Dual Loop versus Single, the facts - **by Gabriel Rouchon, 6-30-10**

The following series of tests represent Part II of the white paper published here

In this article, the question we wanted to answer was:

Are dual (dedicated) loops better than single loops?

Equipment:

- Loop 1: MCR320 Drive radiator, with built-in MCP355 pump, and an Apogee XT waterblock; the loop uses ½ lines. The fans are Gentle Typhoon's (D1225C12B5AP-15) running at 1850 rpm and rated at 28 dB.
- Loop 2: MCR220 Drive radiator, with built-in MCP355 pump,Gentle Typhoon fans (same model) and 1/2" lines.
- All components are connected to the loop with CPC quick-disconnect fittings; they are fairly restrictive, but the time they save in changing setups overshadows any other considerations.
- The CPU is an early Ci7 920, Revision C0/C1 stepping 4.
- The Graphic cards are (2) EVGA GTX470 FTW
- The Motherboard is a Gigabyte EX58-UD3R, and the OS is Windows 7 Ultimate 64 Bit.

Methodology:

- The CPU maximum stable overclock was well established, since we have been using this same 920 ever since its introduction. It is 4095 Mhz (Intel Turbo mode on, and HT enabled), at 1.424v (after droop).
- The GPU's maximum stable overclock was established in the previous graphics tests using Furmark in extreme burn mode at 1920x1050 for a minimum of two hours, and further validated by running 3D Marks Vantage.
- Max stable overclock for 2 cards in SLI was 825 MHz core and 1000 MHz memory, @1.087 Volts.

We conducted two series of tests, reflecting the following hardware configurations:

- Series I, with the Ci7 920@ 4.1Gb and (2) GTX470 FTW in SLI @825/1000, and
- Series II, with the Ci7 920@ 4.1Gb and (1) GTX470 FTW @825/1000.

Within each series, we tested three cooling loop configurations:

- Dual Loop with MCR320 Drive dedicated to cooling the CPU, and MCR220 Drive dedicated to cooling the GPU(s)
- Dual Loop with MCR220 Drive dedicated to cooling the CPU, and MCR320 Drive dedicated to cooling the GPU(s)
- Single loop combining MCR320 Drive, MCR220 Drive, CPU, and parallelly linked GPU's, in series

Within each loop configuration, we simulated three load scenarios consisted in:

- CPU load tests: In order to maintain consistency with previous test data, we ran our usual 8 instances of BurnK6. We logged the temperature results at 2 seconds intervals using CoreTemps. The average temperature of the 4 cores is reported.
- GPU load tests: We used Furmark in extreme burning mode, windowed in 1920x1050, post processing off to enable 100% load to both GPU's in SLI configuration, and logged the temperature results at 2 seconds intervals with GPUZ.
- Combined CPU + GPU load test: We used (7) instances of Burn K6 + Furmark in extreme burning mode - The combination of these two benchmarks placed ~100% load on all four CPU cores, and a load on both GPU cores varying between 98 and 100%.
- Graphics cards were hydraulically linked in parallel, as a result of the findings outlined in part I of this article.

The test results are compiled and summarized in two groups: Temps under load in typical computer use, and Temps under load in Extreme computer use.

- Typical computer use reflects the assumption that at the time of this writing, CPU maximum load and GPU maximum load are in the vast majority of the cases mutually exclusive of each other. In other words, the majority of games placing a heavy load on the GPU's use very few CPU resources, whereas the majority of CPU intensive applications use very little GPU resources.
- The extreme computer use scenario reflects the currently rare occurrences where both CPU and GPU(s) are under maximum load.
- Comparing these two groups provides an insight on the respective device load ratios relative to the heat exchangers and may provide guidance for further system configurations.

Environmental Temperature recording:

- Air temperature: each fan was equipped with a type T Thermocouples (accurate at +/- 0.1c) at the inlet, and the average of the 3 values is reported.
- Coolant temperature was measured at the radiator inlet with a Type T thermocouple (accurate at +/- 0.1°C).

<u>I. Series 1 test results, CPU + SLI configuration:</u>

SERIES I - Dedicated loops: MCR320-Drive to (1) CPU; MCR220-Drive to (2) parallely linked GPU's														
		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (2) GTX470 FTW SLI @825/1000 in Parallel	Avg T CPU	MCR320 Drive Avg T Air CPU loop1	Avg T Water CPU	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg TGPU 1+2	MCR220-Drive Avg T Air GPU loop2	T Water GPU loop 2	ΔT Water GPU to Air	ΔT GPU Air
1	(8) instances of BurnK6		66.9	23.5	27.7	4.2	43.4	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2	Furmark Max Burn		N/A	N/A	N/A	N/A	N/A	53.8	55.0	54.4	26.8	43.3	16.5	27.6
3	Furmark MB + Burnk6 (x7)		67.8	24.7	29.1	4.3	43.1	52.0	54.0	53.0	25.5	41.8	16.3	27.5

SERIES I - Dedicated loops: MCR220-Drive to (1) CPU; MCR320-Drive to (2) parallely linked GPU's														
		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (2) GTX470 FTW SLI @825/1000 in Parallel	Avg T CPU	MCR220 Drive Avg T Air CPU loop1	Avg T Water CPU	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg T GPU 1+2	MCR320-Drive Avg T Air GPU loop2	T Water GPU loop 2	ΔT Water GPU to Air	ΔT GPU Air
4	(8) instances of BurnK6		70.0	25.2	31.1	6.0	44.9	N/A	N/A	N/A	N/A	N/A	N/A	N/A
5	Furmark Max Burn		N/A	N/A	N/A	N/A	N/A	49.0	51.0	50.0	25.8	37.9	12.1	24.2
6	Furmark MB + Burnk6 (x7)		69.3	25.0	31.1	6.1	44.3	47.5	49.5	48.5	24.7	36.7	11.9	23.8

SERIES I - Combined Loop : MCR320-Drive + MCR-220-Drive + (1) CPU + (2) parallely linked GPU's, in series														
		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (2) GTX470 FTW SLI @825/1000 in Parallel	Avg T CPU	MCR320+220 Drive in series Avg T Air	Avg T Water	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg T GPU 1+2	T Air GPU	T Water GPU	ΔT Water GPU to Air	ΔT GPU Air
7	(8) instances of BurnK6		65.8	24.2	27.2	2.9	41.6	31.9	31.6	31.7	N/A	N/A	N/A	7.5
8	Furmark Max Burn		N/A	25.2	33.0	7.7	N/A	45.0	47.0	46.0	N/A	N/A	N/A	20.8
9	Furmark MB + Burnk6 (x7)		72.0	25.0	33.7	8.8	47.1	46.1	48.0	47.1	N/A	N/A	N/A	22.1

Analysis				
Note #	**With radiators and pumps in series,**	**Temperature**		**Dual loop with:**
1	**Under CPU stress only, CPU temps are**	**1.8**	**°C lower than**	**MCR320 drive dedicated to CPU**
2	**Under CPU stress only, CPU temps are**	**3.3**	**°C lower than**	**MCR220 drive dedicated to CPU**
3	**Under GPU stress only, GPU temps are**	**6.8**	**°C lower than**	**MCR220 drive dedicated to GPU's**
4	**Under GPU stress only, GPU temps are**	**3.4**	**°C lower than**	**MCR320 drive dedicated to GPU's**
5	**Under CPU AND GPU stress, CPU temps are**	**4.0**	**°C Higher than**	**MCR320 drive dedicated to CPU**
6	**Under CPU AND GPU stress, CPU temps are**	**2.8**	**°C Higher than**	**MCR220 drive dedicated to CPU**
7	**Under CPU AND GPU stress, GPU temps are**	**5.4**	**°C lower than**	**MCR220 drive dedicated to GPU's**
8	**Under CPU AND GPU stress, GPU temps are**	**1.7**	**°C lower than**	**MCR320 drive dedicated to GPU's**

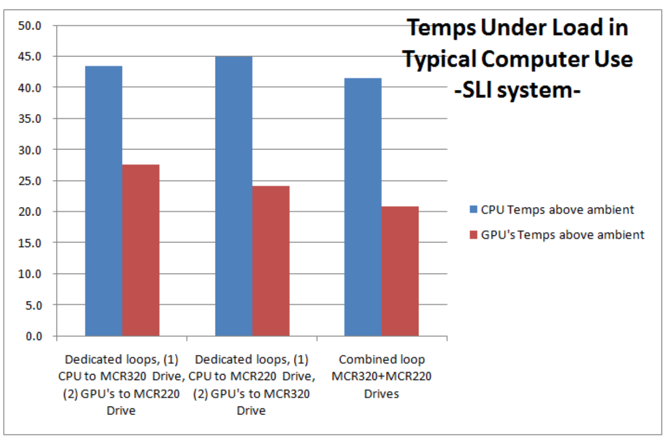

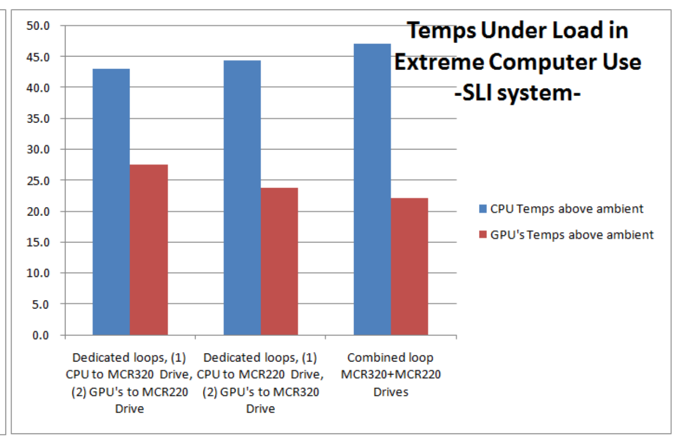

Analysis

Under typical computer use, the above test data suggests as a general rule that users would not benefit from setting up dedicated loops for CPU and GPU. Serializing pumps in the same loop also adds a redundancy factor that dedicated loops cannot provide. With superior reliability and lower temperatures at both CPU and GPU levels, single loops appear to win hands down.

Under extreme computer use, this setup recorded a notable advantage at the CPU temperature level for the dual loop, counterbalanced by the opposite effect at the GPU level. This extreme environment uncovered the critical importance of the respective load ratios generated by CPU class devices vs. GPU class devices, relative to the heat exchangers to which they are connected. Clearly, a CPU generating 150 Watts solely dedicated to a triple radiator will cool substantially better than when mixed with another 400 watts generated by two GPU's even with a second dual radiator in the loop. Jedi Masters would say, "we need to bring balance to the force here", and they would be right.

We could define a simple mathematical method to properly configure a system accounting for loads, but that will be for another article. For now, we can simply illustrate the above in real life testing by removing one of the GPU's from the system. It has for effect to balance the heat load generated by CPU device more evenly against that of the GPU device and demonstrates how load ratios affects the results.

II. Series 2 test results, CPU + Single GPU configuration:

SERIES II - Dedicated loops: MCR320-Drive to (1) CPU; MCR220-Drive to (1) GPU

		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (1) GTX470 FTW @825/1000	Avg T CPU	MCR320 Drive Avg T Air CPU loop1	Avg T Water CPU	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg TGPU 1+2	MCR220-Drive Avg T Air GPU loop2	T Water GPU loop 2	ΔT Water GPU to Air	ΔT GPU Air
10	(8) instances of BurnK6		67.2	23.8	27.9	4.1	**43.4**	N/A	N/A	N/A	N/A	N/A	N/A	N/A
11	Furmark Max Burn		N/A	N/A	N/A	N/A	N/A	44.3	N/A	44.3	25.9	34.5	8.6	**18.4**
12	Furmark MB + Burnk6 (x7)		66.7	23.8	28.0	4.2	**42.9**	42.1	N/A	42.1	24.1	32.6	8.5	**18.0**

SERIES II - Dedicated loops: MCR220-Drive to (1) CPU; MCR320-Drive to (1) GPU

		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (1) GTX470 FTW @825/1000	Avg T CPU	MCR220 Drive Avg T Air CPU loop1	Avg T Water CPU	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg T GPU 1+2	MCR320-Drive Avg T Air GPU loop2	T Water GPU loop 2	ΔT Water GPU to Air	ΔT GPU Air
13	(8) instances of BurnK6		69.3	24.7	30.9	6.2	**44.6**	N/A	N/A	N/A	N/A	N/A	N/A	N/A
14	Furmark Max Burn		N/A	N/A	N/A	N/A	N/A	40.4	N/A	40.4	23.6	29.7	6.2	**16.8**
15	Furmark MB + Burnk6 (x7)		70.0	25.9	31.8	6.0	**44.2**	42.0	N/A	42.0	24.7	31.0	6.3	**17.2**

SERIES II - Combined Loop : MCR320-Drive + MCR-220-Drive + (1) CPU + (1) GPU, in series

		Configuration	CPU DATA (°C)					GPU DATA (°C)						
Test #	Benchmark	Ci7 920@ 4.1Gb + (1) GTX470 FTW @825/1000	Avg T CPU	MCR320+220 Drive in series Avg T Air	Avg T Water	ΔT Water to Air	ΔT CPU Air	Avg T GPU1	Avg T GPU2	Avg T GPU 1+2	T Air GPU	T Water GPU	ΔT Water GPU to Air	ΔT GPU Air
16	(8) instances of BurnK6		67.1	25.3	28.1	2.8	**41.8**	N/A	N/A	N/A	N/A	N/A	N/A	N/A
17	Furmark Max Burn		N/A	25.6	30.1	4.5	N/A	42.0	N/A	42.0	N/A	N/A	N/A	**16.4**
18	Furmark MB + Burnk6 (x7)		69.9	25.8	31.2	5.4	**44.1**	43.0	N/A	43.0	N/A	N/A	N/A	**17.2**

Analysis

Note #	With radiators and pumps in series,	Temperature		Dual loop with:
1	Under CPU stress only, CPU temps are	1.7	°C lower than	MCR320 drive dedicated to CPU
2	Under CPU stress only, CPU temps are	2.9	°C lower than	MCR220 drive dedicated to CPU
3	Under GPU stress only, GPU temps are	2.0	°C lower than	MCR220 drive dedicated to GPU's
4	Under GPU stress only, GPU temps are	0.4	°C lower than	MCR320 drive dedicated to GPU's
5	Under CPU AND GPU stress, CPU temps are	1.2	°C Higher than	MCR320 drive dedicated to CPU
6	Under CPU AND GPU stress, CPU temps are	0.1	°C lower than	MCR220 drive dedicated to CPU
7	Under CPU AND GPU stress, GPU temps are	0.8	°C lower than	MCR220 drive dedicated to GPU's
8	Under CPU AND GPU stress, GPU temps are	0.0	°C lower than	MCR320 drive dedicated to GPU's

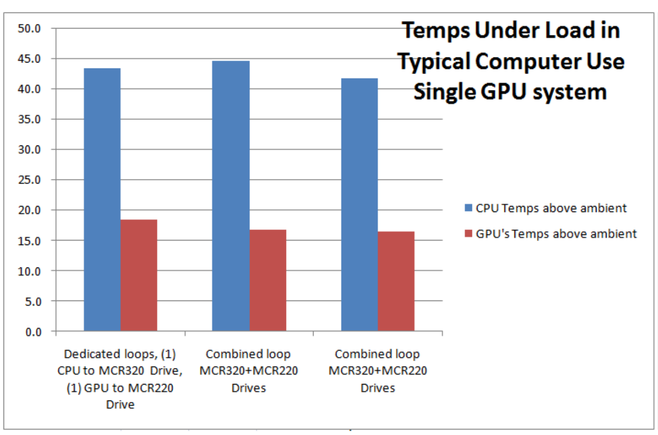

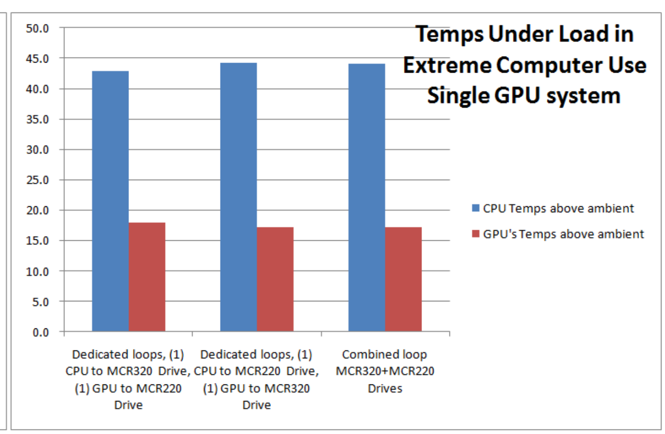

What we see above is that even under extreme use, the dual loop has all but lost its performance advantage against the single loop. Incidentally, the same type of trend could have also been obtained by adding a second CPU in the loop instead of removing a GPU.

Conclusions:

Under extreme performance scenarios, and from a pure performance standpoint, dual loops versus single loop are neither better nor worse, under the strict condition that the load ratios are evenly balanced.
Under the most commonly encountered loads though, single loops do win.

Under both of the above use scenarios, single loops also win from a reliability standpoint because of pump redundancy.

The choice is yours to make.