



HAL
open science

Hermes: Accelerating Long-Latency Load Requests via Perceptron-Based Off-Chip Load Prediction

Rahul Bera, Konstantinos Kanellopoulos, Shankar Balachandran, David Novo, Ataberk Olgun, Mohammad Sadrosadati, Onur Mutlu

► **To cite this version:**

Rahul Bera, Konstantinos Kanellopoulos, Shankar Balachandran, David Novo, Ataberk Olgun, et al.. Hermes: Accelerating Long-Latency Load Requests via Perceptron-Based Off-Chip Load Prediction. MICRO 2022 - 55th IEEE/ACM International Symposium on Microarchitecture, Oct 2022, Chicago, IL, United States. pp.1-18, 10.1109/MICRO56248.2022.00015 . lirmm-03777161

HAL Id: lirmm-03777161

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-03777161>

Submitted on 17 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Hermes: Accelerating Long-Latency Load Requests via Perceptron-Based Off-Chip Load Prediction

Rahul Bera¹ Konstantinos Kanellopoulos¹ Shankar Balachandran² David Novo³
Ataberk Olgun¹ Mohammad Sadrosadati¹ Onur Mutlu¹

¹ETH Zürich ²Intel Processor Architecture Research Lab ³LIRMM, Univ. Montpellier, CNRS

Long-latency load requests continue to limit the performance of modern high-performance processors. To increase the latency tolerance of a processor, architects have primarily relied on two key techniques: sophisticated data prefetchers and large on-chip caches. In this work, we show that: (1) even a sophisticated state-of-the-art prefetcher can only predict half of the off-chip load requests on average across a wide range of workloads, and (2) due to the increasing size and complexity of on-chip caches, a large fraction of the latency of an off-chip load request is spent accessing the on-chip cache hierarchy to solely determine that it needs to go off-chip.

The goal of this work is to accelerate off-chip load requests by removing the on-chip cache access latency from their critical path. To this end, we propose a new technique called Hermes, whose key idea is to: (1) accurately predict which load requests might go off-chip, and (2) speculatively fetch the data required by the predicted off-chip loads directly from the main memory, while also concurrently accessing the cache hierarchy for such loads.

To enable Hermes, we develop a new lightweight, perceptron-based off-chip load prediction technique that learns to identify off-chip load requests using multiple program features (e.g., sequence of program counters, byte offset of a load request). For every load request generated by the processor, the predictor observes a set of program features to predict whether or not the load would go off-chip. If the load is predicted to go off-chip, Hermes issues a speculative load request directly to the main memory controller once the load's physical address is generated. If the prediction is correct, the load eventually misses the cache hierarchy and waits for the ongoing speculative load request to finish, and thus Hermes completely hides the on-chip cache hierarchy access latency from the critical path of the correctly-predicted off-chip load. Our extensive evaluation using a wide range of workloads shows that Hermes provides consistent performance improvement on top of a state-of-the-art baseline system across a wide range of configurations with varying core count, main memory bandwidth, high-performance data prefetchers, and on-chip cache hierarchy access latencies, while incurring only modest storage overhead. The source code of Hermes is freely available at: <https://github.com/CMU-SAFARI/Hermes>.

1. Introduction

Long-latency load requests significantly limit the performance of high-performance out-of-order (OOO) processors. A load request that misses in the on-chip cache hierarchy and goes to the

off-chip main memory (i.e., an *off-chip load*) often stalls the processor core by blocking the instruction retirement from the re-order buffer (ROB), thus limiting the core's performance [88, 91, 92]. To increase the latency tolerance of a core, computer architects primarily rely on two key techniques. First, they employ increasingly sophisticated hardware prefetchers that can learn complex memory address patterns and fetch data required by future load requests before the core demands them [28, 32, 33, 35, 75]. Second, they significantly scale up the size of the on-chip cache hierarchy with each new generation of processors [10, 11, 16].

Key problem. Despite recent advances in processor core design, we observe two key trends in new processor designs that leave a significant opportunity for performance improvement on the table. First, even a sophisticated state-of-the-art prefetcher can only predict half of the long-latency off-chip load requests on average across a wide range of workloads (see §2). This is because even the most sophisticated prefetchers cannot easily learn the irregular access patterns in programs.

Second, a large fraction of the latency of an off-chip load request is spent on accessing the multi-level on-chip cache hierarchy. This is primarily due to the increasing size of the on-chip caches [15, 24, 25]. To cater to workloads with ever increasing data footprints, on-chip caches in recent processors are growing in size and complexity [30, 50, 116]. A larger on-chip cache, on the one hand, improves a core's performance by reducing the fraction of load requests that go off-chip [59, 99, 122]. On the other hand, a larger cache comes with longer cache access latency, which increases the latency of each off-chip load request [18].

Our goal in this work is to accelerate long-latency off-chip load requests by removing on-chip cache access latency from their critical path. To this end, we introduce a new technique called *Hermes*, whose **key idea** is to predict which load requests might go off-chip and start fetching their corresponding data *directly* from the main memory, while also concurrently accessing the cache hierarchy for such a load.¹ By doing so, Hermes hides the on-chip cache access latency under the shadow of the main memory access latency (as illustrated in Fig. 1), thereby significantly reducing the overall latency of an off-chip load request. Hermes works in tandem with any hardware data prefetcher and reduces the long memory access

¹Hence named after Hermes, the Olympian deity [12] who can quickly move between the realms of the divine (i.e., the processor) and the mortals (i.e., the main memory).

latency of off-chip load requests that otherwise could not have been prefetched by sophisticated state-of-the-art prefetchers.




Figure 1: Comparison of the execution timeline of an off-chip load request in a conventional processor and in Hermes.

Key challenge. Although Hermes can potentially improve performance by removing the on-chip cache access latency from the critical path of a correctly-predicted off-chip load request, its performance gain significantly depends on how accurately it can identify the off-chip load requests. This is because a false-positive off-chip prediction (i.e., a load request that is predicted to go off-chip but hits in the cache) generates an unnecessary main memory request, incurring additional main memory bandwidth and latency overheads, which can easily diminish the performance benefit gained by the load latency reduction.

We identify two key challenges in designing an accurate off-chip load prediction mechanism. First, in a system with state-of-the-art high-performance prefetchers, only 1 out of 20 load requests generated by a program on average eventually goes off-chip (see §3.2). Such a small fraction of off-chip loads makes it difficult for an off-chip load predictor to accurately and robustly learn from program behavior to produce highly-accurate predictions. Second, the accuracy of the off-chip prediction of a load can change in the presence of sophisticated prefetching techniques, making it even harder for an off-chip load predictor to learn from both the program’s and the prefetcher’s behavior.

Limitations of prior works. Several prior works [60, 84, 104, 126] propose predicting the cache level that would serve a given load request to enable various performance optimizations (e.g., better instruction scheduling). However, most of these works suffer from two key limitations that make them unsuitable for off-chip load prediction. First, prior predictors suffer from low prediction accuracy (i.e., the fraction of predicted off-chip load requests that actually went off-chip) [84, 126], which increases the bandwidth overhead in main memory. As a result, they often lose the performance benefit gained by the load latency reduction and might lower performance than the baseline system (see §8.1.1 and §8.2.2). Second, prior off-chip load prediction mechanisms often incur impractical metadata overhead (e.g., an operating-system-managed metadata storage inside the physical main memory [60], extending each TLB and cache entry with additional metadata for tracking cache residence and coherence of data [104, 105]), which hinders adoption in commercial processors.

Key mechanism. To enable Hermes, we introduce a new lightweight perceptron-based off-chip predictor, called *POPET*,

that learns to identify off-chip load requests using multiple program features (e.g., sequence of program counters, byte offset of a load request). For every load generated by the processor, POPET observes a set of program features to predict whether or not the load would go off-chip. If the load is predicted to go off-chip, Hermes issues a speculative load request (called a *Hermes request*) directly to the main memory controller once the load’s physical address is generated. This Hermes request is serviced by the main memory controller concurrently with the *regular load request* (i.e., the load issued by the processor that generated the Hermes request) that accesses the on-chip cache hierarchy. If the prediction is correct, the regular load request eventually misses the cache hierarchy and waits for the ongoing Hermes request to finish, and thus Hermes completely hides the on-chip cache hierarchy access latency from the critical path of a correctly-predicted off-chip load.

Results summary. We evaluate Hermes with a diverse set of 110 single-core and 220 multi-core workloads spanning SPEC CPU2006 [22], SPEC CPU2017 [23], PARSEC [20], Ligras [108] graph processing workloads, and commercial workloads [21]. Our evaluation yields five key results that demonstrate Hermes’s effectiveness. First, POPET achieves on average 77.1% accuracy and 74.3% coverage (i.e., the fraction of off-chip load requests of a workload that are successfully predicted), both of which are significantly higher ($1.6\times$ higher accuracy, $3.3\times$ higher coverage) than that of the prior best-accuracy off-chip predictor, HMP [126]. Second, Hermes improves performance on average by (up to) 5.4% (23.4%), 5.1% (25.7%), and 6.2% (32.2%) in single-core, eight-core, and bandwidth-constrained system configuration, on top of the best-performing state-of-the-art data prefetcher Pythia [32]. Third, Hermes consistently improves performance when combined with *any* baseline hardware data prefetcher. When implemented combined with four recently-proposed high-performance prefetchers (SPP [35, 75], Bingo [28], MLOP [106], and SMS [110]) in single-core system, Hermes improves performance on average by (up to) 5.1% (27%), 6.2% (22.4%), 7.6% (26.7%), and 7.7% (25.7%). Fourth, Hermes provides better performance-to-overhead benefit than traditional prefetchers due to its highly-accurate off-chip predictions. For every 1% performance increase, Hermes increases the main memory requests by only 0.5%, whereas Pythia increases them by 2%. Fifth, all of Hermes’s benefits come at a very modest storage overhead of only 4 KB per core, while the state-of-the-art prefetcher Pythia consumes 25.5 KB per core.

We make the following contributions in this paper:

- We identify two key opportunities for performance improvement in modern processors: (1) a significant fraction of the load requests continues to go off-chip even in the presence of sophisticated data prefetchers, and (2) an increasing fraction of the off-chip load latency is spent accessing on-chip caches due to the increasing size of the on-chip cache hierarchy.
- We introduce *Hermes*, a new technique that reduces long memory access latency by predicting off-chip load requests and fetching their corresponding data *directly* from the main

memory, while concurrently accessing the on-chip cache hierarchy for such loads.

- We design a new perceptron-based off-chip load predictor, called POPET, that accurately identifies and predicts the off-chip load requests using multiple program features.
- We show that Hermes significantly improves performance across a wide range of workloads and system configurations with varying core count, main memory bandwidth, high-performance data prefetchers, and on-chip cache access latencies.
- We open-source Hermes and all necessary traces and scripts to reproduce results in <https://github.com/CMU-SAFARI/Hermes>.

2. Motivation

High main memory access latency continues to limit the performance of modern out-of-order (OOO) processors. A load request that misses the on-chip cache hierarchy and goes to off-chip main memory often blocks instruction retirement from the reorder buffer (ROB), preventing the processor from allocating new instructions into the ROB [51, 88, 91, 92], limiting performance.

To tolerate long memory latency, recent high-performance OOO cores have primarily relied on two key techniques. First, modern cores have significantly scaled up their on-chip cache size (e.g., each Intel Alder Lake core [10] employs 4.3MB on-chip cache (including L1, L2 and a per-core last-level cache (LLC) slice), which is $1.88\times$ larger than the on-chip cache in the previous-generation Skylake core [4]). Second, modern cores employ increasingly sophisticated hardware prefetchers [11, 16] that can more effectively predict the addresses of load requests in advance and fetch their corresponding data to on-chip caches before the program demands it, thereby completely or partially hiding the long off-chip load latency for a fraction of off-chip loads [11, 32, 33, 35].

Despite these advances, we observe two key trends in processor design that leave a significant performance improvement opportunity on the table: (1) a large fraction of load requests continues to go off-chip even in the presence of state-of-the-art prefetchers, and (2) an increasing fraction of the latency of an off-chip load request is spent accessing the increasingly larger on-chip caches.

A large fraction of loads is still uncovered by state-of-the-art prefetchers. Over the past decades, researchers have proposed many hardware prefetching techniques that have consistently pushed the limits of performance improvement (e.g., [26, 32, 33, 35, 37, 45, 46, 56, 70, 75-77, 79, 85, 94, 98, 106, 107, 110, 111]). We observe that state-of-the-art prefetchers provide a large performance gain by accurately predicting future load addresses. Yet, a large fraction of off-chip load requests cannot be predicted even by the most advanced prefetchers. These uncovered requests limit the processor’s performance by blocking instruction retirement in the ROB. Fig. 2 shows a stacked graph of total number off-chip load requests in a no-prefetching system and a system with the recently-

proposed hardware data prefetcher Pythia [32], normalized to the no-prefetching system, across 110 workload traces categorized into five workload categories.² Each bar further categorizes load requests into two classes: loads that block instruction retirement from the ROB (called *blocking*) and loads that do not (called *non-blocking*). §7 discusses our evaluation methodology.




Figure 2: The distribution of ROB-blocking and non-blocking load requests (on the left y-axis), and LLC misses per kilo instructions (on the right y-axis) in the absence and presence of a state-of-the-art hardware data prefetcher [32].

We make two key observations from Fig. 2. First, on average, Pythia accurately prefetches nearly half of all off-chip load requests in the no-prefetching system, thereby improving the overall performance (not shown here; see §8.2.1). Second, the remaining half of the off-chip loads are not prefetched even by a sophisticated prefetcher like Pythia. 71.4% of these non-prefetched off-chip loads block instruction retirement from the ROB, significantly limiting performance. We conclude that, state-of-the-art prefetchers, while effective at improving performance, still leave a significant performance improvement opportunity on the table.

An increasing fraction of off-chip load latency is spent accessing the on-chip cache hierarchy. We observe that the on-chip cache hierarchy has not only grown tremendously in size but also in design complexity (e.g., sliced last-level cache organization [30, 50, 74]) in recent processors, in order to cater to workloads with large data footprints. A larger on-chip cache hierarchy, on the one hand, improves a core’s performance by preventing more load requests from going off-chip. On the other hand, all on-chip caches need to be accessed to determine if a load request should be sent off-chip. As a result, on-chip cache access latency significantly contributes to the total latency of an off-chip load. With increasing on-chip cache sizes, and the complexity of the cache hierarchy design and the on-chip network [34, 116], the on-chip cache access latency is increasing in processors [15, 18]. An analysis of the Intel Alder Lake core suggests that the load-to-use latency of an LLC access has increased to 14 ns (which is equivalent to 55 cycles for a core running at 4 GHz) [15, 24, 25].

To demonstrate the effect of long on-chip cache access la-

²We select Pythia as the baseline prefetcher as it provides the highest prefetch coverage and performance benefit among the five contemporary prefetchers considered in this paper (see §7.2 and §8.4.2). Nonetheless, our qualitative observation holds equally true for other prefetchers considered in this work (see §7.2).

tency on the total latency of an off-chip load, Fig. 3 plots the average number of cycles a core stalls due to an off-chip load blocking any instruction from retiring from the ROB, averaged across each workload category in our baseline system with Pythia. Each bar further shows the average number of cycles an off-chip load spends for accessing the on-chip cache hierarchy. Our simulation configuration faithfully models an Intel Alder Lake performance-core with a large ROB, large on-chip caches and publicly-reported cache access latencies (see §7). As Fig. 3 shows, an off-chip load stalls the core for an average of 147.1 cycles. 40.1% of these stall cycles (i.e., 58.9 cycles) can be completely *eliminated* by removing the on-chip cache access latency from the off-chip load’s critical path. We conclude that a large and complex on-chip cache hierarchy is directly responsible for a large fraction of the overall stall cycles caused by an off-chip load request. We envision that this problem will only get exacerbated with new processor designs as on-chip caches continue to grow in size and complexity [18].




Figure 3: The average number of cycles a core stalls due to an off-chip load blocking any instruction from retiring from the ROB across all workload categories. The dark portion in each bar shows the cycles that can be completely eliminated by removing the on-chip cache access latency from an off-chip load’s critical path.

3. Our Goal and Key Idea

Our goal is to improve processor performance by removing the on-chip cache access latency from the critical path of off-chip load requests.

3.1. The Key Idea and Potential Benefits

To this end, we propose a new technique called *Hermes*, whose *key idea* is to predict which load requests might go off-chip and start fetching their corresponding data *directly* from the main memory, while also concurrently accessing the cache hierarchy for such a load.

To understand the potential performance benefits of Hermes, we model an *Ideal Hermes* system in simulation where we reduce the main memory access latency of *every* off-chip load request by the post-L1 on-chip cache hierarchy access latency (which includes L2 and LLC access, and interconnect latency). In other words, in the Ideal Hermes system, we (1) magically and perfectly know if a load request would go off-chip after its physical address is available (i.e., after the translation lookaside buffer access, which happens in parallel with the L1 data cache access in modern processors [29, 36, 95, 121]), and (2) directly

access the off-chip main memory for such a load, eliminating the non-L1-cache related on-chip cache hierarchy access latency from such a load’s total latency. Fig. 4(a) shows the speedup of Ideal Hermes by itself and when combined with Pythia normalized to the no-prefetching system in single-core workloads. We make two key observations from Fig. 4(a). First, Ideal Hermes combined with Pythia outperforms Pythia alone by 8.3% on average across all workloads. Second, Ideal Hermes by itself provides nearly 80% of the performance improvement that Pythia provides. Fig. 4(b) shows the speedup of Ideal Hermes when combined with four other recently-proposed high-performance prefetchers: Bingo [28], SPP [75] (with perceptron filter [35]), MLOP [106], and SMS [110]. Ideal Hermes improves performance by 9.4%, 8.2%, 10.9%, and 13.3% on top of four state-of-the-art prefetchers Bingo, SPP, MLOP, and SMS, respectively. Based on these results, we conclude that Hermes has high potential performance benefit not only when implemented alone but also when combined with a wide variety of high-performance prefetchers.




Figure 4: (a) Speedup of Ideal Hermes by itself and when combined with Pythia in single-core workloads. (b) Speedup of Ideal Hermes when combined with four recently-proposed prefetchers: Bingo [28], SPP [35, 75], MLOP [106], and SMS [110].

3.2. Key Challenge

Even though Hermes has a significant potential to improve performance, Hermes’s performance gain heavily depends on the accuracy (i.e., the fraction of predicted off-chip loads that actually go off-chip) and the coverage (i.e., the fraction of off-chip loads that are successfully predicted) of the off-chip load prediction. A low-accuracy off-chip load predictor generates useless main memory requests, which incur both latency and bandwidth overheads, and causes interference to the useful requests in the main memory. A low-coverage predictor loses opportunity to improve performance.

We identify two key challenges in designing an off-chip load predictor with high accuracy and high coverage. First, only a small fraction of the total loads generated by a workload goes off-chip in presence of a sophisticated data prefetcher. As shown in Figure 5, on average 7.9 loads per kilo instructions miss the LLC and go off-chip in our baseline system with Pythia.

However, these loads constitute only 5.1% of the total loads generated by a workload. This small fraction of off-chip loads makes it difficult for an off-chip load predictor to accurately learn from the workload behavior to produce highly-accurate predictions.




Figure 5: Percentage of loads that miss the LLC and goes off-chip (on the left y-axis) and the LLC MPKI (on the right y-axis) in the baseline system with Pythia.

Second, the off-chip predictability of a workload can change in the presence of modern sophisticated data prefetchers. This is because in the presence of a sophisticated prefetcher, the likelihood of a load request going off-chip not only depends on the program behavior but also on the prefetcher’s ability to successfully prefetch for the load.

In this work, we overcome these two key challenges by designing a new off-chip load prediction technique, called POPET, based on perceptron learning [64, 82, 103]. By learning to identify off-chip loads using multiple program features (e.g., sequence of program counters, byte offset of a load request, page number of the load address), POPET provides both higher accuracy and coverage than a prior cache hit-miss prediction technique [126] and higher accuracy than another off-chip load prediction technique that we develop (see §7.2), in the presence of modern sophisticated prefetchers, without requiring large metadata storage overhead. With small changes to the existing on-chip datapath design, we demonstrate that Hermes with POPET significantly outperforms the baseline system with a state-of-the-art prefetcher across a wide range of workloads and system configurations.

4. Key Related Works

The key idea of load hit-miss prediction (HMP) was proposed in [126] and demonstrated as a method to improve the instruction scheduler efficiency and performance. By predicting which load instructions will miss the L1 data cache,³ HMP enables the instruction scheduler to delay the scheduling of dependent instructions until the data is fetched. This scheduler optimization improves processor performance by scheduling a load-dependent instruction to execute at the time when the data is available. Even though Hermes leverages off-chip load prediction in a very different way than HMP, we compare the accuracy and coverage of our perceptron-based off-chip load predictor POPET, against those of HMP in §8.1.1 and show that Hermes with POPET greatly outperforms Hermes with HMP in §8.2.2.

³Even though HMP was originally proposed to predict loads that miss L1 data cache, it can be extended to predict loads that miss the entire multi-level on-chip cache hierarchy.

Cache-level hit/miss prediction (i.e., predicting which cache level a load might hit) has also been explored in three works: Direct-to-Data cache (D2D [104]), Direct-to-Master cache (D2M [105]), and Level Prediction (LP [60]). All three works employ different mechanisms to track cacheline addresses present in the cache hierarchy along with the cache-level(s) a cacheline is present in. For a given load, these works predict which cache-level the load would likely hit. If the cache-level hit/miss prediction is correct, the processor fetches the data of the correctly-predicted load by only accessing the predicted memory level (L1, L2, LLC, or the main memory) and bypassing all other memory levels. By doing so, a cache-level hit/miss predictor reduces the latency of a correctly-predicted load and improves the processor’s energy efficiency. However, if the predictor incorrectly bypasses a memory level that has more up to date data (e.g., if the data is present in L2, but the predictor suggests fetching the data from LLC), the processor needs to detect such *mispredictions* and access the correct memory level to maintain correct execution, which comes with performance overhead and additional complexity. Hermes differs from all these prior works in three major ways:

Prediction via tracking addresses vs. learning from program context. To make accurate cache-level prediction, D2D, D2M, and LP rely on tracking cacheline addresses present in the on-chip caches by using efficiently-managed metadata structures. For example, D2D and D2M extend each translation look-aside buffer (TLB) entry (called *eTLB*) to keep additional cache-level-related metadata. LP manages the cache-level metadata as an in-memory table and caches the metadata on-chip using a metadata prefetching mechanism. These metadata structures need to be updated for almost all cache operations (e.g., cache insertions, evictions) in order to faithfully track the cache contents and to provide accurate cache-level predictions. In contrast, POPET is built on the insight that the correlation between different types of program context information and off-chip loads can be accurately *learned without tracking cache contents*. As a result, POPET does *not* explicitly track addresses present in the cache hierarchy, but learns to predict off-chip loads by aggregating various program context information.

Lower complexity and hardware overhead. To enable accurate cache-level hit/miss prediction and to recover from a misprediction, prior works need relatively intrusive changes to multiple components on the on-chip datapath. Both D2D and D2M cache designs require extending the TLB and operating system support to store the way and cache-level information for every cacheline within the reach of *eTLB*. LP requires 2-bit metadata for every 64B chunk of *the physical main memory*. LP manages this metadata as an in-memory table in the system-reserved physical main memory space [60], which necessitates support from the operating system. In the worst case, to make a cache-level hit/miss prediction, LP needs to fetch the metadata from main memory, which can take longer than the cache lookup LP is designed to avoid. In contrast, POPET incurs a very modest (4 KB) total storage overhead and small changes to the existing datapath.

No misprediction detection and recovery. LP requires cache level misprediction detection and recovery to prevent the program from using any stale, incoherent data from the memory subsystem. On the other hand, Hermes *never* brings any data to the on-chip cache hierarchy unless the data is required by a LLC miss request (see §6.2.2). Hence, Hermes *does not* require any misprediction detection and recovery mechanism to maintain correct execution.

We design an address tag-tracking based off-chip load predictor (called *TTP*; see §7.2) inspired by these prior works [60, 81, 104] and evaluate it against POPET (see §8). TTP tracks address tags present in the *entire* cache hierarchy in its metadata structure (see §7.2) and predicts that a load would go off-chip if the tag of the load address is not present in its metadata structure. We open-source the implementation of TTP in our repository [13]. Our results show that POPET provides both higher accuracy and higher performance than TTP (see §8.1.1 and §8.2.1).

5. Hermes: Design Overview

Fig. 6 shows a high-level overview of Hermes. POPET is the key component of Hermes that is responsible for making highly-accurate off-chip load predictions. For every demand load request generated by the processor, POPET predicts whether or not the load request would go off-chip (1). If the load is predicted to go off-chip, Hermes issues a speculative memory request (called a *Hermes request*) *directly* to the main memory controller once the load’s physical address is generated to start fetching the corresponding data from the main memory (2). This Hermes request is serviced by the main memory controller concurrently with the *regular load request* (i.e., the load issued by the processor that generated the Hermes request) that accesses the on-chip cache hierarchy. If the prediction is correct, the regular load request to the same address eventually misses the LLC and waits for the ongoing Hermes request to finish, thereby completely hiding the on-chip cache access latency from the critical path of the correctly-predicted off-chip load (3). If a Hermes request returns from the main memory but there has been no regular load request to the same address, Hermes drops the request and does not fill the data into the cache hierarchy. By doing so, Hermes keeps the on-chip cache hierarchy fully coherent even in case of a misprediction. For every regular load request returning to the core, Hermes trains POPET based on whether or not this load has actually gone off-chip (4).

6. Hermes: Detailed Design

We first describe the design of POPET in §6.1, followed by the changes introduced by Hermes to the on-chip cache access datapath in §6.2.

6.1. POPET Design

The purpose of POPET is to accurately predict whether or not a load request generated by the processor will go off-chip. We design POPET using the multi-feature perceptron learning mechanism [35, 47, 61, 62, 64, 65, 103, 113].

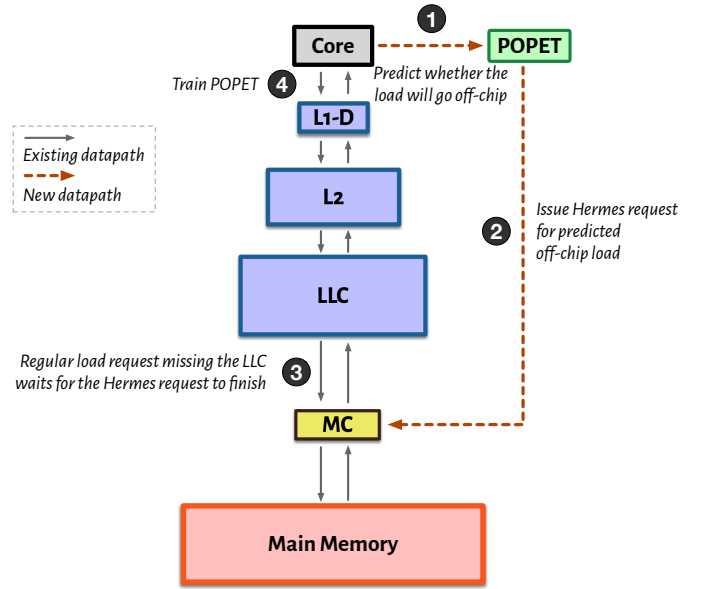


Figure 6: Overview of Hermes

What is perceptron learning? Perceptron learning, whose roots are in [82] and was demonstrated by Rosenblatt [103], is a simplified learning model to mimic biological neurons. Fig. 7 shows a *single-layer perceptron* network where each *input* is connected to the *output* via an *artificial neuron*. Each artificial neuron is represented by a numeric value, called *weight*. The perceptron network as a whole iteratively learns a binary classification function $f(x)$ (shown in Eq. 1), a function that maps the input X (a vector of n values) to a binary output.

$$f(x) = \begin{cases} 1 & \text{if } w_0 + \sum_{i=1}^n w_i x_i > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The perceptron learning algorithm starts by initializing the weight of each neuron and iteratively trains the weights using each input vector from the training dataset in two steps. First, for an input vector X , the perceptron network computes a binary output using Eq. 1 and the current weight values of its neurons. Second, if the computed output differs from the desired output for that input vector provided by the dataset, the weight of each neuron is updated [103]. This iterative process is repeated until the error between the computed and desired output falls below a user-specified threshold.




Figure 7: Overview of a single-layer perceptron model. Each blue circle denotes an input and the green circle denotes the output of the perceptron.

Jimenez et al. [64] have applied the perceptron learning algorithm to design a lightweight, high-performance branch predictor in a processor core. Today, multiple commercial

processors also use perceptron learning for making various microarchitectural predictions (e.g., Samsung Exynos [49], and AMD Ryzen [5]).

We design POPET using an existing microarchitectural perceptron model, known as the *hashed perceptron* [112]. A hashed perceptron model hashes multiple feature values to retrieve weights of each feature from small tables. If the sum of these weights exceeds a threshold, the model makes a positive prediction. Hashed perceptron, as compared to other perceptron models, is lightweight and easy to implement in hardware. Prior works successfully apply hashed perceptron for various microarchitectural predictions, e.g., branch outcome [47, 62, 64], LLC reuse [65, 113], prefetch usefulness [35]. This is the first work that applies hashed perceptron to off-chip load prediction.

Why perceptron? We choose to design POPET based on perceptron learning for two key reasons. First, by learning using multiple program features, perceptron learning can provide highly accurate predictions that could not be otherwise provided by simple history-based learning prediction (e.g., HMP [126], described in §4). Second, perceptron learning can be implemented with low storage overhead, without requiring any impractical metadata support (e.g., extending TLB [104, 105] or in-memory metadata storage [60], described in §4).

POPET overview. POPET is organized as a collection of one-dimensional tables (each called a *weight table*), where each table corresponds to a single program feature. Each table entry stores a *weight* value, implemented using a 5-bit saturating signed integer, that represents the correlation between the corresponding program feature value and the true outcome (i.e., whether a given load actually went off-chip). A weight value saturated near the maximum (i.e., +15) or the minimum (i.e., -16) value represents a strong positive or negative correlation between the program feature value and the true outcome, respectively. A weight value closer to zero signifies a weak correlation. The weights are adjusted during training (step 4 in Fig. 6) to update POPET’s prediction with the true outcome. Each weight table is sized differently based on its corresponding program feature (see Table 3).

6.1.1. Making a Prediction. During load queue (LQ) allocation for a load generated by the core (step 1 in Fig. 6), POPET makes a binary prediction on whether or not the load request would go off-chip. The prediction happens in three stages as shown in Fig. 8. In the first stage, POPET extracts a set of program features from the current load request and a history of prior requests (§6.1.3 shows the list of program features used by POPET). In the second stage, each feature value is hashed and used as an index to retrieve a weight value from the weight table of the corresponding feature. In the third stage, all weight values from individual features are accumulated to generate the *cumulative perceptron weight* (W_σ). If W_σ exceeds a predefined threshold (called the *activation threshold*, τ_{act}), POPET makes a positive prediction (i.e., it predicts that the current load request would go off-chip). Otherwise, POPET makes a negative prediction. The hashed feature values, the cumulative

perceptron weight W_σ , and the predicted outcome are stored in the LQ entry to be reused to train POPET when the load request returns to the processor core (step 4 in Fig. 6).

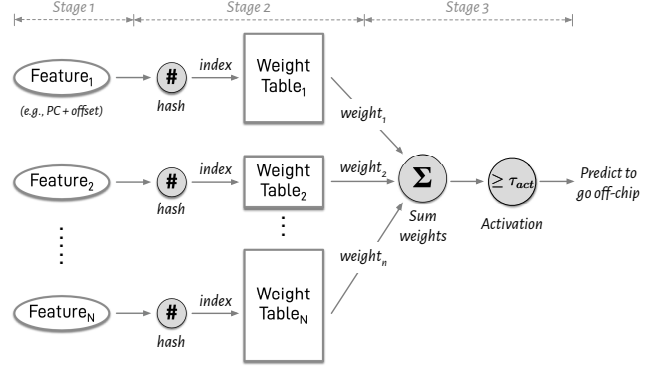


Figure 8: Stages to make a prediction by POPET

6.1.2. Training the Predictor. POPET training is invoked when a demand load request returns to the core and prepares to release its corresponding LQ entry (step 4 in Fig. 6). Every demand load that misses the LLC and goes to the main memory controller is marked as a *true* off-chip load request. This true off-chip outcome, along with the predicted outcome stored in the LQ entry of the demand load, are used to appropriately train the feature weights of POPET. The training happens in two stages. In the first stage, the W_σ (computed during prediction) is retrieved from the LQ entry. If W_σ is neither positively nor negatively saturated (i.e., W_σ lies within a negative and a positive training threshold, T_N and T_P , respectively), the weight training is triggered. This saturation check prevents the individual feature weight values from getting over-saturated, thereby helping POPET to quickly adapt its learning to program phase changes. In the second stage, if the weight training is triggered, the weights for each individual program feature are retrieved from their corresponding weight table using the hashed feature indices stored in the LQ entry. If the true outcome is positive (meaning the load actually went off-chip), the weight value for each feature is incremented by one. If the true outcome is negative, the weight values are decremented by one. This simple weight update mechanism moves each individual feature weight towards the direction of the true outcome, thus gradually increasing the prediction accuracy.

6.1.3. Automated Feature Selection. The selection of the program features used to make the off-chip load prediction is critical to POPET’s performance. A carefully-crafted and selected set of features can significantly improve the accuracy and the coverage of POPET. In this section we propose an automated, offline, performance-driven methodology to find a set of program features for POPET.

We initially select a set of 16 individual program features using our domain expertise that can correlate well with a load going off-chip. Table 1 shows the initial feature set.

The automated feature selection process happens offline during the design time of POPET. The process starts with the initial set of 16 individual program features and iteratively creates a

Table 1: The initial set of program features used for automated feature selection. \oplus represents a bitwise XOR operation.

Features without control-flow information	Features with control-flow information
1. Load virtual address	8. Load PC
2. Virtual page number	9. PC \oplus load virtual address
3. Cacheline offset in page	10. PC \oplus virtual page number
4. First access	11. PC \oplus cacheline offset
5. Cacheline offset + first access	12. PC + first access
6. Byte offset in cacheline	13. PC \oplus byte offset
7. Word offset in cacheline	14. PC \oplus word offset
	15. Last-4 load PCs
	16. Last-4 PCs

list of *feature sets*, each containing n features, at every iteration n in the following way. In the first iteration, we design POPET with each of the 16 initial program features and test its prediction accuracy in 10 randomly-selected workload traces (called *testing workloads*). We select the top-10 features that produce the highest prediction accuracy for the second iteration. In the second iteration, we create 160 two-combination feature sets (meaning, each feature set contains two initial features from Table 1) by combining each of the 16 initial features with each of the 10 winning feature sets from the last iteration, and test the prediction accuracy on the testing workloads. We select the top-10 two-combination feature sets that produce the highest prediction accuracy for the third iteration. This iterative process repeats until the maximum prediction accuracy gets saturated (i.e., the difference in accuracy of two successive iterations is less than 3%).⁴ Table 2 shows the final list of program features selected by the automated feature selection process.

Table 2: POPET configuration parameters

<i>Selected features</i>	<ul style="list-style-type: none"> • PC \oplus cacheline offset • PC \oplus byte offset • PC + first access • Cacheline offset + first access • Last-4 load PCs
<i>Threshold values</i>	$\tau_{act} = -18, T_N = -35, T_P = 40$

Rationale for selected features. Each selected feature correlates with the likelihood of observing an off-chip load request with a different program context information. We explain the rationale for each selected feature below.

(1) PC \oplus cacheline offset. This feature is computed by XOR-ing the load PC value with the cacheline offset of the load address in the virtual page of the load request. The goal of this feature is to learn the likelihood of a load request going off-chip when a given load PC touches a certain cacheline offset in a virtual page. The use of cacheline offset information, instead of load virtual address or virtual page number, enables this feature to apply the learning across different virtual pages.

(2) PC \oplus cacheline byte offset. This feature is computed by XOR-ing the load PC with the byte offset of the load cacheline address. This feature is particularly useful in accurately

predicting off-chip load requests when a program has a streaming access pattern over a linearly allocated data structure. For example, when a program streams through a large array of 4B integers, every 16th load (as a 64B cacheline stores 16 integers) generated by a load PC that is iterating over the array will go off-chip, and the remaining loads will hit in on-chip caches. In this case, this feature learns to identify only those loads that have a byte offset of 0 to go off-chip.

(3) PC + first access. This feature is computed by left-shifting the load PC and adding the *first access* hint at the most-significant bit position. The first access hint is a binary value that represents whether or not a cacheline has been recently touched by the program. The hint is computed using a small 64-entry buffer (called the *page buffer*) that tracks the demanded cachelines from last 64 virtual pages. Each page buffer entry holds two pieces of information: a virtual page tag, and a 64-bit bitmap, where each bit represents one cacheline in the virtual page. During every load request generation, POPET searches the page buffer with the virtual page number of the load address. If a matching entry is found, POPET uses the value of the bit corresponding to the cacheline offset in the matching page buffer entry’s bitmap as the first access hint. If the bit is set (or unset), it signifies that the corresponding cacheline has (not) been recently accessed by the program. If the bit is unset, POPET sets the bit in the page buffer entry’s bitmap. The first access hint provides a crude estimate of a cacheline’s reuse in a short temporal window. However, it alone cannot determine the cacheline’s residency in on-chip caches, as the memory footprint tracked by the page buffer is much smaller than the total cache size.

(4) Cacheline offset + first access. This feature is similar to the PC + first access feature, except that it learns the likelihood of a load request going off-chip when a given cacheline offset is recently touched by the program.

(5) Last-4 load PCs. This feature value is computed as a shifted-XOR of last four load PCs. It represents the execution path of a program and correlates it with the likelihood of observing an off-chip load request whenever the program follows the same execution path.

6.1.4. Parameter Threshold Tuning. POPET has three tunable parameters: negative and positive training thresholds (T_N and T_P , respectively), and the activation threshold (τ_{act}). Properly tuning the values of all these three parameters is also critical to POPET’s performance, since both POPET’s accuracy and coverage are sensitive to parameter values.

We employ a three-step grid search technique to tune each of the three parameters separately. In the first stage, we uniformly sample values from a parameter’s range. For example, τ_{act} can take values in the range $[-80, 75]$.⁵ We uniformly sample values from this range with a grid size of 5. In the second stage, we run Hermes with the randomly-selected 10 test workloads (as mentioned in §6.1.3) for each of the sampled values and

⁴For simplicity, our automated feature selection process optimizes for accuracy. A more comprehensive feature selection process can also include coverage or directly optimize for performance (i.e., execution time).

⁵As POPET uses five program features (see §6.1.3), the sum of all five weights (each represented by a 5-bit saturating signed integer as described in §6.1) can take a maximum and minimum value of 75 and -80 , respectively.

pick the top-10 values that provide the highest performance gain. In the third stage, we run Hermes with all single-core workload traces using the selected 10 parameter values from the second stage. We finally select the value that provides the highest average performance gain. Table 2 shows the selected threshold values of each parameter.

6.2. Hermes Datapath Design

In this section, we describe the key changes introduced to the existing well-optimized on-chip cache access datapath to incorporate Hermes. First, we show how the core issues a Hermes request directly to the main memory controller if POPET predicts the load would go off-chip and how a regular load request that misses the LLC waits for an ongoing Hermes request (see §6.2.1). Second, we discuss how the data fetched from main memory is properly sent back to the core in presence of Hermes while maintaining cache coherence (see §6.2.2).

6.2.1. Issuing a Hermes Request. For every load request predicted to go off-chip, Hermes issues a Hermes request directly to the main memory controller (step ② in Fig. 6) once the load’s physical address is generated. The main memory controller enqueues the Hermes request in its read queue (RQ) and starts fetching the corresponding data from the main memory as dictated by its scheduling policy, while the regular load request is concurrently accessing the on-chip cache hierarchy. If the off-chip prediction is correct, the regular load request eventually misses the LLC and checks the main memory controller’s RQ for any ongoing main memory access to the same load address (step ③). If the address is found, the regular load request waits for the ongoing Hermes request to finish before sending the Hermes-fetched data back to the core.

Hermes’s performance gain depends on the latency to directly issue a Hermes request to the main memory controller (called *Hermes request issue latency*). Although a Hermes request experiences a significantly shorter latency to arrive at the main memory controller than its corresponding regular load request because a Hermes request bypasses the cache hierarchy and on-chip queuing delays, a Hermes request nonetheless pays for a latency to route through the on-chip network. We model two variants of Hermes using an optimistic and a pessimistic estimate of Hermes request issue latency to take into account a wide range of potential differences in on-chip interconnect designs (see §7.2). In §8.4.3, we also evaluate Hermes with a wide range of Hermes request issue latencies (from 0 cycle to 24 cycles) and show that Hermes consistently provides performance benefit even with the most pessimistic Hermes request issue latency.

6.2.2. Returning Data to the Core. For every Hermes request returning from main memory, Hermes checks the RQ of the main memory controller and returns the fetched data back to the LLC if there is a regular load request already waiting for the same load address. If there is no regular load request waiting for the completed Hermes request, Hermes drops the request and does *not* fill the data into the cache hierarchy, which keeps the on-chip cache hierarchy internally coherent.

6.3. Storage Overhead

Table 3 shows the total storage overhead of Hermes. Hermes requires only 4 KB of metadata storage per processor core. POPET consumes 3.2 KB, whereas the metadata stored in LQ for POPET training consumes 0.8 KB.

Table 3: Storage overhead of Hermes

Structure	Description	Size
POPET	<ul style="list-style-type: none"> • Perceptron weight tables <ul style="list-style-type: none"> – PC ⊕ cacheline offset: $1024 \times 5b$ – PC ⊕ byte offset: $1024 \times 5b$ – PC + first access: $1024 \times 5b$ – Cacheline offset + first access: $128 \times 5b$ – Last-4 load PCs: $1024 \times 5b$ • Page buffer: $64 \times 80b$ 	3.2 KB
LQ Metadata	<ul style="list-style-type: none"> Hashed PC: $128 \times 32b$; Last-4 PC: $128 \times 10b$; First access: $128 \times 1b$; perceptron weight: $128 \times 5b$; prediction: $128 \times 1b$ 	0.8 KB
Total		4.0 KB

7. Methodology

We use the ChampSim trace-driven simulator [9] to evaluate Hermes. We faithfully model the latest-generation Intel Alder Lake performance-core [11] with its large ROB, large caches with publicly-reported on-chip cache access latencies [18, 24, 25], and the state-of-the-art prefetcher Pythia [32] at the LLC. Table 4 shows the key microarchitectural parameters. For single-core simulations, we warm up the core using 100M instructions and simulate the next 500M instructions. For multi-programmed simulations, we use 50M and 100M instructions from each workload for warmup and simulation, respectively. If a core finishes early, the workload is replayed until every core has finished executing at least 100M instructions. The source code of Hermes, along with all workload traces and scripts to reproduce our results are freely available at [13].

Table 4: Simulated system parameters

Core	1 and 8 cores, 6-wide fetch/execute/commit, 512-entry ROB, 128/72-entry LQ/SQ, Perceptron branch predictor [61] with 17-cycle misprediction penalty
L1/L2 Caches	Private, 48KB/1.25MB, 64B line, 12/20-way, 16/48 MSHRs, LRU, 5/15-cycle round-trip latency [25]
LLC	3MB/core, 64B line, 12 way, 64 MSHRs/slice, SHiP [122], 55-cycle round-trip latency [24, 25], Pythia prefetcher [32]
Main Memory	1C: 1 channel, 1 rank per channel; 8C: 4 channels, 2 ranks per channel; 8 banks per rank, DDR4-3200 MTPS, 64b data-bus per channel, 2KB row buffer per bank, tRCD=12.5ns, tRP=12.5ns, tCAS=12.5ns
Hermes	Hermes-O/P: 6/18-cycle Hermes request issue latency

7.1. Workloads

We evaluate Hermes using a wide range of memory-intensive workloads spanning SPEC CPU2006 [22], SPEC CPU2017 [23], PARSEC [20], Ligma graph processing workload suite [108], and commercial workloads from the 2nd data value prediction championship (CVP [21]). For SPEC CPU2006 and SPEC CPU2017 workloads, we reuse the instruction traces provided

by the 2nd and the 3rd data prefetching championships (DPC [2, 3]). For PARSEC and Ligra workloads, we reuse the instruction traces open-sourced by Pythia [32]. The CVP workload traces are collected by the Qualcomm Datacenter Technologies and capture complex program behavior from various integer, floating-point, cryptographic, and server applications in the field. We only consider workload traces in our evaluation that have at least 3 LLC misses per kilo instructions (MPKI) in the no-prefetching system. In total, we evaluate Hermes using 110 single-core workload traces from 73 workloads, which are summarized in Table 5. All these traces can be freely downloaded using a script as mentioned in Appendix A.5. For multi-programmed simulations, we create both homogeneous and heterogeneous trace mixes. For an eight-core homogeneous multi-programmed simulation, we run eight copies of each trace from our single-core trace list, one trace in each core. For heterogeneous multi-programmed simulation, we *randomly* select any eight traces from our single-core trace list and run one trace in each core. In total, we evaluate Hermes using 110 homogeneous and 110 heterogeneous eight-core workloads.

Table 5: Workloads used for evaluation

Suite	#Workloads	#Traces	Example Workloads
SPEC06	14	22	gcc, mcf, cactusADM, lbm, ...
SPEC17	11	23	gcc, mcf, pop2, fotonik3d, ...
PARSEC	4	12	canneal, facesim, raytrace, ...
Ligra	11	20	BFS, PageRank, Radix, ...
CVP	33	33	integer, floating-point, server, ...

7.2. Evaluated System Configurations

For a comprehensive analysis, we compare Hermes with various off-chip load prediction mechanisms, as well as in combination with various recently proposed prefetchers. Table 6 compares the storage overhead of all evaluated mechanisms.

(1) *Various off-chip prediction mechanisms.* We compare POPET against two cache hit/miss prediction techniques: (1) HMP, proposed by Yoaz et al. [126], and (2) a simple cacheline tag-tracking based predictor, called TTP, which we design. HMP uses three predictors similar to a hybrid branch predictor: local [125], gshare [83, 125], and gskew [86], each of which individually predicts off-chip loads using a different prediction mechanism. For a given load, HMP consults each individual predictor and selects the majority prediction. We design TTP by taking inspiration from prior cacheline address tracking-based mechanisms [60, 81, 104] (see §4). TTP tracks partial tags of cacheline addresses that are likely to be present in the entire on-chip cache hierarchy in a separate metadata structure. For every cache fill (LLC eviction), the partial tag of the filled (evicted) cacheline address is inserted into (evicted from) TTP’s metadata. To predict whether or not a given load would go off-chip, TTP searches the metadata structure with the partial tag of the load address. If the tag is not present in the metadata structure, TTP predicts the load would go off-chip. We open-source TTP in our repository [13].

(2) *Various data prefetchers.* We evaluate Hermes combined with five recently-proposed high-performance prefetch-

ing techniques: Pythia [32], Bingo [28], SPP [75] (with perceptron filter [35]), MLOP [106], and SMS [110]. As mentioned in Table 4, Pythia is incorporated in our baseline system.

Table 6: Storage overhead of all evaluated mechanisms

HMP [126] with local, gshare, and gskew predictors	11 KB
TTP with a metadata budget similar to the L2 cache	1536 KB
Pythia [32] with the same configuration in [32]	25.5 KB
Bingo [28] with the same configuration in [28]	46 KB
SPP [75] with perceptron-based prefetch filter [35]	39.3 KB
MLOP [106] with the same configuration in [106]	8 KB
SMS [110] with the same configuration in [110]	20 KB
<i>Hermes with POPET (this work)</i>	4 KB

We evaluate two variants of Hermes: **Hermes-O** and **Hermes-P**. These two variants differ only in Hermes request issue latency. *Hermes-O* (i.e., the optimistic Hermes) and *Hermes-P* (i.e., the pessimistic Hermes) use a request issue latency of 6 cycles and 18 cycles, respectively. Unless stated otherwise, *Hermes* represents the optimistic variant *Hermes-O*.

8. Evaluation

8.1. POPET Prediction Analysis

8.1.1. *Accuracy and Coverage of POPET.* Fig. 9 shows the comparison of POPET’s off-chip load prediction accuracy and coverage against those of HMP and TTP in the baseline system. The key takeaway is that POPET has significantly higher accuracy *and* coverage than HMP. POPET provides 77.1% accuracy with 74.3% coverage on average across all single-core workloads, whereas HMP provides 47% accuracy with 22.3% coverage. TTP, with a metadata budget of 1.5 MB, provides the highest coverage (94.8%) but with a significantly lower accuracy (16.6%). POPET’s superior accuracy *and* coverage directly translates to performance benefits both in single-core and eight-core system configuration (see §8.2 and §8.3).




Figure 9: Comparison of (a) accuracy and (b) coverage of POPET against those of HMP [126] and TTP.

8.1.2. *Effect of Different POPET Features.* Fig. 10 shows the accuracy and coverage of POPET using the five selected program features used individually and in various combinations. We make two key observations. First, each program feature individually produces predictions with a wide range of accuracy and coverage. The PC ⊕ cacheline offset feature produces the lowest-quality predictions with only 53.4% accuracy and 14.5% coverage, whereas the cacheline offset + first access feature produces the highest-quality predictions with 70.6% accuracy and 48.1% coverage. Second, by stacking multiple features together, the final POPET design achieves *both* higher accuracy and coverage than those provided by any

single individual program feature. We conclude that POPET is capable of learning from multiple program features to achieve both higher off-chip load prediction accuracy and coverage than any individual program feature can provide.




Figure 10: The accuracy and coverage of POPET using each program feature individually and in various combinations.

8.1.3. Usefulness of all features. To understand the usefulness of multi-feature learning, we analyze per-trace accuracy and coverage of POPET using each individual program feature. Fig. 11(a) shows the line graph of POPET’s prediction accuracy with each of the five program features individually for all single-core workload traces. The traces are sorted in ascending order of POPET accuracy using the feature cacheline offset + first access, since this feature individually has the highest average accuracy (as shown in Fig. 10(a)). The key takeaway from Fig. 11(a) is that there is no *single* program feature that individually provides the highest prediction accuracy across *all* workloads. Out of 110 workload traces, the features PC + first access, cacheline offset + first access, PC ⊕ byte offset, PC ⊕ cacheline offset, and last-4 load PCs provide the highest prediction accuracy in 47, 29, 20, 9, and 5 workload traces, respectively. We observe similar variability in POPET’s coverage, as shown in Fig. 11(b), where no single program feature individually provides the highest coverage across *all* workloads. This large variability of accuracy/coverage with different features in different workloads warrants learning using all features *in unison* to provide higher accuracy and coverage than any individual program feature across a *wide range* of workloads.




Figure 11: Line graph of POPET’s (a) accuracy and (b) coverage using each of the five program features individually across all 110 single-core workloads. No single feature can provide the best accuracy or coverage across all workloads.

8.2. Single-core Performance Analysis

8.2.1. Performance Improvement. Fig. 12 shows performance of Hermes (O and P), Pythia, and Hermes combined with Pythia normalized to the no-prefetching system in single-core workloads. We make three key observations. First, Hermes provides nearly half of the performance benefit of Pythia with only $\frac{1}{5} \times$ the storage overhead. On average, Hermes-O improves performance by 11.5% over a no-prefetching system, whereas Pythia improves performance by 20.3%. Second, Hermes-O (Hermes-P) combined with Pythia outperforms Pythia by 5.4% (4.3%). Third, Hermes combined with Pythia *consistently* outperforms Pythia in *every* workload category.




Figure 12: Speedup in single-core workloads

To better understand Hermes’s performance improvement, Fig. 13 shows the performance line graph of Hermes, Pythia, and Hermes combined with Pythia for every single-core workload trace. The traces are sorted in ascending order of performance gains by Hermes combined with Pythia over the no-prefetching system. We make four key observations from Fig. 13. First, Hermes combined with Pythia outperforms the no-prefetching system in all but three single-core workload traces. The `compute_int_539` and `605.mcf_s-782B` traces experience the highest and the lowest speedup ($2.3 \times$ and $0.8 \times$, respectively). Second, unlike Pythia, Hermes *always* improves performance over the no-prefetching system in *every* workload trace. Third, Hermes outperforms Pythia by 7.9% on average in 51 traces (e.g., `streamcluster-6B`, `Ligra_PageRank-79B`). In the remaining 59 traces, Pythia outperforms Hermes by 26% on average. Fourth, Hermes combined with Pythia consistently outperforms *both* Hermes and Pythia alone in almost every workload trace.

Based on our performance results, we conclude that, Hermes provides significant and consistent performance improvements over a wide range of workloads both by itself and when combined with the state-of-the-art prefetcher Pythia.




Figure 13: Single-core performance of all 110 workloads

8.2.2. Effect of the Off-chip Load Prediction Mechanism. Fig. 14 shows the performance of Hermes with POPET, Hermes-

HMP, Hermes-TTP, and the Ideal Hermes (see §3.1) combined with Pythia normalized to the no-prefetching system in single-core workloads. We make two key observations. First, Hermes with POPET outperforms both Hermes-HMP and Hermes-TTP. On average, Hermes-HMP, Hermes-TTP, and Hermes with POPET combined with Pythia provide 0.8%, 1.7%, and 5.4% performance improvement over Pythia, respectively. Second, Hermes-POPET provides nearly 90% of the performance improvement provided by the Ideal Hermes that employs an ideal off-chip load predictor with 100% accuracy and coverage. We conclude that Hermes provides performance gains due to both the high off-chip load prediction accuracy and coverage of POPET. Thus, designing a good off-chip predictor is critical for Hermes to improve performance.




Figure 14: Speedup of Hermes with three off-chip load predictors (HMP, TTP, and POPET) and the Ideal Hermes.

8.2.3. Effect on Stall Cycles. Fig. 15(a) plots the distribution of the percentage reduction in stall cycles due to off-chip load requests in a system with Hermes over the baseline system in single-core workloads as a box-and-whiskers plot.⁶ The key observation is that Hermes reduces the stall cycles caused by off-chip loads by 16.2% on average (up to 51.8%) across all workloads. PARSEC workloads experience the highest average stall cycle reduction of 23.8%. 90 out of 110 workloads experience at least 10% stall cycle reduction. We conclude that Hermes considerably reduces the stall cycles due to off-chip load requests, which leads to performance improvement.




Figure 15: (a) Reduction in stall cycles caused by off-chip loads. (b) Overhead in the main memory requests.

8.2.4. Overhead in Main Memory Requests. Fig. 15(b) shows the percentage increase in main memory requests in Hermes, Pythia, and Hermes combined with Pythia over the

⁶Each box is lower-bounded by the first quartile (i.e., the middle value between the lowest value and the median value of the data points) and upper-bounded by the third quartile (i.e., the middle value between the median and the highest value of the data points). The inter-quartile range (*IQR*) is the distance between the first and the third quartile (i.e., the length of the box). Whiskers extend an additional $1.5 \times IQR$ on the either side of the box. Any outlier values that falls outside the range of whiskers are marked by dots. The cross marked value within each box represents the mean.

no-prefetching system in all single-core workloads. We make two key observations. First, Hermes increases main memory requests by only 5.5% (on average) over the no-prefetching system, whereas Pythia by 38.5%. This means that, every 1% performance gain (see Fig. 12) comes at a cost of only 0.5% increase in main memory requests in Hermes, whereas nearly 2% increase in main memory requests in Pythia. We attribute this result to the highly-accurate predictions made by POPET, as compared to less-accurate prefetch decisions made by Pythia. Second, Hermes combined with Pythia further increases main memory requests by only 5.9% over Pythia. This means that, every 1% performance benefit by Hermes on top of Pythia comes at a cost of only 1% overhead in main memory requests. We conclude that, Hermes, due to its underlying high-accuracy prediction mechanism, adds considerably lower overhead in main memory requests while providing significant performance improvement both by itself and when combined with Pythia.

8.3. Eight-core Performance Analysis

Fig. 16 shows the performance of Pythia, Hermes-HMP, Hermes-TTP, and Hermes-POPET combined with Pythia normalized to the no-prefetching system in all eight-core workloads. The key takeaway is that due to the highly-accurate predictions by POPET, Hermes-POPET combined with Pythia consistently outperforms Pythia in every workload category. On average, Hermes-HMP, Hermes-TTP, and Hermes-POPET combined with Pythia provide 0.6%, -2.1%, and 5.1% higher performance on top of Pythia, respectively. Due to its inaccurate predictions, TTP generates many unnecessary main memory requests, which reduce the performance of Hermes-TTP combined with Pythia as compared to Pythia alone in the bandwidth-constrained four-core configuration. We conclude that Hermes provides significant and consistent performance improvement in the bandwidth-constrained eight-core system due to its highly-accurate off-chip load prediction.




Figure 16: Speedup in eight-core workloads

8.4. Performance Sensitivity Analysis

8.4.1. Effect of Main Memory Bandwidth. Fig. 17(a) shows the speedup of Hermes, Pythia, and Hermes combined with Pythia over the no-prefetching system in single-core workloads by scaling the main memory bandwidth. We make two key observations. First, Hermes combined with Pythia consistently outperforms Pythia in every main memory bandwidth configuration from $\frac{1}{16} \times$ to $4 \times$ of the baseline system. Hermes combined with Pythia outperforms Pythia alone by 6.2% and 5.5% in the main memory bandwidth configuration




Figure 17: Performance sensitivity to (a) main memory bandwidth, (b) different prefetching techniques, (c) Hermes request issue latency, and (d) on-chip cache hierarchy access latency. The baseline system configuration is highlighted in green. Other highlighted configurations closely match with various commercial processors [6, 7, 17].

with 200 and 12800 million transfers per second (MTPS), respectively. Second, Hermes by itself *outperforms* Pythia in highly-bandwidth-constrained configurations. This is due to the highly-accurate off-chip load predictions made by POPET, which incurs less main memory bandwidth overhead than the aggressive, less-accurate prefetching decisions made by Pythia. Hermes outperforms Pythia by 2.8% and 8.9% in 400 and 200 MTPS configurations, respectively.

8.4.2. Effect of the Baseline Prefetcher. We evaluate Hermes combined with four recently-proposed data prefetchers: Bingo [28], SPP [75] (with perceptron filter [35]), MLOP [106], and SMS [110]. For each experiment, we replace the baseline LLC prefetcher Pythia with a new prefetcher and measure the performance improvement of the prefetcher by itself and Hermes combined with the prefetcher. Fig. 17(b) shows the performance of the baseline prefetcher, and Hermes-P/O combined with the baseline prefetcher, normalized to the no-prefetching system in single-core workloads. The key takeaway is that Hermes combined with *any* baseline prefetcher consistently outperforms the baseline prefetcher by itself for all four evaluated prefetching techniques. Hermes+prefetcher outperforms the prefetcher alone by 6.2%, 5.1%, 7.6%, and 7.7%, for Bingo, SPP, MLOP, and SMS as the baseline prefetcher.

8.4.3. Effect of the Hermes Request Issue Latency. To analyze the performance benefit of Hermes over a wide range of processor designs with simple or complex on-chip datapath, we perform a performance sensitivity study by varying the Hermes request issue latency. Fig. 17(c) shows the performance of Hermes combined with Pythia normalized to the no-prefetching system in single-core workloads as Hermes request issue latency varies from 0 cycles to 24 cycles. The dashed-line represents the performance of Pythia alone. We make two key observations. First, the speedup of Hermes combined with Pythia decreases as the Hermes request issue latency increases. Second, even with a pessimistic Hermes request issue latency of 24 cycles, Hermes combined with Pythia outperforms Pythia. Pythia+Hermes outperforms Pythia by 5.7% and 3.6% with 0-cycle and 24-cycle Hermes request issue latency, respectively.

8.4.4. Effect of the On-chip Cache Hierarchy Access Latency. We evaluate Hermes by varying the on-chip cache hierarchy access latency. For each experiment, we keep the L1 and L2 cache access latencies unchanged and vary the LLC access la-

tency from 25-cycles to 50-cycles, to mimic the access latencies of a wide range of sliced LLC designs with simple or complex on-chip networks. Fig. 17(d) shows the performance of Pythia, and Hermes (O and P) combined with Pythia, normalized to the no-prefetching system in single-core workloads. We make two key observations. First, Hermes combined with Pythia consistently outperforms Pythia for *every* on-chip cache hierarchy latency. Hermes-O combined with Pythia outperforms Pythia alone by 3.6% and 6.2% in system with 40-cycle and 65-cycle on-chip cache hierarchy access latency, respectively. Second, the performance improvement by Hermes combined with Pythia increases as the on-chip cache hierarchy access latency increases. Thus, we posit that Hermes can provide even higher performance benefit in future processors with longer on-chip cache access latencies.

8.4.5. Effect of the Activation Threshold. We evaluate the impact of the activation threshold (τ_{act}) on Hermes’s performance by varying τ_{act} . Fig. 17 shows POPET’s accuracy and coverage (as line graphs on the left y-axis) and the performance of Hermes combined with Pythia over the no-prefetching system (as a bar graph on the right y-axis) across all single-core workloads as τ_{act} varies from -38 to 2 . The key takeaway from Fig. 17 is that POPET’s accuracy (coverage) increases (decreases) as τ_{act} increases. However, Hermes’s performance gain peaks near $\tau_{act} = -26$, which favors higher coverage by trading off accuracy. As POPET’s accuracy directly impacts Hermes’s main memory request overhead (and hence its performance in bandwidth-constrained configurations), we set $\tau_{act} = -18$ in POPET. Doing so simultaneously optimizes both POPET’s accuracy *and* coverage.




Figure 17: Effect of the activation threshold on POPET’s accuracy and coverage (on the left y-axis) and Hermes’s speedup (on the right y-axis) for all single-core workloads.

8.5. Power Overhead

To accurately estimate Hermes’s dynamic power consumption, we model our single-core configuration in McPAT [78] and

compute processor power consumption using statistics from performance simulations. Fig. 18 shows the runtime dynamic power consumed by Hermes, Pythia, and Hermes combined with Pythia, normalized to the no-prefetching system for all single-core workloads. We make two key observations. First, Hermes increases processor power consumption by only 3.6% on average over the no-prefetching system, whereas Pythia increases power consumption by 8.7%. Second, Hermes combined with Pythia incurs only 1.5% additional power overhead on top of Pythia. We conclude that Hermes incurs only a modest power overhead and is more efficient than Pythia alone.




Figure 18: Processor power consumption of Hermes, Pythia, and Hermes combined with Pythia.

9. Other Related Work

To our knowledge, this is the first work that proposes (1) a lightweight perceptron-based off-chip load predictor (POPET) that makes accurate predictions without having large metadata overhead, and (2) a mechanism (Hermes) that takes advantage of such a predictor to improve processor performance by eliminating on-chip cache access latency from the critical path of a correctly-predicted off-chip load. We have already quantitatively and qualitatively compared (1) POPET with HMP [126] and (2) Hermes with multiple state-of-the-art hardware data prefetchers: Pythia [32], Bingo [28], SPP [35, 75], MLOP [106], and SMS [110]. In this section, we discuss other related works.

Hit/miss prediction. Peir et al. [96] propose a Bloom-filter-based hit/miss predictor to optimize instruction scheduling in an out-of-order processor. Memik et al. [84] propose five different heuristics to keep track of cache contents using simple tables. Prior works [81, 100] also explore hit/miss prediction in DRAM caches. Loh and Hill [81] propose *MissMap*, which uses bitvectors to track the residency of cachelines in a large DRAM cache. Adapting *MissMap* to off-chip load prediction poses two key challenges. First, *MissMap* can have high false-positive prediction rate (as discussed in [60]). Second, the size of *MissMap* can grow very large, leading to large latency and storage overheads. POPET, on the other hand, requires only 4 KB of storage overhead, while producing highly-accurate off-chip load predictions.

Cache bypassing. High-performance cache management policies (e.g., [38, 43, 48, 53, 58, 63, 65-68, 72, 73, 80, 97, 101, 102, 114, 115, 120, 122]) dynamically bypass cache levels during a cache fill operation if the incoming cacheline is not expected to be used by the program in the future. This expectation (i.e., reuse prediction) can be provided by either a hardware-based reuse prediction mechanism [43, 58, 63, 65, 72,

73, 80, 97, 101, 102, 120] or a software hint [8, 19, 38, 114, 115], e.g., non-temporal load instructions employed by modern processors [8, 19]. The goal of these cache bypassing techniques is to better utilize the cache space by avoiding the insertion of useless cachelines into the cache. Hermes’s goal is different and orthogonal to these techniques: to reduce the latency of a long-latency off-chip load by eliminating the on-chip cache hierarchy access latency from its critical path. As such, Hermes can be combined with any cache bypassing technique.

Data prefetching. Prior prefetching techniques can be broadly categorized into three classes: (1) precomputation-based prefetchers that pre-execute program code to generate future loads [42, 51, 52, 55, 87-93], (2) temporal prefetchers that predict future load addresses by memorizing long sequence of demanded cacheline addresses [27, 31, 39-41, 44, 54, 57, 69, 71, 109, 117-119, 123, 124], and (3) spatial prefetchers that predict future load addresses by learning program access patterns over different memory regions [26, 28, 32, 33, 35, 37, 46, 56, 70, 75-77, 85, 94, 98, 106, 107, 110, 111]. As we show in §8.4.2, Hermes can be combined with any baseline prefetcher to provide higher performance than just the prefetcher alone.

10. Conclusion

We introduce Hermes, a technique that accelerates long-latency off-chip load requests by eliminating the on-chip cache hierarchy access latency from their critical path. To enable Hermes, we propose a perceptron-learning based off-chip load predictor (POPET) that accurately predicts which load requests might go off-chip. Our extensive evaluations using a wide range of workloads and system configurations show that Hermes provides significant performance benefits over a baseline system with a state-of-the-art prefetcher. As on-chip cache hierarchy continues to grow in size and complexity in future processors, we believe and hope that Hermes’s key observation and off-chip load prediction mechanism would inspire future works to explore a multitude of other memory system optimizations.

Acknowledgments

We thank Anant Nori and Sreenivas Subramoney for their valuable feedback on this work. We thank the anonymous reviewers of MICRO 2022 for their encouraging feedback. We thank the SAFARI Research Group members for providing a stimulating intellectual environment. We acknowledge the generous gifts from our industrial partners: Google, Huawei, Intel, Microsoft, and VMware. This work is supported in part by the Semiconductor Research Corporation and the ETH Future Computing Laboratory. The first author thanks his departed father, whom he lost in COVID-19 pandemic.

References

- [1] “2nd Cache Replacement Championship,” <https://crc2.ece.tamu.edu>.
- [2] “2nd Data Prefetching Championship,” <http://comparch-conf.gatech.edu/dpc2/>.
- [3] “3rd Data Prefetching Championship,” <https://dpc3.compas.cs.stonybrook.edu>.
- [4] “6th Generation Intel® Processor Family,” <https://www.intel.com/content/www/en/processors/core/desktop-6th-gen-core-family-spec-update.html>.
- [5] “AMD Gives More Zen Details: Ryzen,” <https://bit.ly/3ACs9ES>.

- [6] "AMD Ryzen Threadripper 3990X," https://en.wikichip.org/wiki/amd/ryzen_threadripper/3990x.
- [7] "AMD Zen2 EPYC 7702P," <https://en.wikichip.org/wiki/amd/epyc/7702p>.
- [8] "Caching of Temporal vs. Non-Temporal Data - Intel® 64 and IA-32 Architectures Developer's Manual," <https://www.intel.com/content/www/us/en/architecture-and-technology/64-ia-32-architectures-software-developer-vol-1-manual.html>.
- [9] "ChampSim," <https://github.com/ChampSim/ChampSim>.
- [10] "Golden Cove - Microarchitectures - Intel," https://en.wikichip.org/wiki/intel/microarchitectures/golden_cove.
- [11] "Golden Cove Microarchitecture (P-Core) Examined," <https://www.anandtech.com/show/16881/a-deep-dive-into-intels-alder-lake-microarchitectures/3>.
- [12] "Hermes - Wikipedia," <https://en.wikipedia.org/wiki/Hermes>.
- [13] "Hermes GitHub repository," <https://github.com/CMU-SAFARI/Hermes>.
- [14] "Hermes Zenodo repository," <https://doi.org/10.5281/zenodo.6909799>.
- [15] "Intel Core i5-12600K DDR4 Alder Lake CPU Review," <https://www.thefpsreview.com/2021/12/08/intel-core-i5-12600k-ddr4-alder-lake-cpu-review/6/>.
- [16] "Intel Details Golden Cove," <https://fuse.wikichip.org/news/6111/intel-details-golden-cove-next-generation-big-core-for-client-and-server-socs/>.
- [17] "Intel Xeon Gold 6258R," https://en.wikichip.org/wiki/intel/xeon_gold/6258r.
- [18] "L3 Cache Latency Comparison at Base Frequency," <https://www.cpubenchmark.com/cpu/intel-core-i9-10900k/benchmarks/l3-cache-latency-at-base-frequency/nvidia-geforce-rtx-2080-ti?res=1&quality=ultra>.
- [19] "MOVNTI - x86 ISA," <https://www.felixcloutier.com/x86/movnti>.
- [20] "PARSEC," <http://parsec.cs.princeton.edu/>.
- [21] "Second Championship Value Prediction (CVP-2)," <https://www.microarch.org/cvp1/cvp2/rules.html>.
- [22] "SPEC CPU 2006," <https://www.spec.org/cpu2006/>.
- [23] "SPEC CPU 2017," <https://www.spec.org/cpu2017/>.
- [24] "Surprisingly High Latency Discovered in Alder Lake," <https://bit.ly/3RhIk8z>.
- [25] "The Intel 12th Gen Core i9-12900K Review," <https://bit.ly/3wFRM6l>.
- [26] J.-L. Baer and T.-F. Chen, "An Effective On-chip Preloading Scheme to Reduce Data Access Penalty," in *SC*, 1991.
- [27] M. Bakhshalipour, P. Lotfi-Kamran, and H. Sarbazi-Azad, "Domino Temporal Data Prefetcher," in *HPCA*, 2018.
- [28] M. Bakhshalipour, M. Shakerinava, P. Lotfi-Kamran, and H. Sarbazi-Azad, "Bingo Spatial Data Prefetcher," in *HPCA*, 2019.
- [29] A. Basu, M. D. Hill, and M. M. Swift, "Reducing Memory Reference Energy with Opportunistic Virtual Caching," in *ISCA*, 2012.
- [30] B. M. Beckmann and D. A. Wood, "Managing Wire Delay in Large Chip-Multiprocessor Caches," in *MICRO*, 2004.
- [31] M. Bekerman, S. Jourdan, R. Ronen, G. Kirshenboim, L. Rappoport, A. Yoaz, and U. Weiser, "Correlated Load-Address Predictors," *ISCA*, 1999.
- [32] R. Bera, K. Kanellopoulos, A. Nori, T. Shahroodi, S. Subramoney, and O. Mutlu, "Pythia: A Customizable Hardware Prefetching Framework Using Online Reinforcement Learning," in *MICRO*, 2021.
- [33] R. Bera, A. V. Nori, O. Mutlu, and S. Subramoney, "DSPatch: Dual Spatial Pattern Prefetcher," in *MICRO*, 2019.
- [34] M. Besta, S. M. Hassan, S. Yalamanchili, R. Ausavarungnirun, O. Mutlu, and T. Hoefler, "Slim NoC: A Low-diameter On-chip Network Topology for High Energy Efficiency and Scalability," in *ASPLOS*, 2018.
- [35] E. Bhatia, G. Chacon, S. Pugsley, E. Teran, P. V. Gratz, and D. A. Jiménez, "Perceptron-Based Prefetch Filtering," in *ISCA*, 2019.
- [36] M. Cekleov and M. Dubois, "Virtual-Address Caches. Part 1: Problems and Solutions in Uniprocessors," *IEEE Micro*, 1997.
- [37] T.-F. Chen and J.-L. Baer, "Effective Hardware-Based Data Prefetching for High-Performance Processors," in *IEEE TC*, 1995.
- [38] C.-H. Chi and H. Dietz, "Improving Cache Performance by Selective Cache Bypass," in *HICSS*, 1989.
- [39] T. M. Chilimbi and M. Hirzel, "Dynamic Hot Data Stream Prefetching for General-Purpose Programs," in *PLDI*, 2002.
- [40] Y. Chou, "Low-cost Epoch-based Correlation Prefetching for Commercial Applications," in *MICRO*, 2007.
- [41] R. Cooksey, S. Jourdan, and D. Grunwald, "A Stateless, Content-Directed Data Prefetching Mechanism," *ASPLOS*, 2002.
- [42] J. Dundas and T. Mudge, "Improving Data Cache Performance by Pre-executing Instructions Under a Cache Miss," in *ICS*, 1997.
- [43] H. Dybdahl and P. Stenström, "Enhancing Last-Level Cache Performance by Block Bypassing and Early Miss Determination," in *ACSAC*, 2006.
- [44] M. Ferdman and B. Falsafi, "Last-touch Correlated Data Streaming," in *ISPASS*, 2007.
- [45] M. Ferdman, S. Somogyi, and B. Falsafi, "Spatial Memory Streaming with Rotated Patterns," in *In 1st JILP Data Fetching Championship*, 2009.
- [46] J. W. C. Fu, J. H. Patel, and B. L. Janssens, "Stride Directed Prefetching in Scalar Processors," in *MICRO*, 1992.
- [47] E. Garza, S. Mirbagher-Ajorpez, T. A. Khan, and D. A. Jimenez, "Bit-level Perceptron Prediction for Indirect Branches," in *ISCA*, 2019.
- [48] J. Gaur, M. Chaudhuri, and S. Subramoney, "Bypass and Insertion Algorithms for Exclusive Last-Level Caches," in *ISCA*, 2011.
- [49] B. Grayson, J. Rupley, G. Z. Zuraski, E. Quinnell, D. A. Jiménez, T. Nakra, P. Kitchin, R. Hensley, E. Brekelbaum, V. Sinha *et al.*, "Evolution of the Samsung Exynos CPU Microarchitecture," in *ISCA*, 2020.
- [50] N. Hardavellas, M. Ferdman, B. Falsafi, and A. Ailamaki, "Reactive NUCA: Near-Optimal Block Placement and Replication in Distributed Caches," in *ISCA*, 2009.
- [51] M. Hashemi, O. Mutlu, and Y. N. Patt, "Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads," in *MICRO*, 2016.
- [52] M. Hashemi and Y. N. Patt, "Filtered Runahead Execution with a Runahead Buffer," in *MICRO*, 2015.
- [53] Z. Hu, S. Kaxiras, and M. Martonosi, "Timekeeping in the Memory System: Predicting and Optimizing Memory Behavior," in *ISCA*, 2002.
- [54] Z. Hu, M. Martonosi, and S. Kaxiras, "TCP: Tag Correlating Prefetchers," in *HPCA*, 2003.
- [55] S. Iacobovici, L. Spracklen, S. Kadambi, Y. Chou, and S. G. Abraham, "Effective Stream-Based and Execution-Based Data Prefetching," in *ICS*, 2004.
- [56] Y. Ishii, M. Inaba, and K. Hiraki, "Access Map Pattern Matching for Data Cache Prefetch," in *ISCA*, 2009.
- [57] A. Jain and C. Lin, "Linearizing Irregular Memory Accesses for Improved Correlated Prefetching," in *MICRO*, 2013.
- [58] A. Jain and C. Lin, "Back to the Future: Leveraging Belady's Algorithm for Improved Cache Replacement," in *ISCA*, 2016.
- [59] A. Jaleel, K. B. Theobald, S. C. Steely Jr, and J. Emer, "High Performance Cache Replacement Using Re-Reference Interval Prediction (RRIP)," in *ISCA*, 2010.
- [60] M. Jalili and M. Erez, "Reducing Load Latency with Cache Level Prediction," in *HPCA*, 2022.
- [61] D. A. Jiménez and C. Lin, "Dynamic Branch Prediction with Perceptrons," in *HPCA*, 2001.
- [62] D. A. Jiménez, "Fast Path-Based Neural Branch Prediction," in *MICRO*, 2003.
- [63] D. A. Jiménez, "Dead Block Replacement and Bypass with a Sampling Predictor," in *JWAC*, 2010.
- [64] D. A. Jiménez and C. Lin, "Neural Methods for Dynamic Branch Prediction," *TOCS*, 2002.
- [65] D. A. Jiménez and E. Teran, "Multiperspective Reuse Prediction," in *MICRO*, 2017.
- [66] T. L. Johnson, D. A. Connors, and W.-M. W. Hwu, "Run-Time Adaptive Cache Management," in *HICSS*, 1998.
- [67] T. L. Johnson, D. A. Connors, M. C. Merten, and W.-m. W. Hwu, "Run-Time Cache Bypassing," *IEEE Trans. Computers*, 1999.
- [68] T. L. Johnson and W.-m. W. Hwu, "Run-Time Adaptive Cache Hierarchy Management via Reference Analysis," in *ISCA*, 1997.
- [69] D. Joseph and D. Grunwald, "Prefetching using Markov Predictors," in *ISCA*, 1997.
- [70] N. P. Jouppi, "Improving Direct-mapped Cache Performance by the Addition of a Small Fully-associative Cache and Prefetch Buffers," in *ISCA*, 1990.
- [71] M. Karlsson, F. Dahlgren, and P. Stenstrom, "A Prefetching Technique for Irregular Accesses to Linked Data Structures," in *HPCA*, 2000.
- [72] S. M. Khan, Y. Tian, and D. A. Jimenez, "Sampling Dead Block Prediction for Last-Level Caches," in *MICRO*, 2010.
- [73] M. Kharbutli and Y. Solihin, "Counter-Based Cache Replacement and Bypassing Algorithms," *IEEE TC*, 2008.
- [74] C. Kim, D. Burger, and S. W. Keckler, "An Adaptive, Non-Uniform Cache Structure for Wire-Delay Dominated On-Chip Caches," in *ASPLOS*, 2002.
- [75] J. Kim, S. H. Pugsley, P. V. Gratz, A. Reddy, C. Wilkerson, and Z. Chishti, "Path Confidence Based Lookahead Prefetching," in *MICRO*, 2016.
- [76] S. Kondguli and M. Huang, "Division of Labor: A More Effective Approach to Prefetching," in *ISCA*, 2018.
- [77] S. Kumar and C. Wilkerson, "Exploiting Spatial Locality in Data Caches using Spatial Footprints," in *ISCA*, 1998.
- [78] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures," in *MICRO*, 2009.
- [79] H. Litz, G. Ayers, and P. Ranganathan, "CRISP: Critical Slice Prefetching," in *ASPLOS*, 2022.
- [80] H. Liu, M. Ferdman, J. Huh, and D. Burger, "Cache Bursts: A New Approach for Eliminating Dead Blocks and Increasing Cache Efficiency," in *MICRO*, 2008.
- [81] G. H. Loh and M. D. Hill, "Efficiently Enabling Conventional Block Sizes for Very Large Die-Stacked DRAM Caches," in *MICRO*, 2011.
- [82] W. S. McCulloch and W. Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity," *The bulletin of mathematical biophysics*, vol. 5, no. 4, 1943.
- [83] S. McFarling, "Combining Branch Predictors," Digital Western Research Laboratory, Tech. Rep. 36, 1993.
- [84] G. Memik, G. Reinman, and W. H. Mangione-Smith, "Just Say No: Benefits of Early Cache Miss Determination," in *HPCA*, 2003.
- [85] P. Michaud, "Best-Offset Hardware Prefetching," in *HPCA*, 2016.
- [86] P. Michaud, A. Seznec, and R. Uhlig, "Trading Conflict and Capacity Aliasing in Conditional Branch Predictors," in *ISCA*, 1997.
- [87] O. Mutlu, H. Kim, and Y. N. Patt, "Address-Value Delta (AVD) Prediction: Increasing the Effectiveness of Runahead Execution by Exploiting Regular Memory Allocation Patterns," in *MICRO*, 2005.
- [88] O. Mutlu, H. Kim, and Y. N. Patt, "Techniques for Efficient Processing in Runahead Execution Engines," in *ISCA*, 2005.
- [89] O. Mutlu, H. Kim, and Y. N. Patt, "Efficient Runahead Execution: Power-Efficient Memory Latency Tolerance," in *IEEE Micro*, 2006.
- [90] O. Mutlu, H. Kim, J. Stark, and Y. N. Patt, "On Reusing the Results of Pre-Executed Instructions in a Runahead Execution Processor," *IEEE CAL*, 2005.
- [91] O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-order Processors," in *HPCA*, 2003.
- [92] O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead Execution: An Effective

- Alternative to Large Instruction Windows,” in *IEEE Micro*, 2003.
- [93] A. Naithani, S. Ainsworth, T. M. Jones, and L. Eeckhout, “Vector Runahead,” in *ISCA*, 2021.
- [94] S. Pakalapati and B. Panda, “Bouquet of Instruction Pointers: Instruction Pointer Classifier-based Spatial Hardware Prefetching,” in *ISCA*, 2020.
- [95] D. A. Patterson and J. L. Hennessy, *Computer Organization and Design: The Hardware Software Interface*. Morgan Kaufmann, 2016.
- [96] J.-K. Peir, S.-C. Lai, S.-L. Lu, J. Stark, and K. Lai, “Bloom Filtering Cache Misses for Accurate Data Speculation and Prefetching,” in *ICS*, 2002.
- [97] T. Piquet, O. Rochecoste, and A. Seznec, “Exploiting Single-Usage for Effective Memory Management,” in *ACSAC*, 2007.
- [98] S. H. Pugsley, Z. Chishti, C. Wilkerson, P.-f. Chuang, R. L. Scott, A. Jaleel, S.-L. Lu, K. Chow, and R. Balasubramonian, “Sandbox Prefetching: Safe Run-Time Evaluation of Aggressive Prefetchers,” in *HPCA*, 2014.
- [99] M. K. Qureshi, A. Jaleel, Y. N. Patt, S. C. Steely, and J. Emer, “Adaptive Insertion Policies for High Performance Caching,” in *ISCA*, 2007.
- [100] M. K. Qureshi and G. H. Loh, “Fundamental Latency Trade-off in Architecting DRAM Caches: Outperforming Impractical SRAM-Tags with a Simple and Practical Design,” in *MICRO*, 2012.
- [101] J. A. Rivers and E. S. Davidson, “Reducing Conflicts in Direct-Mapped Caches with a Temporality-Based Design,” in *ICPP*, 1996.
- [102] J. A. Rivers, E. S. Tam, G. S. Tyson, E. S. Davidson, and M. Farrens, “Utilizing Reuse Information in Data Cache Management,” in *ICS*, 1998.
- [103] F. Rosenblatt, “The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain,” *Psychological review*, vol. 65, no. 6, 1958.
- [104] A. Sembrant, E. Hagersten, and D. Black-Schaffer, “The Direct-to-Data (D2D) Cache: Navigating the Cache Hierarchy with a Single Lookup,” *ISCA*, 2014.
- [105] A. Sembrant, E. Hagersten, and D. Black-Schaffer, “A Split Cache Hierarchy for Enabling Data-Oriented Optimizations,” in *HPCA*, 2017.
- [106] M. Shakerinava, M. Bakhshalipour, P. Lotfi-Kamran, and H. Sarbazi-Azad, “Multi-Lookahead Offset Prefetching,” 2019.
- [107] M. Shevgoor, S. Koladiya, R. Balasubramonian, C. Wilkerson, S. H. Pugsley, and Z. Chishti, “Efficiently Prefetching Complex Address Patterns,” in *MICRO*, 2015.
- [108] J. Shun and G. E. Blelloch, “Ligra: A Lightweight Graph Processing Framework for Shared Memory,” in *PPoPP*, 2013.
- [109] S. Somogyi, T. F. Wenisch, A. Ailamaki, and B. Falsafi, “Spatio-Temporal Memory Streaming,” in *ISCA*, 2009.
- [110] S. Somogyi, T. F. Wenisch, A. Ailamaki, B. Falsafi, and A. Moshovos, “Spatial Memory Streaming,” in *ISCA*, 2006.
- [111] S. Srinath, O. Mutlu, H. Kim, and Y. N. Patt, “Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers,” in *HPCA*, 2007.
- [112] D. Tarjan and K. Skadron, “Merging Path and Gshare Indexing in Perceptron Branch Prediction,” *TACO*, 2005.
- [113] E. Teran, Z. Wang, and D. A. Jiménez, “Perceptron Learning for Reuse Prediction,” in *MICRO*, 2016.
- [114] G. Tyson, M. Farrens, J. Matthews, and A. Pleszkun, “A Modified Approach to Data Cache Management,” in *MICRO*, 1995.
- [115] G. Tyson, M. Farrens, J. Matthews, and A. R. Pleszkun, “Managing Data Caches using Selective Cache Line Replacement,” *IJPP*, 1997.
- [116] Z. Wang, J. Weng, J. Lowe-Power, J. Gaur, and T. Nowatzki, “Stream Floating: Enabling Proactive and Decentralized Cache Optimizations,” in *HPCA*, 2021.
- [117] T. F. Wenisch, M. Ferdman, A. Ailamaki, B. Falsafi, and A. Moshovos, “Practical Off-chip Meta-data for Temporal Memory Streaming,” in *HPCA*, 2009.
- [118] T. F. Wenisch, M. Ferdman, A. Ailamaki, B. Falsafi, and A. Moshovos, “Making Address-Related Prefetching Practical,” *IEEE micro*, 2010.
- [119] T. F. Wenisch, S. Somogyi, N. Hardavellas, J. Kim, A. Ailamaki, and B. Falsafi, “Temporal Streaming of Shared Memory,” in *ISCA*, 2005.
- [120] W. A. Wong and J.-L. Baer, “Modified LRU Policies for Improving Second-level Cache Behavior,” in *HPCA*, 2000.
- [121] D. A. Wood, S. J. Eggers, G. Gibson, M. D. Hill, and J. M. Pendleton, “An In-Cache Address Translation Mechanism,” *ISCA*, 1986.
- [122] C.-J. Wu, A. Jaleel, W. Hasenplaugh, M. Martonosi, S. C. Steely Jr, and J. Emer, “SHiP: Signature-based Hit Predictor for High Performance Caching,” in *MICRO*, 2011.
- [123] H. Wu, K. Nathella, D. Sunwoo, A. Jain, and C. Lin, “Efficient Metadata Management for Irregular Data Prefetching,” in *ISCA*, 2019.
- [124] H. Wu, K. Nathella, J. Pusdesris, D. Sunwoo, A. Jain, and C. Lin, “Temporal Prefetching Without the Off-Chip Metadata,” in *MICRO*, 2019.
- [125] T.-Y. Yeh and Y. N. Patt, “Two-level Adaptive Training Branch Prediction,” in *MICRO*, 1991.
- [126] A. Yoaz, M. Erez, R. Ronen, and S. Jourdan, “Speculation Techniques for Improving Load Related Instruction Scheduling,” in *ISCA*, 1999.
- [127] A. B. Yoo, M. A. Jette, and M. Grondona, “Slurm: Simple Linux Utility for Resource Management,” in *Workshop on job scheduling strategies for parallel processing*. Springer, 2003, pp. 44–60.

A. Artifact Appendix

A.1. Abstract

We implement Hermes using the ChampSim simulator [9]. In this artifact, we provide the source code of Hermes and necessary instructions to reproduce its key performance results.⁷ We identify four key results to demonstrate Hermes’s novelty:

- Comparison of accuracy and coverage of POPET against HMP and TTP (Fig. 9).
- Workload category-wise performance comparison in single-core workloads (Fig. 12).
- Workload category-wise performance comparison of Hermes with POPET with Hermes-HMP and Hermes-TTP in single-core workloads (Fig. 14).
- Performance sensitivity to different prefetching techniques (Fig. 17(b)).

The artifact can be executed in any machine with a general-purpose CPU and 36 GB disk space. However, we strongly recommend running the artifact on a compute cluster with `slurm` [127] support for bulk experimentation.

A.2. Artifact Check-list (Meta-information)

- **Compilation:** GCC 6.3.0 or above.
- **Data set:** Download traces using the supplied script.
- **Run-time environment:** Perl v5.24.1
- **Metrics:** IPC, predictor’s coverage, and accuracy.
- **Experiments:** Generate experiments using supplied scripts.
- **How much disk space required (approximately)?:** 36 GB
- **How much time is needed to prepare workflow (approximately)?:** ~ 2 hours. Mostly depends on downloading bandwidth.
- **How much time is needed to complete experiments (approximately)?:** ~ 12 hours using a compute cluster with 640 cores.
- **Publicly available?:** Yes.
- **Code licenses (if publicly available)?:** MIT
- **Archived (provide DOI)?:** <https://doi.org/10.5281/zenodo.6909799>

A.3. Description

A.3.1. How to Access. The source code can be downloaded from either GitHub [13] or Zenodo [14].

A.3.2. Hardware Dependencies. Hermes can be run on any system with a general-purpose CPU and at least 36 GB of free disk space.

A.3.3. Software Dependencies.

- `cmake` >= 3.20.2
- `gcc` >= v6.3.0
- `perl` >= v5.24.1
- `xz` >= v5.2.5
- `gzip` >= v1.6
- `megatools` >= v1.11.0
- `md5sum` >= v8.26
- `wget` >= v1.18
- `Microsoft Excel` >= v16.51

⁷This appendix focuses on reproducing four key results mentioned here. Nonetheless, the artifact contains files and necessary scripts to reproduce all results mentioned in the paper.

A.3.4. Data Sets. The ChampSim traces required to evaluate Hermes can be downloaded using the supplied script. Our implementation of Hermes is fully compatible with prior ChampSim traces that are used in previous cache replacement (CRC-2 [1]), data prefetching (DPC-3 [3]) and value-prediction (CVP-2 [21]) championships.

A.4. Installation

1. Clone Hermes from GitHub repository:

```
$ git clone https://github.com/CMU-SAFARI/Hermes.git
```

2. Please make sure to set environment variables as:

```
$ source setvars.sh
```

3. Clone Bloomfilter library inside Hermes home and build:

```
$ cd $HERMES_HOME
$ git clone https://github.com/mavam/libbf.git libbf/
$ cd libbf/
$ mkdir build && cd build/ && cmake ../
$ make clean && make
```

4. Build Hermes for single-core Intel Goldencove configuration as:

```
$ cd $HERMES_HOME
$ ./build_champsim.sh glc multi multi multi multi 1 1 0
```

A.5. Preparing Traces

This section describes the steps to download and verify the necessary ChampSim traces. We recommend the reader to follow the README in the GitHub repository to get up-to-date information.

1. Install the Megatools executable as:

```
$ cd $HERMES_HOME/scripts
$ wget -no-check-certificate
https://megatools.megous.com/builds/builds/
megatools-1.11.0.20220519-linux-x86_64.tar.gz
$ tar -xvf megatools-1.11.0.20220519-linux-x86_64.tar.gz
```

2. Use `download_traces.pl` script to download traces:

```
$ cd $HERMES_HOME/traces
$ perl $HERMES_HOME/scripts/download_traces.pl
-csv artifact_traces.csv -dir ./
```

3. Once the script finishes downloading 110 traces, please verify the checksum as follows. Make sure all traces pass the checksum test.

```
$ cd $HERMES_HOME/traces
$ md5sum -c artifact_traces.md5
```

4. If the traces are downloaded in other path, please update the full path in `MICRO22_AE.tlist` file inside `$HERMES_HOME/experiments` directory appropriately.

A.6. Experimental Workflow

This section describes the steps to generate and execute necessary experiments. We recommend the reader to follow `script/README.md` to know more about each script used in this section.

A.6.1. Launching Experiments. The following instructions enable launching all experiments required to reproduce key results in a local machine. We *strongly* recommend using a compute cluster with `slurm` support to efficiently launch experiments in bulk. To launch experiments using `slurm`, please provide `-local 0` (tested using `slurm v16.05.9`).

1. Create the jobfile for the experiments as follows:

```
$ cd $HERMES_HOME/experiments
```

```
$ perl $HERMES_HOME/scripts/create_jobfile.pl -exe
$HERMES_HOME/bin/glc-perceptron-no-
multi-multi-multi-multi-1core-1ch -tlist
MICRO22_AE.tlist -exp MICRO22_AE.exp -local 1 > jobfile.sh
```

2. Please make sure the paths used in `tlist` and `exp` files are appropriately changed before creating the jobfile.
3. If you are creating jobs for `slurm`, please set the `slurm` partition name appropriately (set as `slurm_part` by default). You might also want to set some key parameters for `slurm` configuration as follows: (a) max memory per node: 2 GB, (b) timeout: 12 hours.
4. Finally, launch the experiments as follows:

```
$ cd $HERMES_HOME/outputs/
$ source ../jobfile.sh
```

A.6.2. Rolling-up Statistics. The `rollup.pl` script parses a set of output files of a given experiment and dumps all statistics in a comma-separated-value (CSV) format. To automate the roll-up process, we use the following set of instructions which enable the creation of four CSV files in `experiments` directory. These CSV files are later used for comparison in Appendix A.7.

```
$ cd $HERMES_HOME/experiments
$ bash automate_rollup.sh
```

A.7. Evaluation

We use three metrics for comparing Hermes with previous works: (1) performance, (2) accuracy and (3) coverage of the off-chip predictor. The performance gain of a simulation configuration is measured with respect to the no-prefetching system using Eq. 2.

$$Perf_x = \frac{IPC_x}{IPC_{nopref}} \quad (2)$$

We measure the accuracy and coverage of each evaluated off-chip prediction mechanism using Eq. 3 and 4.

$$Accuracy_x = \frac{TP}{TP + FP} \quad (3)$$

$$Coverage_x = \frac{TP}{TP + FN} \quad (4)$$

Here, TP represents the number of predicted off-chip requests that *actually go* off-chip, FP represents the number of predicted off-chip requests that *do not go* off-chip, and FN represents the number of requests that are *not predicted* by the off-chip predictor *but go* off-chip.

To easily calculate the metrics, we provide a Microsoft Excel template to post-process the rolled-up CSV files generated in Appendix A.6.2. The template has five sheets. The first sheet (named *metadata*) contains the metadata to identify the category of each workload trace and experiment. Each of the remaining sheets shares the same name of the rolled-up CSV file. Each of these sheets is already populated with our collected results, necessary formulas, pivot tables, and charts to reproduce the results presented in the paper. *Only the blue-highlighted columns* in each sheet need to be populated by *your own* CSV files. Please follow these instructions to reproduce the results from *your own* CSV statistics files:

1. Copy and paste each CSV file into its corresponding sheet's top left corner (i.e., cell A1).

2. If you have copied the CSV file using an CSV application that already automatically separates the columns, then go to Step 4.
3. Immediately after pasting, convert the comma-separated rows into columns by going to Data → Text-to-Columns → Select comma as a delimiter. This replaces the already existing data in the sheet with the newly collected data.
4. *Refresh each pivot table in each sheet* by clicking on them and then clicking Pivot-Table-Analyse → Refresh.

The reader can also use any other data processor (e.g., Python pandas) to reproduce the same result.

A.8. Expected Results

- **Accuracy and coverage comparison.** POPET should show 77.1% accuracy, whereas HMP and TTP should show 46.8% and 16.6% accuracy, respectively. POPET should show 74.4% coverage, whereas HMP and TTP should show 22.3% and 94.8% coverage, respectively.
- **Performance comparison with Pythia.** Hermes-P, Hermes-O, Pythia, Pythia+Hermes-P, and Pythia+Hermes-O should provide 8.9%, 11.5%, 20.5%, 24.7%, and 25.6% geomean performance improvement, respectively, across all single-core workloads.
- **Performance comparison with prior predictors.** Pythia, Pythia+Hermes-HMP, Pythia+Hermes-TTP, and Pythia+Hermes-POPET should provide 20.5%, 21.1%, 22.09%, and 25.6% geomean performance improvement, respectively, across all single-core workload traces.
- **Performance comparison with varying prefetchers.** Both Hermes-P and Hermes-O improves performance than the prefetcher by itself for *every* baseline prefetcher type: Pythia, Bingo, SPP, MLOP, and SMS.

A.9. Experiment Customization

- The configuration of each prefetcher can be customized by changing the ini files inside the config directory.
- The exp files can be customized to run new experiments with different prefetcher combinations. More experiment files can be found inside experiments/extra directory. One can use the same instructions mentioned in Appendix A.6.1 to launch experiments.

A.10. Quick Troubleshooting

Please check the FAQ section in GitHub (<https://github.com/CMU-SAFARI/Hermes#frequently-asked-questions>) for quick troubleshooting tips.

A.11. Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-badging>
- <http://cTuning.org/ae/submission-20201122.html>
- <http://cTuning.org/ae/reviewing-20201122.html>

B. Extended Results

B.1. Performance Sensitivity to ROB Size

Fig. 19 shows the performance of Hermes, Pythia, and Hermes combined with Pythia normalized to the no-prefetching system in single-core workloads as the size of reorder buffer (ROB) varies from 256 entries to 1024 entries. The key takeaway is that Hermes combined with Pythia outperforms Pythia alone in every ROB size configuration. Pythia+Hermes outperforms Pythia by 6.7% and 5.3% in a system with 256-entry and 1024-entry ROB.




Figure 19: Performance sensitivity to reorder buffer size. The baseline configuration is marked in green.

B.2. Performance Sensitivity to LLC Size

Fig. 19 shows the performance of Hermes, Pythia, and Hermes combined with Pythia normalized to the no-prefetching system in single-core workloads as the per-core last-level cache (LLC) size varies from 3 MB to 24 MB. The key takeaway is that Hermes combined with Pythia outperforms Pythia alone in every LLC size configuration. Even in a system with a 12 MB and 24 MB LLC per core, Pythia+Hermes provides 2.5% and 1.3% performance benefit over Pythia alone, respectively.




Figure 20: Performance sensitivity to LLC size. The baseline configuration is marked in green.

B.3. Variation of Prediction Accuracy and Coverage with Different Prefetchers

Fig. 21 shows the off-chip load prediction accuracy and coverage when Hermes is combined with different baseline data prefetchers. We make two key observations. First, POPET’s accuracy and coverage varies widely based on the baseline data prefetcher. When combined with Pythia, Bingo, SPP, MLOP, and SMS, POPET provides accuracy of 77.3%, 78.1%, 73.4%, 79.9%, and 76.0%, while providing coverage of 74.2%, 77.6%, 65.9%, 81.7%, and 84.7%, respectively. Second, in a system without any baseline data prefetcher, POPET provides significantly higher accuracy (88.9%) and coverage (93.6%) than any configuration with a baseline prefetcher. This shows that, the prefetch requests generated by a sophisticated data prefetcher interfere with the off-chip load prediction. This is why POPET’s accuracy and coverage increases in absence of a data prefetcher.




Figure 21: Variation of off-chip load prediction accuracy and coverage with different data prefetchers.

B.4. Overhead in Main Memory Requests with Different Prefetchers

Fig. 22 shows the percentage increase in the main memory requests over the no-prefetching system by different types of data prefetchers alone, and in combination with Hermes in all single-core workloads. Combining Hermes with the baseline prefetcher increases the main memory request overhead by 5.9%, 7.6%, 5.9%, 8.6%, and 15.6% for the baseline prefetchers Pythia, Bingo, SPP, MLOP, and SMS, respectively.




Figure 22: Overhead in main memory requests with different data prefetchers.