IJEEC

International Journal of Electrical Engineering and Computing
Vol.3, No. 2 (2019)

Original Research Papers
UDC 007.5:629.3-592.35]:793.7
DOI 10.7251/IJEEC1902070D

Influence of different factors on the RAID 0 paired magnetic disk arrays

Nikola Davidović[1], Slobodan Obradović[1], Borislav Đorđević[2], Valentina Timčenko[2] and Bojan Škorić[3]

[1] University of East Sarajevo, Faculty of Electrical Engineering, East Sarajevo, RS, Bosnia and Herzegovina,

[2]Institute Mihailo Pupin, Belgrade, Serbia,
[3]VISER, Belgrade, Serbia

nikola.davidovic@etf.ues.rs.ba, slobo.obradovic@gmail.com, borislav.djordjevic@pupin.rs, valentina.timcenko@pupin.rs, bojan.skoric@viser.edu.rs

Abstract— **The rapid technological progress has led to a growing need for more data storage space. The appearance of big data requires larger storage space, faster access and exchange of data as well as data security. RAID (Redundant Array of Independent Disks) technology is one of the most cost-effective ways to satisfy needs for larger storage space, data access and protection. However, the connection of multiple secondary memory devices in RAID 0 aims to improve the secondary memory system in a way to provide greater storage capacity, increase both read data speed and write data speed but it is not fault-tolerant or error-free. This paper provides an analysis of the system for storing the data on the paired arrays of magnetic disks in a RAID 0 formation, with different number of queue entries for overlapped I/O, where queue depth parameter has the value of 1 and 4. The paper presents a range of test results and analysis for RAID 0 series for defined workload characteristics. The tests were carried on in Microsoft Windows Server 2008 R2 Standard operating system, using 2, 3, 4 and 6 paired magnetic disks and controlled by Dell PERC 6/i hardware RAID controller. For the needs of obtaining the measurement results, ATTO Disk Benchmark has been used. The obtained results have been analyzed and compared to the expected behavior.**

Keywords- HDD; secondary memory; magnetic disk; performances; RAID 0; ATTO Disk Benchmark; Windows

I. INTRODUCTION

The technological advancement in the last three decades has led to the fact that the average user needs have grown to TB storage space. The increase in the amount of data being exchanged requires the need for larger storage space, as well as faster access and easier management of data. Such tendencies have led to an increased risk of compromising or even loss of data security.

With years, the improvements of the secondary computer memory performance were slower than the performance improvements of the processor and the main memory. Since computer components highly depend of their physical capabilities and thus can only be improved up to a certain limit, the magnetic disk has practically achieved its maximum. Some further improvement in the performance of the secondary memory can be achieved either by multiple use of parallel disks or by the application of new technologies.

In spite of the emergence of new, the semiconductor memory - Solid State Disk (SSD) technology, magnetic drives - HDD (Hard Disk Drive) still play a dominant role as a secondary computer memory primarily due to its high capacity and cost per MB. One of the biggest disadvantages of using magnetic disk is the speed of data read and write operations.

RAID (initially abbreviated as Redundant Array of Inexpensive Disks [1], now known as the Redundant Array of Independent Disks) is a technology that has been virtually inaccessible to an average user for a long time, primarily because of the high cost of even the simplest RAID system configuration. The first commercial RAID systems came about in 2000, when manufacturers achieved cheaper RAID controls that integrated on the board.

In addition to the fact that the RAID data storage system that is based on the paired disks needs to have an appropriate controller, the cost of the entire system is significantly influenced by the price of the secondary computer memory device. With the decrease of the costs of secondary memory devices, as well as with the need for increasing capacity and providing faster data read and write speeds, the use of RAID technology has finally become meaningful to the average user.

RAID technology is defined through seven (7) different levels of data storage organization on multiple disks, which are combined in a single logical space. Seven RAID levels,

This paper is a revised and expanded version of the paper presented at the 18th International Symposium INFOTEH-JAHORINA 2019 [9] although they have their descriptive names in practice, are most often mentioned under their numbers [1], [2].

In spite of the fact that some manufacturers call their technology other names, it's important to emphasize that all RAID systems are essentially based on some of the seven RAID levels, with the possibility of introducing some minor changes, add-ons or specific modes of implementation. With the standard seven levels of RAID technology, it is possible to realize nested or hybrid RAID.

RAID levels can be realized in two ways: hardware and software RAID. The support for specific RAID levels can be also provided by operating systems (i.e. with some Windows versions), or even provided on the file system level (i.e. using Oracle/Solaris ZFS) [3][4]. In addition, the software RAID can be also found as a stand-alone application [5].

II. Redundant Array of Independent Disks Level 0

RAID 0 offers the highest degree of storage space and provides the best read and write performance when compared to all other RAID levels, but it does not offer redundancy.

Despite the disadvantages it has, and thanks to its unmatched performance, RAID 0 is used in systems where data access speeds and storage space size play the key role. Figures 1 and 2 provide an overview of the stripping procedure and methodology applied in RAID 0.

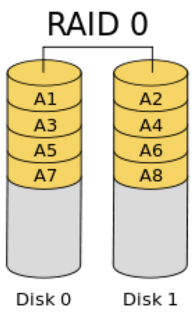

Figure 1. RAID 0 – applies striping to two secondary memory devices

Figure 1 explains the principles of the basic stripping technique and shows one data that is divided into two secondary memory devices. Theoretically, this division should result in two times faster access speeds, read and write. However, in practice, this often does not yield the expected results.

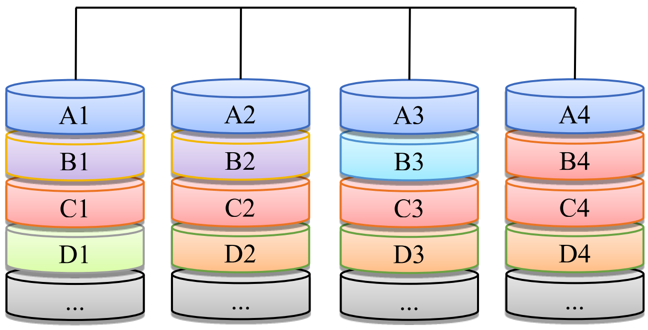

Figure 2. RAID 0 – parallelity and competitiveness – each color is another complete piece of information

Compared with the Figure 1 that explains the principles of the basic stripping technique, Figure 2 provides more detailed example of stripping several data over four different disks that are structured in RAID 0 architecture. Each exemplar data is shown in different color, thus we can see that, for instance, data A is stripped on all the four disks, while data B, being smaller in size, is stripped over two disks from the array.

When compared to all other RAID levels, RAID 0 offers the highest degree of storage space. Theoretically, the addition of each new secondary memory device in a series increases proportionally the performance, but also, since there is no redundancy in the system, there is a statistical risk of losing all data by the failure of only one secondary memory device in the array.

The standard RAID 0 layer is rarely used on servers because it does not offer redundancy but is often used as part of a hybrid RAID system. It is most commonly used on workstations.

Ideally, RAID 0 with N secondary memory devices can be considered to have N times better sequential and random reading and writing data speed (represented by equations 1 and 2) compared to one secondary memory device [6].

$$T_{seq\_rw}^{RAID-0} \approx T_{seq\_rw\_one\_seq\_m\_d}/N \quad (1)$$

$$T_{random\_rw}^{RAID-0} \approx T_{random\_one\_seq\_m\_d}/N \quad (2)$$

Typical characteristic of the RAID 0 level "striping by blocks" works quite similar to ideal model but still shows some inconsistencies.

The size of data for writing or reading from the secondary memory device may be smaller or larger than the defined block (which is essentially the data carrier). Regardless of the size of the data in the block, the secondary memory device accesses the entire block. Therefore, block size estimation is of great importance when designing a RAID system. Block size estimation has a significant role in maximizing the RAID performances.

For estimating the size of a block, it is necessary to take into the account the parallelism and competitiveness (Figure 2). In case when the block size is determined so that the data unit exactly occupies the defined disk memory unit (on the RAID FS - full stripe), we can expect the increase of the speed of data access for N (N is the number of secondary memory devices). When configured in this way, RAID 0 supports parallelism and high sequential performance. However, if the goal is to increase the competitiveness and the random-access speed performance, then the size of the SU (Strip Unit) needs to be adjusted exactly to the size of the data unit.

III. Queue Depth

Queue depth refers to the number of outstanding access operations. In Figure 3, each solid line represents one disk operation, which can be either a read or a write operation. Because three operations overlap in the same period, there's a queue depth of three.

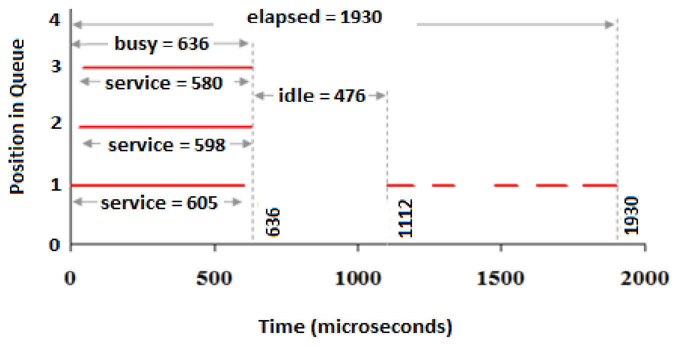

Figure 3. Represents queue depth factor 3

The storage queue depth represents the number of pending input/output (I/O) requests for a volume. In other words, queue depth is the number of I/Os that a device/controller can have at the same time. The other I/Os will be pending in a queue at the OS/app level. Lower queue depth gives lower latencies and a higher queue depth gives better throughput. The device uses the queue depth either for internal parallelism (SSDs) and/or for reordering and merging of related I/Os (HDDs and SSDs).

IV. Test Configuration

A. Hardware configuration

The hardware configuration is shown in Table I. The tests are carried on the Microsoft Windows Server 2008 R2 Standard operating system. No server components or functions are added to the basic installation of the operating system, except for the necessary drivers - the RAID controller and other specific hardware drivers.

TABLE I
Hardware configuration

Hardware	Specification
Server	Dell PowerEdge™ T610
RAM	8 GB, 4 x 2 GB DDR3-SDRAM
CPU Model	Intel Xeon E5530 @ 2,40 GHz (4 cores)
BIOS	Dell Inc. v.2.2.10 (9.11.2010.)
Video adapter	Matrox G200
PCIe x4 storage slot	Dell PERC 6/i
Disk	Hitachi Deskstar 250GB SATA2 x 6
PCIe x4 slot	Dell SAS 5/i
Disk	Hitachi Ultra Star 300GB SAS
OS	Microsoft Windows Server 2008 R2 Standard

The hardware RAID controller, Dell PERC 6/i, which is used for test procedures, supports devices with the second generation SATA/SAS interface (3 Gb/s), while 2 SAS channels allow up to 32 connected devices. It has 256 MB of its own DDR2 cache for quick storage, which can be optionally supported by a battery. It supports operation with RAID levels 0, 1, 5, 6, 10, 50 and 60. The ATTO Disk Benchmark is used for the needs of testing RAID 0 level storage performance.

B. ATTO Disk Benchmark

ATTO Disk Benchmark is a freeware software which helps the measurement of the storage system performance [7]. ATTO identifies the performance levels of the hard drives, solid state drives, RAID arrays, as well as the host connection to the attached storage. One of the advantages of this benchmark is the ability to control the process of data write and read operations, while the drawback is the inability to test the random data access speed. The ATTO Disk Benchmark is compatible with Microsoft Windows and supports the File Allocation Table (FAT) and New Technology File System (NTFS).

Some of the setting options over which ATTO Disk Benchmark can affect system performance or can isolate certain situations in practical work, are:

- *Total length* – this parameter specifies the test file length, which is the total size of data file that is created on the test drive. After finishing the testing procedure, this file is deleted.
- *Force write access* – this option allows to bypass the drive write cache. Otherwise, if this option is not selected, the drive write caching is determined by the drive settings.
- *Direct I/O* – use of the system buffering. File I/O on the test drive is performed with no system buffering or caching.
- *I/O Comparison* – compares the input and output data to detect errors. This option allows the comparison of the data from the test file with the data written on a per block basis.
- *Overlapped I/O* – this option performs queued I/O testing. The factor that specifies the I/O overlapping is the queue depth. Queue depth specifies the number of queue entries for overlapped I/O, i.e. the maximum number of read/write commands that can be executed during one time interval.
- *Neither* – do not perform overlapped I/O or I/O comparisons. The transfer requests are sent one by one.

V. Test Results

Certain restrictions were set in order to get as highest speed as possible for reading and writing during the test procedures.

The first limitation that is set, is to use only a specific part of the magnetic disk for testing, since magnetic discs do not have the same data transfer speed at the beginning and at the end of the disk. The tests are configured to use only the first 10 GB of each magnetic disk of the RAID 0 array. In this way, in the case for 2, 3, 4 and 6 magnetic disks, an entry space of 20 GB, 30 GB, 40 GB and 60 GB is obtained, respectively. Since these 10 GB makes up less than 5% of the disk space, the limitations of the data rate at the beginning (in the middle of the disk) and at the end (disk circumference) of the magnetic disk has been avoided during testing.

Data caching feature could give wrong results so that they would not show the real performance of the magnetic disk, and the performance results would be masked with the cache effect. In addition, a significant effect may also be achieved by caching at the level of a single magnetic disk, controller or operating system. Because of this, when configuring each array, we have used the option to bypass the cache of the disks, as well as to generate the caching on the controller. In ATTO disk benchmark this is enabled by the Force Write Access and Direct I/O options and represents the second limitation.

The third limitation is the number of multiple transfer requests which define the maximum number of read/write commands that can be executed in one-time interval. The queue depth factor specifies the number of queue entries for overlapped I/O. In this way, the ability to test competitiveness is not eliminated.

In the following test we have used 1 GB size NTFS partition, as the test file space was limited to 512 MB. The larger file was selected in order to get better average values for large transfers, which can also be assumed as the sequential data access test. For the allocation unit we have used the standard size of 4kB. When testing the RAID 0 string, the three block sizes were used:

- 8 kB, the smallest block that the controller permits;
- 64 kB, default value;
- 1 MB, the largest block that the controller permits.

The values of SUs in the Table 2 are given in the first column. In the second column of the Table 2, we have presented the number of magnetic disks on which the tests were performed for different SUs. The test procedure starts with the evaluation of 512 bytes and ends with 8192 KB, with a step in which the next data size is twice as the size of the previous one. The results for read (blue) and for write (red) operations are given in MB/s.

A. Queue depth 1

Figure 4 a1) shows the speed of reading the data from a single magnetic disk and a paired series of RAID 0 magnetic disks at a block size of 8 KB for different amounts of data (the measurements are given in Table 2). It can be observed that RAID 0 at a SU size of 8 KB with two and three paired disks has a lower data transfer rate than a single magnetic disk, that is a performance degradation, while for 4 and 6 paired magnetic disks that performance is slightly better. In addition, for a small amount of data, the RAID transmission rates are almost identical to a single magnetic disk. In the case of reading at 8KB SU size, the expected results of RAID 0 with 6 paired magnetic disks are far from expected according to equations (1) and (2). If we look at the results for write operation for a SU size of 8 KB, Figure 4 a2), then it is noticeable that in this configuration RAID 0 is far worse than a single magnetic disk. The reason for such results, in addition to poor selection of SU size, is the lack of competitiveness.

RAID effects are only noticeable when reading and writing are performed in cases where the size of the RAID SU is 64 KB and 1 MB (Figure 4 b) and c), Table 2). Although all set parameters are identical in this case, the influence of the SU size choice is noticeable. For a larger amount of data, the results obtained are much better than in the case for an 8 KB SU size.

TABLE II Tests results - QD 1

SU	MB/s	0.5 KB	1 KB	2 KB	4 KB	8 KB	16 KB	32 KB	64 KB	128 KB	256 KB	512 KB	1024 KB	2048 KB	4096 KB	8192 KB
8 KB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	1769	3539	7097	13461	26166	33512	49756	66466	78486	86231	85948	86731	86872	87013	86175
	3 HDD	1718	3479	7079	13751	26101	48907	63781	90394	106131	123072	125128	124275	123418	127522	125730
	4 HDD	1774	3389	6959	13180	25911	48306	65209	96613	126946	149628	163759	165191	154941	156979	169093
	6 HDD	1765	3328	6959	13685	26295	49648	90112	118617	170365	198266	228698	235987	245146	250874	248551
64 KB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	1735	3505	7097	13312	26360	46923	79149	127254	187122	242828	281343	282068	283558	281575	282563
	3 HDD	1744	3522	7027	13556	26038	47604	80313	125128	226298	282428	359107	402653	421902	421902	419430
	4 HDD	1756	3479	7168	13524	25785	48786	79921	130096	226298	341459	428694	493674	523776	523776	523776
	6 HDD	1765	3381	6976	13461	26101	47836	76920	130419	230790	369914	442925	567516	639123	688296	709521
1 MB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	1698	3454	7062	13312	25416	46589	76204	117817	168924	208889	240625	276262	282068	282563	280594
	3 HDD	1731	3513	7027	13461	25661	45936	75328	118082	169642	207663	248514	280594	421902	416987	417798
	4 HDD	1532	3548	6976	13366	25416	46479	75155	118417	170356	184104	253259	275789	438261	527637	526344
	6 HDD	1723	3531	6959	13212	25661	46923	74642	118082	164517	199255	242828	265777	439158	667749	678152
8 KB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	63	130	279	655	981	1955	3877	7665	14928	15483	20149	23619	25699	26816	27545
	3 HDD	62	126	267	586	1237	1952	3873	7665	15117	28902	29208	37596	38048	41045	41202
	4 HDD	62	125	261	552	1166	1923	3891	7719	14451	28587	31022	39650	40795	50128	53633
	6 HDD	60	129	280	657	1180	1964	3900	7764	15348	29890	52271	56871	71966	74773	80611
64 KB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	60	119	247	504	1044	2226	4576	7665	14911	28401	51451	87154	133883	171524	174308
	3 HDD	61	123	248	501	1024	2155	5112	7867	15187	29289	54784	96733	157440	220029	228942
	4 HDD	61	123	247	500	1015	2092	4507	8478	15294	29587	56193	101680	167249	250874	313348
	6 HDD	61	123	243	496	1005	2055	4519	8202	14979	29822	57932	107589	186737	283060	387166
1 MB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	61	122	241	488	979	1955	3868	7638	14743	27858	49884	82722	133883	134217	133218
	3 HDD	61	122	245	491	972	1957	3810	7620	14628	27710	49554	81221	139266	146886	177771
	4 HDD	61	122	245	491	981	1955	3887	7665	14860	27623	49182	81965	139446	257492	253839
	6 HDD	61	122	245	490	979	1952	3882	7638	14860	27681	50267	83365	139992	255652	270237

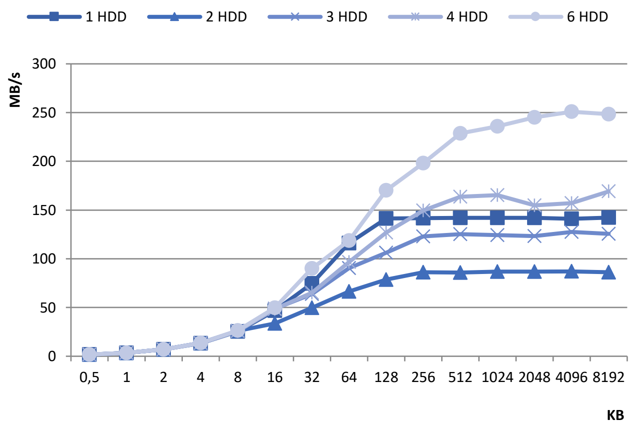

a1) Queue Depth 1, SU 8 KB

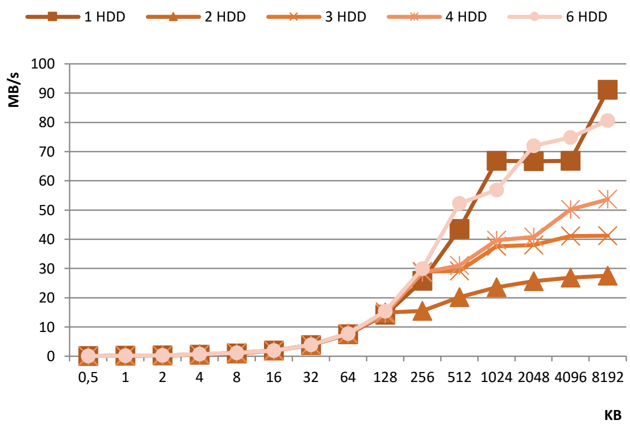

a2) Queue Depth 1, SU 8 KB

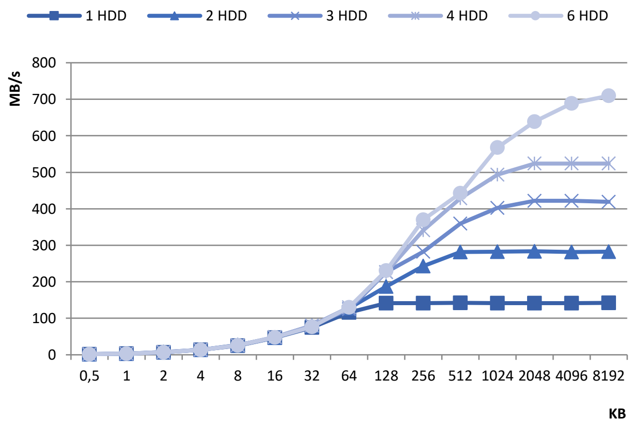

b1) Queue Depth 1, SU 64 KB

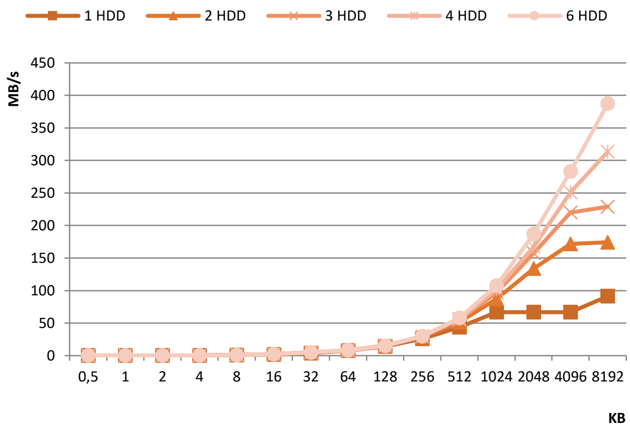

b2) Queue Depth 1, SU 64 KB

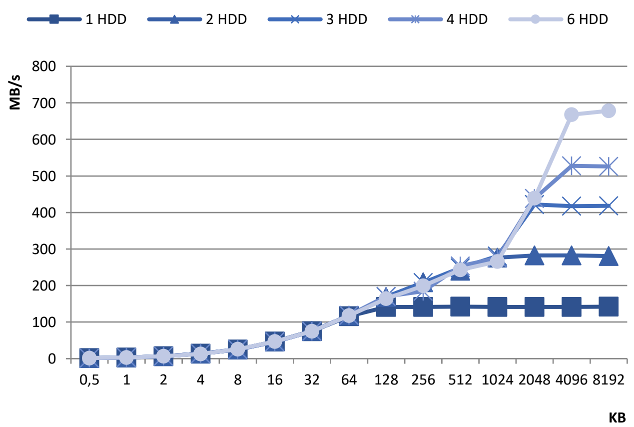

c1) Queue Depth 1, SU 1 MB

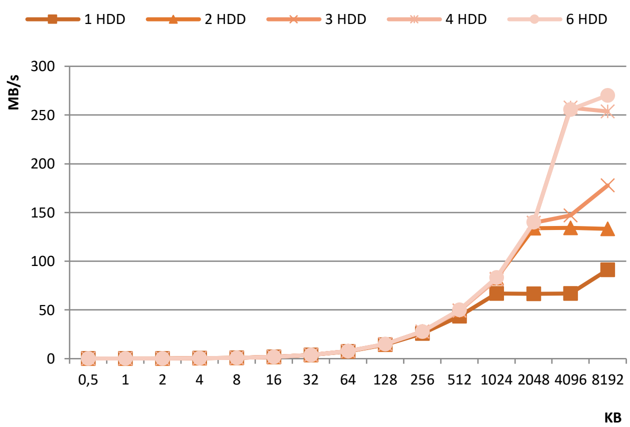

c2) Queue Depth 1, SU 1 MB

Figure 4. Reading speed 1) and writing speed 2) data for single magnetic disk and paired string of 2, 3, 4 and 6 magnetic disk drivers in RAID 0 at the SU size a) 8 KB, b) 64 KB and c) 1 MB

When comparing the values in Table 2 for the SU size of 64KB, for the maximum amount of data transferred, the improvement over one magnetic disk for reading is 198%, 295%, 368% and 499%, while for the writing it is 191%, 250%, 340 % and 425% (for two, three, four and six paired magnetic disks in RAID 0, respectively). Similar but slightly smaller improvements are obtained for 1 MB block values. Comparing the results for 64 KB and 1 MB block size (Table 2), it is noticeable that when increasing the block size to 1 MB there is not a significant improvement in performance compared to the 64 KB block (the performance gain is approximately equal for read and write operations with 64 KB and 1 MB blocks). Although improvements for 64 KB and 1 MB SU sizes are better than when using a single magnetic disk (as well as RAID 0 with an 8 KB SU size), they are smaller than in the ideal case which is represented by equations (1) and (2).

B. Queue depth of value 4

In the Table 3, the values of SUs are given in the first column. In the second column, we have presented the number of magnetic disks on which the tests were performed for different SUs. The test procedure starts with the evaluation of 512 bytes and ends with 8192 KB, with a step in which the next data size is twice as the size of the previous one. The results for read (blue) and for write (red) are provided in MB/s.

The 8kB SU size (Figures 5 a1 and a2) of the disk array characteristics is completely unexpected. When reading, (Figure 5 a1), for workload smaller than 32kB the disk array is 2 to 3 times faster than a single disk. For workloads larger than 32 kB, only RAID 0 with 6 HDD has nearly 2 times the speed of a single disk, while other arrays act as a single disk and even worse. RAID 0 with 4 HDD has a higher speed than single disk of 20% to 25%. RAID 0 with 3 HDD has slower speed than single disk of 8% to 10% and RAID 0 with 2 HDD has slower speed than single disk of 40%.

For writing operation (Figure 5 a2), for workload smaller than 16 kB, the disk fields are 3 to 4 times faster than a single disk. For data blocks larger than 512 kB, only RAID 0 with 6 HDD has a 20% higher speed than for a single disk, while other arrays act as a single disk and even worse. When compared to the speed of a single disk, RAID 0 with 4 HDD has 20% higher speed, RAID 0 with 3 HDD has 40% slower speed and RAID 0 with 2 HDD has 60% slower speed.

For the SU size of 64 kB (Figure 5 b1 and b2), the characteristics of the disk array are almost as the expected. For reading operation (Figure b1), for workload smaller than 32 kB the disk array is twice faster than a single disk. For data blocks larger than 512 kB, all RAID 0 arrays are in accordance with theoretical expectations. RAID 0 with 4 and 6 HDDs have slightly lower read operations speed values than the theoretical ones, expected 4 and 6 times better then single disk and measured 3.7 and 5.5 times better than single disk.

TABLE III TEST RESULTS – QD 4

SU	MB/s	0.5 KB	1 KB	2 KB	4 KB	8 KB	16 KB	32 KB	64 KB	128 KB	256 KB	512 KB	1024 KB	2048 KB	4096 KB	8192 KB
8 KB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	3371	7296	16811	37145	59362	65209	76204	85111	84200	86659	88712	88885	88592	88738	87867
	3 HDD	3278	7260	16646	37328	65048	89219	111890	127875	136593	130745	119974	130944	130625	130308	128900
	4 HDD	3278	6759	17024	36499	65935	95209	103898	125128	90187	167508	179249	178362	178659	179555	175735
	6 HDD	3363	7207	17235	37145	56917	124792	183312	222508	243944	264628	270852	272062	272985	276737	258732
64 KB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	3303	6622	13050	26360	51447	105441	178028	245227	260087	284837	284346	284058	284560	284560	284058
	3 HDD	3215	6793	12987	26360	51697	106594	180215	323459	414686	419900	418607	418612	417798	419430	404422
	4 HDD	3404	6509	13278	25472	51200	104407	186434	337692	480534	485204	523259	522502	521233	521233	518715
	6 HDD	3041	6557	13312	25661	51200	106230	183369	344329	551708	749682	559848	715827	776198	781850	787585
1 MB	1 HDD	1670	3471	6925	13245	25356	46923	74642	115992	141214	141669	142056	141841	141841	141096	142217
	2 HDD	3242	6368	12892	24794	46180	80511	129134	187122	255086	283026	284346	284058	282563	284058	284058
	3 HDD	3011	6509	13019	25036	45732	79533	128817	184957	260087	406237	412940	389036	416197	416987	415374
	4 HDD	3223	6098	13147	25098	46293	80511	128187	188446	263969	418941	516925	514984	528936	528936	527637
	6 HDD	3176	6446	13050	24914	46293	79921	128187	188803	262014	409793	511966	697234	766958	778073	778073
8 KB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	196	441	660	1900	3792	5738	11358	21593	27947	27710	27623	28093	28035	28167	28137
	3 HDD	194	414	966	2022	2904	5748	11397	22253	35959	41487	42077	41975	41975	42074	42173
	4 HDD	187	398	894	1918	3515	5789	11437	15791	41217	51000	54050	55634	55749	55461	56099
	6 HDD	195	421	829	1969	2978	5840	11619	22954	44582	72415	82176	83624	83755	83755	83624
64 KB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	184	375	765	1606	3585	6742	14340	21881	38607	92794	152465	180764	180764	182920	186090
	3 HDD	179	371	757	1575	3323	7602	14928	22215	42625	80166	167375	267766	275789	273913	275318
	4 HDD	176	336	740	1551	3218	6313	13347	25954	43115	81537	160781	290200	354760	371965	371965
	6 HDD	171	169	750	1525	3162	6986	15312	23116	44506	84836	156052	302746	460833	536870	545600
1 MB	1 HDD	61	121	245	487	974	1918	3819	7447	14185	25826	43509	66858	66692	66774	91147
	2 HDD	181	358	718	1437	2899	5789	11457	22329	77710	107216	137100	155389	133549	282068	283060
	3 HDD	180	363	731	1432	2810	5748	11260	22140	78800	103614	158145	226050	305619	238080	412997
	4 HDD	181	366	723	1468	2925	5820	11457	22637	82091	105916	152684	243478	268883	267543	512525
	6 HDD	184	367	735	1468	2920	5563	11338	22520	85296	105278	149669	233422	387166	512525	517465

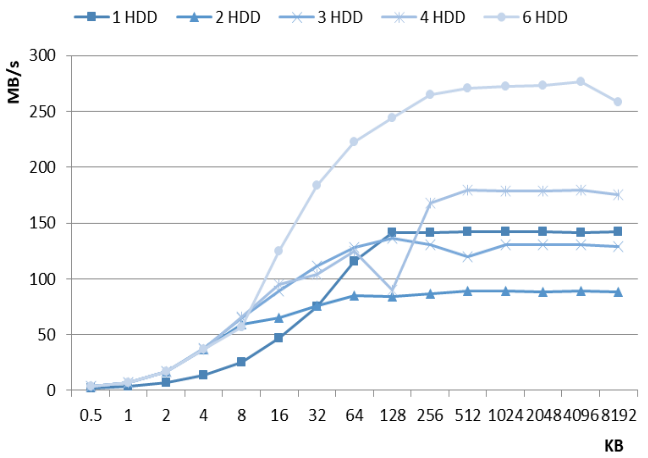

a1) Queue Depth 4, SU 8 KB

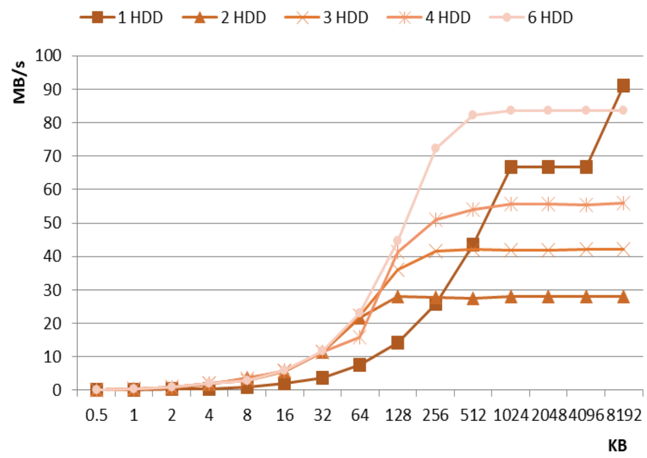

a2) Queue Depth 4, SU 8 KB

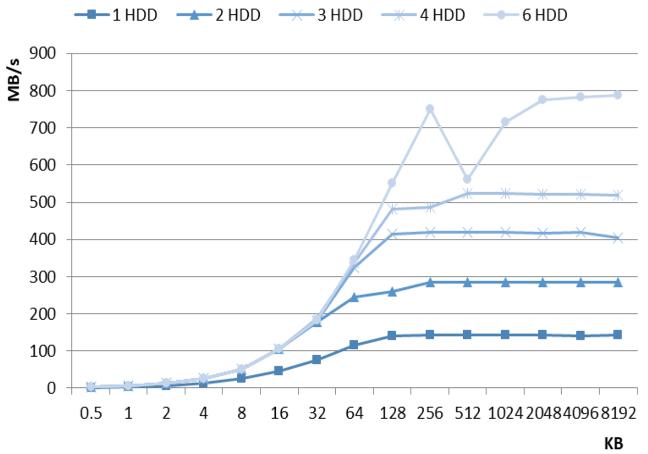

b1) Queue Depth 4, SU 64 KB

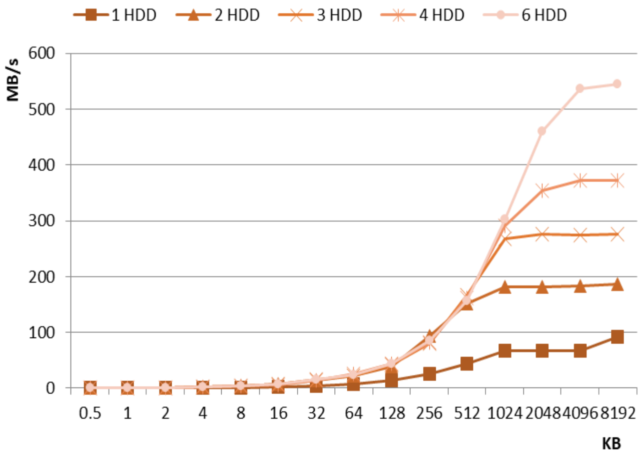

b2) Queue Depth 4, SU 64 KB

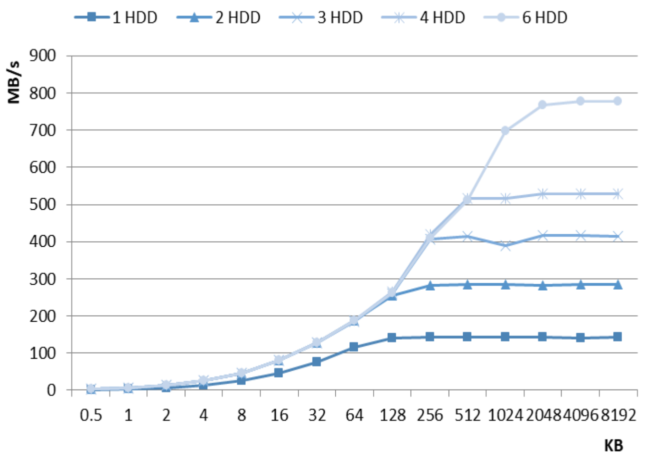

c1) Queue Depth 4, SU 1 MB

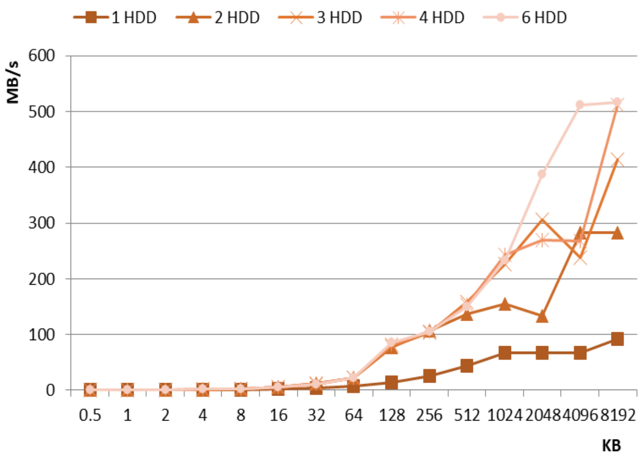

c2) Queue Depth 4, SU 1 MB

Figure 5. Reading speed 1) and writing speed 2) data for single magnetic disk and paired string of 2, 3, 4 and 6 magnetic disk drivers in RAID 0 at the SU size a) 8 KB, b) 64 KB and c) 1 MB

For writing operation (Figure 5 b2 and table 5) for the SU size of 64 kB and workload smaller than 512 kB, the disk arrays are 3 to 4 times faster than a single disk. For data blocks of 1 MB to 4 MB, all arrays behave better than theoretical expectations. When compared to single disk write operation speed, RAID 0 with 6 HDD has up to 8 times higher speed than single disk, RAID 0 with 4 HDD is 5.5 times faster than write operation for single disk, and RAID 0 with 3 HDD has 4 times the speed and 2 HDD is 3 times faster. This is possible because the maximum speeds of the device are not reached at the time of writing and parallelism allows simultaneous writing to 4 disks. Since the data blocks are from 16 to 128 times larger than the SU (64 kB), multiple consecutive SUs will be written on a sector without changing the track and with a possible small rotational delay (1/2r) or without delay on the successive SUs. The transmission time will be approximately $T = b/(rN)$, where b is the number of bytes to transmit, N is the number of bytes per track and r is the rotational speed (rpm). For blocks of 8 MB and higher, the speeds fell theoretically expected for the corresponding disk array as maximum performance was achieved.

Write cycles with small blocks of data can also have a big impact on write-back file caching, which can affect the exceedance of expected RAID accelerations, which cannot be manifest for small reads.

For a SU size of 1 MB, the characteristics of the disk array are given in Figures 5 c1 and c2. For a read operation (Figure 5 c1) and workload smaller than 128 kB the disk array is faster about 2 times than a single disk. For workloads larger than 512 kB, all RAID 0 array behave in accordance with theoretical expectations. RAID 0 with 4 and 6 HDDs have slightly lower read operation values than the theoretical ones, the expected values are 4 and 6 times faster than a single disk, while the measured are 3.75 and 5.5 times faster.

For writing operation (Fig. 5 c2) and workload data blocks smaller than 64 kB, disk array is faster 3 times than a single disk. For data blocks larger than 64 kB, all RAID 0 arrays have a higher speed than a single disk, but the diagrams do not show regularity in the speed increase. RAIDs with 3 and 4 disks are about 4 times faster than single disk speeds.

In all the performed tests, it is noticeable that the writing performance results are far worse than the results obtained for reading operation. One of the reasons for these results is that during the measurement the force write access option was activated. It obviates the use of caching on a disk or controller when writing data, while the caching effects during the reading operation are noticeable.

VI. CONCLUSION

Based on the measured values (Table II and III) magnetic disks have the best performance with larger blocks of workload data (large files) stored on successive disk blocks (SUs), which means in sequential access. That is mostly due to their rotating surfaces. In that way, the write/read head can start and end the transfer in one positioning to the desired cylinder. When hard drives carry very large amounts of data, large files can be scattered across the disk on different cylinders (this is true for random storage and for small files), causing data access to be much slower. Performance degradation can also be caused by the ever-present fragmentation. Due to the seek time, which can significantly affect the total data transfer time and therefore the speed, HDDs can support a relatively small number of I/O operations per second, especially in a random access environment.

The testing results have shown that the RAID 0 operation depends on four factors: queue depth, amount of data, SU size and the number of disks. Inadequate block size selection for a RAID 0 string can significantly degrade the performance of the system. In the case of an adequate configuration, RAID 0 shows a direct gain in performance when increasing the number of disks (equations 1 and 2). The performance gain is approximately the same for data reading and writing operations.

The measurements shown in Tables II and III confirmed that SU size and the workload size significantly affect the transmission rates of RAID 0 and show that the theoretical values were reached for four disk (N=4). The measurement shows that the best features of RAID 0, for this configuration, were achieved by selecting a 64 kB stripe unit (Figures 4 and 5, b1 and b2).

Also, the diagrams in Figure 5 show that for queue depth 4 RAID 0 array and for small block data read operation is better 2 times and write data operation better 3 times than in a single disk, or when there is no overlap of I/O operations (queue depth 1). At queue depth 4 performance improvement occurs with smaller blocks of data than when there is no overlapping of I/O operations.

ACKNOWLEDGMENT

The work presented in this paper has partially been funded by the Ministry of Education, Science and Technological Development of the Republic of Serbia (Grants TR32025, TR32037, III43002 i III44006).

REFERENCES

[1] D. Patterson, G.A. Garth A., R. Katz, “A Case for Redundant Arrays of Inexpensive Disks (RAID)”, [Online]: www.eecs.berkeley.edu/Pubs/TechRpts/1987/CSD-87-391.pdf, dec. 2018.

[2] W. Stallings, Computer Organization and Arhitercture, Pearson Education, Inc., Hoboken, 2016

[3] Oracle Solaris ZFS, [Online]: https://docs.oracle.com/cd/E18752_01/html/819-5461/gavwn.html, february 2019.

[4] Microsoft storage, [Online]: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-fault-tolerance, decembar 2018.

[5] Software RAID, [Online]: https://www.softraid.com/pages/features/software _raid_benefits.html, december 2018.

[6] V. Timcenko, B. Djordjevic, “The comprehensive performance analysis of striped disk array organizations - RAID-0, ” invited paper, in proc. of Proceedings of the 2013 International Conference on Information Systems and Design of Communication, Lisbon, Portugal, 2013.

[7] ATTO benchmark, [Online]: https://www.atto.com/disk-benchmark/, february 2018.

[8] F. Wan, N.J. Dingle, W.J. Knottenbelt, and A.S. Lebrecht, “Simulation and modelling of RAID 0 system performance,” In 22nd Annual European Simulation and Modelling Conference (ESM), 2008.

[9] N. Davidović, B. Đorđević, V. Timčenko, S. Obradović, B. Škorić, “Data storage system on paired magnetic disk arrays - RAID 0”, Proc. 18th International Symposium INFOTEH-JAHORINA, 20-22 March 2019.

[10] N. Davidović, B. Đorđević, V. Timčenko, S. Obradović, B. Škorić, “RAID 0 on paired magnetic disk arrays“, Proc. 6th IcETRAN, Silver Lake, Serbia, June 03 – 06, 2019

Nikola Davidović received his Diploma and M.S. degree in electrical engineering from the Faculty of Electrical Engineering, University of East Sarajevo. He is a Teaching Assistant at the Faculty of Electrical Engineering, University of East Sarajevo. His research interests include computer architecture, operating systems, secondary memory devices, file systems, and storage systems.

Slobodan Obradović received his Diploma and M.S. degree in electrical engineering from the School of Electrical Engineering, University of Belgrade. He received Ph.D. degree from the Faculty of Electrical Engineering, University of East Sarajevo and works as a Professor at the same institution.

Borislav S. Đorđević is a member of the IEEE for one year. He received his Diploma and M.S. degree in electrical engineering and Ph.D. degree in computer science from the School of Electrical Engineering, University of Belgrade, in 1989, 1992 and 2003, respectively. He is a Research Associate at the University of Belgrade, Institute Mihajlo Pupin, Belgrade, Serbia. His research interests include computer architecture, operating systems, hard disks, file systems, storage systems, virtualization techniques, and cloud computing.

ValentinaV. Timčenko received her Diploma and M.S. degree in electrical engineering from the School of Electrical Engineering, University of Belgrade, in 2004 and 2010 respectively. She works as a Research Associate at the University of Belgrade, Institute Mihajlo Pupin, Belgrade, Serbia. She has participated in several research projects and studies concerning NGN design and network management systems. Her scientific and professional commitment includes research, design, and implementation of solutions for telecommunication networks, especially in the area of mobile ad hoc networks. She is currently pursuing Ph.D. degree in computer science at the School of Electrical Engineering, University of Belgrade, Serbia. She is a member of the IEEE for 18 years.

Bojan Škorić, 21. 06. 1983., Beograd, Serbia. Graduated from Fifth Belgrade High School and The School of Electrical and Computer Engineering of Applied Studies. Initial experience in information technology and computer science gained through years of freelancing in technical support services. Worked for companies like P3 Communications GmbH (part of P3 Group A.G., now Umlaut) as system administrator, Informatika a.d., Beograd and NCR Corporation as system engineer. Focused on system infrastructure engineering and cloud services.