

A Theoretical Dissection of Demand Paging and the Comparative Dynamics of Page Replacement Policies in Virtual Memory Management

Mr. Rahul Kumar¹, Parul Kashyap²

¹(Assistant Professor) Department of Computer Science & Engineering Bharat Institute of Technology, Meerut

²(M.Tech Student) Department of Computer Science & Engineering Bharat Institute of Technology, Meerut

Abstract- A key idea in contemporary operating systems that permits effective use of physical memory is virtual memory management. Performance optimization heavily relies on demand paging, a lazy-loading approach that only retrieves pages when required. Demand paging is theoretically broken down in this study, which also critically analyzes the effectiveness and behavior of the main page replacement policies, such as Optimal Replacement, LRU, and FIFO (First-In First-Out). We investigate how each policy affects page fault frequency, memory access time, and system throughput using a conceptual framework. The paper's conclusion provides a clear grasp of when and why particular strategies perform better under different workloads through a comparative review based on theoretical insights rather than actual simulations.

Keywords: Demand Paging, Page Replacement Algorithms, Memory Management, Page Fault, Virtual Memory.

I. INTRODUCTION

The Atlas computer's creators initially proposed the "demand paging" technique of storage allocation, which has an enticingly straightforward conceptual design. Furthermore, it was, to put it mildly, quite alluring to consider the idea that this technology may be used to let programmers overlook the challenges of structuring the information flow between working store and one or more levels of backing storage. Demand paging does have several drawbacks, though. Some were anticipated; for example, Dennis and Glaser² noted a number of years ago that page-turning may be disastrous if used for the incorrect type of information. However, this did not stop the deployment of systems that had significant performance issues due to excessive paging in at

least their early iterations. . Numerous more-or-less ad hoc attempts to enhance system performance have been sparked by these issues. This paper's goal is to examine the different methods that have been put forth or employed in an effort to enhance demand paging systems' performance. It is intended that this will help clarify the connections between the seemingly disparate approaches and facilitate comprehension of demand paging's advantages and disadvantages.

Demand paging

The title of demand paging makes it apparent what its two key features are:

(i) A program requests information at the precise moment it is required, without warning. In actuality, the "demand" is implicit and results from an attempt to use data that isn't already in operational storage.

(ii) Data is moved in and out of working storage in page-sized units. These pages are of the same size, usually 1024 words. As a result, it is believed that the working storage is logically divided into page frames.

All of these features point to one advantage of demand paging, which is that it relieves the programmer of the responsibility of clearly specifying what data needs to be moved to and from working storage and at what times. But as the major body of this paper will make evident, these two traits are also the cause of the performance issues that certain systems really encounter [3] and that several analytical⁴ and simulation studies predict. [5].

Low CPU usage and high channel consumption between working and backup storage, which results from excessive page-turning (excessive, that is, in relation to the amount of processing done), are the typical symptoms of these performance issues. A system that behaves in this way is frequently referred to as "paging itself to death."

There are two main reasons for these inefficiencies.

First, a program will continue to use system resources, particularly working storage, while it waits for a page request to be fulfilled. Reference [6] discusses this, and Figure 1 provides a graphic representation of it. It is obvious that a large portion of resource usage will be ineffective if page transfers occur frequently or take a long period.

The second reason is a little more subtle and results from the nonlinear relationship that has been discovered between the number of page frames allotted to the program and program efficiency (such as Q average number of instructions executed between page demands). [7]. Figure 2 provides an illustration of this relationship. It will be observed that while a program can function very well with a somewhat smaller set of pages than the totality it references, performance quickly deteriorates when the number of pages is lowered below a particular threshold, which is frequently referred to as the program's parachor. [8].

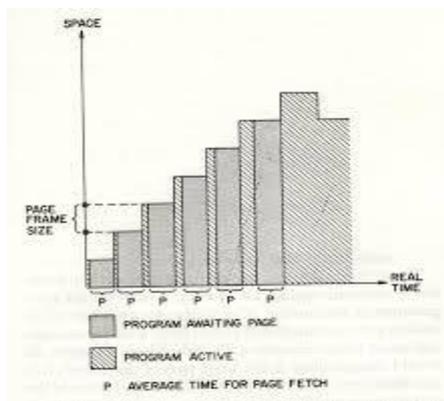


Figure 1: *Storage Demand paging combined with storage utilization*

When deciding whether a demand paging system will be paged to death, three fundamental considerations come into play. They are:

- (i) The hardware properties of the system
- (ii) The load of the application
- (iii) The tactics of the operating system

The first of these is, in a way, the most crucial since an attack targeting this region has the best chance of successfully bringing the paging rate down to a manageable level. For instance, all but the most twisted operating system tactics will be defeated by a significant increase in the working storage's capacity. Nevertheless, increasing working storage space indefinitely is obviously a simple fix.

The challenging issue is figuring out how much working storage must be provided in addition to different backup store types in order to achieve the necessary performance level at the lowest possible cost. The impact of imbalanced system designs, where optimization of all but the most crucial system part has had minimal impact on total system performance, has been demonstrated by Nielsen [8]. It is outside the purview of this study to describe the methods available for identifying the elements of this balance in particular systems, but a few general observations might be appropriate.

Belady[7] has demonstrated how lowering the latency involved in getting information from backing stores affects system performance. These findings imply that the bulk core storage should be used to address the performance issue. [10]. However, because of their very high cost, it is currently challenging to defend the usage of such devices on a cost/performance basis in anything other than the most demanding of situations. On the other hand, the same findings might indicate that flailing arm disks are not a good option for demand-paging settings when it comes to backing storage. In systems that use these devices, the core latency effect seen in Figure 1 essentially prevents effective functioning unless a large quantity of working store capacity is provided. As a result, even if these devices are inexpensive, the performance they can achieve does not result in a profitable cost/performance ratio.

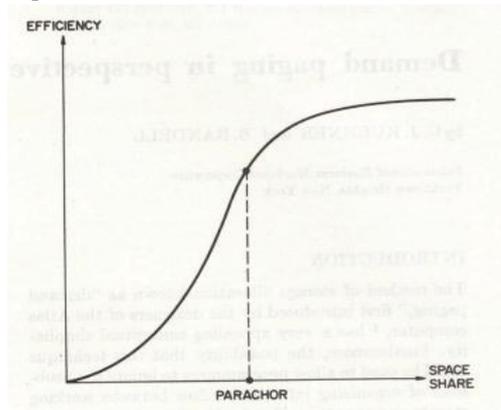


Figure 2: *Space sharing versus efficiency*

It is practically a given that a satisfactory degree of performance may be achieved if one is allowed to change the hardware complement. However, in some ways, these hardware solutions are a brute force strategy meant to create an environment that is insensitive to the challenges brought on by operating system and program load techniques [2]. Attacks on operating system tactics and program load, on the other hand, can be viewed as an effort to increase demand paging's practicality. Thus, the remainder of this study focuses exclusively on the several software solutions that have been suggested to address the demand paging issues. However, two warnings should be issued before starting this conversation. The cost issue comes first.[1]

It is simple to calculate the cost of hardware solutions in terms of purchase or renting money. With software solutions generally, this isn't the case. Comparing different solutions at the same price is desirable, but it is impossible without understanding the cost and level of software expertise required for a certain installation. Second, a system's transition from an undesirable to an acceptable level is not guaranteed by the application of one, two, or all of the strategies outlined below.

Program load

Both user programs and system programs are made up of the data and instructions whose transmission and storage place a strain on the system's channels and storage. Actually, certain operating system components use the very resources that they are responsible for allocating.

Both the volume and the structure of this load have the potential to cause issues. It should be easy to comprehend how the amount of data and instructions affects the restricted storage and channel resources. The consequences of the program load's structure are considerably more subdued. The working set size is the only easily ascertained metric pertaining to the composition of a specific program load.

[11]. The group of unique pages to which storage references were made during a specific execution interval is known as the program's working set. The number of pages that a program and its data take up is thus its greatest working set size. On the other extreme, a single page can also be considered a functional set, but only for a very little period of

time. (Usually, it takes more than one page to execute even one individual instruction.)

Therefore, the working set size depends on how long storage references are monitored. It is evident that the average working set size plotted against the observed execution interval length increases monotonically[4].

It should be kept in mind, nevertheless, that during a specific observation period, various program execution stages may lead to wildly disparate working set sizes and compositions.

Regretfully, the underlying features of the program load structure that influence the working set size measure's value are far removed. These attributes consist of:

- (i) Programming style;
- (ii) Code modularity; and
- (iii) Data layout.

Reducing the average working set size and program load volume is the main objective of all the different paging system efficiency-boosting strategies covered in this section.

Enhancement of program load without recoding

It is not particularly appealing to consider altering the data structures and instruction flow of an already complex program load in order to lower the average working set size.

Therefore, several attempts have been made to accomplish this reduction by merely repackaging the modules that comprise the program layout in virtual memory, which may be larger or smaller than a page. In a brief set of experiments described by Comeau[12], the system loader was shown the modules that made up a whole operating system in a variety of different orders. It was discovered that one change decreased the page transfer rate by more than 80% compared to the previous operating system version.

Assuming one knows how to select a suitable module reordering, this strategy is evidently quite easy to implement. At least some knowledge of the dynamic interplay between software modules and the pattern of referencing activity is necessary for any logical decision.

- (i) Frequency of reference to the various modules;
- (ii) sequence of references to the modules;
- (iii) frequency of reference to the various pages; and

(iv) Sequence of references to pages are examples of possible information types.

During operating system trial runs, such data could be acquired. Immediately, two issues spring to mind. First off, data collected from many runs cannot be guaranteed to be consistent.

Enhancement of program load by redesign and recoding

When hardware or operating system architecture is subpar, a paging system's performance demands a very different kind of program design than is typically needed. In these situations, following a set of programming guidelines is crucial and repulsive for both programmers and compilers. The need for the programmer to maintain constant knowledge of the (at least approximate) location of page borders in respect to his instruction and data layouts is what causes the aversion.

Among the essential rules are:

(i) Avoid making frequent references to a large number of different sites succession. Instead of intermixing references to the much larger combined locality, it is preferable to localize activities for a suitable amount of time before moving on to another location.

(ii) Reduce the amount of space needed for data storage and instruction solely to the extent that doing so enables compliance with command(i). This is predicated on the idea that it makes sense to exchange higher backup storage usage for lower paging activity.

(iii) Steer clear of very modular programs. Program modularity can significantly lower execution performance, even if it is particularly useful for introducing updates and revisions to existing codes. Even the smallest task requests require tens or even hundreds of control transfers if an operating system is made up of literally hundreds of program modules [9]. In addition to being expensive in terms of wasted CPU time, these transfers may, in the worst scenario, need referencing a separate page of data for each transfer. Therefore, it should be evident that there is a clear trade-off between the quantity of dynamic control transfers that immediately result and the effective static level of program modularity.

(iv) Consider the sequence in which data items should be processed while designing data layouts. For instance, it would be foolish to store a $n \times n$ matrix by columns if it were to be periodically scanned by rows (where n is marginally higher than the page size). (For a thorough explanation, go to McKellar and Coffman.[13]).

The importance of following these guidelines increases with the frequency of use of a given program. Because of this, programmers need to exercise extra caution when creating necessary systems. Additionally, it should be mentioned that regulations must be followed when making changes to the software, which can be challenging when the changes are significant and frequent.

To enhance their paging speed, some reasonably successful attempts have been made to rewrite fairly simple algorithms. For instance, list processing systems with dynamic storage allocation of list elements that groups elements that are likely to be referenced quickly together into a single page are described in the articles by Cohen[14] and Bobrow and Murphy[15]. According to Cohen, only a factor of three in speed was lost, despite the fact that access to the backing store was on average 104 slower than access to core storage. This means that each system has made it possible to run very large list processing applications with an efficiency that is surprisingly close to that attained by small applications that fit entirely within core storage.

Another study by Brawn and Gustavson[16] provides proof of the performance gains that can be made, especially when a software has a relatively small number of page frames allotted to it. An extensive series of experiments examining the impact of programming style on the performance obtained from the M44/44X experimental demand paging system is described in this research. Particularly intriguing are the tests that focus on refining a sorting system. The algorithm's first simple implementation underwent a number of modifications.

The order of page references was altered by each modification, but not the quantity. Although the modifications were easy to apply, they required careful consideration of the program's logic and how it would behave in a paging environment.

Strategies for operating systems

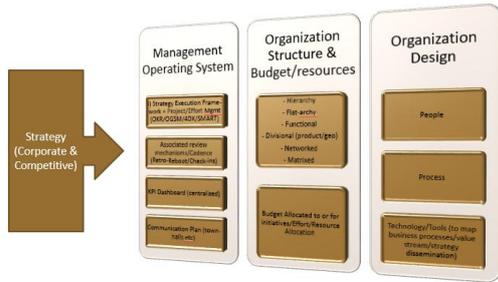


Figure 3: Strategy of OS

Simple operating system techniques will be enough if the hardware and program load are such that programs can typically have more pages in working storage than their paracher while they are running. However, generally speaking, the issue of ensuring a high probability that there is always a job in working storage available for CPU activity and avoiding executing programs in an environment with too little available space must be carefully considered in Figure 3. This is more challenging in a time-sharing setting than it is in a basic multiprogramming setting. In these situations, it is also necessary to provide frequent service to all users in order to promptly respond to basic inquiries.

It is convenient to think of scheduling, allocation, and dispatching as the operating system functions that are pertinent to this paper. In this context, scheduling is defined as the upkeep of an ordered list of the jobs vying for the allocator's services. The allocator manages how much working storage is distributed across the jobs on this list that it decides to support. Lastly, the dispatcher makes the decision if there are multiple jobs waiting for a CPU when one becomes available. A design should ideally ensure that these three mechanisms work in unison under a broad range of stress scenarios.

Unfortunately, there aren't many tried-and-true theoretical methods to goals; Denning's is the most advanced. [11]. But we're going beyond what this paper is about.

This section's remaining content discusses a number of largely distinct methods for enhancing demand paging systems' tactics and, ideally, their functionality. The methods that will be discussed can be separated into two categories: those that are mostly focused on scheduling and those that are focused on allocation.

II. SCHEDULING

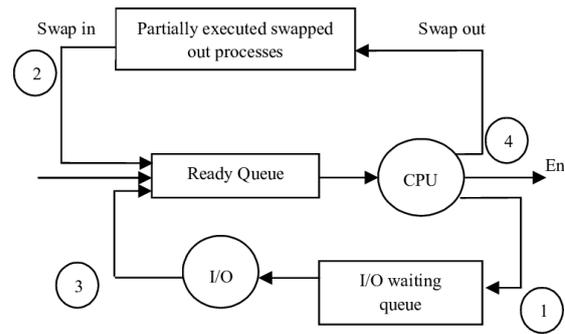


Figure 4: Process of schedulers

The scheduling function in paging systems that are having performance issues needs to be approached somewhat differently than is often done in the literature, which compares and optimizes scheduling methods separately for the system being scheduled. In these situations, scheduling is seen as the primary system that controls function, and allocation is boiled down to making room for jobs as the scheduler specifies. Performance is typically significantly degraded when such an approach is used in a system with a resource class more important than processing time. The issue is that not many "schedulers" sample the system's resources and make decisions based on the findings. On the basis of arbitrary pre-established limitations for time slice, quantum time, amount of occupied storage, etc., they instead try to inject and withdraw jobs from execution. In a demand-paging system, cycling a job into and out of execution has a significant impact on performance. An examination of Figure 4 makes it clear why this would be the case.

In these systems, the scheduling function must be subordinated to the allocation function in order to improve performance through scheduling. With very few, if any, exceptions, the scheduler should introduce and remove jobs from the system based on the condition of the system resources under the allocator's control. One should be cautious about removing a job from the system and forfeiting the investment once it has been implemented for a considerable amount of time and has accrued its paracher of pages. Instead of being the standard operating procedure, the time-slice period need to be an upper bound exception.

The aforementioned scheduling philosophy addresses the question of when and by whom objects should be added to or withdrawn from the list of jobs vying for available storage.

When the situation calls for it, the issue of which jobs should be chosen to be included to the list still

exists. Making arrangements for requests to use the same application to be batched together is a pretty easy and often suggested method of minimizing the amount of paging. This can work rather well, especially in time-sharing systems where a lot of people use a large program—like a compiler or a text editing and filing program—constantly. Reducing the amount of paging has a significant impact on system efficiency, which may even outweigh the delays caused by waiting for a batch to accumulate. In a very large system, the user may not even notice that similar requests are batched together. Limiting the programmer response rate dynamically [17] is a more authoritarian suggestion. According to the principle, even if a system should react to a programmer's request right away, it will benefit both the system and the programmer if the programmer controls how quickly the system responds to him. After a system answer, a straightforward technique would be to lock the terminal keyboard for a duration that is somehow proportionate to the system resources utilized to provide the response. This is meant to encourage users to consider before they type and to reduce the system's load. It is still to be determined whether such a plan would be acceptable or result in a quantitative improvement in performance.

III. ALLOCATION

As previously stated, the allocator is thought to have the responsibility of determining which requests for working storage the scheduler should try to fulfill and when to permit the scheduler to add new jobs into competition for working storage space. In every time interval, this allows the allocator to regulate the degree of multiprogramming (i.e., the number of independent contestants) in working storage. One can add a new, nearly independent mechanism to the original operating system to carry out this task if the changes to the current operating system modules would be too significant to include this scheduling allocation technique. Such a system, known as the a "load-leveler" was incorporated into the M44/44X system.

The load leveler's plan was to periodically check the paging rate and recent CPU utilization, and as required, ask the scheduler to temporarily remove jobs from the list of allocator-served jobs. The jobs that are temporarily put on hold are selected from those that don't see much involvement.

It can be very challenging to temporarily remove a job that has reached the end of its time slice from functioning storage. Such a task might have painstakingly amassed the number of page frames required for efficient advancement. Its initial progress will be small if it has to demand pages one at a time when it becomes eligible for a time slice again. (The simulation findings provided by Fine et al.[5] support this.) The goal of the "pre-paging" strategy (Oppenheimer and Weizer[18]) is to prevent this circumstance. Before a job is assigned to a CPU, preparing entails retrieving a set of pages into working storage that are thought to be the job's working set. Determining what to prepare (i.e., Q the makeup of the working set) is the evident challenge with pre paging. The task of creating a page replacement algorithm is quite similar to the difficulty of forecasting future program activity. [19].

Lessons to be learned

It is practical to try to condense the aforementioned conversation by taking into account the implications of the demand and paging notions independently. The concept of running software alerting the system At one extreme of the range of potential resource allocation strategies is the requirement for a certain resource after it has reached the point where it cannot advance further until that resource is supplied to it. The opposite extreme is to provide a program all of the resources it will require from the start of its existence in the system. (Of course, this necessitates full knowledge of all the resources that a program will need.) Inefficiencies occur in the former scenario when resources are used inefficiently while a program waits for its needs to be met. When resources are provided to a program for longer than necessary, they become inaccessible to other users, which leads to inefficiency in the latter scenario. The range of approaches that can be used to regulate resource deallocation is discussed in a manner similar to that of the resource allocation remarks made above. [11],[19]. The cost of inefficient resource use, the time required to meet resource demands, the likelihood of receiving precise advance notice of resource requirements, and the frequency of such demands are all important considerations when selecting allocation and deallocation strategies for a given system.

It should be mentioned that neither paging systems nor the specific resource of storage space are

specifically responsible for this issue. Experience has demonstrated, however, that in a dynamic storage allocation system, a poor strategy decision can have severe consequences. This seems to be caused by the non-linear relationship between program efficiency and space allocation, as well as the fact that storage is a very important resource in the majority of systems.

(Paging systems provide experimental support for this connection, but there is no reason to assume that dynamic allocation systems, which distribute blocks of working storage of varying sizes, are exempt from its application.)

Regarding paging, it is evident that selecting a page size is one of the most noticeable challenges that emerges during the construction of a paging system. A modest page size is preferred from the perspective of preventing the transfer of a lot of information that might not be used to functioning storage. However, the overhead resulting from extremely large page tables will be severe if a page size is selected that is too small. It should be obvious that a suitable page size can be chosen and that a very basic page replacement strategy for working storage allocation will be sufficient if the fundamental hardware and operating system strategies are sufficient for the application load.

When this isn't the case, there is an issue, and programmers must constantly be mindful of the underlying paging environment when creating or altering applications. Under these conditions, the page size becomes a straight jacket, and it would be better to provide a variety of storage area sizes, if only for the convenience of programmers. Indeed, some evidence [20] indicates that this type of environment allows for greater working storage usage than one where only a single page size is allowed. However, a definitive comparison is quite challenging because so much depends on compilation techniques and programming style.

CONCLUSION

All of the software strategies for enhancing demand paging systems' performance mentioned above make an effort to accomplish one or more of the following tasks in their own unique ways:

- (i) Give prior notice of page requirements
- (ii) Decrease the programs' working set and, thus, their parameter

(iii) make sure that programs are not overly space squeezed (ideally, each program is shown just over its parachor of pages).

(iv) clearly identify the working set;

It becomes evident that these strategies aim to eliminate both the "demand" and the "paging" aspects, rather than merely "tuning" the fundamental demand-paging theory, when one tries to abstract the common denominator of them. For instance, pre paging aims to choose a relevant section of the application to preload in order to prevent both "paging" and the "demand" page exception with the associated overhead and delay. The end result is a storage management system that combines slow swapping with rigorous paging.

Therefore, it's critical to keep in mind that demand paging is only one of several potential methods for dynamic storage allocation. Like any approach, it has a restricted range of effective applications. Demand paging's advantages include the ease of dealing with consistent units of allocation and the potential effectiveness of putting into working storage those information units that are truly referenced.

The demand paging systems that have been put into place thus far have generally demonstrated that these benefits come at the expense of quicker backup storage and/or more working storage than were previously thought to be required for the achieved throughput and response times.

Although demand paging is an effective method of managing virtual memory, how page replacement is handled has a significant impact on how well it performs. Even little strategic inefficiencies can significantly impair performance when memory resources are scarce and system loads are high.

Demand paging systems can be greatly improved by combining a well-planned program design, coordinated scheduling, and intelligent prediction (such as prepaging or working set tracking). Future research might combine these methods in hybrid tactics or use real-time program behavior monitoring for adaptive page management.

REFERENCES

- [1] Bairstow, J. N. (1968). Time-sharing. *Electronic Design*, 16(9), C1-C22.

- [2] Belady, L. A. (1966). A study of replacement algorithms for a virtual storage computer. *IBM Systems Journal*, 5(2), 78–101.
- [3] Belady, L. A., & Kuehner, C. J. (1968). *Dynamic space sharing in computer systems* (Report RC 2064). IBM Thomas J. Watson Research Center.
- [4] Bobrow, D. G., & Murphy, D. L. (1967). Structure of a LISP system using two-level storage. *Communications of the ACM*, 10(3), 155–159.
- [5] Brawn, B., & Gustavson, F. (1968). *An evaluation of program performance on the M44/44X system, Part I* (Report RC 2083). IBM T. J. Watson Research Center.
- [6] Cohen, J. A. (1967). Use of fast and slow memories in list processing languages. *Communications of the ACM*, 10(2), 82–86.
- [7] Comeau, L. W. (1967). *A study of the effect of user program optimization in a paging system*. ACM Symposium on Operating System Principles, Gatlinburg, Tennessee, October 1–4.
- [8] Dennis, J. B., & Glaser, E. L. (1965). The structure of on-line information processing systems. In *Proceedings of the 2nd Congress on Information Systems Sciences* (pp. 5–14). Spartan Books.
- [9] Denning, P. J. (1968). The working set model for program behavior. *Communications of the ACM*, 11(5), 323–333.
- [10] Fine, G. H., Jackson, C. W., & McIsaac, P. V. (1966). Dynamic program behavior under paging. In *Proceedings of the 21st ACM National Conference* (pp. 223–228). Thompson Book Co.
- [11] Kilburn, T., Edwards, D. B. G., Lanigan, M. J., & Sumner, F. H. (1962). One-level storage system. *IRE Transactions on Electronic Computers*, EC-11(2), 223–235.
- [12] Lauer, H. C. (1967). Bulk core in a 360/67 time-sharing system. In *AFIPS Conference Proceedings*, 31 (FJCC, pp. 601–609).
- [13] McKellar, A. C., & Coffman, E. G. (1968). *The organization of matrices and matrix operations in a paged multiprogramming environment* (Technical Report No. 59). Department of Electrical Engineering, Princeton University.
- [14] Nielsen, N. R. (1967). The simulation of time-sharing systems. *Communications of the ACM*, 10(7), 397–412.
- [15] Oppenheimer, G., & Weizer, N. (1968). Resource management for a medium scale time-sharing operating system. *Communications of the ACM*, 11(5), 313–322.
- [16] Patrick, R. L. (1967). Time-sharing tally sheet. *Datamation*, 13(11), 42–47.
- [17] Randell, B. (1968). *A note on storage fragmentation and program segmentation* (Report RC 2102). IBM T. J. Watson Research Center.
- [18] Randell, B., & Kuehner, C. J. (1968). Dynamic storage allocation systems. *Communications of the ACM*, 11(5), 297–306.
- [19] Shemer, J. E., & Shippey, G. A. (1966). *Statistical analysis of paged and segmented computer systems* (Technical Report). General Electric Company, Phoenix, Arizona.
- [20] Smith, J. L. (1967). Multiprogramming under a page on demand strategy. *Communications of the ACM*, 10(10), 636–646.