Notes on BlueGene Architectures

H. Simma

September 5, 2012

1 Basic Parameters

	BG/L	BG/P	BG/Q
core	"processor"		"processor"
f_{clk}	700 MHz	850 MHz	1.6 GHz
arch	440 PowerPC	450 PowerPC	A2
execution			in order
threads	1	1...4	4
DP/clk	4	4	8
	(2 way FMA)	(2 way FMA)	(4 way FMA)
perf (DP)			
registers	32 x 128 bit		32 x 256 bit
L1 inst/core		8 kB	16 KB private
L1 data/core	32 KB	32 KB	16 KB private
L2	14 pref. streams	14 pref. streams	32 MB shared
L3	4 MB shared	8 MB shared	—
node	"compute card"	"compute card"	"compute chip"
	(=2 chips)	(=1 chip)	(=1 chip)
cores	2x2	4	16 (+1+1)
perf (DP)	5.6 GF	13.6 GF	204 GF
memory size	0.5 GB	2 GB	16 GB
memory BW	5.6 GB/s	13.6 GB/s	42 GB/s
rack	1024 compute cards	1024 compute cards	1024 compute chips
	(=32 "node cards")	(=32 "node cards")	(=32 "node boards")
			(=2 midplanes)
cores	4 K	4 K	16 K
perf (DP)	5.7 TF	13.9 TF	209 TF
power	23 kW	29 kW	99 kW
power efficiency	0.24 GF/W	0.5 GF/W	2.1 GF/W
network		4d	5d
BW	0.175 GB/s	0.425 MB/s	2 GB/s
	(*6 links *2 trx)	(*6 lins *2 trx)	(*10 links *2 trx)

2 System Architecture

BG/P

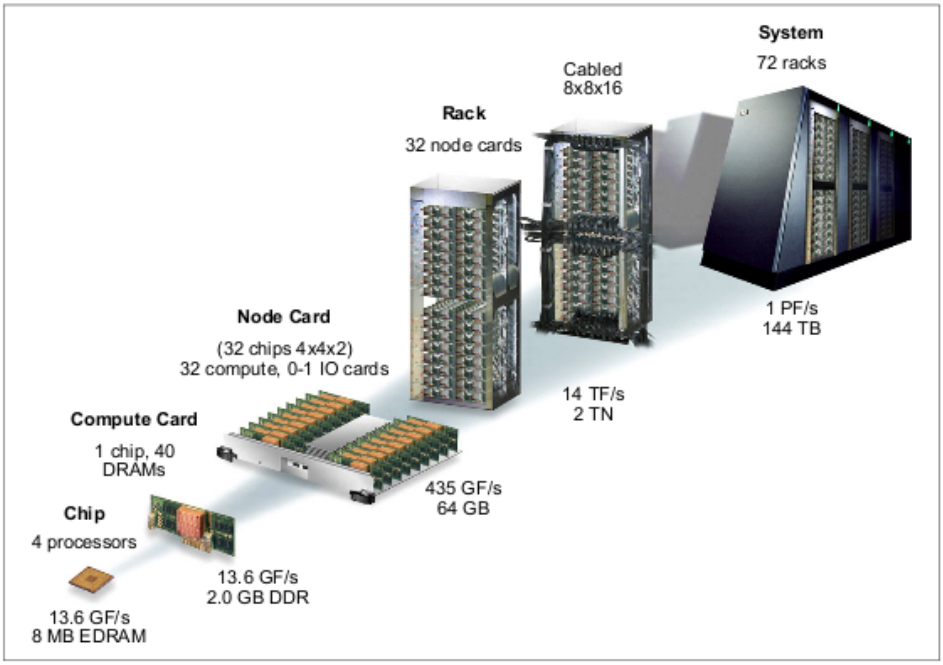

BG/Q

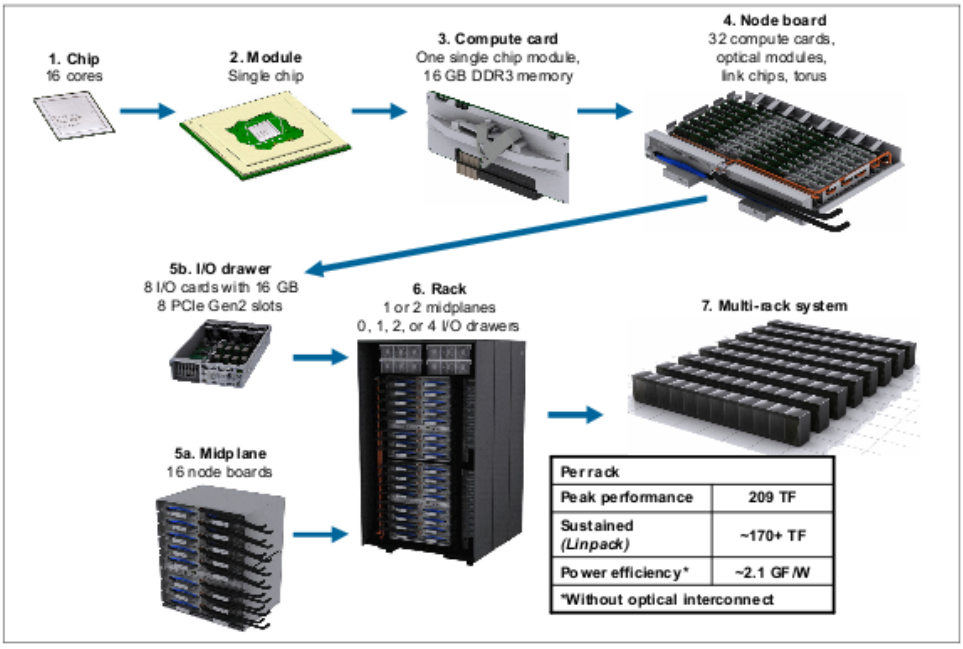

3 Memory System

	BG/L []	BG/P [2]	BG/Q [4]
registers	32 x 128 bit		32 x 256 bit
bandwidth		8 B/clk	
latency		4 clk (3 for int)	6 clk
private	L1	L1	L1
size/core	32 KB	32 KB	16 KB
assocoativity		64-way, 16 sets	8-way set
replacement		RR	pseudo LRU
line size		32 B	64 B
bandwidth		(16 ld + 16 st) B/clk (theo) 5.6 B/clk (st) 4.6 B/clk (ld)	(16 ld + 32 st) B/clk
latency		12 clk (11 for int)	24 clk
private	L2	L2P	L1P
associativity		15-way, fully	16 streams
replacement		RR	depth stealing and RR
line size	—	128 B	128 B
bandwidth		4.6 B/clk (ld)	(32 B ld + 32 B st)/ 2 clk 563 GB/s
latency		50 clk	72 ... 82 clk
shared	L3	L3	L2
size / node	4 MB	8 MB	32 MB / 16 slices
associativity		8-way 2 banks interleaved	16-way, 1024 set (each slice) 4 (8 sub-) banks per slice
replacement		LRU	LRU
line size		128 B	128 B
bandwidth		40 B/clk (single) 3.7 B/clk (4 core)	32 (r/w) B/clk
latency		104 clk	223 clk
external	DDR	DDR	DDR
		4 banks	2 channels

4 Node Architecture

BG/P

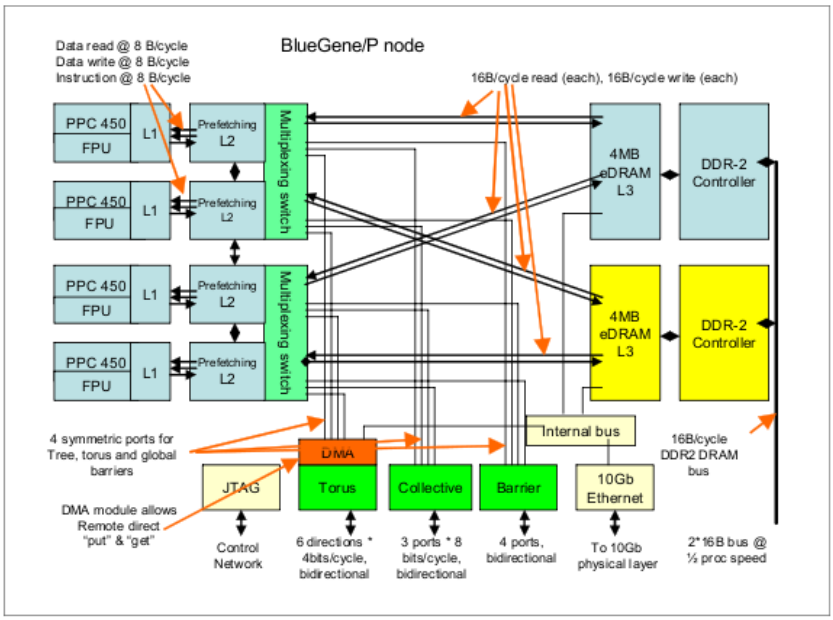

BG/Q

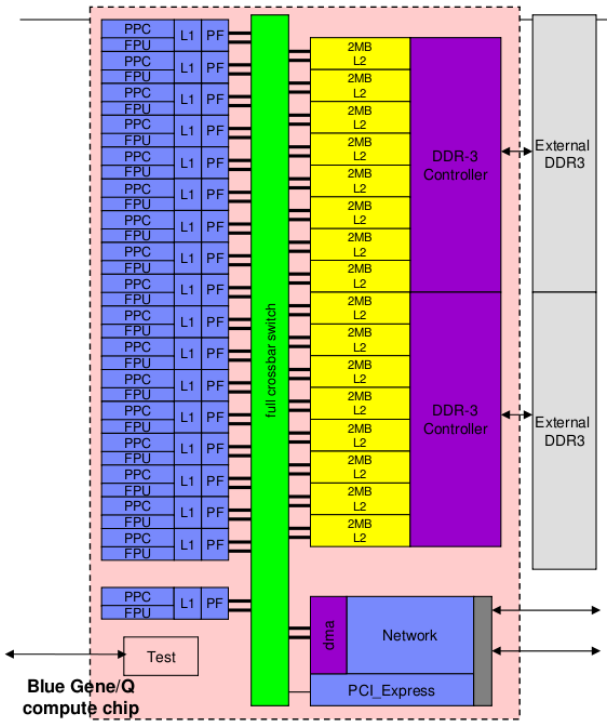

5 FPU Architecture

BG/P: DH

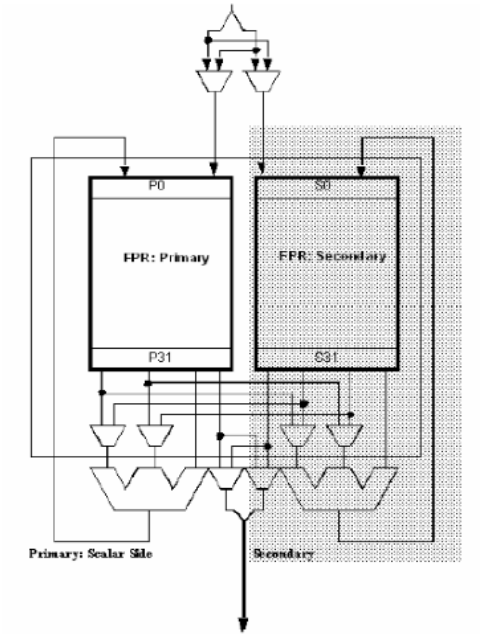

BG/Q: QPX

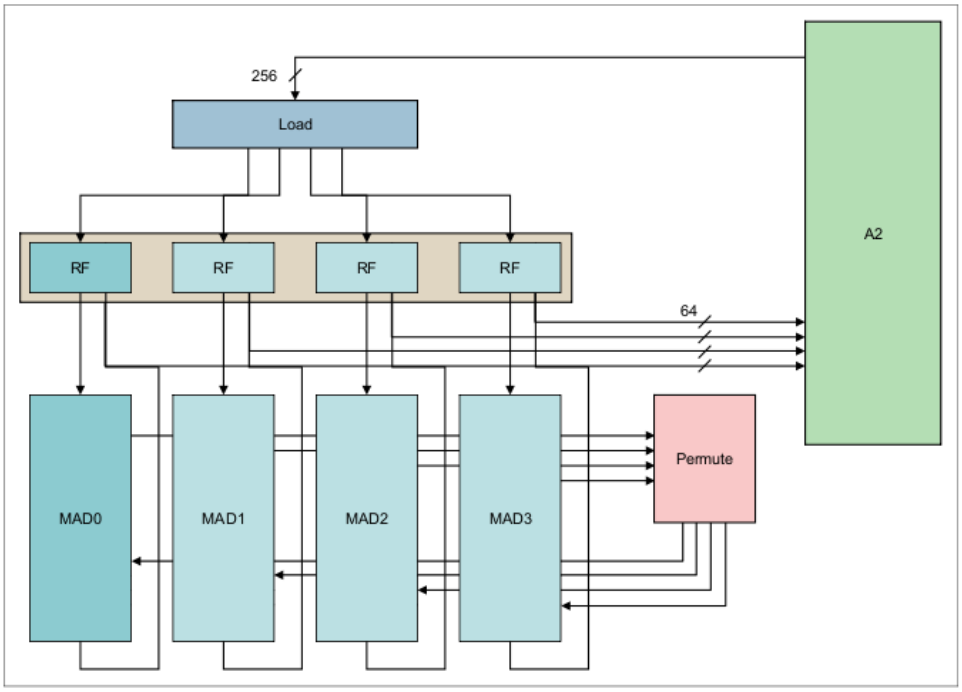

References

[1] B. Steinmacher-Burow, slides from talk at Lattice 2007, Regensburg

[2] IBM Redbooks "BlueGene/P Application Developpment"
SG 24-7287-00 (2007)

[3]

[4] IBM Redbooks "BlueGene/Q Application Developpment"
SG 24-7948-00 (2012)

[5] http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6109225

[6] http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6086530

[7]

[8] http://pic.dhe.ibm.com/infocenter/compbg/v121v141/topic/
com.ibm.xlcpp121.bg.doc/compiler_ref/vec_intrin_cpp.html

[9] http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers
/JUQUEEN/Documentation/Documention_node.html