

# Hardware-Level Thread Migration to Reduce On-Chip Data Movement via Reinforcement Learning

Quintin Fettes, *Student Member, IEEE*, Avinash Karanth, *Senior Member, IEEE*, Razvan Bunescu, *Member, IEEE*, Ahmed Louri, *Fellow, IEEE*, Kyle Shiflett, *Student Member, IEEE*,

**Abstract**—As the number of processing cores and associated threads in chip multiprocessors (CMPs) continues to scale out, on-chip memory access latency dominates application execution time due to increased data movement. Although tiled CMP architectures with distributed shared caches provide a scalable design, increased physical distance between requesting and responding cores has led to both increased on-chip memory access latency and excess energy consumption. Near data processing is a promising approach that can migrate threads closer to data, however prior hand-engineered rules for fine-grained hardware-level thread migration are either too slow to react to changes in data access patterns, or unable to exploit the large variety of data access patterns.

In this paper, we propose to use reinforcement learning (RL) to learn relatively complex data access patterns to improve on hardware-level thread migration techniques. By utilizing the recent history of memory access locations as input, each thread learns to recognize the relationship between prior access patterns and future memory access locations. This leads to the unique ability of the proposed technique to make fewer, more effective migrations to intermediate cores that minimize distance to multiple distinct memory access locations. By allowing a low-overhead RL agent to learn a policy from real interaction with parallel programming benchmarks in a parallel simulator, we show that a migration policy which recognizes more complex data access patterns can be learned. The proposed approach reduces on-chip data movement and energy consumption by an average of 41%, while reducing execution time by 43% when compared to a simple baseline with no thread migration; furthermore, energy consumption and execution time are reduced by an additional 10% when compared to a hand-engineered fine-grained migration policy.

**Index Terms**—Chip Multiprocessors, Thread Migration, Reinforcement Learning, Data Movement.

## I. INTRODUCTION

Manycore architectures suffer from both energy and performance penalties that exacerbate with the increase in the number of processing cores and associated threads [1]. As on-chip and off-chip memory access points become increasingly scattered around the on-chip interconnect fabric, memory access latency and energy consumption will increase as a result of longer round-trip request and response distances for data traversal. As technology scales into the sub-nanometer regime and provides opportunities to integrate hundreds of processing cores on a single chip, the cost of moving data

This article was presented in the International Conference on Hardware/Software Codesign and System Synthesis 2020 and appears as part of the ESWEEK-TCAD special issue.

in terms of energy and performance is currently dominating overall chip costs [1].

Near Data Processing is a powerful technique for reducing on-chip data movement during task execution [2]. Rather than repeatedly moving data from its on-chip storage location to the locus of computation, threads are moved closer to the locus of the data they require; if the cost (latency, energy) of moving the thread is lesser than the cost of moving the data, then the overall on-chip data movement is significantly reduced [3], [4], [5], [6], [7], [8], [9], [10], [11], [2]. By reducing the distance data travels on the chip, both energy consumed by the on-chip network and memory access latency can be reduced.

Prior work advanced a deadlock free, fine-grain thread migration scheme that can take advantage of local memory access patterns to migrate threads at the granularity of a few cycles [6], [12], [8], [13], [9], [11]. Such a model used a distributed non-uniform memory architecture (NUMA) where each core has a slice of shared cache, and every cache line corresponds to a unique core where the data is located on chip. By enforcing a single copy of the data to be present on-chip, prior work avoids the complexity and performance drawbacks of cache coherence protocols. However, the performance benefits come at a cost of increased memory access latency since the data must be repeatedly requested from the remote core every time it is required because the data cannot be cached locally. To mitigate the loss in performance, prior work used simple hand-engineered rules which recognize one type of data access pattern to migrate the threads. However, prior thread migration techniques are either too slow to react to changes in data access patterns or too simple to exploit the large variety of data access patterns.

In this paper, we propose to use reinforcement learning (RL) to learn data access patterns to improve hardware-level, deadlock-free thread migration techniques. In the proposed scheme, we use a history of recent memory accesses and their on-chip location to train a low-overhead RL policy to make migration decisions. By using the recent history of memory access locations as input, each thread learns to recognize the relationship between prior access patterns and future memory access locations. This enables the proposed technique to make fewer, more effective migrations to intermediate cores that minimize distance to multiple remote memory locations simultaneously. By training RL agents on real interactions within parallel programming benchmarks in a parallel simulator, we show that a migration policy based on more complex data

access patterns can be learned. The main contributions of the paper are as follows:

**Recognition of complex access patterns:** In [11], [13], simple manually engineered rules trigger thread migration whenever a fixed number of consecutive accesses to remote data is observed. In this paper, we show that a more intelligent, RL-trained policy can further reduce on-chip data movement. The proposed RL formulation relaxes constraints on where the threads can migrate and this allows the threads to move to intermediate cores that place the thread closer to multiple remote data locations. In Section III-A we illustrate how the proposed RL policy learns to combine many local memory access statistics in a more sophisticated way to improve both energy and execution time.

**Low-overhead RL to make migration decisions:** As each thread can theoretically make a decision on whether or not to migrate on every memory access, it was crucial to design a RL solution that had minimal overhead. While tabular RL algorithms would be the quickest to determine the optimal core to migrate, the size of the state-action table can be prohibitively large to store in memory. Therefore, we use the six most recent memory accesses and the current location of the thread on chip as features for approximate Q-Learning [14], [15], as well as other approximations to reduce the cost of computing migrations.

**RL as an improvement operator for existing solutions:** In the proposed formulation, there are approximately  $1.76 \times 10^{16}$  possible states. We use the simple policy from [9], [11] to collect experience for pre-training the RL agent. Effectively, this forces the RL agent to first learn the value of a known, good policy, and then fine tune it to form a better, more complex policy. In this sense, the proposed algorithm can be thought of as a policy improvement operator on the previous policy. This step proved crucial to good performance.

**Less on-chip traffic than existing methods:** We compared the performance of our proposed algorithm to a baseline that never migrates, as well as against our implementation of the NUMA architecture described in [9], [11], [13] using the Snipersim simulator [16]. Relative to the baseline, the RL approach was able to save 41.1% energy and reduce execution time by 43.1% on a set of Splash2 [17] and Parsec [18] benchmarks. Likewise, when compared to prior work from [9], the proposed RL algorithm saved an additional 10.2% energy and reduced execution time by an additional 9.6% on the same set of benchmarks.

## II. BACKGROUND

In this section, we provide background on the key aspects of the prior thread-migration based framework described in [9]; this prior work is referred to as TM throughout this paper. We also provide background on reinforcement learning (RL) and the RL algorithm used to learn the migration policy.

### A. Prior Hardware-Level Thread Migration

1) *Remote Access Cache:* In the NUMA architecture used in [9], [11], [6], [8], [13], the address space is divided among

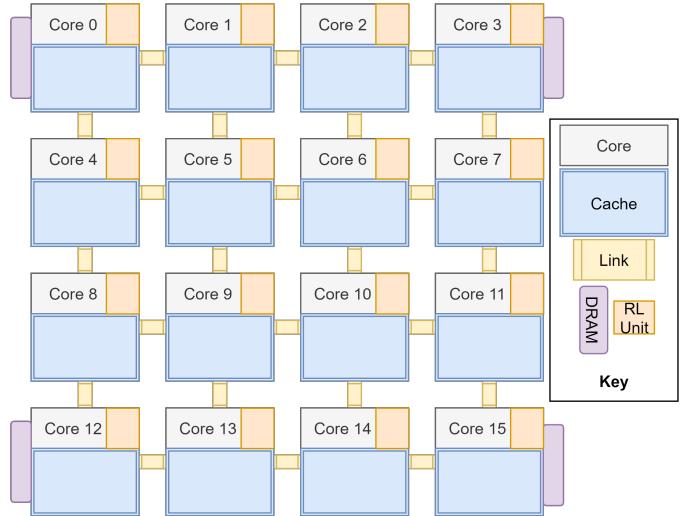


Fig. 1: Example  $4 \times 4$  mesh architecture using NUMA design. Each core is connected to its immediate neighbors by links. Each core views all cache slices as a single, logical cache; however, accessing data on non-local cache slice will require the core to go on the network.

the cache slices such that each cache line is assigned a unique home core. Suppose a thread  $T$  running on core  $C$  is reading or writing data with address  $A$  whose home core is  $H$ . If  $H = C$ ,  $T$  behaves normally such that the value is read/written without interacting with the on-chip interconnection network. On the other hand, if  $H \neq C$ ,  $T$  must issue a request on the on-chip interconnection network and wait for a response for the memory operation to be completed. When compared to a directory-based cache coherence protocol, an important difference is that the remote cache block is *never cached locally at  $C$* . Thus, accessing the cache line at  $H$  incurs round trip costs every time it is accessed. An example of a  $4 \times 4$  mesh-based NUMA architecture is shown in Figure 1.

2) *Fine-grained Deadlock-free Thread Migration:* Prior work in [9] seeks to exploit data locality by moving threads to the locus of data. Instead of remotely requesting memory operations,  $T$  could choose to migrate within the on-chip interconnection network to core  $H$  and execute the memory operation locally. If  $T$  is able to execute a sufficient number of memory operations on the remote core, the migration will have been “worth it”. In [9], a thread  $T$  chooses to migrate if it executes  $\zeta = 3$  consecutive memory accesses to a remote core. After  $\zeta$  consecutive accesses to the remote core, the hardware interrupts the execution of the thread and sends it on the on-chip interconnection network to the remote core.

If a Thread  $T_1$  is migrating to core  $H$  and a thread  $T'$  is already executing there,  $T'$  must be evicted. As described thus far, there are situations in which repeated evictions could lead to deadlock scenarios. To prevent deadlock, each core is made to have two thread contexts for which the core can multiplex execution [6], [12], [9], [13]. One context on each core  $H$  is marked as the *native context*  $H.nctxt$  for the thread initialized to that core; the other context is marked as the *guest context*  $H.gctxt$ . Only the thread initialized to a core can execute on

**Algorithm 1** Fine-Grained Thread Migration [9]

```

1: Input: a predetermined migration threshold  $\zeta$ 
2: Input: thread T executing on core C
3: Input: address A whose home core is H

4: procedure ONMEMORYACCESS( $T, C, H, \zeta$ )
5:   if  $T.last == H$  then
6:      $T.d = T.d + 1$ 
7:   else
8:      $T.d = 1$ 
9:    $T.last = H$ 

10:  if  $T.d == \zeta$  then
11:     $T.d = 0$ 
12:    if  $H \neq C$  then
13:      if  $T.ncore == H$  then
14:        Migrate T to  $H.nctxt$ 
15:      else
16:        if  $T'$  occupies  $H.gctxt$  then
17:          Evict  $T'$  to  $T'.ncore.nctxt$ 
18:        Migrate T to  $H.gctxt$ 
19:    else
20:      T stays at core C

```

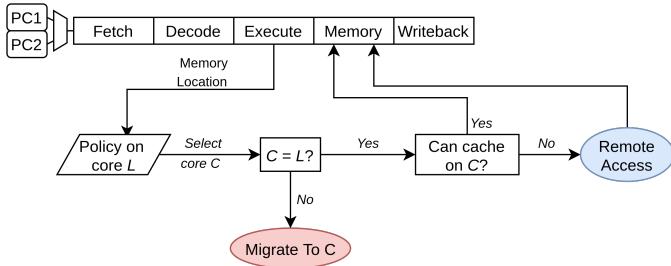


Fig. 2: This shows the migration decision process in a 5-stage pipeline. At the execution stage, migration policy (hand-engineered or RL) identifies the home core of a memory access instruction (load/store). This outputs a core, which can indicate either a migration or a decision to stay.

its native context; we will use  $T.ncore$  to denote the core that can run  $T$  in its native context. As a result, for every thread  $T$  there exists a native context, on the core  $T.ncore$ , which can only be occupied by  $T$ ; this guarantees that  $T$  can always return to an unoccupied thread context where it can resume execution after being evicted from a guest context. Furthermore, to prevent livelock scenarios in which threads are repeatedly migrated before making any progress, each thread is required to execute at least one instruction upon migrating to a new core. The migration policy implemented by TM is shown in Algorithm 1. There, the depth counter attribute  $T.d$  keeps track of the number of consecutive accesses made by thread  $T$  to the same core, whereas  $T.last$  denotes the core at which the most recently accessed memory address was located.

**B. Reinforcement Learning**

Reinforcement Learning (RL) is a type of machine learning which attempts to maximize a cumulative reward signal that an agent receives after performing actions in an environment. The reward is a scalar value describing progress towards a goal after being in a state and taking an action, and can be sparse and noisy. The formal setting for RL is the Markov Decision Process (MDP), represented as a tuple  $(S, A, P, R, \gamma)$  where:

- 1)  $S$  is the set of all states  $s$ .
- 2)  $A$  is the set of all actions  $a$ .
- 3)  $P : S \times A \rightarrow S$  is the model dynamics; it gives the probability of transitioning from the current state  $s$  to the next state  $s'$  upon taking the action  $a$ .
- 4)  $R : (S \times A) \times S \rightarrow \mathbb{R}$  yields rewards for state transitions.
- 5)  $\gamma$  is a discount factor that controls how far into the future the agent will optimize the reward.

The goal of the agent is to learn an optimal policy  $\pi^* : S \rightarrow A$  which maps states to actions such that the long term expected reward is maximized. It can be shown that the optimal policy is  $\pi^*(s) = \arg \max_a Q^*(s, a)$ , where the action-value function  $Q^*$  satisfies the Bellman optimality equation [19]:

$$Q^*(s, a) = \sum_{s', r} P(s', r | s, a) (r + \gamma \max_a Q^*(s', a')) \quad (1)$$

When the state space is small and discrete, the tabular Q-learning algorithm [19] is guaranteed to find the optimal action-value function  $Q^*(s, a)$  by storing the value of every possible state-action pair in a table, provided that the agent visits every possible state-action pair a sufficient number of times. One common approach to ensure this is to use an  $\epsilon$ -greedy policy, which chooses the action that maximizes the action-value function with probability  $1 - \epsilon$ , but reserves a small probability  $\epsilon$  for selecting random actions [20].

**Algorithm 2** Double Q-Learning with Linear Function Approximation and Offline Learning [14], [15]

```

1: Input: A Linear Function  $Q : S \times A \times \mathbb{R}^d \rightarrow \mathbb{R}$ 
2: Initialize Replay Memory  $D$  to capacity  $M$ 
3: Initialize online  $Q$  parameters  $\theta$  randomly
4: Initialize target  $\hat{Q}$  parameter  $\theta^- = \theta$ 
5: for each Episode do
6:   Observe state vector  $s_0 \in \mathbb{R}^d$ 
7:   for each  $t$  in Episode do
8:     Select  $a_t$  via an  $\epsilon$ -greedy policy on  $Q(s_t, \cdot; \theta)$ 
9:     Execute  $a_t$  then observe  $r_{t+1}$  and  $s_{t+1}$ 
10:    Store  $(s_t, a_t, r_{t+1}, s_{t+1})$  in  $D$ 
11:   for each Update Step do
12:     Sample a minibatch of  $(s_j, a_j, r_{j+1}, s_{j+1})$  from  $D$ 
13:     Select each  $a'_j = \arg \max_a Q(s_{j+1}, a; \theta)$ 
14:      $y_j = \begin{cases} r_{j+1} & \text{if terminal} \\ r_{j+1} + \gamma \hat{Q}(s_{j+1}, a'_j; \theta^-) & \text{else} \end{cases}$ 
15:     Calculate the TD error  $\delta_j = (y_j - Q(s_j, a_j; \theta))^2$ 
16:     Perform a gradient descent step on  $\delta$  w.r.t.  $\theta$ 
17:     Every  $C$  steps set  $\theta^- = \theta$ 

```

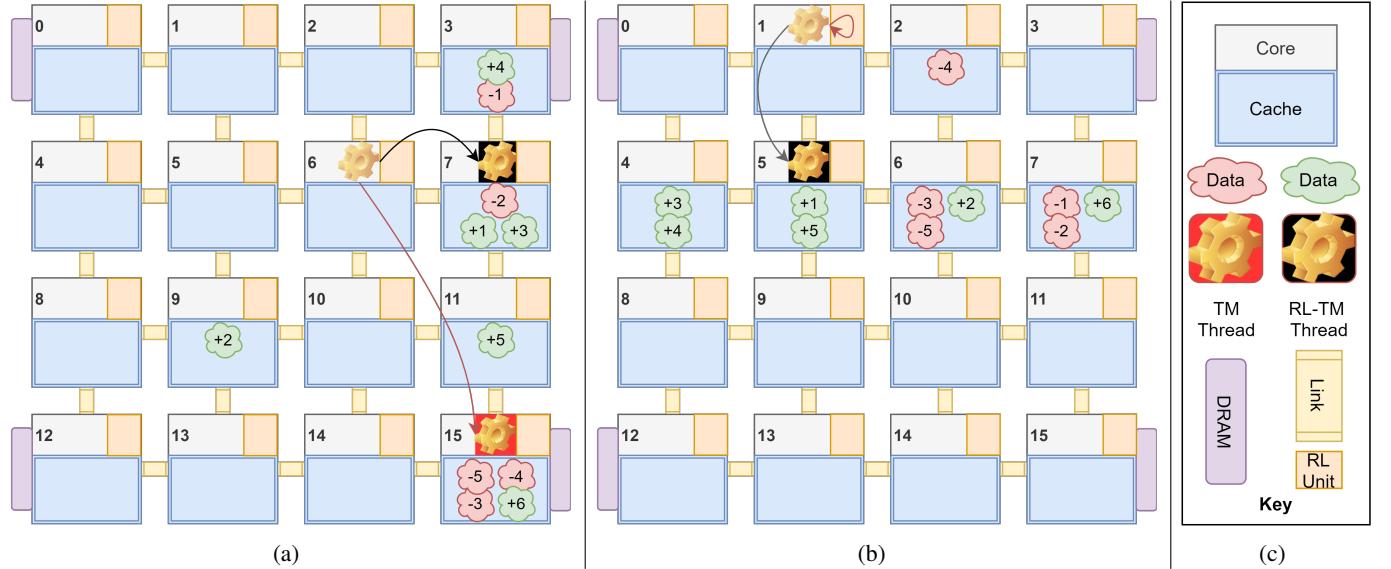


Fig. 3: Figures (a) and (b) show two examples of thread migrations made by the reinforcement learning policy RL-TM and the baseline hand engineered policy TM. The transparent gear shows the thread at its current location. The black lines show migrations made by RL-TM, whereas the red lines correspond to migration made by TM. RL-TM migrates in both scenarios: in (a) from Core 6 to Core 7, and in (b) from Core 1 to Core 5. TM migrates the thread only in (a), from Core 6, after observing three consecutive accesses to Core 15, whereas in (b) it never observes three consecutive accesses to the same core. The legend key is shown in (c). A detailed walk-through of each scenario is provided in the text.

1) *Double Q-Learning with Linear Function Approximation*: When the state space is large, as we will show later is the case for our proposed problem formulation, it is difficult or impossible to store the entire state-action table in memory and visit every state-action pair in a reasonable amount of time. In these cases, the action-value function can be approximated by some differentiable function  $Q(s, a; \theta)$  where  $\theta$  is a set of real-valued parameters, and the states are represented as vectors  $s \subseteq \mathbb{R}^d$ ; this allows the agent to generalize to unseen state-action pairs and learn more quickly. However, Q-Learning with function approximation is known to be unstable when the parameters are updated via gradient descent [20]. To improve stability during learning, [14] added a *target network*, i.e. a copy of the action-value function that updates more slowly, and *experience replay*, i.e. a buffer that stores old experience then samples mini-batches from it to get a better estimate of the gradient. We adopt a version of this algorithm that exclusively updates parameters offline and uses linear function approximation in lieu of a deep neural network. This is to keep the simulator simple and fast, and to minimize the overhead of computing the optimal action, respectively. The target network is used in a Double Q-learning setting [21], [15] to alleviate maximization bias and speed up learning. The full procedure that we used to learn the optimal policy is described in Algorithm 2. As seen in lines 13-14 of algorithm 2, the online network parameterized by  $\theta$  is used to select the next actions, while the offline target network parameterized by  $\theta^-$  is used to provide their corresponding state-action values, effectively de-correlating action selection from action evaluation. Finally, when the buffer is full, new transitions are added by replacing older transitions at random. In Section III-B we explain the

use of this algorithm for training thread migration RL agents.

### III. REINFORCEMENT LEARNING FOR HARDWARE LEVEL THREAD MIGRATION

This section discusses the proposed reinforcement learning-based thread migration (RL-TM). First, we use a walkthrough example of RL-TM that shows the advantages of RL-TM over the prior TM approach. Second, we introduce a formulation of hardware-level thread migration as a low-overhead RL problem. Last, we will discuss the RL overhead and implementation complexity.

#### A. Walk-through Example of Proposed RL-TM vs. Prior TM

Figure 3 illustrates the different migration decisions made by the proposed RL-TM and the prior TM algorithms on two real examples. The red  $-k$  clouds indicate memory accesses  $k$  time steps in the past, e.g.  $-1$  for the most recent memory access,  $-2$  for the memory access before it. Similarly, the green  $+k$  clouds indicate memory accesses  $k$  time steps in the future, e.g.  $+1$  indicates the upcoming memory access. At the time of the migration decision, RL-TM is aware of the location of the upcoming memory access as well as the prior 5 memory accesses, whereas TM operates using location of the upcoming and previous 2 memory accesses.

Consider the example of the thread executing on Core 6 in Figure 3a. Upon observing the location of the access labeled  $+1$ , but prior to performing the actual memory access, RL-TM migrates to Core 7 which is one hop away from Core 6 to its right. This moves the thread controlled by RL-TM closer to five future accesses, while moving it further from only a single future access. On the other hand, the same thread, when

controlled by TM, it observes 3 consecutive accesses (labeled  $-5, -4, -3$ ) to Core 15 and immediately migrates to Core 15 prior to performing the access labeled  $-3$ . As a result, under TM the thread made a more costly three hop migration to Core 15, moving it closer to three future memory accesses, but also farther from five future memory accesses.

In Figure 3b, based on the five prior accesses labeled with red clouds and the upcoming access labeled  $+1$ , RL-TM migrates the thread one hop away, from Core 1 to Core 5, which ends up closer to all 6 future memory accesses. On the other hand, TM does not migrate the thread, because no three consecutive accesses are made to any single core. This is a good example of the relative complexity of our policy. While TM was unable to identify a pattern which indicated a migration would be beneficial, RL-TM was able to consider a weighted combination of prior memory access locations to make a migration which ultimately reduced overall on-chip data movement. In both examples, by exploiting learned patterns that correlate previous with future access locations, RL-TM is more effective than TM at reducing the total distance of future accesses and also at lowering the cost of migrations.

### B. Threads as Reinforcement Learning Agents

Each thread is associated with an RL agent that attempts to minimize the overall on-chip traffic. An agent collects local information about recent memory accesses from its operating core and then uses that information to make migration decisions prior to completing each individual memory access. As shown in Figure 2, the RL-TM algorithm identifies the home core for a memory access instruction in the execution stage and passes this information to the RL agent. This location along with the previous 5 memory access locations forms an access history of length  $k = 6$  which is used as input to the RL policy to select a core. If the core is the current location of the thread, the thread does not migrate and either executes the memory access locally or issues a remote request on the interconnection network. If the core is non-local, execution is interrupted, the thread context is serialized, and the thread is sent to the remote core where it resumes execution of the instruction.

Formally, there are  $N$  possible *actions*, where  $N$  is the number of cores in the network. For a thread running on core  $m$ , selecting action  $n$  corresponds to migrating to core  $n$  if  $m \neq n$ , or staying at the same core otherwise. The *state* contains information about the current location of the thread and the core at which the data for the  $k$  most recent memory accesses is located. Importantly, the most recent memory access has not been completed yet when the state vector is formed; the agent has identified the location of the upcoming access and is making a decision on whether or not it should migrate before completing that access, as shown in Figure 2. The location of each of the  $k$  most recent memory accesses  $ma_1, ma_2, \dots, ma_k$  and the current thread location  $tl$  are encoded as one-hot vectors of length  $N$ . The resulting  $k+1$  one-hot vectors are then concatenated to form the overall state vector, i.e.  $s_t = [ma_1, ma_2, \dots, ma_k, tl]$ .

The *reward* is computed after each completed memory access as a sum of a *data* cost and a *thread* migration cost:

$$r(s_t, a_t) = -dhops(s_t.ma_1) - tsize * thops(a_t, s_t.tl) \quad (2)$$

where *dhops* is the round-trip number of hops the current memory access needs to reach the requesting core, *tsize* is the size of the thread in flits, and *thops* is the number of hops the thread travels to migrate from the current core  $s_t.tl$  to the core indicated by action  $a_t$ . The first term thus corresponds to the network traffic cost incurred from accessing data, whereas the second term estimates the network traffic cost caused by thread migration. Note that when  $a_t = s_t.cl$ , no migration happens, and thus  $thops(a_t, s_t.tl) = 0$ . Because the reward decreases as the amount of traffic increases, an agent seeking to maximize the reward is encouraged to make migration decisions that reduce the distance traveled by data in a way that optimally offsets the cost of migrating the thread.

During training the agents follow Algorithm 2. One episode is considered to be a complete run of a single benchmark application, so each iteration of the outer for loop on lines 5-17 begins when a benchmark application starts. A new timestep begins when an agent (operating from the perspective of a thread) identifies the location of a memory access, but before the memory access has been completed; for this reason, each agent has its own set of timesteps and follows lines 7-10 asynchronously. At line 8, each agent uses previously observed memory access locations and the location of the upcoming memory access as the state  $s_t$ . The state  $s_t$  is used as input to the function  $Q$  to select the action  $a_t$  that determines the core to move to (or stay at). At line 9, the agent executes the migration decision then waits until the memory access location for the next memory access instruction is known; the reward  $r_{t+1}$  is computed and the new location is used to compute the next state  $s_{t+1}$ . Finally, at line 10, the agent stores the full transition information,  $(s_t, a_t, r_{t+1}, s_{t+1})$  in the experience replay buffer. At this point, one iteration of the first inner for-loop has completed and the agent is ready to make another migration decision for timestep  $t + 1$ . After the benchmark application has finished, lines 11-17 are executed for a fixed number of steps. In these steps, the algorithm repeatedly samples minibatches of the transitions stored by each agent to update the parameters  $\theta$ . Because all agents try to solve the same optimization problem, i.e. minimize traffic due to on-chip data movement, they use the same policy. Therefore, they share the same action-value function  $Q(S, A; \theta)$  and experience replay buffer  $D$ , which requires the simulator API to globally communicate transition, action selection, and update information during training. Allowing the agents to share experience and use the same policy has the added benefit that training will have a lower sample complexity and thus proceed much faster. Furthermore, for the first 10 episodes of training, the behavior policy in line 8 of Algorithm 2 is replaced with the deterministic policy from [9], which migrates a thread to a core  $C$  if and only if it made three consecutive accesses to that core. After 10 episodes, the agent follows its own  $\epsilon$ -greedy policy shown on line 8, to improve on the initial deterministic policy.

**Algorithm 3** Low-Overhead RL at Test Time

---

- 1: *Input:* A Linear Function  $Q : S \times A \rightarrow \mathbb{R}$
- 2: Load previously trained  $Q$  parameters  $\theta$
- 3: Observe state vector  $s_0 \in \mathbb{R}^d$
- 4: **for each**  $t$  **do**
- 5:     Select action  $a_t$  using an  $\epsilon$ -greedy policy on  $Q(s_t, \cdot; \theta)$
- 6:     Execute  $a_t$  then observe next state  $s_{t+1}$

---

At test time, each thread performs thread migration according to Algorithm 3. In step 2, a copy of the trained parameters are loaded into the L1-D cache for each core so that threads can access the trained policy without the need for non-local communication. Additionally, the specialized hardware described in Section III-C is added to each core to facilitate low-overhead computation of migration decisions, which results in a very low overhead RL policy at test time. Similarly to training, each agent has its own set of timesteps and executes lines 4-6 of Algorithm 3 asynchronously. At line 5, each agent uses the location of previously observed memory access locations and the upcoming access location to form the state vector  $s_t$ ;  $s_t$  is then used as input to the function  $Q$  to select the action  $a_t$ , which determines the core to migrate to (or stay at). Finally, at line 6, the agent executes  $a_t$  and waits to observe the location of the next memory access to form the next state  $s_{t+1}$ . This loop continues indefinitely as long as threads need to execute memory access instructions. While using a fixed policy may lead to lower adaptability at test time, prior work has shown that a single, well-selected thread migration policy can work well across a wide range of application benchmarks [9], [11], [13].

**C. Overhead of Reinforcement Learning at Test Time**

The total number of parameters requiring storage at test time is calculated as follows. There are  $k + 1$  one-hot vectors used as input, each of length  $N$ , where  $N$  denotes the number of cores. The first  $k$  vectors represent the location of the  $k$  most recent memory accesses by a thread, whereas the last one-hot vector represents the thread's current location. In our experiments,  $k$  and  $N$  were 6 and 16, respectively, as seen in Tables II and I. Given that a distinct set of parameters is needed for each of the  $N$  actions for computing their action value, this implies that the total number of parameters is  $(k + 1)N^2$ . However, at test time the policy is copied into each core; the one-hot vector that represents the location of the thread is not needed anymore. Thus, only  $kN^2$  parameters must be stored at each core. To further reduce the storage and computation overhead of Algorithm 3, the 32b floats used at training time are converted to 8b integers at test time. In total, the action-value parameters are stored in approximately 1.5kb of the L1-D cache at each core.

To further reduce the overhead, during testing we restricted the agents to make migration decisions only on every 6<sup>th</sup> memory access. Figure 4 shows the energy and execution time of RL-TM as the period of migration decisions is varied. While the performance on splash2-fft is slightly better for a period of 5, the 0.3% performance gain would be offset by

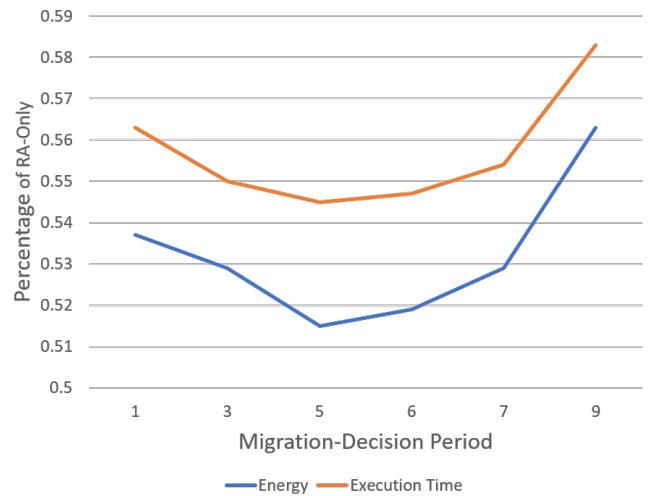


Fig. 4: Energy and Execution Time performance of RL-TM relative to the baseline RA-Only on splash2-fft as the migration period is varied.

TABLE I: System Configuration

Parameter	Settings
Cores, $N$	16
Thread Contexts per Core	2
Process Size	45nm
L1/L2 cache per core	32KB/128KB
L1/L2 Associativity	8/8 way set associative
Cache Block Size	64b
Interconnection Network Shape	2D Mesh
Interconnection Network Routing	XY routing
Hop Latency	2 cycles
Flit Size	128b
Thread Context Size	3.1kb

the increased overhead cost (Section III-D) of making more frequent migration decisions. Additionally, we found that it was possible to reduce the number of cores considered at test time with a simple heuristic: a minimal bounding box was formed to contain all cores corresponding to the last  $k$  memory accesses. On average, this reduced the number of actions considered from  $N = 16$  (the number of cores) to  $N_0 = 9.87$ , which is a 38.31% decrease in computation.

Because all  $k$  input vectors use one-hot encoding, multiplication with the corresponding parameter vectors simply requires looking up the parameters corresponding to the index of the non-zero element in the one-hot vector. Given that there are  $k$  input vectors, this implies that  $k$  parameters must be summed up for each of the  $N_0$  pruned actions. The resulting  $N_0$  action-values must then be compared to find the maximum. Finally, selecting actions on every 6<sup>th</sup> access reduces the total number of selections to  $\frac{1}{6}$  of the number of memory accesses. Overall, the time complexity consists of  $kN_0$  8b integer lookups,  $(k - 1)N_0$  8b integer additions, and  $N_0 - 1$  8b integer comparisons.

TABLE II: RL Hyperparameters

Parameter	Settings
Training Steps	1.5e4 per episode
Memory Access Location History	6 accesses
Learning Rate, $\alpha$	$2.5e - 5$
Gamma, $\gamma$	0.99
Optimizer	Adam [22]
Adam Epsilon	$1e - 8$
Experience Replay Size	1.5e6
Update Batch Size	32
Target Network Update Frequency, C	1e3
Exploration Rate, $\epsilon$	$0.1 * 0.9^E$
Train/Test Decision Frequency	Every 1/6 access(es)
Length of Access History, $k$	6

#### D. RL-TM Unit Implementation

The RL unit first uses adders to sum  $k$  8b integers for an average of 10 actions on each migration decision; these sums represent the action-value for each of the pruned actions. To sum the  $k$  values, we utilize carry-save adders (CSA) organized in a Wallace tree [23] to ensure a low propagation delay. Furthermore, a Brent-Kung adder (BKA) [24] is utilized to minimize power and area overhead. The 10 sets of  $k$  weights are assigned to a unique adder unit to compute the sums asynchronously; the resulting 10 sums are passed to a comparator tree to find the maximizing action, which is returned and used to make a migration decision. The RL unit implementation is shown in Figure 5. After modeling the hardware unit in Verilog HDL and synthesizing via Synopsys DC Ultra using the NanGate 45nm open cell library, results show a timing overhead of 1.7ns (< 5 cycles) and a power overhead of 5.23pJ, while the area overhead is  $7813\mu\text{m}^2$ . Including the RL units, the total area of the chip is  $746.14\text{mm}^2$ ; thus, collectively the 16 RL units occupy 15.9% of the total area of the chip. Because migration decisions are computed every 6 memory accesses in this design, the timing and energy overhead per memory access is effectively 0.28ns and 0.87pJ, respectively. Using the specialized hardware added to each core, this yields a timing and energy overhead of 1 cycle and 0.87pJ on average for every migration decision. In context, 1.4% of the total cycles are spent computing migration decisions. Overall, the resulting energy and timing penalty are negligible relative to the total energy consumption and execution time. While already small, these overhead numbers will decrease significantly as the architecture is modernized into the sub-10nm scale.

## IV. PERFORMANCE EVALUATION

This section discusses the network simulator, its configuration and its use in training and evaluation of the proposed RL-TM scheme.

#### A. Simulation Setup

Evaluation was performed using the Snipersim multi-core simulator [16] on a set of Splash2 [17] and Parsec [18] benchmarks. Snipersim is designed to use a directory-based memory hierarchy when organized in a network on chip

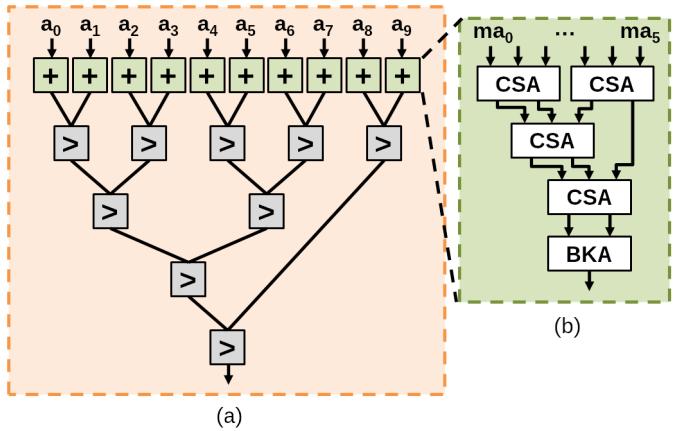


Fig. 5: RL-TM unit block diagram showing: (a) The adder and comparator tree organization (b) The Wallace tree adder.

mesh architecture. Accordingly, the simulator was modified to enable the remote access memory architecture described throughout this document. Furthermore, the simulator does not support hardware-level thread migration out of the box. To enable these experiments, the simulator was also modified to compute the cost in latency and energy associated with migrating threads over the interconnection network. For a fair comparison to [9], which was originally evaluated on a different simulator, we implemented their migration algorithm on the modified copy of Snipersim using the same simulation parameters. The relevant chip design specifications can be seen in Table I; notably, the system uses only 16 cores due to the time complexity of making migration decisions frequently at every core during simulation. As simulator technology improves along with computational power, training thread migration models on a larger number of cores will become more feasible, with the results in Section V indicating that it will scale effectively.

For the RL-enabled thread migration described in previous sections, henceforth referred to as RL-TM, a trained RL policy will be queried on every memory access by every thread on the chip as seen in Figure 2. The migration policy of the TM model compares the number of consecutive accesses to memory on the same core to a predetermined threshold; RL-TM replaces that simple manually engineered policy with a policy learned using a more complex combination of network statistics via Algorithm 2.

#### B. RL-TM Evaluation Procedure

We use a leave-one-out evaluation scenario, where at every step a single Splash2 or Parsec benchmark is used for testing, whereas all the remaining benchmarks are used to gather experience and train the RL policy; this process is repeated so that each benchmark is used for testing once. By never using the test benchmark during training, the leave-one-out evaluation results reflect how well a trained policy generalizes to an unseen benchmark. The hyperparameters used during training can be seen in Table II.

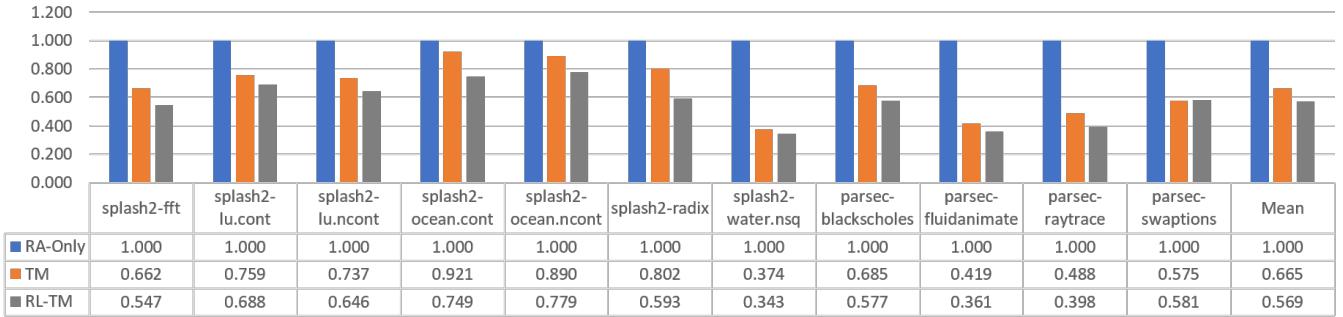


Fig. 6: Shown is execution time of the thread migration frameworks discussed. All values are normalized to the value of the RA-Only migration framework.

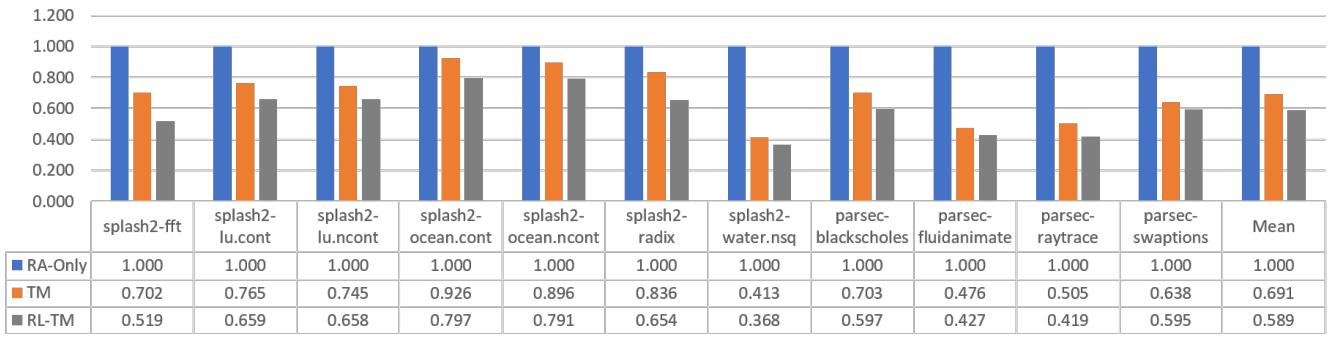


Fig. 7: The energy consumption of the three thread migration policies.

## V. RESULTS AND ANALYSIS

We compare the proposed RL-TM to both a Remote-Access-Only (RA-Only) baseline which can never migrate threads and TM. A breakdown of the execution time for each of the benchmarks can be viewed in Figure 6. It can be seen that the RL-enabled thread migration algorithm reduces the total execution time by an average of 43.1% relative to RA-Only, and by an additional 9.6% relative to TM. Additionally, the comparison of energy consumption is presented in Figure 7. Here, it's shown that the RL-enabled thread migration algorithm reduced the overall energy consumption by 41.1% relative to the RA-Only baseline and an additional 10.2% compared to TM. Figures 9 and 10 show that across all benchmarks, the number of migrations and evictions using RL-TM were reduced by 79.3% and 54.0%, respectively. The average distance (in hops) each memory access must travel round-trip on chip is shown in Figure 11. RL-TM reduces the average access distance by 23.7% relative to RA-Only, and an additional 1.1% relative to TM. Finally, Figure 12 shows the normalized energy consumption of RL-TM when it is trained using applications from one benchmark suite evaluated using application from the other benchmark; results are normalized to the energy consumption of RL-TM when trained using benchmark applications from both benchmark suites, excluding the application being evaluated. Figure 12a shows an average increase of 2.6% in energy consumption, whereas Figure 12b shows a mean increase of 3.4%. While it's possible that similarity among applications within the same benchmark suite (see splash2-lu.cont and splash2-lu.ncont) helps boost performance to some degree, it's also possible that

the relative drop in performance when restricting the training set is simply due to a less diverse training set.

### A. Analysis of Learned Policy

Results in Figures 6 and 7 show that RL-TM is a more effective thread migration algorithm than prior work while also maintaining enough generality to work across multiple benchmark suites using the same set of hyperparameters. Thus, the results indicate that using the hyperparameters shown in Table II would lead to effective thread migrations on a wide variety of parallel applications with little or no added design complexity, allowing the methodology to be extended to new applications easily, and could improve with a wider variety of training data. The rest of this section will analyze and explain the performance of RL-TM.

To understand why RL-TM is working, it is useful to first look at the number of migrations and evictions shown in Figure 9 and Figure 10. These show that the overall number of migrations and evictions for RL-TM are significantly reduced when compared to TM. Additionally, Figure 11 shows that RL-TM manages to reduce the average access distance by a smaller amount. Together, these three results show that RL-TM is able to achieve better data locality with a much smaller number of migrations. Migrations have a very large overhead due to the large thread context that must be sent over the interconnection network, so making less migrations significantly reduces timing and energy overhead for thread migration schemes. Interestingly, Figure 8 shows that RL-TM slightly improves over the RA-Only baseline by 2% increased local access rate, but makes 14.3% less local accesses than

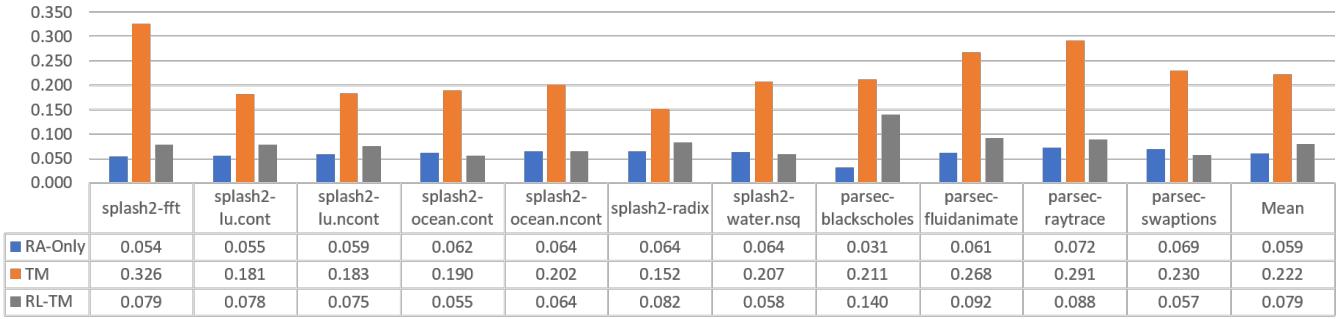


Fig. 8: The percentage of cache accesses which were local to the accessing thread's core.

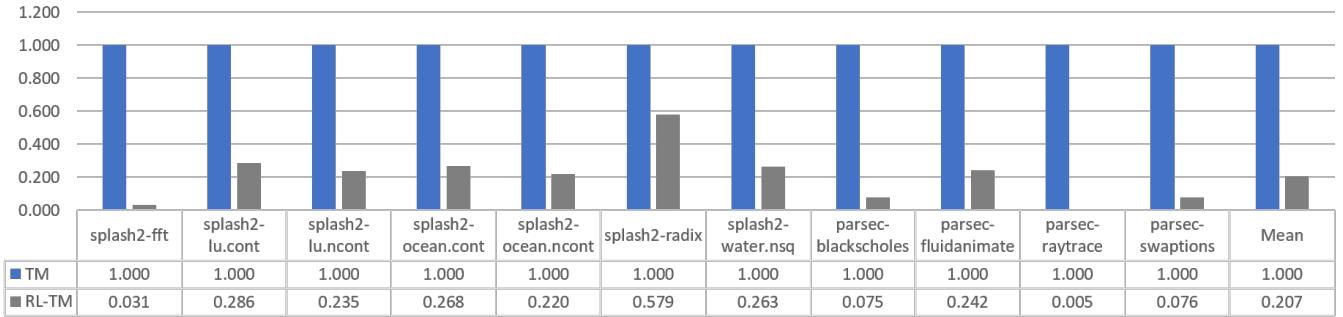


Fig. 9: The number of migrations performed by RL-TM relative to the TM algorithm of [9]; RA-Only is omitted because it never migrates threads.

TM. This displays the other strength of RL-TM; where TM is only ever able to migrate to a core where it has made multiple consecutive memory accesses, RL-TM is also able to migrate to intermediate cores that are close to multiple caches from which it needs data. In other words, while RL-TM less frequently moves directly to cores that can cache the data it needs, the benefit of moving to intermediate cores allows it to migrate less frequently while maintaining better data locality than RA-Only and TM. This reliance on intermediate cores suggests that RL-TM will become increasingly beneficial as the number of cores grows when compared to either baseline, as RL-TM will minimize expected distance to data, which will be increasingly distributed among a larger number of cores. While this indicates that a trained RL policy could become even more useful with higher core counts, the performance of the algorithm under practical constraints imposed by a growing state-action space remains an empirical question and is left for future work.

#### B. Need for bootstrapping

In the proposed RL formulation, there are approximately  $1.76 \times 10^{16}$  possible states. Rather than throwing out the valuable policy used by TM, we use it as the initial policy for the RL agents. Effectively, this forces the RL agent to learn the value of a known, good solution, and then fine-tune it into a better, more complex policy. This bootstrapping step was instrumental for obtaining good performance with RL-TM, which effectively makes RL-TM a policy improvement algorithm. A similar approach was taken in [25], where a policy gradient reinforcement learning algorithm was used to

improve a simpler search algorithm to play the game of Go. This view suggests that collecting data from other manually engineered policies, which make good migrations based on other simple access patterns, and using that data to pre-train RL-TM could further improve performance.

## VI. RELATED WORK

Most similar and a direct inspiration to our work is the work in [6], [8], [9]. There, the authors focused on introducing the concept of a directoryless, shared-memory protocol that relied on informed, fast, provably deadlock-free thread migration to reduce cache miss rates. The authors used a program counter to track the number of consecutive accesses to remote cores; after a predetermined, hand selected number of remote accesses to the same core, a thread would interrupt execution and migrate via the chip interconnection network to the location of the remote data. Our work adopts the idea of a NUMA memory design that relies on migrations to improve data locality, but uses a history of memory access locations to train a low-overhead RL policy to make decisions to migrate to any core in the network, rather than just the core associated with the remote access. In [11], this idea was improved upon by introducing a stack based core architecture. This allowed the system to dynamically decide to migrate differing amounts of the thread context, thus reducing the overhead incurred by migrating a thread over the interconnection network. [3] introduces the idea of computation migration; this is shown as an alternative to bringing data to the locus of computation, where programmers annotate the location where procedures should be run to improve data locality. Rather than copying

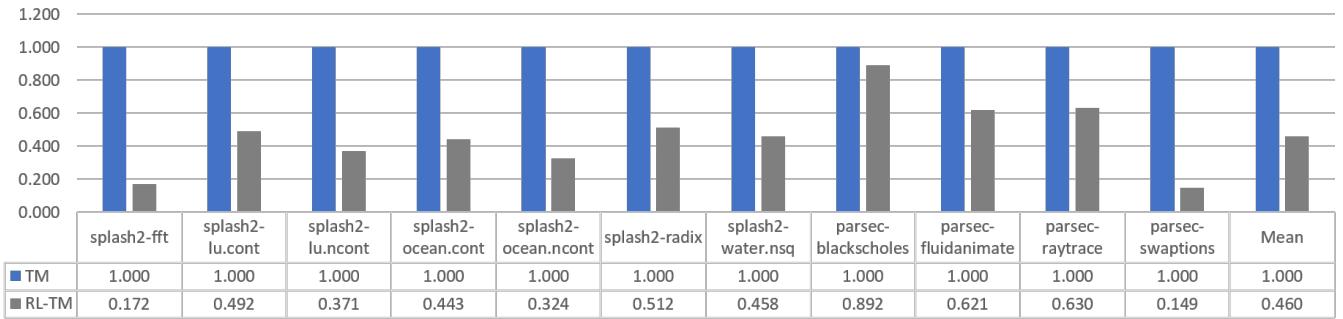


Fig. 10: The number of evictions caused by RL-TM relative to the algorithm of [9]; RA-Only is omitted because it never migrates threads.

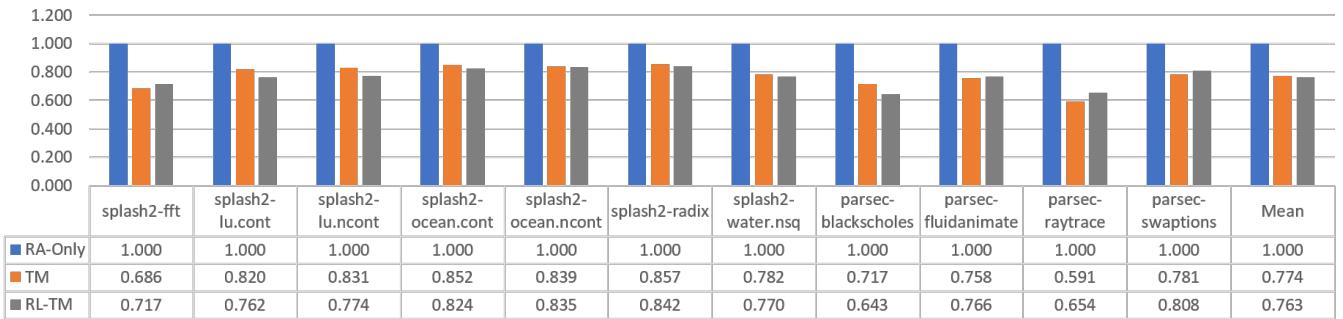


Fig. 11: Shown is the normalized mean memory access distance (in hops) for each method; for each algorithm on each benchmark, this is the expected distance each memory access will have to travel on-chip to be completed.

memory to multiple locations and using a coherence protocol, [26] explore systems to reactively place pages and threads on chip according to dynamic measurements made at run-time at the programming language level. [27] introduces a thread migration technique in a SMP multiprocessor system which is performed by the operating system at timeslice granularity; OS-based migration protocols such as [27] were effective for achieving long-term goals like load-balancing, but due to the high overhead of such methods, optimizing data locality for short bursts of data accesses is better suited for methods like those presented in this paper. [28] used a programming language level method to identify idle cycles at compile time which could be utilized to hide the overhead of these relatively high-overhead programming language level and OS-level migration schemes. Similarly, [29] attempts to alleviate the overhead of migrating call-stacks by compiling functions to a machine-independent string so that a programmer can specify when to move function execution on a heterogeneous system. [10] introduces a system that makes migration decisions both via compile-time optimization and via an OS-level system that analyzes the run-time behavior of threads. On the other hand, [4] and [5] utilize both dynamic voltage and frequency scaling and thread migration to save energy while sidestepping the area/performance overhead penalties of using DVFS alone; likewise, all methods rely on a central arbiter or handshake that makes the decision to migrate threads to a higher or lower frequency core based on their current workload, and can't be considered truly low-overhead or fine-grained. In an attempt to further reduce the latency of the

actual migration, [10] introduces the Teleportation Register File and special migration hardware to reduce the amount of time it takes to migrate threads on a heterogeneous chip in a register based ISA, rather than the stack based ISA used by [11].

Additional work has been done which focuses on the heterogeneous nature of modern System-on-Chips (SoCs). [30] attempted to map tasks to the appropriate type of processing element in heterogeneous System-on-Chips; this idea is improved in [31] by dynamically scheduling tasks which communicate to a single processing element. Similarly, [32], to decomposes applications to tasks then dynamically maps these tasks onto processing elements at runtime to minimize communication between tasks. [33] utilizes compile-time optimization to identify dependent instructions, then partitions them into tasks; a RL-based task scheduler then attempts to place tasks on their preferred type of processing element such that the distance to other tasks which they communicate with is minimized. While similar in spirit to thread migration, the aforementioned task mapping algorithms dynamically provide only initial mappings for tasks, and they could benefit from thread migration algorithms, such as the proposed algorithm RL-TM, which allow threads to move after their initial mapping; this is especially true in parallel applications which require frequent communications between many threads.

In addition, RL has been applied to many other tasks on the Network-on-Chip (NoC). Power management is frequently addressed by training RL agent(s) to make dynamic voltage and frequency scaling decisions for on-chip components [34],

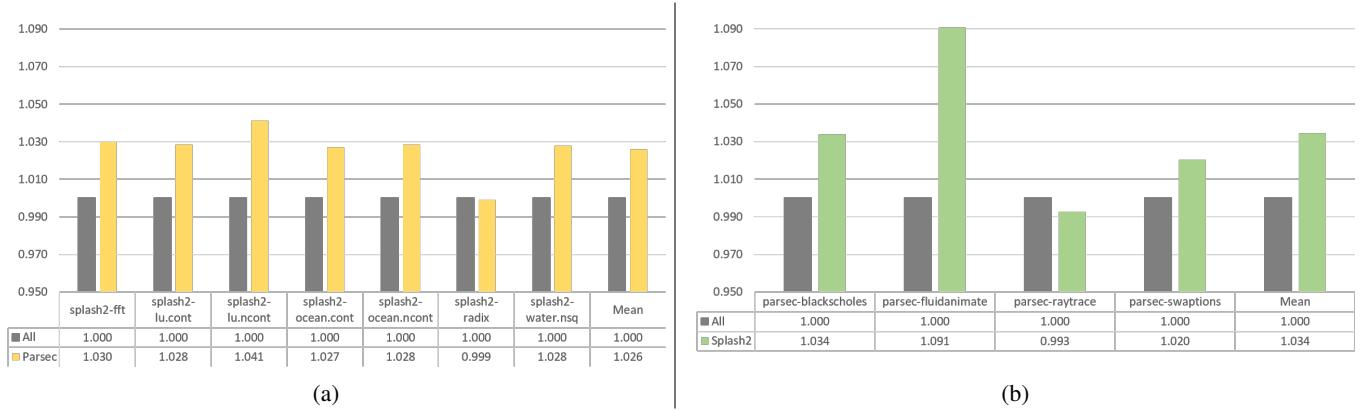


Fig. 12: Figures (a) and (b) show the energy consumption of RL-TM when restricting the training set to include applications from only one benchmark suite, normalized to the energy consumption when RL-TM is allowed to train on all benchmark applications except the one being evaluated. In Figure 12a, RL-TM was trained using the 4 Parsec benchmarks then evaluated on the 7 Splash2 benchmarks. In Figure 12b, RL-TM was trained on Splash2 benchmarks and evaluated on Parsec benchmarks.

[35], [36]. Additionally, RL policies have been used to save energy by learning to power-gate NoC components with the goal of reducing static power consumption [34], [37]. [37] each use RL to provide adaptive error detection and correction in the NoC. [38] used RL to learn an adaptive scheduling policy for the memory controller in a chip multiprocessor. Another common application of RL in the NoC is learning dynamic/adaptive routing schemes [39], [40], [41]; each of these works attempt to alleviate network congestion by introducing a dynamic routing algorithm learned via RL. On the other hand, our proposed RL-TM algorithm attempts to dynamically assign threads to cores for any parallel processing application, and is orthogonal to all of the above RL-for-NoC work.

## VII. CONCLUSIONS

Researchers predict the number of cores on multicores will reach the thousands in the near future. The proposed RL-TM offers a proof of concept that low-overhead approximate RL can be used to train a policy which effectively utilizes intermediate cores to make migrations to minimize the expected access distance for future memory accesses. By stripping out test-time learning which is characteristic of most contemporary RL algorithms and making several other improvements, RL-TM requires negligible overhead to compute migration decisions. Compared with a simple baseline that never migrates threads, the proposed methodology reduces execution time by 43.1% and energy consumption by 41.1%. When compared to a similar algorithm which utilizes a human engineered rule, RL-TM reduces execution time by an additional 9.6% and energy consumption by an additional 10.2%. Data collected from experiments shows that this decrease can largely be attributed to a 79.3% decrease in the number of migrations made, a 54.0% reduction in thread evictions and the use of intermediate cores as migration locations to improve data locality. RL-TM is able to do this by bootstrapping from the simpler human engineered policy. Finally, the experiments on both Splash2 and Parsec benchmarks suggest that our proposed methodology generalizes to many parallel applications. Future

work could diversify the set of benchmarks even further as simulators become more efficient and computers to run experiments become faster. Altogether, RL-TM is a general approach to hardware-level thread migration which effectively reduces on-chip data movement to both reduce execution time and save energy.

## ACKNOWLEDGMENT

This research was partially supported by NSF awards, CCF-1513606, CCF-1513923, CCF-1702980, CCF-1703013, CCF-1812495, CCF-1901192, and CCF-1936794. We sincerely thank all the anonymous reviewers for their excellent feedback on this work.

## REFERENCES

- [1] M. Ottavi, S. Pontarelli, D. Gizopoulos, C. Bolchini, M. K. Michael, L. Anghel, M. Tahoori, A. Paschalidis, P. Reviriego, O. Bringmann, V. Izosimov, H. Manhaeve, C. Strydis, and S. Hamdioui, “Dependable Multicore Architectures at Nanoscale: The View From Europe,” *IEEE Design Test*, vol. 32, no. 2, pp. 17–28, Apr. 2015, conference Name: IEEE Design Test.
- [2] A. Pattnaik, X. Tang, O. Kayiran, A. Jog, A. Mishra, M. T. Kandemir, A. Sivasubramaniam, and C. R. Das, “Opportunistic Computing in GPU Architectures,” in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA ’19. New York, NY, USA: ACM, 2019, pp. 210–223, event-place: Phoenix, Arizona. [Online]. Available: <http://doi.acm.org/10.1145/3307650.3322212>
- [3] W. C. Hsieh, P. Wang, and W. E. Weihl, “Computation migration: enhancing locality for distributed-memory parallel systems,” *ACM SIGPLAN Notices*, vol. 28, no. 7, pp. 239–248, Jul. 1993. [Online]. Available: <https://doi.org/10.1145/173284.155357>
- [4] K. K. Rangan, G.-Y. Wei, and D. Brooks, “Thread Motion: Fine-grained Power Management for Multi-core Systems,” in *Proceedings of the 36th Annual International Symposium on Computer Architecture*, ser. ISCA ’09. New York, NY, USA: ACM, 2009, pp. 302–313, event-place: Austin, TX, USA. [Online]. Available: <http://doi.acm.org/10.1145/1555754.1555793>
- [5] Q. Cai, J. González, G. Magklis, P. Chaparro, and A. González, “Thread shuffling: Combining DVFS and thread migration to reduce energy consumptions for multi-core systems,” in *IEEE/ACM International Symposium on Low Power Electronics and Design*, Aug. 2011, pp. 379–384, iSSN: null.
- [6] M. Lis, K. Shim, M. H. Cho, O. Khan, and S. Devadas, “Directoryless Shared Memory Coherence using Execution Migration,” *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, Dec. 2011.

[7] C. Ravishankar, S. Ananthanarayanan, S. Garg, and A. Jennings, "Analysis and evaluation of greedy thread swapping based dynamic power management for MPSoC platforms," in *Thirteenth International Symposium on Quality Electronic Design (ISQED)*, Mar. 2012, pp. 617–624, iSSN: 1948-3287.

[8] M. Lis, K. S. Shim, B. Cho, I. Lebedev, and S. Devadas, "Hardware-level thread migration in a 110-core shared-memory multiprocessor," in *2013 IEEE Hot Chips 25 Symposium (HCS)*, Aug. 2013, pp. 1–27.

[9] K. S. Shim, M. Lis, O. Khan, and S. Devadas, "Thread Migration Prediction for Distributed Shared Caches," *IEEE Computer Architecture Letters*, vol. 13, no. 1, pp. 53–56, Jan. 2014. [Online]. Available: <http://ieeexplore.ieee.org/document/6313584/>

[10] E. Forbes, Z. Zhang, R. Widjalaksono, B. Dwiel, R. B. R. Chowdhury, V. Srinivasan, S. Lipa, E. Rotenberg, W. R. Davis, and P. D. Franzon, "Under 100-cycle thread migration latency in a single-ISA heterogeneous multi-core processor," in *2015 IEEE Hot Chips 27 Symposium (HCS)*, Aug. 2015, pp. 1–1, iSSN: null.

[11] K. S. Shim, M. Lis, O. Khan, and S. Devadas, "The Execution Migration Machine: Directoryless Shared-Memory Architecture," *Prof. Devadas via Phoebe Ayres*, Sep. 2015. [Online]. Available: <https://dspace.mit.edu/handle/1721.1/108136>

[12] K. S. Shim, M. Lis, M. H. Cho, O. Khan, and Srinivas, "System-level Optimizations for Memory Access in the Execution Migration Machine (EM 2)," in *Second Workshop on Computer Architecture and Operating System co-design (CAOS)*, 2011.

[13] M. N. M. N. Lis, "Hardware-level fine-grained thread migration," Thesis, Massachusetts Institute of Technology, 2014. [Online]. Available: <https://dspace.mit.edu/handle/1721.1/93066>

[14] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015. [Online]. Available: <https://www.nature.com/articles/nature14236>

[15] H. v. Hasselt, A. Guez, and D. Silver, "Deep Reinforcement Learning with Double Q-Learning," in *Thirtieth AAAI Conference on Artificial Intelligence*, Mar. 2016. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12389>

[16] T. E. Carlson, W. Heirman, S. Eyerman, I. Hur, and L. Eeckhout, "An Evaluation of High-Level Mechanistic Core Models," *ACM Transactions on Architecture and Code Optimization (TACO)*, 2014.

[17] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 Programs: Characterization and Methodological Considerations," in *Proceedings of the 22Nd Annual International Symposium on Computer Architecture*, ser. ISCA '95. New York, NY, USA: ACM, 1995, pp. 24–36, event-place: S. Margherita Ligure, Italy. [Online]. Available: <http://doi.acm.org/10.1145/223982.223990>

[18] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. dissertation, Princeton University, January 2011.

[19] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3, pp. 279–292, May 1992. [Online]. Available: <https://doi.org/10.1007/BF00992698>

[20] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[21] H. V. Hasselt, "Double Q-learning," in *Advances in Neural Information Processing Systems 23*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. Curran Associates, Inc., 2010, pp. 2613–2621. [Online]. Available: <http://papers.nips.cc/paper/3964-double-q-learning.pdf>

[22] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]*, Jan. 2017, arXiv: 1412.6980. [Online]. Available: <http://arxiv.org/abs/1412.6980>

[23] C. S. Wallace, "A Suggestion for a Fast Multiplier," *IEEE Transactions on Electronic Computers*, vol. EC-13, no. 1, pp. 14–17, Feb. 1964, conference Name: IEEE Transactions on Electronic Computers.

[24] Brent and Kung, "A Regular Layout for Parallel Adders," *IEEE Transactions on Computers*, vol. C-31, no. 3, pp. 260–264, Mar. 1982, conference Name: IEEE Transactions on Computers.

[25] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis, "Mastering the game of Go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, Oct. 2017. [Online]. Available: <https://www.nature.com/articles/nature24270>

[26] M. Becchi and P. Crowley, "Dynamic thread assignment on heterogeneous multiprocessor architectures," in *Proceedings of the 3rd conference on Computing frontiers*, ser. CF '06. Ischia, Italy: Association for Computing Machinery, May 2006, pp. 29–40. [Online]. Available: <https://doi.org/10.1145/1128022.1128029>

[27] B. Weissman, B. Gomes, J. W. Quittek, and M. Holtkamp, "Efficient fine-grain thread migration with active threads," in *Proceedings of the First Merged International Parallel Processing Symposium and Symposium on Parallel and Distributed Processing*, Mar. 1998, pp. 410–414.

[28] H. Jiang and V. Chaudhary, "MigThread: thread migration in DSM systems," in *Proceedings. International Conference on Parallel Processing Workshop*, Aug. 2002, pp. 581–588, iSSN: 1530-2016.

[29] R. Veldema and M. Philipsen, "Near Overhead-free Heterogeneous Thread-migration," in *2005 IEEE International Conference on Cluster Computing*, Sep. 2005, pp. 1–10, iSSN: 2168-9253.

[30] J. Huang, A. Raabe, C. Buckl, and A. Knoll, "A workflow for runtime adaptive task allocation on heterogeneous MPSoCs," in *2011 Design, Automation Test in Europe*, Mar. 2011, pp. 1–6, iSSN: 1558-1101.

[31] A. K. Singh, T. Srikanthan, A. Kumar, and W. Jigang, "Communication-aware heuristics for run-time task mapping on NoC-based MPSoC platforms," *Journal of Systems Architecture*, vol. 56, no. 7, pp. 242–255, Jul. 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1383762110000330>

[32] A. Masood, E. U. Munir, M. M. Rafique, and S. U. Khan, "HETS: Heterogeneous Edge and Task Scheduling Algorithm for Heterogeneous Computing Systems," in *2015 IEEE 17th International Conference on High Performance Computing and Communications, 2015 IEEE 7th International Symposium on Cyberspace Safety and Security, and 2015 IEEE 12th International Conference on Embedded Software and Systems*, Aug. 2015, pp. 1865–1870.

[33] Y. Xiao, S. Nazarian, and P. Bogdan, "Self-Optimizing and Self-Programming Computing Systems: A Combined Compiler, Complex Networks, and Machine Learning Approach," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 6, pp. 1416–1427, Jun. 2019, conference Name: IEEE Transactions on Very Large Scale Integration (VLSI) Systems.

[34] H. Zheng and A. Louri, "An Energy-Efficient Network-on-Chip Design Using Reinforcement Learning," in *Proceedings of the 56th Annual Design Automation Conference 2019*, ser. DAC '19. New York, NY, USA: ACM, 2019, pp. 47:1–47:6, event-place: Las Vegas, NV, USA. [Online]. Available: <https://doi.acm.org/10.1145/3316781.3317768>

[35] R. Jain, P. R. Panda, and S. Subramoney, "Cooperative Multi-Agent Reinforcement Learning-Based Co-optimization of Cores, Caches, and On-chip Network," *ACM Transactions on Architecture and Code Optimization*, vol. 14, no. 4, pp. 32:1–32:25, Nov. 2017. [Online]. Available: <https://doi.org/10.1145/3132170>

[36] Q. Fettes, M. Clark, R. Bunescu, A. Karanth, and A. Louri, "Dynamic Voltage and Frequency Scaling in NoCs with Supervised and Reinforcement Learning Techniques," *IEEE Transactions on Computers*, vol. 68, no. 3, pp. 375–389, Mar. 2019, conference Name: IEEE Transactions on Computers.

[37] K. Wang, A. Louri, A. Karanth, and R. Bunescu, "IntelliNoC: a holistic design framework for energy-efficient and reliable on-chip communication for manycores," in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA '19. Phoenix, Arizona: Association for Computing Machinery, Jun. 2019, pp. 589–600. [Online]. Available: <https://doi.org/10.1145/3307650.3322274>

[38] E. Ipek, O. Mutlu, J. F. Martinez, and R. Caruana, "Self-Optimizing Memory Controllers: A Reinforcement Learning Approach," *ACM SIGARCH Computer Architecture News*, vol. 36, no. 3, pp. 39–50, Jun. 2008. [Online]. Available: <https://doi.org/10.1145/1394608.1382172>

[39] Y. Xiang, J. Meng, and D. Ma, "A Q-routing based self-regulated routing scheme for network-on-chip," in *2017 IEEE 9th International Conference on Communication Software and Networks (ICCSN)*, May 2017, pp. 177–181, iSSN: 2472-8489.

[40] M. A. Kinsky, S. Khadka, and M. Isakov, "PreNoc: Neural Network based Predictive Routing for Network-on-Chip Architectures," in *Proceedings of the on Great Lakes Symposium on VLSI 2017*, ser. GLSVLSI '17. Banff, Alberta, Canada: Association for Computing Machinery, May 2017, pp. 65–70. [Online]. Available: <https://doi.org/10.1145/3060403.3060406>

[41] S.-C. Kao, C.-H. H. Yang, P.-Y. Chen, X. Ma, and T. Krishna, "Reinforcement learning based interconnection routing for adaptive traffic optimization," in *Proceedings of the 13th IEEE/ACM International Symposium on Networks-on-Chip*, ser. NOCS '19. New York, New York: Association for Computing Machinery, Oct. 2019, pp. 1–2. [Online]. Available: <https://doi.org/10.1145/3313231.3352369>