1 Test of PCIe SSD Card Intel P3608

1.1 Introduction

Before results of our test are proposed we describe some preliminaries related to disk operations. In the case of disk operations (read and write) there is a minimal chunk of data addressed on a disk. We call it the *data block* and its size is usually 2 kB in the current files systems. In our test we consider two basic disk operations: *sequential* and *random* accesses. Sequential accesses are especially utilized when bigger files are processed (read or written) where neighbour data blocks are consecutively accessed. On the other hand, random accesses occur when small files are processed or a DBMS searches for a record in an index: in this case two consecutive disk operations access distant data blocks of a disk. Therefore we test both these disk operations and we use 4 or 8 kB data blocks to test random accesses. As usual random accesses are 2–3 orders of magnitude slower compared to sequential accesses. Moreover, Solid State Disk (*SSD*) provides significantly faster random accesses compared to Hard Drive Disk (*HDD*).

1.2 Random/Sequential Operations Test

In the first test, we compare sequential and random operations for various disks. We compare three non-volatile memory data storages: an SDHC card, Samsung 840 Pro Series SSD[1] (further only Samsung SSD) and Intel P3608[2] (further only Intel SSD) with two HDDs: SATA disks with RAID 1 (further only SATA disk) and SAS disks with RAID 1+0[3] (further only SAS disk). Let us note that only Intel SSD is connected via PCIe and we used the CrystalDiskMark application[4] in this test. The results of the experiment are presented in Table 1. In this table we show the throughput of the operations as well as the number of data blocks read per second (IOPS).

Although random accesses have been tested with 4 kB and 512 kB data blocks, we should understand that the operation with 512 kB data block does not represent a typical random operation. As we can see, the throughput of sequential reads and writes is comparable for the SATA and Samsung SSD disks, however Samsung SSD overcomes the SATA disk. The SDHC card provides the lowest throughput of the sequential operations; the throughput is 2 orders of magnitude lower compared to other devices. Intel SSD provides 3× higher throughput of the sequential operations than Samsung SSD. When we consider

[1]http://www.samsung.com/global/business/semiconductor/Downloads/DataSheet-Samsung_SSD_840_PRO_Rev12.pdf

[2]http://www.intel.com/content/www/us/en/solid-state-drives/solid-state-drives-dc-p3608-series.html

[3]David A. Patterson et al. : A Case for Redundant Arrays of Inexpensive Disks (RAID). In Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data. Url: http://dl.acm.org/citation.cfm?id=50214

[4]http://crystalmark.info/software/CrystalDiskMark/index-e.html

Table 1: The throughput (MB/s) of the various disks for the sequential and random operations

	SDHC	SATA RAID1	SAS RAID10	Samsung SSD	Intel SSD	Ratio
Sequential Read	12.1	159.3	416.8	516.4	**1,611.0**	**3.1**
Sequential Write	11.0	177.9	249.6	494.3	**1,313.4**	**2.7**
Random Read 4kB	3.7	2.5	4.2	30.3	**41.3**	**1.4**
IOPS	902.6	613.3	1 020.8	7 392.4	**10,079.8**	
Random Write 4kB	1.4	12.1	15.9	63.8	**195.3**	**3.1**
IOPS	344.3	2 946.4	3 868.8	15 574.8	**47,669.8**	
Random Read 512kB	11.8	173.5	239.0	465.5	**1,211.8**	**2.6**
Random Write 512kB	1.5	213.1	237.3	475.7	**1,317.9**	**2.8**

the 4 kB data block, the throughput of random operations is 1–2 orders of magnitude lower than sequential operations for all disks; Samsung SSD provides one order of magnitude higher throughput than both HDDs and SDHC and it provides 75% of the throughput of Intel SSD. On the other hand, Intel SSD provides 3× higher throughput of the random writes than Samsung SSD. In summary, although we can see that the throughput of Intel SSD is up-to 3× higher compared to Samsung SSD for most of the operations, the throughput of random reads is only 40 MB/s, i.e. it is comparable with Samsung SSD and this operation therefore remains the most problematic operation.

1.3 Random Reads Test

In the second test, we show the throughput of random reads for various data block sizes from 4 kB–16 MB. This test has been executed only for Intel SSD and we used the Diskspd utility[5] to measure the performance. The parameters of the application have been set as follows for the 8 kB data block:

```
diskspd64.exe -b8K -d5 -o1 -t1 -W -r -h -wWWW -c1G <data_file>
```

The results of the experiment are presented in Table 2 and Figure 1. As we can see, the throughput of the random reads significantly increases with the increasing data block size. 90% of the maximal throughput is reached for the 4 MB data block, however we must keep in mind that such a data block does not represent real sequential reads.

[5]https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223

Table 2: The throughput (MB/s) of the random reads with the increasing data block size

Block Size (kB)	**Throughput (MB/s)**
4	29.47
8	50.36
16	80.82
32	128.75
64	180.86
128	287.65
256	529.27
512	911.52
1,024	1185.70
2,048	1427.82
4,096	1628.97
8,192	1678.52
16,384	1646.33

1.4 Database Index Performance Test

In the last test, we focus on the throughput of operations insert and point query over two data structures often used in relational DBMSs: the sequential array (also called the Heap table) and the B-tree. Let us note that the sequential array is often used as a basic storage for data, whereas the B-tree is often used as the basic index. We insert 50 millions randomly generated records with the 5 attributes. A prototype a DBMS called RadegastDB[6] has been used in the experiment, the 8 kB data block has been set for both data structures. We use 10 different configurations of the data structures: in the first configuration the cache buffer size (i.e. the main memory area of the DBMS) can include only 10% of data blocks, whereas in the last configuration all data blocks are stored in the cache buffer.

The results of the experiment are put forward in Table 3 and Figures 2 and 3. Evidently, there are only sequential accesses in the case of the sequential array, whereas there are only random accesses in the case of the B-tree.

The throughput of the insert operation in the sequential array is 3× higher when all data blocks are stored in the cache buffer compared to the configuration where only 10% of data blocks are stored in the cache buffer. The throughput of the query operation is only 1.5× higher for the cache buffer including all data blocks. It reflects the high performance of sequential operations of Intel SSD

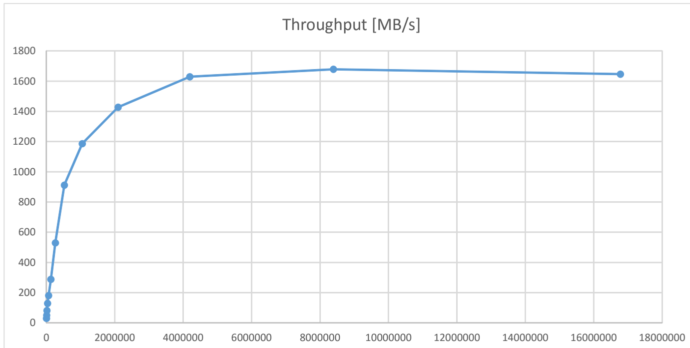

Figure 1: Intel SSD – Random reads test

Table 3: P3608 - Throughput (op./s) of the operations insert and query in a sequential array and B-tree

	Sequential Array		**B-tree**	
	Insert	**Point Query**	**Insert**	**Point Query**
10 %	2,245,475	6,269	7,442	7,979
20 %	2,848,353	7,280	10,293	9,106
30 %	3,155,967	7,739	14,708	10,306
40 %	3,458,293	8,003	21,337	12,051
50 %	3,624,353	8,240	32,590	14,105
60 %	4,146,624	8,669	52,589	17,085
70 %	4,441,408	8,887	88,679	23,629
80 %	5,131,362	9,329	138,397	33,469
90 %	5,504,124	9,732	263,531	63,963
100 %	6,830,601	9,917	317,373	273,569

presented in the previous tests.

A different behaviour occurs in the case of the B-tree, since random accesses are common during insert as well as query operations. The throughput of the insert operation is 2 orders of magnitude higher when all data blocks are stored in the cache buffer compared with the configuration where 10% of data blocks are stored in the cache buffer. The query operation provides the similar results.

[6] http://db.cs.vsb.cz

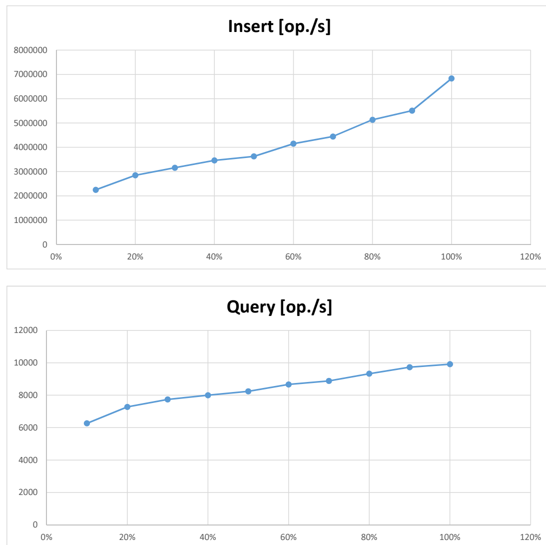

Figure 2: Intel SSD – Test of the sequential array

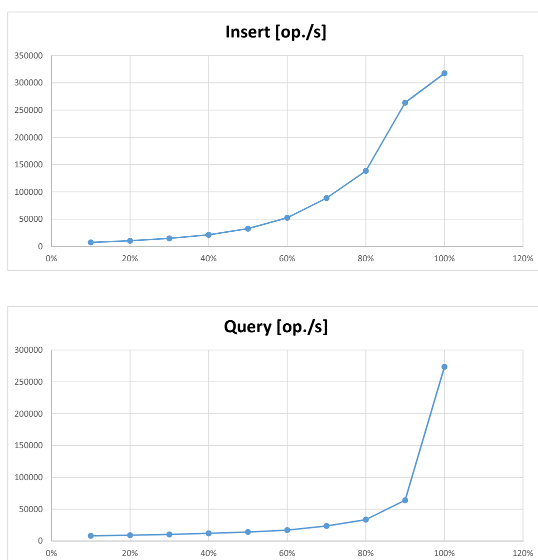

Figure 3: Intel SSD – Test of the B-tree

2 Summary

Although Intel SSD provides the throughput up-to 1.6 GB/s of the sequential operations and thus it overcomes Samsung SSD 3×, the random read still remains an inefficient operation since the throughput is only 41 MB/s. On the other hand the throughput of the random write is 3× higher compared to other disks – 195 MB/s.