Improving buffer and image copies on Turnip

Karmjit Mahil

2025-09-30

Vulkan buffer/image functions

We'll be focusing on these:

Function	Source	Destination
vkCmdCopyBufferToImage()	VkBuffer	VkImage
vkCmdCopyImage()	VkImage	VkImage
vkCmdCopyImageToBuffer()	VkImage	VkBuffer
vkCmdCopyBuffer()	VkBuffer	VkBuffer

On Adreno

The 2D hardware, the A2D, is used to accomplish these tasks.

We still have to fallback to the 3D path in some cases:

- Sample count greater than 1 which isn't supported by the A2D.
- `NaN` canonicalization issue with float 16 formats (prior to discovering the `HALF_PRECISON` bit). E.g. expceted: `0xFFFF` but got `0x7E00`.
- Plane 0 of NV12 having a different UBWC compression.

A2D Buffer to Image Transfer

```
Setup the transfer:
  write RB_A2D_PIXEL_CNTL (8c01)
    RB_A2D_PIXEL_CNTL: 0
  - Unknown what this does - Always 0

  write RB_A2D_BLT_CNTL (8c00)
    RB_A2D_BLT_CNTL: { ROTATE = ROTATE_0 | COLOR_FORMAT = FMT6_4_4_4_4_UNORM | MASK = 0xf | IFMT = R2D_UNORM8 | RASTER_MODE = TYPE_TILED }
  write GRAS_A2D_BLT_CNTL (8400)
    GRAS_A2D_BLT_CNTL: { ROTATE = ROTATE_0 | COLOR_FORMAT = FMT6_4_4_4_4_UNORM | MASK = 0xf | IFMT = R2D_UNORM8 | RASTER_MODE = TYPE_TILED }
  write TPL1_A2D_BLT_CNTL (b2d2)
    TPL1_A2D_BLT_CNTL: { START_OFFSET_TEXELS = 0 | TYPE = A6XX_TEX_IMG_BUFFER }
  write SP_A2D_OUTPUT_INFO (a9bf)
    SP_A2D_OUTPUT_INFO: { IFMT_TYPE = OUTPUT_IFMT_2D_FLOAT | COLOR_FORMAT = FMT6_4_4_4_4_UNORM | MASK =0xf }

  - COLOR_FORMAT: HW format mostly based on the destination format and UBWC.
  - IFMT: Internal format based on the destination format.
  - MASK: Component mask. Always 0xF (0b1111)
  - RASTER_MODE: Always TYPE_TILED. TYPE_WRITER doesn't seem to be useful for Vulkan.
  - GRAS_A2D_BLT_CNTL: Always the same as RB_A2D_BLT_CNTL.
  - Always use A6XX_TEX_2D unless we're dealing with an unaligned source buffer in which case we have to use A6XX_TEX_IMG_BUFFER for lesser alignment requirements.
```


```
Setup source and destination:
Setup the transfer source:
 write TPL1_A2D_SRC_TEXTURE_INFO (b2c0)
   TPL1_A2D_SRC_TEXTURE_INFO: { COLOR_FORMAT = FMT6_4_4_4_4_UNORM | TILE_MODE = TILE6_LINEAR | COLOR_SWAP = XYZW | SAMPLES = MSAA_ONE | UNK20 | UNK22 }
   TPL1_A2D_SRC_TEXTURE_SIZE: { WIDTH = 64 | HEIGHT = 64 }
   TPL1_A2D_SRC_TEXTURE_BASE: 0x1000af000
   TPL1_A2D_SRC_TEXTURE_PITCH: { PITCH = 128 }
 write GRAS_A2D_SRC_XMIN (8401)
   GRAS_A2D_SRC_XMIN: 0
   GRAS_A2D_SRC_XMAX: 63
   GRAS_A2D_SRC_YMIN: 0
   GRAS_A2D_SRC_YMAX: 63

Setup the transfer destination:
 write RB_A2D_DEST_BUFFER_INFO (8c17)
   RB_A2D_DEST_BUFFER_INFO: { COLOR_FORMAT = FMT6_4_4_4_4_UNORM | TILE_MODE = TILE6_3 | COLOR_SWAP = WZYX | FLAGS | SAMPLES = MSAA_ONE }
   RB_A2D_DEST_BUFFER_BASE: 0x1000a7000
   RB_A2D_DEST_BUFFER_PITCH: 256
 write RB_A2D_DEST_FLAG_BUFFER_BASE (8c20)
   RB_A2D_DEST_FLAG_BUFFER_BASE: 0x1000a6000
   RB_A2D_DEST_FLAG_BUFFER_PITCH: { PITCH = 64 | ARRAY_PITCH = 1024 }
 write GRAS_A2D_DEST_TL (8405)
   GRAS_A2D_DEST_TL: { X = 0 | Y = 0 }
   GRAS_A2D_DEST_BR: { X = 63 | Y = 63 }

 - COLOR_FORMAT: HW format mostly derived from the VkFormat for a VkImage.
 - TILE_MODE: Linear, Tiled (TILE6_3), or TILE6_2 (seems to be GMEM specific)
 - COLOR_SWAP: akin to a format swizzle, where WZYX means no swizzle.
 - FLAG+FLAG_BUFFER: Buffer relating to UBWC.
```


```
Execute the transfer.
  opcode: CP_BLIT (2c) (2 dwords)
    { OP = BLIT_OP_SCALE }
```


HALF_PRECISION

Previously we had to resort to using the 3D path for these:

- `VK_FORMAT_R16_SFLOAT`
- `VK_FORMAT_R16G16_SFLOAT`
- `VK_FORMAT_R16G16B16A16_SFLOAT`

An implicit conversion to 32 bits was suspected, as the hardware would canonicalize `NaN`s, leading to a non bit-exact copy.

Turns out we were missing something. We **discovered** the `HALF_PRECISION` present since the Adreno 600 series.

Mesa MR **!35709** started making use of this bit, and dropped the 3D path used to work around the issue.

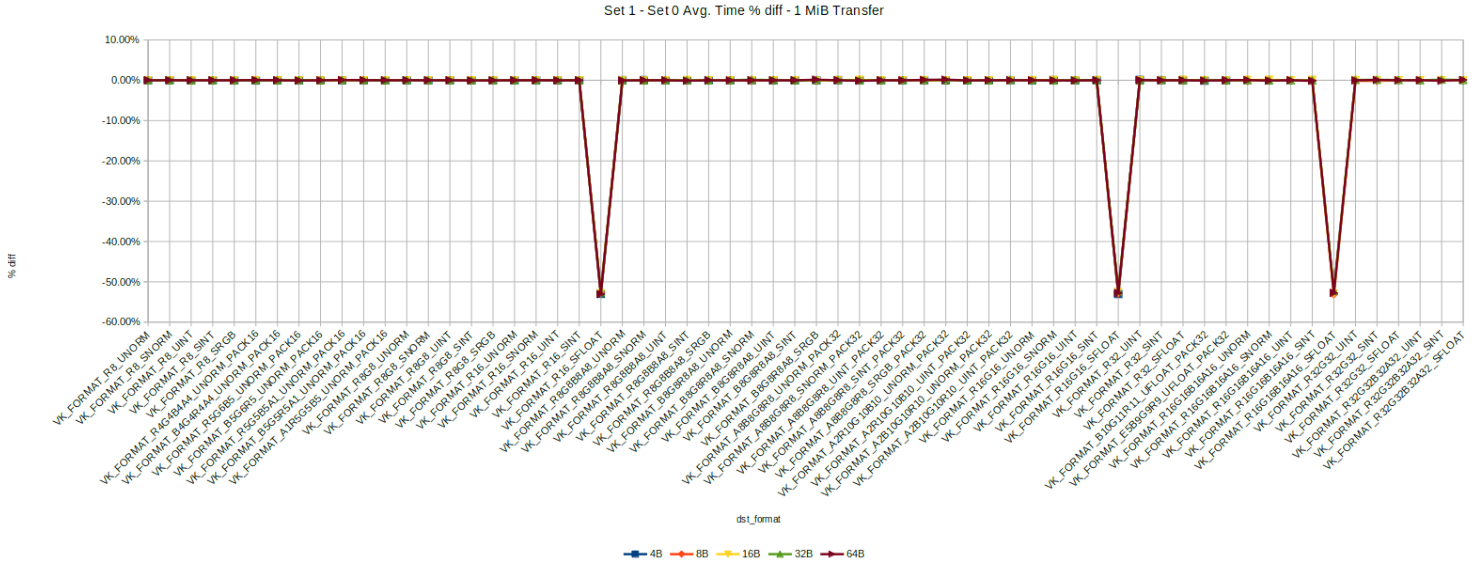
Set 1 - Set 0 Avg. Time % diff - 1 MiB Transfer
% diff
10.00%
0.00%
-10.00%
-20.00%
-30.00%
-40.00%
-50.00%
-60.00%
dst_format
4B
8B
16B
32B
64B

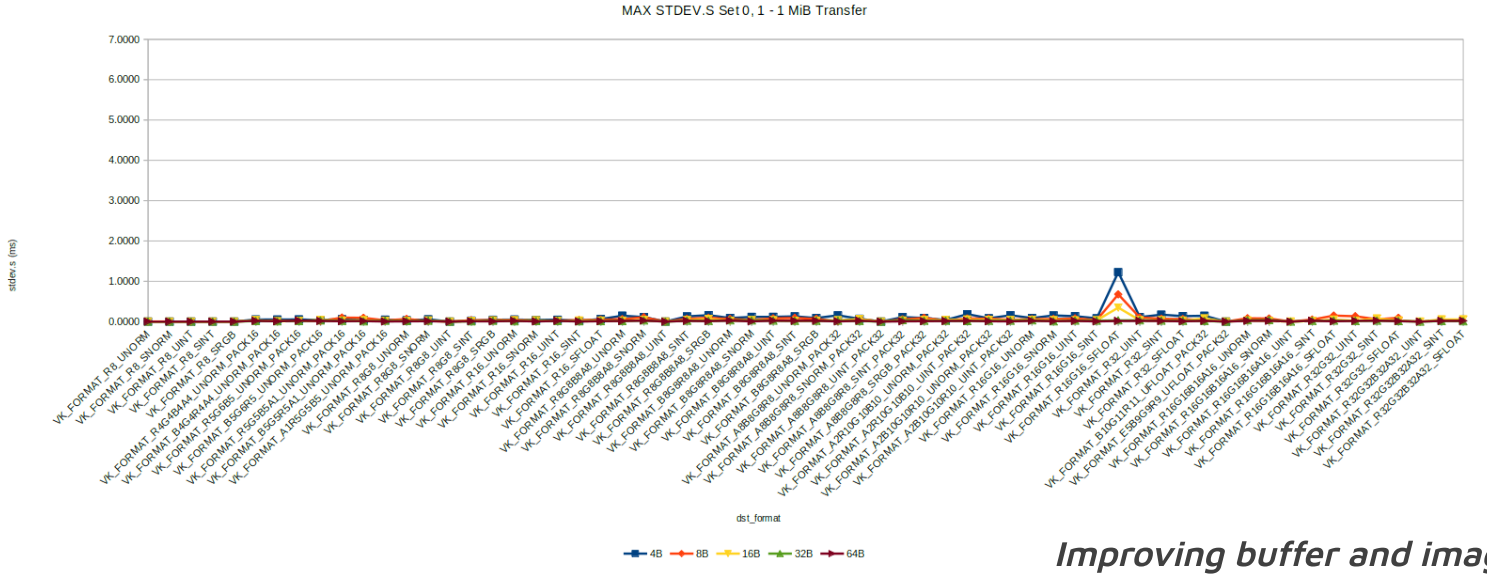
MAX STDEV.S Set 0, 1 - 1 MiB Transfer
stdev.s (ms)
7.0000
6.0000
5.0000
4.0000
3.0000
2.0000
1.0000
0.0000
dst_format
4B
8B
16B
32B
64B

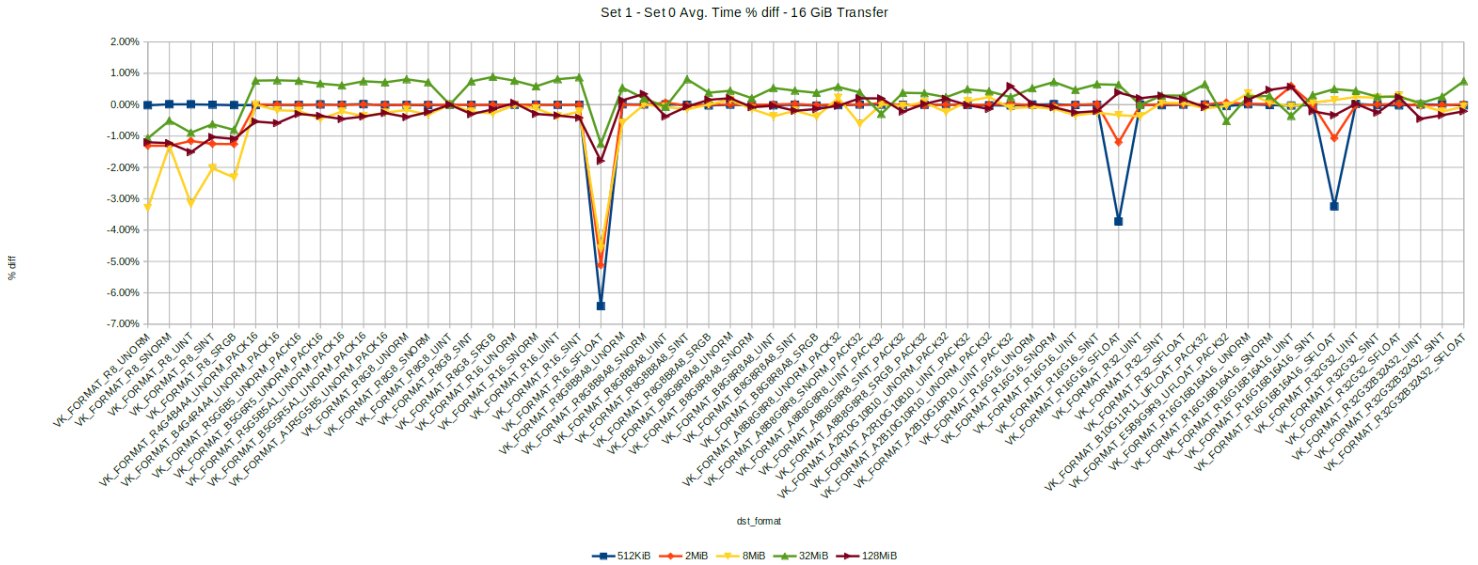
Set 1 - Set 0 Avg. Time % diff - 16 GiB Transfer
% diff
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

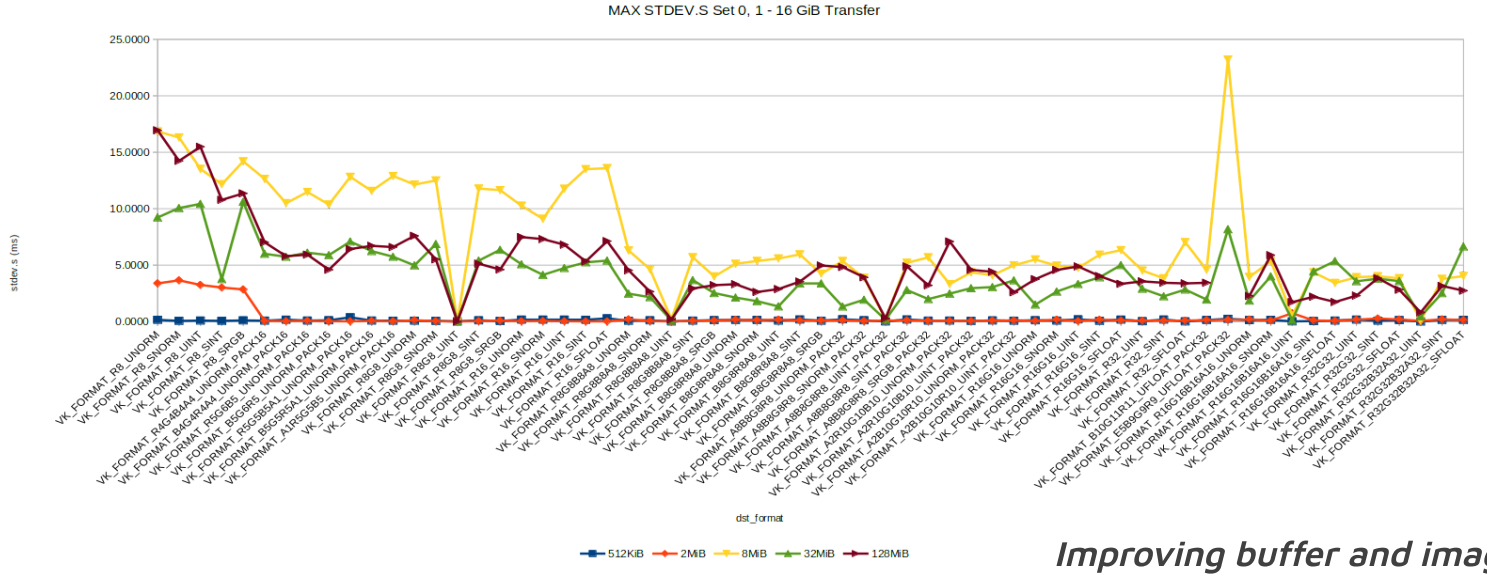
MAX STDEV.S Set 0, 1 - 16 GiB Transfer
stdev.s (ms)
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

Raw Copy

Seems to have been introduced in A750.

- `COPY` bit in `RB_A2D_BLT_CNTL`.
- `INT` used in `RB_A2D_BLT_CNTL -> IFMT`.
- `RAW_COPY` bit in `TPL1_A2D_SRC_TEXTURE_INFO`.
- `UINT` used in `SP_A2D_OUTPUT_INFO -> COLOR_FORMAT`.
- Seems like the source and destination `COLOR_FORMAT` also gets overwritten, unsure if that's necessary.

`R8G8B8A8_UNORM` -> `R8G8B8A8_UNORM`:

```
write RB_A2D_BLT_CNTL (8c00)
        RB_A2D_BLT_CNTL: { ROTATE = ROTATE_0 | COLOR_FORMAT = FMT6_8_8_8_8_UNORM | MASK = 0xf | IFMT = R2D_INT8 | RASTER_MODE = TYPE_TILED | COPY }
write TPL1_A2D_BLT_CNTL (b2d2)
        TPL1_A2D_BLT_CNTL: { RAW_COPY | START_OFFSET_TEXELS = 0 | TYPE = A6XX_TEX_2D }
write SP_A2D_OUTPUT_INFO (a9bf)
        SP_A2D_OUTPUT_INFO: { HALF_PRECISION | IFMT_TYPE = OUTPUT_IFMT_2D_UINT | COLOR_FORMAT = FMT6_8_8_8_8_UINT | MASK = 0xf }

write TPL1_A2D_SRC_TEXTURE_INFO (b2c0)
        TPL1_A2D_SRC_TEXTURE_INFO: { COLOR_FORMAT = FMT6_8_8_8_8_UNORM | TILE_MODE = TILE6_3 | COLOR_SWAP = WZYX | FLAGS | SAMPLES = MSAA_ONE | UNK20 | UNK22 }
        ...

write RB_A2D_DEST_BUFFER_INFO (8c17)
        RB_A2D_DEST_BUFFER_INFO: { COLOR_FORMAT = FMT6_8_8_8_8_UNORM | TILE_MODE = TILE6_3 | COLOR_SWAP = WZYX | FLAGS | SAMPLES = MSAA_ONE }
        ...
```


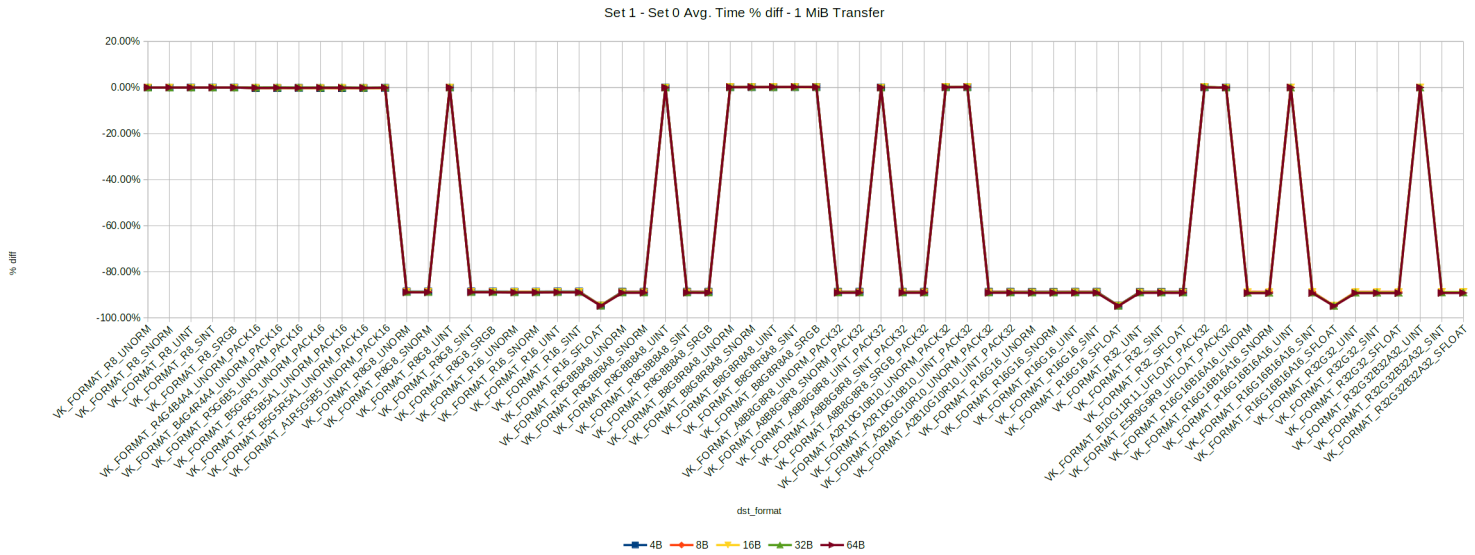
Set 1 - Set 0 Avg. Time % diff - 1 MiB Transfer
% diff
20.00%
0.00%
-20.00%
-40.00%
-60.00%
-80.00%
-100.00%
dst_format
4B
8B
16B
32B
64B

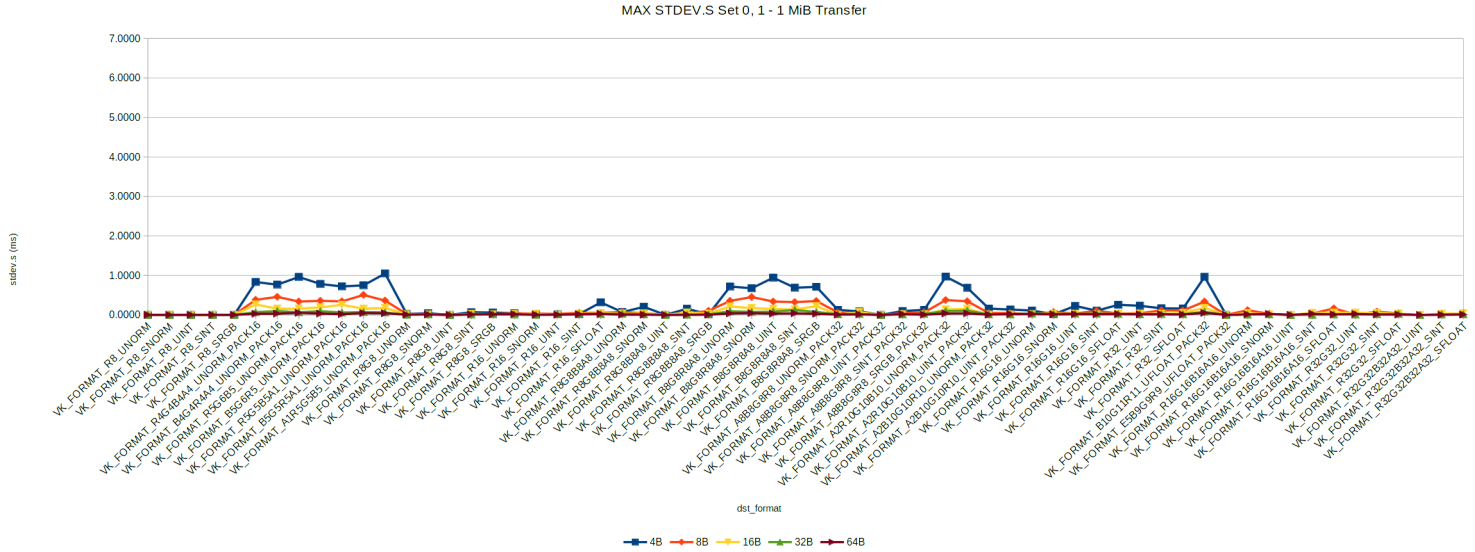
MAX STDEV.S Set 0, 1 - 1 MiB Transfer
stdev.s (ms)
7.0000
6.0000
5.0000
4.0000
3.0000
2.0000
1.0000
0.0000
dst_format
4B
8B
16B
32B
64B

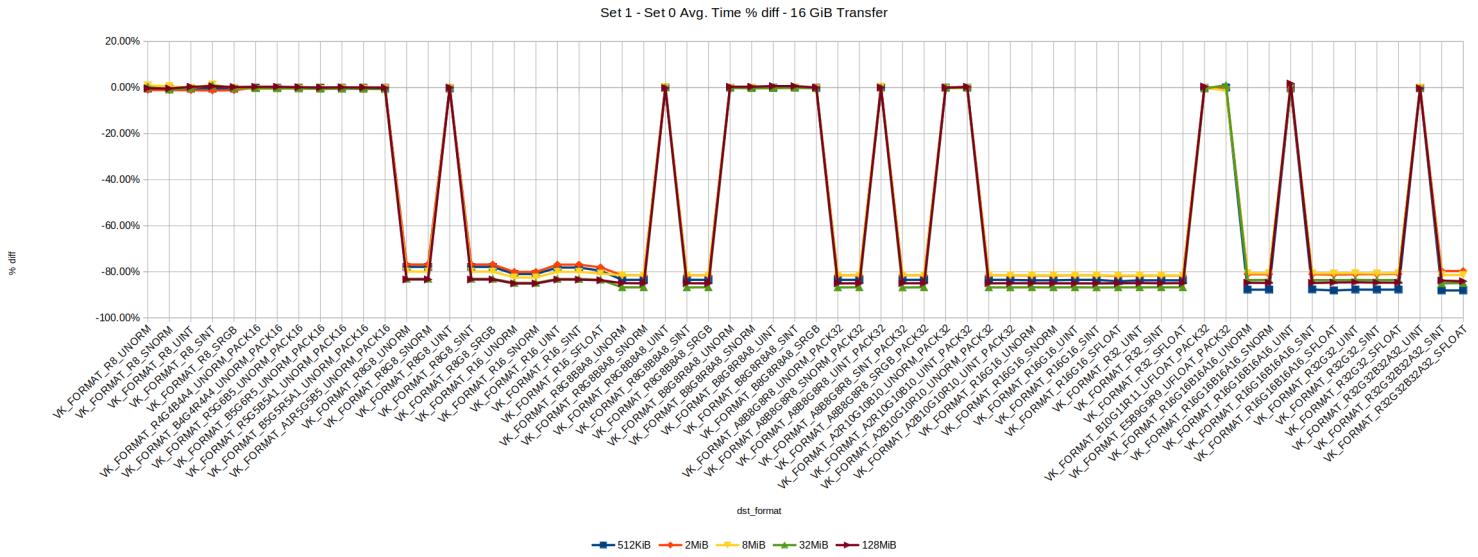
Set 1 - Set 0 Avg. Time % diff - 16 GiB Transfer
% diff
20.00%
0.00%
-20.00%
-40.00%
-60.00%
-80.00%
-100.00%
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

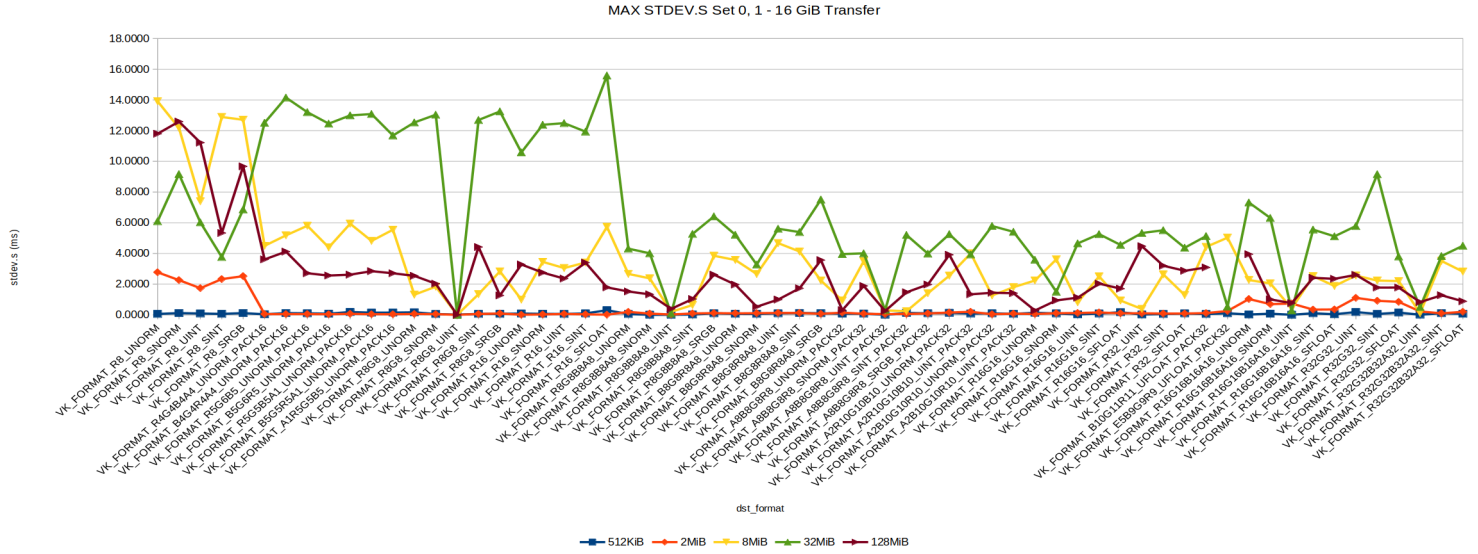
MAX STDEV.S Set 0, 1 - 16 GiB Transfer
stdev.s (ms)
18.0000
16.0000
14.0000
12.0000
10.0000
8.0000
6.0000
4.0000
2.0000
0.0000
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

Wow! That's too good to be true.

Yep, something was wrong.

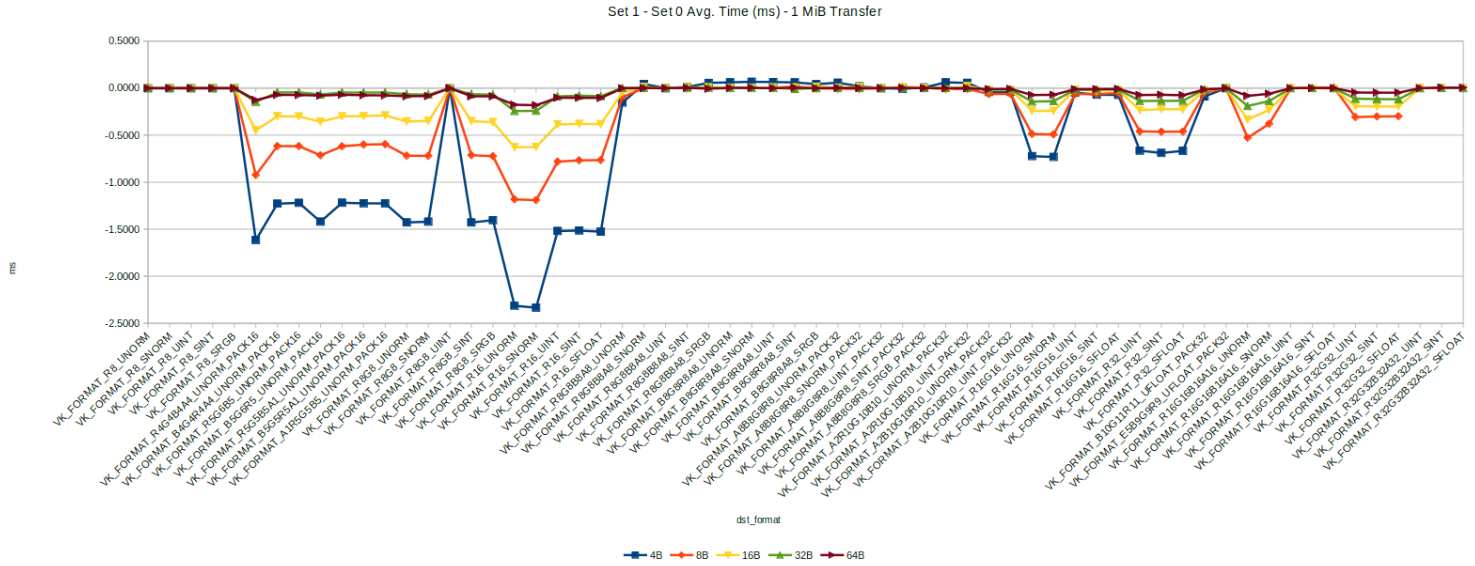
Set 1 - Set 0 Avg. Time (ms) - 1 MiB Transfer
ms
dst_format
4B
8B
16B
32B
64B

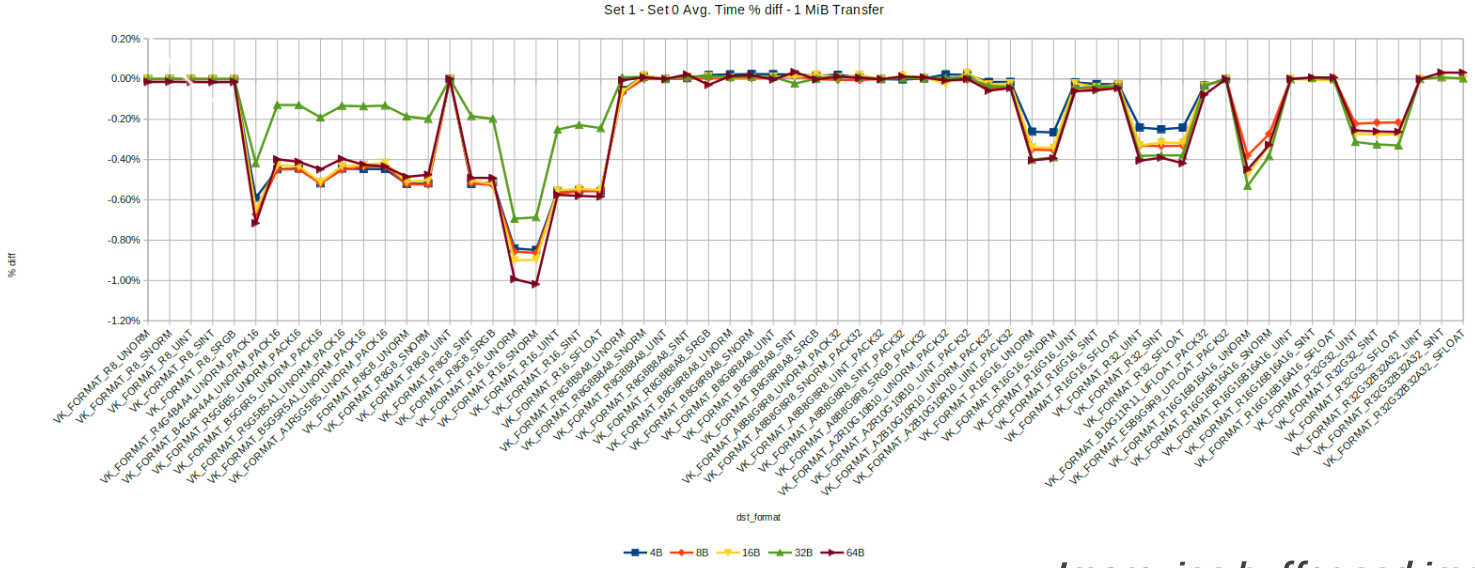
Set 1 - Set 0 Avg. Time % diff - 1 MiB Transfer
% diff
dst_format
4B
8B
16B
32B
64B

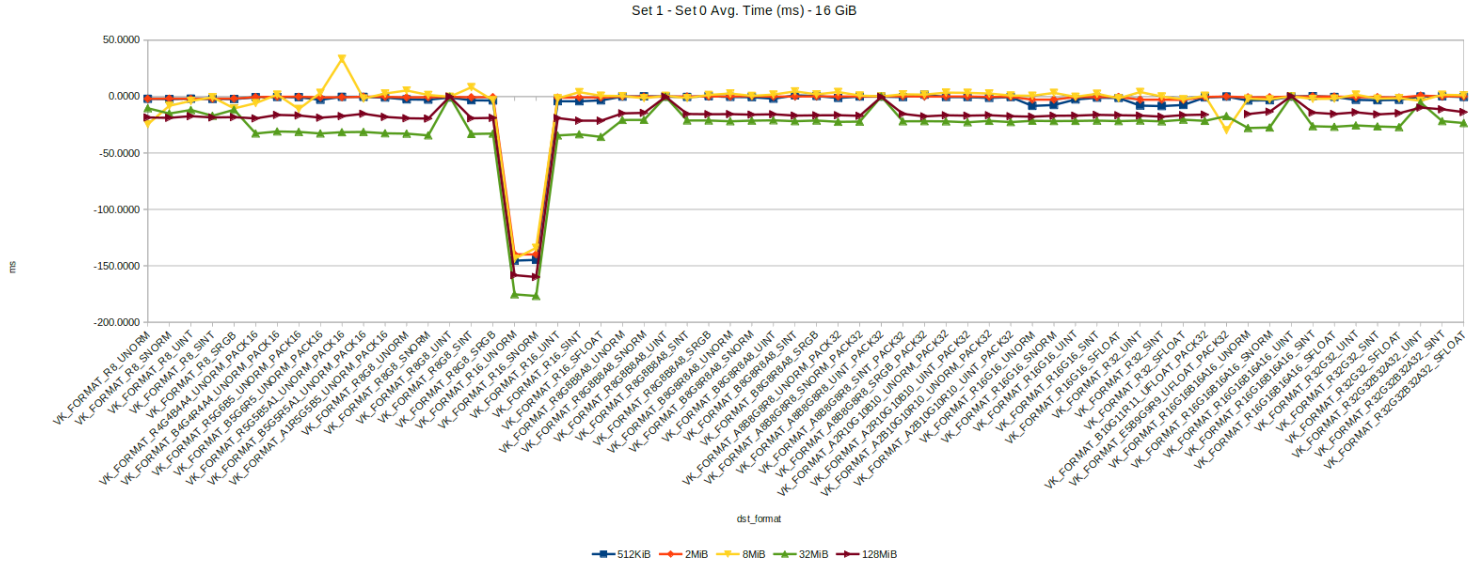
Set 1 - Set 0 Avg. Time (ms) - 16 GiB
ms
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

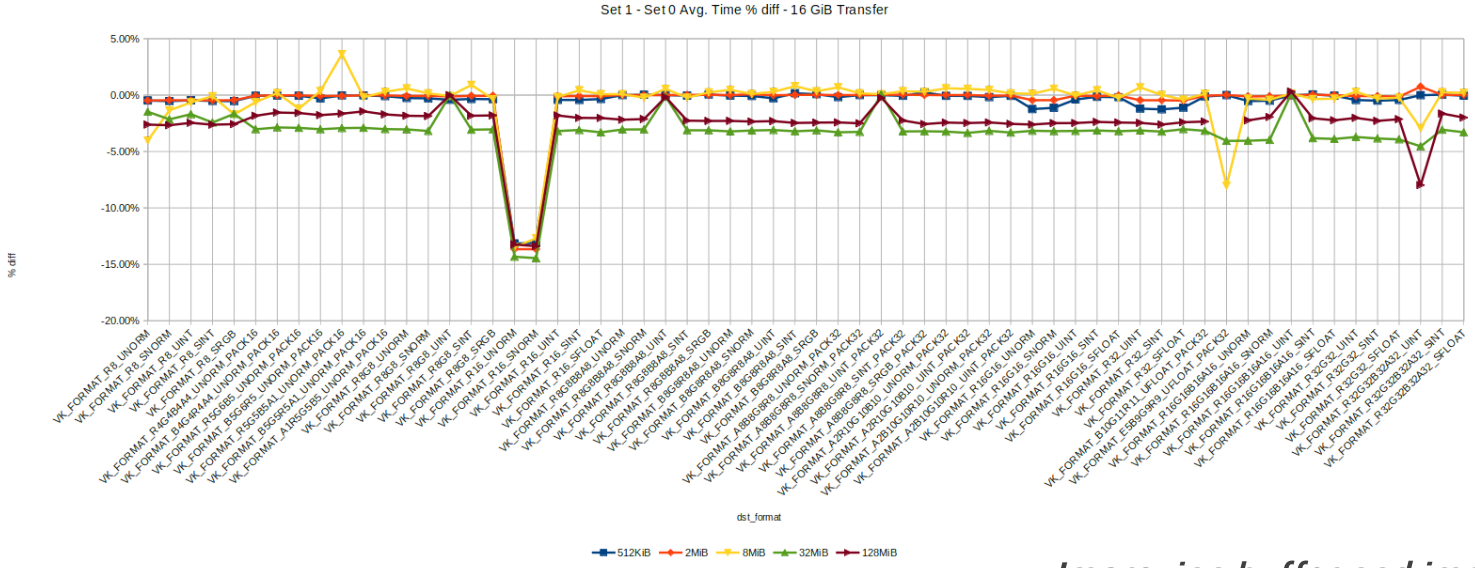
Set 1 - Set 0 Avg. Time % diff - 16 GiB Transfer
% diff
dst_format
512KiB
2MiB
8MiB
32MiB
128MiB

Got tired of looking at graphs yet!

Here's a table

Set 0															
Avg. Time / Transfer	1MiB									16GiB					
	4B	8B	16B	32B	64B	128B	256B	512B	1KiB	128KiB	512KiB	2MiB	8MiB	32MiB	128MiB
VK_FORMAT_R8_UNORM	31.0640	15.5321	7.7664	3.8836	1.9429	0.9729	0.5041	0.3133	0.1942	490.5568	545.9696	520.9930	776.7085	1077.0194	919.1565
VK_FORMAT_R8_SNORM	31.0641	15.5320	7.7664	3.8836	1.9430	0.9729	0.5040	0.3132	0.1942	491.1631	546.5480	520.1989	777.7412	1091.3565	933.2852
VK_FORMAT_R8_UINT	31.0643	15.5321	7.7664	3.8837	1.9430	0.9729	0.5041	0.3133	0.1942	492.1802	546.5615	520.1917	774.1917	1097.3262	935.4896
VK_FORMAT_R8_SINT	31.0642	15.5321	7.7664	3.8836	1.9430	0.9729	0.5041	0.3133	0.1942	487.7746	546.8095	520.2641	783.9694	1099.4875	938.3098
VK_FORMAT_R8_SRGB	31.0641	15.5323	7.7664	3.8836	1.9430	0.9729	0.5041	0.3133	0.1941	488.9044	546.6190	520.1694	779.9983	1099.7950	941.4540
VK_FORMAT_R4G4B4A4_UNORM_PACK16	31.0640	15.5322	7.7664	3.8836	1.9427	0.9723	0.4872	0.2567	0.1459	398.3234	385.8897	379.4442	639.8959	772.7255	588.8592
VK_FORMAT_B4G4R4A4_UNORM_PACK16	31.0639	15.5322	7.7664	3.8836	1.9427	0.9723	0.4872	0.2568	0.1459	391.5929	385.1233	378.8151	634.3394	772.4067	585.4321
VK_FORMAT_R5G6B5_UNORM_PACK16	31.0640	15.5322	7.7664	3.8836	1.9426	0.9722	0.4872	0.2568	0.1459	396.9756	386.8827	380.1365	638.5672	773.4560	774.7016
VK_FORMAT_B5G6R5_UNORM_PACK16	31.0641	15.5321	7.7664	3.8836	1.9427	0.9723	0.4872	0.2569	0.1459	398.3790	387.3002	380.4963	633.6237	771.5116	772.1277
VK_FORMAT_R5G5B5A1_UNORM_PACK16	31.0638	15.5320	7.7664	3.8836	1.9426	0.9723	0.4872	0.2568	0.1459	394.9953	386.0385	339.8869	620.2707	589.4318	735.4613
VK_FORMAT_B5G5R5A1_UNORM_PACK16	31.0640	15.5321	7.7664	3.8836	1.9426	0.9723	0.4872	0.2568	0.1459	394.9964	385.9554	295.1178	581.3880	757.4355	777.4042
VK_FORMAT_A1R5G5B5_UNORM_PACK16	31.0642	15.5321	7.7664	3.8836	1.9426	0.9723	0.4872	0.2568	0.1459	399.8698	339.7562	366.4480	351.2881	710.4247	760.6067
VK_FORMAT_R8G8_UNORM	31.0639	15.5322	7.7665	3.8836	1.9427	0.9723	0.4872	0.2568	0.1459	405.3013	306.7798	297.8062	458.3913	645.0249	651.8841
VK_FORMAT_R8G8_SNORM	31.0640	15.5321	7.7664	3.8836	1.9426	0.9723	0.4872	0.2568	0.1460	405.3270	306.5173	297.8378	458.1295	647.0073	651.4298
VK_FORMAT_R8G8_UINT	31.0642	15.5321	7.7664	3.8837	1.9427	0.9723	0.4872	0.2569	0.1459	405.5821	306.8350	297.6722	460.9095	647.0229	649.3826
VK_FORMAT_R8G8_SINT	31.0642	15.5322	7.7665	3.8836	1.9427	0.9723	0.4872	0.2568	0.1459	407.2299	306.4993	298.1558	456.7470	645.0172	649.7256
VK_FORMAT_R8G8_SRGB	31.0640	15.5322	7.7664	3.8837	1.9427	0.9723	0.4872	0.2568	0.1460	405.3601	306.4549	297.8302	458.0999	644.4163	650.6576
VK_FORMAT_R16_UNORM	31.0641	15.5327	7.7667	3.8839	1.9432	0.9732	0.5051	0.3045	0.1862	460.6756	419.0644	408.7538	603.4853	774.9791	781.0233
VK_FORMAT_R16_SNORM	31.0642	15.5324	7.7668	3.8839	1.9432	0.9732	0.5053	0.3045	0.1863	461.7034	419.0619	408.7399	601.8040	777.3703	781.6036
VK_FORMAT_R16_UINT	31.0640	15.5323	7.7665	3.8836	1.9427	0.9725	0.4875	0.2674	0.1560	396.4313	313.0557	347.5513	593.5239	770.1108	774.2870
VK_FORMAT_R16_SINT	31.0639	15.5323	7.7665	3.8836	1.9428	0.9725	0.4875	0.2676	0.1559	396.0708	312.6789	347.4119	600.5733	774.3710	780.3791
VK_FORMAT_R16_SFLOAT	31.0640	15.5321	7.7664	3.8836	1.9428	0.9725	0.4875	0.2675	0.1560	396.5400	312.9687	347.7192	589.8664	771.5277	775.2672
VK_FORMAT_R8G8B8A8_UNORM	31.0641	15.5322	7.7665	3.8836	1.9425	0.9719	0.4868	0.2446	0.1245	282.8212	286.9747	296.0009	299.2406	448.7891	458.7730
VK_FORMAT_R8G8B8A8_SNORM	31.0643	15.5321	7.7664	3.8835	1.9424	0.9720	0.4867	0.2446	0.1245	287.7820	287.4357	296.4341	299.3737	449.7971	458.5804
VK_FORMAT_R8G8B8A8_UINT	31.0639	15.5320	7.7664	3.8835	1.9424	0.9720	0.4867	0.2446	0.1246	287.0116	287.6349	296.6890	299.6292	448.6193	458.7108
VK_FORMAT_R8G8B8A8_SINT	31.0640	15.5321	7.7664	3.8836	1.9424	0.9719	0.4867	0.2446	0.1245	287.6286	288.2261	297.5125	299.6061	448.2470	457.3677
VK_FORMAT_R8G8B8A8_SRGB	31.0641	15.5324	7.7664	3.8836	1.9424	0.9719	0.4867	0.2446	0.1245	287.3852	287.0276	296.4852	299.8227	451.3247	460.6554
VK_FORMAT_B8G8R8A8_UNORM	31.0639	15.5322	7.7664	3.8836	1.9424	0.9719	0.4867	0.2446	0.1245	287.5882	287.0890	296.7953	299.2929	449.3909	457.6375
VK_FORMAT_B8G8R8A8_SNORM	31.0640	15.5320	7.7664	3.8836	1.9425	0.9719	0.4867	0.2446	0.1245	286.0432	287.0597	296.2085	299.5716	447.4572	457.5975
VK_FORMAT_B8G8R8A8_UINT	31.0641	15.5323	7.7664	3.8835	1.9424	0.9719	0.4867	0.2446	0.1245	287.8162	287.3891	296.1350	299.8109	448.8153	457.0829
VK_FORMAT_B8G8R8A8_SINT	31.0641	15.5322	7.7664	3.8836	1.9425	0.9719	0.4867	0.2446	0.1246	285.9975	287.9374	296.5663	299.7532	449.8639	458.1305
VK_FORMAT_B8G8R8A8_SRGB	31.0639	15.5324	7.7665	3.8837	1.9424	0.9719	0.4868	0.2446	0.1245	285.3791	286.8558	295.6961	300.0146	448.5461	457.9832
VK_FORMAT_A8B8G8R8_UNORM_PACK32	31.0639	15.5322	7.7664	3.8836	1.9424	0.9719	0.4868	0.2446	0.1245	285.8586	286.2479	297.2270	299.2676	448.0109	459.6933
VK_FORMAT_A8B8G8R8_SNORM_PACK32	31.0639	15.5321	7.7664	3.8836	1.9425	0.9719	0.4868	0.2446	0.1246	283.3270	285.8444	296.5747	299.6974	448.0743	457.2787
VK_FORMAT_A8B8G8R8_UINT_PACK32	31.0643	15.5325	7.7664	3.8836	1.9424	0.9719	0.4867	0.2446	0.1245	286.4077	286.9478	296.0867	298.7118	450.1341	457.4819
VK_FORMAT_A8B8G8R8_SINT_PACK32	31.0642	15.5322	7.7664	3.8836	1.9424	0.9719	0.4867	0.2446	0.1245	283.4157	286.6408	296.1850	299.5171	448.6629	457.9683
VK_FORMAT_A8B8G8R8_SRGB_PACK32	31.0641	15.5321	7.7664	3.8836	1.9425	0.9720	0.4868	0.2446	0.1245	285.0571	286.7673	296.7792	300.2768	449.9206	458.7667
VK_FORMAT_A2R10G10B10_UNORM_PACK32	31.0641	15.5321	7.7664	3.8836	1.9426	0.9721	0.4869	0.2448	0.1302	285.5032	294.5107	301.6842	299.0707	441.9302	456.9738
VK_FORMAT_A2R10G10B10_UINT_PACK32	31.0641	15.5322	7.7665	3.8836	1.9426	0.9720	0.4869	0.2448	0.1302	286.3530	294.6674	301.7694	299.7411	441.8451	456.7015
VK_FORMAT_A2B10G10R10_UNORM_PACK32	31.0641	15.5323	7.7664	3.8836	1.9426	0.9721	0.4869	0.2449	0.1302	285.1402	295.1270	300.9697	299.8142	450.9626	459.4404
VK_FORMAT_A2B10G10R10_UINT_PACK32	31.0643	15.5321	7.7664	3.8836	1.9426	0.9721	0.4869	0.2449	0.1301	284.3376	294.8654	301.6741	299.2926	449.8596	456.7190
VK_FORMAT_R16G16_UNORM	31.0641	15.5324	7.7667	3.8839	1.9429	0.9725	0.4874	0.2597	0.1495	311.1997	325.4433	301.7332	296.6454	443.2972	449.6157
VK_FORMAT_R16G16_SNORM	31.0643	15.5325	7.7667	3.8839	1.9429	0.9726	0.4875	0.2598	0.1497	312.4833	325.2794	301.7480	296.8895	441.2288	447.4345
VK_FORMAT_R16G16_UINT	31.0640	15.5320	7.7664	3.8836	1.9426	0.9720	0.4869	0.2449	0.1302	284.6482	284.5921	295.8831	299.0362	442.8103	451.7877
VK_FORMAT_R16G16_SINT	31.0640	15.5322	7.7664	3.8836	1.9426	0.9721	0.4869	0.2449	0.1301	284.0720	285.1017	295.6914	298.5541	437.5465	448.0142
VK_FORMAT_R16G16_SFLOAT	31.0640	15.5324	7.7665	3.8836	1.9426	0.9721	0.4869	0.2449	0.1302	286.4791	284.7865	296.0938	298.4999	442.4913	449.5331
VK_FORMAT_R32_UINT	31.0644	15.5325	7.7667	3.8839	1.9429	0.9725	0.4874	0.2587	0.1490	309.7895	325.3467	301.8269	296.9022	442.0904	449.3085
VK_FORMAT_R32_SINT	31.0642	15.5324	7.7667	3.8839	1.9430	0.9725	0.4874	0.2588	0.1490	311.9538	325.9670	302.4669	296.6993	443.9686	448.9185
VK_FORMAT_R32_SFLOAT	31.0643	15.5324	7.7666	3.8839	1.9430	0.9725	0.4874	0.2587	0.1491	311.0632	325.6610	302.3176	296.8058	444.2167	449.2521
VK_FORMAT_B10G11R11_UFLOAT_PACK32	31.0640	15.5322	7.7664	3.8836	1.9426	0.9721	0.4869	0.2449	0.1302	286.2417	293.7774	301.5615	337.0344	449.6064	459.6660
VK_FORMAT_E5B9G9R9_UFLOAT_PACK32	31.0644	15.5323	7.7668	3.8839	1.9429	0.9725	0.4874	0.2586	0.1489	305.7982	378.8646	368.0606	680.9350	802.3032	797.7667
VK_FORMAT_R16G16B16A16_UNORM	#N/A	15.5323	7.7667	3.8839	1.9428	0.9723	0.4871	0.2451	0.1335	354.7968	333.5201	317.2837	318.9849	458.2741	441.5656
VK_FORMAT_R16G16B16A16_SNORM	#N/A	15.5325	7.7667	3.8839	1.9428	0.9723	0.4871	0.2451	0.1334	357.3045	332.9956	318.4473	319.2799	459.6654	442.8376
VK_FORMAT_R16G16B16A16_UINT	#N/A	15.5325	7.7665	3.8835	1.9424	0.9718	0.4866	0.2443	0.1232	345.1678	331.0779	318.0020	318.6520	459.2296	443.3169
VK_FORMAT_R16G16B16A16_SINT	#N/A	15.5321	7.7665	3.8836	1.9424	0.9719	0.4866	0.2443	0.1232	345.2863	334.1853	317.7875	319.2137	458.3290	442.5131
VK_FORMAT_R16G16B16A16_SFLOAT	#N/A	15.5323	7.7664	3.8836	1.9424	0.9718	0.4866	0.2443	0.1232	347.2361	335.0440	319.0809	320.8202	468.2718	443.7742
VK_FORMAT_R32G32_UINT	#N/A	15.5325	7.7666	3.8839	1.9427	0.9722	0.4870	0.2448	0.1261	354.6209	332.6457	318.0095	319.3542	456.5614	442.2408
VK_FORMAT_R32G32_SINT	#N/A	15.5323	7.7666	3.8839	1.9427	0.9722	0.4870	0.2448	0.1261	356.1847	333.1014	318.5604	319.3332	458.4851	442.0305
VK_FORMAT_R32G32_SFLOAT	#N/A	15.5326	7.7667	3.8839	1.9427	0.9722	0.4870	0.2448	0.1261	354.3011	333.0595	318.0221	318.3156	458.6576	441.9156
VK_FORMAT_R32G32B32A32_UINT	#N/A	#N/A	7.7666	3.8838	1.9427	0.9720	0.4868	0.2444	0.1234	330.2286	328.9527	299.8719	300.3736	441.8532	436.8818
VK_FORMAT_R32G32B32A32_SINT	#N/A	#N/A	7.7666	3.8838	1.9427	0.9720	0.4868	0.2444	0.1234	331.7535	330.8332	299.4377	300.2069	439.1075	435.6886
VK_FORMAT_R32G32B32A32_SFLOAT	#N/A	#N/A	7.7666	3.8838	1.9427	0.9720	0.4868	0.2444	0.1234	330.8337	328.3990	300.2775	300.3520	440.6635	435.4826

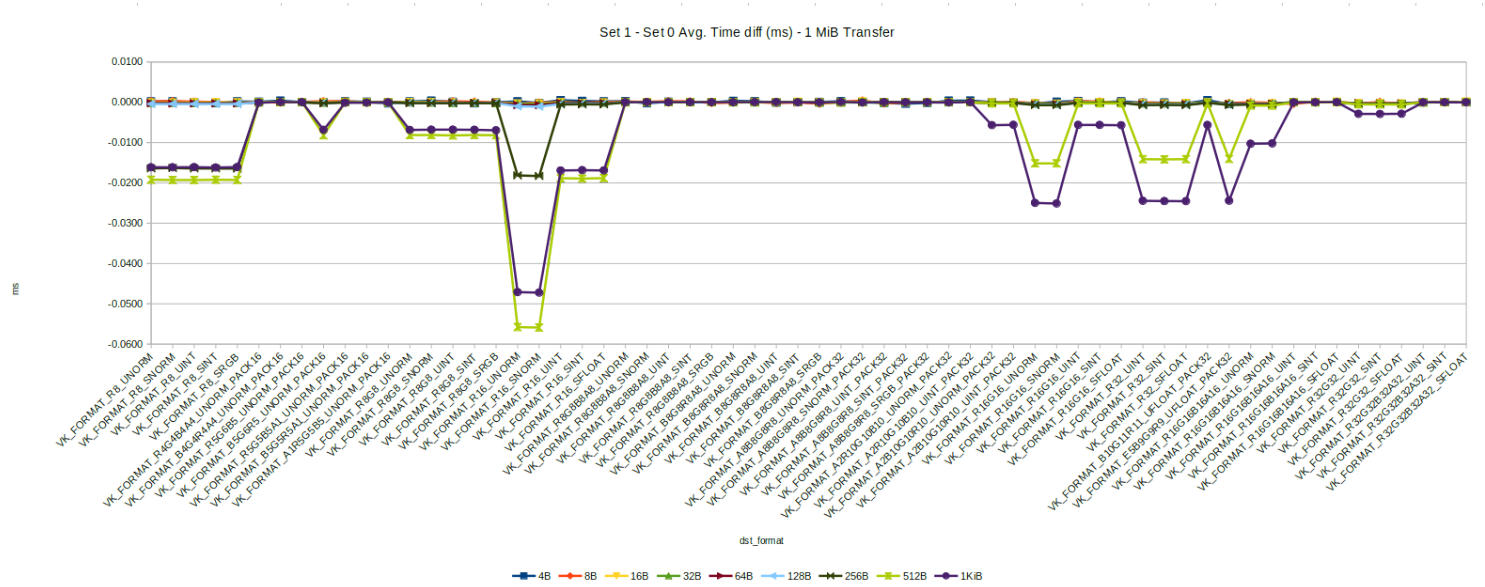
Set 1 - Set 0 Avg. Time diff (ms) - 1 MiB Transfer
ms
0.0100
0.0000
-0.0100
-0.0200
-0.0300
-0.0400
-0.0500
-0.0600
dst_format
4B
8B
16B
32B
64B
128B
256B
512B
1KiB

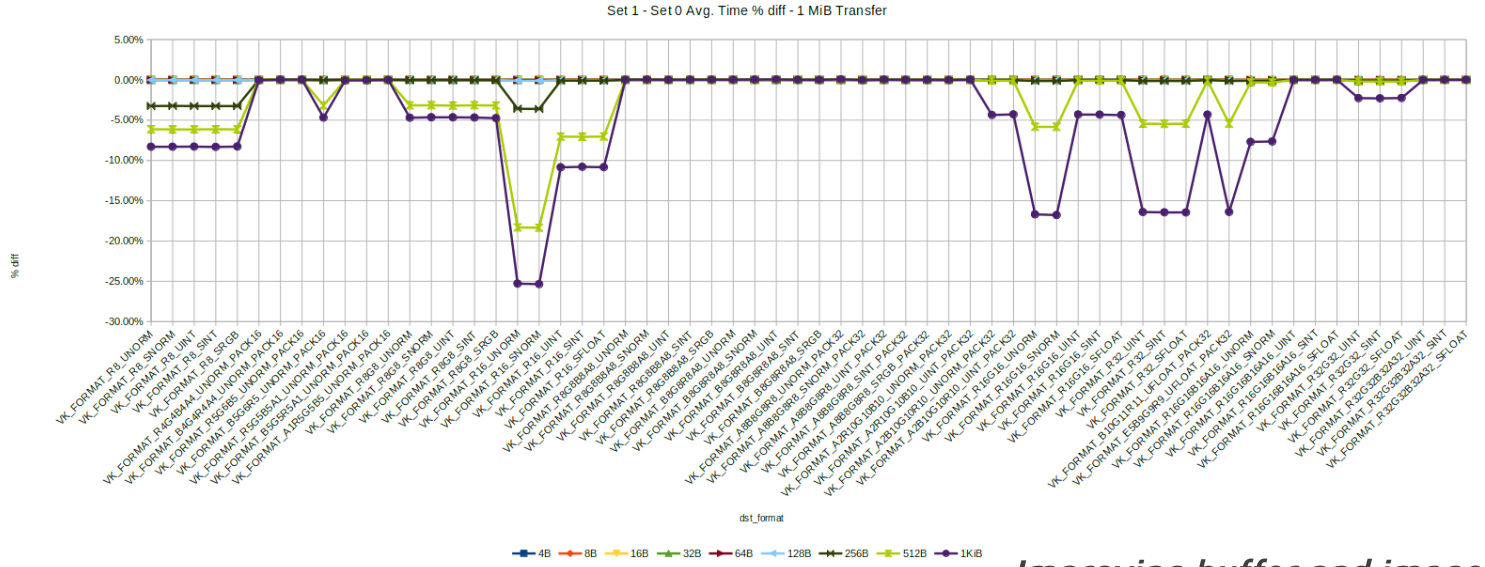
Set 1 - Set 0 Avg. Time % diff - 1 MiB Transfer
% diff
5.00%
0.00%
-5.00%
-10.00%
-15.00%
-20.00%
-25.00%
-30.00%
dst_format
4B
8B
16B
32B
64B
128B
256B
512B
1KiB

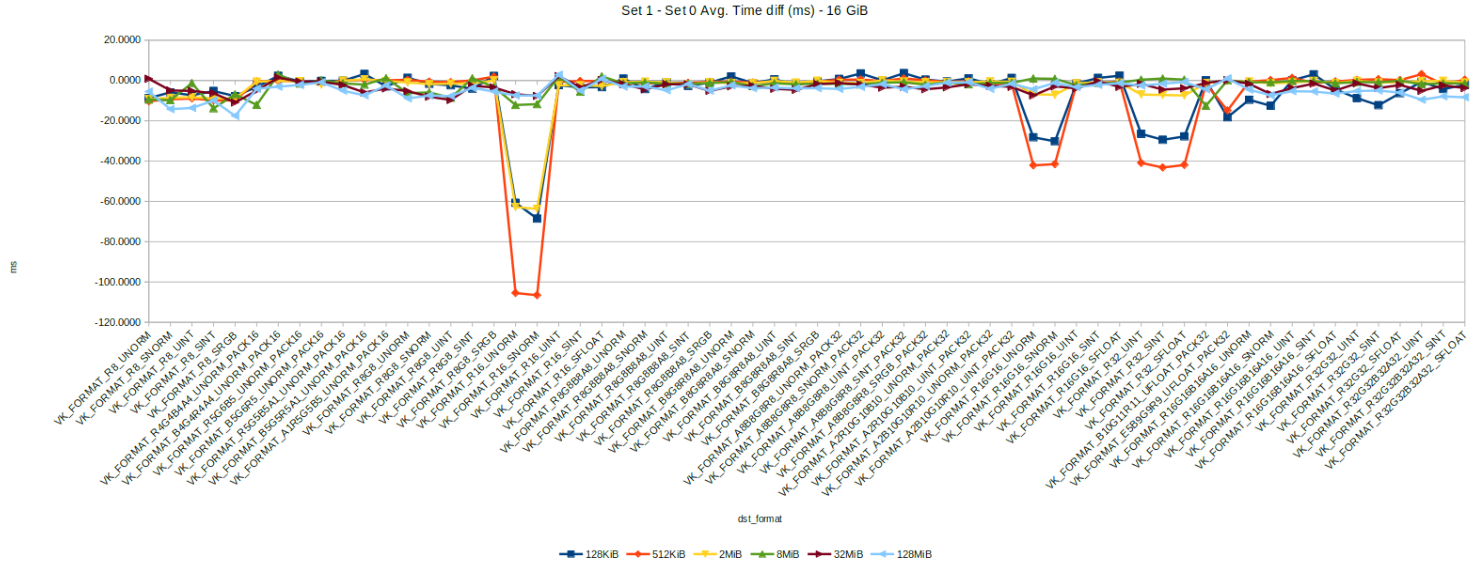
Set 1 - Set 0 Avg. Time diff (ms) - 16 GiB
ms
20.0000
0.0000
-20.0000
-40.0000
-60.0000
-80.0000
-100.0000
-120.0000
dst_format
128KiB
512KiB
2MiB
8MiB
32MiB
128MiB

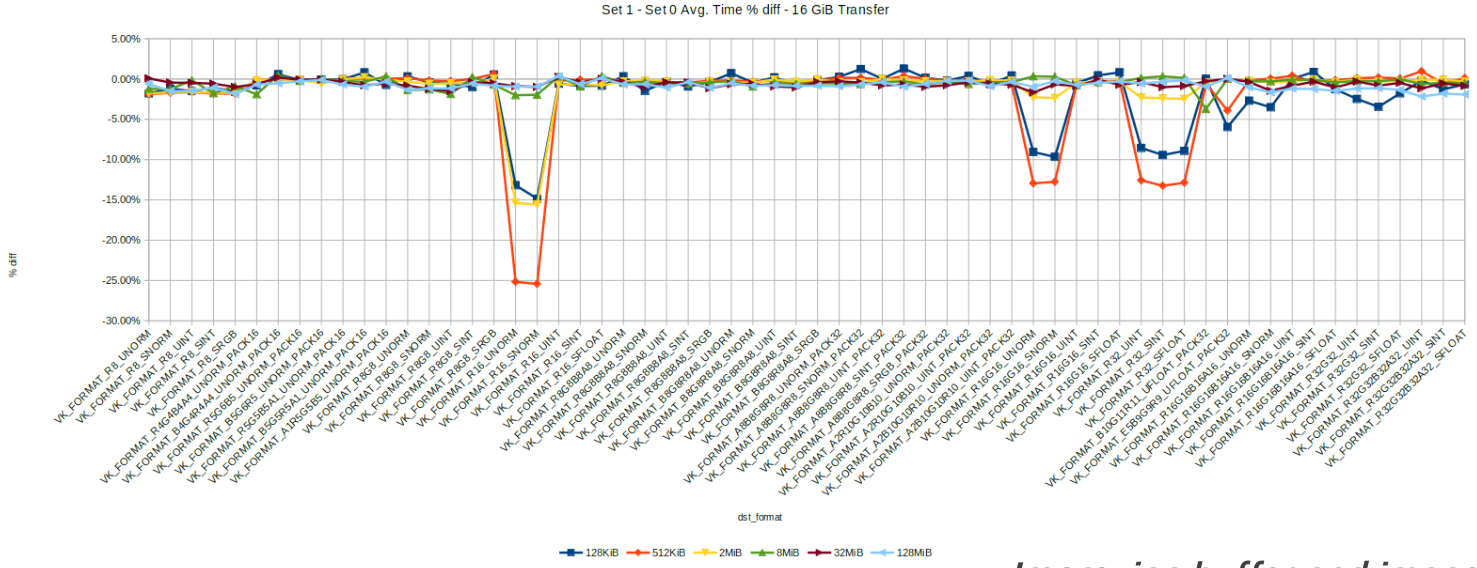
Set 1 - Set 0 Avg. Time % diff - 16 GiB Transfer
% diff
5.00%
0.00%
-5.00%
-10.00%
-15.00%
-20.00%
-25.00%
-30.00%
dst_format
128KiB
512KiB
2MiB
8MiB
32MiB
128MiB

Some other things to look into:

- Somehow with increasing chunk_size for the 16 GiB transfer the gpu time also increased.
- Certain chunk_size with specific formats really benefited from raw copy. How come not all cases?
- Raw copy could be used with depth stencil formats.
- In some cases we could be using the A2D for formats like NV12.
- In place compression/decompression.


igalia