Parallel Computing: Technology Trends I. Foster et al. (Eds.) © 2020 The authors and IOS Press. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0). doi:10.3233/APC200066

AITuning: Machine Learning-Based Tuning Tool for Run-Time Communication Libraries

Alessandro Fanfarillo^{a,1}, Davide Del Vento^a

^aNational Center for Atmospheric Research, Boulder, Colorado, USA

Abstract. In this work, we address the problem of tuning communication libraries by using a deep reinforcement learning approach. Reinforcement learning is a machine learning technique incredibly effective in solving game-like situations. In fact, tuning a set of parameters in a communication library in order to get better performance in a parallel application can be expressed as a game: *Find the right combination/path that provides the best reward*. Even though AITuning has been designed to be utilized with different run-time libraries, we focused this work on applying it to the OpenCoarrays run-time communication library, built on top of MPI-3. This work not only shows the potential of using a reinforcement learning algorithm for tuning communication libraries, but also demonstrates how the MPI Tool Information Interface, introduced by the MPI-3 standard, can be used effectively by run-time libraries to improve the performance without human intervention.

Keywords. MPI Machine Learning Reinforcement Learning Coarray Fortran

1. Motivaton

Tuning a general-purpose communication library is tightly related to the communication pattern utilized by the application, the network interconnect, the computer architecture, and the problem size. Profilers and other performance analysis tools have improved substantially in recent years and they are now able to provide the user with very accurate and descriptive interpretations of the various bottlenecks in a parallel application. However, most users in the scientific computing community do not have the time or expertise to study and tune the parameters of the communication libraries used by their codes. In fact, optimizing the parameters of communication libraries requires technical knowledge and time to try different configurations. For example, most Message Passing Interface (MPI) implementations offer hundreds of parameters that can provide significant speedup if they are set to their optimal value (which varies depending on the application), compared to the default configuration.

Furthermore, general-purpose communication libraries, like MPI, express several parallel programming models (e.g. one-sided, message-passing, task-based, etc...), and

¹Corresponding Author: elfanfa@ucar.edu

the optimal setting of a parameter used for a programming model might impact the performance when used on a different application, using a different programming model.

On the other hand, run-time communication libraries usually express fewer parallel programming models than general-purpose parallel programming libraries, and thus the communication pattern exposed by a run-time library can be interpreted and modeled much more easily.

In this work, we explore the use of machine learning techniques to optimize a particular run-time communication library, namely the OpenCoarrays run-time (used by the GNU Fortran compiler to implement the coarray support) and particularly its implementation on top of MPI-3.

Another important goal of this work is to demonstrate how the MPI Tool Information Interface, introduced by the MPI-3 standard, can be used effectively for automatic performance improvements when used by run-time libraries based on MPI-3, such as OpenCoarrays.

2. Related Work

The problem of tuning and auto-tuning communication libraries, like MPI, has been tackled several times in the past, using many different approaches.

In [10], Miceli et al. propose AutoTune, an extension of Periscope [1], an automatic distributed performance analysis tool. This framework tries to optimize a parallel application under many aspects including MPI tuning, thread affinity, and CPU frequency.

In [15], Sikora et al. extend again Periscope as part of the AutoTune project to implement autotuning capabilities for MPI applications. The output of the framework proposed is a set of tuning recommendation that can be integrated into the production version of the code. This tool provides the user with evolutionary algorithms able to heuristically guide the search of the most significant tuning parameters in MPI by executing a reasonable number of experiments.

Pellegrini et al. in [12] propose the use of two machine learning algorithms (decision trees and neural networks), to implement a predictive model that analyzes any MPI input program, and according to gained knowledge of the architecture, produces the value of a set of a predefined runtime parameters that provide optimal speedup. The overall approach proposed by Pellegrini et al. is similar to what we describe in this work, but our machine learning approach and modelization is completely different because it makes use of deep reinforcement learning techniques.

3. (Deep) Reinforcement Learning

The idea behind Reinforcement Learning (RL) is to have a learner called *agent* which interacts with an *environment* through *actions*. The *environment* responds to the *actions* and it presents new situations to the *agent*. The *environment* also gives rise to *rewards*: a numerical representation that the *agent* tries to maximize. The final goal of Reinforcement Learning is to find a *policy*, that maximizes the overall reward for the agent. A *policy* is a mapping from states to probabilities of selecting a certain *action*. Reinforcement Learning methods specify how the *agent* changes its *policy* as a result of experience.

Q-Learning is a reinforcement learning technique. It belongs to the class of modelfree methods and tries to estimate the Q-value function using the update equation expressed in 1.

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha [r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)]$$
(1)

Q-learning is just the Bellman optimality equation applied iteratively to evaluate and improve the Q-value function in a model-free problem, using a greedy policy. In other words, the best update rule to estimate the optimal action-value function Q for a given state, is the quantity that leads to the optimal policy. The optimal policy is the one given by the Bellman optimality equation, which is the max Q among all possible actions in the next state.

The Q-learning algorithm can be implemented by just keeping track of the Q-values of all the visited states in a table, but this is prohibitive for real problem with a large number of states.

Alternatively, one could estimate the Q-value of the states, using various techniques. One of these is called "Deep Q-Learning" and it involves the use of a deep neural network for the estimate. Unfortunately, applying non-linear function approximators to modelfree algorithms, such as Q-learning, could cause the Q-network to diverge [16], however there have been works to fix the divergence issue such as the gradient temporal-difference methods like [7] and [8].

The most famous and meaningful example of successful application of deep reinforcement learning is probably [11], where a convolutional neural network has been used to interpret the state of an Atari video game to produce the values of Q for all the possible actions allowed by the game. In the Atari work [11], the stability of the Q-learning algorithm, while using neural networks, is guaranteed by two mechanisms: experience replay and fixed Q-targets. Experience replay is random sampling over the entire experience accumulated and applying an optimization step on the neural network using the samples. This mechanism makes sure to break the temporal correlation of the experience observed by the network, resulting in a better stability and convergence of the algorithm. Q-targets means that the Q values used to compute the updates of the Q-learning algorithm belong to a neural network trained on old values. In [11], the authors use two neural networks, an they switch between the two after a certain number of steps to compute the Q-value for the targets in the Q-learning algorithm.

4. Potential in Communication Library Introspection

Understanding the performance issues of an MPI code is an operation that requires lowlevel information; for example, knowing how much time is spent in an MPI_Recv can help to understand whether the application suffers of poor load balancing or just high communication costs. Such a low-level information is usually hidden into the internal variables of the MPI implementation. For example, a typical information that can be useful to know is *how many messages are in the Unexpected Message Queue waiting to be received*?.

With the new tools information interface introduced in MPI-3, MPI provides a standard way to access performance data contained inside the MPI implementation (called *performance variables*) and internal variables that control the behavior of the implementation (called *control variables*). An example of a *control variable* is the one that defines the threshold, associated with the message size, that decides whether a message should be sent using the eager or rendezvous protocol.

Although the performance variables are common to any MPI implementation (e.g., Unexpected Message Queue length), the MPI Forum does not specify a direct way to get the status of these variables. The intent of the MPI Tool Information Interface (from now on MPI_T, see Section 4.1) is to enable an MPI implementation to expose implementation-specific details; for this reason is not possible to define variables that all MPI implementations must provide. This approach is called *introspection*. The most common use case for the MPI_T is to provide performance information and control variables to profilers and debuggers in order to help the users understanding issues and bottlenecks in MPI applications.

It is possible to write applications that take advantage of the information provided by MPI_T, but introducing such low-level concepts in user code is not advisable. We believe that the best opportunities to improve the performance of an MPI application using MPI_T are in the run-time communication libraries built on top of MPI. In fact, MPI_T has been already successfully used by run-time communication libraries to select the best algorithm based on the support provided by the MPI implementation. For example, Fanfarillo and Hammond in [5] use the MPI_T to select the best algorithm to implement *events* in OpenCoarrays [4], with a remarkable performance enhancement.

4.1. MPI Tool Information Interface (MPI_T)

MPI.T provides a standard interface to access *performance variables* and *control variables*. For both types of variables, there are several common concepts. In order to access a variable, an handle must be created first. With the handle the MPI implementation can provide low-overhead access to the internal variable.

Control variables allow the use to influence how the MPI implementation works. In order to use a *control variable*, the variable needs to be discovered. MPI provides functions to implement *introspection*, discover how many control variables are available, getting their details and modifying their values. During this work, we found out that it is important to modify all the *control variables* values before calling MPI_Init.

Performance variables are usually expressed in terms of queue lengths, waiting times, re-transmission attempts. For example, in a load imbalanced situations, where some processes make send requests before that the corresponding receives have been posted, the length of the unexpected message queue will be longer on some processes than on others. Another typical symptom of load imbalance is the longer time spent in a receive, waiting for the data to arrive. By combining the data with an understanding of how the implementation works, profilers are able to provide clues to the programmer on how to determine the source of the performance problem. The way *performance variables* are accessed is similar to the way *control variables* are managed but *performance variables* require an additional step: the creation of a session. A session enables different parts of the code. In other words, a session provides a way to isolate the use of a *performance variable* to a specific part of the code. In order to read the value associated with a *performance variable* the creation of handle and session should be performed after calling MPI_Init.

4.2. OpenCoarrays

OpenCoarrays [4] is an open-source software project for developing, porting and tuning transport layers that support coarray Fortran compilers. It targets compilers that conform to the coarray parallel programming feature set specified in the Fortran 2008 standard. It also supports several features defined in the Fortran 2018 standard including: *events* for fine-grain synchronization between parallel entities, *failed images* to manage failures, collective/reduction (called *collective*), and a partial implementation of *teams*, used to create independent subgroups of parallel entities. Currently, it is used as the run-time communication library by the GNU Fortran (GFortran) compiler.

OpenCoarrays defines an application binary interface (ABI) that translates highlevel communication and synchronization requests into low-level calls to a user-specified communication run-time library. This design decision liberates compiler teams from hardwiring communication-library choice into their compilers and it frees Fortran programmers to express parallel algorithms once, and reuse identical CAF source with whichever communication library is most efficient for a given hardware platform.

Since the first release of OpenCoarrays (August 2014), the widest coverage of coarray features was provided by a MPI based run-time library (LIBCAF_MPI). Because of the one-sided nature of coarrays, the run-time library uses almost exclusively MPI one-sided communication routines based on passive synchronization.

5. AITuning Design

AITuning has been designed as a separate component from run-time communication libraries. Its purpose is to guide the automatic tuning process of the libraries utilizing machine learning techniques. It is written in C++ and it is structured to be completely agnostic of run-time libraries, communication libraries, and machine learning algorithms and paradigms (although RL approaches are well suited for this problem).

5.1. Architecture

²The Controller class exposes a set of methods identified by the prefix AITuning_* that can be called by the run-time library. The method AITuning_start(string layer) takes a string representing the communication layer to be used. This method needs to be called before the initialization of the communication library (in this case MPI_Init_thread). In order to plug AITuning in OpenCoarrays without changing the source code of the latter, we decided to use the MPI Profiling Interface. We created wrappers for the MPI functions that AITuning needs to interact with (e.g. MPI_Init and MPI_Finalize) and called the AITuning_* methods from there.

As explained in Section 4.1, *control variables* and *performance variables* needed to be set before and after the actual call to MPI_Init_thread, respectively. Once the layer has been passed to the Controller object, a specific CollectionCreator is instantiated using the CollectionCreator object. The actual collection (in our case MPICHCollection-Creator) has predefined lists of control and performance variables that we decided and used for a specific AI component.

²A class diagram of the architecture is available on https://github.com/NCAR/AITuning

In order to make AITuning general enough to handle any kind of control and performance variables, we decided to declare the classes ControlVariable and PerformanceVariable as abstract. In fact, besides the default control and performance variables defined in a specific Collection object (related to a specific communication library implementation), it is possible to define UserDefined Performance Variables. This class of variables allows the user to defined specific performance variables, like the time spent to run the entire application, the time spent to execute a MPI_Win_flush and similar. Since they all inherit from the abstract class Performance Variable, they can be stored in the CollectionPerformanceVar object. In order to read performance variables, specific objects of the class Probes should be used. This class makes sure that the performance variables read using MPI_T or any other way (user defined included), respect certain criteria, like datatype, precision, and range.

All the performance variables keep track of the values detected during the program execution. At the end of the execution, in a wrapper of MPI_Finalize, statistics of the values get collected (e.g. average, max, min, median) and they will form the "state" representation passed to the AI component.

The entire machine learning process is performed in the MPI_Finalize wrapper, at the end of the program. The AI components receives a representation of the state of the application, which represents the state of the environment in a reinforcement learning setting. The reward gets computed in the AI component, based on previous data (in particular total_execution_time) and the reinforcement learning algorithm gets trained on the new data and produces a new action, defined as a "change" for a control variable. The new values for the control variables will be used during the next execution of the same application. A detailed description of the training process and AI component is provided in Section 5.2.

Not all the performance variables are the same; a variable like total_time cannot be passed to the RL algorithm as an absolute value. In fact, the same application has very different execution times when run on a different numbers of processes. In AITuning it is possible to declare a performance variable as "Relative". During the first run, the performance variable declared as relative will maintain in memory the absolute value of the quantity they represent. During the other runs, all the values of a relative performance variable are express as the difference between the absolute value obtained during the first run and the current absolute value. For example, if we consider the total execution time as performance variables, a positive value can be seen as a performance improvement, since during the first run the execution time was higher that the new value. This representation allowed us to write easy reward functions based on the results of relative variables.

5.2. Training

As first step, all the values of the performance variables are "standardized" against a reference run. To do so, a first run (or set of runs) is used as a reference for performance variables related to time and to a specific run in a consistent way. For this reason, when AITuning is active, the first run of the application is used to record the performance variables of the application when using a vanilla MPI implementation. The user communicates the first run by setting an environment variable AITUNING_FIRST_RUN = 1.

For every run other than the first, the algorithm produces a new action in the form of a "change" on a control variable. Each control variable has a fixed "step" to be used to change the absolute value of the control variable. For example, the MPICH control variable MPIR_CVAR_ASYNC_PROGRESS which controls the use of a helper thread to implement MPI asynchronous progress, can assume only two values: 0 and 1. On the other hand, the variable MPIR_CVAR_CH3_EAGER_MAX_MSG_SIZE assumes a numerical value representing the message size threshold to switch from the eager to the rendezvous protocol: in this case AITuning will change its value in predefined steps of 1024.

In every run, the neural network in charge of estimating the Q-value produces an estimate of the Q-value given a certain state provided by the performance variables. At the end of the run, the new reward gets computed and the neural network gets retrained based on the outcome. In order to make the Q-learning stable, we used the replay technique described in Section 3. We pick a random subset of the whole experience accumulated every 200 runs, and we train the neural network on that. We have not implemented the Q-target technique.

5.3. Control and Performance Variables for MPICH

For now, we focused our efforts only MPICH-3.2.1 because of the small number of control and performance variables exposed by the implementation, which made our reinforcement learning algorithm design and training faster. The control variables chosen for MPICH-3.2.1 are ASYNC_PROGRESS, CH3_ENABLE_HCOLL, CH3_RMA_DELAY_ISSUING_FOR_PIGGYBACKING,

CH3_RMA_OP_PIGGYBACK_LOCK_DATA_SIZE, POLLS_BEFORE_YIELD,

CH3_EAGER_MAX_MSG_SIZE. The only performance variable chosen from MPICH-3.2.1 was unexpected_recvq_length, representing the length of the unexpected message queue. We use several user-defined performance variables related to the average and maximum time needed to complete MPI_Win_Flush, MPI_Put,MPI_Get, and total application time. We also added the number of processes used in the run as input parameter.

5.4. Inference

AITuning will be shipped along with OpenCoarrays already trained for several MPI implementations and transport layers (e.g. GASNet). When the user decides to activate AITuning, he/she will compile OpenCoarrays using the PMPI wrapper. At this point, we recommend the user to run their application for at least 20 times. During these 20 runs, the RL algorithm will "explore" the new application and produce the right combination of parameters. During this exploration phase, AITuning may produce a configuration that penalizes the performance. At the end of the 20 runs, AITuning analyzes the results, discards the runs where the performance was penalized, and applies the median over the values of the control variables of the runs that provided good results within 5% from the best (creating an ensemble). Further runs of the same applications with different data input but same number of images will not require additional runs.

6. Experimental Evaluation

In order to train AITuning properly on MPICH-3.2.1, we decided to use two different supercomputers: Cheyenne (NCAR) an SGI machine with InfiniBand network interconnect and Edison (NERSC) a Cray XC30 with Aries interconnect. For the training we

decided to use four main codes parallelized with Coarrays Fortran: 1) CloverLeaf [9], 2) Lattice-Boltzmann code [13], 3) Skeleton Particle-in-cell [2], 4) Parallel Research Kernels [3]. We have run the aforementioned codes using a different number of processes going from 64 to 2048 for a total of 5000 runs.

6.1. Intermediate Complexity Atmospheric Research

The Intermediate Complexity Atmospheric Research (ICAR) [6] model developed at NCAR, is a simplified atmospheric model designed primarily for climate downscaling, atmospheric sensitivity tests, and educational uses. ICAR is a quasi-dynamical downscaling approach that uses simplified wind dynamics to perform high-resolution meteorological simulations 100 to 1000 times faster than a traditional atmospheric model and can therefore be used to better characterize uncertainty across numerical weather prediction models and climate models, and in dynamical downscaling.

In [14], Rouson et al. developed a mini-app of the ICAR model using coarray Fortran, showing great performance improvements. Since then, lead developer of ICAR, Ethan Gutmann, developed a fully functional version of ICAR based on coarray Fortran, which we used for testing AITuning. The version of ICAR we used is a full atmospheric model; the code include computation, communication and IO parts.

6.2. Results Evaluation

In Figure 1, we report the results obtained for ICAR running on Cheyenne using the default "vanilla" configuration set in MPICH-3.2.1, the optimized configuration found by AITuning after running ICAR 20 times, and an human optimized version based on reasonable guesses. The "default" bars represent the total time needed to complete a test case on ICAR using the default settings and in both cases, with 256 and 512 images, it provides the worst performance. On the other hand, the "optimized" version produced by AITuning always leads to the best performance. In both the 256 and 512 images cases, the manual optimization increased the eager limit by an order of magnitude higher than the default while leaving all the other setting as in the default configuration. For the case with 256 images, the optimized version provides 13% performance improvement compared to the vanilla version. For the case with 512 images, the optimized version provide 25% performance improvement over the vanilla version, mostly because of the higher communication cost imposed by the higher number of processes and same problem size (strong scaling).

The most influential tuning parameter for the ICAR test case resulted to be the presence of the asynchronous progress thread. We also noticed that some parameters have a different influence based on the number of processes being used. In particular, the value of MPICH_POLLS_BEFORE_YIELD played a much more relevant role in the case with 512 images than in the case with 256 images. This is not surprising because ICAR attempt to overlap computation with communication by using coarray "puts" instead of "gets". For the 256 case, the optimal configuration found by AITuning had MPICH_POLLS_BEFORE_YIELD set to the default value 1000, meaning that it was found not relevant. On the other hand, for the 512 images case, AITuning found a value of 1100. We manually changed the value of MPICH_POLLS_BEFORE_YIELD by keeping the configurations found by AITuning the same for both cases and found that in the case with



Figure 1. Performance comparison between default and optimized configurations

512 images, a value of MPICH_POLLS_BEFORE_YIELD between 1200 and 1500 provides the best performance, so it seems there is still room for improvement.

7. Conclusions and Future Work

In this work, we presented AITuning, a machine learning-base tuning tool for run-time libraries. AITuning has been released under open-source license and it is currently available on github ³. It currently works with the OpenCoarrays library, but its structure allows it to be extended to any run-time communication library, based on any communication layer. To the best of our knowledge, this paper is a unique contribution because it is the first attempt to try to find the optimal tuning parameters used a deep reinforcement learning algorithm and MPI_T. We tested AITuning and our RL algorithm, carefully designed for MPICH-3.2.1, using a real atmospheric code: ICAR. AITuning was able to produce a configuration of parameters that lead to 13% and 25% performance improvement for the case running on 256 and 512 images, respectively, compared to the default configuration. It also improves performance compared to an expertly tuned configuration, marginally for 256 images and substantially for 512.

In the future, we plan to extend our analysis to other MPI implementations with a higher number of control and performance variables. Furthermore, we will explore more options on the RL algorithm, and potentially other machine learning approaches. In our brief preliminary tests, it has been clear that whatever technique is chosen, it must be very robust to the noise of run-to-run variability. However, finding the best learning algorithm for AITuning is beyond the scope of this paper and left for a future work.

Finally, to better evaluate the results of the tool, we plan to test it on a larger number of machines and on a larger and more diverse set of applications.

³https://github.com/NCAR/AITuning

References

- Shajulin Benedict, Ventsislav Petkov, and Michael Gerndt. Periscope: An online-based distributed performance analysis tool. In Matthias S. Müller, Michael M. Resch, Alexander Schulz, and Wolfgang E. Nagel, editors, *Tools for High Performance Computing 2009*, pages 1–16, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [2] Viktor K. Decyk. Skeleton pic codes for parallel computers. *Computer Physics Communications*, 87(1):87 – 94, 1995. Particle Simulation Methods.
- [3] R. F. Van der Wijngaart and T. G. Mattson. The parallel research kernels. In 2014 IEEE High Performance Extreme Computing Conference (HPEC), pages 1–6, Sept 2014.
- [4] Alessandro Fanfarillo, Tobias Burnus, Valeria Cardellini, Salvatore Filippone, Dan Nagle, and Damian Rouson. Opencoarrays: Open-source transport layers supporting coarray fortran compilers. In Proceedings of the 8th International Conference on Partitioned Global Address Space Programming Models, PGAS '14, pages 4:1–4:11, New York, NY, USA, 2014. ACM.
- [5] Alessandro Fanfarillo and Jeff Hammond. Caf Events Implementation Using MPI-3 Capabilities. In Proceedings of the 23rd Message Passinge Interface Users and Developers Conference, EUROMPI '16. ACM, 2016.
- [6] Ethan Gutmann, Idar Barstad, Martyn Clark, Jeffrey Arnold, and Roy Rasmussen. The intermediate complexity atmospheric research model (icar). *Journal of Hydrometeorology*, 17(3):957–973, 2016.
- [7] Hamid R. Maei, Csaba Szepesvári, Shalabh Bhatnagar, Doina Precup, David Silver, and Richard S. Sutton. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Proceedings of the 22Nd International Conference on Neural Information Processing Systems*, NIPS'09, pages 1204–1212, USA, 2009. Curran Associates Inc.
- [8] Hamid Reza Maei, Csaba Szepesvári, Shalabh Bhatnagar, and Richard S. Sutton. Toward off-policy learning control with function approximation. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 719–726, USA, 2010. Omnipress.
- [9] A. C. Mallinson, S. A. Jarvis, W. P. Gaudin, and J. A. Herdman. Experiences at scale with pgas versions of a hydrodynamics application. In *Proceedings of the 8th International Conference on Partitioned Global Address Space Programming Models*, PGAS '14, pages 9:1–9:11, New York, NY, USA, 2014. ACM.
- [10] Renato Miceli, Gilles Civario, Anna Sikora, Eduardo César, Michael Gerndt, Houssam Haitof, Carmen Navarrete, Siegfried Benkner, Martin Sandrieser, Laurent Morin, and François Bodin. Autotune: A plugin-driven approach to the automatic tuning of parallel applications. In Pekka Manninen and Per Öster, editors, *Applied Parallel and Scientific Computing*, pages 328–342, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [11] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602, 2013.
- [12] Simone Pellegrini, Jie Wang, Thomas Fahringer, and Hans Moritsch. Optimizing mpi runtime parameter settings by using machine learning. In Matti Ropo, Jan Westerholm, and Jack Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 196–206, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [13] C. Rosales. Porting to the intel xeon phi: Opportunities and challenges. In 2013 Extreme Scaling Workshop (xsw 2013), pages 1–7, Aug 2013.
- [14] Damian Rouson, Ethan D. Gutmann, Alessandro Fanfarillo, and Brian Friesen. Performance portability of an intermediate-complexity atmospheric research model in coarray fortran. In *Proceedings of the Second Annual PGAS Applications Workshop*, PAW17, pages 4:1–4:4, New York, NY, USA, 2017. ACM.
- [15] Anna Sikora, Eduardo César, Isaías Comprés, and Michael Gerndt. Autotuning of mpi applications using ptf. In Proceedings of the ACM Workshop on Software Engineering Methods for Parallel and High Performance Applications, SEM4HPC '16, pages 31–38, New York, NY, USA, 2016. ACM.
- [16] J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690, May 1997.