Energy and Performance Evaluation of Lossless File Data Compression on Server Systems

Appears in the Proceedings of the Israeli Experimental Systems Conference (ACM SYSTOR '09)

Rachita Kothiyal, Vasily Tarasov, Priya Sehgal, and Erez Zadok
Stony Brook University
{rachita,vass,psehgal,ezk}@fsl.cs.sunysb.edu

ABSTRACT

Data compression has been claimed to be an attractive solution to save energy consumption in high-end servers and data centers. However, there has not been a study to explore this. In this paper, we present a comprehensive evaluation of energy consumption for various file compression techniques implemented in software. We apply various compression tools available on Linux to a variety of data files, and we try them on server class and workstation class systems. We compare their energy and performance results against raw reads and writes. Our results reveal that software based data compression cannot be considered as a universal solution to reduce energy consumption. Various factors like the type of the data file, the compression tool being used, the read-to-write ratio of the workload, and the hardware configuration of the system impact the efficacy of this technique. In some cases, however, we found compression to save substantial energy and improve performance.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Performance Attributes; E.4 [**Coding and Information Theory**]: Data Compaction and Compression; H.3.4 [**Information Storage and Retrieval**]: Systems and Software—*performance evaluation (efficiency and effectiveness)*

General Terms

Measurement, Performance

Keywords

Data compression, Energy, Performance evaluation, Storage

1. INTRODUCTION

Until recently, power management research was mostly directed towards battery powered portable computers and mobile devices [4, 33, 38, 39, 52, 58, 59]. The motivation behind these efforts has been to enhance user satisfaction by reducing the frequency of battery recharges. However, the growing costs of power and cooling have now caused researchers to look at the same issue on desktops and commercial servers [10, 11, 18, 26, 28, 41, 49, 68]. Data centers and servers primarily deal with data. Data compression has been suggested an effective way of saving energy in such systems. To the best of our knowledge, there has not been a study evaluating these claims.

In this paper, we study several compression algorithms, implemented in software, applied to various types of data files, and evaluate all in terms of performance and energy metrics. We compress these data files at the CPU, and write the compressed file to the disk; the compressed file is read from the disk, and decompressed at the CPU. We use four different types of files for our experiments: *zero*, *text*, *binary*, and *random*. These file types exhibit different levels of data redundancy, with zero being the highest and random being the lowest. Our benchmarks include four popular compression utilities on Linux: gzip, lzop, bzip2, and compress. File compression is known to be computationally intensive, but can reduce the amount of I/O being incurred due to a reduction in file size. The aim of this study is to evaluate each of the compression tools, and determine if the savings due to reduced I/O (both in time and energy) are worth the added overhead at the CPU and memory. To be able to view the effects of compression/decompression on energy and performance simultaneously, we use the energy-delay product metric [23] for our analysis.

Our results reveal that software based data compression cannot be considered a universal solution to reduce energy consumption in data centers and server class machines; it greatly depends on the type of data files being compressed, the compression algorithm applied, the workload of the system, and the hardware configuration. As we expected, compressing zero files was found to almost always save energy, compared to raw reads and writes, no matter what compression algorithm was used. We realize that such high levels of redundancy are not common in real-life settings, but we include it in our study to evaluate the best-case scenarios. Second to zero files, we observed that text files exhibited the most potential for energy savings by compression, followed by binary files. Although some utilities always performed better than plain writes and reads for text files, other tools required some number of reads for every write to result in energy savings. This is because compression typically consumes more CPU than decompression. To represent the possible savings in such cases, we developed a simple read-write model: it calculates the minimum number of decompressions required to offset the extra energy expended by a single compression. This number can be useful in deciding whether or not a workload whose read-to-write ratio is known would benefit from compressing its data files using a particular compression tool. Finally, also as expected, random files showed no energy or performance benefits upon compression. Again, we included random files to be able to evaluate the worst-case scenarios for compression.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SYSTOR '09, May 4–6, Haifa, Israel
Copyright 2009 ACM X-XXXXX-XX-X/XX/XX ...$5.00.

The rest of the paper is organized as follows. Section 2 provides some background and discusses related work in the area. In Section 3 we talk about the various metrics used for evaluating our results. We describe the details of the experimental methodology in Section 4. We present the actual experimental results obtained from the various benchmarks and the read-write model of evaluation in Section 5. We conclude in Section 6.

2. BACKGROUND AND RELATED WORK

Section 2.1 provides an overview of some existing power management techniques in computing systems. In Section 2.2, we present various power management solutions for primary storage media. In Section 2.3, we address compression techniques implemented at various levels and their energy impact. We also draw out important distinctions between our work and other research in this area.

2.1 Power Management Approaches

Energy management techniques can be implemented at several levels in a computer system. The fundamental idea behind these approaches has been to transition a component to a lower power mode or to turn it off completely when not in use. Lorch et al. discuss software techniques to utilize the power saving provisions provided by the various hardware components, such as the CPUs, disks, displays, wireless communication devices, main memory, etc. [38]. Dynamic Voltage and Frequency Scaling (DVFS) techniques have been widely employed for reducing CPU power consumption [10, 58, 59]. DVFS allows processors to dynamically switch to different operating voltages and frequencies. Choosing a lower voltage would translate to a reduction in power consumption. However, since voltage cannot be changed independent of the frequency, it would also result in some degree of performance degradation. Several processors support *Clock Gating* as a means to halt idle components, and save power [15, 22, 25, 48].

Su et al. proposed and evaluated several CPU cache designs based on *Gray codes* and cache organization [52]. As Gray codes require only one bit modification to represent consecutive numbers, Su et al. were able to obtain significant energy savings because of reduced bit switching. They also found that cache sub-banking [53] (i.e., organizing cache into banks), was an effective way to reduce energy consumption of caches. Power Aware Page Allocation [33] reduces the memory energy consumption by adding energy awareness to the operating system's virtual memory page allocator. The authors explored various page allocation policies to harness the power management features of emerging DRAM devices.

The OS has also been used to monitor the usage of hardware resources, in order to transition the components to low power states during periods of inactivity [4, 19, 39]. Zeng et al. propose an Energy-Centric Operating System (ECOSystem), which allows energy to be managed as a first-class resource by the OS [67].

2.2 Energy Saving Techniques for Storage

One of the earliest ideas for energy conservation in disks was to spin them down when idle. The controls on when to spin them down have ranged from simple threshold-based policies to intelligent prediction policies [16, 17, 36, 61]. Techniques such as Massive Array of Idle Disks (MAID) [11], Popular Data Concentration (PDC) [44], and write offloading [41] are based on the idea of directing the requests to a subset of the disks or special logging devices. This increases the idle time between requests, hence justifying the spin down of the unused disks. GreenFS [28], is a stackable file system for client systems. It services I/O requests from remote servers in addition to adding a flash layer to the storage hierarchy. In enterprise settings, with existing backup server infrastructure already in place, the energy cost of network transfers for small transfers is much smaller than spinning up and writing to the local disk. This allows the hard disks to be powered down for longer, and hence save more energy. Many vendors, (e.g., NetApp, EMC, etc.,) provide a large NVRAM to cache disk writes.

Analogous to DVFS for CPUs, Gurumurthi et al. [49] proposed disks which can dynamically change their rotation speeds based on the request traffic, thereby lowering their power consumption. Zhu et al. [68] considered storage cache replacement techniques to selectively keep some blocks from the disk in the main memory cache, to increase the disk's idle times; this allows disks to remain in low power mode for longer.

Another approach taken by many researchers, distinct from the disk spin-down policy, has been to reduce the energy consumed by head seek operations. Essary et al. present a Predictive Data Grouping technique [18] which attempts to co-locate related data blocks on disk through remapping and replication. Huang et al. proposed a file system, FS2 [26] which dynamically replicates data so that the nearest copy of the data can be served on a request. As the mechanical movement of the disk head is reduced by these techniques, it results in power savings. Interestingly, the increased proximity of the data to the disk head also reduces the seek and rotational delays, which translates to better performance.

2.3 Saving Energy using Compression

Compression has been widely used to reduce traffic and latencies on communication channels (Data bus, network, etc.) [6, 9, 27, 32, 62], and save storage space [2, 46]. Over the last decade, compression has been implemented at various levels of the memory hierarchy and proved to be a successful method of saving energy. For example, several encoding schemes have been proposed for compressing the contents of the CPU instruction cache [7, 8, 34, 65]. These techniques, called *code compression*, map the program instructions into a set of much shorter instructions, thereby reducing the memory requirements and bus traffic. A decompressor, typically between the cache and the CPU, translates the compressed instructions to the normal program instructions before execution on the CPU. Various compression algorithms have been employed on CPU data caches as well [30, 31, 54, 56].

Benini et al. propose a hardware implementation of the compression-decompression logic between the main memory and the CPU cache for embedded processor systems [5]. On a cache write-back, compressed data is written to main memory, while decompressed data is written from main memory to the cache. IBM's Memory Expansion Technology (MXT) [55] has made main memory data compression commercially available to a wide range of systems. Kandemir et al. extend compression to multi-bank memory systems, by compressing infrequently used data, and transitioning those banks to lower power mode after a threshold idle time [29]. Sadler et al. employ lossless compression on data communication in sensor networks to reduce energy expenditure [50].

The work most closely related to ours, albeit in a different environment of embedded and mobile devices, is that of Barr et al. [3]. Because the energy cost of a single bit wireless transmission is many times that of a single 32-bit computation, they apply lossless compression techniques to data before transmitting. In their work, Barr et al. analyze various data compression algorithms from the energy perspective. They found `lzop` and `compress` to be the most energy efficient. Unlike their work, our study is focused on server class machines and file compression. Our goal, however, is to investigate potential energy savings in the storage stack, rather than from transmission over a network. Furthermore, we found that

only lzop can be widely applicable in such environments.

Another related work by Xu et al. [64] explores data compression as a means to reduce battery consumption of hand-held devices when downloading data from proxy servers over a wireless LAN. They assume that the data from the proxy server is available in compressed format and hence focus their study only on the energy costs related to decompression. Our target systems differ from theirs in that our systems would have to incur the costs of both reads and writes. Hence, we take into account the energy costs of both compression and decompression in our analysis.

Data compression for storage can be implemented at both hardware [13, 42] and software levels. However, in this work we focus our analysis on software implementations only, so as to minimize variations due to hardware changes.

3. METRICS

The increasing number of studies in the area of green technologies have revealed a problem of lack of agreement on a proper metric for comparing energy efficiency of computer systems. The choice of an appropriate metric depends on several factors: which component of a system the metric will be applied to, what are the purposes of comparison and how different are the systems. For our study, the metric must be generic enough to express the energy efficiency of the system as a whole. We would also like the metric to be usable for different purposes of comparison. Therefore, we present in this section, not one, but several metrics based on a simple view of a computer system. This family of metrics allows us to describe the energy efficiency profile of a system from several angles, which offers enough scope for a broad analysis.

We define a *system* as any device capable of performing computational work. The work is provided to the system by a user as a list of *tasks*. A task is a logically independent unit of work that the user wants the system to perform. The rational metric representing the performance of such a system is its computational power: the number of tasks the system is able to perform in a unit of time. For the purposes of this paper, however, it is more convenient to use the inverse value of computational power: the time required to finish a single task. We denote this value as T and measure it in seconds per task. Notice that the notion of a system and a task are highly conceptual here. Depending on the specific scenario, the system can be a CPU that is executing instructions, a disk drive performing I/O requests, a server executing compression algorithms on a piece of data and writing the results to a disk, and more.

While performing computational work, the system consumes electrical energy. In other words, the system converts electrical energy (typically measured in Joules) to computational work (measured in tasks accomplished). In terms of power consumption, we are mostly interested in the effectiveness of this conversion: the number of tasks the system is able to perform by using a unit of energy—or in its inverse form—the energy consumed by the system to perform a single task. We denote the latter value as E and measure it in Joules per task. Figure 1 provides the system view we used in our study.

Many projects use the plain metric E to compare energy efficiency of different systems [12, 24, 58, 59]. However, this metric ignores the amount of time it takes to complete a task, T. For example Gonzaleze et al. [23] showed that it is fairly easy to improve a processor's energy efficiency E, but it typically leads to degraded performance of the chip. Sometimes, it is reasonable to ignore T. For instance, when each system already has the desired performance characteristics [12]. However, in some cases we would like to have a unified metric that gives us a solid understanding of both the system's energy efficiency and its performance. For such

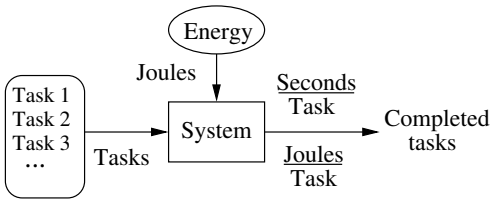

Figure 1: System view for energy efficiency estimations.

a metric we need to take into account both quantities. It is useful to know, for example, how many tasks per Joule per second the system can produce:

$$\frac{Tasks}{Joules \times Seconds}$$

This metric has a clear physical meaning: given its value, one can multiply it by the amount of energy and time, and obtain the number of tasks the system is able to perform under these constraints. The inverse of this metric can be written in the following form:

$$\frac{Joules \times Seconds}{Tasks} = \frac{Joules}{Tasks/Seconds} = \frac{Joules}{Throughput}$$

Again, this number has a natural meaning: how many Joules we pay for the speed of execution of a task, as tasks per seconds in the denominator is the throughput of the system. This metric is widely known as *energy-delay* [23, 37, 69]. We denote it as ET.

We believe that omitting any of the metrics represented above (T, E, ET) takes away valuable information about the system. T gives a good understanding of performance, but does not convey power consumption. E provides reliable information about energy efficiency, but ignores the performance. The ET metric has an intuitive underlying physics and is valuable to compare systems in the general case, but is not applicable when one is interested in energy savings or performance only. For these reasons we adopt all three metrics in this paper. The metrics we used are summarized in Table 1. By convention we omit tasks unit from the table, as all units are implicitly per task.

Metric	Notation	Unit
Time	T	$Seconds$
Energy	E	$Joules$
Energy-delay	ET	$Joules \times Seconds$

Table 1: The metrics and corresponding units we used to evaluate performance and energy efficiency of a system.

Another well-known metric of energy efficiency is $energy \times delay^2$ [40]. However, it is specific to the situations where the voltage applied varies from system to system, for instance for comparison of different DVFS levels. Therefore, we do not use this metric in the paper.

4. EXPERIMENTAL METHODOLOGY

This section details the setup used for our evaluations. We describe our testbed and the instruments used for energy measurement in Section 4.1. In Section 4.2 we present the types of files and the compression tools we examined. We describe the various benchmarks and the motivation behind their selection in Section 4.3.

4.1 Experimental Setup

We used two different machines for our experiments. The first was a Dell PowerEdge SC1425 rack-mountable server, with 2 dual-core Intel® Xeon™ CPUs at 2.8GHz, 1GB RAM, 73GB primary

hard disk (SCSI Seagate ST373207LW, 10000 RPM) and a dedicated 20GB partition on a separate hard disk (SCSI Seagate ST373207LW, 10000 RPM) for the tests. The server was running the Fedora Core 6 (kernel 2.6.20-1.2952.fc6) distribution of Linux. In order to simplify our evaluation, we enabled only one processor unit by using the `maxcpus = 1` boot time parameter in Linux.

The second machine was a desktop system, with an Intel® Pentium® CPU at 1.7GHz, 1GB RAM, 20GB primary hard disk (WDC WD200BB-00AUA1, 7200 RPM) and a 20GB test partition on a separate disk (Maxtor 6E040L0, 7200 RPM). It was running the same 2.6.20-1.2952.fc6 Linux kernel as the server.

As our goal is to study the energy impact of data compression on the entire system, and not on a component in isolation, we measure the total energy of the machine. Hence, we used a WattsUP Pro ES [57] power meter to measure the energy consumption of the system under test, instead of a current clamp attached to a digital multimeter [20, 21], which can provide component level energy measurements. The WattsUP Pro ES is a plug-in style power meter, which allows power measurements by plugging in the AC supply of the test machine in the meter's receptacle. It calculates the cumulative energy in Watt-hours (1 Watt-hour = 3,600 Joules) every second, and stores it in its non-volatile memory. It has a 1 second time resolution and a 0.1 Watt-hour (360 Joules) resolution for energy measurements; it has an accuracy of $\pm 1.5\% + 3$ counts of the displayed value. We used a `wattsup` Linux utility [60] to download the recorded data from the meter over a USB interface to the test machine.

4.2 File Types and Compression Tools

Power consumption in the evaluated systems depends on the effectiveness of compression, which is typically measured by *Compression Ratio* (CR) defined as:

$$CR = \frac{Originalfilesize}{Compressedfilesize}$$

Compression ratio is heavily affected by the type of input data file. Hence, we include the file type as one of the dimensions for our evaluation. In order to have a representative set of possible data files, we chose to run the tests on four types of files of size 2GB each: zero, text, binary, and random. These files denote the best-to-worst cases of compression, in order. We chose the file size to be 2GB to ensure that each test ran for a considerable amount of time, thereby reducing the scope of errors and high standard deviations arising out of even slight differences in recorded values across multiple iterations of the test. Also, the 2GB file, being larger than the system RAM (1GB), forces I/O to take place. We created the *zero* file by writing zeroes to the file. We generated the *text* file by concatenating source files from the Linux kernel and other open source projects. We created the *binary* test file by combining object files from the Linux kernel, Linux libraries and other open source executables. We created the *random* file by reading from `/dev/urandom`. All the files were generated before running the benchmarks, so the time/energy required for the generation was not included in the measurements.

Another factor influencing the compression effectiveness is the compression algorithm itself. This constitutes the second dimension of our analysis. We examined four popular compression utilities available on Linux: compress, gzip, lzop [43], and bzip2. They have significant differences in implementation and cover a wide range of compression algorithms. Bar and Asanovic discuss these tools and their algorithms in detail [3]. The compress utility, regarded as the oldest, implements the Lempel-Ziv-Welch (LZW) algorithm which is a variant of the LZ78 algorithm. It uses *m* bits (9–16) to encode the input symbols, and stores the string-to-code mapping in a dictionary. Although based on the same LZ77 algorithm, gzip and lzop differ significantly in their implementation. As lzop was designed with the main goal to improve compression/decompression speed, it tends to be generally faster than gzip. The bzip2 utility is based on the Burrows Wheeler Transform (BWT); it achieves better compression ratio than the Lempel-Ziv based tools, at the expense of compression speed. The block size for compression (commonly 100k–900k) can be specified at command invocation. A larger block size typically increases the compression ratio, while increasing the memory footprint.

4.3 Benchmarks

Writing an uncompressed file involves reading the input and writing it to disk. We will refer to this as a *plain-write* in the rest of the paper. Writing a compressed file involves reading the uncompressed input, compressing it, and writing the compressed file to disk. We shall call this *compress-write*. Similarly, we use the term *plain-read* to denote reading the uncompressed file from the disk; and we use *decompress-read* to indicate reading the compressed file and decompressing it. Each of the operations described above (plain-write, compress-write, plain-read, and decompress-read), constitute a task which we defined in Section 3.

The aim of this study is to compare a *plain-write* to *compress-write*, and a *plain-read* to *decompress-read*, in terms of both energy consumption and performance. We therefore broadly have four types of benchmarks: *plain-write, compress-write, plain-read, and decompress-read.* As mentioned above, we used four different compression tools, each of which can be invoked with tunable parameters. For example, gzip allows the user to specify an *effort* parameter in the range 1–9 to choose between speed of compression and compression ratio; a choice of 1 would result in fast compression, but poorer compression ratio; and a 9 would give the best compression ratio, but would be slower than 1. Table 2 lists the various parameter values considered for the compress-write benchmarks. For each of the compression tools we chose the default invocation, and the options which provide the best and worst case of compression speed (if not already covered by the default option). Table 3 lists the compression ratios achieved by compressing different types of files using various compression applications.

Invocation	Implications
gzip –1	Favors speed over compression ratio
gzip –6	Default
gzip –9	Favors compression ratio over speed
lzop –1	Favors speed over compression ratio
lzop –3	Default
lzop –9	Favors compression ratio over speed
bzip2 –1	Use 100K block size
bzip2 –9	Use 900K block size (default)
compress –b 10	Use 10 bit codes
compress –b 16	Use 16 bit codes (default)

Table 2: Parameters used for invocation of various compression tools for compress-write benchmark.

We used the Auto-pilot test suite infrastructure [63] to run the benchmarks. Auto-pilot measures the time required to run a benchmark and reports it in terms of Elapsed, System, User, and Wait times. We developed an Auto-pilot script plug-in to measure the energy consumed while running the benchmark. The plug-in relies on the Linux utility described in Section 4.1 to communicate with the meter. The plug-in uses the utility to send a command to clear

Tool	File Type			
	Text	Binary	Rand	Zero
None	1	1	1	1
gz-1	4.16	1.81	0.95	~10^2
gz-6	4.79	1.81	0.95	~10^3
gz-9	4.84	1.81	0.95	~10^3
lzo-1	3.52	1.53	0.95	~10^2
lzo-3	3.51	1.53	0.95	~10^2
lzo-9	4.37	1.81	0.95	~10^2
bz-1	5.09	1.81	0.95	~10^6
bz-9	6.11	2.09	0.95	~10^7
c-10	1.17	1	0.8	~10^2
c-16	2.07	1.17	0.8	~10^5

Table 3: Compression ratios achieved by various compression utilities on 2GB files.

the meter's internal memory before starting the benchmark. After the benchmark has finished execution, we invoke the utility to send a command to read the data from the meter, and extract the total energy expended (in Joules) while running the benchmark. Since the benchmark themselves run for a significant time, any energy measurement errors due to the measurement tool itself are negligible.

We ran all tests at least five times and computed the 95% confidence intervals for the mean elapsed, system, user, and wait times using the Student's-t distribution. In each case, unless otherwise noted, the half widths of the intervals were less than 5% of the mean. In all bar graphs, we show the half widths using an error bar. Wait time is elapsed time less system and user time and mostly measures time performing I/O, though it can also be affected by process scheduling.

We ran the tests on a dedicated hard disk, with the partition formatted with the Ext2 file system and mounted using the default options. To ensure that writes to the partition were flushed to the disk during our measurements, we unmounted the partition at the end of each test iteration.

5. EVALUATION

In this section, we evaluate the effect of compression and decompression on energy savings and performance, based on the metrics: energy (E), time (T), and energy-delay (ET), as discussed in Section 3. Section 5.1 explains the terms we use later. Section 5.2 presents our read-write model. Section 5.3 analyzes the results of the compression utilities for text files, on both server class and desktop machines. As we found similar results on both the classes of machines, Sections 5.4, 5.5, and 5.6 evaluate the results on server class machines only, for binary, random, and zero files, respectively. Finally, Section 5.7 summarizes the evaluations.

5.1 Terminology

Figures 2, 4, 5, and 6 show the metrics plotted for text, binary, random, and zero files, respectively. In all these figures, the x-axis denotes *alg-mode-level*, where *alg* is the type of the compression/decompression algorithm: gzip, lzop, bzip, or compress; *mode* is either Compression or Decompression; *level* is passed as a parameter to the compression/decompression algorithm to control the compression ratio (CR). Similarly, we use the notation *alg-level*, to refer to a given tool operating at a specific compression level.

The time result figures show the total time required to compress-write or decompress-read a 2GB file using the compression utilities discussed above, compared to plain-writes and plain-reads, respectively. The y-axis on this graph denotes the elapsed time, which constitutes of the *system* time, *user* time, and *wait* time.

The second type of metric plotted is the energy results. These results compare the total energy required in performing a plain-write/plain-read versus a compress-write/decompress-read. On the y-axis we have the total energy, constituting of *active* and *passive* energy. Passive energy is the energy that is consumed by an idle system, for the elapsed period, without any other activity. For calculating the passive energy, we first need to estimate the average power consumption of the idle system. To compute this, we let the system idle ten times for 10 minutes each, computed the average idle power, and we verified that the standard deviations were small. We then divided the total energy measured by the duration of the idleness, yielding the average idle power of the system. Passive energy can be obtained by multiplying the average idle power with the elapsed time. Active energy is the extra energy required, apart from the passive counterpart, to complete the required task. In our graphs, we represent energy in units of Kilojoules, where $1Kilojoule = 10^3 Joules$.

The *energy-delay product* (ET) metric, as discussed in Section 3, compares the ET results of compression/ decompression versus pure writes/reads. Similar to the energy results, the total ET also consists of an active and passive component. We have plotted the ET results in units of Kilojoule-seconds.

5.2 Read-Write Model

The best case for compression would be when compress-write outperforms plain-write, and decompress-read fares better than plain-read, in terms of a metric. However, there might be scenarios when only one of these comparisons favor compression. For example, for a given compression tool, compress-write might require more energy than plain-write, but expends less energy for a decompress-read than a plain-read. Notice that the metric we consider in this example is energy, but the argument applies to the other metrics as well (e.g., time or energy-delay). Compression might still achieve energy savings in such a case if the number of reads is more than a "break-even" value to amortize the extra energy consumed by a single compress-write.

Workloads are characterized by a *read-to-write* ratio (n), which represents the distribution of read and write I/O requests. There have been extensive studies to characterize workloads based on this parameter [47, 35]. Given a workload, with knowledge about its read-to-write ratio and the type of file data it handles, we can use this break-even value (n_{be}) to decide if compressing the data files would be beneficial. We formalize this by the following model.

For a given metric M, let M_w, M_c, M_r, and M_d be the measured values of M on a plain-write, compress-write, plain-read, and decompress-read, respectively. Let n_{be}^M represent the break-even read-to-write ratio to obtain energy savings. Assuming we first need to write once before reading, the following inequality must hold to compensate the excess energy expended during the write:

$$M_c - M_w \leq n_{be}^M \times (M_r - M_d)$$

Solving for n_{be}^M, we get

$$n_{be}^M \geq \frac{(M_c - M_w)}{(M_r - M_d)}$$

where $M \in \{T, E, ET\}$

We calculate and present the n_{be}^M values for the T, E, and ET metrics for the various compression tools and test files in Tables 4 and 5 of Section 5.3.

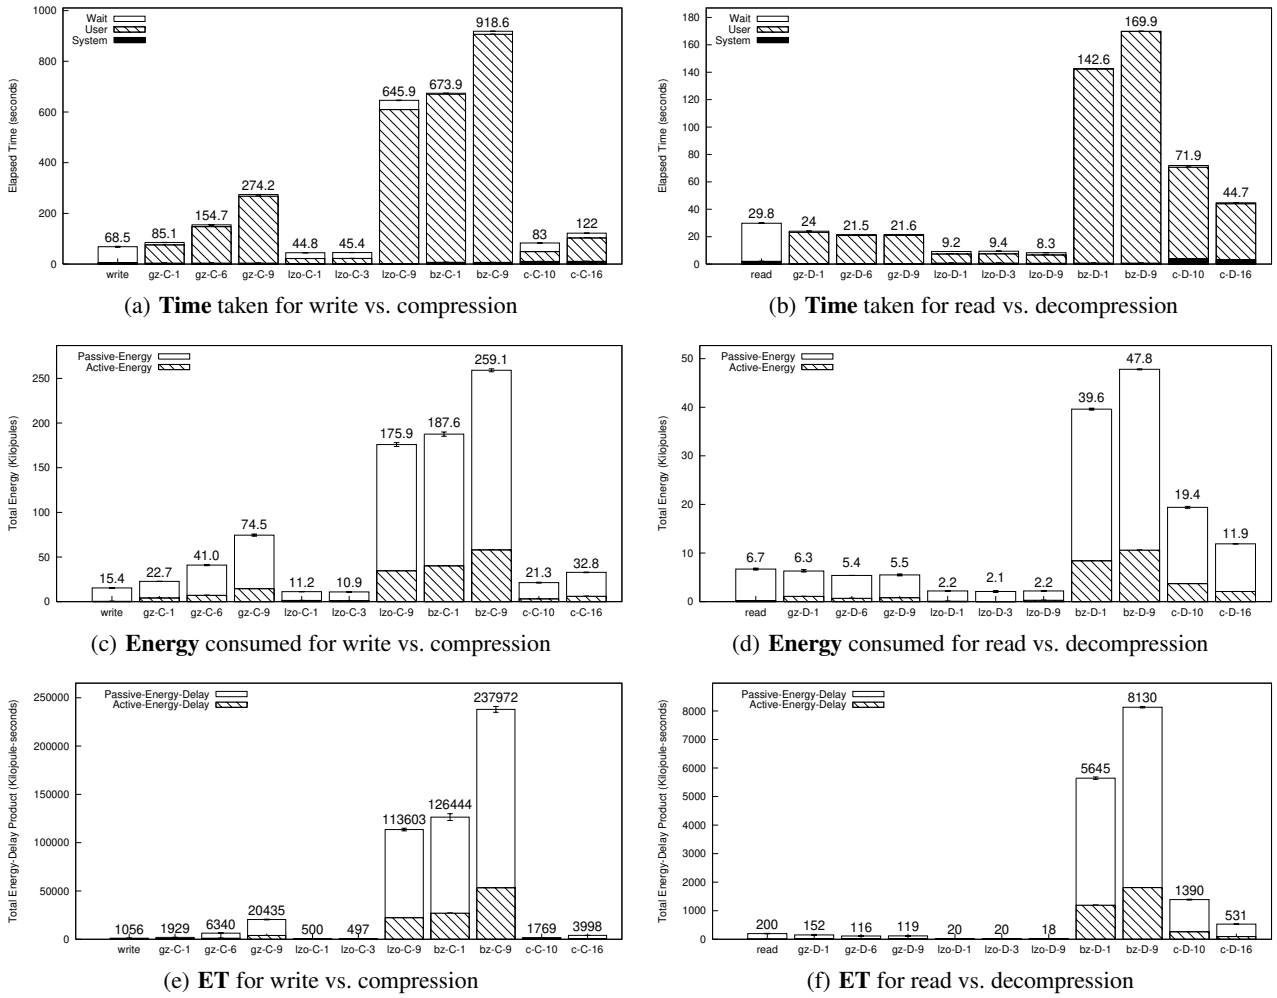

(a) **Time** taken for write vs. compression

(b) **Time** taken for read vs. decompression

(c) **Energy** consumed for write vs. compression

(d) **Energy** consumed for read vs. decompression

(e) **ET** for write vs. compression

(f) **ET** for read vs. decompression

Figure 2: Text file: The time, energy and energy-delay product (ET) for compressing/decompressing a 2GB text file. In (a) and (b), the values at the x-axis are of the form *alg-mode-level*, where *alg* is the type of the compression/decompression algorithm: gzip, lzop, bzip, or compress; *mode* is Compression or Decompression; *level* is the parameter passed to the compression/decompression algorithm. The energy results in (c) and (d) represent the total energy (kilojoules) required for compressing/decompressing the file with the corresponding algorithm. Section 5.1 describes the terms *Active-Energy* and *Passive-Energy*. In (e) and (f) we use kilojoule-seconds to denote the energy-delay product to compress-write/decompress-read the same file.

We define the energy savings (E_{sav}) for decompress-read and compress-write vs. plain-read and plain-write for a given value of read-to-write ratio, n:

$$E_{sav} = (n \times (E_r - E_d)) + (E_w - E_c)$$

where, E_w, E_c, E_r, and E_d is the energy expended in plain-write, compress-write, plain-read, and decompress-read, respectively.

Note that a negative value of E_{sav} means energy loss. Figure 3 presents the values of E_{sav} for n ranging from 0 to 30 for different compression algorithms applied on a text file. The value of n for which E_{sav} becomes zero is the n_{be}^{E}.

5.3 Text File Analysis

As we observe in Figures 2(a), 2(b), 2(c), and 2(d), lzo-1 and lzo-3 always outperform pure writes and reads in terms of both time and energy consumption for text files. Both lzo-1 and lzo-3 save approximately 29% energy compared to plain-writes, and 68% compared to plain-reads. Conversely, gz-6, for example, requires more energy to compress-write than plain-write, but it saves during the decompress-read of the same file compared to a plain-read. As discussed in Section 5.2, we achieve significant energy savings without compromising performance when the read-to-write ratio exceeds the *break-even* value. Both lzo-1 and lzo-3 save more energy than their counterparts (e.g., gz-6) because lzop implements a fast compression algorithm at the cost of lower compression ratio, thereby delivering the most energy savings.

Figure 3 demonstrates the dependency of energy savings or loss on the read-to-write ratio on a server system. The y-axis denotes the energy savings, E_{sav}, (in Kilojoules). $E_{sav} = 0$ indicates neither energy savings nor energy losses. A positive value of E_{sav} means some energy savings, whereas a negative value denotes energy loss. The plots for gz-1, gz-6, gz-9, and lzo-9 cross the $E_{sav} = 0$ line, denoting that there exists a read-to-write ratio when the corresponding algorithm becomes beneficial in terms of energy. For example, gz-1 crosses the $E_{sav} = 0$ line when n is equal to 20.2. This means

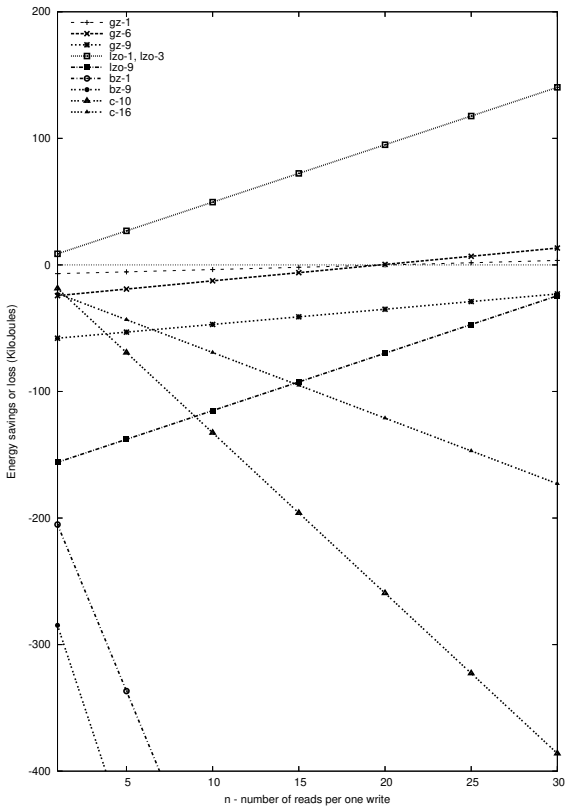

Figure 3: Energy savings or loss (E_{sav}) compared to plain read and write as a function of the number of reads per one write for text files on the server system. The lzo-1 and lzo-3 results are so close, that we put them on the same line.

Tool	Text	Binary	Rand	Zero
gz-1	20.2 / 18.3	×	×	∀
gz-6	19.7 / 62.5	×	×	∀
gz-9	49.0 /~10^3	×	×	∀
lzo-1	∀	2.1 / 4.0	×	∀
lzo-3	∀	2.0 / 4.3	×	∀
lzo-9	35.4 /~10^2	~10^2 /~10^3	×	5.4 / 19.94
bz-1	×	×	×	0.28 / ∀
bz-9	×	×	×	1.28 / 2.36
c-10	×	×	×	∀
c-16	×	×	×	∀

Table 4: The number of reads for each write required to benefit from compression for E (n_{be}^{E}) and ET (n_{be}^{ET}) metrics on a server system, separated by a /. For break-even values greater than 100 we only report magnitudes. ∀ denotes that it is beneficial to use compression for any read-to-write ratio. The symbol × denotes scenarios when no savings can be made. Values less than 1 represent that just one read can compensate for multiple writes.

Tool	Text	Binary	Rand	Zero
gz-1	× / 9.1	×	×	∀
gz-6	33.2 / 38.4	×	×	∀
gz-9	78.6 /~10^2	×	×	∀
lzo-1	∀	3.8 / 6.7	×	∀
lzo-3	∀	3.5 / 5.9	×	∀
lzo-9	31.6 /~10^2	~10^2 /~10^3	×	4.2 / 15.3
bz-1	×	×	×	1.3 / 1.4
bz-9	×	×	×	1.7 / 2.3
c-10	×	×	×	∀
c-16	×	×	×	∀

Table 5: Same results as in Table 4, but for the desktop system.

that if for every write the system experiences 21 or more reads, we can save energy. The lines for lzo-1 and lzo-3 (coincided because of the proximity of results) are always above the $E_{sav} = 0$ line, indicating that these compression tools save energy for any ratio. Conversely, the plots for bz-1, bz-9, c-10, and c-16 never cross the $E_{sav} = 0$ line, implying that these tools expend so much more energy compared to plain-write and plain-read that they are unable to amortize the energy losses for any read-to-write ratio. Both bz-1 and bz-9 consume more time and energy during compress-write as well as decompress-read, because of the algorithmic complexity of bzip, which uses the Burrows-Wheeler transform, the move-to-front transform, and Huffman coding.

Table 4 contains the calculations of n_{be}^{E} and n_{be}^{ET} for the compression tools discussed in this paper. The break-even value varies depending on the compression tools used and their CR values. Often, the higher the CR, the slower the utility operates, consuming more energy, thereby raising the break-even ratio. Although the break-even ratio for some utilities (gz-9, lzo-9, etc.) is greater than the read-to-write ratio on a common server system [35, 47], we can consider them beneficial for a read-intensive workload (e.g., public FTP mirrors). In this case, tools with a higher CR can be applicable, especially if storage space and network traffic are a great concern.

Figures 2(e) and 2(f) show the combined ET metric, which highlights both energy consumption and performance. In these figures, we observe the same trend as in the energy and time results: lzo-1 and lzo-3 performing the best among all, followed by gz-1 and gz-6, considering both read and write workloads.

These results indicate a strong linear relationship between the time to complete the compression/decompression and the energy consumed during these operations. To show this linearity we use the correlation coefficient, R, which measures how linearly energy and time are related [14, 45]. A value of $R = 1$ means that energy and time both lie on the same line (i.e., they are linear in relationship). A value close to 1 implies a stronger linear relationship between the two values. We calculated the value of R to be greater than 0.991 for all four sets of files in our benchmarks, which corroborates the strong linear relationship between time and energy. Linear relationship is the result of the fact that compression tools consume equal amount of energy per unit time to perform compression or decompression. Due to this linearity, we omit the time metric for the remaining types of files in the subsections below.

Apart from the four compression utilities mentioned before, we also used the PPMd compression utility [51] on a 2GB text file and found that it is the largest consumer of energy and time. This can be attributed to the fact that PPMd is based on the PPM algorithm, which is known to produce the best compression ratio, at the expense of considerably greater time and memory resources. We find PPMd to yield a compression ratio of 7.6, but consuming about 10 times more energy than a plain write. Unlike all other compression utilities, which often decompress faster than they compress, PPMd has to perform similar operations during compression as well as decompression. Hence, it is equally slow and energy exhaustive during both compression and decompression of files. It

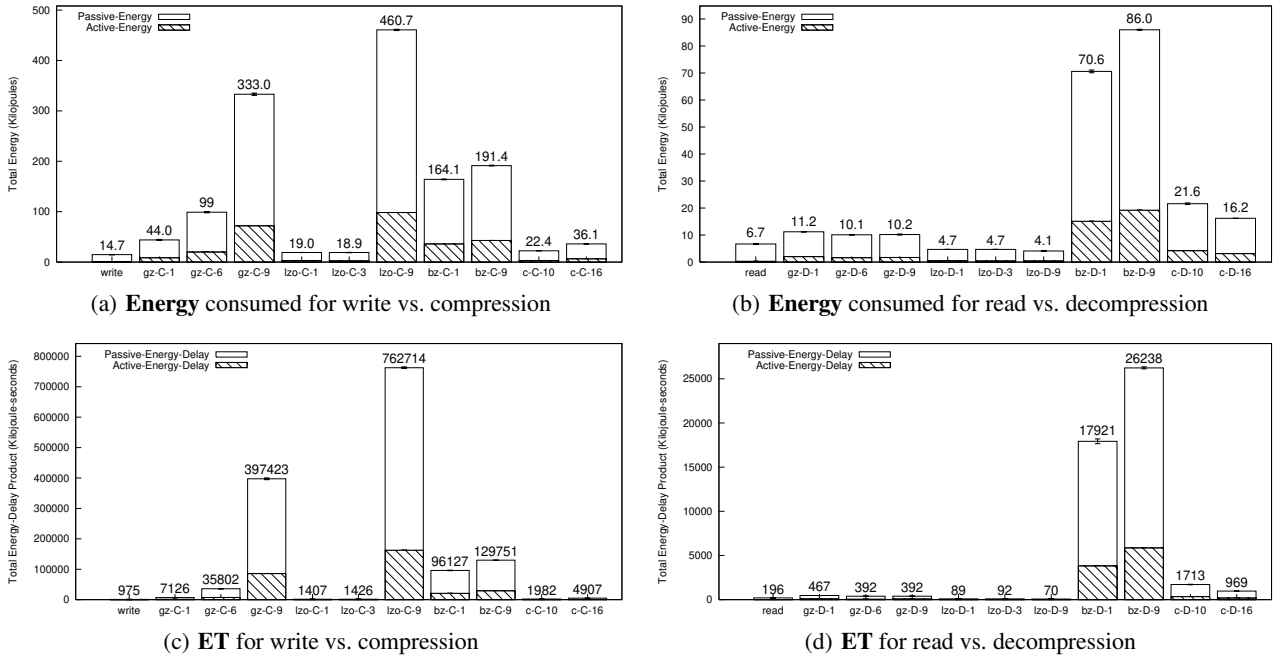

(a) **Energy** consumed for write vs. compression

(b) **Energy** consumed for read vs. decompression

(c) **ET** for write vs. compression

(d) **ET** for read vs. decompression

Figure 4: Binary file: The energy and energy-delay product (ET) for compress-write/decompress-read a 2GB binary file.

consumes approximately 30 times more energy during decompression as compared to a normal read. As PPMd does not save energy during either compression or decompression for all types of files, we do not present the analysis of its results here.

We also observed that the compress tool takes significantly more system time as compared to other compression utilities. We ran `strace` and observed that compress performs multiple read and write system calls in units of 1024 bytes, instead of a more optimal unit such as 4KB (page frame size), thereby increasing system time.

Analysis on desktop machines.

We ran the same set of tests on a slower desktop class machine (described in Section 4.1) and found similar results as that on the server class machine. Even on desktop machines, lzo-1 and lzo-3 proved to be the most energy-efficient. But, the break-even ratio for the desktop class machine, shown in Table 5, differs from that computed for a server class machine. Generally we noticed that it is harder to realize energy savings on the slower desktop system, because its CPU is slower. Based on these observations, we conclude that CPU speeds and hardware configuration considerably impact the break-even ratio. Hence, a compression utility which seems promising in terms of energy savings on one hardware configuration, might turn out to be more expensive on the other, for a given set of workload. In Table 5, we see that although gz-1 never saves energy for text files, it does provide savings, after a break-even value of 9.1, when we factor in performance along with energy. This happens because in spite of the decompression and the plain read taking the same amount of energy, the former finishes significantly faster than the latter.

5.4 Binary File Analysis

In the case of a 2GB binary file, as shown in Figure 4(a), the energy consumption during compress-write using both lzo-1 and lzo-3 is greater than plain-write. Conversely, both of them save energy during decompress-read, seen in Figure 4(b). Hence, similar to the discussion in Sections 5.2 and 5.3, we compute that lzo-1 and lzo-3 save energy only when $n_{be}^{E} \geq 2$. However, if we consider the energy-delay metric, shown in Figures 4(c) and 4(d), the value of the break-even ratio changes to 4 and 4.3, for lzo-1 and lzo-3, respectively. We also see that lzo-9 consumes significantly more energy during compression compared to pure-write, that it is difficult to recoup the over-consumption of power through multiple decompress-read workloads. This is evident from the large value of n_{be}^{ET} ($\sim 10^3$) in Table 4. As all other compression utilities have a greater energy consumption and ET values than plain-writes, they cannot be considered as good candidates for compression with energy savings in mind.

5.5 Random File Analysis

It is evident from Figure 5 that no compression utility saves energy during compression or decompression. Consequently, the ET values for the compression utilities is also greater than that of plain-read and plain-write. The reason is that compression utilities find it difficult to discover repeated patterns in a random file, which inherently has a high entropy [1]. Therefore, the tools waste a lot of CPU time and energy trying to compress, but do not gain much in terms of CR, as shown in Table 3. Hence, modern storage systems should recognize high entropy files, such as multimedia, already compressed files, encrypted files, etc., and write them directly to the disk without compression.

5.6 Zero File Analysis

As expected, all the compression utilities, except lzo-9, bz-1, and bz-9 consume less energy than writes and reads (Figure 6). The energy consumption by lzo-9 compress-write is almost twice that of plain write, but it recovers from the energy losses for $n_{be}^{E} \geq 5.4$ (Table 4). Similarly, the break-even ratios for bz-1 and bz-9 are 0.28 and 1.28, respectively. Considering the ET metric, the break-even ratio for lzo-9 rises to 20. Most of the compression utilities

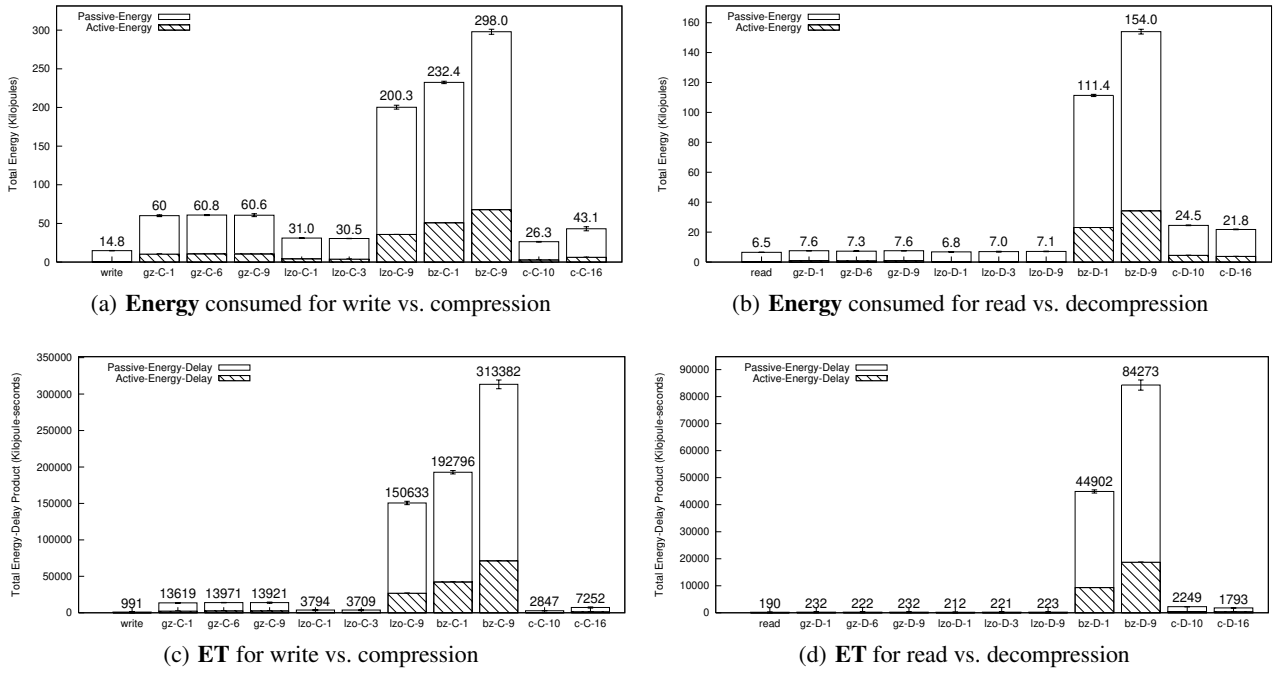

(a) **Energy** consumed for write vs. compression

(b) **Energy** consumed for read vs. decompression

(c) **ET** for write vs. compression

(d) **ET** for read vs. decompression

Figure 5: Random file: The energy and energy-delay product (ET) for compress-write/decompress-read a 2GB random file.

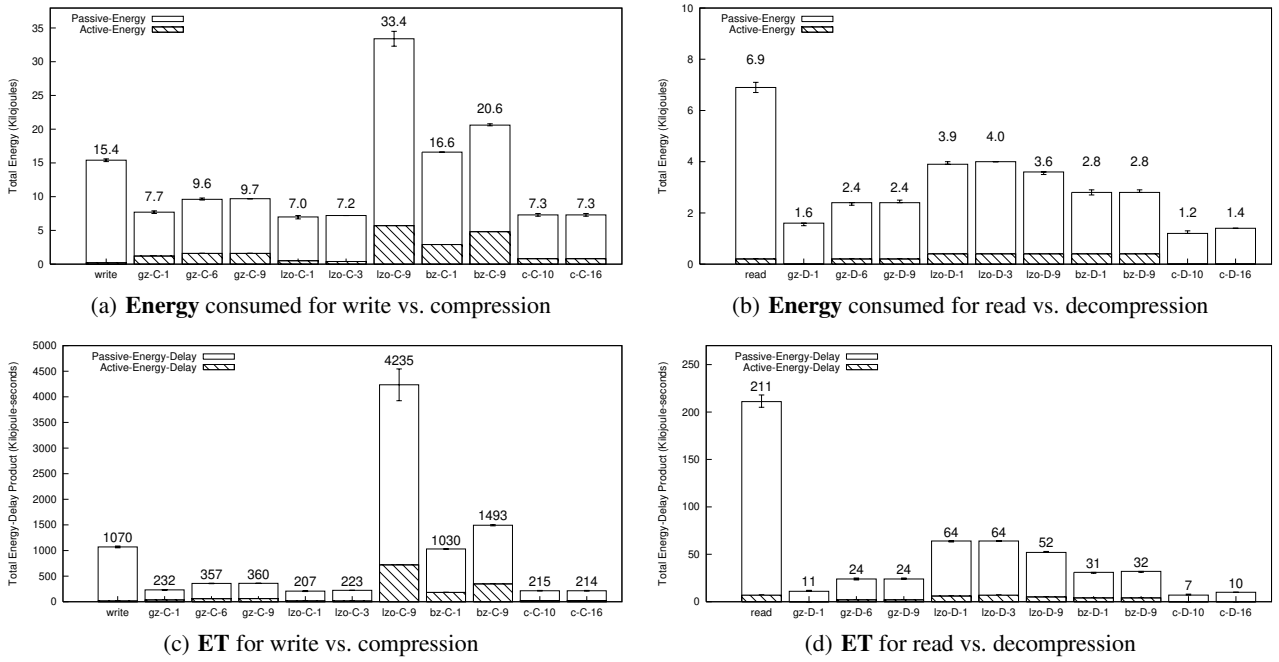

(c) **ET** for write vs. compression

(d) **ET** for read vs. decompression

Figure 6: Zero file: The energy and energy-delay product (ET) for compress-write/decompress-read a 2GB zero file.

achieve a high CR without increasing the ET because of the large frequency of repeated patterns. From this observation, we can suggest that if a storage system consists of files with a large number of repeated patterns (e.g., log files) compression is definitely a better alternative in terms of saving on energy without performance degradation.

5.7 Summary of Evaluation

The strong linear dependency between energy and time found in all experiments indicates that energy consumption per unit of time (power) is independent of the compression algorithm. This leads to the conclusion that the fastest algorithm is the most energy-efficient one. The time required to accomplish the task consists of the time

required to perform an I/O and the time required to compress (or decompress) the data. The time for completing I/O operations, in turn, depends on the amount of data to be written, whereas the compression time depends on the algorithm used. This means that an optimal compression tool should have a high compression ratio and low compression time. There is a clear trade-off between compression ratio and the speed of compression.

The file data type affects the compression ratio dramatically. None of the compression tools we considered provided advantages in energy consumption for the files with random content. For zero files, however, almost all tools provided benefits for both reading and writing. For text and binary files, we observe situations when compress-write is less energy-efficient than plain-write, but decompress-read is more energy-efficient than plain-read. We calculated the break-even ratio of reads to writes (n_{be}) in such cases. For most of the compression tools, this value is significantly higher than the read-to-write ratio on common server systems, which has been found to be typically 2–4 [47, 35]. Some notable exceptions were lzo-1 and lzo-3, which are always beneficial for text files. They also save energy in case of binary files, if the read-to-write ratio is at least 2 (or about 4 in terms of ET metric). We recommend the use of lzo-1 and lzo-3 in all cases, except the situations where disk space is a greater concern than energy or performance. We also recommend that future systems recognize high-entropy files (e.g., encrypted, random, etc.) and avoid compressing them at all. Lastly, while most of the results reported in this section were pertaining to server-class systems, we found similar trends for desktop class systems too; however, we noted that on slower systems with slower CPUs, it is harder to save energy with compression.

6. CONCLUSIONS

In this paper, our research contribution was to investigate the validity of the assumption that data compression is an effective technique to reduce energy consumption in server systems. We evaluated several compression tools on Linux on a variety of data files and compared them against raw reads and writes based on performance and energy metrics. Our experimental results suggest that no generalized conclusion regarding the efficacy of compression can be drawn. It greatly depends on the data redundancy of the file, the compression algorithm being used, the read-to-write ratio of the workload, and the hardware configuration of the system. We found that compressing zero files is beneficial for almost all the compression tools. Random files are better-off not being compressed at all. Text files, when compressed with lzop using options 1 and 3, will always save energy, irrespective of the workload's read-to-write ratio. We developed a simple read-write model to evaluate energy savings in cases where only compression or decompression saves energy. When applied to text and binary files, it reveals that only gzip and lzop can offer energy savings; in most cases the break-even read-to-write ratio is significantly greater (more than 20) than that found in common workloads. Other than on zero files, bzip2 and the compress utility never save any energy.

Future Work.

We intend to extend this study to a wider range of systems, including systems with multiple cores and multiple CPUs, CPUs with Dynamic Voltage and Frequency Scaling (DVFS), different disk speeds, etc. We also plan on conducting our study on real server workloads. Compression significantly reduces the storage requirement for data, and hence can result in lesser spinning disks. We plan to extend our current model to factor in the additional power savings thus achieved. We are currently also working on extending gzipfs [66], a stackable compression file system, to include the various compression algorithms we wish to compare. In the future, we plan to explore and evaluate data de-duplication as an energy saving technique.

Another interesting direction would be to include archivers in the evaluation. Archivers generally work by combining multiple files into one. In scenarios where we have several small files, with similar content or format, compressing their archive would typically result in better compression. Decompression on an archive to read one file will, however, be more expensive than if the file was individually compressed.

7. ACKNOWLEDGMENTS

This work was made possible in part thanks to a 2008 IBM Faculty Award and a 2009 NetApp award. The work is also sponsored in part by the Stony Brook Advanced Energy Research and Technology Center (AERTC, `www.aertc.org`).

8. REFERENCES

[1] K. Y. Alina Oprea, Michael K. Reiter. Space-efficient block storage integrity. In *Proceedings of the NDSS Symposium*, 2005.

[2] P. A. Alsberg. Space and time savings through large data base compression and dynamic restructuring. *Proceedings of the IEEE*, 63(8):1114–1122, 1975.

[3] K. C. Barr and K. Asanovi. Energy-aware lossless data compression. *ACM Transactions on Computer Systems*, 24(3):250–291, 2006.

[4] L. Benini, A. Bogliolo, S. Cavallucci, and B. Ricco. Monitoring system activity for OS-directed dynamic power management. In *Proceedings of the 1998 international symposium on Low power electronics and design (ISLPED '98)*, pages 185–190, New York, NY, USA, 1998. ACM.

[5] L. Benini, D. Bruni, A. Macii, and E. Macii. Hardware-assisted data compression for energy minimization in systems with embedded processors. *Design, Automation and Test in Europe Conference and Exhibition (DATE '02)*, pages 449–453, 2002.

[6] L. Benini, D. Bruni, B. Ricco, A. Macii, and E. Macii. An adaptive data compression scheme for memory traffic minimization in processor-based systems. *IEEE International Symposium on Circuits and Systems (ISCAS '02)*, 4:IV–866–IV–869, 2002.

[7] L. Benini, A. Macii, E. Macii, and M. Poncino. Selective instruction compression for memory energy reduction in embedded systems. In *Proceedings of the 1999 international symposium on Low power electronics and design (ISLPED '99)*, pages 206–211, New York, NY, USA, 1999. ACM.

[8] L. Benini, A. Macii, and A. Nannarelli. Cached-code compression for energy minimization in embedded processors. In *Proceedings of the 2001 international symposium on Low power electronics and design (ISLPED '01)*, pages 322–327, New York, NY, USA, 2001. ACM.

[9] L. Benini, F. Menichelli, and M. Olivieri. A class of code compression schemes for reducing power consumption in embedded microprocessor systems. *IEEE Transactions on Computers*, 53(4):467–482, 2004.

[10] P. Bohrer, E. N. Elnozahy, T. Keller, M. Kistler, and L. C. Mcdowell. The case for Power Management in Web Servers, 2002. `www.research.ibm.com/people/l/lefurgy/Publications/pac2002.pdf`.

[11] D. Colarelli and D. Grunwald. Massive arrays of idle disks

for storage archives. In *Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, pages 1–11, 2002.

[12] CompuGreen, LLC. The Green500 List. `www.green500.org`, 2008.

[13] D. J. Craft. A fast hardware data compression algorithm and some algorithmic extensions. *IBM Journal of Research and Development*, 42(6):733–745, 1998.

[14] D. C. Montgomery. *Engineering Statistics*. Wiley, 3 edition, 2004.

[15] G. Debnath, K. Debnath, and R. Fernando. The Pentium processor-90/100, microarchitecture and low power circuit design. In *Proceedings of the 8th International Conference on VLSI Design (VLSID '95)*, page 185, Washington, DC, USA, 1995. IEEE Computer Society.

[16] F. Douglis, P. Krishnan, and B. Bershad. Adaptive Disk Spin-down Policies for Mobile Computers. In *Proceedings of the 2nd Symposium on Mobile and Location-Independent Computing*, pages 121–137, Berkeley, CA, USA, 1995. USENIX Association.

[17] F. Douglis, P. Krishnan, and B. Marsh. Thwarting the Power-Hungry Disk. In *Proceedings of the 1994 Winter USENIX Conference*, pages 293–306, 1994.

[18] D. Essary and A. Amer. Predictive data grouping: Defining the bounds of energy and latency reduction through predictive data grouping and replication. *ACM Transactions on Storage (TOS)*, 4(1):1–23, May 2008.

[19] J. Flinn and M. Satyanarayanan. Energy-Aware adaptation for Mobile Applications. In *Proceedings of the seventeenth ACM symposium on Operating systems principles*, volume 33, pages 48–63, 1999.

[20] Fluke 289 Digital Multimeter. `http://assets.fluke.com/manuals/287_289_umeng0100.pdf`.

[21] Fluke i410 AC/DC Current Clamp. `http://assets.fluke.com/manuals/i4101010iseng0200.pdf`.

[22] S. Gary, P. Ippolito, G. Gerosa, C. Dietz, J. Eno, and H. Sanchez. PowerPC 603, A Microprocessor for Portable Computers. *IEEE Design and Test*, 11(4):14–23, 1994.

[23] R. Gonzalez and M. Horowitz. Energy Dissipation in General Purpose Microprocessors. *IEEE Journal of Solid-state Circuits*, 31(9):1277–1284, September 1996.

[24] B. Gordon and T. Meng. A low power subband video decoder architecture. In *Proceeding of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume ii, pages 409–412, 1994.

[25] M. K. Gowan, L. L. Biro, and D. B. Jackson. Power considerations in the design of the Alpha 21264 microprocessor. In *Proceedings of the 35th annual conference on Design automation (DAC '98)*, pages 726–731, New York, NY, USA, 1998. ACM.

[26] H. Huang, W. Hung, and K. Shin. FS2: Dynamic data replication in free disk space for improving disk performance and energy consumption. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP '05)*, pages 263–276, Brighton, UK, October 2005. ACM Press.

[27] E. Jeannot, B. Knutsson, and M. Bjorkman. Adaptive online data compression. In *Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing (HPDC '02)*, pages 379–388, Edinburgh, Scotland, July 2002. IEEE Computer Society.

[28] N. Joukov and J. Sipek. GreenFS: Making enterprise computers greener by protecting them better. In *Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems 2008 (EuroSys 2008)*, Glasgow, Scotland, April 2008. ACM. (**Won best paper award**).

[29] M. Kandemir, O. Ozturk, M. Irwin, and I. Kolcu. Using data compression to increase energy savings in multi-bank memories. In *Proceedings of the 10th International Euro-Par Conference on Parallel Processing (Euro-Par '04)*, volume 3149 of *Lecture Notes in Computer Science*, pages 310–317. Springer-Verlag Berlin Heidelberg, 2004.

[30] G. Keramidas, K. Aisopos, and S. Kaxiras. Dynamic Dictionary-Based Data Compression for Level-1 Caches. In *Proceedings of the 19th International Conference on Architecture of Computing Systems (ARCS '06)*, pages 114–129, 2006.

[31] N. Kim, T. Austin, and T. Mudge. Low-Energy Data Cache using Sign Compression and Cache Line Bisection. In *Proceedings of the 2nd Annual Workshop on Memory Performance Issues (WMPI '02)*, 2002.

[32] C. Krintz and B. Cadler. Reducing delay with dynamic selection of compression formats. In *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC '01)*, pages 266–277, 2001.

[33] A. R. Lebeck, X. Fan, H. Zeng, and C. Ellis. Power Aware Page Allocation. In *Proceedings of Architectural Support for Programming Languages and Operating Systems*, pages 105–116, 2000.

[34] H. Lekatsas, J. Henkel, and W. Wolf. Code compression for low power embedded system design. In *Proceedings of the 37th conference on Design automation (DAC '00)*, pages 294–299, New York, NY, USA, 2000. ACM.

[35] A. W. Leung, S. Pasupathy, G. Goodson, and E. L. Miller. Measurement and analysis of large-scale network file system workloads. In *Proceedings of the USENIX Annual Technical Conference (ATC '08)*, pages 213–226, Berkeley, CA, 2008. USENIX Association.

[36] K. Li, R. Kumpf, P. Horton, and T. Anderson. A Quantitative Analysis of Disk Drive Power Management in Portable Computers. In *Proceedings of the 1994 Winter USENIX Conference*, pages 279–291, 1994.

[37] W.-C. Lin and C.-H. Chen. An energy-delay efficient power management scheme for embedded system in multimedia applications. In *Proceedings of The IEEE Asia Pacific Conference on Circuit and System (APCCAS)*, pages 869–872, 2004.

[38] J. R. Lorch and A. J. Smith. Software strategies for portable computer energy management. *IEEE Personal Communications*, 5:48–63, 1998.

[39] Y. Lu, L. Benini, and G. D. Micheli. Operating-System Directed Power Reduction. In *Proceedings of the 2000 international symposium on Low power electronics and design*, pages 37–42, 2000.

[40] R. Manohar and N. Nystrom. Implications of voltage scaling in asynchronous architectures. Technical Report CSL-TR-2001-1013, Departament of Computer Science, Cornell University, 2001.

[41] D. Narayanan, A. Donnelly, and A. Rowstron. Write off-loading: practical power management for enterprise storage. In *Proceedings of the 6th USENIX Conference on File and Storage Technologies (FAST 2008)*, 2008.

[42] J. L. Nunez and S. Jones. Gbit/s lossless data compression hardware. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 11(3):499–510, 2003.

[43] M. Oberhumer. lzop data compression utility. www.lzop.org/.

[44] E. Pinheiro and R. Bianchini. Energy Conservation Techniques for Disk Array-Based Servers. In *Proceedings of the 18th International Conference on Supercomputing (ICS 2004)*, pages 68–78, 2004.

[45] R. Jain. *The Art of Computer System Performance Analysis*. Wiley, 1991.

[46] T. Raita. An automatic system for file compression. *The Computer Journal*, pages 80–86, 1987.

[47] D. Roselli, J. R. Lorch, and T. E. Anderson. A comparison of file system workloads. In *Proc. of the Annual USENIX Technical Conference*, pages 41–54, San Diego, CA, June 2000. USENIX Association.

[48] S. B. Furber. *ARM System Architecture*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1996.

[49] S. Gurumurthi and A. Sivasubramaniam and M. Kandemir and H. Franke. DRPM: Dynamic Speed Control for Power Management in Server Class Disks. In *Proceedings of the 30th annual international symposium on Computer architecture*, pages 169–181, 2003.

[50] C. M. Sadler and M. Martonosi. Data compression algorithms for energy-constrained devices in delay tolerant networks. In *Proceedings of the 4th international conference on Embedded networked sensor systems (SenSys '06)*, pages 265–278, New York, NY, USA, 2006. ACM.

[51] D. Shkarin. PPMd data compression utility. www.compression.ru/ds/.

[52] C. Su and A. M. Despain. Cache design trade-offs for power and performance optimization: a case study. In *Proceedings of the 1995 international symposium on Low power design*, pages 63–68, 1995.

[53] C. Su and A. M. Despain. Cache designs for energy efficiency. In *Proceedings of the 28th Hawaii International Conference on System Sciences (HICSS'95)*, page 306, Washington, DC, USA, 1995. IEEE Computer Society.

[54] K. Tanaka and T. Kawahara. Leakage energy reduction in cache memory by data compression. *ACM SIGARCH Computer Architecture News*, 35(5):17–24, December 2007.

[55] R. B. Tremaine, P. A. Franaszek, J. T. Robinson, C. O. Schulz, T. B. Smith, M. E. Wazlowski, and P. M. Bland. IBM Memory Expansion Technology (MXT). *IBM Journal of Research and Development*, 45(2):271–286, 2001.

[56] L. Villa, M. Zhang, and K. Asanovi. Dynamic zero compression for cache energy reduction. In *Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture (MICRO 33)*, pages 214–220, New York, NY, USA, 2000. ACM.

[57] Watts up? PRO ES Power Meter. www.wattsupmeters.com/secure/products.php.

[58] M. Weiser, B. Welch, A. Demers, and S. Shenker. Scheduling for reduced CPU energy. In *Proceedings of the 1st USENIX conference on Operating Systems Design and Implementation*, 1994.

[59] M. Weiser, B. Welsh, A. Demers, and S. Shenker. Scheduling for reduced CPU energy. *Mobile Computing*, 353:449–471, 1996.

[60] D. Wheeler. Linux utility for wattsup pro es power meter. www.wattsupmeters.com/forum/index.php?topic=7.0.

[61] J. Wilkes. Predictive power conservation. Technical Report HPL-CSP-92-5, Hewlett-Packard Laboratories, February 1992.

[62] Y. Wiseman, K. Schwan, and P. Widener. Efficient end to end data exchange using configurable compression. *ACM SIGOPS Operating Systems Review*, 39(3):4–23, 2005.

[63] C. P. Wright, N. Joukov, D. Kulkarni, Y. Miretskiy, and E. Zadok. Auto-pilot: A platform for system software benchmarking. In *Proceedings of the Annual USENIX Technical Conference, FREENIX Track*, pages 175–187, Anaheim, CA, April 2005. USENIX Association.

[64] R. Xu, Z. Li, C. Wang, and P. Ni. Impact of Data Compression on Energy Consumption of Wireless-Networked Handheld Devices. In *Proceedings of the 23rd International Conference on Distributed Computing Systems (ICDCS '03)*, page 302, 2003.

[65] Y. Yoshida, B. Song, H. Okuhata, T. Onoye, and I. Shirakawa. An object code compression approach to embedded processors. In *Proceedings of the 1997 international symposium on Low power electronics and design (ISLPED '97)*, pages 265–268, New York, NY, USA, 1997. ACM.

[66] E. Zadok, J. M. Anderson, I. Bădulescu, and J. Nieh. Fast indexing: Support for size-changing algorithms in stackable file systems. In *Proceedings of the Annual USENIX Technical Conference*, pages 289–304, Boston, MA, June 2001. USENIX Association.

[67] H. Zeng, C. S. Ellis, A. R. Lebeck, and A. Vahdat. ECOSystem: Managing Energy as a First Class Operating System Resource. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS X)*, pages 123–132, 2002.

[68] Q. Zhu, F. M. David, C. F. Devaraj, Z. Li, Y. Zhou, and P. Cao. Reducing Energy Consumption of Disk Storage Using Power-Aware Cache Management. In *Proceedings of the 10th International Symposium on High-Performance Computer Architecture*, pages 118–129, 2004.

[69] V. Zyuban and P. Kogge. Optimization of high-performance superscalar architectures for energy efficiency. In *Proceedings of the 2000 International Symposium on Low Power Electronics and Design (ISLPED '00)*, pages 84–89, 2000.