BE-T-B_T1-SW-013

Тесты демонстрации производительности для микропроцессора Baikal-T1

Список изменений

Revision	Date	Description
1.0	15.03.2017	Initial version
1.1	08.08.2017	Added SPEC CPU2006 Int results
2.0	15.07.2018	Migration to gcc8.1 compiler. Added SPEC CPU2006 FP results.
2.1	18.07.2018	Iperf results for XGbE added.
2.2	20.07.2018	FIO results for PCIE, SATA, USB added.

Содержание

1. Список тестов демонстрации производительности

Test name	Description	Comments
Coremark	List processing (find and sort), Matrix (mathematics) manipulation (common matrix operations), state machine (determine if an input stream contains valid numbers), and CRC.	OpenSource Multi-threading Integer Performance Results http://www.coremark.org/benchmark
Dhrystone	Developed in 1984 by R.P. Wecker, Dhrystone is a benchmark program written in C that tests a system's integer performance. The program is CPU bound, performing no I/O functions or operating system calls. Dhrystones per second is the metric used to measure the number of times the program can run in a second. Original versions of the benchmark gave performance ratings in terms of Dhrystones per second. This was later changed to VAX MIPS by dividing Dhrystones per second by 1757, the DEC VAX 11/780 result.	Integer Performance OpenSource Results http://www.roylongbottom.org.uk/dhrystone%20results.htm
Whetstone	The Whetstone benchmark measure computing power in units of Millions of Whetstone Instructions Per Second (MWIPS). Test Loop In Tables 1. floating point MFLOP 1 2. floating point MFLOP 2 3. if then else IF MOPS 4. fixed point FIXPT MOPS 5. sin,cos etc. COS MOPS 6. floating point MFLOP 3 7. assignments EQUAL MOPS 8. exp,sqrt etc. EXP MOPS	OpenSource Multi-threading Floating-Point Performance (Double and Single precision) **Results** http://freespace.virgin.net/roy.longbottom/whetstone%20results.htm#anchorandroidC
Stream	This program measure memory transfer rates in MB/s for simple operations (copy, scale, add, and triad) **http://www.cs.virginia.edu/stream/**	Multi-threading OpenSource **Results**

		http://ssvb.github.io/2011/09/13/origenboard-memory-performance.html
SPEC CPU 2006 INT	List of Benchmarks 400.perlbench (C) Programming Language 401.bzip2 (C) Compression 403.gcc (C) C Compiler 429.mcf (C) Combinatorial Optimization 445.gobmk (C) Artificial Intelligence: Go 456.hmmer (C) Search Gene Sequence 458.sjeng (C) Artificial Intelligence: chess 462.libquantum (C) Physics / Quantum Computing 464.h264ref (C) Video Compression 471.omnetpp (C++) Discrete Event Simulation 473.astar (C++) Path-finding Algorithms 483.xalancbmk (C++) XML Processing	**Results** http://spec.org/cpu2006/results/
Iperf	**Iperf** is a test for network performance measurement. Iperf has “client” and “server” functionality, and can create data streams to measure the throughput between the two ends.	OpenSource
fio	Flexible I/O Tester. **fio** is a tool that will spawn a number of threads or processes doing a particular type of I/O action as specified by the user. The typical use of fio is to write a job file matching the I/O load one wants to simulate.	OpenSource https://github.com/axboe/fio
hdparm	hdparm is a command line program for Linux to set and view ATA hard disk drive hardware parameters and test performance. It can set parameters such as drive caches, sleep mode, power management, acoustic management, and DMA settings. GParted and Parted Magic both include hdparm. Perform timings of device (-t) and cache (-T) reads for benchmark and comparison purposes.	OpenSource

2. Оборудование для демонстрации производительности

Процессор	CPU: MIPS32 P5600 @ 1200 MHz (Rev 1.0) FPU: Present Cores: 2 Timer: 600 MHz ECC: L1 L2 (80800ff0) PLLs: CPU: 1200MHz SATA: 600MHz ETH:1250MHz PCIE:1200MHz DDR3: 400MHz AXI: 600MHz
Плата	Baikal-T1 BFK3.1
Память	Kingston KVR16S11/8. 1.5V 1G x 64-bit (8GB) DDR3-1600 CL11 SDRAM
OS	Linux 4.4.135
SDK	Baikal-T1 SDK 4.14 (https://www.baikalelectronics.ru/products/T1/?type=razrabotka)
Disks	SSD Intel DC S3510 Series 120 GB SSDSC2BB120G601 ORICO 2.5 inch USB2.0 Hard Drive Enclosure (2588US) Plextor M8Pe 128GB PCIe Gen 3 x4

3. Результаты тестов демонстрации производительности

OS: Linux 4.4.135

При сборке бенчмарков используются следующие компиляторы с опциями.
Compilers:

- GCC 8.1 -Ofast -funroll-all-loops -mmsa -EL -mtune=p5600 -static
- Для получения максимальной производительности на coremark

  в качестве компилятора использовался Mentor тулчейн 4.9.1(2014.1-22)
  https://sourcery.mentor.com/GNUToolchain/subscription3537?lite=MIPS
  с плагином (tree_switch_shortcut_elf-2014.11-21.so)
  http://community.imgtec.com/developers/mips/tools/benchmarks/ в качестве компоновщика
  https://sourcery.mentor.com/GNUToolchain/release2935

  GCC4.9.1 -static -O3 -funroll-all-loops -fgcse-sm -fgcse-las -finline-functions -finline-limit=1000 -msoft-float -EL -march=74kc -falign-functions=16 -mno-dsp -fplugin=./tree_switch_shortcut_elf-2014.11-21.so

Benchmark	Baikal-T1 MIPS P5600 1200 MHz		Compiler / Software
Coremark	GCC 8.1 Coremarks 10653 (2 threads) Coremarks/MHz 8.88 Coremarks/Mhz/core 4.44	Mentor GCC 4.9.1 with plugin Coremarks 13142 (2 threads) Coremarks/MHz 10.95 Coremarks/Mhz/core 5.47	GCC 8.1 MentorGCC 4.9
Dhrystone	4432 VAX MIPS (1 thread) 3.69 DMIPS/Mhz		GCC 8.1
Whetstone	1679 MWIPS (2 threads) 0.69 MWIPS/Mhz/core		GCC 8.1
STREAM	Copy: 3307 MB/s Scale: 3300 MB/s Add: 2479 MB/s Triad: 2482 MB/s		GCC 8.1
SPEC CPU2006	INT(geomean) 5.6 FP(geomean) 3.8 FP(geomean without milc, gamess) 4.8		GCC 8.1
iperf	1 Gb Ethernet Bandwidth: 940 Mbits/sec (TCP) 10 Gb Ethernet Bandwidth: 3050 Mbits/sec (TCP)		iperf2.0.8 iperf3.2 GCC 8.1
FIO	PCEe (SSD Plextor) WRITE: bw=454MiB/s ; READ: bw=654MiB/s SATA (Intel SSD 120 GB) WRITE: bw=139MiB/s; READ: bw=366MiB/s USB2.0 (Intel SSD 120 GB over ORICO USB2.0 Hard Drive Adapter) WRITE: bw=39.1MiB/s ; READ: bw=37.9MiB/s		fio3.7

4. Лог запуска тестов

Coremark

```
developer@baikal:~ $ ./coremark_49.exe
2K performance run parameters for coremark.
CoreMark Size    : 666
Total ticks      : 16740
Total time (secs): 16.740000
Iterations/Sec   : 13142.174432
Iterations       : 220000
Compiler version : GCC4.9.1
Compiler flags   : -O3 -funroll-all-loops -fgcse-sm -fgcse-las -finline-
functions -finline-limit=1000 -msoft-float -EL -G4
-fplugin=./tree_switch_shortcut_elf-2014.11-21.so -march=74kc -falign-
functions=16 -mno-dsp -DMULTITHREAD=2 -DUSE_FORK -DPERFORMANCE_RUN=1  -lrt
Parallel Fork : 2
Memory location  : Please put data memory location here
                   (e.g. code in flash, data on heap etc)
seedcrc          : 0xe9f5
[0]crclist       : 0xe714
[1]crclist       : 0xe714
[0]crcmatrix     : 0x1fd7
[1]crcmatrix     : 0x1fd7
[0]crcstate      : 0x8e3a
[1]crcstate      : 0x8e3a
[0]crcfinal      : 0x33ff
[1]crcfinal      : 0x33ff
Correct operation validated. See readme.txt for run and reporting rules.
CoreMark 1.0 : 13142.174432 / GCC4.9.1 -O3 -funroll-all-loops -fgcse-sm -fgcse-
las -finline-functions -finline-limit=1000 -msoft-float -EL -G4
-fplugin=./tree_switch_shortcut_elf-2014.11-21.so -march=74kc -falign-
functions=16 -mno-dsp -DMULTITHREAD=2 -DUSE_FORK -DPERFORMANCE_RUN=1  -lrt /
Heap / 2:Fork

developer@baikal:~ $ ./coremark_81.exe
2K performance run parameters for coremark.
CoreMark Size    : 666
Total ticks      : 20650
Total time (secs): 20.650000
Iterations/Sec   : 10653.753027
Iterations       : 220000
Compiler version : GCC8.1.0
Compiler flags   : -static -lrt -Ofast -funroll-all-loops -fgcse-sm -fgcse-las
-finline-functions -finline-limit=1000 --fast-math -EL -march=p5600 -mtune=p5600
-falign-functions=16
Parallel Fork : 2
```

```
Memory location  : Please put data memory location here
                   (e.g. code in flash, data on heap etc)
seedcrc          : 0xe9f5
[0]crclist       : 0xe714
[1]crclist       : 0xe714
[0]crcmatrix     : 0x1fd7
[1]crcmatrix     : 0x1fd7
[0]crcstate      : 0x8e3a
[1]crcstate      : 0x8e3a
[0]crcfinal      : 0x33ff
[1]crcfinal      : 0x33ff
Correct operation validated. See readme.txt for run and reporting rules.
CoreMark 1.0 : 10653.753027 / GCC8.1.0 -static -lrt -Ofast -funroll-all-loops
-fgcse-sm -fgcse-las -finline-functions -finline-limit=1000 --fast-math -EL
-march=p5600 -mtune=p5600 -falign-functions=16 / Heap / 2:Fork
```

Whetstone

```
developer@baikal: $ # ./whetsmp.exe
 ####################################################
  get_nprocs() - CPUs 2, Configured CPUs 2
  get_phys_pages() and size - RAM Size  3.61 GB, Page Size 16384 Bytes
  uname() - Linux, localhost.localdomain, 4.4.135-bfk3-06728-gc5040bf
  #51 SMP Thu Jun 21 19:55:09 MSK 2018, mips

 Multithreading Single Precision Whetstones 32-Bit Version 1.0

             Using 2 threads - Wed Jul 18 08:55:31 2018
 Calibrate
    0.01336 Seconds          1   Passes (x 100)
    0.06249 Seconds          5   Passes (x 100)
    0.30946 Seconds         25   Passes (x 100)

 Use 807  passes (x 100)

 MFLOPS 1     411    408
 MFLOPS 2     474    471
 IFMOPS      8110   7160
 FIXPMOPS 22352471363444
 COSMOPS       14     14
 MFLOPS 3     265    265
 EQUMOPS     9616   9091
 EXPMOPS       11     11
 millisec    9608   9595
 MWIPS        840    841

         MWIPS MFLOPS MFLOPS MFLOPS    Cos    Exp  Fixpt     If  Equal
 Thread             1      2      3   MOPS   MOPS   MOPS   MOPS   MOPS

 Total    1681    819    945    530     27     223598691  15270  18707

 MWIPS    1679 Based on time for last thread to finish


 Results Of Calculations Thread 1

 MFLOPS 1   -1.12356138229370117      MFLOPS 2   -1.13133072853088379
 IFMOPS      1.00000000000000000      FIXPMOPS   12.00000000000000000
 COSMOPS     0.49911013245582581      MFLOPS 3    0.99999982118606567
 EQUMOPS     3.00000000000000000      EXPMOPS     0.93536460399627686
```

Dhrystone

```
developer@baikal:$ ./dhry.exe
 ###################################################
  get_nprocs() - CPUs 2, Configured CPUs 2
  get_phys_pages() and size - RAM Size  3.61 GB, Page Size 16384 Bytes
  uname() - Linux, localhost.localdomain, 4.4.135-bfk3-06728-gc5040bf
  #51 SMP Thu Jun 21 19:55:09 MSK 2018, mips

Dhrystone Benchmark, Version 2.1 (Language: C or C++)

Optimisation    Opt 2 32 Bit
Register option not selected

       10000 runs   0.00 seconds
      100000 runs   0.01 seconds
     1000000 runs   0.13 seconds
     2000000 runs   0.26 seconds
     4000000 runs   0.51 seconds
     8000000 runs   1.03 seconds
    16000000 runs   2.05 seconds

Final values (* implementation-dependent):

Int_Glob:      O.K.  5  Bool_Glob:     O.K.  1
Ch_1_Glob:     O.K.  A  Ch_2_Glob:     O.K.  B
Arr_1_Glob[8]: O.K.  7  Arr_2_Glob8/7: O.K.    16000010
Ptr_Glob->             Ptr_Comp:       *    4854664
  Discr:       O.K.  0  Enum_Comp:     O.K.  2
  Int_Comp:    O.K.  17 Str_Comp:      O.K.  DHRYSTONE PROGRAM, SOME STRING
Next_Ptr_Glob->        Ptr_Comp:       *    4854664 same as above
  Discr:       O.K.  0  Enum_Comp:     O.K.  1
  Int_Comp:    O.K.  18 Str_Comp:      O.K.  DHRYSTONE PROGRAM, SOME STRING
Int_1_Loc:     O.K.  5  Int_2_Loc:     O.K.  13
Int_3_Loc:     O.K.  7  Enum_Loc:      O.K.  1
Str_1_Loc:                             O.K.  DHRYSTONE PROGRAM, 1'ST STRING
Str_2_Loc:                             O.K.  DHRYSTONE PROGRAM, 2'ND STRING

Microseconds for one run through Dhrystone:         0.13
Dhrystones per Second:                           7787511
VAX  MIPS rating =                                 4432.28
```

STREAM

```
developer@baikal:$ ./stream.exe
-------------------------------------------------------------
STREAM version $Revision: 5.10 $
-------------------------------------------------------------
This system uses 8 bytes per array element.
-------------------------------------------------------------
Array size = 48000000 (elements), Offset = 0 (elements)
Memory per array = 366.2 MiB (= 0.4 GiB).
Total memory required = 1098.6 MiB (= 1.1 GiB).
Each kernel will be executed 10 times.
 The *best* time for each kernel (excluding the first iteration)
 will be used to compute the reported bandwidth.
-------------------------------------------------------------
Number of Threads requested = 1
Number of Threads counted = 1
-------------------------------------------------------------
Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 242435 microseconds.
   (= 242435 clock ticks)
Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.
-------------------------------------------------------------
WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.
-------------------------------------------------------------
Function    Best Rate MB/s  Avg time     Min time     Max time
Copy:            3307.8     0.232516     0.232175     0.233148
Scale:           3300.8     0.233006     0.232674     0.233518
Add:             2479.9     0.464971     0.464536     0.465537
Triad:           2482.6     0.464478     0.464036     0.465782
-------------------------------------------------------------
Solution Validates: avg error less than 1.000000e-13 on all three arrays
-------------------------------------------------------------
```

5. Результаты SPECCPU2006

Результаты производительности микропроцессора БАЙКАЛ-Т1 на бенчмарках из пакета SPECCPU2006 представлены на диаграммах (см. ниже рисунок 1, 2).

SPECCPU2006 INT	Ref results
400. perlbench	5.8
401, bzip2	3.7
403. gcc	5.0
429. mcf	4.9
445. gobmk	5.6
456. hmmer	8.8
458. sjeng	5.5
462. libquantum	12.1
464. h264ref	5.7
471. omnetpp	4.1
473. astar	4.2
483. xalancbmk	5.6
SPECCPU2006 INT (geomean)	**5.6**

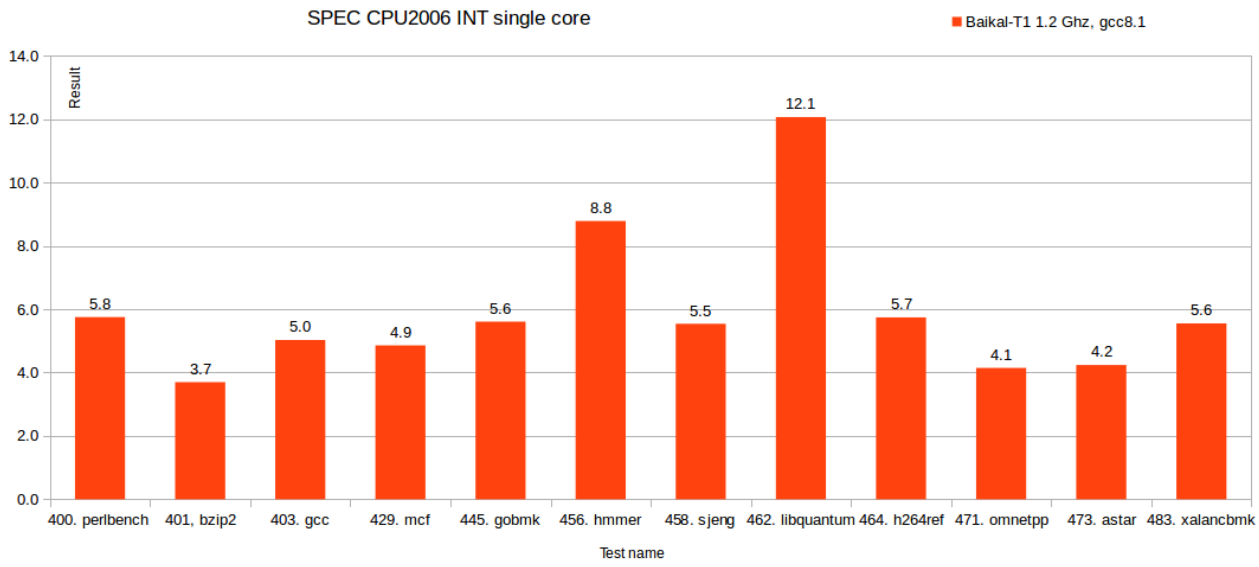

Рисунок 1. Результаты на бенчмарках из пакета SPECCPU2006 INT (ref workload).

SPECCPU 2006 FP	Ref results
410. bwaves	5.5
416. gamess	*1.5*
433. milc	*0.2*
434. zeusmp	6.0
435. gromacs	4.3
436. cactusADM	4.0
437. leslie3d	5.2
444. namd	5.2
447. dealII	8.2
450. soplex	4.4
453. povray	5.8
454. calculix	4.1
459. GemsFDTD	4.4
465. tonto	3.1
470. lbm	4.1
481. wrf	4.9
482. sphinx3	5.6
SPECfp (geomean)	3.8
SPECfp (geomean without milc, gamess)	4.8

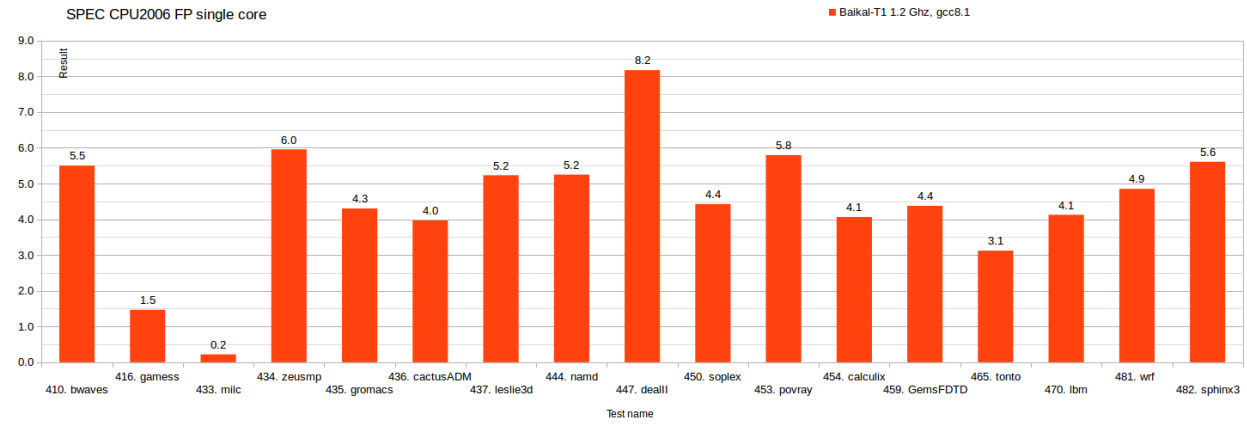

Рисунок 2. Результаты на бенчмарках из пакета SPECCPU2006 FP (ref workload).

SPECPU2006 options

```
# Optimization

OPTIMIZE = -Ofast -static -funroll-all-loops -mmsa -ffast-math
-falign-functions=16 -falign-loops=32 -march=p5600 -mtune=p5600
-EL

COPTIMIZE = $(OPTIMIZE)

CXXOPTIMIZE = $(OPTIMIZE) -std=c++03 -fpermissive

FOPTIMIZE   = $(OPTIMIZE)  -DSPEC_CPU_LINUX  -ffixed-form

############################################################# #
32/64 bit Portability Flags - all


# Portability Flags

400.perlbench=default=default=default:

CPORTABILITY =  -fno-strict-aliasing -fno-store-merging -fsigned-
char -DSPEC_CPU_LINUX -mno-mips16 -mno-interlink-mips16 -std=gnu89

462.libquantum=default=default=default:

CPORTABILITY   =  -DSPEC_CPU_LINUX

483.xalancbmk=default=default=default:

CXXPORTABILITY = -DSPEC_CPU_LINUX

481.wrf=default=default=default:

CPORTABILITY   = -DSPEC_CPU_CASE_FLAG -DSPEC_CPU_LINUX

436.cactusADM=default=default=default:

CPORTABILITY   = -DSPEC_CPU_LINUX

416.gamess=default=default=default:

FPORTABILITY = -DSPEC_CPU_LP64 -fno-strict-aliasing -mno-mips16
-mno-interlink-mips16 -std=legacy
```

6. Производительность Gigabit Ethernet

Результаты получены на бенчмарке `iperf` (ver. 2.0.8b, https://iperf.fr), предназначенном для тестирования пропускной способности интернет канала (GbE интерфейс, протокол TCP).

Байкал-T1 в качестве «клиента»:

```
root@baikal:~# iperf -c 192.168.68.10
connect failed: Connection refused
root@baikal:~# iperf -c 192.168.68.10
------------------------------------------------------------
Client connecting to 192.168.68.10, TCP port 5001
TCP window size: 43.8 KByte (default)
------------------------------------------------------------
[  3] local 192.168.68.27 port 51434 connected with 192.168.68.10 port 5001
[ ID] Interval       Transfer     Bandwidth
[  3]  0.0-10.0 sec  1.10 GBytes   942 Mbits/sec
root@baikal:~# iperf -c 192.168.68.10 -t 100
------------------------------------------------------------
Client connecting to 192.168.68.10, TCP port 5001
TCP window size: 43.8 KByte (default)
------------------------------------------------------------
[  3] local 192.168.68.27 port 51436 connected with 192.168.68.10 port 5001
[ ID] Interval       Transfer     Bandwidth
[  3]  0.0-100.0 sec  11.0 GBytes   941 Mbits/sec
root@baikal:~# iperf -c 192.168.68.10 -t 300
------------------------------------------------------------
Client connecting to 192.168.68.10, TCP port 5001
TCP window size: 48.1 KByte (default)
------------------------------------------------------------
[  3] local 192.168.68.27 port 51438 connected with 192.168.68.10 port 5001
[ ID] Interval       Transfer     Bandwidth
[  3]  0.0-300.0 sec  32.9 GBytes   941 Mbits/sec
```

Байкал-T1 в качестве «сервера»:

```
root@baikal:~# iperf -s
------------------------------------------------------------
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
------------------------------------------------------------
[  4] local 192.168.68.27 port 5001 connected with 192.168.68.10 port 51069
[ ID] Interval       Transfer     Bandwidth
[  4]  0.0-10.0 sec  1.10 GBytes   940 Mbits/sec
[  5] local 192.168.68.27 port 5001 connected with 192.168.68.10 port 51070
[  5]  0.0-300.0 sec  32.8 GBytes   940 Mbits/sec
```

7. Производительность 10 Gigabit Ethernet

Используются две платы BFK 3.1, соединённые непосредственно друг с другом без использования какого-либо промежуточного сетевого оборудования (топология «точка-точка»). Соединение осуществляется при помощи оптического либо медного кабеля с XGbE-трансиверами на концах.

XGbE PHY	Marvell Alaska X 88X2222
Трансивер XGbE для оптического кабеля	Intel FTLX8571D3BCV-IT
Трансивер XGbE для медного кабеля	Juniper Networks 740-030429 Rev 01

Нижеприведённые результаты получены на стандартном тесте `iperf3` (https://iperf.fr) версии 3.2+. Тест `iperf3` реализует клиент-серверную модель. Таким образом, одной плате BFK 3.1 назначается условная роль сервера, а второй — роль клиента. Роль определяется ключами, передаваемыми исполняемому файлу `iperf3` при запуске (см. ниже).

Предварительные настройки

На стороне клиента и сервера:

```
sysctl -w net.ipv4.tcp_timestamps=0
sysctl -w net.ipv4.tcp_sack=0
sysctl -w net.ipv4.tcp_rmem="10000000 10000000 10000000"
sysctl -w net.ipv4.tcp_wmem="10000000 10000000 10000000"
sysctl -w net.ipv4.tcp_mem="10000000 10000000 10000000"
sysctl -w net.core.rmem_max=4194303
sysctl -w net.core.wmem_max=4194303
sysctl -w net.core.rmem_default=4194303
sysctl -w net.core.wmem_default=4194303
sysctl -w net.core.optmem_max=4194303
sysctl -w net.core.netdev_max_backlog=300000
echo 2 > /proc/irq/82/smp_affinity
echo 2 > /proc/irq/83/smp_affinity
echo 2 > /proc/irq/84/smp_affinity
echo 2 > /proc/irq/85/smp_affinity
```

Только на стороне сервера:

```
ifconfig eth0 10.0.4.1 netmask 255.255.255.0 mtu 9000 txqueuelen 8000 up
```

Только на стороне клиента:

```
ifconfig eth0 10.0.4.2 netmask 255.255.255.0 mtu 9000 txqueuelen 8000 up
```

Запуск

На стороне сервера:

```
/mnt/disk1/opt/iperf/iperf3 -A 1 -V -s
```

На стороне клиента:

```
/mnt/disk1/opt/iperf/iperf3x -A 1 -V -l 1M -c 10.0.4.1 -O 1
```

Результаты

В результате серии тестов установлено, что скорость обмена данными составляет ~3,05 Гбит/с при использовании как медного, так и оптического кабеля. Ниже приводятся примеры журналов сервера и клиента.

Сервер

```
iperf 3.2
Linux baikal-BFK3-0 4.4.135-bfk3 #4 SMP Mon Jul 2 19:43:46 MSK 2018 mips
-----------------------------------------------------------
Server listening on 5201
-----------------------------------------------------------
Time: Fri, 06 Jul 2018 13:06:18 GMT
Accepted connection from 10.0.4.2, port 52280
      Cookie: wv64rl7yk32wjydjip5kc2g7uuvdmqhduku7
      TCP MSS: 0 (default)
[  5] local 10.0.4.1 port 5201 connected to 10.0.4.2 port 52282
Starting Test: protocol: TCP, 1 streams, 1048576 byte blocks, omitting 1 seconds, 10 second test,
tos 0
[ ID] Interval           Transfer     Bitrate
[  5]   0.00-1.00   sec   346 MBytes  2.89 Gbits/sec                  (omitted)
[  5]   0.00-1.00   sec   366 MBytes  3.06 Gbits/sec
[  5]   1.00-2.00   sec   367 MBytes  3.07 Gbits/sec
[  5]   2.00-3.00   sec   367 MBytes  3.08 Gbits/sec
[  5]   3.00-4.00   sec   364 MBytes  3.05 Gbits/sec
[  5]   4.00-5.01   sec   365 MBytes  3.06 Gbits/sec
[  5]   5.01-6.00   sec   360 MBytes  3.02 Gbits/sec
[  5]   6.00-7.00   sec   365 MBytes  3.06 Gbits/sec
```

```
[  5]   7.00-8.01   sec   367 MBytes  3.08 Gbits/sec
[  5]   8.01-9.01   sec   367 MBytes  3.07 Gbits/sec
[  5]   9.01-10.00  sec   362 MBytes  3.06 Gbits/sec
- - - - - - - - - - - - - - - - - - - - - - - - -
Test Complete. Summary Results:
[ ID] Interval           Transfer     Bitrate
[  5] (sender statistics not available)
[  5]   0.00-10.00  sec  3.56 GBytes  3.06 Gbits/sec                  receiver
CPU Utilization: local/receiver 37.2% (0.0%u/37.1%s), remote/sender 51.5% (0.6%u/50.9%s)
snd_tcp_congestion cubic
rcv_tcp_congestion cubic
```

Клиент

```
iperf 3.2
Linux baikal-BFK3-0 4.4.135-bfk3 #4 SMP Mon Jul 2 19:43:46 MSK 2018 mips
Control connection MSS 8960
Time: Fri, 06 Jul 2018 13:06:19 GMT
Connecting to host 10.0.4.1, port 5201
      Cookie: wv64rl7yk32wjydjip5kc2g7uuvdmqhduku7
      TCP MSS: 8960 (default)
[  5] local 10.0.4.2 port 52282 connected to 10.0.4.1 port 5201
Starting Test: protocol: TCP, 1 streams, 1048576 byte blocks, omitting 1 seconds, 10 second test,
tos 0
[ ID] Interval           Transfer     Bitrate         Retr  Cwnd
[  5]   0.00-1.00   sec   356 MBytes  2.99 Gbits/sec    0   3.32 MBytes       (omitted)
[  5]   0.00-1.00   sec   365 MBytes  3.06 Gbits/sec    0   3.67 MBytes
[  5]   1.00-2.00   sec   367 MBytes  3.08 Gbits/sec    0   3.67 MBytes
[  5]   2.00-3.00   sec   367 MBytes  3.08 Gbits/sec    0   3.67 MBytes
[  5]   3.00-4.00   sec   363 MBytes  3.05 Gbits/sec    0   3.67 MBytes
[  5]   4.00-5.00   sec   365 MBytes  3.06 Gbits/sec    0   3.67 MBytes
[  5]   5.00-6.00   sec   361 MBytes  3.03 Gbits/sec    0   4.45 MBytes
[  5]   6.00-7.01   sec   369 MBytes  3.07 Gbits/sec    0   4.45 MBytes
[  5]   7.01-8.00   sec   364 MBytes  3.08 Gbits/sec    0   4.45 MBytes
[  5]   8.00-9.00   sec   365 MBytes  3.06 Gbits/sec    0   4.45 MBytes
[  5]   9.00-10.00  sec   368 MBytes  3.08 Gbits/sec    0   4.67 MBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
Test Complete. Summary Results:
[ ID] Interval           Transfer     Bitrate         Retr
```

```
[  5]   0.00-10.00  sec  3.57 GBytes  3.06 Gbits/sec    0             sender
[  5]   0.00-10.00  sec  3.56 GBytes  3.06 Gbits/sec                  receiver
CPU Utilization: local/sender 54.7% (0.2%u/54.5%s), remote/receiver 37.2% (0.0%u/37.1%s)
snd_tcp_congestion cubic
rcv_tcp_congestion cubic
```

8. Производительность PCIe

Подсистема PCIe тестировалась на карте PCIe NVMe SSD Plextor M8PeGN. Карта инициализировалась на скорости GEN3. Для замеров скорости записи и чтения применялась утилита fio v3.7 (лог. запуска см. ниже). Для использования драйвера в конфигурацию ядра Linux нужно добавить опции:

```
CONFIG_BLK_DEV_NVME=y
CONFIG_NVMEM=y
```

Выдержка из `dmesg`:

```
pci 0000:01:00.0: Link Capability is GEN3, x4
pci 0000:01:00.0: Link Status is GEN1, x4
pci 0000:01:00.0: retrain link to GEN3
pci 0000:01:00.0: Link Status is GEN3, x4
bus: 'pci': driver_probe_device: matched device 0000:01:00.0 with
driver nvme
bus: 'pci': really_probe: probing driver nvme with device
0000:01:00.0
devices_kset: Moving 0000:01:00.0 to end of list
device: 'nvme0': device_add
PM: Adding info for No Bus:nvme0
driver: 'nvme': driver_bound: bound to device '0000:01:00.0'
bus: 'pci': really_probe: bound device 0000:01:00.0 to driver nvme
```

Запуск утилиты fio:

```
# fio --filename=/dev/nvme0n1 --direct=1 --rw=read
--ioengine=vsync --bs=1M --iodepth=1 --numjobs=2 --size=10G
--group_reporting --name pcie
pcie: (g=0): rw=read, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB,
(T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
   READ: bw=654MiB/s (685MB/s), 654MiB/s-654MiB/s (685MB/s-
```

```
685MB/s), io=20.0GiB (21.5GB), run=31330-31330msec
Disk stats (read/write):
  nvme0n1: ios=162911/0, merge=0/0, ticks=177360/0,
in_queue=177330, util=94.57%
```

```
# fio --filename=/dev/nvme0n1 --direct=1 --rw=write
--ioengine=vsync --bs=1M --iodepth=1 --numjobs=2 --size=10G
--group_reporting --name pcie
pcie: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-
1024KiB, (T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
  WRITE: bw=454MiB/s (477MB/s), 454MiB/s-454MiB/s (477MB/s-
477MB/s), io=20.0GiB (21.5GB), run=45062-45062msec
Disk stats (read/write):
  nvme0n1: ios=0/163068, merge=0/0, ticks=0/395670,
in_queue=396500, util=95.14%
```

```
# hdparm -tT /dev/nvme0n1p1
Timing buffer-cache reads: 1254 MB in 0.50 seconds = 2518422 kB/s
Timing buffered disk reads: 1505 MB in 3.00 seconds = 513684 kB/s
```

9. Производительность SATA

Для замеров скорости записи и чтения на SATA интерфейсе применялась утилита fio v3.7 (лог. запуска ниже). В качестве оборудования использовался SSD-накопитель Intel 120 GB подключенный SATA-кабелем к БФК3.1.

```
# hdparm -tT /dev/sdb
Timing buffer-cache reads:  1148 MB in 0.50 seconds = 2305930 kB/s
Timing buffered disk reads:  791 MB in 3.00 seconds = 269785 kB/s
```

```
# fio --filename=/dev/sdb --direct=1 --rw=read --ioengine=vsync
--bs=1M --iodepth=1 --numjobs=2 --size=4G --group_reporting --name
```

sata_ssd

```
sata_ssd: (g=0): rw=read, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-
1024KiB, (T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
   READ: bw=366MiB/s (384MB/s), 366MiB/s-366MiB/s (384MB/s-
384MB/s), io=8192MiB (8590MB), run=22377-22377msec


Disk stats (read/write):
  sdb: ios=8108/0, merge=0/0, ticks=43710/0, in_queue=43700,
util=99.63%
```

 # fio --filename=/dev/sdb --direct=1 --rw=write --ioengine=vsync --bs=1M --iodepth=1 --numjobs=2 --size=4G --group_reporting --name sata_ssd

```
sata_ssd: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-
1024KiB, (T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
  WRITE: bw=139MiB/s (146MB/s), 139MiB/s-139MiB/s (146MB/s-
146MB/s), io=8192MiB (8590MB), run=58894-58894msec


Disk stats (read/write):
  sdb: ios=0/8191, merge=0/0, ticks=0/116230, in_queue=116240,
util=99.89%
```

10. Производительность USB

Для замеров скорости записи и чтения на USB2.0 интерфейсе применялась утилита fio v3.7 (лог. запуска ниже). В качестве оборудования использовался SSD-накопитель Intel 120 GB вставленный в USB2.0-адаптер ORICO подключенный USB-кабелем к БФК3.1.

 # fio --filename=/dev/sdb --direct=1 --rw=write --ioengine=vsync --bs=1M --iodepth=1 --numjobs=2 --size=2G --group_reporting --name usb

```
usb: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB,
```

```
(T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
  WRITE: bw=39.1MiB/s (40.0MB/s), 39.1MiB/s-39.1MiB/s (40.0MB/s-
40.0MB/s), io=4096MiB (4295MB), run=104853-104853msec
--
Disk stats (read/write):
  sdb: ios=0/36790, merge=0/0, ticks=0/1458840, in_queue=1459200,
util=99.96%
```

 # fio --filename=/dev/sdb --direct=1 --rw=read --ioengine=vsync --bs=1M --iodepth=1 --numjobs=2 --size=2G --group_reporting --name usb

```
usb: (g=0): rw=read, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB,
(T) 1024KiB-1024KiB, ioengine=vsync, iodepth=1
...
Run status group 0 (all jobs):
   READ: bw=37.9MiB/s (39.7MB/s), 37.9MiB/s-37.9MiB/s (39.7MB/s-
39.7MB/s), io=4096MiB (4295MB), run=108165-108165msec


Disk stats (read/write):
  sdb: ios=36857/0, merge=0/0, ticks=1514290/0, in_queue=1514310,
util=99.97%
```

hdparm -tT /dev/sdb

```
Timing buffer-cache reads:  1314 MB in 0.50 seconds = 2640402 kB/s
Timing buffered disk reads:  113 MB in 3.01 seconds = 38333 kB/s
```