

For vertex shaders, TGSI_SEMANTIC_POSITION indicates the vertex shader output register which contains the homogeneous vertex position in the clip space coordinate system. After clipping, the X, Y and Z components of the vertex will be divided by the W value to get normalized device coordinates.

For fragment shaders, TGSI_SEMANTIC_POSITION is used to indicate that fragment shader input (or system value, depending on which one is supported by the driver) contains the fragment's window position. The X component starts at zero and always increases from left to right. The Y component starts at zero and always increases but Y=0 may either indicate the top of the window or the bottom depending on the fragment coordinate origin convention (see TGSI_PROPERTY_FS_COORD_ORIGIN). The Z coordinate ranges from 0 to 1 to represent depth from the front to the back of the Z buffer. The W component contains the interpolated reciprocal of the vertex position W component (corresponding to `gl_FragCoord`, but unlike `d3d10` which interpolates the same $1/w$ but then gives back the reciprocal of the interpolated value).

Fragment shaders may also declare an output register with TGSI_SEMANTIC_POSITION. Only the Z component is writable. This allows the fragment shader to change the fragment's Z position.

TGSI_SEMANTIC_COLOR

For vertex shader outputs or fragment shader inputs/outputs, this label indicates that the register contains an R,G,B,A color.

Several shader inputs/outputs may contain colors so the semantic index is used to distinguish them. For example, `color[0]` may be the diffuse color while `color[1]` may be the specular color.

This label is needed so that the flat/smooth shading can be applied to the right interpolants during rasterization.

TGSI_SEMANTIC_BCOLOR

Back-facing colors are only used for back-facing polygons, and are only valid in vertex shader outputs. After rasterization, all polygons are front-facing and COLOR and BCOLOR end up occupying the same slots in the fragment shader, so all BCOLORs effectively become regular COLORS in the fragment shader.

TGSI_SEMANTIC_FOG

Vertex shader inputs and outputs and fragment shader inputs may be labeled with TGSI_SEMANTIC_FOG to indicate that the register contains a fog coordinate. Typically, the fragment shader will use the fog coordinate to compute a fog blend factor which is used to blend the normal fragment color with a constant fog color. But fog coord really is just an ordinary `vec4` register like regular semantics.

TGSI_SEMANTIC_PSIZE

Vertex shader input and output registers may be labeled with TGSI_SEMANTIC_PSIZE to indicate that the register contains a point size in the form (S, 0, 0, 1). The point size controls the width or diameter of points for rasterization. This label cannot be used in fragment shaders.

When using this semantic, be sure to set the appropriate state in the *Rasterizer* first.

TGSI_SEMANTIC_TEXCOORD

Only available if PIPE_CAP_TGSI_TEXCOORD is exposed !

Vertex shader outputs and fragment shader inputs may be labeled with this semantic to make them replaceable by sprite coordinates via the `sprite_coord_enable` state in the *Rasterizer*. The semantic index permitted with this semantic is limited to ≤ 7 .

If the driver does not support `TEXCOORD`, sprite coordinate replacement applies to inputs with the `GENERIC` semantic instead.

The intended use case for this semantic is `gl_TexCoord`.

TGSI_SEMANTIC_PCOORD

Only available if `PIPE_CAP_TGSI_TEXCOORD` is exposed !

Fragment shader inputs may be labeled with `TGSI_SEMANTIC_PCOORD` to indicate that the register contains sprite coordinates in the form $(x, y, 0, 1)$, if the current primitive is a point and point sprites are enabled. Otherwise, the contents of the register are undefined.

The intended use case for this semantic is `gl_PointCoord`.

TGSI_SEMANTIC_GENERIC

All vertex/fragment shader inputs/outputs not labeled with any other semantic label can be considered to be generic attributes. Typical uses of generic inputs/outputs are texcoords and user-defined values.

TGSI_SEMANTIC_NORMAL

Indicates that a vertex shader input is a normal vector. This is typically only used for legacy graphics APIs.

TGSI_SEMANTIC_FACE

This label applies to fragment shader inputs (or system values, depending on which one is supported by the driver) and indicates that the register contains front/back-face information.

If it is an input, it will be a floating-point vector in the form $(F, 0, 0, 1)$, where F will be positive when the fragment belongs to a front-facing polygon, and negative when the fragment belongs to a back-facing polygon.

If it is a system value, it will be an integer vector in the form $(F, 0, 0, 1)$, where F is `0xffffffff` when the fragment belongs to a front-facing polygon and `0` when the fragment belongs to a back-facing polygon.

TGSI_SEMANTIC_EDGEFLAG

For vertex shaders, this semantic label indicates that an input or output is a boolean edge flag. The register layout is $[F, x, x, x]$ where F is `0.0` or `1.0` and $x = \text{don't care}$. Normally, the vertex shader simply copies the edge flag input to the edgeflag output.

Edge flags are used to control which lines or points are actually drawn when the polygon mode converts triangles/quads/polygons into points or lines.

TGSI_SEMANTIC_STENCIL

For fragment shaders, this semantic label indicates that an output is a writable stencil reference value. Only the Y component is writable. This allows the fragment shader to change the fragments `stencilref` value.

TGSI_SEMANTIC_VIEWPORT_INDEX

For geometry shaders, this semantic label indicates that an output contains the index of the viewport (and scissor) to use. This is an integer value, and only the X component is used.

If PIPE_CAP_TGSI_VS_LAYER_VIEWPORT or PIPE_CAP_TGSI_TES_LAYER_VIEWPORT is supported, then this semantic label can also be used in vertex or tessellation evaluation shaders, respectively. Only the value written in the last vertex processing stage is used.

TGSI_SEMANTIC_LAYER

For geometry shaders, this semantic label indicates that an output contains the layer value to use for the color and depth/stencil surfaces. This is an integer value, and only the X component is used. (Also known as rendertarget array index.)

If PIPE_CAP_TGSI_VS_LAYER_VIEWPORT or PIPE_CAP_TGSI_TES_LAYER_VIEWPORT is supported, then this semantic label can also be used in vertex or tessellation evaluation shaders, respectively. Only the value written in the last vertex processing stage is used.

TGSI_SEMANTIC_CULLDIST

Used as distance to plane for performing application-defined culling of individual primitives against a plane. When components of vertex elements are given this label, these values are assumed to be a float32 signed distance to a plane. Primitives will be completely discarded if the plane distance for all of the vertices in the primitive are < 0 . If a vertex has a cull distance of NaN, that vertex counts as “out” (as if its < 0); The limits on both clip and cull distances are bound by the PIPE_MAX_CLIP_OR_CULL_DISTANCE_COUNT define which defines the maximum number of components that can be used to hold the distances and by the PIPE_MAX_CLIP_OR_CULL_DISTANCE_ELEMENT_COUNT which specifies the maximum number of registers which can be annotated with those semantics.

TGSI_SEMANTIC_CLIPDIST

Note this covers clipping and culling distances.

When components of vertex elements are identified this way, these values are each assumed to be a float32 signed distance to a plane.

For clip distances: Primitive setup only invokes rasterization on pixels for which the interpolated plane distances are ≥ 0 .

For cull distances: Primitives will be completely discarded if the plane distance for all of the vertices in the primitive are < 0 . If a vertex has a cull distance of NaN, that vertex counts as “out” (as if its < 0);

Multiple clip/cull planes can be implemented simultaneously, by annotating multiple components of one or more vertex elements with the above specified semantic. The limits on both clip and cull distances are bound by the PIPE_MAX_CLIP_OR_CULL_DISTANCE_COUNT define which defines the maximum number of components that can be used to hold the distances and by the PIPE_MAX_CLIP_OR_CULL_DISTANCE_ELEMENT_COUNT which specifies the maximum number of registers which can be annotated with those semantics. The properties NUM_CLIPDIST_ENABLED and NUM_CULLDIST_ENABLED are used to divide up the $2 \times \text{vec4}$ space between clipping and culling.

TGSI_SEMANTIC_SAMPLEID

For fragment shaders, this semantic label indicates that a system value contains the current sample id (i.e. `gl_SampleID`) as an unsigned int. Only the X component is used. If per-sample shading is not enabled, the result is (0, undef, undef, undef).

Note that if the fragment shader uses this system value, the fragment shader is automatically executed at per sample frequency.

TGSI_SEMANTIC_SAMPLEPOS

For fragment shaders, this semantic label indicates that a system value contains the current sample's position as `float4(x, y, undef, undef)` in the render target (i.e. `gl_SamplePosition`) when per-fragment shading is in effect. Position values are in the range [0, 1] where 0.5 is the center of the fragment.

Note that if the fragment shader uses this system value, the fragment shader is automatically executed at per sample frequency.

TGSI_SEMANTIC_SAMPLEMASK

For fragment shaders, this semantic label can be applied to either a shader system value input or output.

For a system value, the sample mask indicates the set of samples covered by the current primitive. If MSAA is not enabled, the value is (1, 0, 0, 0).

For an output, the sample mask is used to disable further sample processing.

For both, the register type is `uint[4]` but only the X component is used (i.e. `gl_SampleMask[0]`). Each bit corresponds to one sample position (up to 32x MSAA is supported).

TGSI_SEMANTIC_INVOCATIONID

For geometry shaders, this semantic label indicates that a system value contains the current invocation id (i.e. `gl_InvocationID`). This is an integer value, and only the X component is used.

TGSI_SEMANTIC_INSTANCEID

For vertex shaders, this semantic label indicates that a system value contains the current instance id (i.e. `gl_InstanceID`). It does not include the base instance. This is an integer value, and only the X component is used.

TGSI_SEMANTIC_VERTEXID

For vertex shaders, this semantic label indicates that a system value contains the current vertex id (i.e. `gl_VertexID`). It does (unlike in d3d10) include the base vertex. This is an integer value, and only the X component is used.

TGSI_SEMANTIC_VERTEXID_NOBASE

For vertex shaders, this semantic label indicates that a system value contains the current vertex id without including the base vertex (this corresponds to d3d10 vertex id, so `TGSI_SEMANTIC_VERTEXID_NOBASE + TGSI_SEMANTIC_BASEVERTEX == TGSI_SEMANTIC_VERTEXID`). This is an integer value, and only the X component is used.

TGSI_SEMANTIC_BASEVERTEX

For vertex shaders, this semantic label indicates that a system value contains the base vertex (i.e. `gl_BaseVertex`). Note that for non-indexed draw calls, this contains the first (or start) value instead. This is an integer value, and only the X component is used.

TGSI_SEMANTIC_PRIMID

For geometry and fragment shaders, this semantic label indicates the value contains the primitive id (i.e. `gl_PrimitiveID`). This is an integer value, and only the X component is used. FIXME: This right now can be either a ordinary input or a system value...

TGSI_SEMANTIC_PATCH

For tessellation evaluation/control shaders, this semantic label indicates a generic per-patch attribute. Such semantics will not implicitly be per-vertex arrays.

TGSI_SEMANTIC_TESSCOORD

For tessellation evaluation shaders, this semantic label indicates the coordinates of the vertex being processed. This is available in XYZ; W is undefined.

TGSI_SEMANTIC_TESSOUTER

For tessellation evaluation/control shaders, this semantic label indicates the outer tessellation levels of the patch. Isoline tessellation will only have XY defined, triangle will have XYZ and quads will have XYZW defined. This corresponds to `gl_TessLevelOuter`.

TGSI_SEMANTIC_TESSINNER

For tessellation evaluation/control shaders, this semantic label indicates the inner tessellation levels of the patch. The X value is only defined for triangle tessellation, while quads will have XY defined. This is entirely undefined for isoline tessellation.

TGSI_SEMANTIC_VERTICESIN

For tessellation evaluation/control shaders, this semantic label indicates the number of vertices provided in the input patch. Only the X value is defined.

TGSI_SEMANTIC_HELPER_INVOCATION

For fragment shaders, this semantic indicates whether the current invocation is covered or not. Helper invocations are created in order to properly compute derivatives, however it may be desirable to skip some of the logic in those cases. See `gl_HelperInvocation` documentation.

TGSI_SEMANTIC_BASEINSTANCE

For vertex shaders, the base instance argument supplied for this draw. This is an integer value, and only the X component is used.

TGSI_SEMANTIC_DRAWID

For vertex shaders, the zero-based index of the current draw in a `glMultiDraw*` invocation. This is an integer value, and only the X component is used.

TGSI_SEMANTIC_WORK_DIM

For compute shaders started via `opengl` this retrieves the `work_dim` parameter to the `clEnqueueNDRangeKernel` call with which the shader was started.

TGSI_SEMANTIC_GRID_SIZE

For compute shaders, this semantic indicates the maximum (x, y, z) dimensions of a grid of thread blocks.

TGSI_SEMANTIC_BLOCK_ID

For compute shaders, this semantic indicates the (x, y, z) coordinates of the current block inside of the grid.

TGSI_SEMANTIC_BLOCK_SIZE

For compute shaders, this semantic indicates the maximum (x, y, z) dimensions of a block in threads.

TGSI_SEMANTIC_THREAD_ID

For compute shaders, this semantic indicates the (x, y, z) coordinates of the current thread inside of the block.

TGSI_SEMANTIC_SUBGROUP_SIZE

This semantic indicates the subgroup size for the current invocation. This is an integer of at most 64, as it indicates the width of lanemasks. It does not depend on the number of invocations that are active.

TGSI_SEMANTIC_SUBGROUP_INVOCATION

The index of the current invocation within its subgroup.

TGSI_SEMANTIC_SUBGROUP_EQ_MASK

A bit mask of `bit index == TGSI_SEMANTIC_SUBGROUP_INVOCATION`, i.e. $1 \ll \text{subgroup_invocation}$ in arbitrary precision arithmetic.

TGSI_SEMANTIC_SUBGROUP_GE_MASK

A bit mask of `bit index >= TGSI_SEMANTIC_SUBGROUP_INVOCATION`, i.e. $((1 \ll (\text{subgroup_size} - \text{subgroup_invocation})) - 1) \ll \text{subgroup_invocation}$ in arbitrary precision arithmetic.

TGSI_SEMANTIC_SUBGROUP_GT_MASK

A bit mask of bit index $>$ TGSI_SEMANTIC_SUBGROUP_INVOCATION, i.e. $((1 \ll (\text{subgroup_size} - \text{subgroup_invocation} - 1)) - 1) \ll (\text{subgroup_invocation} + 1)$ in arbitrary precision arithmetic.

TGSI_SEMANTIC_SUBGROUP_LE_MASK

A bit mask of bit index \leq TGSI_SEMANTIC_SUBGROUP_INVOCATION, i.e. $(1 \ll (\text{subgroup_invocation} + 1)) - 1$ in arbitrary precision arithmetic.

TGSI_SEMANTIC_SUBGROUP_LT_MASK

A bit mask of bit index $<$ TGSI_SEMANTIC_SUBGROUP_INVOCATION, i.e. $(1 \ll \text{subgroup_invocation}) - 1$ in arbitrary precision arithmetic.

Declaration Interpolate

This token is only valid for fragment shader INPUT declarations.

The Interpolate field specifies the way input is being interpolated by the rasteriser and is one of TGSI_INTERPOLATE_*.

The Location field specifies the location inside the pixel that the interpolation should be done at, one of TGSI_INTERPOLATE_LOC_*. Note that when per-sample shading is enabled, the implementation may choose to interpolate at the sample irrespective of the Location field.

The CylindricalWrap bitfield specifies which register components should be subject to cylindrical wrapping when interpolating by the rasteriser. If TGSI_CYLINDRICAL_WRAP_X is set to 1, the X component should be interpolated according to cylindrical wrapping rules.

Declaration Sampler View

Follows Declaration token if file is TGSI_FILE_SAMPLER_VIEW.

DCL SVIEW[#], resource, type(s)

Declares a shader input sampler view and assigns it to a SVIEW[#] register.

resource can be one of BUFFER, 1D, 2D, 3D, 1DArray and 2DArray.

type must be 1 or 4 entries (if specifying on a per-component level) out of UNORM, SNORM, SINT, UINT and FLOAT.

For TEX* style texture sample opcodes (as opposed to SAMPLE* opcodes which take an explicit SVIEW[#] source register), there may be optionally SVIEW[#] declarations. In this case, the SVIEW index is implied by the SAMP index, and there must be a corresponding SVIEW[#] declaration for each SAMP[#] declaration. Drivers are free to ignore this if they wish. But note in particular that some drivers need to know the sampler type (float/int/unsigned) in order to generate the correct code, so cases where integer textures are sampled, SVIEW[#] declarations should be used.

NOTE: It is NOT legal to mix SAMPLE* style opcodes and TEX* opcodes in the same shader.

Declaration Resource

Follows Declaration token if file is TGSI_FILE_RESOURCE.

DCL RES[#], resource [, WR] [, RAW]

Declares a shader input resource and assigns it to a RES[#] register.

resource can be one of BUFFER, 1D, 2D, 3D, CUBE, 1DArray and 2DArray.

If the RAW keyword is not specified, the texture data will be subject to conversion, swizzling and scaling as required to yield the specified data type from the physical data format of the bound resource.

If the RAW keyword is specified, no channel conversion will be performed: the values read for each of the channels (X,Y,Z,W) will correspond to consecutive words in the same order and format they're found in memory. No element-to-address conversion will be performed either: the value of the provided X coordinate will be interpreted in byte units instead of texel units. The result of accessing a misaligned address is undefined.

Usage of the STORE opcode is only allowed if the WR (writable) flag is set.

Properties

Properties are general directives that apply to the whole TGSI program.

FS_COORD_ORIGIN

Specifies the fragment shader TGSI_SEMANTIC_POSITION coordinate origin. The default value is UPPER_LEFT.

If UPPER_LEFT, the position will be (0,0) at the upper left corner and increase downward and rightward. If LOWER_LEFT, the position will be (0,0) at the lower left corner and increase upward and rightward.

OpenGL defaults to LOWER_LEFT, and is configurable with the GL_ARB_fragment_coord_conventions extension.

DirectX 9/10 use UPPER_LEFT.

FS_COORD_PIXEL_CENTER

Specifies the fragment shader TGSI_SEMANTIC_POSITION pixel center convention. The default value is HALF_INTEGER.

If HALF_INTEGER, the fractionary part of the position will be 0.5 If INTEGER, the fractionary part of the position will be 0.0

Note that this does not affect the set of fragments generated by rasterization, which is instead controlled by half_pixel_center in the rasterizer.

OpenGL defaults to HALF_INTEGER, and is configurable with the GL_ARB_fragment_coord_conventions extension.

DirectX 9 uses INTEGER. DirectX 10 uses HALF_INTEGER.

FS_COLOR0_WRITES_ALL_CBUFS

Specifies that writes to the fragment shader color 0 are replicated to all bound cbufs. This facilitates OpenGL's fragColor output vs fragData[0] where fragData is directed to a single color buffer, but fragColor is broadcast.

VS_PROHIBIT_UCPS

If this property is set on the program bound to the shader stage before the fragment shader, user clip planes should have no effect (be disabled) even if that shader does not write to any clip distance outputs and the rasterizer's `clip_plane_enable` is non-zero. This property is only supported by drivers that also support shader clip distance outputs. This is useful for APIs that don't have UCPs and where clip distances written by a shader cannot be disabled.

GS_INVOCATIONS

Specifies the number of times a geometry shader should be executed for each input primitive. Each invocation will have a different `TGSI_SEMANTIC_INVOCATIONID` system value set. If not specified, assumed to be 1.

VS_WINDOW_SPACE_POSITION

If this property is set on the vertex shader, the `TGSI_SEMANTIC_POSITION` output is assumed to contain window space coordinates. Division of X,Y,Z by W and the viewport transformation are disabled, and 1/W is directly taken from the 4-th component of the shader output. Naturally, clipping is not performed on window coordinates either. The effect of this property is undefined if a geometry or tessellation shader are in use.

TCS_VERTICES_OUT

The number of vertices written by the tessellation control shader. This effectively defines the patch input size of the tessellation evaluation shader as well.

TES_PRIM_MODE

This sets the tessellation primitive mode, one of `PIPE_PRIM_TRIANGLES`, `PIPE_PRIM_QUADS`, or `PIPE_PRIM_LINES`. (Unlike in GL, there is no separate isolines settings, the regular lines is assumed to mean isolines.)

TES_SPACING

This sets the spacing mode of the tessellation generator, one of `PIPE_TESS_SPACING_*`.

TES_VERTEX_ORDER_CW

This sets the vertex order to be clockwise if the value is 1, or counter-clockwise if set to 0.

TES_POINT_MODE

If set to a non-zero value, this turns on point mode for the tessellator, which means that points will be generated instead of primitives.

NUM_CLIPDIST_ENABLED

How many clip distance scalar outputs are enabled.

NUM_CULLDIST_ENABLED

How many cull distance scalar outputs are enabled.

FS_EARLY_DEPTH_STENCIL

Whether depth test, stencil test, and occlusion query should run before the fragment shader (regardless of fragment shader side effects). Corresponds to GLSL `early_fragment_tests`.

NEXT_SHADER

Which shader stage will MOST LIKELY follow after this shader when the shader is bound. This is only a hint to the driver and doesn't have to be precise. Only set for VS and TES.

CS_FIXED_BLOCK_WIDTH / HEIGHT / DEPTH

Threads per block in each dimension, if known at compile time. If the block size is known all three should be at least 1. If it is unknown they should all be set to 0 or not set.

MUL_ZERO_WINS

The MUL TGSI operation (FP32 multiplication) will return 0 if either of the operands are equal to 0. That means that $0 * \text{Inf} = 0$. This should be set the same way for an entire pipeline. Note that this applies not only to the literal MUL TGSI opcode, but all FP32 multiplications implied by other operations, such as MAD, FMA, DP2, DP3, DP4, DST, LOG, LRP, and possibly others. If there is a mismatch between shaders, then it is unspecified whether this behavior will be enabled.

FS_POST_DEPTH_COVERAGE

When enabled, the input for TGSI_SEMANTIC_SAMPLEMASK will exclude samples that have failed the depth/stencil tests. This is only valid when FS_EARLY_DEPTH_STENCIL is also specified.

Texture Sampling and Texture Formats

This table shows how texture image components are returned as (x,y,z,w) tuples by TGSI texture instructions, such as *TEX*, *TXD*, and *TXP*. For reference, OpenGL and Direct3D conventions are shown as well.

Texture Components	Gallium	OpenGL	Direct3D 9
R	(r, 0, 0, 1)	(r, 0, 0, 1)	(r, 1, 1, 1)
RG	(r, g, 0, 1)	(r, g, 0, 1)	(r, g, 1, 1)
RGB	(r, g, b, 1)	(r, g, b, 1)	(r, g, b, 1)
RGBA	(r, g, b, a)	(r, g, b, a)	(r, g, b, a)
A	(0, 0, 0, a)	(0, 0, 0, a)	(0, 0, 0, a)
L	(l, l, l, 1)	(l, l, l, 1)	(l, l, l, 1)
LA	(l, l, l, a)	(l, l, l, a)	(l, l, l, a)
I	(i, i, i, i)	(i, i, i, i)	N/A
UV	XXX TBD	(0, 0, 0, 1) ¹	(u, v, 1, 1)
Z	XXX TBD	(z, z, z, 1) ²	(0, z, 0, 1)
S	(s, s, s, s)	unknown	unknown

¹ http://www.opengl.org/registry/specs/ATI/envmap_bumpmap.txt

² the default is (z, z, z, 1) but may also be (0, 0, 0, z) or (z, z, z, z) depending on the value of GL_DEPTH_TEXTURE_MODE.

A screen is an object representing the context-independent part of a device.

Flags and enumerations

XXX some of these don't belong in this section.

PIPE_CAP_*

Capability queries return information about the features and limits of the driver/GPU. For floating-point values, use *get_paramf*, and for boolean or integer values, use *get_param*.

The integer capabilities:

- **PIPE_CAP_NPOT_TEXTURES:** Whether *NPOT* textures may have repeat modes, normalized coordinates, and mipmaps.
- **PIPE_CAP_TWO_SIDED_STENCIL:** Whether the stencil test can also affect back-facing polygons.
- **PIPE_CAP_MAX_DUAL_SOURCE_RENDER_TARGETS:** How many dual-source blend RTs are supported. *Blend* for more information.
- **PIPE_CAP_ANISOTROPIC_FILTER:** Whether textures can be filtered anisotropically.
- **PIPE_CAP_POINT_SPRITE:** Whether point sprites are available.
- **PIPE_CAP_MAX_RENDER_TARGETS:** The maximum number of render targets that may be bound.
- **PIPE_CAP_OCCLUSION_QUERY:** Whether occlusion queries are available.
- **PIPE_CAP_QUERY_TIME_ELAPSED:** Whether **PIPE_QUERY_TIME_ELAPSED** queries are available.
- **PIPE_CAP_TEXTURE_SHADOW_MAP:** indicates whether the fragment shader hardware can do the depth texture / Z comparison operation in **TEX** instructions for shadow testing.
- **PIPE_CAP_TEXTURE_SWIZZLE:** Whether swizzling through sampler views is supported.

- `PIPE_CAP_MAX_TEXTURE_2D_LEVELS`: The maximum number of mipmap levels available for a 2D texture.
- `PIPE_CAP_MAX_TEXTURE_3D_LEVELS`: The maximum number of mipmap levels available for a 3D texture.
- `PIPE_CAP_MAX_TEXTURE_CUBE_LEVELS`: The maximum number of mipmap levels available for a cube-map.
- `PIPE_CAP_TEXTURE_MIRROR_CLAMP`: Whether mirrored texture coordinates with clamp are supported.
- `PIPE_CAP_BLEND_EQUATION_SEPARATE`: Whether alpha blend equations may be different from color blend equations, in *Blend* state.
- `PIPE_CAP_SM3`: Whether the vertex shader and fragment shader support equivalent opcodes to the Shader Model 3 specification. XXX oh god this is horrible
- `PIPE_CAP_MAX_STREAM_OUTPUT_BUFFERS`: The maximum number of stream buffers.
- `PIPE_CAP_PRIMITIVE_RESTART`: Whether primitive restart is supported.
- `PIPE_CAP_INDEP_BLEND_ENABLE`: Whether per-rendertarget blend enabling and channel masks are supported. If 0, then the first rendertarget's blend mask is replicated across all MRTs.
- `PIPE_CAP_INDEP_BLEND_FUNC`: Whether per-rendertarget blend functions are available. If 0, then the first rendertarget's blend functions affect all MRTs.
- `PIPE_CAP_MAX_TEXTURE_ARRAY_LAYERS`: The maximum number of texture array layers supported. If 0, the array textures are not supported at all and the `ARRAY` texture targets are invalid.
- `PIPE_CAP_TGSI_FS_COORD_ORIGIN_UPPER_LEFT`: Whether the TGSI property `FS_COORD_ORIGIN` with value `UPPER_LEFT` is supported.
- `PIPE_CAP_TGSI_FS_COORD_ORIGIN_LOWER_LEFT`: Whether the TGSI property `FS_COORD_ORIGIN` with value `LOWER_LEFT` is supported.
- `PIPE_CAP_TGSI_FS_COORD_PIXEL_CENTER_HALF_INTEGER`: Whether the TGSI property `FS_COORD_PIXEL_CENTER` with value `HALF_INTEGER` is supported.
- `PIPE_CAP_TGSI_FS_COORD_PIXEL_CENTER_INTEGER`: Whether the TGSI property `FS_COORD_PIXEL_CENTER` with value `INTEGER` is supported.
- `PIPE_CAP_DEPTH_CLIP_DISABLE`: Whether the driver is capable of disabling depth clipping (through `pipe_rasterizer_state`)
- `PIPE_CAP_SHADER_STENCIL_EXPORT`: Whether a stencil reference value can be written from a fragment shader.
- `PIPE_CAP_TGSI_INSTANCEID`: Whether `TGSI_SEMANTIC_INSTANCEID` is supported in the vertex shader.
- `PIPE_CAP_VERTEX_ELEMENT_INSTANCE_DIVISOR`: Whether the driver supports per-instance vertex attribs.
- `PIPE_CAP_FRAGMENT_COLOR_CLAMPED`: Whether fragment color clamping is supported. That is, is the `pipe_rasterizer_state::clamp_fragment_color` flag supported by the driver? If not, the state tracker will insert clamping code into the fragment shaders when needed.
- `PIPE_CAP_MIXED_COLORBUFFER_FORMATS`: Whether mixed colorbuffer formats are supported, e.g. `RGBA8` and `RGBA32F` as the first and second colorbuffer, resp.
- `PIPE_CAP_VERTEX_COLOR_UNCLAMPED`: Whether the driver is capable of outputting unclamped vertex colors from a vertex shader. If unsupported, the vertex colors are always clamped. This is the default for DX9 hardware.

- `PIPE_CAP_VERTEX_COLOR_CLAMPED`: Whether the driver is capable of clamping vertex colors when they come out of a vertex shader, as specified by the `pipe_rasterizer_state::clamp_vertex_color` flag. If unsupported, the vertex colors are never clamped. This is the default for DX10 hardware. If both clamped and unclamped CAPs are supported, the clamping can be controlled through `pipe_rasterizer_state`. If the driver cannot do vertex color clamping, the state tracker may insert clamping code into the vertex shader.
- `PIPE_CAP_GLSL_FEATURE_LEVEL`: Whether the driver supports features equivalent to a specific GLSL version. E.g. for GLSL 1.3, report 130.
- `PIPE_CAP_QUADS_FOLLOW_PROVOKING_VERTEX_CONVENTION`: Whether quads adhere to the `flatshade_first` setting in `pipe_rasterizer_state`.
- `PIPE_CAP_USER_VERTEX_BUFFERS`: Whether the driver supports user vertex buffers. If not, the state tracker must upload all data which is not in hw resources. If user-space buffers are supported, the driver must also still accept HW resource buffers.
- `PIPE_CAP_VERTEX_BUFFER_OFFSET_4BYTE_ALIGNED_ONLY`: This CAP describes a hw limitation. If true, `pipe_vertex_buffer::buffer_offset` must always be aligned to 4. If false, there are no restrictions on the offset.
- `PIPE_CAP_VERTEX_BUFFER_STRIDE_4BYTE_ALIGNED_ONLY`: This CAP describes a hw limitation. If true, `pipe_vertex_buffer::stride` must always be aligned to 4. If false, there are no restrictions on the stride.
- `PIPE_CAP_VERTEX_ELEMENT_SRC_OFFSET_4BYTE_ALIGNED_ONLY`: This CAP describes a hw limitation. If true, `pipe_vertex_element::src_offset` must always be aligned to 4. If false, there are no restrictions on `src_offset`.
- `PIPE_CAP_COMPUTE`: Whether the implementation supports the compute entry points defined in `pipe_context` and `pipe_screen`.
- `PIPE_CAP_USER_CONSTANT_BUFFERS`: Whether user-space constant buffers are supported. If not, the state tracker must put constants into HW resources/buffers. If user-space constant buffers are supported, the driver must still accept HW constant buffers also.
- `PIPE_CAP_CONSTANT_BUFFER_OFFSET_ALIGNMENT`: Describes the required alignment of `pipe_constant_buffer::buffer_offset`.
- `PIPE_CAP_START_INSTANCE`: Whether the driver supports `pipe_draw_info::start_instance`.
- `PIPE_CAP_QUERY_TIMESTAMP`: Whether `PIPE_QUERY_TIMESTAMP` and the `pipe_screen::get_timestamp` hook are implemented.
- `PIPE_CAP_TEXTURE_MULTISAMPLE`: Whether all MSAA resources supported for rendering are also supported for texturing.
- `PIPE_CAP_MIN_MAP_BUFFER_ALIGNMENT`: The minimum alignment that should be expected for a pointer returned by `transfer_map` if the resource is `PIPE_BUFFER`. In other words, the pointer returned by `transfer_map` is always aligned to this value.
- `PIPE_CAP_TEXTURE_BUFFER_OFFSET_ALIGNMENT`: Describes the required alignment for `pipe_sampler_view::u.buf.offset`, in bytes. If a driver does not support offset/size, it should return 0.
- `PIPE_CAP_BUFFER_SAMPLER_VIEW_RGBA_ONLY`: Whether the driver only supports R, RG, RGB and RGBA formats for `PIPE_BUFFER` sampler views. When this is the case it should be assumed that the swizzle parameters in the sampler view have no effect.
- `PIPE_CAP_TGSI_TEXCOORD`: This CAP describes a hw limitation. If true, the hardware cannot replace arbitrary shader inputs with sprite coordinates and hence the inputs that are desired to be replaceable must be declared with `TGSI_SEMANTIC_TEXCOORD` instead of `TGSI_SEMANTIC_GENERIC`. The rasterizer's `sprite_coord_enable` state therefore also applies to the `TEXCOORD` semantic. Also, `TGSI_SEMANTIC_PCOORD` becomes available, which labels a fragment shader input that will always be replaced with sprite coordinates.

- `PIPE_CAP_PREFER_BLIT_BASED_TEXTURE_TRANSFER`: Whether it is preferable to use a blit to implement a texture transfer which needs format conversions and swizzling in state trackers. Generally, all hardware drivers with dedicated memory should return 1 and all software rasterizers should return 0.
- `PIPE_CAP_QUERY_PIPELINE_STATISTICS`: Whether `PIPE_QUERY_PIPELINE_STATISTICS` is supported.
- `PIPE_CAP_TEXTURE_BORDER_COLOR QUIRK`: Bitmask indicating whether special considerations have to be given to the interaction between the border color in the sampler object and the sampler view used with it. If `PIPE_QUIRK_TEXTURE_BORDER_COLOR_SWIZZLE_R600` is set, the border color may be affected in undefined ways for any kind of permutational swizzle (any swizzle `XYZW` where `X/Y/Z/W` are not `ZERO`, `ONE`, or `R/G/B/A` respectively) in the sampler view. If `PIPE_QUIRK_TEXTURE_BORDER_COLOR_SWIZZLE_NV50` is set, the border color state should be swizzled manually according to the swizzle in the sampler view it is intended to be used with, or herein undefined results may occur for permutational swizzles.
- `PIPE_CAP_MAX_TEXTURE_BUFFER_SIZE`: The maximum accessible size with a buffer sampler view, in texels.
- `PIPE_CAP_MAX_VIEWPORTS`: The maximum number of viewports (and scissors since they are linked) a driver can support. Returning 0 is equivalent to returning 1 because every driver has to support at least a single viewport/scissor combination.
- `PIPE_CAP_ENDIANNES`:: The endianness of the device. Either `PIPE_ENDIAN_BIG` or `PIPE_ENDIAN_LITTLE`.
- `PIPE_CAP_MIXED_FRAMEBUFFER_SIZES`: Whether it is allowed to have different sizes for fb color/zs attachments. This controls whether `ARB_framebuffer_object` is provided.
- `PIPE_CAP_TGSI_VS_LAYER_VIEWPORT`: Whether `TGSI_SEMANTIC_LAYER` and `TGSI_SEMANTIC_VIEWPORT_INDEX` are supported as vertex shader outputs. Note that the viewport will only be used if multiple viewports are exposed.
- `PIPE_CAP_MAX_GEOMETRY_OUTPUT_VERTICES`: The maximum number of vertices output by a single invocation of a geometry shader.
- `PIPE_CAP_MAX_GEOMETRY_TOTAL_OUTPUT_COMPONENTS`: The maximum number of vertex components output by a single invocation of a geometry shader. This is the product of the number of attribute components per vertex and the number of output vertices.
- `PIPE_CAP_MAX_TEXTURE_GATHER_COMPONENTS`: Max number of components in format that texture gather can operate on. 1 == `RED`, `ALPHA` etc, 4 == All formats.
- `PIPE_CAP_TEXTURE_GATHER_SM5`: Whether the texture gather hardware implements the SM5 features, component selection, shadow comparison, and run-time offsets.
- `PIPE_CAP_BUFFER_MAP_PERSISTENT_COHERENT`: Whether `PIPE_TRANSFER_PERSISTENT` and `PIPE_TRANSFER_COHERENT` are supported for buffers.
- `PIPE_CAP_TEXTURE_QUERY_LOD`: Whether the `LODQ` instruction is supported.
- `PIPE_CAP_MIN_TEXTURE_GATHER_OFFSET`: The minimum offset that can be used in conjunction with a texture gather opcode.
- `PIPE_CAP_MAX_TEXTURE_GATHER_OFFSET`: The maximum offset that can be used in conjunction with a texture gather opcode.
- `PIPE_CAP_SAMPLE_SHADING`: Whether there is support for per-sample shading. The `context->set_min_samples` function will be expected to be implemented.
- `PIPE_CAP_TEXTURE_GATHER_OFFSETS`: Whether the `TG4` instruction can accept 4 offsets.

- `PIPE_CAP_TGSI_VS_WINDOW_SPACE_POSITION`: Whether `TGSI_PROPERTY_VS_WINDOW_SPACE_POSITION` is supported, which disables clipping and viewport transformation.
- `PIPE_CAP_MAX_VERTEX_STREAMS`: The maximum number of vertex streams supported by the geometry shader. If stream-out is supported, this should be at least 1. If stream-out is not supported, this should be 0.
- `PIPE_CAP_DRAW_INDIRECT`: Whether the driver supports taking draw arguments { `count`, `instance_count`, `start`, `index_bias` } from a `PIPE_BUFFER` resource. See `pipe_draw_info`.
- `PIPE_CAP_MULTI_DRAW_INDIRECT`: Whether the driver supports `pipe_draw_info::indirect_stride` and `::indirect_count`
- `PIPE_CAP_MULTI_DRAW_INDIRECT_PARAMS`: Whether the driver supports taking the number of indirect draws from a separate parameter buffer, see `pipe_draw_indirect_info::indirect_draw_count`.
- `PIPE_CAP_TGSI_FS_FINE_DERIVATIVE`: Whether the fragment shader supports the FINE versions of DDX/DDY.
- `PIPE_CAP_VENDOR_ID`: The vendor ID of the underlying hardware. If it's not available one should return `0xFFFFFFFF`.
- `PIPE_CAP_DEVICE_ID`: The device ID (PCI ID) of the underlying hardware. `0xFFFFFFFF` if not available.
- `PIPE_CAP_ACCELERATED`: Whether the renderer is hardware accelerated.
- `PIPE_CAP_VIDEO_MEMORY`: The amount of video memory in megabytes.
- `PIPE_CAP_UMA`: If the device has a unified memory architecture or on-card memory and GART.
- `PIPE_CAP_CONDITIONAL_RENDER_INVERTED`: Whether the driver supports inverted condition for conditional rendering.
- `PIPE_CAP_MAX_VERTEX_ATTRIB_STRIDE`: The maximum supported vertex stride.
- `PIPE_CAP_SAMPLER_VIEW_TARGET`: Whether the sampler view's target can be different than the underlying resource's, as permitted by `ARB_texture_view`. For example a 2d array texture may be reinterpreted as a cube (array) texture and vice-versa.
- `PIPE_CAP_CLIP_HALFZ`: Whether the driver supports the `pipe_rasterizer_state::clip_halfz` being set to true. This is required for enabling `ARB_clip_control`.
- `PIPE_CAP_VERTEXID_NOBASE`: If true, the driver only supports `TGSI_SEMANTIC_VERTEXID_NOBASE` (and not `TGSI_SEMANTIC_VERTEXID`). This means state trackers for APIs whose vertexIDs are offset by basevertex (such as GL) will need to lower `TGSI_SEMANTIC_VERTEXID` to `TGSI_SEMANTIC_VERTEXID_NOBASE` and `TGSI_SEMANTIC_BASEVERTEX`, so drivers setting this must handle both these semantics. Only relevant if geometry shaders are supported. (`BASEVERTEX` could be exposed separately too via `PIPE_CAP_DRAW_PARAMETERS`).
- `PIPE_CAP_POLYGON_OFFSET_CLAMP`: If true, the driver implements support for `pipe_rasterizer_state::offset_clamp`.
- `PIPE_CAP_MULTISAMPLE_Z_RESOLVE`: Whether the driver supports blitting a multisampled depth buffer into a single-sampled texture (or depth buffer). Only the first sampled should be copied.
- `PIPE_CAP_RESOURCE_FROM_USER_MEMORY`: Whether the driver can create a `pipe_resource` where an already-existing piece of (malloc'd) user memory is used as its backing storage. In other words, whether the driver can map existing user memory into the device address space for direct device access. The create function is `pipe_screen::resource_from_user_memory`. The address and size must be page-aligned.
- `PIPE_CAP_DEVICE_RESET_STATUS_QUERY`: Whether `pipe_context::get_device_reset_status` is implemented.

- `PIPE_CAP_MAX_SHADER_PATCH_VARYINGS`: How many per-patch outputs and inputs are supported between tessellation control and tessellation evaluation shaders, not counting in `TESSINNER` and `TESSOUTER`. The minimum allowed value for OpenGL is 30.
- `PIPE_CAP_TEXTURE_FLOAT_LINEAR`: Whether the linear minification and magnification filters are supported with single-precision floating-point textures.
- `PIPE_CAP_TEXTURE_HALF_FLOAT_LINEAR`: Whether the linear minification and magnification filters are supported with half-precision floating-point textures.
- `PIPE_CAP_DEPTH_BOUNDS_TEST`: Whether `bounds_test`, `bounds_min`, and `bounds_max` states of `pipe_depth_stencil_alpha_state` behave according to the `GL_EXT_depth_bounds_test` specification.
- `PIPE_CAP_TGSI_TXQS`: Whether the `TXQS` opcode is supported
- `PIPE_CAP_FORCE_PERSAMPLE_INTERP`: If the driver can force per-sample interpolation for all fragment shader inputs if `pipe_rasterizer_state::force_persample_interp` is set. This is only used by GL3-level sample shading (`ARB_sample_shading`). GL4-level sample shading (`ARB_gpu_shader5`) doesn't use this. While GL3 hardware has a state for it, GL4 hardware will likely need to emulate it with a shader variant, or by selecting the interpolation weights with a conditional assignment in the shader.
- `PIPE_CAP_SHAREABLE_SHADERS`: Whether shader CSOs can be used by any `pipe_context`.
- `PIPE_CAP_COPY_BETWEEN_COMPRESSED_AND_PLAIN_FORMATS`: Whether copying between compressed and plain formats is supported where a compressed block is copied to/from a plain pixel of the same size.
- `PIPE_CAP_CLEAR_TEXTURE`: Whether `clear_texture` will be available in contexts.
- `PIPE_CAP_DRAW_PARAMETERS`: Whether `TGSI_SEMANTIC_BASEVERTEX`, `TGSI_SEMANTIC_BASEINSTANCE`, and `TGSI_SEMANTIC_DRAWID` are supported in vertex shaders.
- `PIPE_CAP_TGSI_PACK_HALF_FLOAT`: Whether the `UP2H` and `PK2H` TGSI opcodes are supported.
- `PIPE_CAP_TGSI_FS_POSITION_IS_SYSVAL`: If state trackers should use a system value for the `POSITION` fragment shader input.
- `PIPE_CAP_TGSI_FS_FACE_IS_INTEGER_SYSVAL`: If state trackers should use a system value for the `FACE` fragment shader input. Also, the `FACE` system value is integer, not float.
- `PIPE_CAP_SHADER_BUFFER_OFFSET_ALIGNMENT`: Describes the required alignment for `pipe_shader_buffer::buffer_offset`, in bytes. Maximum value allowed is 256 (for GL conformance). 0 is only allowed if shader buffers are not supported.
- `PIPE_CAP_INVALIDATE_BUFFER`: Whether the use of `invalidate_resource` for buffers is supported.
- `PIPE_CAP_GENERATE_MIPMAP`: Indicates whether `pipe_context::generate_mipmap` is supported.
- `PIPE_CAP_STRING_MARKER`: Whether `pipe->emit_string_marker()` is supported.
- `PIPE_CAP_SURFACE_REINTERPRET_BLOCKS`: Indicates whether `pipe_context::create_surface` supports reinterpreting a texture as a surface of a format with different block width/height (but same block size in bits). For example, a compressed texture image can be interpreted as a non-compressed surface whose texels are the same number of bits as the compressed blocks, and vice versa. The width and height of the surface is adjusted appropriately.
- `PIPE_CAP_QUERY_BUFFER_OBJECT`: Driver supports `context::get_query_result_resource` callback.
- `PIPE_CAP_PCI_GROUP`: Return the PCI segment group number.
- `PIPE_CAP_PCI_BUS`: Return the PCI bus number.
- `PIPE_CAP_PCI_DEVICE`: Return the PCI device number.

- `PIPE_CAP_PCI_FUNCTION`: Return the PCI function number.
- `PIPE_CAP_FRAMEBUFFER_NO_ATTACHMENT`: If non-zero, rendering to framebuffers with no surface attachments is supported. The `context->is_format_supported` function will be expected to be implemented with `PIPE_FORMAT_NONE` yielding the MSAA modes the hardware supports. N.B., The maximum number of layers supported for rasterizing a primitive on a layer is obtained from `PIPE_CAP_MAX_TEXTURE_ARRAY_LAYERS` even though it can be larger than the number of layers supported by either rendering or textures.
- `PIPE_CAP_ROBUST_BUFFER_ACCESS_BEHAVIOR`: Implementation uses bounds checking on resource accesses by shader if the context is created with `PIPE_CONTEXT_ROBUST_BUFFER_ACCESS`. See the `ARB_robust_buffer_access_behavior` extension for information on the required behavior for out of bounds accesses and accesses to unbound resources.
- `PIPE_CAP_CULL_DISTANCE`: Whether the driver supports the `arb_cull_distance` extension and thus implements proper support for culling planes.
- `PIPE_CAP_PRIMITIVE_RESTART_FOR_PATCHES`: Whether primitive restart is supported for patch primitives.
- `PIPE_CAP_TGSI_VOTE`: Whether the `VOTE_*` ops can be used in shaders.
- `PIPE_CAP_MAX_WINDOW_RECTANGLES`: The maximum number of window rectangles supported in `set_window_rectangles`.
- `PIPE_CAP_POLYGON_OFFSET_UNITS_UNSCALED`: If true, the driver implements support for `pipe_rasterizer_state::offset_units_unscaled`.
- `PIPE_CAP_VIEWPORT_SUBPIXEL_BITS`: Number of bits of subpixel precision for floating point viewport bounds.
- `PIPE_CAP_MIXED_COLOR_DEPTH_BITS`: Whether there is non-fallback support for color/depth format combinations that use a different number of bits. For the purpose of this cap, Z24 is treated as 32-bit. If set to off, that means that a B5G6R5 + Z24 or RGBA8 + Z16 combination will require a driver fallback, and should not be advertised in the GLX/EGL config list.
- `PIPE_CAP_TGSI_ARRAY_COMPONENTS`: If true, the driver interprets the UsageMask of input and output declarations and allows declaring arrays in overlapping ranges. The components must be a contiguous range, e.g. a UsageMask of xy or yzw is allowed, but xz or yw isn't. Declarations with overlapping locations must have matching semantic names and indices, and equal interpolation qualifiers. Components may overlap, notably when the gaps in an array of `dvec3` are filled in.
- `PIPE_CAP_STREAM_OUTPUT_INTERLEAVE_BUFFERS`: Whether interleaved stream output mode is able to interleave across buffers. This is required for `ARB_transform_feedback3`.
- `PIPE_CAP_TGSI_CAN_READ_OUTPUTS`: Whether every TGSI shader stage can read from the output file.
- `PIPE_CAP_GLSL_OPTIMIZE_CONSERVATIVELY`: Tell the GLSL compiler to use the minimum amount of optimizations just to be able to do all the linking and lowering.
- `PIPE_CAP_TGSI_FS_FBFETCH`: Whether a fragment shader can use the `FBFETCH` opcode to retrieve the current value in the framebuffer.
- `PIPE_CAP_TGSI_MUL_ZERO_WINS`: Whether TGSI shaders support the `TGSI_PROPERTY_MUL_ZERO_WINS` shader property.
- `PIPE_CAP_DOUBLES`: Whether double precision floating-point operations are supported.
- `PIPE_CAP_INT64`: Whether 64-bit integer operations are supported.
- `PIPE_CAP_INT64_DIVMOD`: Whether 64-bit integer division/modulo operations are supported.
- `PIPE_CAP_TGSI_TEX_TXF_LZ`: Whether `TEX_LZ` and `TXF_LZ` opcodes are supported.

- `PIPE_CAP_TGSI_CLOCK`: Whether the `CLOCK` opcode is supported.
- `PIPE_CAP_POLYGON_MODE_FILL_RECTANGLE`: Whether the `PIPE_POLYGON_MODE_FILL_RECTANGLE` mode is supported for `pipe_rasterizer_state::fill_front` and `pipe_rasterizer_state::fill_back`.
- `PIPE_CAP_SPARSE_BUFFER_PAGE_SIZE`: The page size of sparse buffers in bytes, or 0 if sparse buffers are not supported. The page size must be at most 64KB.
- `PIPE_CAP_TGSI_BALLOT`: Whether the `BALLOT` and `READ_*` opcodes as well as the `SUBGROUP_*` semantics are supported.
- `PIPE_CAP_TGSI_TES_LAYER_VIEWPORT`: Whether `TGSI_SEMANTIC_LAYER` and `TGSI_SEMANTIC_VIEWPORT_INDEX` are supported as tessellation evaluation shader outputs.
- `PIPE_CAP_CAN_BIND_CONST_BUFFER_AS_VERTEX`: Whether a buffer with just `PIPE_BIND_CONSTANT_BUFFER` can be legally passed to `set_vertex_buffers`.
- `PIPE_CAP_ALLOW_MAPPED_BUFFERS_DURING_EXECUTION`: As the name says.
- `PIPE_CAP_POST_DEPTH_COVERAGE`: whether `TGSI_PROPERTY_FS_POST_DEPTH_COVERAGE` is supported.
- `PIPE_CAP_BINDLESS_TEXTURE`: Whether bindless texture operations are supported.
- `PIPE_CAP_NIR_SAMPLERS_AS_DEREF`: Whether `NIR tex` instructions should reference texture and sampler as `NIR derefs` instead of by indices.
- `PIPE_CAP_QUERY_SO_OVERFLOW`: Whether the `PIPE_QUERY_SO_OVERFLOW_PREDICATE` and `PIPE_QUERY_SO_OVERFLOW_ANY_PREDICATE` query types are supported. Note that for a driver that does not support multiple output streams (i.e., `PIPE_CAP_MAX_VERTEX_STREAMS` is 1), both query types are identical.
- `PIPE_CAP_MEMOBJ`: Whether operations on memory objects are supported.
- `PIPE_CAP_LOAD_CONSTBUF`: True if the driver supports `TGSI_OPCODE_LOAD` use with constant buffers.
- `PIPE_CAP_TGSI_ANY_REG_AS_ADDRESS`: Any `TGSI register` can be used as an address for indirect register indexing.

PIPE_CAPF_*

The floating-point capabilities are:

- `PIPE_CAPF_MAX_LINE_WIDTH`: The maximum width of a regular line.
- `PIPE_CAPF_MAX_LINE_WIDTH_AA`: The maximum width of a smoothed line.
- `PIPE_CAPF_MAX_POINT_WIDTH`: The maximum width and height of a point.
- `PIPE_CAPF_MAX_POINT_WIDTH_AA`: The maximum width and height of a smoothed point.
- `PIPE_CAPF_MAX_TEXTURE_ANISOTROPY`: The maximum level of anisotropy that can be applied to anisotropically filtered textures.
- `PIPE_CAPF_MAX_TEXTURE_LOD_BIAS`: The maximum *LOD* bias that may be applied to filtered textures.
- `PIPE_CAPF_GUARD_BAND_LEFT`, `PIPE_CAPF_GUARD_BAND_TOP`, `PIPE_CAPF_GUARD_BAND_RIGHT`, `PIPE_CAPF_GUARD_BAND_BOTTOM`: **TODO**

PIPE_SHADER_CAP_*

These are per-shader-stage capability queries. Different shader stages may support different features.

- `PIPE_SHADER_CAP_MAX_INSTRUCTIONS`: The maximum number of instructions.
- `PIPE_SHADER_CAP_MAX_ALU_INSTRUCTIONS`: The maximum number of arithmetic instructions.
- `PIPE_SHADER_CAP_MAX_TEX_INSTRUCTIONS`: The maximum number of texture instructions.
- `PIPE_SHADER_CAP_MAX_TEX_INDIRECTIONS`: The maximum number of texture indirections.
- `PIPE_SHADER_CAP_MAX_CONTROL_FLOW_DEPTH`: The maximum nested control flow depth.
- `PIPE_SHADER_CAP_MAX_INPUTS`: The maximum number of input registers.
- `PIPE_SHADER_CAP_MAX_OUTPUTS`: The maximum number of output registers. This is valid for all shaders except the fragment shader.
- `PIPE_SHADER_CAP_MAX_CONST_BUFFER_SIZE`: The maximum size per constant buffer in bytes.
- `PIPE_SHADER_CAP_MAX_CONST_BUFFERS`: Maximum number of constant buffers that can be bound to any shader stage using `set_constant_buffer`. If 0 or 1, the pipe will only permit binding one constant buffer per shader.

If a value greater than 0 is returned, the driver can have multiple constant buffers bound to shader stages. The `CONST` register file is accessed with two-dimensional indices, like in the example below.

```
DCL CONST[0][0..7] # declare first 8 vectors of constbuf 0
DCL CONST[3][0] # declare first vector of constbuf 3
MOV OUT[0], CONST[0][3] # copy vector 3 of constbuf 0
```

- `PIPE_SHADER_CAP_MAX_TEMPS`: The maximum number of temporary registers.
- `PIPE_SHADER_CAP_TGSI_CONT_SUPPORTED`: Whether the continue opcode is supported.
- `PIPE_SHADER_CAP_INDIRECT_INPUT_ADDR`: Whether indirect addressing of the input file is supported.
- `PIPE_SHADER_CAP_INDIRECT_OUTPUT_ADDR`: Whether indirect addressing of the output file is supported.
- `PIPE_SHADER_CAP_INDIRECT_TEMP_ADDR`: Whether indirect addressing of the temporary file is supported.
- `PIPE_SHADER_CAP_INDIRECT_CONST_ADDR`: Whether indirect addressing of the constant file is supported.
- `PIPE_SHADER_CAP_SUBROUTINES`: Whether subroutines are supported, i.e. `BGNSUB`, `ENDSUB`, `CAL`, and `RET`, including `RET` in the main block.
- `PIPE_SHADER_CAP_INTEGERS`: Whether integer opcodes are supported. If unsupported, only float opcodes are supported.
- `PIPE_SHADER_CAP_INT64_ATOMICS`: Whether int64 atomic opcodes are supported. The device needs to support `add`, `sub`, `swap`, `cmpswap`, `and`, `or`, `xor`, `min`, and `max`.
- **`PIPE_SHADER_CAP_FP16`: Whether half precision floating-point opcodes are supported.** If unsupported, half precision ops need to be lowered to full precision.
- `PIPE_SHADER_CAP_MAX_TEXTURE_SAMPLERS`: The maximum number of texture samplers.
- `PIPE_SHADER_CAP_PREFERRED_IR`: Preferred representation of the program. It should be one of the `pipe_shader_ir` enum values.
- `PIPE_SHADER_CAP_MAX_SAMPLER_VIEWS`: The maximum number of texture sampler views. Must not be lower than `PIPE_SHADER_CAP_MAX_TEXTURE_SAMPLERS`.

- `PIPE_SHADER_CAP_TGSI_DROUND_SUPPORTED`: Whether double precision rounding is supported. If it is, `DTRUNC/DCEIL/DFLR/DROUND` opcodes may be used.
- `PIPE_SHADER_CAP_TGSI_DFRACEXP_DLDEXP_SUPPORTED`: Whether `DFRACEXP` and `DLDEXP` are supported.
- `PIPE_SHADER_CAP_TGSI_LDEXP_SUPPORTED`: Whether `LDEXP` is supported.
- `PIPE_SHADER_CAP_TGSI_FMA_SUPPORTED`: Whether `FMA` and `DFMA` (doubles only) are supported.
- `PIPE_SHADER_CAP_TGSI_ANY_INOUT_DECL_RANGE`: Whether the driver doesn't ignore `tgsi_declaration_range::Last` for shader inputs and outputs.
- `PIPE_SHADER_CAP_MAX_UNROLL_ITERATIONS_HINT`: This is the maximum number of iterations that loops are allowed to have to be unrolled. It is only a hint to state trackers. Whether any loops will be unrolled is not guaranteed.
- `PIPE_SHADER_CAP_MAX_SHADER_BUFFERS`: Maximum number of memory buffers (also used to implement atomic counters). Having this be non-0 also implies support for the `LOAD`, `STORE`, and `ATOM*` `TGSI` opcodes.
- `PIPE_SHADER_CAP_SUPPORTED_IRS`: Supported representations of the program. It should be a mask of `pipe_shader_ir` bits.
- `PIPE_SHADER_CAP_MAX_SHADER_IMAGES`: Maximum number of image units.
- `PIPE_SHADER_CAP_LOWER_IF_THRESHOLD`: `IF` and `ELSE` branches with a lower cost than this value should be lowered by the state tracker for better performance. This is a tunable for the `GLSL` compiler and the behavior is specific to the compiler.
- `PIPE_SHADER_CAP_TGSI_SKIP_MERGE_REGISTERS`: Whether the merge registers `TGSI` pass is skipped. This might reduce code size and register pressure if the underlying driver has a real backend compiler.

PIPE_COMPUTE_CAP_*

Compute-specific capabilities. They can be queried using `pipe_screen::get_compute_param`.

- `PIPE_COMPUTE_CAP_IR_TARGET`: A description of the target of the form `processor-arch-manufacturer-os` that will be passed on to the compiler. This CAP is only relevant for drivers that specify `PIPE_SHADER_IR_LLVM` or `PIPE_SHADER_IR_NATIVE` for their preferred IR. Value type: null-terminated string. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_GRID_DIMENSION`: Number of supported dimensions for grid and block coordinates. Value type: `uint64_t`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_GRID_SIZE`: Maximum grid size in block units. Value type: `uint64_t []`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_BLOCK_SIZE`: Maximum block size in thread units. Value type: `uint64_t []`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_THREADS_PER_BLOCK`: Maximum number of threads that a single block can contain. Value type: `uint64_t`. Shader IR type dependent. This may be less than the product of the components of `MAX_BLOCK_SIZE` and is usually limited by the number of threads that can be resident simultaneously on a compute unit.
- `PIPE_COMPUTE_CAP_MAX_GLOBAL_SIZE`: Maximum size of the `GLOBAL` resource. Value type: `uint64_t`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_LOCAL_SIZE`: Maximum size of the `LOCAL` resource. Value type: `uint64_t`. Shader IR type dependent.

- `PIPE_COMPUTE_CAP_MAX_PRIVATE_SIZE`: Maximum size of the PRIVATE resource. Value type: `uint64_t`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_INPUT_SIZE`: Maximum size of the INPUT resource. Value type: `uint64_t`. Shader IR type dependent.
- `PIPE_COMPUTE_CAP_MAX_MEM_ALLOC_SIZE`: Maximum size of a memory object allocation in bytes. Value type: `uint64_t`.
- `PIPE_COMPUTE_CAP_MAX_CLOCK_FREQUENCY`: Maximum frequency of the GPU clock in MHz. Value type: `uint32_t`
- `PIPE_COMPUTE_CAP_MAX_COMPUTE_UNITS`: Maximum number of compute units Value type: `uint32_t`
- `PIPE_COMPUTE_CAP_IMAGES_SUPPORTED`: Whether images are supported non-zero means yes, zero means no. Value type: `uint32_t`
- `PIPE_COMPUTE_CAP_SUBGROUP_SIZE`: The size of a basic execution unit in threads. Also known as wavefront size, warp size or SIMD width.
- `PIPE_COMPUTE_CAP_ADDRESS_BITS`: The default compute device address space size specified as an unsigned integer value in bits.
- `PIPE_COMPUTE_CAP_MAX_VARIABLE_THREADS_PER_BLOCK`: Maximum variable number of threads that a single block can contain. This is similar to `PIPE_COMPUTE_CAP_MAX_THREADS_PER_BLOCK`, except that the variable size is not known a compile-time but at dispatch-time.

PIPE_BIND_*

These flags indicate how a resource will be used and are specified at resource creation time. Resources may be used in different roles during their lifecycle. Bind flags are cumulative and may be combined to create a resource which can be used for multiple things. Depending on the pipe driver's memory management and these bind flags, resources might be created and handled quite differently.

- `PIPE_BIND_RENDER_TARGET`: A color buffer or pixel buffer which will be rendered to. Any surface/resource attached to `pipe_framebuffer_state::cbufs` must have this flag set.
- `PIPE_BIND_DEPTH_STENCIL`: A depth (Z) buffer and/or stencil buffer. Any depth/stencil surface/resource attached to `pipe_framebuffer_state::zbuf` must have this flag set.
- `PIPE_BIND_BLENDABLE`: Used in conjunction with `PIPE_BIND_RENDER_TARGET` to query whether a device supports blending for a given format. If this flag is set, surface creation may fail if blending is not supported for the specified format. If it is not set, a driver may choose to ignore blending on surfaces with formats that would require emulation.
- `PIPE_BIND_DISPLAY_TARGET`: A surface that can be presented to screen. Arguments to `pipe_screen::flush_front_buffer` must have this flag set.
- `PIPE_BIND_SAMPLER_VIEW`: A texture that may be sampled from in a fragment or vertex shader.
- `PIPE_BIND_VERTEX_BUFFER`: A vertex buffer.
- `PIPE_BIND_INDEX_BUFFER`: An vertex index/element buffer.
- `PIPE_BIND_CONSTANT_BUFFER`: A buffer of shader constants.
- `PIPE_BIND_STREAM_OUTPUT`: A stream output buffer.
- `PIPE_BIND_CUSTOM`:
- `PIPE_BIND_SCANOUT`: A front color buffer or scanout buffer.

- `PIPE_BIND_SHARED`: A sharable buffer that can be given to another process.
- `PIPE_BIND_GLOBAL`: A buffer that can be mapped into the global address space of a compute program.
- `PIPE_BIND_SHADER_BUFFER`: A buffer without a format that can be bound to a shader and can be used with load, store, and atomic instructions.
- `PIPE_BIND_SHADER_IMAGE`: A buffer or texture with a format that can be bound to a shader and can be used with load, store, and atomic instructions.
- `PIPE_BIND_COMPUTE_RESOURCE`: A buffer or texture that can be bound to the compute program as a shader resource.
- `PIPE_BIND_COMMAND_ARGS_BUFFER`: A buffer that may be sourced by the GPU command processor. It can contain, for example, the arguments to indirect draw calls.

PIPE_USAGE_*

The `PIPE_USAGE` enums are hints about the expected usage pattern of a resource. Note that drivers must always support read and write CPU access at any time no matter which hint they got.

- `PIPE_USAGE_DEFAULT`: Optimized for fast GPU access.
- `PIPE_USAGE_IMMUTABLE`: Optimized for fast GPU access and the resource is not expected to be mapped or changed (even by the GPU) after the first upload.
- `PIPE_USAGE_DYNAMIC`: Expect frequent write-only CPU access. What is uploaded is expected to be used at least several times by the GPU.
- `PIPE_USAGE_STREAM`: Expect frequent write-only CPU access. What is uploaded is expected to be used only once by the GPU.
- `PIPE_USAGE_STAGING`: Optimized for fast CPU access.

Methods

XXX to-do

`get_name`

Returns an identifying name for the screen.

The returned string should remain valid and immutable for the lifetime of `pipe_screen`.

`get_vendor`

Returns the screen vendor.

The returned string should remain valid and immutable for the lifetime of `pipe_screen`.

`get_device_vendor`

Returns the actual vendor of the device driving the screen (as opposed to the driver vendor).

The returned string should remain valid and immutable for the lifetime of `pipe_screen`.

get_param

Get an integer/boolean screen parameter.

param is one of the *PIPE_CAP_** names.

get_paramf

Get a floating-point screen parameter.

param is one of the *PIPE_CAPF_** names.

context_create

Create a pipe_context.

priv is private data of the caller, which may be put to various unspecified uses, typically to do with implementing swapbuffers and/or front-buffer rendering.

is_format_supported

Determine if a resource in the given format can be used in a specific manner.

format the resource format

target one of the PIPE_TEXTURE_x flags

sample_count the number of samples. 0 and 1 mean no multisampling, the maximum allowed legal value is 32.

bindings is a bitmask of *PIPE_BIND_** flags.

Returns TRUE if all usages can be satisfied.

can_create_resource

Check if a resource can actually be created (but don't actually allocate any memory). This is used to implement OpenGL's proxy textures. Typically, a driver will simply check if the total size of the given resource is less than some limit.

For PIPE_TEXTURE_CUBE, the pipe_resource::array_size field should be 6.

resource_create

Create a new resource from a template. The following fields of the pipe_resource must be specified in the template:

target one of the pipe_texture_target enums. Note that PIPE_BUFFER and PIPE_TEXTURE_X are not really fundamentally different. Modern APIs allow using buffers as shader resources.

format one of the pipe_format enums.

width0 the width of the base mip level of the texture or size of the buffer.

height0 the height of the base mip level of the texture (1 for 1D or 1D array textures).

depth0 the depth of the base mip level of the texture (1 for everything else).

array_size the array size for 1D and 2D array textures. For cube maps this must be 6, for other textures 1.

last_level the last mip map level present.

nr_samples the nr of msaa samples. 0 (or 1) specifies a resource which isn't multisampled.

usage one of the *PIPE_USAGE_** flags.

bind bitmask of the *PIPE_BIND_** flags.

flags bitmask of PIPE_RESOURCE_FLAG flags.

resource_changed

Mark a resource as changed so derived internal resources will be recreated on next use.

When importing external images that can't be directly used as texture sampler source, internal copies may have to be created that the hardware can sample from. When those resources are reimported, the image data may have changed, and the previously derived internal resources must be invalidated to avoid sampling from old copies.

resource_destroy

Destroy a resource. A resource is destroyed if it has no more references.

get_timestamp

Query a timestamp in nanoseconds. The returned value should match PIPE_QUERY_TIMESTAMP. This function returns immediately and doesn't wait for rendering to complete (which cannot be achieved with queries).

get_driver_query_info

Return a driver-specific query. If the **info** parameter is NULL, the number of available queries is returned. Otherwise, the driver query at the specified **index** is returned in **info**. The function returns non-zero on success. The driver-specific query is described with the `pipe_driver_query_info` structure.

get_driver_query_group_info

Return a driver-specific query group. If the **info** parameter is NULL, the number of available groups is returned. Otherwise, the driver query group at the specified **index** is returned in **info**. The function returns non-zero on success. The driver-specific query group is described with the `pipe_driver_query_group_info` structure.

get_disk_shader_cache

Returns a pointer to a driver-specific on-disk shader cache. If the driver failed to create the cache or does not support an on-disk shader cache NULL is returned. The callback itself may also be NULL if the driver doesn't support an on-disk shader cache.

Thread safety

Screen methods are required to be thread safe. While gallium rendering contexts are not required to be thread safe, it is required to be safe to use different contexts created with the same screen in different threads without locks. It is also required to be safe using screen methods in a thread, while using one of its contexts in another (without locks).

Resources and derived objects

Resources represent objects that hold data: textures and buffers.

They are mostly modelled after the resources in Direct3D 10/11, but with a different transfer/update mechanism, and more features for OpenGL support.

Resources can be used in several ways, and it is required to specify all planned uses through an appropriate set of bind flags.

TODO: write much more on resources

Transfers

Transfers are the mechanism used to access resources with the CPU.

OpenGL: OpenGL supports mapping buffers and has inline transfer functions for both buffers and textures

D3D11: D3D11 lacks transfers, but has special resource types that are mappable to the CPU address space

TODO: write much more on transfers

Resource targets

Resource targets determine the type of a resource.

Note that drivers may not actually have the restrictions listed regarding coordinate normalization and wrap modes, and in fact efficient OpenCL support will probably require drivers that don't have any of them, which will probably be advertised with an appropriate cap.

TODO: document all targets. Note that both 3D and cube have restrictions that depend on the hardware generation.

PIPE_BUFFER

Buffer resource: can be used as a vertex, index, constant buffer (appropriate bind flags must be requested).

Buffers do not really have a format, it's just bytes, but they are required to have their type set to a R8 format (without a specific "just byte" format, R8_UINT would probably make the most sense, but for historic reasons R8_UNORM is ok too). (This is just to make some shared buffer/texture code easier so format size can be queried.) width0 serves as size, most other resource properties don't apply but must be set appropriately (depth0/height0/array_size must be 1, last_level 0).

They can be bound to stream output if supported. TODO: what about the restrictions lifted by the several later GL transform feedback extensions? How does one advertise that in Gallium?

They can be also be bound to a shader stage (for sampling) as usual by creating an appropriate sampler view, if the driver supports PIPE_CAP_TEXTURE_BUFFER_OBJECTS. This supports larger width than a 1d texture would (TODO limit currently unspecified, minimum must be at least 65536). Only the "direct fetch" sample opcodes are supported (TGSI_OPCODE_TXF, TGSI_OPCODE_SAMPLE_I) so the sampler state (coord wrapping etc.) is mostly ignored (with SAMPLE_I there's no sampler state at all).

They can be also be bound to the framebuffer (only as color render target, not depth buffer, also there cannot be a depth buffer bound at the same time) as usual by creating an appropriate view (this is not usable in OpenGL). TODO there's no CAP bit currently for this, there's also unspecified size etc. limits TODO: is there any chance of supporting GL pixel buffer object acceleration with this?

OpenGL: vertex buffers in GL 1.5 or GL_ARB_vertex_buffer_object

- Binding to stream out requires GL 3.0 or GL_NV_transform_feedback
- Binding as constant buffers requires GL 3.1 or GL_ARB_uniform_buffer_object
- Binding to a sampling stage requires GL 3.1 or GL_ARB_texture_buffer_object

D3D11: buffer resources - Binding to a render target requires D3D_FEATURE_LEVEL_10_0

PIPE_TEXTURE_1D / PIPE_TEXTURE_1D_ARRAY

1D surface accessed with normalized coordinates. 1D array textures are supported depending on PIPE_CAP_MAX_TEXTURE_ARRAY_LAYERS.

- **If PIPE_CAP_NPOT_TEXTURES is not supported**, width must be a power of two
- height0 must be 1
- depth0 must be 1
- array_size must be 1 for PIPE_TEXTURE_1D
- Mipmaps can be used
- Must use normalized coordinates

OpenGL: GL_TEXTURE_1D in GL 1.0

- PIPE_CAP_NPOT_TEXTURES is equivalent to GL 2.0 or GL_ARB_texture_non_power_of_two

D3D11: 1D textures in D3D_FEATURE_LEVEL_10_0

PIPE_TEXTURE_RECT

2D surface with OpenGL GL_TEXTURE_RECTANGLE semantics.

- depth0 must be 1

- array_size must be 1
- last_level must be 0
- Must use unnormalized coordinates
- Must use a clamp wrap mode

OpenGL: GL_TEXTURE_RECTANGLE in GL 3.1 or GL_ARB_texture_rectangle or GL_NV_texture_rectangle

OpenCL: can create OpenCL images based on this, that can then be sampled arbitrarily

D3D11: not supported (only PIPE_TEXTURE_2D with normalized coordinates is supported)

PIPE_TEXTURE_2D / PIPE_TEXTURE_2D_ARRAY

2D surface accessed with normalized coordinates. 2D array textures are supported depending on PIPE_CAP_MAX_TEXTURE_ARRAY_LAYERS.

- **If PIPE_CAP_NPOT_TEXTURES is not supported**, width and height must be powers of two
- depth0 must be 1
- array_size must be 1 for PIPE_TEXTURE_2D
- Mipmaps can be used
- Must use normalized coordinates
- No special restrictions on wrap modes

OpenGL: GL_TEXTURE_2D in GL 1.0

- PIPE_CAP_NPOT_TEXTURES is equivalent to GL 2.0 or GL_ARB_texture_non_power_of_two

OpenCL: can create OpenCL images based on this, that can then be sampled arbitrarily

D3D11: 2D textures

- PIPE_CAP_NPOT_TEXTURES is equivalent to D3D_FEATURE_LEVEL_9_3

PIPE_TEXTURE_3D

3-dimensional array of texels. Mipmap dimensions are reduced in all 3 coordinates.

- **If PIPE_CAP_NPOT_TEXTURES is not supported**, width, height and depth must be powers of two
- array_size must be 1
- Must use normalized coordinates

OpenGL: GL_TEXTURE_3D in GL 1.2 or GL_EXT_texture3D

- PIPE_CAP_NPOT_TEXTURES is equivalent to GL 2.0 or GL_ARB_texture_non_power_of_two

D3D11: 3D textures

- PIPE_CAP_NPOT_TEXTURES is equivalent to D3D_FEATURE_LEVEL_10_0

PIPE_TEXTURE_CUBE / PIPE_TEXTURE_CUBE_ARRAY

Cube maps consist of 6 2D faces. The 6 surfaces form an imaginary cube, and sampling happens by mapping an input 3-vector to the point of the cube surface in that direction. Cube map arrays are supported depending on PIPE_CAP_CUBE_MAP_ARRAY.

Sampling may be optionally seamless if a driver supports it (PIPE_CAP_SEAMLESS_CUBE_MAP), resulting in filtering taking samples from multiple surfaces near to the edge.

- Width and height must be equal
- depth0 must be 1
- array_size must be a multiple of 6
- **If PIPE_CAP_NPOT_TEXTURES is not supported**, width and height must be powers of two
- Must use normalized coordinates

OpenGL: GL_TEXTURE_CUBE_MAP in GL 1.3 or EXT_texture_cube_map

- PIPE_CAP_NPOT_TEXTURES is equivalent to GL 2.0 or GL_ARB_texture_non_power_of_two
- Seamless cube maps require GL 3.2 or GL_ARB_seamless_cube_map or GL_AMD_seamless_cubemap_per_texture
- Cube map arrays require GL 4.0 or GL_ARB_texture_cube_map_array

D3D11: 2D array textures with the D3D11_RESOURCE_MISC_TEXTURECUBE flag

- PIPE_CAP_NPOT_TEXTURES is equivalent to D3D_FEATURE_LEVEL_10_0
- Cube map arrays require D3D_FEATURE_LEVEL_10_1

Surfaces

Surfaces are views of a resource that can be bound as a framebuffer to serve as the render target or depth buffer.

TODO: write much more on surfaces

OpenGL: FBOs are collections of surfaces in GL 3.0 or GL_ARB_framebuffer_object

D3D11: render target views and depth/stencil views

Sampler views

Sampler views are views of a resource that can be bound to a pipeline stage to be sampled from shaders.

TODO: write much more on sampler views

OpenGL: texture objects are actually sampler view and resource in a single unit

D3D11: shader resource views

Formats in gallium

Gallium format names mostly follow D3D10 conventions, with some extensions.

Format names like $X_n Y_n Z_n W_n$ have the X component in the lowest-address n bits and the W component in the highest-address n bits; for B8G8R8A8, byte 0 is blue and byte 3 is alpha. Note that platform endianness is not considered in this definition. In C:

```
struct x8y8z8w8 { uint8_t x, y, z, w; };
```

Format aliases like XYZWstrq are $(s+t+r+q)$ -bit integers in host endianness, with the X component in the s least-significant bits of the integer. In C:

```
uint32_t xyzw8888 = (x << 0) | (y << 8) | (z << 16) | (w << 24);
```

Format suffixes affect the interpretation of the channel:

- SINT: N bit signed integer $[-2^{(N-1)} \dots 2^{(N-1)} - 1]$
- SNORM: N bit signed integer normalized to $[-1 \dots 1]$
- SSCALED: N bit signed integer $[-2^{(N-1)} \dots 2^{(N-1)} - 1]$
- FIXED: Signed fixed point integer, $(N/2 - 1)$ bits of mantissa
- FLOAT: N bit IEEE754 float
- NORM: Normalized integers, signed or unsigned per channel
- UINT: N bit unsigned integer $[0 \dots 2^N - 1]$
- UNORM: N bit unsigned integer normalized to $[0 \dots 1]$
- USCALED: N bit unsigned integer $[0 \dots 2^N - 1]$

The difference between SINT and SSCALED is that the former are pure integers in shaders, while the latter are floats; likewise for UINT versus USCALED.

There are two exceptions for FLOAT. R9G9B9E5_FLOAT is nine bits each of red green and blue mantissa, with a shared five bit exponent. R11G11B10_FLOAT is five bits of exponent and five or six bits of mantissa for each color channel.

For the `NORM` suffix, the signedness of each channel is indicated with an `S` or `U` after the number of channel bits, as in `R5SG5SB6U_NORM`.

The `SRGB` suffix is like `UNORM` in range, but in the `sRGB` colorspace.

Compressed formats are named first by the compression format string (`DXT1`, `ETC1`, etc), followed by a format-specific subtype. Refer to the appropriate compression spec for details.

Formats used in video playback are named by their `FOURCC` code.

Format names with an embedded underscore are subsampled. `R8G8_B8G8` is a single 32-bit block of two pixels, where the `R` and `B` values are repeated in both pixels.

References

DirectX Graphics Infrastructure documentation on `DXGI_FORMAT` enum: <http://msdn.microsoft.com/en-us/library/windows/desktop/bb173059%28v=vs.85%29.aspx>

`FOURCC` codes for `YUV` formats: <http://www.fourcc.org/yuv.php>

A Gallium rendering context encapsulates the state which effects 3D rendering such as blend state, depth/stencil state, texture samplers, etc.

Note that resource/texture allocation is not per-context but per-screen.

Methods

CSO State

All Constant State Object (CSO) state is created, bound, and destroyed, with triplets of methods that all follow a specific naming scheme. For example, `create_blend_state`, `bind_blend_state`, and `destroy_blend_state`.

CSO objects handled by the context object:

- *Blend*: `*_blend_state`
- *Sampler*: Texture sampler states are bound separately for fragment, vertex, geometry and compute shaders with the `bind_sampler_states` function. The `start` and `num_samplers` parameters indicate a range of samplers to change. NOTE: at this time, `start` is always zero and the CSO module will always replace all samplers at once (no sub-ranges). This may change in the future.
- *Rasterizer*: `*_rasterizer_state`
- *Depth, Stencil, & Alpha*: `*_depth_stencil_alpha_state`
- *Shader*: These are create, bind and destroy methods for vertex, fragment and geometry shaders.
- *Vertex Elements*: `*_vertex_elements_state`

Resource Binding State

This state describes how resources in various flavours (textures, buffers, surfaces) are bound to the driver.

- `set_constant_buffer` sets a constant buffer to be used for a given shader type. `index` is used to indicate which buffer to set (some apis may allow multiple ones to be set, and binding a specific one later, though drivers are mostly restricted to the first one right now).
- `set_framebuffer_state`
- `set_vertex_buffers`

Non-CSO State

These pieces of state are too small, variable, and/or trivial to have CSO objects. They all follow simple, one-method binding calls, e.g. `set_blend_color`.

- `set_stencil_ref` sets the stencil front and back reference values which are used as comparison values in stencil test.
- `set_blend_color`
- `set_sample_mask`
- `set_min_samples` sets the minimum number of samples that must be run.
- `set_clip_state`
- `set_polygon_stipple`
- `set_scissor_states` sets the bounds for the scissor test, which culls pixels before blending to render targets. If the *Rasterizer* does not have the scissor test enabled, then the scissor bounds never need to be set since they will not be used. Note that scissor `xmin` and `ymin` are inclusive, but `xmax` and `ymax` are exclusive. The inclusive ranges in `x` and `y` would be `[xmin..xmax-1]` and `[ymin..ymax-1]`. The number of scissors should be the same as the number of set viewports and can be up to `PIPE_MAX_VIEWPORTS`.
- `set_viewport_states`
- `set_window_rectangles` sets the window rectangles to be used for rendering, as defined by `GL_EXT_window_rectangles`. There are two modes - include and exclude, which define whether the supplied rectangles are to be used for including fragments or excluding them. All of the rectangles are ORed together, so in exclude mode, any fragment inside any rectangle would be culled, while in include mode, any fragment outside all rectangles would be culled. `xmin/ymin` are inclusive, while `xmax/ymax` are exclusive (same as scissor states above). Note that this only applies to draws, not clears or blits. (Blits have their own way to pass the requisite rectangles in.)
- `set_tess_state` configures the default tessellation parameters:
 - `default_outer_level` is the default value for the outer tessellation levels. This corresponds to GL's `PATCH_DEFAULT_OUTER_LEVEL`.
 - `default_inner_level` is the default value for the inner tessellation levels. This corresponds to GL's `PATCH_DEFAULT_INNER_LEVEL`.
- `set_debug_callback` sets the callback to be used for reporting various debug messages, eventually reported via `KHR_debug` and similar mechanisms.

Samplers

`pipe_sampler_state` objects control how textures are sampled (coordinate wrap modes, interpolation modes, etc). Note that samplers are not used for texture buffer objects. That is, `pipe_context::bind_sampler_views()` will not bind a sampler if the corresponding sampler view refers to a `PIPE_BUFFER` resource.

Sampler Views

These are the means to bind textures to shader stages. To create one, specify its format, swizzle and LOD range in sampler view template.

If texture format is different than template format, it is said the texture is being cast to another format. Casting can be done only between compatible formats, that is formats that have matching component order and sizes.

Swizzle fields specify the way in which fetched texel components are placed in the result register. For example, `swizzle_r` specifies what is going to be placed in first component of result register.

The `first_level` and `last_level` fields of sampler view template specify the LOD range the texture is going to be constrained to. Note that these values are in addition to the respective `min_lod`, `max_lod` values in the `pipe_sampler_state` (that is if `min_lod` is 2.0, and `first_level` 3, the first mip level used for sampling from the resource is effectively the fifth).

The `first_layer` and `last_layer` fields specify the layer range the texture is going to be constrained to. Similar to the LOD range, this is added to the array index which is used for sampling.

- `set_sampler_views` binds an array of sampler views to a shader stage. Every binding point acquires a reference to a respective sampler view and releases a reference to the previous sampler view.
- `create_sampler_view` creates a new sampler view. `texture` is associated with the sampler view which results in sampler view holding a reference to the texture. Format specified in template must be compatible with texture format.
- `sampler_view_destroy` destroys a sampler view and releases its reference to associated texture.

Shader Resources

Shader resources are textures or buffers that may be read or written from a shader without an associated sampler. This means that they have no support for floating point coordinates, address wrap modes or filtering.

There are 2 types of shader resources: buffers and images.

Buffers are specified using the `set_shader_buffers` method.

Images are specified using the `set_shader_images` method. When binding images, the `level`, `first_layer` and `last_layer` `pipe_image_view` fields specify the mipmap level and the range of layers the image will be constrained to.

Surfaces

These are the means to use resources as color render targets or depthstencil attachments. To create one, specify the mip level, the range of layers, and the bind flags (either `PIPE_BIND_DEPTH_STENCIL` or `PIPE_BIND_RENDER_TARGET`). Note that layer values are in addition to what is indicated by the geometry shader output variable `XXX_FIXME` (that is if `first_layer` is 3 and geometry shader indicates index 2, the 5th layer of the resource will be used). These `first_layer` and `last_layer` parameters will only be used for 1d array, 2d array, cube, and 3d textures otherwise they are 0.

- `create_surface` creates a new surface.
- `surface_destroy` destroys a surface and releases its reference to the associated resource.

Stream output targets

Stream output, also known as transform feedback, allows writing the primitives produced by the vertex pipeline to buffers. This is done after the geometry shader or vertex shader if no geometry shader is present.

The stream output targets are views into buffer resources which can be bound as stream outputs and specify a memory range where it's valid to write primitives. The pipe driver must implement memory protection such that any primitives written outside of the specified memory range are discarded.

Two stream output targets can use the same resource at the same time, but with a disjoint memory range.

Additionally, the stream output target internally maintains the offset into the buffer which is incremented everytime something is written to it. The internal offset is equal to how much data has already been written. It can be stored in device memory and the CPU actually doesn't have to query it.

The stream output target can be used in a draw command to provide the vertex count. The vertex count is derived from the internal offset discussed above.

- `create_stream_output_target` create a new target.
- `stream_output_target_destroy` destroys a target. Users of this should use `pipe_so_target_reference` instead.
- `set_stream_output_targets` binds stream output targets. The parameter `offset` is an array which specifies the internal offset of the buffer. The internal offset is, besides writing, used for reading the data during the `draw_auto` stage, i.e. it specifies how much data there is in the buffer for the purposes of the `draw_auto` stage. `-1` means the buffer should be appended to, and everything else sets the internal offset.

NOTE: The currently-bound vertex or geometry shader must be compiled with the properly-filled-in structure `pipe_stream_output_info` describing which outputs should be written to buffers and how. The structure is part of `pipe_shader_state`.

Clearing

Clear is one of the most difficult concepts to nail down to a single interface (due to both different requirements from APIs and also driver/hw specific differences).

`clear` initializes some or all of the surfaces currently bound to the framebuffer to particular RGBA, depth, or stencil values. Currently, this does not take into account color or stencil write masks (as used by GL), and always clears the whole surfaces (no scissoring as used by GL `clear` or explicit rectangles like `d3d9` uses). It can, however, also clear only depth or stencil in a combined depth/stencil surface. If a surface includes several layers then all layers will be cleared.

`clear_render_target` clears a single color rendertarget with the specified color value. While it is only possible to clear one surface at a time (which can include several layers), this surface need not be bound to the framebuffer. If `render_condition_enabled` is false, any current rendering condition is ignored and the clear will be unconditional.

`clear_depth_stencil` clears a single depth, stencil or depth/stencil surface with the specified depth and stencil values (for combined depth/stencil buffers, it is also possible to only clear one or the other part). While it is only possible to clear one surface at a time (which can include several layers), this surface need not be bound to the framebuffer. If `render_condition_enabled` is false, any current rendering condition is ignored and the clear will be unconditional.

`clear_texture` clears a non-PIPE_BUFFER resource's specified level and bounding box with a clear value provided in that resource's native format.

`clear_buffer` clears a PIPE_BUFFER resource with the specified clear value (which may be multiple bytes in length). Logically this is a memset with a multi-byte element value starting at `offset` bytes from resource start, going for `size` bytes. It is guaranteed that `size % clear_value_size == 0`.

Uploading

For simple single-use uploads, use `pipe_context::stream_uploader` or `pipe_context::const_uploader`. The latter should be used for uploading constants, while the former should be used for uploading everything else. `PIPE_USAGE_STREAM` is implied in both cases, so don't use the uploaders for static allocations.

Usage:

Call `u_upload_alloc` or `u_upload_data` as many times as you want. After you are done, call `u_upload_unmap`. If the driver doesn't support persistent mappings, `u_upload_unmap` makes sure the previously mapped memory is unmapped.

Gotchas: - Always fill the memory immediately after `u_upload_alloc`. Any following call to `u_upload_alloc` and `u_upload_data` can unmap memory returned by previous `u_upload_alloc`. - Don't interleave calls using `stream_uploader` and `const_uploader`. If you use one of them, do the upload, unmap, and only then can you use the other one.

Drawing

`draw_vbo` draws a specified primitive. The primitive mode and other properties are described by `pipe_draw_info`.

The `mode`, `start`, and `count` fields of `pipe_draw_info` specify the the mode of the primitive and the vertices to be fetched, in the range between `start` to `start`+`count-1`, inclusive.

Every instance with `instanceID` in the range between `start_instance` and `start_instance`+`instance_count-1`, inclusive, will be drawn.

If `index_size != 0`, all vertex indices will be looked up from the index buffer.

In indexed draw, `min_index` and `max_index` respectively provide a lower and upper bound of the indices contained in the index buffer inside the range between `start` to `start`+`count-1`. This allows the driver to determine which subset of vertices will be referenced during the draw call without having to scan the index buffer. Providing a over-estimation of the the true bounds, for example, a `min_index` and `max_index` of 0 and 0xffffffff respectively, must give exactly the same rendering, albeit with less performance due to unreferenced vertex buffers being unnecessarily DMA'ed or processed. Providing a underestimation of the true bounds will result in undefined behavior, but should not result in program or system failure.

In case of non-indexed draw, `min_index` should be set to `start` and `max_index` should be set to `start`+`count-1`.

`index_bias` is a value added to every vertex index after lookup and before fetching vertex attributes.

When drawing indexed primitives, the primitive restart index can be used to draw disjoint primitive strips. For example, several separate line strips can be drawn by designating a special index value as the restart index. The `primitive_restart` flag enables/disables this feature. The `restart_index` field specifies the restart index value.

When primitive restart is in use, array indexes are compared to the restart index before adding the `index_bias` offset.

If a given vertex element has `instance_divisor` set to 0, it is said it contains per-vertex data and effective vertex attribute address needs to be recalculated for every index.

$$\text{attribAddr} = \text{stride} * \text{index} + \text{src_offset}$$

If a given vertex element has `instance_divisor` set to non-zero, it is said it contains per-instance data and effective vertex attribute address needs to be recalculated for every `instance_divisor`-th instance.

$$\text{attribAddr} = \text{stride} * \text{instanceID} / \text{instance_divisor} + \text{src_offset}$$

In the above formulas, `src_offset` is taken from the given vertex element and `stride` is taken from a vertex buffer associated with the given vertex element.

The calculated `attribAddr` is used as an offset into the vertex buffer to fetch the attribute data.

The value of `instanceID` can be read in a vertex shader through a system value register declared with `INSTANCEID` semantic name.

Queries

Queries gather some statistic from the 3D pipeline over one or more draws. Queries may be nested, though not all state trackers exercise this.

Queries can be created with `create_query` and deleted with `destroy_query`. To start a query, use `begin_query`, and when finished, use `end_query` to end the query.

`create_query` takes a query type (`PIPE_QUERY_*`), as well as an index, which is the vertex stream for `PIPE_QUERY_PRIMITIVES_GENERATED` and `PIPE_QUERY_PRIMITIVES_EMITTED`, and allocates a query structure.

`begin_query` will clear/reset previous query results.

`get_query_result` is used to retrieve the results of a query. If the `wait` parameter is `TRUE`, then the `get_query_result` call will block until the results of the query are ready (and `TRUE` will be returned). Otherwise, if the `wait` parameter is `FALSE`, the call will not block and the return value will be `TRUE` if the query has completed or `FALSE` otherwise.

`get_query_result_resource` is used to store the result of a query into a resource without synchronizing with the CPU. This write will optionally wait for the query to complete, and will optionally write whether the value is available instead of the value itself.

`set_active_query_state` Set whether all current non-driver queries except `TIME_ELAPSED` are active or paused.

The interface currently includes the following types of queries:

`PIPE_QUERY_OCCLUSION_COUNTER` counts the number of fragments which are written to the framebuffer without being culled by *Depth, Stencil, & Alpha* testing or shader `KILL` instructions. The result is an unsigned 64-bit integer. This query can be used with `render_condition`.

In cases where a boolean result of an occlusion query is enough, `PIPE_QUERY_OCCLUSION_PREDICATE` should be used. It is just like `PIPE_QUERY_OCCLUSION_COUNTER` except that the result is a boolean value of `FALSE` for cases where `COUNTER` would result in 0 and `TRUE` for all other cases. This query can be used with `render_condition`.

In cases where a conservative approximation of an occlusion query is enough, `PIPE_QUERY_OCCLUSION_PREDICATE_CONSERVATIVE` should be used. It behaves like `PIPE_QUERY_OCCLUSION_PREDICATE`, except that it may return `TRUE` in additional, implementation-dependent cases. This query can be used with `render_condition`.

`PIPE_QUERY_TIME_ELAPSED` returns the amount of time, in nanoseconds, the context takes to perform operations. The result is an unsigned 64-bit integer.

`PIPE_QUERY_TIMESTAMP` returns a device/driver internal timestamp, scaled to nanoseconds, recorded after all commands issued prior to `end_query` have been processed. This query does not require a call to `begin_query`. The result is an unsigned 64-bit integer.

`PIPE_QUERY_TIMESTAMP_DISJOINT` can be used to check the internal timer resolution and whether the timestamp counter has become unreliable due to things like throttling etc. - only if this is `FALSE` a timestamp query (within the `timestamp_disjoint` query) should be trusted. The result is a 64-bit integer specifying the timer resolution in Hz, followed by a boolean value indicating whether the timestamp counter is discontinuous or disjoint.

`PIPE_QUERY_PRIMITIVES_GENERATED` returns a 64-bit integer indicating the number of primitives processed by the pipeline (regardless of whether stream output is active or not).

`PIPE_QUERY_PRIMITIVES_EMITTED` returns a 64-bit integer indicating the number of primitives written to stream output buffers.

`PIPE_QUERY_SO_STATISTICS` returns 2 64-bit integers corresponding to the result of `PIPE_QUERY_PRIMITIVES_EMITTED` and the number of primitives that would have been written to stream output buffers if they had infinite space available (`primitives_storage_needed`), in this order. XXX the 2nd value is equivalent to `PIPE_QUERY_PRIMITIVES_GENERATED` but it is unclear if it should be increased if stream output is not active.

`PIPE_QUERY_SO_OVERFLOW_PREDICATE` returns a boolean value indicating whether a selected stream output target has overflowed as a result of the commands issued between `begin_query` and `end_query`. This query can be used with `render_condition`. The output stream is selected by the stream number passed to `create_query`.

`PIPE_QUERY_SO_OVERFLOW_ANY_PREDICATE` returns a boolean value indicating whether any stream output target has overflowed as a result of the commands issued between `begin_query` and `end_query`. This query can be used with `render_condition`, and its result is the logical OR of multiple `PIPE_QUERY_SO_OVERFLOW_PREDICATE` queries, one for each stream output target.

`PIPE_QUERY_GPU_FINISHED` returns a boolean value indicating whether all commands issued before `end_query` have completed. However, this does not imply serialization. This query does not require a call to `begin_query`.

`PIPE_QUERY_PIPELINE_STATISTICS` returns an array of the following 64-bit integers: Number of vertices read from vertex buffers. Number of primitives read from vertex buffers. Number of vertex shader threads launched. Number of geometry shader threads launched. Number of primitives generated by geometry shaders. Number of primitives forwarded to the rasterizer. Number of primitives rasterized. Number of fragment shader threads launched. Number of tessellation control shader threads launched. Number of tessellation evaluation shader threads launched. If a shader type is not supported by the device/driver, the corresponding values should be set to 0.

Gallium does not guarantee the availability of any query types; one must always check the capabilities of the *Screen* first.

Conditional Rendering

A drawing command can be skipped depending on the outcome of a query (typically an occlusion query, or streamout overflow predicate). The `render_condition` function specifies the query which should be checked prior to rendering anything. Functions always honoring `render_condition` include (and are limited to) `draw_vbo` and `clear`. The `blit`, `clear_render_target` and `clear_depth_stencil` functions (but not `resource_copy_region`, which seems inconsistent) can also optionally honor the current render condition.

If `render_condition` is called with `query = NULL`, conditional rendering is disabled and drawing takes place normally.

If `render_condition` is called with a non-null `query` subsequent drawing commands will be predicated on the outcome of the query. Commands will be skipped if `condition` is equal to the predicate result (for non-boolean queries such as `OCCLUSION_QUERY`, zero counts as `FALSE`, non-zero as `TRUE`).

If `mode` is `PIPE_RENDER_COND_WAIT` the driver will wait for the query to complete before deciding whether to render.

If `mode` is `PIPE_RENDER_COND_NO_WAIT` and the query has not yet completed, the drawing command will be executed normally. If the query has completed, drawing will be predicated on the outcome of the query.

If `mode` is `PIPE_RENDER_COND_BY_REGION_WAIT` or `PIPE_RENDER_COND_BY_REGION_NO_WAIT` rendering will be predicated as above for the non-REGION modes but in the case that an occlusion query returns

a non-zero result, regions which were occluded may be omitted by subsequent drawing commands. This can result in better performance with some GPUs. Normally, if the occlusion query returned a non-zero result subsequent drawing happens normally so fragments may be generated, shaded and processed even where they're known to be obscured.

Flushing

`flush`

`PIPE_FLUSH_END_OF_FRAME`: Whether the flush marks the end of frame.

`PIPE_FLUSH_DEFERRED`: It is not required to flush right away, but it is required to return a valid fence. If `fence_finish` is called with the returned fence and the context is still unflushed, and the `ctx` parameter of `fence_finish` is equal to the context where the fence was created, `fence_finish` will flush the context.

`flush_resource`

Flush the resource cache, so that the resource can be used by an external client. Possible usage: - flushing a resource before presenting it on the screen - flushing a resource if some other process or device wants to use it This shouldn't be used to flush caches if the resource is only managed by a single `pipe_screen` and is not shared with another process. (i.e. you shouldn't use it to flush caches explicitly if you want to e.g. use the resource for texturing)

Resource Busy Queries

`is_resource_referenced`

Blitting

These methods emulate classic blitter controls.

These methods operate directly on `pipe_resource` objects, and stand apart from any 3D state in the context. Blitting functionality may be moved to a separate abstraction at some point in the future.

`resource_copy_region` blits a region of a resource to a region of another resource, provided that both resources have the same format, or compatible formats, i.e., formats for which copying the bytes from the source resource unmodified to the destination resource will achieve the same effect of a textured quad blitter.. The source and destination may be the same resource, but overlapping blits are not permitted. This can be considered the equivalent of a CPU `memcpy`.

`blit` blits a region of a resource to a region of another resource, including scaling, format conversion, and up-/downsampling, as well as a destination clip rectangle (scissors) and window rectangles. It can also optionally honor the current render condition (but either way the blit itself never contributes anything to queries currently gathering data). As opposed to manually drawing a textured quad, this lets the pipe driver choose the optimal method for blitting (like using a special 2D engine), and usually offers, for example, accelerated stencil-only copies even where `PIPE_CAP_SHADER_STENCIL_EXPORT` is not available.

Transfers

These methods are used to get data to/from a resource.

`transfer_map` creates a memory mapping and the transfer object associated with it. The returned pointer points to the start of the mapped range according to the box region, not the beginning of the resource. If `transfer_map` fails, the returned pointer to the buffer memory is `NULL`, and the pointer to the transfer object remains unchanged (i.e. it can be non-`NULL`).

`transfer_unmap` remove the memory mapping for and destroy the transfer object. The pointer into the resource should be considered invalid and discarded.

`texture_subdata` and `buffer_subdata` perform a simplified transfer for simple writes. Basically `transfer_map`, data write, and `transfer_unmap` all in one.

The `box` parameter to some of these functions defines a 1D, 2D or 3D region of pixels. This is self-explanatory for 1D, 2D and 3D texture targets.

For `PIPE_TEXTURE_1D_ARRAY` and `PIPE_TEXTURE_2D_ARRAY`, the `box::z` and `box::depth` fields refer to the array dimension of the texture.

For `PIPE_TEXTURE_CUBE`, the `box::z` and `box::depth` fields refer to the faces of the cube map ($z + \text{depth} \leq 6$).

For `PIPE_TEXTURE_CUBE_ARRAY`, the `box::z` and `box::depth` fields refer to both the face and array dimension of the texture ($\text{face} = z \% 6$, $\text{array} = z / 6$).

transfer_flush_region

If a transfer was created with `FLUSH_EXPLICIT`, it will not automatically be flushed on write or unmap. Flushes must be requested with `transfer_flush_region`. Flush ranges are relative to the mapped range, not the beginning of the resource.

texture_barrier

This function flushes all pending writes to the currently-set surfaces and invalidates all read caches of the currently-set samplers. This can be used for both regular textures as well as for framebuffer read via `FBFETCH`.

memory_barrier

This function flushes caches according to which of the `PIPE_BARRIER_*` flags are set.

resource_commit

This function changes the commit state of a part of a sparse resource. Sparse resources are created by setting the `PIPE_RESOURCE_FLAG_SPARSE` flag when calling `resource_create`. Initially, sparse resources only reserve a virtual memory region that is not backed by memory (i.e., it is uncommitted). The `resource_commit` function can be called to commit or uncommit parts (or all) of a resource. The driver manages the underlying backing memory.

The contents of newly committed memory regions are undefined. Calling this function to commit an already committed memory region is allowed and leaves its content unchanged. Similarly, calling this function to uncommit an already uncommitted memory region is allowed.

For buffers, the given `box` must be aligned to multiples of `PIPE_CAP_SPARSE_BUFFER_PAGE_SIZE`. As an exception to this rule, if the size of the buffer is not a multiple of the page size, changing the commit state of the last (partial) page requires a `box` that ends at the end of the buffer (i.e., $\text{box->x} + \text{box->width} == \text{buffer->width}$).

PIPE_TRANSFER

These flags control the behavior of a transfer object.

PIPE_TRANSFER_READ Resource contents read back (or accessed directly) at transfer create time.

PIPE_TRANSFER_WRITE Resource contents will be written back at `transfer_unmap` time (or modified as a result of being accessed directly).

PIPE_TRANSFER_MAP_DIRECTLY a transfer should directly map the resource. May return NULL if not supported.

PIPE_TRANSFER_DISCARD_RANGE The memory within the mapped region is discarded. Cannot be used with `PIPE_TRANSFER_READ`.

PIPE_TRANSFER_DISCARD_WHOLE_RESOURCE Discards all memory backing the resource. It should not be used with `PIPE_TRANSFER_READ`.

PIPE_TRANSFER_DONTBLOCK Fail if the resource cannot be mapped immediately.

PIPE_TRANSFER_UNSYNCHRONIZED Do not synchronize pending operations on the resource when mapping. The interaction of any writes to the map and any operations pending on the resource are undefined. Cannot be used with `PIPE_TRANSFER_READ`.

PIPE_TRANSFER_FLUSH_EXPLICIT Written ranges will be notified later with *transfer_flush_region*. Cannot be used with `PIPE_TRANSFER_READ`.

PIPE_TRANSFER_PERSISTENT Allows the resource to be used for rendering while mapped. `PIPE_RESOURCE_FLAG_MAP_PERSISTENT` must be set when creating the resource. If `COHERENT` is not set, `memory_barrier(PIPE_BARRIER_MAPPED_BUFFER)` must be called to ensure the device can see what the CPU has written.

PIPE_TRANSFER_COHERENT If `PERSISTENT` is set, this ensures any writes done by the device are immediately visible to the CPU and vice versa. `PIPE_RESOURCE_FLAG_MAP_COHERENT` must be set when creating the resource.

Compute kernel execution

A compute program can be defined, bound or destroyed using `create_compute_state`, `bind_compute_state` or `destroy_compute_state` respectively.

Any of the subroutines contained within the compute program can be executed on the device using the `launch_grid` method. This method will execute as many instances of the program as elements in the specified N-dimensional grid, hopefully in parallel.

The compute program has access to four special resources:

- `GLOBAL` represents a memory space shared among all the threads running on the device. An arbitrary buffer created with the `PIPE_BIND_GLOBAL` flag can be mapped into it using the `set_global_binding` method.
- `LOCAL` represents a memory space shared among all the threads running in the same working group. The initial contents of this resource are undefined.
- `PRIVATE` represents a memory space local to a single thread. The initial contents of this resource are undefined.
- `INPUT` represents a read-only memory space that can be initialized at `launch_grid` time.

These resources use a byte-based addressing scheme, and they can be accessed from the compute program by means of the `LOAD/STORE TGSI` opcodes. Additional resources to be accessed using the same opcodes may be specified by the user with the `set_compute_resources` method.

In addition, normal texture sampling is allowed from the compute program: `bind_sampler_states` may be used to set up texture samplers for the compute stage and `set_sampler_views` may be used to bind a number of sampler views to it.

Mipmap generation

If `PIPE_CAP_GENERATE_MIPMAP` is true, `generate_mipmap` can be used to generate mipmaps for the specified texture resource. It replaces texel image levels `base_level+1` through `last_level` for layers range from `first_layer`

through `last_layer`. It returns `TRUE` if mipmap generation succeeds, otherwise it returns `FALSE`. Mipmap generation may fail when it is not supported for particular texture types or formats.

Device resets

The state tracker can query or request notifications of when the GPU is reset for whatever reason (application error, driver error). When a GPU reset happens, the context becomes unusable and all related state should be considered lost and undefined. Despite that, context notifications are single-shot, i.e. subsequent calls to `get_device_reset_status` will return `PIPE_NO_RESET`.

- `get_device_reset_status` queries whether a device reset has happened since the last call or since the last notification by callback.
- `set_device_reset_callback` sets a callback which will be called when a device reset is detected. The callback is only called synchronously.

Bindless

If `PIPE_CAP_BINDLESS_TEXTURE` is `TRUE`, the following `pipe_context` functions are used to create/delete bindless handles, and to make them resident in the current context when they are going to be used by shaders.

- `create_texture_handle` creates a 64-bit unsigned integer texture handle that is going to be directly used in shaders.
- `delete_texture_handle` deletes a 64-bit unsigned integer texture handle.
- `make_texture_handle_resident` makes a 64-bit unsigned texture handle resident in the current context to be accessible by shaders for texture mapping.
- `create_image_handle` creates a 64-bit unsigned integer image handle that is going to be directly used in shaders.
- `delete_image_handle` deletes a 64-bit unsigned integer image handle.
- `make_image_handle_resident` makes a 64-bit unsigned integer image handle resident in the current context to be accessible by shaders for image loads, stores and atomic operations.

Using several contexts

Several contexts from the same screen can be used at the same time. Objects created on one context cannot be used in another context, but the objects created by the screen methods can be used by all contexts.

Transfers

A transfer on one context is not expected to synchronize properly with rendering on other contexts, thus only areas not yet used for rendering should be locked.

A flush is required after `transfer_unmap` to expect other contexts to see the uploaded data, unless:

- Using persistent mapping. Associated with coherent mapping, unmapping the resource is also not required to use it in other contexts. Without coherent mapping, `memory_barrier(PIPE_BARRIER_MAPPED_BUFFER)` should be called on the context that has mapped the resource. No flush is required.
- Mapping the resource with `PIPE_TRANSFER_MAP_DIRECTLY`.

CSO, Constant State Objects, are a core part of Gallium's API.

CSO work on the principle of reusable state; they are created by filling out a state object with the desired properties, then passing that object to a context. The context returns an opaque context-specific handle which can be bound at any time for the desired effect.

Blend

This state controls blending of the final fragments into the target rendering buffers.

Blend Factors

The blend factors largely follow the same pattern as their counterparts in other modern and legacy drawing APIs.

Dual source blend factors are supported for up to 1 MRT, although you can advertise > 1 MRT, the stack cannot handle them for a few reasons. There is no definition on how the 1D array of shader outputs should be mapped to something that would be a 2D array (location, index). No current hardware exposes > 1 MRT, and we should revisit this issue if anyone ever does.

Logical Operations

Logical operations, also known as logicops, lops, or rops, are supported. Only two-operand logicops are available. When logicops are enabled, all other blend state is ignored, including per-render-target state, so logicops are performed on all render targets.

Warning: The `blend_enable` flag is ignored for all render targets when logical operations are enabled.

For a source component s and destination component d , the logical operations are defined as taking the bits of each channel of each component, and performing one of the following operations per-channel:

- CLEAR: 0
- NOR: $\neg(s \vee d)$
- AND_INVERTED: $\neg s \wedge d$
- COPY_INVERTED: $\neg s$
- AND_REVERSE: $s \wedge \neg d$
- INVERT: $\neg d$
- XOR: $s \oplus d$
- NAND: $\neg(s \wedge d)$
- AND: $s \wedge d$
- EQUIV: $\neg(s \oplus d)$
- NOOP: d
- OR_INVERTED: $\neg s \vee d$
- COPY: s
- OR_REVERSE: $s \vee \neg d$
- OR: $s \vee d$
- SET: 1

Note: The logical operation names and definitions match those of the OpenGL API, and are similar to the ROP2 and ROP3 definitions of GDI. This is intentional, to ease transitions to Gallium.

Members

These members affect all render targets.

dither

Whether dithering is enabled.

Note: Dithering is completely implementation-dependent. It may be ignored by drivers for any reason, and some render targets may always or never be dithered depending on their format or usage flags.

logicop_enable

Whether the blender should perform a logicop instead of blending.

logicop_func

The logicop to use. One of PIPE_LOGICOP.

independent_blend_enable If enabled, blend state is different for each render target, and for each render target set in the respective member of the `rt` array. If disabled, blend state is the same for all render targets, and only the first member of the `rt` array contains valid data.

rt Contains the per-rendertarget blend state.

alpha_to_coverage If enabled, the fragment's alpha value is used to override the fragment's coverage mask. The coverage mask will be all zeros if the alpha value is zero. The coverage mask will be all ones if the alpha value is one. Otherwise, the number of bits set in the coverage mask will be proportional to the alpha value. Note that this step happens regardless of whether multisample is enabled or the destination buffer is multisampled.

alpha_to_one If enabled, the fragment's alpha value will be set to one. As with `alpha_to_coverage`, this step happens regardless of whether multisample is enabled or the destination buffer is multisampled.

Per-rendertarget Members

blend_enable If blending is enabled, perform a blend calculation according to blend functions and source/destination factors. Otherwise, the incoming fragment color gets passed unmodified (but `colormask` still applies).

rgb_func The blend function to use for `rgb` channels. One of `PIPE_BLEND`.

rgb_src_factor The blend source factor to use for `rgb` channels. One of `PIPE_BLENDFACTOR`.

rgb_dst_factor The blend destination factor to use for `rgb` channels. One of `PIPE_BLENDFACTOR`.

alpha_func The blend function to use for the alpha channel. One of `PIPE_BLEND`.

alpha_src_factor The blend source factor to use for the alpha channel. One of `PIPE_BLENDFACTOR`.

alpha_dst_factor The blend destination factor to use for alpha channel. One of `PIPE_BLENDFACTOR`.

colormask Bitmask of which channels to write. Combination of `PIPE_MASK` bits.

Depth, Stencil, & Alpha

These three states control the depth, stencil, and alpha tests, used to discard fragments that have passed through the fragment shader.

Traditionally, these three tests have been clumped together in hardware, so they are all stored in one structure.

During actual execution, the order of operations done on fragments is always:

- Alpha
- Stencil
- Depth

Depth Members

enabled Whether the depth test is enabled.

writemask Whether the depth buffer receives depth writes.

func The depth test function. One of `PIPE_FUNC`.

Stencil Members

enabled Whether the stencil test is enabled. For the second stencil, whether the two-sided stencil is enabled. If two-sided stencil is disabled, the other fields for the second array member are not valid.

func The stencil test function. One of PIPE_FUNC.

valuemask Stencil test value mask; this is ANDed with the value in the stencil buffer and the reference value before doing the stencil comparison test.

writemask Stencil test writemask; this controls which bits of the stencil buffer are written.

fail_op The operation to carry out if the stencil test fails. One of PIPE_STENCIL_OP.

zfail_op The operation to carry out if the stencil test passes but the depth test fails. One of PIPE_STENCIL_OP.

zpass_op The operation to carry out if the stencil test and depth test both pass. One of PIPE_STENCIL_OP.

Alpha Members

enabled Whether the alpha test is enabled.

func The alpha test function. One of PIPE_FUNC.

ref_value Alpha test reference value; used for certain functions.

Rasterizer

The rasterizer state controls the rendering of points, lines and triangles. Attributes include polygon culling state, line width, line stipple, multisample state, scissoring and flat/smooth shading.

Linkage

clamp_vertex_color

If set, TGSI_SEMANTIC_COLOR registers are clamped to the [0, 1] range after the execution of the vertex shader, before being passed to the geometry shader or fragment shader.

OpenGL: glClampColor(GL_CLAMP_VERTEX_COLOR) in GL 3.0 or GL_ARB_color_buffer_float

D3D11: seems always disabled

Note the PIPE_CAP_VERTEX_COLOR_CLAMPED query indicates whether or not the driver supports this control. If it's not supported, the state tracker may have to insert extra clamping code.

clamp_fragment_color

Controls whether TGSI_SEMANTIC_COLOR outputs of the fragment shader are clamped to [0, 1].

OpenGL: glClampColor(GL_CLAMP_FRAGMENT_COLOR) in GL 3.0 or ARB_color_buffer_float

D3D11: seems always disabled

Note the PIPE_CAP_FRAGMENT_COLOR_CLAMPED query indicates whether or not the driver supports this control. If it's not supported, the state tracker may have to insert extra clamping code.

Shading

flatshade

If set, the provoking vertex of each polygon is used to determine the color of the entire polygon. If not set, fragment colors will be interpolated between the vertex colors.

The actual interpolated shading algorithm is obviously implementation-dependent, but will usually be Gourard for most hardware.

Note: This is separate from the fragment shader input attributes CONSTANT, LINEAR and PERSPECTIVE. The flatshade state is needed at clipping time to determine how to set the color of new vertices.

Draw can implement flat shading by copying the provoking vertex color to all the other vertices in the primitive.

flatshade_first

Whether the first vertex should be the provoking vertex, for most primitives. If not set, the last vertex is the provoking vertex.

There are a few important exceptions to the specification of this rule.

- `PIPE_PRIMITIVE_POLYGON`: The provoking vertex is always the first vertex. If the caller wishes to change the provoking vertex, they merely need to rotate the vertices themselves.
- `PIPE_PRIMITIVE_QUAD`, `PIPE_PRIMITIVE_QUAD_STRIP`: The option only has an effect if `PIPE_CAP_QUADS_FOLLOW_PROVOKING_VERTEX_CONVENTION` is true. If it is not, the provoking vertex is always the last vertex.
- `PIPE_PRIMITIVE_TRIANGLE_FAN`: When set, the provoking vertex is the second vertex, not the first. This permits each segment of the fan to have a different color.

Polygons

light_twoside

If set, there are per-vertex back-facing colors. The hardware (perhaps assisted by *Draw*) should be set up to use this state along with the front/back information to set the final vertex colors prior to rasterization.

The frontface vertex shader color output is marked with TGSI semantic COLOR[0], and backface COLOR[1].

front_ccw Indicates whether the window order of front-facing polygons is counter-clockwise (TRUE) or clockwise (FALSE).

cull_mode Indicates which faces of polygons to cull, either `PIPE_FACE_NONE` (cull no polygons), `PIPE_FACE_FRONT` (cull front-facing polygons), `PIPE_FACE_BACK` (cull back-facing polygons), or `PIPE_FACE_FRONT_AND_BACK` (cull all polygons).

fill_front Indicates how to fill front-facing polygons, either `PIPE_POLYGON_MODE_FILL`, `PIPE_POLYGON_MODE_LINE` or `PIPE_POLYGON_MODE_POINT`.

fill_back Indicates how to fill back-facing polygons, either `PIPE_POLYGON_MODE_FILL`, `PIPE_POLYGON_MODE_LINE` or `PIPE_POLYGON_MODE_POINT`.

poly_stipple_enable Whether polygon stippling is enabled.

poly_smooth Controls OpenGL-style polygon smoothing/antialiasing

offset_point If set, point-filled polygons will have polygon offset factors applied

offset_line If set, line-filled polygons will have polygon offset factors applied

offset_tri If set, filled polygons will have polygon offset factors applied

offset_units Specifies the polygon offset bias

offset_units_unscaled Specifies the unit of the polygon offset bias. If false, use the GL/D3D1X behaviour. If true, `offset_units` is a floating point offset which isn't scaled (D3D9). Note that GL/D3D1X behaviour has different formula whether the depth buffer is unorm or float, which is not the case for D3D9.

offset_scale Specifies the polygon offset scale

offset_clamp Upper (if > 0) or lower (if < 0) bound on the polygon offset result

Lines

line_width The width of lines.

line_smooth Whether lines should be smoothed. Line smoothing is simply anti-aliasing.

line_stipple_enable Whether line stippling is enabled.

line_stipple_pattern 16-bit bitfield of on/off flags, used to pattern the line stipple.

line_stipple_factor When drawing a stippled line, each bit in the stipple pattern is repeated N times, where $N = \text{line_stipple_factor} + 1$.

line_last_pixel Controls whether the last pixel in a line is drawn or not. OpenGL omits the last pixel to avoid double-drawing pixels at the ends of lines when drawing connected lines.

Points

sprite_coord_enable

The effect of this state depends on `PIPE_CAP_TGSI_TEXCOORD` !

Controls automatic texture coordinate generation for rendering sprite points.

If `PIPE_CAP_TGSI_TEXCOORD` is false: When bit `k` in the `sprite_coord_enable` bitfield is set, then generic input `k` to the fragment shader will get an automatically computed texture coordinate.

If `PIPE_CAP_TGSI_TEXCOORD` is true: The bitfield refers to inputs with `TEXCOORD` semantic instead of generic inputs.

The texture coordinate will be of the form $(s, t, 0, 1)$ where `s` varies from 0 to 1 from left to right while `t` varies from 0 to 1 according to the state of `'sprite_coord_mode'` (see below).

If any bit is set, then `point_smooth` MUST be disabled (there are no round sprites) and `point_quad_rasterization` MUST be true (sprites are always rasterized as quads). Any mismatch between these states should be considered a bug in the state-tracker.

This feature is implemented in the *Draw* module but may also be implemented natively by GPUs or implemented with a geometry shader.

sprite_coord_mode

Specifies how the value for each shader output should be computed when drawing point sprites. For PIPE_SPRITE_COORD_LOWER_LEFT, the lower-left vertex will have coordinates (0,0,0,1). For PIPE_SPRITE_COORD_UPPER_LEFT, the upper-left vertex will have coordinates (0,0,0,1). This state is used by *Draw* to generate texcoords.

point_quad_rasterization

Determines if points should be rasterized according to quad or point rasterization rules.

(Legacy-only) OpenGL actually has quite different rasterization rules for points and point sprites - hence this indicates if points should be rasterized as points or according to point sprite (which decomposes them into quads, basically) rules. Newer GL versions no longer support the old point rules at all.

Additionally Direct3D will always use quad rasterization rules for points, regardless of whether point sprites are enabled or not.

If this state is enabled, point smoothing and antialiasing are disabled. If it is disabled, point sprite coordinates are not generated.

Note: Some renderers always internally translate points into quads; this state still affects those renderers by overriding other rasterization state.

point_tri_clip Determines if clipping of points should happen after they are converted to “rectangles” (required by d3d) or before (required by OpenGL, though this rule is ignored by some IHVs). It is not valid to set this to enabled but have point_quad_rasterization disabled.

point_smooth Whether points should be smoothed. Point smoothing turns rectangular points into circles or ovals.

point_size_per_vertex Whether the vertex shader is expected to have a point size output. Undefined behaviour is permitted if there is disagreement between this flag and the actual bound shader.

point_size The size of points, if not specified per-vertex.

Other Members

scissor Whether the scissor test is enabled.

multisample Whether *MSAA* is enabled.

half_pixel_center When true, the rasterizer should use (0.5, 0.5) pixel centers for determining pixel ownership (e.g., OpenGL, D3D10 and higher):




When false, the rasterizer should use (0, 0) pixel centers for determining pixel ownership (e.g., D3D9 or earlier):





bottom_edge_rule Determines what happens when a pixel sample lies precisely on a triangle edge.

When true, a pixel sample is considered to lie inside of a triangle if it lies on the *bottom edge* or *left edge* (e.g., OpenGL drawables):



When false, a pixel sample is considered to lie inside of a triangle if it lies on the *top edge* or *left edge* (e.g., OpenGL FBOs, D3D):



Where:

- a *top edge* is an edge that is horizontal and is above the other edges;
- a *bottom edge* is an edge that is horizontal and is below the other edges;
- a *left edge* is an edge that is not horizontal and is on the left side of the triangle.

Note: Actually all graphics APIs use a top-left rasterization rule for pixel ownership, but their notion of top varies with the axis origin (which can be either at $y = 0$ or at $y = \text{height}$). Gallium instead always assumes that top is always at $y=0$.

See also:

- <http://msdn.microsoft.com/en-us/library/windows/desktop/cc627092.aspx>
- <http://msdn.microsoft.com/en-us/library/windows/desktop/bb147314.aspx>

clip_halfz When true clip space in the z axis goes from $[0..1]$ (D3D). When false $[-1, 1]$ (GL)

depth_clip When false, the near and far depth clipping planes of the view volume are disabled and the depth value will be clamped at the per-pixel level, after polygon offset has been applied and before depth testing.

clip_plane_enable For each k in $[0, \text{PIPE_MAX_CLIP_PLANES})$, if bit k of this field is set, clipping half-space k is enabled, if it is clear, it is disabled. The clipping half-spaces are defined either by the user clip planes in

`pipe_clip_state`, or by the clip distance outputs of the shader stage preceding the fragment shader. If any clip distance output is written, those half-spaces for which no clip distance is written count as disabled; i.e. user clip planes and shader clip distances cannot be mixed, and clip distances take precedence.

Sampler

Texture units have many options for selecting texels from loaded textures; this state controls an individual texture unit's texel-sampling settings.

Texture coordinates are always treated as four-dimensional, and referred to with the traditional (S, T, R, Q) notation.

Members

wrap_s How to wrap the S coordinate. One of `PIPE_TEX_WRAP_*`.

wrap_t How to wrap the T coordinate. One of `PIPE_TEX_WRAP_*`.

wrap_r How to wrap the R coordinate. One of `PIPE_TEX_WRAP_*`.

The wrap modes are:

- `PIPE_TEX_WRAP_REPEAT`: Standard coord repeat/wrap-around mode.
- `PIPE_TEX_WRAP_CLAMP_TO_EDGE`: Clamp coord to edge of texture, the border color is never sampled.
- `PIPE_TEX_WRAP_CLAMP_TO_BORDER`: Clamp coord to border of texture, the border color is sampled when coords go outside the range [0,1].
- `PIPE_TEX_WRAP_CLAMP`: The coord is clamped to the range [0,1] before scaling to the texture size. This corresponds to the legacy OpenGL `GL_CLAMP` texture wrap mode. Historically, this mode hasn't acted consistently across all graphics hardware. It sometimes acts like `CLAMP_TO_EDGE` or `CLAMP_TO_BORDER`. The behaviour may also vary depending on linear vs. nearest sampling mode.
- `PIPE_TEX_WRAP_MIRROR_REPEAT`: If the integer part of the coordinate is odd, the coord becomes (1 - coord). Then, normal texture `REPEAT` is applied to the coord.
- `PIPE_TEX_WRAP_MIRROR_CLAMP_TO_EDGE`: First, the absolute value of the coordinate is computed. Then, regular `CLAMP_TO_EDGE` is applied to the coord.
- `PIPE_TEX_WRAP_MIRROR_CLAMP_TO_BORDER`: First, the absolute value of the coordinate is computed. Then, regular `CLAMP_TO_BORDER` is applied to the coord.
- `PIPE_TEX_WRAP_MIRROR_CLAMP`: First, the absolute value of the coord is computed. Then, regular `CLAMP` is applied to the coord.

min_img_filter The image filter to use when minifying texels. One of `PIPE_TEX_FILTER_*`.

mag_img_filter The image filter to use when magnifying texels. One of `PIPE_TEX_FILTER_*`.

The texture image filter modes are:

- `PIPE_TEX_FILTER_NEAREST`: One texel is fetched from the texture image at the texture coordinate.
- `PIPE_TEX_FILTER_LINEAR`: Two, four or eight texels (depending on the texture dimensions; 1D/2D/3D) are fetched from the texture image and linearly weighted and blended together.

min_mip_filter The filter to use when minifying mipmapped textures. One of `PIPE_TEX_MIPFILTER_*`.

The texture mip filter modes are:

- `PIPE_TEX_MIPFILTER_NEAREST`: A single mipmap level/image is selected according to the texture LOD (lambda) value.
- `PIPE_TEX_MIPFILTER_LINEAR`: The two mipmap levels/images above/below the texture LOD value are sampled from. The results of sampling from those two images are blended together with linear interpolation.
- `PIPE_TEX_MIPFILTER_NONE`: Mipmap filtering is disabled. All texels are taken from the level 0 image.

compare_mode If set to `PIPE_TEX_COMPARE_R_TO_TEXTURE`, the result of texture sampling is not a color but a true/false value which is the result of comparing the sampled texture value (typically a Z value from a depth texture) to the texture coordinate's R component. If set to `PIPE_TEX_COMPARE_NONE`, no comparison calculation is performed.

compare_func The inequality operator used when `compare_mode=1`. One of `PIPE_FUNC_x`.

normalized_coords If set, the incoming texture coordinates (nominally in the range [0,1]) will be scaled by the texture width, height, depth to compute texel addresses. Otherwise, the texture coords are used as-is (they are not scaled by the texture dimensions). When `normalized_coords=0`, only a subset of the texture wrap modes are allowed: `PIPE_TEX_WRAP_CLAMP`, `PIPE_TEX_WRAP_CLAMP_TO_EDGE` and `PIPE_TEX_WRAP_CLAMP_TO_BORDER`.

lod_bias Bias factor which is added to the computed level of detail. The normal level of detail is computed from the partial derivatives of the texture coordinates and/or the fragment shader `TEX/TXB/TXL` instruction.

min_lod Minimum level of detail, used to clamp LOD after bias. The LOD values correspond to mipmap levels where `LOD=0` is the level 0 mipmap image.

max_lod Maximum level of detail, used to clamp LOD after bias.

border_color Color union used for texel coordinates that are outside the [0,width-1], [0, height-1] or [0, depth-1] ranges. Interpreted according to sampler view format, unless the driver reports `PIPE_CAP_TEXTURE_BORDER_COLOR_QUIRK`, in which case special care has to be taken (see description of the cap).

max_anisotropy Maximum anisotropy ratio to use when sampling from textures. For example, if `max_anisotropy=4`, a region of up to 1 by 4 texels will be sampled. Set to zero to disable anisotropic filtering. Any other setting enables anisotropic filtering, however it's not unexpected some drivers only will change their filtering with a setting of 2 and higher.

seamless_cube_map If set, the bilinear filter of a cube map may take samples from adjacent cube map faces when sampled near a texture border to produce a seamless look.

Shader

One of the two types of shaders supported by Gallium.

Members

tokens A list of `tgsl_tokens`.

Vertex Elements

This state controls the format of the input attributes contained in `pipe_vertex_buffers`. There is one `pipe_vertex_element` array member for each input attribute.

Input Formats

Gallium supports a diverse range of formats for vertex data. Drivers are guaranteed to support 32-bit floating-point vectors of one to four components. Additionally, they may support the following formats:

- Integers, signed or unsigned, normalized or non-normalized, 8, 16, or 32 bits wide
- Floating-point, 16, 32, or 64 bits wide

At this time, support for varied vertex data formats is limited by driver deficiencies. It is planned to support a single uniform set of formats for all Gallium drivers at some point.

Rather than attempt to specify every small nuance of behavior, Gallium uses a very simple set of rules for padding out unspecified components. If an input uses less than four components, it will be padded out with the constant vector $(0, 0, 0, 1)$.

Fog, point size, the facing bit, and edgeflags, all are in the standard format of $(x, 0, 0, 1)$, and so only the first component of those inputs is used.

Position

Vertex position may be specified with two to four components. Using less than two components is not allowed.

Colors

Colors, both front- and back-facing, may omit the alpha component, only using three components. Using less than three components is not allowed.

Members

src_offset The byte offset of the attribute in the buffer given by `vertex_buffer_index` for the first vertex.

instance_divisor The instance data rate divisor, used for instancing. 0 means this is per-vertex data, n means per-instance data used for n consecutive instances ($n > 0$).

vertex_buffer_index The vertex buffer this attribute lives in. Several attributes may live in the same vertex buffer.

src_format The format of the attribute data. One of the `PIPE_FORMAT` tokens.

Along with the interface definitions, the following drivers, state trackers, and auxiliary modules are shipped in the standard Gallium distribution.

Drivers

Intel i915

Driver for Intel i915 and i945 chipsets.

LLVM Softpipe

A version of *Softpipe* that uses the Low-Level Virtual Machine to dynamically generate optimized rasterizing pipelines.

nVidia nv30

Driver for the nVidia nv30 and nv40 families of GPUs.

nVidia nv50

Driver for the nVidia nv50 family of GPUs.

nVidia nvc0

Driver for the nVidia nvc0 / fermi family of GPUs.

VMware SVGA

Driver for VMware virtualized guest operating system graphics processing.

ATI r300

Driver for the ATI/AMD r300, r400, and r500 families of GPUs.

ATI/AMD r600

Driver for the ATI/AMD r600, r700, Evergreen and Northern Islands families of GPUs.

AMD radeonsi

Driver for the AMD Southern Islands family of GPUs.

freedreno

Driver for Qualcomm Adreno a2xx, a3xx, and a4xx series of GPUs.

Softpipe

Reference software rasterizer. Slow but accurate.

Trace

Wrapper driver. Trace dumps an XML record of the calls made to the *Context* and *Screen* objects that it wraps.

Rbug

Wrapper driver. *Remote Debugger* driver used with stand alone rbug-gui.

State Trackers

Clover

Tracker that implements the Khronos OpenCL standard.

Direct Rendering Infrastructure

Tracker that implements the client-side DRI protocol, for providing direct acceleration services to X11 servers with the DRI extension. Supports DRI1 and DRI2. Only GL is supported.

GLX

MesaGL

Tracker implementing a GL state machine. Not usable as a standalone tracker; Mesa should be built with another state tracker, such as *Direct Rendering Infrastructure* or EGL.

VDPAU

Tracker for Video Decode and Presentation API for Unix.

WGL

Xorg DDX

Tracker for Xorg X11 servers. Provides device-dependent modesetting and acceleration as a DDX driver.

XvMC

Tracker for X-Video Motion Compensation.

Auxiliary

OS

The OS module contains the abstractions for basic operating system services:

- memory allocation
- simple message logging
- obtaining run-time configuration option
- threading primitives

This is the bare minimum required to port Gallium to a new platform.

The OS module already provides the implementations of these abstractions for the most common platforms. When targeting an embedded platform no implementation will be provided – these must be provided separately.

CSO Cache

The CSO cache is used to accelerate preparation of state by saving driver-specific state structures for later use.

Draw

Draw is a software *TCL* pipeline for hardware that lacks vertex shaders or other essential parts of pre-rasterization vertex preparation.

Gallivm

Indices

Indices provides tools for translating or generating element indices for use with element-based rendering.

Pipe Buffer Managers

Each of these managers provides various services to drivers that are not fully utilizing a memory manager.

Remote Debugger

Runtime Assembly Emission

TGSI

The TGSI auxiliary module provides basic utilities for manipulating TGSI streams.

Translate

Util

Driver specific documentation.

Freedreno

Freedreno driver specific docs.

IR3 NOTES

Some notes about ir3, the compiler and machine-specific IR for the shader ISA introduced with adreno a3xx. The same shader ISA is present, with some small differences, in adreno a4xx.

Compared to the previous generation a2xx ISA (ir2), the a3xx ISA is a “simple” scalar instruction set. However, the compiler is responsible, in most cases, to schedule the instructions. The hardware does not try to hide the shader core pipeline stages. For a common example, a common (cat2) ALU instruction takes four cycles, so a subsequent cat2 instruction which uses the result must have three intervening instructions (or nops). When operating on vec4’s, typically the corresponding scalar instructions for operating on the remaining three components could typically fit. Although that results in a lot of edge cases where things fall over, like:

```
ADD TEMP [0], TEMP [1], TEMP [2]
MUL TEMP [0], TEMP [1], TEMP [0].wzyx
```

Here, the second instruction needs the output of the first group of scalar instructions in the wrong order, resulting in not enough instruction spots between the `add r0.w, r1.w, r2.w` and `mul r0.x, r1.x, r0.w`. Which is why the original (old) compiler which merely translated nearly literally from TGSI to ir3, had a strong tendency to fall over.

So the current compiler instead, in the frontend, generates a directed-acyclic-graph of instructions and basic blocks, which go through various additional passes to eventually schedule and do register assignment.

For additional documentation about the hardware, see wiki: [a3xx ISA](#).

External Structure

ir3_shader A single vertex/fragment/etc shader from gallium perspective (ie. maps to a single TGSI shader), and manages a set of shader variants which are generated on demand based on the shader key.

ir3_shader_key The configuration key that identifies a shader variant. Ie. based on other GL state (two-sided-color, render-to-alpha, etc) or render stages (binning-pass vertex shader) different shader variants are generated.

ir3_shader_variant The actual hw shader generated based on input TGSI and shader key.

ir3_compiler Compiler frontend which generates ir3 and runs the various backend stages to schedule and do register assignment.


The IR

The ir3 IR maps quite directly to the hardware, in that instruction opcodes map directly to hardware opcodes, and that dst/src register(s) map directly to the hardware dst/src register(s). But there are a few extensions, in the form of *meta* instructions. And additionally, for normal (non-const, etc) src registers, the IR3_REG_SSA flag is set and `reg->instr` points to the source instruction which produced that value. So, for example, the following TGSI shader:

```

VERT
DCL IN[0]
DCL IN[1]
DCL OUT[0], POSITION
DCL TEMP[0], LOCAL
  1: DP3 TEMP[0].x, IN[0].xyzz, IN[1].xyzz
  2: MOV OUT[0], TEMP[0].xxxx
  3: END
    
```

eventually generates:



(after scheduling, etc, but before register assignment).

Internal Structure

ir3_block Represents a basic block.

TODO: currently blocks are nested, but I think I need to change that to a more conventional arrangement before implementing proper flow control. Currently the only flow control handles is *if/else* which gets flattened out and results chosen with *sel* instructions.

ir3_instruction Represents a machine instruction or *meta* instruction. Has pointers to dst register (`regs[0]`) and src register(s) (`regs[1..n]`), as needed.

ir3_register Represents a src or dst register, flags indicate const/relative/etc. If IR3_REG_SSA is set on a src register, the actual register number (name) has not been assigned yet, and instead the `instr` field points to src instruction.

In addition there are various util macros/functions to simplify manipulation/traversal of the graph:

foreach_src(srcreg, instr) Iterate each instruction's source `ir3_registers`

foreach_src_n(srcreg, n, instr) Like `foreach_src`, also setting `n` to the source number (starting with 0).

foreach_ssa_src(srcinstr, instr) Iterate each instruction's SSA source `ir3_instructions`. This skips non-SSA sources (consts, etc), but includes virtual sources (such as the address register if *relative addressing* is used).

foreach_ssa_src_n(srcinstr, n, instr) Like `foreach_ssa_src`, also setting `n` to the source number.

For example:

```
foreach_ssa_src_n(src, i, instr) {
    unsigned d = delay_calc_srcn(ctx, src, instr, i);
    delay = MAX2(delay, d);
}
```

TODO probably other helper/util stuff worth mentioning here

Meta Instructions

input Used for shader inputs (registers configured in the command-stream to hold particular input values, written by the shader core before start of execution. Also used for connecting up values within a basic block to an output of a previous block.

output Used to hold outputs of a basic block.

flow TODO

phi TODO

fanin Groups registers which need to be assigned to consecutive scalar registers, for example *sam* (texture fetch) src instructions (see *register groups*) or array element dereference (see *relative addressing*).

fanout The counterpart to **fanin**, when an instruction such as *sam* writes multiple components, splits the result into individual scalar components to be consumed by other instructions.

Flow Control

TODO


Register Groups

Certain instructions, such as texture sample instructions, consume multiple consecutive scalar registers via a single src register encoded in the instruction, and/or write multiple consecutive scalar registers. In the simplest example:

```
sam (f32) (xyz) r2.x, r0.z, s#0, t#0
```


for a 2d texture, would read `r0.zw` to get the coordinate, and write `r2.xyz`.

Before register assignment, to group the two components of the texture src together:



The frontend sets up the SSA ptrs from `sam` source register to the `fanin` meta instruction, which in turn points to the instructions producing the `coord.x` and `coord.y` values. And the *grouping* pass sets up the `left` and `right` neighbor pointers to the `fanin`'s sources, used later by the *register assignment* pass to assign blocks of scalar registers.

And likewise, for the consecutive scalar registers for the destination:



Relative Addressing

Most instructions support addressing indirectly (relative to address register) into const or gpr register file in some or all of their src/dst registers. In this case the register accessed is taken from $r\langle a0.x + n \rangle$ or $c\langle a0.x + n \rangle$, ie. address register ($a0.x$) value plus n , where n is encoded in the instruction (rather than the absolute register number).

Note that `cat5` (texture sample) instructions are the notable exception, not supporting relative addressing of src or dst.

Relative addressing of the const file (for example, a uniform array) is relatively simple. We don't do register assignment of the const file, so all that is required is to schedule things properly. Ie. the instruction that writes the address register must be scheduled first, and we cannot have two different address register values live at one time.

But relative addressing of gpr file (which can be as src or dst) has additional restrictions on register assignment (ie. the array elements must be assigned to consecutive scalar registers). And in the case of relative dst, subsequent instructions now depend on both the relative write, as well as the previous instruction which wrote that register, since we do not know at compile time which actual register was written.

Each instruction has an optional `address` pointer, to capture the dependency on the address register value when relative addressing is used for any of the src/dst register(s). This behaves as an additional virtual src register, ie. `foreach_ssa_src()` will also iterate the address register (last).


Note that `nop's` for timing constraints, type specifiers (ie. `add.f` vs `add.u`), etc, omitted for brevity in examples

```

mov a0.x, hr1.y
sub r1.y, r2.x, r3.x
add r0.x, r1.y, c<a0.x + 2>

```

results in:




The scheduling pass has some smarts to schedule things such that only a single `a0.x` value is used at any one time.

To implement variable arrays, values are stored in consecutive scalar registers. This has some overlap with *register groups*, in that `fanin` and `fanout` are used to help group things for the *register assignment* pass.

To use a variable array as a src register, a slight variation of what is done for const array src. The instruction src is a *fanin* instruction that groups all the array members:

```
mova a0.x, hr1.y
sub r1.y, r2.x, r3.x
add r0.x, r1.y, r<a0.x + 2>
```

results in:



TODO better describe how actual deref offset is derived, ie. based on array base register.

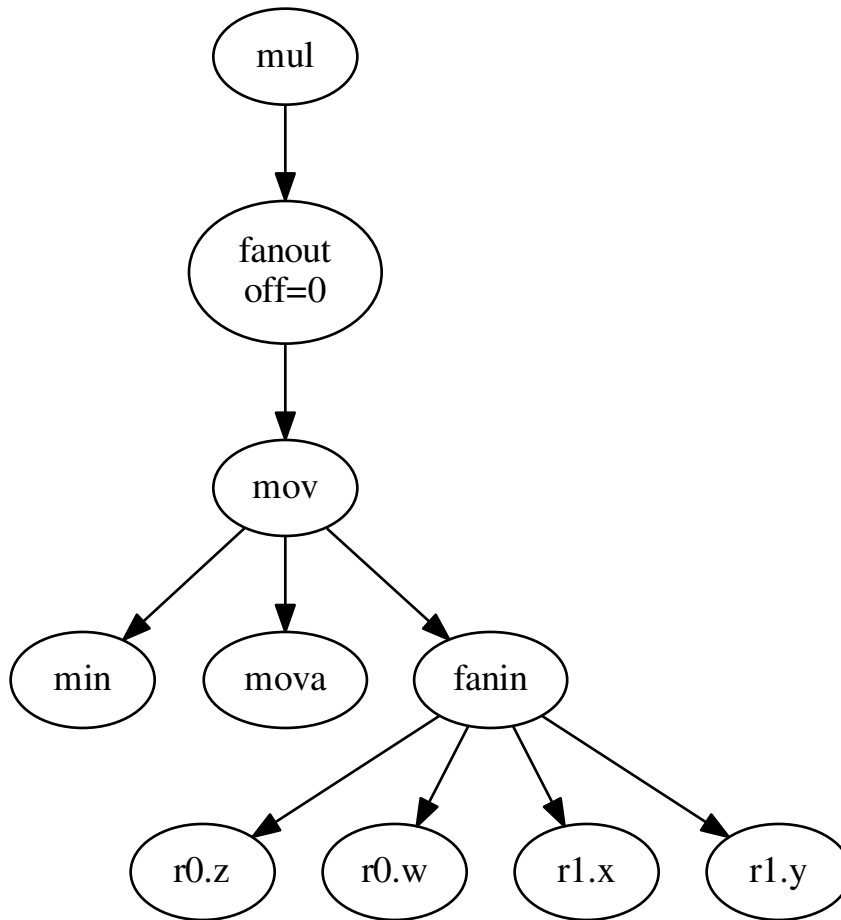
To do an indirect write to a variable array, a `fanout` is used. Say the array was assigned to registers `r0.z` through `r1.y` (hence the constant offset of 2):

Note that only `cat1` (`mov`) can do indirect write.

```

mova a0.x, hrl.y
min r2.x, r2.x, c0.x
mov r<a0.x + 2>, r2.x
mul r0.x, r0.z, c0.z
  
```

In this case, the `mov` instruction does not write all elements of the array (compared to usage of `fanout` for `sum` instructions in *grouping*). But the `mov` instruction does need an additional dependency (via `fanin`) on instructions that last wrote the array element members, to ensure that they get scheduled before the `mov` in *scheduling* stage (which also serves to group the array elements for the *register assignment* stage).



Note that there would in fact be `fanout` nodes generated for each array element (although only the reachable ones will be scheduled, etc).

Shader Passes

After the frontend has generated the use-def graph of instructions, they are run through various passes which include *scheduling* and *register assignment*. Because inserting `mov` instructions after scheduling would also require inserting additional `nop` instructions (since it is too late to reschedule to try and fill the bubbles), the earlier stages try to ensure that (at least given an infinite supply of registers) that *register assignment* after *scheduling* cannot fail.

Note that we essentially have ~256 scalar registers in the architecture (although larger register usage will at some thresholds limit the number of threads which can run in parallel). And at some point we will have to deal with spilling.

Flatten

In this stage, simple if/else blocks are flattened into a single block with `phi` nodes converted into `sel` instructions. The a3xx ISA has very few predicated instructions, and we would prefer not to use branches for simple if/else.

Copy Propagation

Currently the frontend inserts `movs` in various cases, because certain categories of instructions have limitations about const regs as sources. And the CP pass simply removes all simple `movs` (ie. `src-type` is same as `dst-type`, no `abs/neg` flags, etc).

The eventual plan is to invert that, with the front-end inserting no `movs` and CP legalize things.

Grouping

In the grouping pass, instructions which need to be grouped (for `fanins`, etc) have their `left / right` neighbor pointers setup. In cases where there is a conflict (ie. one instruction cannot have two unique left or right neighbors), an additional `mov` instruction is inserted. This ensures that there is some possible valid *register assignment* at the later stages.

Depth

In the depth pass, a depth is calculated for each instruction node within it's basic block. The depth is the sum of the required cycles (delay slots needed between two instructions plus one) of each instruction plus the max depth of any of it's source instructions. (*meta* instructions don't add to the depth). As an instruction's depth is calculated, it is inserted into a per block list sorted by deepest instruction. Unreachable instructions and inputs are marked.

TODO: we should probably calculate both hard and soft depths (?) to try to coax additional instructions to fit in places where we need to use sync bits, such as after a texture fetch or SFU.

Scheduling

After the *grouping* pass, there are no more instructions to insert or remove. Start scheduling each basic block from the deepest node in the depth sorted list created by the *depth* pass, recursively trying to schedule each instruction after it's source instructions plus delay slots. Insert `nops` as required.

Register Assignment

TODO

OpenSWR

The Gallium OpenSWR driver is a high performance, highly scalable software renderer targeted towards visualization workloads. For such geometry heavy workloads there is a considerable speedup over `llvmpipe`, which is to be expected as the geometry frontend of `llvmpipe` is single threaded.

This rasterizer is x86 specific and requires AVX or above. The driver fits into the gallium framework, and reuses `gallium` for doing the TGSI to vectorized `llvm-IR` conversion of the shader kernels.

Usage

Requirements

- An x86 processor with AVX or above
- LLVM version 3.9 or later
- C++14 capable compiler

Building

To build with GNU automake, select building the swr driver at configure time, for example:

```
configure --with-gallium-drivers=swrast,swr
```

Using

On Linux, building with autotools will create a drop-in alternative for libGL.so into:

```
lib/gallium/libGL.so  
lib/gallium/libswrAVX.so  
lib/gallium/libswrAVX2.so
```

Alternatively, building with SCons will produce:

```
build/linux-x86_64/gallium/targets/libgl-xlib/libGL.so  
build/linux-x86_64/gallium/drivers/swr/libswrAVX.so  
build/linux-x86_64/gallium/drivers/swr/libswrAVX2.so
```

To use it set the LD_LIBRARY_PATH environment variable accordingly.

IMPORTANT: Mesa will default to using llvmpipe or softpipe as the default software renderer. To select the OpenSWR driver, set the GALLIUM_DRIVER environment variable appropriately:

```
GALLIUM_DRIVER=swr
```

To verify OpenSWR is being used, check to see if a message like the following is printed when the application is started:

```
SWR detected AVX2
```

FAQ

Why another software rasterizer?

Good question, given there are already three (swrast, softpipe, llvmpipe) in the Mesa tree. Two important reasons for this:

- Architecture - given our focus on scientific visualization, our workloads are much different than the typical game; we have heavy vertex load and relatively simple shaders. In addition, the core counts of machines we run on are much higher. These parameters led to design decisions much different than llvmpipe.

- Historical - Intel had developed a high performance software graphics stack for internal purposes. Later we adapted this graphics stack for use in visualization and decided to move forward with Mesa to provide a high quality API layer while at the same time benefiting from the excellent performance the software rasterizer gives us.

What's the architecture?

SWR is a tile based immediate mode renderer with a sort-free threading model which is arranged as a ring of queues. Each entry in the ring represents a draw context that contains all of the draw state and work queues. An API thread sets up each draw context and worker threads will execute both the frontend (vertex/geometry processing) and backend (fragment) work as required. The ring allows for backend threads to pull work in order. Large draws are split into chunks to allow vertex processing to happen in parallel, with the backend work pickup preserving draw ordering.

Our pipeline uses just-in-time compiled code for the fetch shader that does vertex attribute gathering and AOS to SOA conversions, the vertex shader and fragment shaders, streamout, and fragment blending. SWR core also supports geometry and compute shaders but we haven't exposed them through our driver yet. The fetch shader, streamout, and blend is built internally to swr core using LLVM directly, while for the vertex and pixel shaders we reuse bits of llvmpipe from `gallium/auxiliary/gallium` to build the kernels, which we wrap differently than llvmpipe's `auxiliary/draw` code.

What's the performance?

For the types of high-geometry workloads we're interested in, we are significantly faster than llvmpipe. This is to be expected, as llvmpipe only threads the fragment processing and not the geometry frontend. The performance advantage over llvmpipe roughly scales linearly with the number of cores available.

While our current performance is quite good, we know there is more potential in this architecture. When we switched from a prototype OpenGL driver to Mesa we regressed performance severely, some due to interface issues that need tuning, some differences in shader code generation, and some due to conformance and feature additions to the core swr. We are looking to recovering most of this performance back.

What's the conformance?

The major applications we are targeting are all based on the Visualization Toolkit (VTK), and as such our development efforts have been focused on making sure these work as best as possible. Our current code passes vtk's rendering tests with their new "OpenGL2" (really OpenGL 3.2) backend at 99%.

piglit testing shows a much lower pass rate, roughly 80% at the time of writing. Core SWR undergoes rigorous unit testing and we are quite confident in the rasterizer, and understand the areas where it currently has issues (example: line rendering is done with triangles, so doesn't match the strict line rendering rules). The majority of the piglit failures are errors in our driver layer interfacing Mesa and SWR. Fixing these issues is one of our major future development goals.

Why are you open sourcing this?

- Our customers prefer open source, and allowing them to simply download the Mesa source and enable our driver makes life much easier for them.
- The internal gallium APIs are not stable, so we'd like our driver to be visible for changes.
- It's easier to work with the Mesa community when the source we're working with can be used as reference.

What are your development plans?

- Performance - see the performance section earlier for details.
- Conformance - see the conformance section earlier for details.
- Features - core SWR has a lot of functionality we have yet to expose through our driver, such as MSAA, geometry shaders, compute shaders, and tessellation.
- AVX512 support

What is the licensing of the code?

- All code is under the normal Mesa MIT license.

Will this work on AMD?

- If using an AMD processor with AVX or AVX2, it should work though we don't have that hardware around to test. Patches if needed would be welcome.

Will this work on ARM, MIPS, POWER, <other non-x86 architecture>?

- Not without a lot of work. We make extensive use of AVX and AVX2 intrinsics in our code and the in-tree JIT creation. It is not the intention for this codebase to support non-x86 architectures.

What hardware do I need?

- Any x86 processor with at least AVX (introduced in the Intel SandyBridge and AMD Bulldozer microarchitectures in 2011) will work.
- You don't need a fire-breathing Xeon machine to work on SWR - we do day-to-day development with laptops and desktop CPUs.

Does one build work on both AVX and AVX2?

Yes. The build system creates two shared libraries, `libswrAVX.so` and `libswrAVX2.so`, and `swr_create_screen()` loads the appropriate one at runtime.

Profiling

OpenSWR contains built-in profiling which can be enabled at build time to provide insight into performance tuning.

To enable this, uncomment the following line in `rasterizer/core/knobs.h` and rebuild:

```
//#define KNOB_ENABLE_RDTSC
```

Running an application will result in a `rdtsc.txt` file being created in current working directory. This file contains profile information captured between the `KNOB_BUCKETS_START_FRAME` and `KNOB_BUCKETS_END_FRAME` (see `knobs` section).

The resulting file will contain sections for each thread with a hierarchical breakdown of the time spent in the various operations. For example:

Thread 0 (API)							
%Tot	%Par	Cycles	CPE	NumEvent	CPE2	NumEvent2	Bucket
0.00	0.00	28370	2837	10	0	0	└─>
↔APIClearRenderTarget							
0.00	41.23	11698	1169	10	0	0	─>└─>
↔APIDrawWakeAllThreads							
0.00	18.34	5202	520	10	0	0	─>└─>
↔APIGetDrawContext							
98.72	98.72	12413773688	29957	414380	0	0	APIDraw
0.36	0.36	44689364	107	414380	0	0	─>└─>
↔APIDrawWakeAllThreads							
96.36	97.62	12117951562	9747	1243140	0	0	─>└─>
↔APIGetDrawContext							
0.00	0.00	19904	995	20	0	0	APIStoreTiles
0.00	7.88	1568	78	20	0	0	─>└─>
↔APIDrawWakeAllThreads							
0.00	25.28	5032	251	20	0	0	─>└─>
↔APIGetDrawContext							
1.28	1.28	161344902	64	2486370	0	0	APIGetDrawContext
0.00	0.00	50368	2518	20	0	0	APISync
0.00	2.70	1360	68	20	0	0	─>└─>
↔APIDrawWakeAllThreads							
0.00	65.27	32876	1643	20	0	0	─>└─>
↔APIGetDrawContext							
Thread 1 (WORKER)							
%Tot	%Par	Cycles	CPE	NumEvent	CPE2	NumEvent2	Bucket
83.92	83.92	13198987522	96411	136902	0	0	FEProcessDraw
24.91	29.69	3918184840	167	23410158	0	0	─> FEFetchShader
11.17	13.31	1756972646	75	23410158	0	0	─>└─>
↔FEVertexShader							
8.89	10.59	1397902996	59	23410161	0	0	─> FEPAAssemble
19.06	22.71	2997794710	384	7803387	0	0	─>└─>
↔FEClipTriangles							
11.67	61.21	1834958176	235	7803387	0	0	─>└─>
↔FEBinTriangles							
0.00	0.00	0	0	187258	0	0	─>└─>
↔FECullZeroAreaAndBackface							
0.00	0.00	0	0	60051033	0	0	─>└─>
↔FECullBetweenCenters							
0.11	0.11	17217556	2869592	6	0	0	└─>
↔FEProcessStoreTiles							
15.97	15.97	2511392576	73665	34092	0	0	└─>
↔WorkerWorkOnFifoBE							
14.04	87.95	2208687340	9187	240408	0	0	─>└─>
↔WorkerFoundWork							
0.06	0.43	9390536	13263	708	0	0	─>└─>
↔BELoadTiles							
0.00	0.01	293020	182	1609	0	0	─> BEClear
12.63	89.94	1986508990	949	2093014	0	0	─>└─>
↔BERasterizeTriangle							
2.37	18.75	372374596	177	2093014	0	0	─>└─>
↔BETriangleSetup							
0.42	3.35	66539016	31	2093014	0	0	─>└─>
↔BESetupSetup							
0.00	0.00	0	0	21766	0	0	─>└─>
↔BETrivialReject							

1.05	8.33	165410662	79	2071248	0	0	-> <u> </u>
↔BERasterizePartial							
6.06	48.02	953847796	1260	756783	0	0	-> <u> </u>
↔BEPixelBackend							
0.20	3.30	31521202	41	756783	0	0	-> <u> </u>
↔BESetup							
0.16	2.69	25624304	33	756783	0	0	-> <u> </u>
↔BEBarycentric							
0.18	2.92	27884986	36	756783	0	0	-> <u> </u>
↔BEEarlyDepthTest							
0.19	3.20	30564174	41	744058	0	0	-> <u> </u>
↔BEPixelShader							
0.26	4.30	41058646	55	744058	0	0	-> <u> </u>
↔BEOutputMerger							
1.27	20.94	199750822	32	6054264	0	0	-> <u> </u>
↔BEEndTile							
0.33	2.34	51758160	23687	2185	0	0	-> <u> </u>
↔BEStoreTiles							
0.20	60.22	31169500	28807	1082	0	0	-> <u> </u>
↔B8G8R8A8_UNORM							
0.00	0.00	302752	302752	1	0	0	-> <u> </u>
↔WorkerWaitForThreadEvent							

Knobs

OpenSWR has a number of environment variables which control its operation, in addition to the normal Mesa and gallium controls.

KNOB_ENABLE_ASSERT_DIALOGS Type: bool **Default: true**

Use dialogs when asserts fire. Asserts are only enabled in debug builds

KNOB_SINGLE_THREADED Type: bool **Default: false**

If enabled will perform all rendering on the API thread. This is useful mainly for debugging purposes.

KNOB_DUMP_SHADER_IR Type: bool **Default: false**

Dumps shader LLVM IR at various stages of jit compilation.

KNOB_USE_GENERIC_STORETILE Type: bool **Default: false**

Always use generic function for performing StoreTile. Will be slightly slower than using optimized (jitted) path

KNOB_FAST_CLEAR Type: bool **Default: true**

Replace 3D primitive execute with a SWRClearRT operation and defer clear execution to first backend op on hostile, or hostile store

KNOB_MAX_NUMA_NODES Type: uint32_t **Default: 0**

Maximum # of NUMA-nodes per system used for worker threads 0 == ALL NUMA-nodes in the system N == Use at most N NUMA-nodes for rendering

KNOB_MAX_CORES_PER_NUMA_NODE Type: uint32_t **Default: 0**

Maximum # of cores per NUMA-node used for worker threads. 0 == ALL non-API thread cores per NUMA-node N == Use at most N cores per NUMA-node

KNOB_MAX_THREADS_PER_CORE Type: uint32_t **Default: 1**

Maximum # of (hyper)threads per physical core used for worker threads. 0 == ALL hyper-threads per core N == Use at most N hyper-threads per physical core

KNOB_MAX_WORKER_THREADS Type: uint32_t **Default: 0**

Maximum worker threads to spawn. IMPORTANT: If this is non-zero, no worker threads will be bound to specific HW threads. They will all be “floating” SW threads. In this case, the above 3 KNOBS will be ignored.

KNOB_BUCKETS_START_FRAME Type: uint32_t **Default: 1200**

Frame from when to start saving buckets data. NOTE: KNOB_ENABLE_RDTSC must be enabled in core/knobs.h for this to have an effect.

KNOB_BUCKETS_END_FRAME Type: uint32_t **Default: 1400**

Frame at which to stop saving buckets data. NOTE: KNOB_ENABLE_RDTSC must be enabled in core/knobs.h for this to have an effect.

KNOB_WORKER_SPIN_LOOP_COUNT Type: uint32_t **Default: 5000**

Number of spin-loop iterations worker threads will perform before going to sleep when waiting for work

KNOB_MAX_DRAWS_IN_FLIGHT Type: uint32_t **Default: 160**

Maximum number of draws outstanding before API thread blocks.

KNOB_MAX_PRIMS_PER_DRAW Type: uint32_t **Default: 2040**

Maximum primitives in a single Draw(). Larger primitives are split into smaller Draw calls. Should be a multiple of (3 * vectorWidth).

KNOB_MAX_TESS_PRIMS_PER_DRAW Type: uint32_t **Default: 16**

Maximum primitives in a single Draw() with tessellation enabled. Larger primitives are split into smaller Draw calls. Should be a multiple of (vectorWidth).

KNOB_MAX_FRAC_ODD_TESS_FACTOR Type: float **Default: 63.0f**

(DEBUG) Maximum tessellation factor for fractional-odd partitioning.

KNOB_MAX_FRAC_EVEN_TESS_FACTOR Type: float **Default: 64.0f**

(DEBUG) Maximum tessellation factor for fractional-even partitioning.

KNOB_MAX_INTEGER_TESS_FACTOR Type: uint32_t **Default: 64**

(DEBUG) Maximum tessellation factor for integer partitioning.

KNOB_BUCKETS_ENABLE_THREADVIZ Type: bool **Default: false**

Enable threadviz output.

KNOB_TOSS_DRAW Type: bool **Default: false**

Disable per-draw/dispatch execution

KNOB_TOSS_QUEUE_FE Type: bool **Default: false**

Stop per-draw execution at worker FE NOTE: Requires KNOB_ENABLE_TOSS_POINTS to be enabled in core/knobs.h

KNOB_TOSS_FETCH Type: bool **Default: false**

Stop per-draw execution at vertex fetch NOTE: Requires KNOB_ENABLE_TOSS_POINTS to be enabled in core/knobs.h

KNOB_TOSS_IA Type: bool **Default: false**

Stop per-draw execution at input assembler NOTE: Requires `KNOB_ENABLE_TOSS_POINTS` to be enabled in `core/knobs.h`

`KNOB_TOSS_VS` Type: bool **Default: false**

Stop per-draw execution at vertex shader NOTE: Requires `KNOB_ENABLE_TOSS_POINTS` to be enabled in `core/knobs.h`

`KNOB_TOSS_SETUP_TRIS` Type: bool **Default: false**

Stop per-draw execution at primitive setup NOTE: Requires `KNOB_ENABLE_TOSS_POINTS` to be enabled in `core/knobs.h`

`KNOB_TOSS_BIN_TRIS` Type: bool **Default: false**

Stop per-draw execution at primitive binning NOTE: Requires `KNOB_ENABLE_TOSS_POINTS` to be enabled in `core/knobs.h`

`KNOB_TOSS_RS` Type: bool **Default: false**

Stop per-draw execution at rasterizer NOTE: Requires `KNOB_ENABLE_TOSS_POINTS` to be enabled in `core/knobs.h`

GLSL GL Shading Language. The official, common high-level shader language used in GL 2.0 and above.

layer This term is used as the name of the “3rd coordinate” of a resource. 3D textures have z-slices, cube maps have faces, 1D and 2D array textures have array members (other resources do not have multiple layers). Since the functions only take one parameter no matter what type of resource is used, use the term “layer” instead of a resource type specific one.

LOD Level of Detail. Also spelled “LoD.” The value that determines when the switches between mipmaps occur during texture sampling.

MSAA Multi-Sampled Anti-Aliasing. A basic anti-aliasing technique that takes multiple samples of the depth buffer, and uses this information to smooth the edges of polygons.

NPOT Non-power-of-two. Usually applied to textures which have at least one dimension which is not a power of two.

TCL Transform, Clipping, & Lighting. The three stages of preparation in a rasterizing pipeline prior to the actual rasterization of vertices into fragments.

CHAPTER 12

Indices and tables

- `genindex`
- `modindex`
- `search`

A

ADD (TGSI opcode), 9
AND (TGSI opcode), 22
ARL (TGSI opcode), 8
ARR (TGSI opcode), 15
ATOMAND (TGSI opcode), 43
ATOMCAS (TGSI opcode), 42
ATOMIMAX (TGSI opcode), 44
ATOMIMIN (TGSI opcode), 44
ATOMOR (TGSI opcode), 43
ATOMUADD (TGSI opcode), 42
ATOMUMAX (TGSI opcode), 43
ATOMUMIN (TGSI opcode), 43
ATOMXCHG (TGSI opcode), 42
ATOMXOR (TGSI opcode), 43

B

BALLOT (TGSI opcode), 44
BARRIER (TGSI opcode), 42
BFI (TGSI opcode), 27
BGNLOOP (TGSI opcode), 28
BGNSUB (TGSI opcode), 28
BREV (TGSI opcode), 27
BRK (TGSI opcode), 29

C

CAL (TGSI opcode), 28
CASE (TGSI opcode), 29
CEIL (TGSI opcode), 18
CLOCK (TGSI opcode), 20
CMP (TGSI opcode), 16
CONT (TGSI opcode), 28
COS (TGSI opcode), 12

D

D2F (TGSI opcode), 33
D2I (TGSI opcode), 33
D2I64 (TGSI opcode), 37
D2U (TGSI opcode), 34

D2U64 (TGSI opcode), 37
DABS (TGSI opcode), 30
DADD (TGSI opcode), 30
DCEIL (TGSI opcode), 31
DDIV (TGSI opcode), 32
DDX, DDX_FINE (TGSI opcode), 12
DDY, DDY_FINE (TGSI opcode), 12
DEFAULT (TGSI opcode), 29
DFLR (TGSI opcode), 31
DFMA (TGSI opcode), 32
DFRAC (TGSI opcode), 31
DFRACEXP (TGSI opcode), 31
DIV (TGSI opcode), 16
DLDEXP (TGSI opcode), 31
DMAD (TGSI opcode), 32
DMAX (TGSI opcode), 32
DMIN (TGSI opcode), 32
DMUL (TGSI opcode), 32
DP2 (TGSI opcode), 17
DP3 (TGSI opcode), 9
DP4 (TGSI opcode), 9
DRAW_USE_LLVM (environment variable), 5
DRCP (TGSI opcode), 32
DROUND (TGSI opcode), 31
DRSQ (TGSI opcode), 33
DSEQ (TGSI opcode), 30
DSGE (TGSI opcode), 30
DSLTL (TGSI opcode), 30
DSNE (TGSI opcode), 30
DSQRT (TGSI opcode), 33
DSSG (TGSI opcode), 31
DST (TGSI opcode), 9
DTRUNC (TGSI opcode), 31

E

ELSE (TGSI opcode), 29
EMIT (TGSI opcode), 28
ENDIF (TGSI opcode), 29
ENDLOOP (TGSI opcode), 28
ENDPRIM (TGSI opcode), 28

ENDSUB (TGSI opcode), 29
ENDSWITCH (TGSI opcode), 29
EX2 (TGSI opcode), 11
EXP (TGSI opcode), 8

F

F2D (TGSI opcode), 33
F2I (TGSI opcode), 20
F2I64 (TGSI opcode), 37
F2U (TGSI opcode), 20
F2U64 (TGSI opcode), 37
FBFETCH (TGSI opcode), 41
FD_MESA_DEBUG (environment variable), 6
FLR (TGSI opcode), 11
FMA (TGSI opcode), 11
FRC (TGSI opcode), 11
FSEQ (TGSI opcode), 26
FSGE (TGSI opcode), 25
FSLT (TGSI opcode), 25
FSNE (TGSI opcode), 26

G

GALLIUM_DUMP_CPU (environment variable), 5
GALLIUM_PRINT_OPTIONS (environment variable), 5
GALLIUM_RBUG (environment variable), 5
GALLIUM_TRACE (environment variable), 5
GLSL, 123

I

I2D (TGSI opcode), 33
I2F (TGSI opcode), 20
I2I64 (TGSI opcode), 37
I642D (TGSI opcode), 38
I642F (TGSI opcode), 38
I64ABS (TGSI opcode), 34
I64DIV (TGSI opcode), 36
I64MAX (TGSI opcode), 35
I64MIN (TGSI opcode), 35
I64MOD (TGSI opcode), 37
I64NEG (TGSI opcode), 34
I64SGE (TGSI opcode), 35
I64SHR (TGSI opcode), 36
I64SLT (TGSI opcode), 35
I64SSG (TGSI opcode), 34
I915_DEBUG (environment variable), 6
I915_DUMP_CMD (environment variable), 6
I915_NO_HW (environment variable), 6
IABS (TGSI opcode), 27
IBFE (TGSI opcode), 27
IDIV (TGSI opcode), 22
IF (TGSI opcode), 29
IMAX (TGSI opcode), 23
IMIN (TGSI opcode), 23
IMSB (TGSI opcode), 28

IMUL_HI (TGSI opcode), 21
INEG (TGSI opcode), 26
INTERP_CENTROID (TGSI opcode), 30
INTERP_OFFSET (TGSI opcode), 30
INTERP_SAMPLE (TGSI opcode), 30
ISGE (TGSI opcode), 25
ISHR (TGSI opcode), 24
ISLT (TGSI opcode), 25
ISSG (TGSI opcode), 24

K

KILL (TGSI opcode), 16
KILL_IF (TGSI opcode), 16
KNOB_BUCKETS_ENABLE_THREADVIZ (environment variable), 121
KNOB_BUCKETS_END_FRAME (environment variable), 121
KNOB_BUCKETS_START_FRAME (environment variable), 121
KNOB_DUMP_SHADER_IR (environment variable), 120
KNOB_ENABLE_ASSERT_DIALOGS (environment variable), 120
KNOB_FAST_CLEAR (environment variable), 120
KNOB_MAX_CORES_PER_NUMA_NODE (environment variable), 120
KNOB_MAX_DRAWS_IN_FLIGHT (environment variable), 121
KNOB_MAX_FRAC_EVEN_TESS_FACTOR (environment variable), 121
KNOB_MAX_FRAC_ODD_TESS_FACTOR (environment variable), 121
KNOB_MAX_INTEGER_TESS_FACTOR (environment variable), 121
KNOB_MAX_NUMA_NODES (environment variable), 120
KNOB_MAX_PRIMS_PER_DRAW (environment variable), 121
KNOB_MAX_TESS_PRIMS_PER_DRAW (environment variable), 121
KNOB_MAX_THREADS_PER_CORE (environment variable), 120
KNOB_MAX_WORKER_THREADS (environment variable), 121
KNOB_SINGLE_THREADED (environment variable), 120
KNOB_TOSS_BIN_TRIS (environment variable), 122
KNOB_TOSS_DRAW (environment variable), 121
KNOB_TOSS_FETCH (environment variable), 121
KNOB_TOSS_IA (environment variable), 121
KNOB_TOSS_QUEUE_FE (environment variable), 121
KNOB_TOSS_RS (environment variable), 122
KNOB_TOSS_SETUP_TRIS (environment variable), 122

KNOB_TOSS_VS (environment variable), 122
 KNOB_USE_GENERIC_STORETILE (environment variable), 120
 KNOB_WORKER_SPIN_LOOP_COUNT (environment variable), 121

L

layer, 123
 LDEXP (TGSI opcode), 12
 LG2 (TGSI opcode), 11
 LIT (TGSI opcode), 8
 LOAD (TGSI opcode), 41
 LOD, 123
 LOD (TGSI opcode), 41
 LODQ (TGSI opcode), 19
 LOG (TGSI opcode), 9
 LP_DEBUG (environment variable), 6
 LP_NUM_THREADS (environment variable), 6
 LRP (TGSI opcode), 10
 LSB (TGSI opcode), 28

M

MAD (TGSI opcode), 10
 MAX (TGSI opcode), 10
 MEMBAR (TGSI opcode), 42
 MIN (TGSI opcode), 10
 MOD (TGSI opcode), 18
 MOV (TGSI opcode), 8
 MSA, 123
 MUL (TGSI opcode), 9

N

NOP (TGSI opcode), 29
 NOT (TGSI opcode), 22
 NPOT, 123

O

OR (TGSI opcode), 23

P

PK2H (TGSI opcode), 12
 PK2US (TGSI opcode), 12
 PK4B (TGSI opcode), 13
 PK4UB (TGSI opcode), 13
 POPC (TGSI opcode), 28
 POW (TGSI opcode), 12

R

RCP (TGSI opcode), 8
 READ_FIRST (TGSI opcode), 45
 READ_INVOC (TGSI opcode), 45
 RESQ (TGSI opcode), 41
 RET (TGSI opcode), 28

ROUND (TGSI opcode), 11
 RSQ (TGSI opcode), 8

S

SEQ (TGSI opcode), 13
 SGE (TGSI opcode), 10
 SGT (TGSI opcode), 13
 SHL (TGSI opcode), 24
 SIN (TGSI opcode), 13
 SLE (TGSI opcode), 13
 SLT (TGSI opcode), 10
 SNE (TGSI opcode), 13
 SQRT (TGSI opcode), 8
 SSG (TGSI opcode), 15
 ST_DEBUG (environment variable), 6
 STORE (TGSI opcode), 41
 SWITCH (TGSI opcode), 29

T

TCL, 123
 TEX (TGSI opcode), 14
 TEX2 (TGSI opcode), 14
 TEX_LZ (TGSI opcode), 17
 TG4 (TGSI opcode), 19
 TGSI_PRINT_SANITY (environment variable), 5
 TRUNC (TGSI opcode), 18
 TXB (TGSI opcode), 16
 TXB2 (TGSI opcode), 16
 TXD (TGSI opcode), 14
 TXF (TGSI opcode), 18
 TXL (TGSI opcode), 17
 TXL2 (TGSI opcode), 17
 TXP (TGSI opcode), 14
 TXQ (TGSI opcode), 19
 TXQS (TGSI opcode), 19

U

U2D (TGSI opcode), 33
 U2F (TGSI opcode), 20
 U2I64 (TGSI opcode), 37
 U642D (TGSI opcode), 38
 U642F (TGSI opcode), 38
 U64ADD (TGSI opcode), 34
 U64DIV (TGSI opcode), 36
 U64MAX (TGSI opcode), 36
 U64MIN (TGSI opcode), 35
 U64MOD (TGSI opcode), 36
 U64MUL (TGSI opcode), 34
 U64SEQ (TGSI opcode), 34
 U64SGE (TGSI opcode), 35
 U64SHL (TGSI opcode), 36
 U64SHR (TGSI opcode), 36
 U64SLT (TGSI opcode), 35
 U64SNE (TGSI opcode), 35

UADD (TGSI opcode), 21
UARL (TGSI opcode), 18
UBFE (TGSI opcode), 27
UCMP (TGSI opcode), 24
UDIV (TGSI opcode), 22
UIF (TGSI opcode), 29
UMAD (TGSI opcode), 21
UMAX (TGSI opcode), 23
UMIN (TGSI opcode), 23
UMOD (TGSI opcode), 22
UMSB (TGSI opcode), 28
UMUL (TGSI opcode), 21
UMUL_HI (TGSI opcode), 21
UP2H (TGSI opcode), 15
UP2US (TGSI opcode), 15
UP4B (TGSI opcode), 15
UP4UB (TGSI opcode), 15
USEQ (TGSI opcode), 26
USGE (TGSI opcode), 25
USHR (TGSI opcode), 24
USLT (TGSI opcode), 25
USNE (TGSI opcode), 26

V

VOTE_ALL (TGSI opcode), 44
VOTE_ANY (TGSI opcode), 44
VOTE_EQ (TGSI opcode), 44

X

XOR (TGSI opcode), 23