

FAST: Quick Application Launch on Solid-State Drives

Yongsoo Joo[†], Junhee Ryu[‡], Sangsoo Park[†], and Kang G. Shin^{†*}

[†]Ewha Womans University, 11-1 Daehyun-dong Seodaemun-gu, Seoul 120-750, Korea

[‡]Seoul National University, 599 Kwanak-Gu Kwanak Rd., Seoul 151-744, Korea

* University of Michigan, 2260 Hayward St., Ann Arbor, MI 48109, USA

Abstract

Application launch performance is of great importance to system platform developers and vendors as it greatly affects the degree of users' satisfaction. The single most effective way to improve application launch performance is to replace a hard disk drive (HDD) with a solid state drive (SSD), which has recently become affordable and popular. A natural question is then whether or not to replace the traditional HDD-aware application launchers with a new SSD-aware optimizer.

We address this question by analyzing the inefficiency of the HDD-aware application launchers on SSDs and then proposing a new SSD-aware application prefetching scheme, called the *Fast Application Starter* (FAST). The key idea of FAST is to overlap the computation (CPU) time with the SSD access (I/O) time during an application launch. FAST is composed of a set of user-level components and system debugging tools provided by the Linux OS (operating system). In addition, FAST uses a system-call wrapper to automatically detect application launches. Hence, FAST can be easily deployed in any recent Linux versions without kernel recompilation. We implemented FAST on a desktop PC with a SSD running Linux 2.6.32 OS and evaluated it by launching a set of widely-used applications, demonstrating an average of 28% reduction of application launch time as compared to PC without a prefetcher.

1 Introduction

Application launch performance is one of the important metrics for the design or selection of a desktop or a laptop PC as it critically affects the user-perceived performance. Unfortunately, application launch performance has not kept up with the remarkable progress of CPU performance that has thus far evolved according to Moore's law. As frequently-used or popular applications get "heavier" (by adding new functions) with each new

release, their launch takes longer even if a new, powerful machine equipped with high-speed multi-core CPUs and several GBs of main memory is used. This undesirable trend is known to stem from the poor random access performance of hard disk drives (HDDs). When an application stored in a HDD is launched, up to thousands of block requests are sent to the HDD, and a significant portion of its launch time is spent on moving the disk head to proper track and sector positions, i.e., *seek* and *rotational* latencies. Unfortunately, the HDD seek and rotational latencies have not been improved much over the last few decades, especially compared to the CPU speed improvement. In spite of the various optimizations proposed to improve the HDD performance in launching applications, users must often wait tens of seconds for the completion of launching frequently-used applications, such as Windows Outlook.

A quick and easy solution to eliminate the HDD's seek and rotational latencies during an application launch is to replace the HDD with a solid state drive (SSD). A SSD consists of a number of NAND flash memory modules, and does not use any mechanical parts, unlike disk heads and arms of a conventional HDD. While the HDD access latency—which is the sum of seek and rotational latencies—ranges up to a few tens of milliseconds (ms), depending on the seek distance, the SSD shows a rather uniform access latency of about a few hundred microseconds (*us*). Replacing a HDD with a SSD is, therefore, the single most effective way to improve application launch performance.

Until recently, using SSDs as the secondary storage of desktops or laptops has not been an option for most users due to the high cost-per-bit of NAND flash memories. However, the rapid advance of semiconductor technology has continuously driven the SSD price down, and at the end of 2009, the price of an 80 GB SSD has fallen below 300 US dollars. Furthermore, SSDs can be installed in existing systems without additional hardware or software support because they are usually equipped with the

same interface as HDDs, and OSES see a SSD as a block device just like a HDD. Thus, end-users begin to use a SSD as their system disk to install the OS image and applications.

Although a SSD can significantly reduce the application launch time, it does not give users ultimate satisfaction for all applications. For example, using a SSD reduces the launch time of a heavy application from tens of seconds to several seconds. However, users will soon become used to the SSD launch performance, and will then want the launch time to be reduced further, just as they see from light applications. Furthermore, users will keep on adding functions to applications, making them heavier with each release and their launch time greater. According to a recent report [24], the growth of software is rapid and limited only by the ability of hardware. These call for the need to further improve application launch performance on SSDs.

Unfortunately, most previous optimizers for application launch performance are intended for HDDs and have not accounted for the SSD characteristics. Furthermore, some of them may rather be detrimental to SSDs. For example, running a disk defragmentation tool on a SSD is not beneficial at all because changing the physical location of data in the SSD does not affect its access latency. Rather, it generates unnecessary write and erase operations, thus shortening the SSD’s lifetime.

In view of these, the first step toward SSD-aware optimization may be to simply disable the traditional optimizers designed for HDDs. For example, Windows 7 disables many functions, such as disk defragmentation, application prefetch, Superfetch, and Readyboost when it detects a SSD being used as a system disk [27]. Let’s consider another example. Linux is equipped with four disk I/O schedulers: NOOP, anticipatory, deadline, and completely fair queueing. The NOOP scheduler almost does nothing to improve HDD access performance, thus providing the worst performance on a HDD. Surprisingly, it has been reported that NOOP shows better performance than the other three sophisticated schedulers on a SSD [11].

To the best of our knowledge, this is the first attempt to focus entirely on improving application launch performance on SSDs. Specifically, we propose a new application prefetching method, called the *Fast Application Starter* (FAST), to improve application launch time on SSDs. The key idea of FAST is to overlap the computation (CPU) time with the SSD access (I/O) time during each application launch. To achieve this, we monitor the sequence of block requests in each application, and launch the application simultaneously with a prefetcher that generates I/O requests according to the *a priori* monitored application’s I/O request sequence. FAST consists of a set of user-level components, a system-call wrap-

per, and system debugging tools provided by the Linux OS. FAST can be easily deployed in most recent Linux versions without kernel recompilation. We have implemented and evaluated FAST on a desktop PC with a SSD running Linux 2.6.32, demonstrating an average of 28% reduction of application launch time as compared to PC without a prefetcher.

This paper makes the following contributions:

- Qualitative and quantitative evaluation of the inefficiency of traditional HDD-aware application launch optimizers on SSDs;
- Development of a new SSD-aware application prefetching scheme, called FAST; and
- Implementation and evaluation of FAST, demonstrating its superiority and deployability.

While FAST can be also applied to HDDs, its performance improvements are only limited to high I/O requirements of application launches on HDDs. We observed that existing application prefetchers outperformed FAST on HDDs by effectively optimizing disk head movements, which will be discussed further in Section 5.

The paper is organized as follows. In Section 2, we review other related efforts and discuss their performance in optimizing application launch on SSDs. Section 3 describes the key idea of FAST and presents an upper bound for its performance. Section 4 details the implementation of FAST on the Linux OS, while Section 5 evaluates its performance using various real-world applications. Section 6 discusses the applicability of FAST to smartphones and Section 7 compares FAST with traditional I/O prefetching techniques. We conclude the paper with Section 8.

2 Background

2.1 Application Launch Optimization

Application-level optimization. Application developers are usually advised to optimize their applications for fast startup. For example, they may be advised to postpone loading non-critical functions or libraries so as to make applications respond as fast as possible [2, 30]. They are also advised to reduce the number of symbol relocations while loading libraries, and to use dynamic library loading. There have been numerous case studies—based on in-depth analyses and manual optimizations—of various target applications/platforms, such as Linux desktop suite platform [8], a digital TV [17], and a digital still camera [33]. However, such an approach requires the experts’ manual optimizations for each and every application. Hence, it is economically infeasible for general-purpose systems with many (dynamic) application programs.

Snapshot technique. A snapshot boot technique has also been suggested for fast startup of embedded systems [19], which is different from the traditional hibernate shutdown function in that a snapshot of the main memory after booting an OS is captured only once, and used repeatedly for every subsequent booting of the system. However, applying this approach for application launch is not practical for the following reasons. First, the page cache in main memory is shared by all applications, and separating only the portion of the cache content that is related to a certain application is not possible without extensive modification of the page cache. Furthermore, once an application is updated, its snapshot should be invalidated immediately, which incurs runtime overhead.

Prediction-based prefetch. Modern desktops are equipped with large (up to several GBs) main memory, and often have abundant free space available in the main memory. Prediction-based prefetching, such as Superfetch [28] and Preload [12], loads an application’s code blocks in the free space even if the user does not explicitly express his intent to execute that particular application. These techniques monitor and analyze the users’ access patterns to predict which applications to be launched in future. Consequently, the improvement of launch performance depends strongly on prediction accuracy.

Sorted prefetch. The Windows OS is equipped with an application prefetcher [36] that prefetches application code blocks in a sorted order of their logical block addresses (LBAs) to minimize disk head movements. A similar idea has also been implemented for Linux OS [15, 25]. We call these approaches *sorted prefetch*. It monitors HDD activities to maintain a list of blocks accessed during the launch of each application. Upon detection of an application launch, the application prefetcher immediately pauses its execution and begins to fetch the blocks in the list in an order sorted by their LBAs. The application launch is resumed after fetching all the blocks, and hence, no page miss occurs during the launch.

Application defragmentation. The block list information can also be used in a different way to further reduce the seek distance during an application launch. Modern OSes commonly support a HDD defragmentation tool that reorganizes the HDD layout so as to place each file in a contiguous disk space. In contrast, the defragmentation tool can relocate the blocks in the list of each application by their access order [36], which helps reduce the total HDD seek distance during the launch.

Data pinning on flash caches. Recently, flash cache has been introduced to exploit the advantage of SSDs at a cost comparable to HDDs. A flash cache can be integrated into traditional HDDs, which is called a *hybrid HDD* [37]. Also, a PCI card-type flash cache is available

[26], which is connected to the mother board of a desktop or laptop PC. As neither seek nor rotational latency is incurred while accessing data in the flash cache, we can accelerate application launch by storing the code blocks of frequently-used applications, which is called a *pinned set*. Due to the small capacity of flash cache, how to determine the optimal pinned set subject to the capacity constraint is a key to making performance improvement, and a few results of addressing this problem have been reported [16, 18, 22]. We expect that FAST can be integrated with the flash cache for further improvement of performance, but leave it as part of our future work.

2.2 SSD Performance Optimization

SSDs have become affordable and begun to be deployed in desktop and laptop PCs, but their performance characteristics have not yet been understood well. So, researchers conducted in-depth analyses of their performance characteristics, and suggested ways to improve their runtime performance. Extensive experiments have been carried out to understand the performance dynamics of commercially-available SSDs under various workloads, without knowledge of their internal implementations [7]. Also, SSD design space has been explored and some guidelines to improve the SSD performance have been suggested [10]. A new write buffer management scheme has also been suggested to improve the random write performance of SSDs [20]. Traditional I/O schedulers optimized for HDDs have been revisited in order to evaluate their performance on SSDs, and then a new I/O scheduler optimized for SSDs has been proposed [11, 21].

2.3 Launch Optimization on SSDs

As discussed in Section 2.1, various approaches have been developed and deployed to improve the application launch performance on HDDs. On one hand, many of them are effective on SSDs as well, and orthogonal to FAST. For example, application-level optimization and prediction-based prefetch can be used together with FAST to further improve application launch performance.

On the other hand, some of them exploit the HDD characteristics to reduce the seek and rotational delay during an application launch, such as the sorted prefetch and the application defragmentation. Such methods are ineffective for SSDs because the internal structure of a SSD is very different from that of a HDD. A SSD typically consists of multiple NAND flash memory modules, and does not have any mechanical moving part. Hence, unlike a HDD, the access latency of a SSD is irrelevant to the LBA distance between the last and the current block

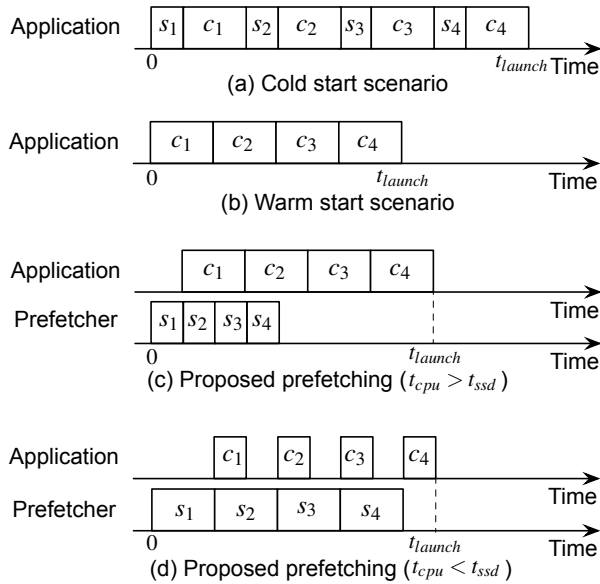


Figure 1: Various application launch scenarios ($n = 4$).

requests. Thus, prefetching the application code blocks according to the sorted order of their LBAs or changing their physical locations will not make any significant performance improvement on SSDs. As the sorted prefetch has the most similar structure to FAST, we will quantitatively compare its performance with FAST in Section 5.

3 Application Prefetching on SSDs

This section illustrates the main idea of FAST with examples and derives a lower bound of the application launch time achievable with FAST.

3.1 Cold and Warm Starts

We focus on the performance improvement in case of a cold start, or the first launch of an application upon system bootup, representing the worst-case application launch performance. Figure 1(a) shows an example cold start scenario, where s_i is the i -th block request generated during the launch and n the total number of block requests. After s_i is completed, the CPU proceeds with the launch process until another page miss takes place. Let c_i denote this computation.

The opposite extreme is a warm start in which all the code blocks necessary for launch have been found in the page cache, and thus, no block request is generated, as shown in Figure 1(b). This occurs when the application is launched again shortly after its closure. The warm start represents an upper-bound of the application launch performance improvement achievable with optimization of

the secondary storage.

Let the time spent for s_i and c_i be denoted by $t(s_i)$ and $t(c_i)$, respectively. Then, the computation (CPU) time, t_{cpu} , is expressed as

$$t_{cpu} = \sum_{i=1}^n t(c_i), \quad (1)$$

and the SSD access (I/O) time, t_{ssd} , is expressed as

$$t_{ssd} = \sum_{i=1}^n t(s_i). \quad (2)$$

3.2 The Proposed Application Prefetcher

The rationale behind FAST is that the I/O request sequence generated during an application launch does not change over repeated launches of the application in case of cold-start. The key idea of FAST is to overlap the SSD access (I/O) time with the computation (CPU) time by running the application prefetcher concurrently with the application itself. The application prefetcher replays the I/O request sequence of the original application, which we call an *application launch sequence*. An application launch sequence S can be expressed as (s_1, \dots, s_n) .

Figure 1(c) illustrates how FAST works, where $t_{cpu} > t_{ssd}$ is assumed. At the beginning, the target application and the prefetcher start simultaneously, and compete with each other to send their first block request to the SSD. However, the SSD always receives the same block request s_1 regardless of which process gets the bus grant first. After s_1 is fetched, the application can proceed with its launch by the time $t(c_1)$, while the prefetcher keeps issuing the subsequent block requests to the SSD. After completing c_1 , the application accesses the code block corresponding to s_2 , but no page miss occurs for s_2 because it has already been fetched by the prefetcher. It is the same for the remaining block requests, and thus, the resulting application launch time t_{launch} becomes

$$t_{launch} = t(s_1) + t_{cpu}. \quad (3)$$

Figure 1(d) shows another possible scenario where $t_{cpu} < t_{ssd}$. In this case, the prefetcher cannot complete fetching s_2 before the application finishes computation c_1 . However, s_2 can be fetched by $t(c_1)$ earlier than that of the cold start, and this improvement is accumulated for all of the remaining block requests, resulting in t_{launch} :

$$t_{launch} = t_{ssd} + t(c_n). \quad (4)$$

Note that n ranges up to a few thousands for typical applications, and thus, $t(s_1) \ll t_{cpu}$ and $t(c_n) \ll t_{ssd}$. Consequently, Eqs. (3) and (4) can be combined into a single equation as:

$$t_{launch} \approx \max(t_{ssd}, t_{cpu}), \quad (5)$$




Figure 2: A worst-case example ($t_{cpu} = t_{ssd}$).

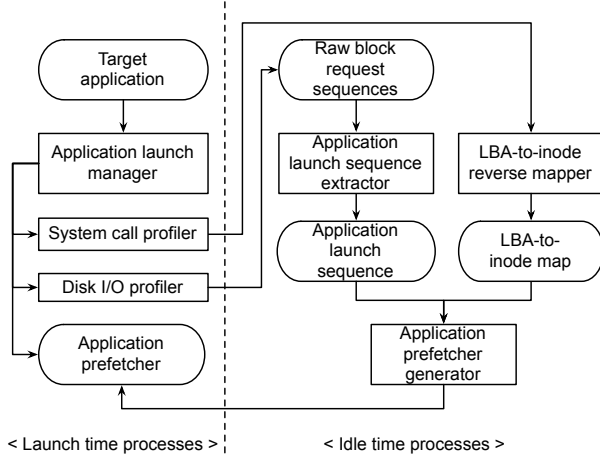


Figure 3: The proposed application prefetching.

which represents a lower bound of the application launch time achievable with FAST.

However, FAST may not achieve application launch performance close to Eq. (5) when there is a significant variation of I/O intensiveness, especially if the beginning of the launch process is more I/O intensive than the other. Figure 2 illustrates an extreme example of such a case, where the first half of this example is SSD-bound and the second half is CPU-bound. In this example, t_{cpu} is equal to t_{ssd} , and thus the expected launch time $t_{expected}$ is given to be $t_{ssd} + t(c_8)$, according to Eq. (4). However, the actual launch time t_{actual} is much larger than $t_{expected}$. The CPU usage in the first half of the launch time is kept quite low despite the fact that there are lots of remaining CPU computations (i.e., c_5, \dots, c_8) due to the dependency between s_i and c_i . We will provide a detailed analysis for this case using real applications in Section 5.

4 Implementation

We chose the Linux OS to demonstrate the feasibility and the superior performance of FAST. The implementation of FAST consists of a set of components: an application launch manager, a system-call profiler, a disk I/O profiler, an application launch sequence extractor, a LBA-to-inode reverse mapper, and an application prefetcher generator. Figure 3 shows how these components interact with each other. In what follows, we detail the implementation of each of these components.

4.1 Application Launch Sequence

4.1.1 Disk I/O Profiler

The disk I/O profiler is used to track the block requests generated during an application launch. We used Blktrace [3], a built-in Linux kernel I/O-tracing tool that monitors the details of I/O behavior for the evaluation of I/O performance. Blktrace can profile various I/O events: inserting an item into the block layer, merging the item with a previous request in the queue, remapping onto another device, issuing a request to the device driver, and a completion signal from the device. From these events, we collect the trace of device-completion events, each of which consists of a device number, a LBA, the I/O size, and completion time.

4.1.2 Application Launch Sequence Extractor

Ideally, the application launch sequence should include all of the block requests that are generated every time the application is launched in the cold start scenario, without including any block requests that are not relevant to the application launch. We observed that the raw block request sequence captured by Blktrace does not vary from one launch to another, i.e., deterministic for multiple launches of the same application. However, we observed that other processes (e.g., OS and application daemons) sometimes generate their own I/O requests simultaneously with the application launch. To handle this case, the application launch sequence extractor collects two or more raw block request sequences to extract a common sequence, which is then used as a launch sequence of the corresponding application. The implementation of the application launch sequence extractor is simple: it searches for and removes any block requests appearing in some of the input sequences. This procedure makes all the input sequences the same, so we use any of them as an application launch sequence.

4.2 LBA-to-Inode Map

4.2.1 LBA-to-Inode Reverse Mapper

Our goal is to create an application prefetcher that generates exactly the same block request sequence as the obtained application launch sequence, where each block request is represented as a tuple of starting LBA and size. Since the application prefetcher is implemented as a user-level program, every disk access should be made via system calls with a file name and an offset in that file. Hence, we must obtain the file name and the offset of each block request in an application launch sequence.

Most file systems, including EXT3, do not support such a reverse mapping from LBA to file name and offset. However, for a given file name, we can easily find

the LBA of all of the blocks that belong to the file and their relative offset in the file. Hence, we can build a LBA-to-inode map by gathering this information for every file. However, building such a map of the entire file system is time-consuming and impractical because a file system, in general, contains tens of thousands of files and their block locations on the disk change very often.

Therefore, we build a separate LBA-to-inode map for each application, which can significantly reduce the overhead of creating a LBA-to-inode map because (1) the number of applications and the number of files used in launching each application are very small compared to the number of files in the entire file system; and (2) most of them are shared libraries and application code blocks, so their block locations remain unchanged unless they are updated or disk defragmentation is performed.

We implement the LBA-to-inode reverse mapper that receives a list of file names as input and creates a LBA-to-inode map as output. A LBA-to-inode map is built using a red-black tree in order to reduce the search time. Each node in the red-black tree has the LBA of a block as its key, and a block type as its data by default. According to the block type, different types of data are added to the node. A block type includes a super block, a group descriptor, an inode block bitmap, a data block bitmap, an inode table, and a data block. For example, a node for a data block has a block type, a device number, an inode number, an offset, and a size. Also, for a data block, a table is created to keep the mapping information between an inode number and its file name.

4.2.2 System-Call Profiler

The system-call profiler obtains a full list of file names that are accessed during an application launch,¹ and passes it to the LBA-to-inode reverse mapper. We used `strace` for the system-call profiler, which is a debugging tool in Linux. We can specify the argument of `strace` so that it may monitor only the system calls that have a file name as their argument. As many of these system calls are rarely called during an application launch, we monitor only the following system calls that frequently occur during application launches: `open()`, `creat()`, `execve()`, `stat()`, `stat64()`, `lstat()`, `lstat64()`, `access()`, `truncate()`, `truncate64()`, `statfs()`, `statfs64()`, `readlink()`, and `unlink()`.

¹Files mounted on pseudo file systems such as `procfs` and `sysfs` are not processed because they never generate any disk I/O request.

Table 1: System calls to replay access of blocks in an application launch sequence

Block type	System call
Inode table	<code>open()</code>
Data block: a directory	<code>opendir()</code> and <code>readdir()</code>
Data block: a regular file	<code>read()</code> or <code>posix_fadvise()</code>
Data block: a symbolic link file	<code>readlink()</code>

4.3 Application Prefetcher

4.3.1 Application Prefetcher Generator

The application prefetcher is a user-level program that replays the disk access requests made by a target application. We implemented the application prefetcher generator to automatically create an application prefetcher for each target application. It performs the following operations.

1. Read s_i one-by-one from S of the target application.
2. Convert s_i into its associated data items stored in the LBA-to-inode map, e.g.,
 $(dev, LBA, size) \rightarrow (datablk, filename, offset, size)$ OR
 $(dev, LBA, size) \rightarrow (inode, start_inode, end_inode)$.
3. Depending on the type of block, generate an appropriate system call using the converted disk access information.
4. Repeat Steps 1–3 until processing all s_i .

Table 1 shows the kind of system calls used for each block type. There are two system calls that can be used to replay the disk access for data blocks of a regular file. If we use `read()`, data is first moved from the SSD to the page cache, and then copying takes place from the page cache to the user buffer. The second step is unnecessary for our purpose, as the process that actually manipulates the data is not the application prefetcher but the target application. Hence, we chose `posix_fadvise()` that performs only the first step, from which we can avoid the overhead of `read()`. We use the `POSIX_FADV_WILLNEED` parameter, which informs the OS that the specified data will be used in the near future. When to issue the corresponding disk access after `posix_fadvise()` is called depends on the OS implementation. We confirmed that the current version of Linux we used issues a block request immediately after receiving the information through `posix_fadvise()`, thus meeting our need. A symbolic-linked file name is stored in data block pointers in an inode entry when the length of the file name is less than or equal to 60 bytes (c.f., the space of

```

int main(void) {
    ...
    readlink("/etc/fonts/conf.d/90-ttf-arphic-uming-embolden.conf", linkbuf, 256);
    int fd423;
    fd423 = open("/etc/fonts/conf.d/90-ttf-arphic-uming-embolden.conf", O_RDONLY);
    posix_fadvise(fd423, 0, 4096, POSIX_FADV_WILLNEED);
    posix_fadvise(fd351, 286720, 114688, POSIX_FADV_WILLNEED);
    int fd424;
    fd424 = open("/usr/share/fontconfig/conf.avail/90-ttf-arphic-uming-embolden.conf", O_RDONLY);
    posix_fadvise(fd424, 0, 4096, POSIX_FADV_WILLNEED);
    int fd425;
    fd425 = open("/root/.gnupg/trustdb.gpg", O_RDONLY);
    posix_fadvise(fd425, 0, 4096, POSIX_FADV_WILLNEED);
    dirp = opendir("/var/cache/");
    if (dirp) while (readdir (dirp));
    ...
    return 0;
}

```

Figure 4: An example application prefetcher.

data block pointers is 60 bytes, 4*12 for direct, 4 for single indirect, another 4 for double indirect, and last 4 for triple indirect data block pointer). If the length of linked file name is more than 60 bytes, the name is stored in the data blocks pointed to by data block pointers in the inode entry. We use `readlink()` to replay the data block access of symbolic-link file names that are longer than 60 bytes.

Figure 4 is an example of automatically-generated application prefetcher. Unlike the target application, the application prefetcher successively fetches all the blocks as soon as possible to minimize the time between adjacent block requests.

4.3.2 Implicitly-Prefetched Blocks

In the EXT3 file system, the inode of a file includes pointers of up to 12 data blocks, so these blocks can be found immediately after accessing the inode. If the file size exceeds 12 blocks, indirect, double indirect, and triple indirect pointer blocks are used to store the pointers to the data blocks. Therefore, requests for indirect pointer blocks may occur in the cold start scenario when the application is accessing files larger than 12 blocks. We cannot explicitly load those indirect pointer blocks in the application prefetcher because there is no such system call. However, the `posix_fadvise()` call for a data block will first make a request for the indirect block when needed, so it can be fetched in a timely manner by running the application prefetcher.

The following types of block request are not listed in Table 1: a superblock, a group descriptor, an inode entry bitmap, a data block bitmap. We found that requests to these types of blocks seldom occur during an application launch, so we did not consider their prefetching.

Table 2: Variables and parameters used by the application launch manager

Type	Description
n_{init}	A counter to record the number of application launches done in the initial launch phase
n_{pref}	A counter to record the number of launches done in the application prefetch phase after the last check of the miss ratio of the application prefetcher
N_{rawseq}	The number of raw block request sequences that are to be captured at the launch profiling phase
N_{chk}	The period to check the miss ratio of the application prefetcher
R_{miss}	A threshold value for the prefetcher miss ratio that is used to determine if an update of the application or shared libraries has taken place
T_{idle}	A threshold value for the idle time period that is used to determine if an application launch is completed
$T_{timeout}$	The maximum amount of time allowed for the disk I/O profiler to capture block requests

4.4 Application Launch Manager

The role of the application launch manager is to detect the launch of an application and to take an appropriate action. We can detect the beginning of an application launch by monitoring `execve()` system call, which is implemented using a system-call wrapper. There are three phases with which the application launch manager deals: a launch profiling phase, a prefetcher generation phase, and an application prefetch phase. The application launch manager uses a set of variables and parameters for each application to decide when to change its phase. These are summarized in Table 2.

Here we describe the operations performed in each phase:

(1) Launch profiling. If no application prefetcher is found for that application, the application launch manager regards the current launch as the first launch of this application, and enters the initial launch phase. In this phase, the application launch manager performs the following operations in addition to the launch of the target application:

1. Increase n_{init} of the current application by 1.
2. If $n_{init} = 1$, run the system call profiler.
3. Flush the page cache, dentries (directory entries), and inodes in the main memory to ensure a cold start scenario, which is done by the following command:

```
echo 3 > /proc/sys/vm/drop_caches
```
4. Run the disk I/O profiler. Terminate the disk I/O profiler when any of the following conditions are

met: (1) if no block request occurs during the last T_{idle} seconds or (2) the elapsed time since the start of the disk I/O profiler exceeds $T_{timeout}$ seconds.

5. If $n_{init} = N_{rawseq}$, enter the prefetcher generation phase after the current launch is completed.

(2) Prefetcher generation. Once application launch profiling is done, it is ready to generate an application prefetcher using the information obtained from the first phase. This can be performed either immediately after the application launch is completed, or when the system is idle. The following operations are performed:

1. Run the application launch sequence extractor.
2. Run the LBA-to-inode reverse mapper.
3. Run the application prefetcher generator.
4. Reset the values of n_{init} and n_{pref} to 0.

(3) Application prefetch. If the application prefetcher for the current application is found, the application launch manager runs the prefetcher simultaneously with the target application. It also periodically checks the miss ratio of the prefetcher to determine if there has been any update of the application or shared libraries. Specifically, the following operations are performed:

1. Increase n_{pref} of the current application by 1.
2. If $n_{pref} = N_{chk}$, reset the value of n_{pref} to 0 and run the disk I/O profiler. Its termination conditions are the same as those in the first phase.
3. Run the application prefetcher simultaneously with the target application.
4. If a raw block request sequence is captured, use it to calculate the miss ratio of the application prefetcher. If it exceeds R_{miss} , delete the application prefetcher.

The miss ratio is defined as the ratio of the number of block requests not issued by the prefetcher to the total number of block requests in the application launch sequence.

5 Performance Evaluation

5.1 Experimental Setup

Experimental platform. We used a desktop PC equipped with an Intel i7-860 2.8 GHz CPU, 4GB of PC12800 DDR3 SDRAM and an Intel 80GB SSD (X25-M G2 Mainstream). We installed a Fedora 12 with Linux kernel 2.6.32 on the desktop, in which we set NOOP as the default I/O scheduler. For benchmark applications, we chose frequently used user-interactive applications, for which application launch performance matters much. Such an application typically uses graphical

user interfaces and requires user interaction immediately after completing its launch. Applications like gcc and gzip are not included in our set of benchmarks as launch performance is not an issue for them. Our benchmark set consists of the following Linux applications: Acrobat reader, Designer-qt4, Eclipse, F-Spot, Firefox, Gimp, Gnome, Houdini, Kdevdesigner, Kdevelop, Konqueror, Labview, Matlab, OpenOffice, Skype, Thunderbird, and XilinxISE. In addition to these, we used Wine [1], which is an implementation of the Windows API running on the Linux OS, to test Access, Excel, Powerpoint, Visio, and Word—typical Windows applications.

Test scenarios. For each benchmark application, we measured its launch time for the following scenarios.

- *Cold start:* The application is launched immediately after flushing the page cache, using the method described in Section 4.4. The resulting launch time is denoted by t_{cold} .
- *Warm start:* We first run the application prefetcher only to load all the blocks in the application launch sequence to the page cache, and then launch the application. Let t_{warm} denote the resulting launch time.
- *Sorted prefetch:* To evaluate the performance of the sorted prefetch [15, 25, 36] on SSDs, we modify the application prefetcher to fetch the block requests in the application launch sequence in the sorted order of their LBAs. After flushing the page cache, we first run the modified application prefetcher, then immediately run the application. Let t_{sorted} denote the resulting launch time.
- *FAST:* We flush the page cache, and then run the application simultaneously with the application prefetcher. The resulting launch time is denoted by t_{FAST} .
- *Prefetcher only:* We flush the page cache and run the application prefetcher. The completion time of the application prefetcher is denoted by t_{ssd} . It is used to calculate a lower bound of the application launch time $t_{bound} = \max(t_{ssd}, t_{cpu})$, where $t_{cpu} = t_{warm}$ is assumed.

Launch-time measurement. We start an application launch by clicking an icon or inputting a command, and can accurately measure the launch start time by monitoring when `execve()` is called. Although it is difficult to clearly define the completion of a launch, a reasonable definition is the first moment the application becomes responsive to the user [2]. However, it is difficult to accurately and automatically measure that moment. So, as an alternative, we measured the completion time of the last block request in an application launch sequence using `Blktrace`, assuming that the launch will be completed




Figure 5: The size of application launch sequences.

very soon after issuing the last block request. For the warm start scenario, we executed `posix_fadvise()` with `POSIX_FADV_DONTNEED` parameter to evict the last block request from the page cache. For the sorted prefetch and the FAST scenarios, we modified the application prefetcher so that it skips prefetching of the last block request.

5.2 Experimental Results

Application launch sequence generation. We captured 10 raw block request sequences during the cold start launch of each application. We ran the application launch sequence extractor with a various number of input block request sequences, and observed the size of the resulting application launch sequences. Figure 5 shows that for all the applications we tested, there is no significant reduction of the application launch sequence size while increasing the number of inputs from 2 to 10. Hence, we set the value of N_{rawseq} in Table 2 to 2 in this paper. We used the size of the first captured input sequence as the number of inputs one in Figure 5 (the application launch sequence extractor requires at least two input sequences). For some applications, there are noticeable differences in size between the number of inputs one and two. This is because the first raw input request sequence includes a set of bursty I/O requests generated by OS and user daemons that are irrelevant to the application launch. Figure 5 shows that such I/O requests can be effectively excluded from the resulting application launch sequence using just two input request sequences.

The second and third columns of Table 3 summarize the total number of block requests and accessed blocks of the thus-obtained application launch sequences, respectively. The last column shows the total number of files used during the launch of each application.

Testing of the application prefetcher. Application prefetchers are automatically generated for the benchmark applications using the application launch sequences in Table 3. In order to see if the application prefetchers fetch all the blocks used by an application, we first flushed the page cache, and launched each applica-

Table 3: Collected launch sequences ($N_{rawseq} = 2$)

Application	# of block requests	# of fetched blocks	# of used files
Access	1296	106 992	555
Acrobat reader	960	73 784	178
Designer-qt4	2400	138 608	410
Eclipse	4163	155 216	787
Excel	1610	169 112	583
F-Spot	1180	49 968	304
Firefox	1566	60 944	433
Gimp	1939	66 928	799
Gnome	4739	228 872	538
Houdini	4836	290 320	724
Kdevdesigner	1537	44 904	467
Kdevelop	1970	63 104	372
Konqueror	1780	62 216	296
Labview	2927	154 768	354
Matlab	6125	267 312	742
OpenOffice	1425	104 600	308
Powerpoint	1405	120 808	576
Skype	892	41 560	197
Thunderbird	1533	64 784	429
Visio	1769	168 832	662
Word	1715	181 496	613
Xilinx ISE	4718	328 768	351

tion immediately after running the application prefetcher. During the application launch, we captured all the block requests generated using `Blktrace`, and counted the number of missed block requests. The average number of missed block requests was 1.6% of the number of block requests in the application launch sequence, but varied among repeated launches, e.g., from 0% to 6.1% in the experiments we performed.

By examining the missed block requests, we could categorize them into three types: (1) files opened by OS daemons and user daemons at boot time; (2) journaling data or swap partition accesses; and (3) files dynamically created or renamed at every launch (e.g., `tmpfile()`). The first type occurs because we force the page cache to be flushed in the experiment. In reality, they are highly likely to reside in the page cache, and thus, this type of misses will not be a problem. The second type is irrelevant to the application, and observed even during idle time. The third type occurs more or less often, depending on the application. FAST does not prefetch this type of block requests as they change at every launch.

Experiments for the test scenarios. We measured the launch time of the benchmark applications for each test scenario listed in Section 5.1. Figure 6 shows that the average launch time reduction of FAST is 28% over the cold start scenario. The performance of FAST varies considerably among applications, ranging from 16% to 46% reduction of launch time. In particular, FAST shows




Figure 6: Measured application launch time (normalized to t_{cold}).




Figure 7: Usage of CPU and SSD (sampling rate = 1 KHz).

performance very close to t_{bound} for some applications, such as Eclipse, Gnome, and Houdini. On the other hand, the gap between t_{bound} and t_{FAST} is relatively larger for such applications as Acrobat reader, Firefox, OpenOffice, and Labview.

Launch time behavior. We conducted experiments to see if the application prefetcher works well as expected when it is simultaneously run with the application. We chose Firefox because it shows a large gap between t_{bound} and t_{FAST} . We monitored the generated block requests during the launch of Firefox with the application prefetcher, and observed that the first 12 of the entire 1566 block requests were issued by Firefox, which took about 15 ms. As the application prefetcher itself should be launched as well, FAST cannot prefetch these block requests until finishing its launch. However, we observed that all the remaining block requests were issued by FAST, meaning that they are successfully prefetched before the CPU needs them.

CPU and SSD usage patterns. We performed another experiment to observe the CPU and SSD usage patterns in each test scenario. We chose two applications, Eclipse and Firefox, representing the two groups of applications

of which t_{FAST} is close to and far from t_{bound} , respectively. We modified the OS kernel to sample the number of CPU cores having runnable processes and to count the number of cores in the I/O wait state. Figure 7 shows the CPU and SSD usage of the two applications, where the entire CPU is regarded as busy if at least one of its cores is active. Similarly, the SSD is assumed busy if there are one or more cores in the I/O wait state. In the cold start scenario, there is almost no overlap between CPU computation and SSD access for both applications. In the warm start scenario, the CPU stays fully active until the launch is completed as there is no wait. One exception we observed is the time period marked with Circle (a), during which the CPU seems to be in the event-waiting state. FAST is shown to be successful in overlapping CPU computation with SSD access as we intended. However, CPU usage is observed to be low at the beginning of launch for both applications, which can be explained with the example in Figure 2. As Eclipse shows a shorter such time period (Circle (b)) than Firefox (Circle (c)), t_{FAST} can reach closer to t_{bound} . In the case of Firefox, however, the ratio of t_{cpu} to t_{ssd} is close to 1:1, allowing FAST to achieve more reduction




Figure 8: Simultaneous launch of multiple applications.

of launch time for Firefox than for Eclipse.

Performance of sorted prefetch. Figure 6 shows that the sorted prefetch reduces the application launch time by an average of 7%, which is less efficient than FAST, but non-negligible. One reason for this improvement is the difference in I/O burstiness between the cold start and the sorted prefetch. Most SSDs (including the one we used) support the native command queueing (NCQ) feature, which allows up to 31 block requests to be sent to a SSD controller. Using this information, the SSD controller can read as many NAND flash chips as possible, effectively increasing read throughput. The average queue depth in the cold start scenario is close to 1, meaning that for most of time there is only one outstanding request in case of SSD. In contrast, in the sorted prefetch scenario, the queue depth will likely grow larger than 1 because the prefetcher may successively issue asynchronous I/O requests using `posix_fadvise()`, at small inter-issue intervals.

On the other hand, we could not find a clear evidence that sorting block requests in their LBA order is advantageous in case of SSD. Rather, the execution time of the sorted prefetcher was slightly longer than its unsorted version for most of the applications we tested. Also, the sorted prefetch shows worse performance than the cold start for Excel, Powerpoint, Skype, and Word. Although these observations were consistent over repeated tests, a further investigation is necessary to understand such a behavior.

Simultaneous launch of applications. We performed experiments to see how well FAST can scale up for launching multiple applications. We launched multiple applications starting from the top of Table 3, adding five at a time, and measured the launch completion time of all launched applications². Figure 8 shows that FAST could reduce the launch completion time for all the tests, whereas the sorted prefetch does not scale beyond 10 applications. Note that the FAST improvement decreased from 20% to 7% as the number of applications increased from 5 to 20.

Runtime and space overhead. We analyzed the runtime overhead of FAST for seven possible combinations of running processes, and summarized the results in Ta-

²Except for Gnome that cannot be launched with other applications, and Houdini whose license had expired.

Table 4: Runtime overhead (application: Firefox)

Running processes	Runtime (sec)
1. Application only (cold start scenario)	0.86
2. <code>strace + blktrace + application</code>	1.21
3. <code>blktrace + application</code>	0.88
4. Prefetcher generation	5.01
5. Prefetcher + application	0.56
6. Prefetcher + <code>blktrace</code> + application	0.59
7. Miss ratio calculation	0.90

ble 4. Cases 2 and 3 belong to the *launch profiling* phase, which was described in Section 4.4. During this phase, Case 2 occurs only once, and Case 3 occurs N_{rawseq} times. Case 4 corresponds to the *prefetcher generation* phase (the right side of Figure 3), and shows a relatively long runtime. However, we can hide it from users by running it in background. Also, since we primarily focused on functionality in the current implementation, there is room for further optimization. Cases 5, 6, and 7 belong to the *application prefetch* phase, and repeatedly occur until the application prefetcher is invalidated. Cases 6 and 7 occur only when n_{pref} reaches N_{chk} , and Case 7 can be run in background.

FAST creates temporary files such as system call log files and I/O traces, but these can be deleted after FAST completes creating application prefetchers. However, the generated prefetchers occupy disk space as far as application prefetching is used. In addition, application launch sequences are stored to check the miss ratio of the corresponding application prefetcher. In our experiment, the total size of the application prefetchers and application launch sequences for all 22 applications was 7.2 MB.

FAST applicability. While previous examples clearly demonstrated the benefits of FAST for a wide range of applications, FAST does not guarantee improvements for all cases. One such a scenario is when a target application is too small to offset the overhead of loading the prefetcher. We tested FAST with the Linux utility `uname`, which displays the name of the OS. It generated 3 I/O requests whose total size was 32 KB. The measured t_{cold} was 2.2 ms, and t_{FAST} was 2.3 ms, 5% longer than the cold start time.

Another possible scenario is when the target application experiences a major update. In this scenario, FAST may fetch data that will not be used by the newly updated application until it detects the application update and enters a new launch profiling phase. We modified the application prefetcher so that it fetches the same size of data from the same file but from another offset that is not used by the application. We tested the modified prefetcher with Firefox. Even in this case, FAST reduced application launch time by 4%, because FAST

could still prefetch some of the metadata used by the application. Assuming most of the file names are changed after the update, we ran Firefox with the prefetcher for Gimp, which fetches a similar number of blocks as Firefox. In this experiment, the measured application launch time was 7% longer than the cold start time, but the performance degradation was not drastic due to the internal parallelism of the SSD we used (10 channels).

Configuring application launch manager. The application launch manager has a set of parameters to be configured, as shown in Table 2. If N_{rawseq} is set too large, users will experience the cold-start performance during the initialization phase. If it is set too small, unnecessary blocks may be included in the application prefetcher. Figure 5 shows that setting it between 2 and 4 is a good choice. The proper value of N_{chk} will depend on the runtime overhead of `Blktrace`; if FAST is placed in the OS kernel, the miss ratio of the application prefetcher may be checked upon every launch ($N_{chk} = 1$) without noticeable overhead. Also, setting R_{miss} to 0.1 is reasonable, but it needs to be adjusted after gaining enough experience in using FAST. To find the proper value of T_{idle} , we investigated the SSD’s maximum idle time during the cold-start of applications, and found it to range from 24 ms (Thunderbird) to 826 ms (Xilinx ISE). Hence, setting T_{idle} to 2 seconds is proper in practice. As the maximum cold-start launch time is observed to be less than 10 seconds, 30 seconds may be reasonable for $T_{timeout}$. All these values may need to be adjusted, depending on the underlying OS and applications.

Running FAST on HDDs. To see how FAST works on a HDD, we replaced the SSD with a Seagate 3.5” 1 TB HDD (ST31000528AS) and measured the launch time of the same set of benchmark applications. Although FAST worked well as expected by hiding most of CPU computation from the application launch, the average launch time reduction was only 16%. It is because the application launch on a HDD is mostly I/O bound; in the cold start scenario, we observed that about 85% of the application launch time was spent on accessing the HDD. In contrast, the sorted prefetch was shown to be more effective; it could reduce the application launch time by an average of 40% by optimizing disk head movements.

We performed another experiment by modifying the sorted prefetch so that the prefetcher starts simultaneously with the original application, like FAST. However, the resulting launch time reduction was only 19%, which is worse than that of the unmodified sorted prefetch. The performance degradation is due to the I/O contention between the prefetcher and the application.




Figure 9: Measured application launch time on iPhone 4 (CPU: 1 GHz, SDRAM: 512 MB, NAND flash: 32 GB).

6 Applicability of FAST to Smartphones

The similarity between modern smartphones and PCs with SSDs in terms of the internal structure and the usage pattern, as summarized below, makes smartphones a good candidate to which we can apply FAST:

- Unlike other mobile embedded systems, smartphones run different applications at different times, making application launch performance matter more;
- Smartphones use NAND flash as their secondary storage, of which the performance characteristics are basically the same as the SSD; and
- Smartphones often use slightly customized (if not the same) OSes and file systems that are designed for PCs, reducing the effort to port FAST to smartphones.

Furthermore, a smartphone has the characteristics that enhance the benefit of using FAST as follows:

- Users tend to launch and quit applications more frequently on smartphones than on PCs;
- Due to relatively smaller main memory of a smartphone, users will experience cold start performance more frequently; and
- Its relatively slower CPU and flash storage speed may increase the absolute reduction of application launch time by applying FAST.

Although we have not yet implemented FAST on a smartphone, we could measure the launch time of some smartphone applications by simply using a stopwatch. We randomly chose 14 applications installed on the iPhone 4 to compare their cold and warm start times, of which the results are plotted in Figure 9. The average cold start time of the smartphone applications is 6.1 seconds, which is more than twice of the average cold start time of the PC applications (2.4 seconds) shown in Figure 6. Figure 9 also shows that the average warm start time is 63% of the cold start time (almost the same ratio as in Figure 6), implying that we can achieve similar benefits from applying FAST to smartphones.

7 Comparison of FAST with Traditional Prefetching

FAST is a special type of prefetching optimized for application launch, whereas most of the traditional prefetching schemes focus on runtime performance improvement. We compare FAST with the traditional prefetching algorithms by answering the following three questions that are inspired by previous work [32].

7.1 What to Prefetch

FAST prefetches the blocks appeared in the application launch sequence. While many prediction-based prefetching schemes [9, 23, 39] suffer from the low hit ratio of the prefetched data, FAST can achieve near 100% hit ratio. This is because the application launch sequence changes little over repeated launches of an application, as observed by previous work [4, 18, 34].

Sequential pattern detection schemes like readahead [13, 31] can achieve a fairly good hit ratio when activated, but they are applicable only when such a pattern is detected. By contrast, FAST guarantees stable performance improvement for every application launch.

One way to enhance the prefetch hit ratio for a complicated disk I/O pattern is to analyze the application source code to extract its access pattern. Using the thus-obtained pattern, prefetching can be done by either inserting prefetch codes into the application source code [29, 38] or converting the source code into a computation thread and a prefetch thread [40]. However, such an approach does not work well for application launch optimization because many of the block requests generated during an application launch are not from the application itself but from other sources, such as loading shared libraries, which cannot be analyzed by examining the application source code. Furthermore, both require modification of the source code, which is usually not available for most commercial applications. Even if the source code is available, modifying and recompiling every application would be very tedious and inconvenient. In contrast, FAST does not require application source code and is thus applicable for any commercial application.

Another relevant approach [6] is to deploy a shadow process that speculatively executes the copy of the original application to get hints for the future I/O requests. It does not require any source modification, but consumes non-negligible CPU and memory resources for the shadow process. Although it is acceptable when CPU is otherwise stalled waiting for the I/O completion, employing such a shadow process in FAST may degrade application launch performance as there is not enough CPU idle period as shown in Figure 7.

7.2 When to Prefetch

FAST is not activated until an application is launched, which is as conservative as demand paging. Thus, unlike prediction-based application prefetching schemes [12, 28], there is no cache-pollution problem or additional disk I/O activity during idle period. However, once activated, FAST aggressively performs prefetching: it keeps on fetching subsequent blocks in the application launch sequence *asynchronously* even in the absence of page misses. As the prefetched blocks are mostly (if not all) used by the application, the performance improvement of FAST is comparable to that of the prediction-based schemes when their prediction is accurate.

7.3 What to Replace

FAST does not modify the replacement algorithm of page cache in main memory, so the default page replacement algorithm is used to determine which page to evict in order to secure free space for the prefetched blocks.

In general, prefetching may significantly affect the performance of page replacement. Thus, previous work [5, 14, 35] emphasized the need for integrated prefetching and caching. However, FAST differs from the traditional prefetching schemes since it prefetches only those blocks that will be referenced before the application launch completes (e.g., in next few seconds). If the page cache in the main memory is large enough to store all the blocks in the application launch sequence, which is commonly the case, FAST will have minimal effect on the optimality of the page replacement algorithm.

8 Conclusion

We proposed a new I/O prefetching technique called *FAST* for the reduction of application launch time on SSDs. We implemented and evaluated FAST on the Linux OS, demonstrating its deployability and performance superiority. While the HDD-aware application launcher showed only 7% of launch time reduction on SSDs, FAST achieved a 28% reduction with no additional overhead, demonstrating the need for, and the utility of, a new SSD-aware optimizer. FAST with a well-designed entry-level SSD can provide end-users the fastest application launch performance. It also incurs fairly low implementation overhead and has excellent portability, facilitating its wide deployment in various platforms.

Acknowledgments

We deeply appreciate Prof. Heonshik Shin for his support and providing research facility. We also thank our

shepherd Arkady Kanevsky, and the anonymous reviewers for their invaluable comments that improved this paper. This research was supported by WCU (World Class University) program through National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (R33-10085), and RP-Grant 2010 of Ewha Womans University. Sangsoo Park is the corresponding author (email: sangsoo.park@ewha.ac.kr).

References

- [1] *Wine User Guide*. <http://www.winehq.org/docs/wineusr-guide/index>, Last accessed on: 17 November 2010.
- [2] APPLE INC. *Launch Time Performance Guidelines*. <http://developer.apple.com/documentation/Performance/Conceptual/LaunchTime/LaunchTime.pdf>, 2006.
- [3] AXBOE, J. *Block IO Tracing*. <http://www.kernel.org/git/?p=linux/kernel/git/axboe/blktrace.git;a=blob;f=README>, 2006.
- [4] BHADKAMKAR, M., GUERRA, J., USECHE, L., BURNETT, S., LIPTAK, J., RANGASWAMI, R., AND HRISTIDIS, V. BORG: Block-reORGanization for self-optimizing storage systems. In *Proc. FAST* (2009), pp. 183–196.
- [5] CAO, P., FELTEN, E. W., KARLIN, A. R., AND LI, K. A study of integrated prefetching and caching strategies. In *Proc. SIGMETRICS* (1995), pp. 188–197.
- [6] CHANG, F., AND GIBSON, G. A. Automatic I/O hint generation through speculative execution. In *Proc. OSDI* (1999), pp. 1–14.
- [7] CHEN, F., KOUFATY, D. A., AND ZHANG, X. Understanding intrinsic characteristics and system implications of flash memory based solid state drives. In *Proc. SIGMETRICS* (2009), pp. 181–192.
- [8] COLITTI, L. Analyzing and improving GNOME startup time. In *Proc. SANE* (2006), pp. 1–11.
- [9] CUREWITZ, K. M., KRISHNAN, P., AND VITTER, J. S. Practical prefetching via data compression. *SIGMOD Rec.* 22, 2 (1993), 257–266.
- [10] DIRIK, C., AND JACOB, B. The performance of PC solid-state disks (SSDs) as a function of bandwidth, concurrency, device architecture, and system organization. In *Proc. ISCA* (2009), pp. 279–289.
- [11] DUNN, M., AND REDDY, A. L. N. A new I/O scheduler for solid state devices. Tech. Rep. TAMU-ECE-2009-02, Department of Electrical and Computer Engineering, Texas A&M University, 2009.
- [12] ESFAHBOD, B. Preload—An adaptive prefetching daemon. Master’s thesis, Graduate Department of Computer Science, University of Toronto, Canada, 2006.
- [13] FENGGUANG, W., HONGSHENG, X., AND CHENFENG, X. On the design of a new Linux readahead framework. *SIGOPS Oper. Syst. Rev.* 42, 5 (2008), 75–84.
- [14] GILL, B. S., AND MODHA, D. S. SARC: Sequential prefetching in adaptive replacement cache. In *Proc. USENIX* (2005), pp. 293–308.
- [15] HUBERT, B. On faster application startup times: Cache stuffing, seek profiling, adaptive preloading. In *Proc. OLS* (2005), pp. 245–248.
- [16] INTEL. *Intel Turbo Memory with User Pinning*. Intel, <http://www.intel.com/design/flash/nand/turbomemory/index.htm>, Last accessed on: 17 November 2010.
- [17] JO, H., KIM, H., JEONG, J., LEE, J., AND MAENG, S. Optimizing the startup time of embedded systems: A case study of digital TV. *IEEE Trans. Consumer Electron.* 55, 4 (2009), 2242–2247.
- [18] JOO, Y., CHO, Y., LEE, K., AND CHANG, N. Improving application launch times with hybrid disks. In *Proc. CODES+ISSS* (2009), pp. 373–382.
- [19] KAMINAGA, H. Improving Linux startup time using software resume. In *Proc. OLS* (2006), pp. 25–34.
- [20] KIM, H., AND AHN, S. BPLRU: A buffer management scheme for improving random writes in flash storage. In *Proc. FAST* (2008), pp. 1–14.
- [21] KIM, J., OH, Y., KIM, E., CHOI, J., LEE, D., AND NOH, S. H. Disk schedulers for solid state drives. In *Proc. EMSOFT* (2009), pp. 295–304.
- [22] KIM, Y.-J., LEE, S.-J., ZHANG, K., AND KIM, J. I/O performance optimization techniques for hybrid hard disk-based mobile consumer devices. *IEEE Trans. Consumer Electron.* 53, 4 (2007), 1469–1476.
- [23] KOTZ, D., AND ELLIS, C. S. Practical prefetching techniques for parallel file systems. In *Proc. PDIS* (1991), pp. 182–189.
- [24] LARUS, J. Spending Moore’s dividend. *Commun. ACM* 52, 5 (2009), 62–69.
- [25] LICHOTA, K. *Prefetch: Linux solution for prefetching necessary data during application and system startup*. <http://code.google.com/p/prefetch/>, 2007.
- [26] MATTHEWS, J., TRIKA, S., HENSGEN, D., COULSON, R., AND GRIMSRUD, K. Intel® Turbo Memory: Nonvolatile disk caches in the storage hierarchy of mainstream computer systems. *ACM Trans. Storage* 4, 2 (2008), 1–24.
- [27] MICROSOFT. *Support and Q&A for Solid-State Drives*. Microsoft, <http://blogs.msdn.com/e7/archive/2009/05/05/support-and-q-a-for-solid-state-drives-and.aspx>, 2009.
- [28] MICROSOFT. *Windows PC Accelerators*. <http://www.microsoft.com/whdc/system/sysperf/perfaccel.msp>, Last accessed on: 17 November 2010.
- [29] MOWRY, T. C., DEMKE, A. K., AND KRIEGER, O. Automatic compiler-inserted I/O prefetching for out-of-core applications. In *Proc. OSDI* (1996), pp. 3–17.
- [30] NEELAKANTH NADGIR. *Reducing Application Startup Time in the Solaris 8 OS*. http://developers.sun.com/solaris/articles/reducing_app.html, 2002.
- [31] PAI, R., PULAVARTY, B., AND CAO, M. Linux 2.6 performance improvement through readahead optimization. In *Proc. OLS* (2004), pp. 105–116.
- [32] PAPATHANASIOU, A. E., AND SCOTT, M. L. Energy efficient prefetching and caching. In *Proc. USENIX* (2004), pp. 22–22.
- [33] PARK, C., KIM, K., JANG, Y., AND HYUN, K. Linux bootup time reduction for digital still camera. In *Proc. OLS* (2006), pp. 239–248.
- [34] PARUSH, N., PELLEG, D., BEN-YEHUDA, M., AND TA-SHMA, P. Out-of-band detection of boot-sequence termination events. In *Proc. ICAC* (2009), pp. 71–72.
- [35] PATTERSON, R. H., GIBSON, G. A., GINTING, E., STODOLSKY, D., AND ZELENKA, J. Informed prefetching and caching. In *Proc. SOSP* (1995), pp. 79–95.

- [36] RUSSINOVICH, M. E., AND SOLOMON, D. *Microsoft Windows Internals*, 4th ed. Microsoft Press, 2004, pp. 458–462.
- [37] SAMSUNG SEMICONDUCTOR. *Samsung Hybrid Hard Drive*. <http://www.samsung.com/global/business/semiconductor/support/brochures/downloads/hdd/hd.datasheet.200708.pdf>, 2007.
- [38] VANDEBOGART, S., FROST, C., AND KOHLER, E. Reducing seek overhead with application-directed prefetching. In *Proc. USENIX* (2009).
- [39] VELLANKI, V., AND CHERVENAK, A. L. A cost-benefit scheme for high performance predictive prefetching. In *Proc. SC* (1999).
- [40] YANG, C.-K., MITRA, T., AND CHIUEH, T.-C. A decoupled architecture for application-specific file prefetching. In *Proc. USENIX* (2002), pp. 157–170.