

New Performance Modeling Methods for Parallel Data Processing Applications

JANKI BHIMANI, Northeastern University, USA

NINGFANG MI, Northeastern University, USA

MIRIAM LEESER, Northeastern University, USA

ZHENGYU YANG, Northeastern University, USA

Predicting the performance of an application running on parallel computing platforms is increasingly becoming important because of its influence on development time and resource management. However, predicting the performance with respect to parallel processes is complex for iterative and multi-stage applications. This research proposes a performance approximation approach *FiM* to predict the calculation time with *FiM-Cal* and communication time with *FiM-Com*, of an application running on a distributed framework. *FiM-Cal* consists of two key components that are coupled with each other: 1) Stochastic Markov Model to capture non-deterministic runtime that often depends on parallel resources, e.g., number of processes. 2) Machine Learning Model that extrapolates the parameters for calibrating our Markov model when we have changes in application parameters such as dataset. Along with the parallel calculation time, parallel computing platforms consume some data transfer time to communicate among different nodes. *FiM-Com* consists of a simulation queuing model to quickly estimate communication time. Our new modeling approach considers different design choices along multiple dimensions, namely (i) process-level parallelism, (ii) distribution of cores on multi-processor platform, (iii) application related parameters, and (iv) characteristics of datasets. The major contribution of our prediction approach is that *FiM* can provide an accurate prediction of parallel processing time for the datasets which have a much larger size than that of the training datasets. We evaluate our approach with NAS Parallel Benchmarks and real iterative data processing applications. We compare the predicted results (e.g., end-to-end execution time) with actual experimental measurements on a real distributed platform. We also compare our work with an existing prediction technique based on machine learning. We rank the number of processes according to the actual and predicted results from *FiM* and calculate the correlation between the actual and predicted rankings. Our results show that *FiM* obtains a high correlation in the range of 0.80 to 0.99, which indicates considerable accuracy of our technique. Such prediction provides data analysts a useful insight of optimal configuration of parallel resources (e.g., number of processes and number of cores) and also helps system designers to investigate the impact of changes in application parameters on system performance.

Additional Key Words and Phrases: Performance Modeling, Queuing Theory, Markov Model, Distributed Systems, Execution Time, Parallel Calculation, Communication Network, Prediction

ACM Reference Format:

Janki Bhimani, Ningfang Mi, Miriam Leeser, and Zhengyu Yang. 2019. New Performance Modeling Methods for Parallel Data Processing Applications. 1, 1 (January 2019), 24 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Authors' addresses: Janki Bhimani, Northeastern University, 360 Huntington Ave, Boston, MA, 02115, USA, bhimani@ece.neu.edu; Ningfang Mi, Northeastern University, 360 Huntington Ave, Boston, MA, 02115, USA, ningfang@ece.neu.edu; Miriam Leeser, Northeastern University, 360 Huntington Ave, Boston, MA, 02115, USA, mel@coe.neu.edu; Zhengyu Yang, Northeastern University, 360 Huntington Ave, Boston, MA, 02115, USA, yangzy1988@coe.neu.edu.

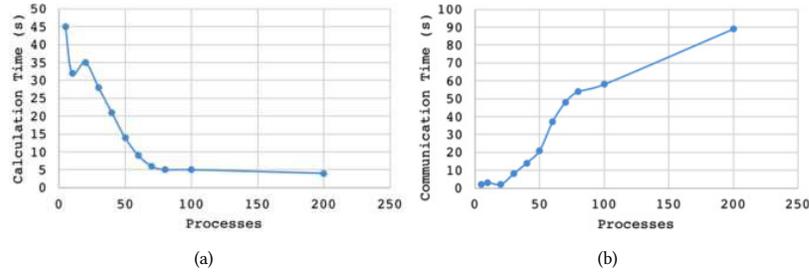


Fig. 1. Latency variation across different number of parallel processes for (a) calculation time, and (b) communication time

1 INTRODUCTION

High-Performance Computing (HPC) systems are ubiquitous in processing data for myriad applications involving huge datasets. How to achieve the best performance with an optimal configuration of parallel resources (e.g., number of processes and number of cores) is a challenging research problem. Currently, researchers run their application codes on a representative dataset, fix application parameters, and try different configurations of parallel resources to determine the optimal one. However, if we want to find optimal application parameters, then the investigation needs to consider all possible combinations of application parameters and parallel resources. Such an investigation becomes very expensive, requiring a significant amount of time and hardware resources. Besides, on parallel computing platforms, using more parallel resources does not always guarantee performance improvement. Hence, it is beneficial if we can approximate the optimal performance in terms of parallel resources and application parameters. The prediction of expected performance before the porting of an actual implementation on a hardware platform can save significant time and hardware resources spent in experimentally finding the optimal performance. The emergence of huge datasets as workloads and parallel computing has emphasized the importance of predictive analysis, and performance bottleneck identification. Designing efficient prediction tools thus becomes critically important to system designers and application programmers [24].

As motivation, we show an example with an iterative K-means clustering application running on a framework using Message Passing Protocol (MPI [19; 35]) in Fig. 1. We observe that the calculation time decreases when we have more parallel Message Passing Interface (MPI) processes; however, the time to communicate data increases. Such an observation implies that speeding up parallel calculation time may not guarantee overall application speed-up. Also, the decrease in calculation time levels off after 70 parallel processes. Increasing the number of parallel processes further consumes more system resources, but may not improve overall application runtime. Thus, the capability of predicting such an optimal point (e.g., 70 in Fig. 1) is vital to system designers for making good design choices.

In this paper, we develop a new performance modeling approach, named FiM, to estimate both computation and communication times of iterative, multi-stage data processing applications using MPI. Each node is a CPU which has multiple cores, and each core can support multiple MPI processes. One of the nodes accesses an application dataset and determines the distribution of processing among the processes of other nodes. All these nodes then perform parallel computations using multiple processes. Such a parallel phase is known as one stage in our model. At the end of each stage, all processes synchronize to decide on the termination or launch next stage. In this research, we concentrate on predicting parallel processing time of such an iterative and multi-staged application running with global synchronizations. One of the key innovations in our work is that FiM relies only on small datasets for training but can

105 predict the execution times for larger datasets. More specifically, this paper aims to answer the following questions
106 through our prediction models.
107

- 108
- 109 • Can we quickly estimate the parallel calculation and communication times of an application to identify the
110 optimal number of processes?
- 111 • Can we use small datasets as training to predict performance of applications operating in parallel on large
112 datasets?
- 113
- 114 • How does the number of processes impact calculation and communication time?
- 115 • Is the application communication or compute bound?
116
- 117

118 To answer these questions, we introduce FiM which consists of two main components: (1) FiM-Cal, and (2) FiM-Com.
119 The goal of FiM-Cal is to predict the calculation time by using a stochastic Markov model and a machine learning
120 model. The stochastic Markov model is built using the probabilistic technique to estimate the impact of an increase
121 in the number of parallel processes. We first develop the base case of the parallel paradigm and then derive a generic
122 model that applies to any number of parallel processes as well as any number of dependent stages (e.g., iterations) of an
123 application. The base case of the Markov model is calibrated using the minimum number of system parameters. The
124 machine learning model is then designed to extrapolate the calibrated parameters for the Stochastic Markov model to
125 adapt to changes in application parameters such as datasets.
126

127 The goal of FiM-Com is to predict the communication time using a set of simulation queuing models. Here our
128 motive is to get a quick estimate using a simplified prediction model. Such an estimate of communication time along
129 with calculation time can provide instant insight to users. Thus, our FiM approach can use the minimum possible
130 calibration parameters to quickly predict the expected computation and communication time as well as the optimal
131 number of processes for platform configuration. While comparing actual and predicted time, the worst prediction
132 error of the overall application runtime by FiM is observed to be less than 20%. The source code of our calculation and
133 communication time prediction model is available at GitHub (<https://github.com/bhimanijanki/FiM>)
134

135 The remainder of this paper is organized as follows. We present the two FiM components in Section 2 and Section 3.2,
136 respectively. We evaluate our models on a distributed memory platform, see Section 4. In Section 6, we discuss some
137 related work. Section 7 presents our conclusions and future works.
138

142 2 FIM-CAL: CALCULATION PREDICTION

143 In this section, we present *FiM-Cal*, an analytical approach to predicting the calculation time of an application
144 running on a distributed multi-process platform. FiM-Cal consists of two key components: a stochastic Markov model
145 and a machine learning model. We first use the stochastic Markov model to represent the computational processing of
146 an application in a parallel MPI framework. Then we design a machine learning algorithm to estimate the parameters
147 related to the system for calibrating our stochastic Markov model. This parameter extrapolation enables our model to
148 predict an application's calculation time when we have a different number of parallel processes or variable application
149 parameters (such as dataset size) without any system state instrumentation. Table 1 lists the notations used in this
150 paper for FiM-Cal. Fig. 2 shows the overall workflow of our proposed FiM-Cal. We will introduce the details of each
151 component in Fig. 2 in the remainder of this section.
152
153
154
155
156

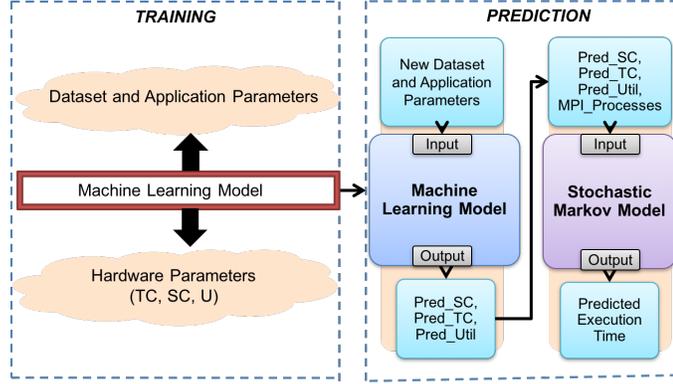


Fig. 2. Prediction procedure of FiM-Cal

Table 1. Notations used in *FiM-Cal*

Notation	Description
S_j^i	Markov chain state with i active and j passive processes
P_{s_i}	Stage completion probability of i^{th} stage
P_{ij}	State transition probability of moving from i active to j active processes
P_{act}	Probability P_{11} of 1 process model
P_{p2a}	Probability P_{01} of 1 process model
P_{a2p}	Probability P_{10} of 1 process model
P_{pass}	Probability P_{00} of 1 process model
F	Frequency (GHz)
TC	Total cycles
SC	Total stall cycles
U	Utilization per process
T_i	Total time taken by stage i
α	Sensitivity constant
β	Regression constant
y_i	Dependent variables
\vec{X}_i	Vector of independent variables

2.1 Stochastic Markov Model

Our stochastic Markov model is designed to model computational processing for an application running on a system with parallel multi-core CPUs deployed using MPI. Such a stochastic model allows us to capture a non-deterministic runtime that often depends on parallel resources, e.g., number of processes. If there exists a global synchronization call in an application, then all processes wait until the barrier. The processing of an application is partitioned into multiple stages with respect to this global synchronization, such that each stage corresponds to a parallel phase until all processes have completed their tasks and are in an active state to proceed to the next stage. In this section, we first

introduce our base case which models a single stage for a single-process system and then show its extension to the generic case with multiple processes and multiple stages.

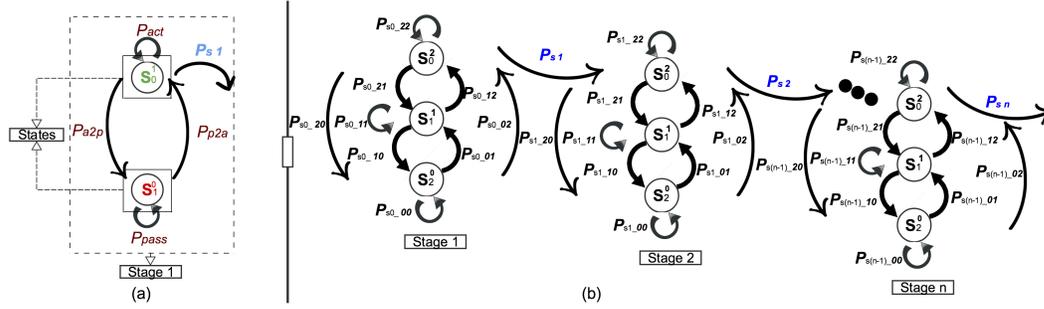


Fig. 3. Modeling (a) Base Case: single process with single stage, (b) Generic Case: two processes with multiple stages

2.1.1 Base Case. The base case model is built to represent a single process for a single stage application, see Fig. 3 (a). In our model, each process is considered to be either in *active* or *passive* state. As shown in Fig. 3(a), when only a single process runs in the computing platform, we have two states for a stage such that state S_0^1 represents that one process is active while S_1^1 represents that one process is passive. We also introduce transition probabilities (e.g., P_{act} , P_{a2p} , P_{p2a} , P_{pass}) of switching between two states or staying in the same state, as well as the stage completion probability (P_{s1}) of transferring from one stage to another. In the active state, the process performs constructive work and typically changes from the active state to the passive state when it is blocked by an event that would create a latency stall. Such a latency stall might be caused due to a cache miss that takes many CPU cycles. In this work, we do not model memory latencies, contentions, inter-dependencies and deadlocks individually for each process but rather treat the combined effect as a process remaining passive.

The probabilities in the base case model are parameterized by instrumenting the system details, which is further used to derive the probabilities of the generic case. In order to parameterize the probabilities of the base case, we use the `perf` tool [16] to instrument the required data, including the hardware clock rate (F), system CPU utilization factor per process (U), total number of cycles required for execution (TC) and stall cycles (SC). In particular, we run an application with a single stage on a single process and use the `perf stat` command to collect and report the required data as listed above.

In the base case (i.e., single-process and single-stage), the probability to remain in the active state (P_{act}) is primarily determined by the proportion of time that the process is performing useful work. Therefore, we use Eq. 1 to get P_{act} ,

$$P_{act} = U \quad (1)$$

where U is utilization per process. As shown in Fig. 3 (a), when the process is in an active state (i.e., S_0^1), there are three possible events for its next transition: (1) remain in S_0^1 with probability P_{act} , (2) transition to with probability P_{a2p} , and (3) complete the stage with probability P_{s1} . Now, if there are TC total cycles to be processed for the given dataset, then processing is completed only after completing the last cycle. This gives the probability of completion as $1/TC$.

Probability P_{a2p} for the process to transit to the passive state can then be calculated as shown in Eq. 2.

$$P_{a2p} = 1 - (P_{act}) - (1/TC) \quad (2)$$

The probability of the process remaining passive (P_{pass}) is primarily determined by the ratio of stall cycles (SC) to total cycles (TC) as shown in Eq. 3.

$$P_{pass} = SC/TC \quad (3)$$

We can determine the probability of switching from passive to active (P_{p2a}) by applying the control flow equation to the passive state (S_1^0) as shown in Eq. 4.

$$P_{p2a} = 1 - (P_{pass}) \quad (4)$$

We finally get the stage completion probability (P_{s1}) by applying the control flow equation to the active state (S_0^1) as shown in Eq. 5.

$$P_{s1} = 1 - (P_{act}) - (P_{a2p}) \quad (5)$$

Note that, for the base case with one stage and only one possible active phase, P_{s1} is the same as $1/TC$ because all active cycles can be spent only in one active state (S_0^1). Later, in generic cases, we discuss the calculation of P_{s1} , which is then not the same as $1/TC$.

2.1.2 Generic Cases. Now, we consider generic cases where we can have multiple processes operating on an application with multiple stages. This generic behavior can be modeled as an extension of the base case. The processing of an application may have multiple inter-dependent parallel stages. For example, an iterative application with 500 iterations can be divided into 500 parallel stages such that each stage represents an iteration and is entered only after the completion of all prior stages. Thus, the first stage corresponds to the parallel calculation phase by all processes in the first iteration and is followed by the remaining stages in the same order. In general a simple non-iterative and single stage application can be treated as having 1 iteration and 1 stage while predicting its runtime using our proposed technique.

Fig. 3 (b) shows an underlying Markov model for an application with two processes using n stages, each with 2 parallel processes. The entire workflow of an iterative, multi-stage, multi-process application can be mapped with a chain of n parallel stages, and $P_{s1}, P_{s2}, \dots, P_{sn}$, are the completion probabilities for all n stages, see Fig. 3 (b). Note that these stage completion probabilities are non-uniform and dependent on all the completion probabilities of prior stages as well as intra-state transition probabilities of that stage. Also for every stage, all its state transition probabilities (P_{ij}) depend on the completion probability of the prior stage. Thus, the value of P_{ij} in a stage is different from that of P_{ij} in another stage even for the same i and j . Furthermore, a single stage can only complete when all of its processes are active, i.e., not being blocked by any events. A calculation phase of an application is completed when tasks assigned to all processes are completed in the last stage. Thus, the completion probability of an application is P_{sn} .

To model an iterative, multi-stage paradigm with multiple processes, we use multiple states within each stage to represent activities (active or passive) of all processes. Consider t processes with i active processes and j passive processes, where $0 \leq i \leq t$, $0 \leq j \leq t$ and $t = i + j$. Each stage consists of a total of $M = t + 1$ states. Thus, the transition probabilities of jumping from any one of these M states to other states or itself can be divided into 3 types: 1) probability to remain in the same state (e.g., P_{22}, P_{11} and P_{00} in Fig. 3 (b)), 2) probability to increase active processes (e.g., P_{01}, P_{12}, P_{02} in Fig. 3 (b)), and 3) probability to increase passive processes (P_{10}, P_{21}, P_{20} in Fig. 3 (b)). Given M states, we have M probabilities to remain in the same state, $\sum_{i=1}^{M-1} i$ probabilities to increase active processes,

and $\sum_{i=1}^{M-1} i$ probabilities to increase passive processes. Thus, the total number of probabilities to be calculated for a single stage with M states is equal to $M + 2 \sum_{i=1}^{M-1} i = M^2$.

2.1.3 Solving the Generic Model. We solve such a generic Markov model and derive its probabilities by relating them to the preliminary transition probabilities of the base case. That is, once we have transition probabilities for $M = 2$ (base case), we can calculate all probabilities for a generic case with $M > 2$. In [27], a mathematical relation between the transition probabilities of a Markov model with two states and a Markov model with more than two states has been derived. We use their method to relate transition probabilities of the cases with $M = 2$ and $M > 2$. Initially, base case probabilities are calculated.

Specifically, let us consider to derive generic model probabilities for an application with $M = 3$ from the base case probabilities of $M = 2$. For $M = 3$, the application must have two parallel processes (see Fig. 3 (b)). At any given time during the execution, there are three states ($M = 3$): 1) both these processes can be active (state S_0^2), 2) one process can be active and the other passive (state S_1^1), 3) both can be passive (state S_2^0). The probability of 2 active processes to continue as active is $P_{22} = P_{act} * P_{act}$. Similarly, the probability of 2 passive processes to continue as passive is $P_{00} = P_{pass} * P_{pass}$. If one process that was active in previous stage remains active and the other process that was passive remains passive, or if the active process becomes passive and the other process that was passive becomes active, then in both of the above cases the model stays in state S_1^1 . Thus, the probability of such state transition is $P_{11} = (P_{act} * P_{pass}) + (P_{a2p} * P_{p2a})$. P_{01} is state transition probability from S_2^0 (i.e., both the passive processes) to S_1^1 (i.e., one active process and the other passive process).

Same as above, the state transition probabilities for any M can be derived. These derived equations can be mathematically reduced to a more generic form. Eq.s 6 to 8 give the state transition probabilities for $M > 2$, where i corresponds to the number of active processes in the previous state and j corresponds to the number of active processes in the targeted state. For example, in Fig. 3 (b), P_{21} indicates the state transition probability of moving from a state with 2 active processes (S_0^2) to a state with 1 active process (S_1^1). Substituting appropriate i and j in Eq.s 6 to 8 for $M = 3$, all the probabilities explained and derived above for 2 processes such as P_{00} , P_{11} , P_{22} , P_{01} , P_{02} , P_{12} , P_{10} , P_{20} and P_{21} can be obtained. These equations represent the stochastic process of a Markov chain and can be calculated by mathematical induction after solving the Markov chain with a finite number of states. Particularly, as shown in Eq.s 6 to 8, we use the probabilities (P_{act} , P_{a2p} , P_{pass} and P_{p2a}) that are obtained by the base case (Sec. 2.1.1) to calculate the state transition probabilities in generic cases.

If $i == j$,

$$P_{ii}(t) = \sum_{k=0}^{\min\{i,t-i\}} \binom{i}{k} \cdot \binom{t-i}{t-i-k} \cdot (P_{act}^{i-k}) \cdot (P_{a2p}^k) \cdot (P_{pass}^{t-i-k}) \cdot (P_{p2a}^k) \quad (6)$$

If $i < j$,

$$P_{ij}(t) = \sum_{k=0}^{\min\{i,t-j\}} \binom{i}{k} \cdot \binom{t-i}{t-j-k} \cdot (P_{act}^{i-k}) \cdot (P_{a2p}^k) \cdot (P_{pass}^{t-j-k}) \cdot (P_{p2a}^{j-i+k}) \quad (7)$$

If $i > j$,

$$P_{ij}(t) = \sum_{k=0}^{\min\{j,t-i\}} \binom{i}{i-j+k} \cdot \binom{t-i}{t-i-k} \cdot (P_{act}^{j-k}) \cdot (P_{a2p}^{i-j+k}) \cdot (P_{pass}^{t-i-k}) \cdot (P_{p2a}^k) \quad (8)$$

Additionally, after capturing the state transition probabilities of the first stage, we calculate the stage completion probability P_{s1} using Eq. 9 and then use P_{s1} as an incoming probability for calculating the state transition probabilities of stage 2, and so on. This chaining process captures an iterative and multi-stage application running with multiple

processes. Finally, P_{s_n} for the n^{th} stage is calculated by Eq. 9.

$$\begin{aligned}
 P_{s_n} &= P(X_s = n \mid X_{s-1} = n - 1) \\
 &= 1 - \sum_{j=0}^{j=i} (P_{ij} \mid X_{s-1} = n - 1) \\
 &\text{for } i = \text{Max}(\#Processes)
 \end{aligned} \tag{9}$$

We further use Eq. 10 to calculate the time (T_n) spent in performing parallel calculation for n stages, given the completion probability (P_{s_n}) and CPU frequency (F).

$$T_n = \frac{1}{(P_{s_n}) \cdot (F)} \tag{10}$$

Consequently, our stochastic Markov model can predict the computation time required to process any particular dataset using different levels of parallelism such as different number of processes in MPI. Next, we present our machine learning technique which assists to extrapolate the data (such as, F , U , TC and SC) required for calibrating the base case model.

2.2 Machine Learning Model

Our stochastic Markov model allows us to predict the calculation time of an application when we have a different number of processes in the system. However, the required hardware parameters (i.e., TC , SC , U) need to be instrumented for every new dataset and a new setting of application parameters. This limits the scope of the model to predict for a particular set of datasets and fixed application parameters. Most analytical models suffer from this lack of flexibility. Therefore, we develop our machine learning model to avoid additional instrumentation for a new dataset or a new set of application parameters. For a new application, we need first to train to derive a new model. However, for new sets of application parameters and new datasets of the same application, the derived model can be used to predict the runtime. To reduce the complexity of the machine learning model, we also assume that the application calculation time is dependent on the fewest possible hardware parameters. Our evaluation results are shown in Section 4 demonstrate the feasibility of this assumption by showing the fairly accurately predicted results obtained by our approach which is good to give a quick approximation. Here, we introduce a two-stage machine learning model that emulates hardware behaviors without performing actual instrumentation for required hardware-related data. Such a hybrid emulation of hardware is the key to allowing the approach to be able to predict parameters for datasets with sizes much larger than those of the training datasets.

2.2.1 Regression Mapping. The focus of regression is to find the relationship between a dependent variable (such as the hardware parameters which we want to emulate) and one or more independent variables (such as application parameters and datasets). This analysis estimates the conditional expectation of a dependent variable given values of all related independent variables. We find that the generalized linear regression model performs the best when compared to others (quadratic, Poisson model, and gradient decent) for modeling all desired hardware parameters (U , TC and SC). We show the validation of a linear regression model in Sec.4.3.

The linear regression equation for learning variable y_i is shown in Eq. 11, where \vec{X}_i is a vector of p independent variables related to application parameters and datasets, $\vec{\beta}_i$ consists of a vector of $p + 1$ constants, and n is total number of scalar dependent variables. Suppose for the K-means application, elements of \vec{X}_i would consist of the number of desired clusters (K), the number of iterations (I) and size (N). For our model, we have three scalar dependent variables,

i.e., U , TC and SC , which can be predicted after building this linear regression model. Thus, we have three equations for y_1 , y_2 and y_3 with $n = 3$.

$$\begin{aligned} y_i &= \beta_{i0} + \beta_{i1}x_{i1} + \dots + \beta_{ip}x_{ip} \\ &= \vec{\beta}_i(1 + \vec{X}_i^T), \text{ for } i = 1, 2, \dots, n \end{aligned} \quad (11)$$

This linear regression model is used to find values for constants $\vec{\beta}_i$ using the training data for which both dependent variables and independent variables are known. We obtain the regression curve and regression constants $\vec{\beta}_i$ by building our machine learning model in MATLAB.

Algorithm 1: Calibration of α

```

1 Input:  $\epsilon$ ,  $\tau$   $y_i$ ,  $\vec{X}_i$ , Output:  $\alpha$ 
2 Initialize:  $\vec{\beta}_i$ ,  $TC$ ,  $SC$ ,  $U$  using regression mapping,  $\alpha = 0$  and iter = 0
3 if  $0 \leq U \leq 1$ ,  $TC > 0$ ,  $SC > 0$ ,  $SC < TC$  then
4   Predict comp. time using stochastic Markov model
5   Calculate RMS error (Actual, Predicted)
6   if iter == 0 then
7     Calculate error (Actual - Predicted)
8     Decide OP = + or -, depending on positive or negative error
9   if RMS error  $< \tau$  then
10    return
11  else
12     $\alpha = \alpha \{OP\} \epsilon$ 
13    iter ++
14    Calculate TC, SC and U using Eq. 12
15    goto line 3
16 else
17   Neglect bad values
18   goto line 4

```

2.2.2 Iterative Improvement Model. We found non-negligible errors between the actual and predicted calculation times when we pair the linear regression model described above with our stochastic Markov model, to predict calculation time of large datasets based on small training datasets. To handle this issue, we develop an iterative improvement model which uses a sensitivity parameter α to tune the constant factors $\vec{\beta}_i$ with respect to the predicted results (i.e., calculation time) of our stochastic Markov model as shown in Eq. 12. Note that in this equation the constants in vector $\vec{\beta}_i$ are obtained from the regression between hardware parameters and application parameters, but constant α is obtained by using both Markov model and regression model as described in Algorithm 1.

$$\begin{aligned} y_i &= \alpha(\beta_{i0} + \beta_{i1}x_{i1} + \dots + \beta_{ip}x_{ip}) \\ &\text{for } i = 1, 2, \dots, n \end{aligned} \quad (12)$$

Initially, we use hardware parameters (U , TC and SC) and application parameters of training data with regression mapping (Eq. 11) to obtain constants of vector $\vec{\beta}_i$. The hardware parameters are passed to our stochastic Markov model to predict calculation time. In the first iteration, the error between actual and predicted time is used to decide the adjust

direction of α , see lines 6 to 8. That is, if the actual value is greater than the predicted one, then the algorithm increases α and vice versa. In the following iterations, Algorithm 1 increases or decreases α by a small value ϵ (e.g., $\epsilon = 1e-5$) and the hardware parameters (U , TC and SC) are predicted using constants of vector $\vec{\beta}_i$ and α with Eq. 12 (see line 14). Then the computation time is predicted using U , TC and SC as the inputs to our stochastic Markov model (see line 4). The adjustment process continues until the root mean square (RMS) error becomes smaller than a predefined threshold (e.g., $\tau = 0.01$), see line 9. The value of τ needs to be set such that we avoid underfitting as well as overfitting the training data. From our experiments, we observe that $\tau = 0.01$ is a good value on an average across different choices of applications and training datasets. Thus, the algorithm adjusts the value of α until the predicted calculation time becomes close to the actual time. We tune α while training the model and then use the exact tuned α value throughout the prediction stage. Note that our machine learning model is used to calculate the constants of vector $\vec{\beta}_i$ and α in the training phase, which are after that used for extrapolation of hardware parameters.

In summary, Fig. 2 shows the overall procedure of our prediction model *FiM-Cal*, which includes the training and prediction phases. In the prediction phase, our machine learning model extrapolates the dependent variables (such as hardware parameters - SC , TC , U) for new datasets and new sets of application parameters. These predicted hardware details can then be used as an input to our stochastic Markov model to predict the calculation time.

3 FIM-COM: COMMUNICATION PREDICTION

Apart from computation time [7; 8], applications spend some time in the communication of data to various processes. With the increase in the number of processes, this communication time keeps increasing [37]. Therefore, it is essential to roughly estimate communication time, i.e., the runtime of *data transfer*. For applications running on a distributed multi-process system, the data transfer time for processes lying on the same node is different from that between processes lying on different nodes. The framework considered in this work consider all inter-node communication.

3.1 Communication Patterns

It is non-trivial to predict communication time, due to various communication patterns and the non-deterministic latency of the communication network. There exist many, more complex models to predict accurate network communication time [10; 20; 23; 30], but here our motive is to get a quick estimate using a simplified prediction model. In this work, we investigate three types of collective communication patterns including downlink (scatter and broadcast) and uplink (gather). We build different queuing models to capture those patterns when running a data processing application such as K-means clustering on a multi-process system with MPI. Because the data transfer time along with computation time is significant, our model aims to provide very quick rough estimates of communication time to make fast decisions.

Specifically, we consider communications from the one to many nodes as the *downlink*. Such a downlink communication can have either the scatter pattern or the broadcast pattern. Under the scatter pattern, the data is distributed among the multiple nodes by the one node such that each node gets a unique part of the data. Under the broadcast pattern, the data is broadcast to all nodes, and each node receives the same copy of the data. The scatter communication is usually undertaken by communicating data to one after another process in MPI. We also consider communications from multiple nodes to the one node as the uplink gather pattern. When gathering, each node sends their own data to a shared buffer, and the managing node reads those data one-by-one from the buffer.

521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572

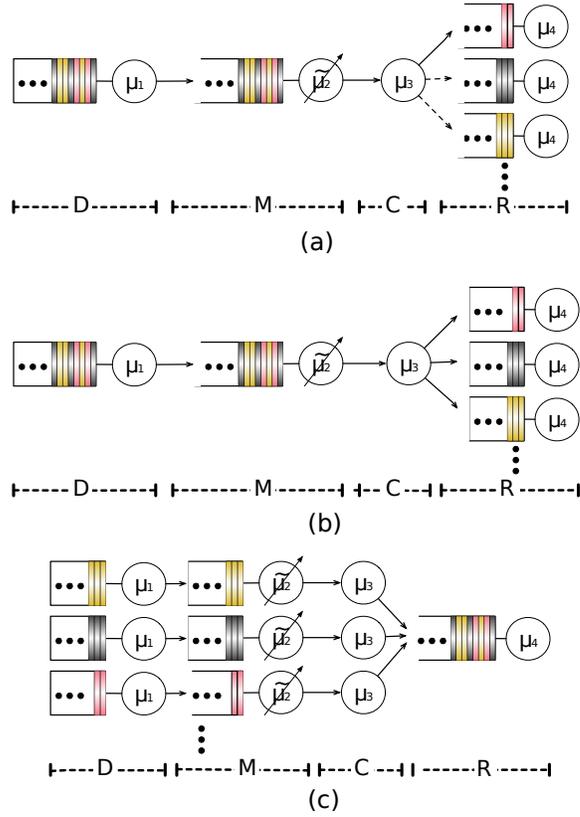


Fig. 4. Queuing models for the (a) scatter, (b) broadcast, (c) gather communication patterns

3.2 Communication Models

We develop a set of queuing models to capture these three collective communication patterns. Such queuing models can give a high-level abstraction of real communication systems in terms of packet arrival rates, delay/waiting time and packet transfer rate, which helps to estimate the overall data communication time. Note that we try to keep these queuing models simple to enable fast estimation.

Fig. 4 presents our queuing models to represent the scatter, broadcast and gather patterns, respectively. Typically, communication time depends on various network properties such as initialization cost, maximum network bandwidth, network load and the amount of metadata (e.g., headers, acknowledgments, and flags). Therefore, in each of these models, we use four components, Data transfer (D), Metadata (M), Cost of network initialization (C) and Receive (R) in Fig. 4, in series to capture different properties of network communication. Each job in the queue of D, M, C and R represents a data packet e.g., a pixel for K-means clustering.

First, D is used to capture the actual data transfer rate through the communication network, where mean service rate μ_1 indicates the available network bandwidth. We assume that all data packets are available in memory. Thus, all jobs are assumed to wait in the queue of part D before data transfer begins.

The second part M, in the models of Fig. 4 captures the effect of metadata on overall communication time. Such metadata can include headers, acknowledgments, addresses, offsets, padding and flags that need to be transferred over the network along with the original data packets as part of the network protocol. The time consumed by such metadata transfers depends on the number of actual data packets and instantaneous network load. To capture this, we design component M with a variable mean service rate $\tilde{\mu}_2$, which depends on the instantaneous queue length of part M. The arrivals to this queue are the departures from part D.

Part C in the model represents network initialization. Usually, the cost for network initialization is a network latency that is added to overall communication time. We use a delay server with mean service rate μ_3 to emulate such network latency. The last part of our models (R in Fig. 4) is used to estimate the data fetching time on the receiver side. We use part R to capture the receiving, while the first three parts (D, M and C) capture the sending. We assume that the receiving process is homogeneous across different nodes. Therefore, multiple servers and queues are used to represent multiple receivers, and the mean service rate of each server is μ_4 .

We present the corresponding models for scatter and broadcast in Fig. 4 (a) and (b), to predict communication times for the downlink, i.e., from the one node to multiple nodes. Under the scatter pattern, the one node sends different pieces of data to other nodes one-by-one. We use the round robin policy to distribute jobs into queues in part R, and thus there is only one solid arrow in R to indicate data transfer to one process at a time, see Fig. 4 (a). In Fig. 4 (b), there are multiple solid arrows to R, which represents that we duplicate jobs (data points) and distribute them to all queues of R in parallel, to capture the broadcast pattern. The third model, shown in Fig. 4 (c), represents the gather pattern of the uplink communication, where multiple nodes send their data to one node. We have multiple servers running in parallel in parts D, M and C. The First-In-First-Out (FIFO) discipline is used in the last queue to emulate data fetching by the managing node.

3.3 Communication Model Calibration

The main tasks in communication model calibration are 1) to determine an appropriate queuing model, and 2) to derive service processes for all servers in the model based on its collective communication pattern. Our models only need to be calibrated when we have a new application or run on a new hardware platform. We then leverage the knowledge of actual transfer processes on a new platform to calibrate the service rates of all servers in the models. The calibrated models can be used directly to predict communication times for new datasets, new application parameters (e.g., number of iterations in K-means) and different number of MPI processes.

In particular, we calibrate μ_1 using the maximum network bandwidth obtained from the provided network configuration of a cluster (e.g., 10 Gb/s backplane). We observe from conducting experiments on multiple applications for both 10 Gb/s and faster 56Gb/s TCP/IP backplane that the actual data transfer rates are often within 90% to 100% of the maximum network bandwidth, with a uniform distribution. Therefore, we calculate the lower bound (ζ_{min}) and the upper bound (ζ_{max}) of transfer rate per unit job (i.e., each data point) and derive the service rate μ_1 in part D using Eq. 13.

$$\mu_1 = Uniform(\zeta_{min}, \zeta_{max}) \quad (13)$$

To get the transfer time of the metadata, we first measure the actual total communication time for sending a unit job (a single data point) and then deduct the measured runtime of the other three parts, i.e., network bandwidth, network initialization, and receiver's data fetch time. We find that the derived transfer rates of metadata are logarithmic to the instantaneous network load (refer Section 4.3). Therefore, we use Eq. 14 to generate $\tilde{\mu}_2$ for metadata where χ represents the number of jobs currently waiting in the queue of M. The tilde on μ_2 represents the variable mean service rate of

component M , which depends on the instantaneous queue length of part M (χ) to symbolize variance as explained above.

$$\tilde{\mu}_2 = \log(\chi) \quad (14)$$

In part C , μ_3 captures the network initialization latency using a delay server. Specifically, the sending node first activates itself (`comselfcreate`) and then activates the receiving nodes (`comworldcreate`). The MPE tool [19] is used to collect the `mpilog` log files. We observe that the initialization latency follows an exponential distribution (refer Section 4.3). Thus, we use Eq. 15 to get the service rate μ_3 , where η is the mean of the service rate distribution.

$$\mu_3 = \exp(\eta) \quad (15)$$

Finally, we observe that the data receiving rate is exponentially distributed (refer Section 4.3). Additionally, we consider a multiplicative factor ϕ to model the number of receivers that simultaneously fetch data. For example, ϕ is equal to one under the scatter pattern since only one node can receive data at any time. For the broadcast pattern, ϕ is equal to the number of nodes. Thus, Eq. 16 is used to draw μ_4 for all the servers in part R , where ψ is the mean service rate distribution.

$$\mu_4 = \phi * \exp(\psi) \quad (16)$$

4 EVALUATION

Table 2. Platform configurations labelled as C1 to C5 (2-E52670 - Two multi-core, hyper-threaded Intel Xeon E5 2670 CPU's @ 2.60 GHz and 256 GByte of RAM) (2-E52650 - Two multi-core, hyper-threaded Intel Xeon E5 2650 CPU's @ 2.00 GHz and 128 GByte of RAM)

	C1	C2	C3	C4	C5
CPU	2-E52670	2-E52670 2-E52650	2-E52670 4-E52650	2-E52670 6-E52650	2-E52670 8-E52650
Cores	32	64	96	128	160
Network	10 Gb/s Ethernet backplane TCP/IP				
Shared FS	NFS				
OS	Linux				

We evaluate our FiM prediction approach (a combination of FiM-Cal and FiM-Com), by comparing the predicted results (e.g., end-to-end execution time) with actual experimental measurements on a real distributed platform. We also compare our prediction approach with an existing work, named RBASP [5], which is a regression-based approach to extrapolate execution time. In our evaluation, we consider the end-to-end execution time of data processing applications as the sum of the runtime spent in communicating the required data to parallel processing units, performing the calculations in parallel, and transferring the results back to the managing node. We assume that no overlap exists between the calculation and communication in the application implementation. We use the Discovery Cluster at Northeastern University [1] to build our experimental platform. Table 2 describes five parallel platform configurations we used in our evaluation, where each CPU belongs to different nodes.

We evaluate our approach with six NAS Parallel Benchmarks (NPB - version *NPB3.3.1-MPI*) [2], with the large size dataset of *Problem Class C*. Table 3 lists the six benchmarks we used in our evaluation. For each benchmark, we train our model using three small size datasets of *Problem Class S*. Note that we use the trained model to predict the performance for datasets of *Problem Class C* which are much larger than the datasets of *Problem Class S*. We reprogram the NPB benchmarks to implement iterative, multi-stage paradigm versions in MPI using C. The time spent for computation in all iterations (or multiple phases) is the total computation time. For each iteration, we measure the time from the start

Table 3. NPB Benchmarks

BT	Block Tri-diagonal solver	Compute Bound
EP	Embarrassingly Parallel	Compute Bound
SP	Scalar Penta-diagonal solver	I/O Bound
LU	Lower-Upper Gauss-Seidel solver	I/O Bound
IS	Integer Sort	Memory Bound
CG	Conjugate Gradient	Memory Bound

of parallel processes to the completion of all the processes. We also evaluate FiM with two iterative data processing applications: K-means [6] and Pagerank. For each of these applications, we run experiments on 15 different datasets and choose one dataset as a representative to show the results, i.e., N_L (the large dataset with 13 million data points). For both applications, we train our model using three small datasets, i.e., N_S (the small datasets each with 3 thousand data points).

K-means Clustering (KM): Our K-means clustering implementation [6] takes color images as input datasets. We cluster pixels in an image based on five features, including three *RGB* channels and the position (x, y) of each pixel. We choose random data points to initialize each cluster centroid and then use the Euclidean distance to calculate the nearest cluster for each data point. The parameters of K-means include a number of desired clusters (K), number of iterations (I), and size of input dataset (N).

Pagerank (PR): The Pagerank application takes a network of directed vertexes and edges as an input dataset. The output of the Pagerank application is a probability distribution representing the weights of each vertex (page). We choose a damping factor of 0.85 and initialize all vertices (pages) with the same probability weights. The parameters of this application include a number of vertexes (V), number of iterations (I), size of input dataset (N) and network nature (dense or sparse).

4.1 Performance Evaluation

In our evaluation, we consider a regression-based approach named RBASP [5] to compare with our FiM approach. RBASP is well known for its simplicity and accuracy in extrapolating execution time of multi-process applications. In our FiM approach, we combine our stochastic Markov modeling with machine learning regression to predict calculation time and use our queuing models to predict communication time. Our prediction model with this combination of popular techniques helps to predict accurately in most cases and gives a quick estimation. We choose RBASP to compare our prediction accuracy because similar to FiM, it also gives a quick prediction and can estimate runtime of datasets larger than the training datasets. RBASP is a pure regression-based approach, unlike FiM it directly estimates total time (calculation + communication) using regression. RBASP model predicts the execution time (y) of a given parallel application on p processes by using several instrumented runs of an application on q processes, where $q \in \{1, \dots, p_0\}$ and $p_0 < p$. By varying the values of independent variables (x_1, x_2, \dots, x_n) , this model aims to calculate coefficients $(\beta_0, \dots, \beta_n)$ by the linear regression fit for $\log_2(y)$ (Eq. 17), where $g(q)$ can be either a linear function or a quadratic function.

$$\log_2(y) = \beta_0 + \beta_1 x_1 + \dots + \beta_n x_n + g(q) \quad (17)$$

While reproducing the RBASP model, we use $p_0 = 1, 2, 4$ as the training set and predict the performance with two forms of $g(q)$ function as suggested in [5]. The RBASP approach directly predicts the execution time using regression, which requires performing training with data points processed for a different number of multiple processes. In contrast,

FiM extrapolates the hardware parameters for a given computing platform and then uses these hardware parameters as the inputs to the stochastic Markov model for predicting execution time for a different number of processes. FiM does not need to be trained again when we change the number of processes used in a given computing platform.

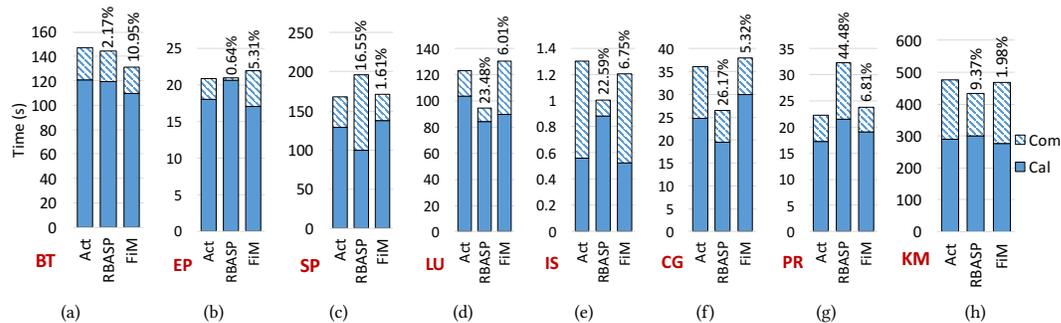


Fig. 5. Actual and predicted execution time using FiM and RBASP with the relative prediction error listed on top of each bar

Fig. 5 shows the predicted results using RBASP and FiM for six NPB benchmarks and two iterative data processing applications. We run these experiments on the C5 platform (see its configuration in Table 2), using the actual optimal number of processes listed in Table 4. As shown in Fig. 5, our FiM approach achieves a pretty good agreement between the predicted and actual results across all the six benchmarks and two applications. We also observe that RBASP has lower relative prediction error than FiM for only BT and EP. Both BT and EP are compute-intensive benchmarks, so a pure regression technique is sufficient to predict their execution time. However, prediction using only regression is not good enough for I/O and memory intensive applications as observed from RBASP prediction results in Fig. 5 for the rest of the applications. As shown in Fig. 5, for all the remaining benchmarks and applications, FiM performs better. Many real-time applications are I/O or memory intensive, for whom simple regression model is not sufficient to give accurate predictions. For the results shown in Fig. 5, the calculation time reported includes all background delay that occurs due to I/O latency to access the data that is required to perform the computations. The time required for data transfer among different parallel processes is communication time. We also observe that the prediction error of FiM remains less than 20% for individual prediction of calculation time and communication time. Furthermore, RBASP's prediction is limited to the fixed application parameters on which the model is trained because if the application parameters are changed, then RBASP needs to be re-trained. In contrast, FiM can predict execution time without any prior training for new application parameters because we do not regress the execution time directly, but instead, we train our model to learn the change in system counters like total cycles, and stalled cycles when application parameters are changed. Thus, this also advanced our approach to avoid over-fitting to the execution time of the training datasets.

Table 4 lists the best, average and worst prediction errors as well as the actual and predicted optimal numbers of processes using RBASP and FiM. Apart from having a lower average error, also having a tight prediction error range is important for these prediction approaches because such a range can be used to provide a quick approximation before conducting actual experiments. We observe that FiM has a relatively tight prediction error range from the best to the worst, compared to RBASP. Despite the lower prediction error under the best case, RBASP obtains higher prediction errors in the worst case for all benchmarks and applications. This is because the pure regression model used in RBASP

Table 4. Summary of results for all applications (Rel. Er.-Relative Error) (Act.-Actual) (App.-Application) (Opt.-Optimal)

App.	Best Rel. Er. %		Average Rel. Er. %		Worst Rel. Er. %		Opt. # MPI Process		
	RBASP	FiM	RBASP	FiM	RBASP	FiM	RBASP	FiM	Act.
BT	1.98	1.04	2.17	10.95	66.45	18.07	136	360	360
EP	0.44	2.13	0.64	5.31	76.74	15.32	309	252	248
SP	1.98	0.21	16.55	1.61	94.21	19.14	212	64	56
LU	0.14	0.06	23.48	6.01	79.15	26.88	82	28	20
IS	1.73	1.3	22.59	6.75	90.94	34.36	70	70	70
CG	2.63	2.99	26.17	5.32	77.65	40.73	156	102	102
PR	2.02	0.91	44.48	6.81	57.65	31.46	21	56	64
KM	1.61	0.42	9.37	1.98	30.43	10.87	64	192	192

has poor adaptability to changes in the values of attributes (e.g., number of processes). FiM provides tighter error bounds, which is very important for such a quick estimation modeling technique. We further observe that the optimal number of processes predicted by FiM is very close to the actual one. We also rank the number of processes according to the actual and predicted results from FiM and calculate the correlation¹ between the actual and predicted rankings. We obtain a high correlation in the range of 0.80 to 0.99, which indicates considerable accuracy of our FiM estimation technique.

We further use the Chi-square goodness-of-fit test [29] to evaluate how well the predicted runtime distribution matches with the actual measured runtime distribution with varying application parameters and the number of processes. Our null hypothesis is that the predicted data are not consistent with actual distribution. The significance level of our test is 0.1, meaning that the deviation of the predicted value is not more than 10%. The p -value² that we obtain for our test is equal to 0.08. Since the P -value (0.08) is less than the significance level (0.1), we cannot accept the null hypothesis. This indicates reasonably high confidence of less than 10% error between the actual and predicted values.

4.1.1 Individual Prediction Results. We carefully study the effectiveness of our model for predicting execution time with the increasing number of processes. Fig. 6 shows the results for FiM per data point (e.g., per pixel in K-means) when operated with a dataset consisting of 13 million data points after getting trained using three datasets with less than 1000 data points. These experiments are performed using the C5 hardware configuration of Table 2. We conduct all experiments for 1000 times and measure the average runtime as well as the minimum and maximum actual runtime obtained in these 1000 runs which are shown by the range bars. From Fig. 6 (a), we observe that the actual calculation time keeps decreasing when we have more simultaneous processes running. We also observe that 192 is the optimal number of processes. Increasing the number of processes further above 192 does not give any performance improvement. Our model can accurately predict this optimal number of processes.

Fig. 6 (b) shows the results of FiM-Com for the scatter pattern. For scatter of "n" data points over "p" processes, "n/p" data points need to be transferred to each process. For scatter, each process receives a unique subset of the dataset and thus, with the increase in the number of processes, the number of data points to be transferred in total does not increase. Hence, the runtime of the scatter communication depends mainly on the size of input datasets and not on the number of processes. So, we evaluate the FiM-Com scatter model under different datasets of different sizes. The

¹A correlation between actual and predicted ranks describes the degree of agreement between them. Correlation ranges between -1 and 1 with 1 being the best; higher correlation signifies better accuracy of predicted results.

²The p -value is a statistical measure of the deviation of the actual distribution from the hypothesis.

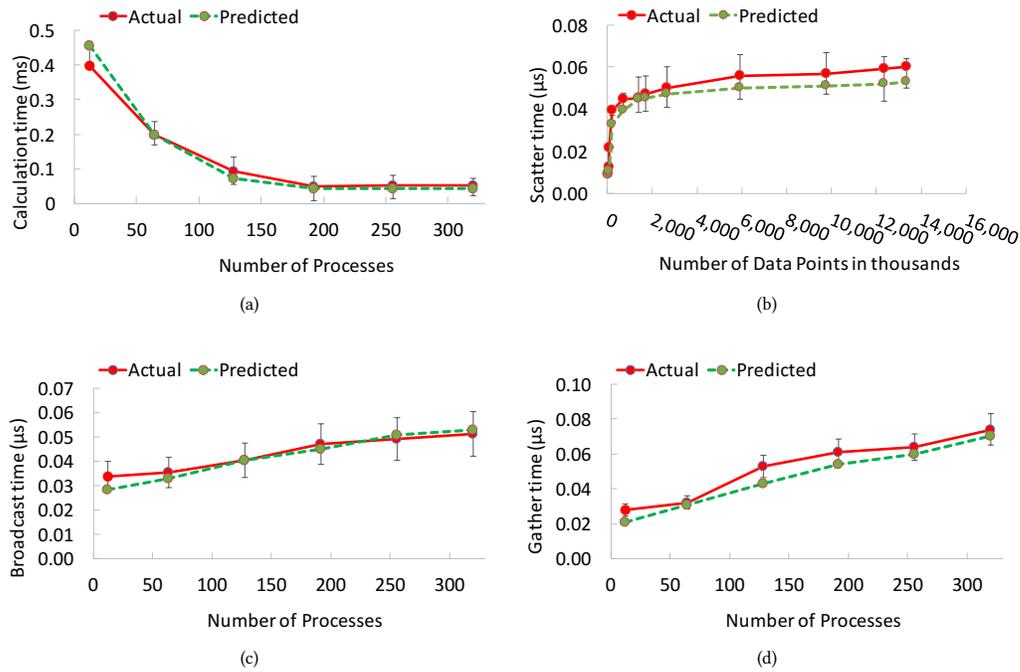


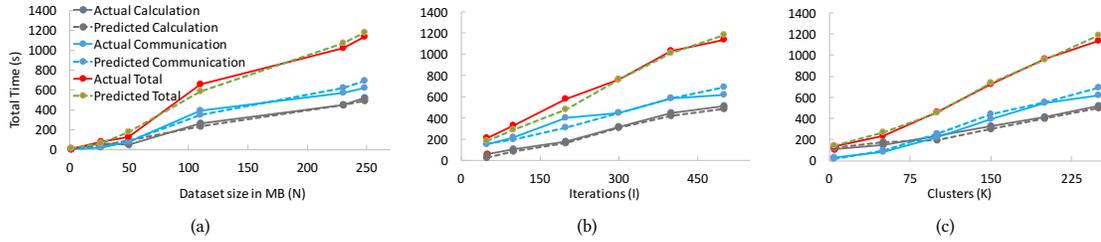
Fig. 6. Actual and predicted time - (a) Calculation (b) Scatter (c) Broadcast (d) Gather

results are shown in Fig. 6 (b) are obtained using 192 processes and for datasets with the different number of data points mentioned on the x-axis in thousands. The smallest dataset consists of one thousand data points and the largest dataset consists of 13 million data points. We observe that, when we have relatively small datasets, the scatter time per data point increases rapidly with the increase in dataset size until 600 thousand data points scattered to 192 processes. However, for larger datasets, the scattering time per data point remains almost constant or increases very slowly. Each data point is of 20 Bytes, thus for less than or equal to 600 thousand data points scattered over 192 processes, less than or equal to 64KB is required to be transferred to each process. Note that the eager protocol message size is set to 64KB with environment variable `MP_EAGER_LIMIT`. Accordingly, the default receive buffer size is increased with environment variable `MP_BUFFER_MEM`. Hence, for a data transfer smaller than 64KB, the communication network follows the eager protocol. However, for a large data transfer, the communication network follows the rendezvous protocol [35] to perform sender-receiver handshake. Our scatter model is designed to capture the effect of this protocol shift.

Fig. 6 (c) shows the results of FiM-Com for broadcast as a function of the number of processes. We see that data transfer of each data point consumes more time as the number of processes increases. This is because for broadcast, with each additional process, the number of data points to be transmitted also increases unlike scatter. The increase in data packets incurs an additional load which in turn increases the average broadcast time of each data point. Lastly, the results under the uplink gather pattern are plotted in Fig. 6 (d) with respect to the number of processes for each data point. The gathering time increases linearly with an increase in the number of processes due to the increased

885 congestion in the communication network, with more data points to be received by the data collecting node. From
 886 Fig. 6, we observe that our FiM models accurately predict the calculation and communication times under different
 887 communication patterns. Also, the predicted results mostly remain in the range bars (i.e., the minimum to the maximum)
 888 of the actual observations. All these graphs show that FiM estimates are very accurate.

890 4.2 Sensitivity Analysis



903 Fig. 7. Sensitivity analysis - (a) Dataset size (b) Number of Iterations (c) Number of Clusters

904
 905
 906 One significant contribution of our modeling techniques is to accurately predict for large datasets by using only
 907 small datasets to train and calibrate the models. In our experiments, we use three small datasets as the training ones to
 908 collect data for model calibration. For a new application, we need first to train to derive a new model. However, for
 909 new sets of application parameters and new datasets (including both small and large ones) of the same application, the
 910 derived model can be used to predict the runtime. Therefore, we perform a sensitivity analysis of different dataset sizes
 911 and application parameters. We argue that if the user has the flexibility to choose application parameters for achieving
 912 optimal performance, our model then can provide useful guidance by helping the user to decide appropriate parameters.

913
 914 For example, K-means processing with more iterations and more clusters can provide better clustering results and
 915 accuracy but consume more time. Therefore, it would be useful to use FiM to estimate the execution time with respect
 916 to the increase in the number of iterations and number of clusters to determine how much extra latency is needed to
 917 achieve better accuracy. The results of execution time for the K-means algorithm as a function of (a) dataset size, (b)
 918 the number of iterations, and (c) number of clusters are plotted in Fig. 7. Fig. 7 also shows the actual and predicted
 919 calculation and communication time. We experiment with different hardware configurations as shown in Table 2, and
 920 present the results of the C5 configuration here. In these experiments, we also use the predicted optimal number of
 921 processes listed in Table 4. For each plot in Fig. 7, we do a sensitivity analysis on one parameter and fix the remaining
 922 two parameters with dataset size of 250 MB, 500 iterations and 250 clusters. We observe that our models can accurately
 923 predict the execution time even when the datasets become large, see Fig. 7 (a). Note that we only use small datasets
 924 to train our models. Fig. 7 (b) shows a linear increase of execution time with increasing number of iterations. Fig. 7
 925 (c) further shows execution time with respect to the increase in number of clusters. Summarizing from Fig. 7, we can
 926 see that FiM consistently achieves predicted results in good agreement with actual ones under different application
 927 parameters like dataset size, number of iterations and number of clusters.

928
 929 We further evaluate the prediction of our models under different hardware configurations. A distributed computing
 930 platform deployed using MPI offers different choices in the number of parallel processes and the distribution of cores
 931 (e.g., C1-C5 listed in Table 2). Fig. 8 shows the actual and predicted execution times with respect to the number of
 932
 933
 934
 935
 936

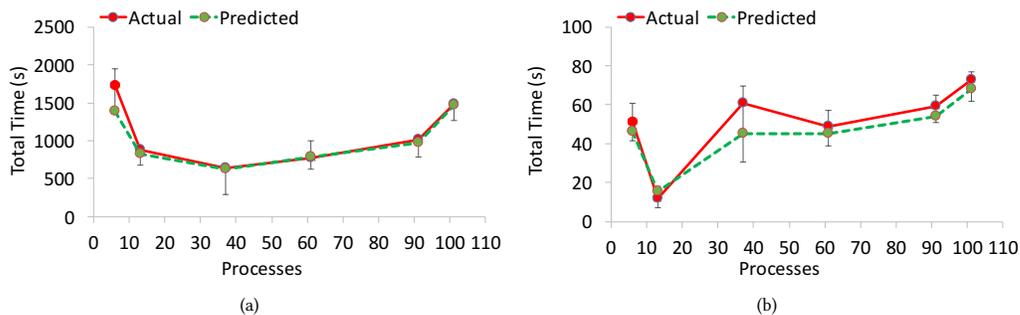


Fig. 8. Sensitivity analysis w.r.t. number of MPI processes for (a) K-means and (b) Pagerank

parallel processes for (a) K-means and (b) Pagerank, respectively. We can see that the best performance (i.e., the shortest execution time) is achieved in the middle range of processes, e.g., 38 for K-means and 12 for Pagerank. FiM is able to predict the optimal performance for both applications accurately. Predicted results match well with actual ones across the different number of processes.

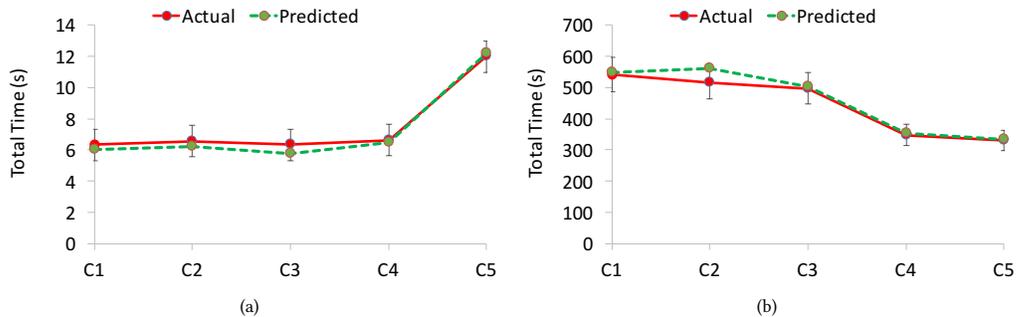


Fig. 9. Sensitivity analysis w.r.t. distribution of cores under (a) a small dataset with 40 vertices and (b) a large dataset with 4039 vertices for Pagerank application

Fig. 9 plots the predicted results for Pagerank under five different hardware configurations listed in Table 2. In order to be able to predict across such heterogeneous platforms, we calibrate our model for data compute and data transfer among each type of available hardware. We observe that the first three configurations (C1, C2, and C3) achieve a shorter execution time for a small dataset (see Fig. 9 (a)). When we have large datasets, C4 and C5 with more distributed cores are better (see Fig. 9 (b)). FiM can accurately predict such performance trends, i.e., Pagerank becomes more scalable on the higher number of distributed cores for larger datasets. These estimation results can thus provide us with insightful data regarding the scalability of an application on a multi-core computing platform.

Here, we evaluate our technique by individually varying each variable parameter. From the results presented above, we see that this sensitivity analysis helps to observe that prediction accuracy of our model remains intact with changing application parameters, number of parallel processes and different hardware platforms.

4.3 Validation

In this section, we present the validation for considering the linear regression model to extrapolate hardware variables such as TC , SC and U . In particular, we show results using K-means clustering as a representative. Recall, for K-means, we chose linear regression to extrapolate hardware variables as a function of application parameters (I , N , and K), see Section 2. There are a variety of regression models that can be used. It is not straightforward to choose the right regression model that is best suited to our requirement. Even a complex model might over-fit the training data, generating a substantial prediction error on other datasets. On the other hand, a simple model may under-estimate the learning trends and produce incorrect predicted results [39]. Considering these two cases, we first investigate the learning trend of three hardware variables (TC , SC , and U) under different settings of application parameters. We choose the linear regression model after examining other regression models such as a piece-wise linear model, Poisson models and quadratic models with different degrees of the polynomial.

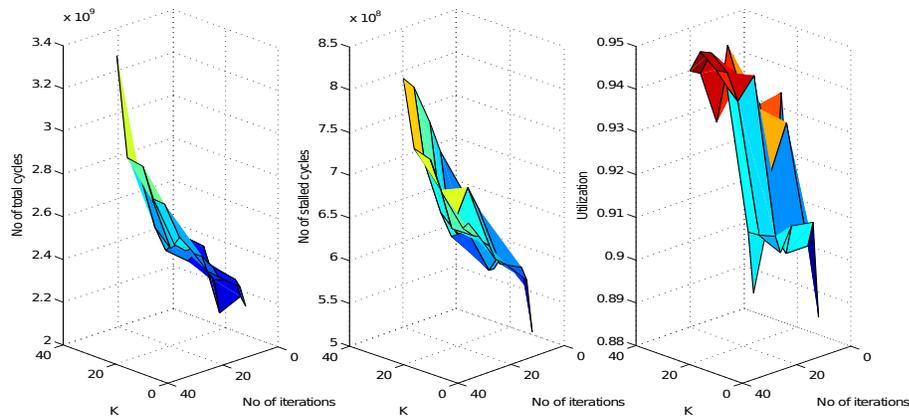


Fig. 10. Regression mapping (K - number of clusters, I - number of iterations)

We investigate the learning trend of three hardware variables under different settings of application parameters. We show the hardware variables such as TC , SC , and U for K-means clustering as a representative. Fig. 10 depicts the resulting surface of each hardware variable as a function of application parameters (e.g., I and K for the K-means application). Similar results can be obtained for other combinations of application parameters, such as (I , N) and (N , K). In Fig. 10, linear surfaces can be found for different hardware variables. Thus, we conclude that hardware variables TC , SC and U , linearly depend on the increment in application parameters (such as K , I and N). These results confirm the use of the linear regression mapping in our machine learning approach.

We also investigate the calibration of our communication models, i.e., the training results of the service rates for each server in the queuing models, see Section 3.2. In Fig. 11, we plot the CDFs of the measured communication service times (i.e., $1/\mu_2$, $1/\mu_3$ and $1/\mu_4$) for the three components (i.e., metadata transfer (M), network initialization (C) and data fetch time (R)) in the communication. The predicted service times drawn from our calibrated service processes (i.e., Eq. 14, Eq. 15 and Eq. 16) are also plotted in Fig. 11. We can see that the predicted service time distributions well capture the actual service times.

The Chi-square goodness-of-fit test [29] is used to evaluate how the three derived statistical distributions (i.e., uniform, logarithmic and exponential) for μ_2 , μ_3 and μ_4 fit the actual measurements. The chi-squared test is used to

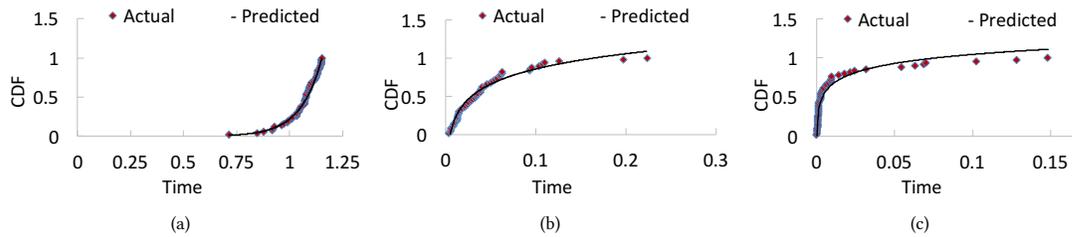


Fig. 11. Cumulative distributed functions (CDFs) of the actual and predicted service times (per job) for (a) metadata transfer, (b) network initialization, (c) data fetching

determine whether there is a significant difference between the expected distribution and the observed distribution. Our null hypothesis is that the derived statistical distributions are not consistent with actual distribution. The significance level of our test is 0.15, meaning that the maximum deviation of the predicted distribution is not more than 15%. The p -values³ that we obtain for service rates μ_2 , μ_3 and μ_4 are equal to 0.08, 0.11 and 0.14, respectively. Since the P -values are less than the significance level (0.15), we reject the null hypothesis. This indicates reasonably high confidence of not more than 15% error between the actual and predicted values.

5 ACKNOWLEDGEMENTS

This work was partially supported by National Science Foundation Career Award CNS-1452751 and AFOSR grant FA9550-14-1-0160. A preliminary version of the paper, titled "FiM: Performance Prediction for Parallel Computation in Iterative Data Processing Applications" was published in IEEE International Conference on Cloud Computing (CLOUD), 2017 [8].

6 RELATED WORK

Modeling helps to shorten the development cycle by providing the necessary insights to obtain optimal performance. Performance modeling can be approached in several different ways, including empirical evaluation [17; 18], simulation [9; 26; 28] and analytical modeling [4; 15]. Empirical evaluation is a technique for gaining knowledge about a system from observations of experiments. This requires the exact implementation as well as similar hardware to the target because results are based on observed ground truth. Empirical techniques were popular in the past when computer hardware was stable over long periods. They give result fast and accurately for the datasets similar to training.

Simulators model hardware such as memory hierarchy, communication buses, parallel ports, and accelerators [31; 33; 40]. They evaluate each block of the given codes, similar to the manner that it is executed on the target machine. Thus, they require source codes while performing prediction. Although simulators like SimpleScalar [3] and CACTI [38] can predict with high accuracy, they also consume a long time to give predicted results. Their slow running time and infrastructure cost are major drawbacks.

Analytical modeling is the technique of building a set of equations to show the high-level abstraction of the behavior of an application and a hardware architecture. This type of modeling can be evaluated quickly and easily as it is reduced to the form of the set of equations to be solved. Analytical models can be flexible and scalable but are comparatively less accurate than empirical and simulation-based models because they lack accurate hardware machine models. The major drawback of analytical modeling is that it limits the scope of prediction. These models will give high prediction errors if tested with parameters which were not captured when building the model equations. An innovative idea is the combination of the above models, e.g., COMPASS [25]. It gives good accuracy but requires the compilation of source code for each new test dataset as well as the conversion of source codes to ASPEN [36].

³The p -value is a statistical measure of deviation of the actual distribution from the hypothesis.

COMPASS requires the hardware machine model to be formed not only for training but also for testing, which is quite time consuming, especially for large datasets.

Predicting the performance of any parallel processing platform consists of two major phases, i.e., calculation and communication. The Markov chain modeling using probabilistic distribution assists in predicting the calculation load of multi-process and multi-core architectures [12; 27; 32]. These approaches model systems in the form of equations, where changes to the code or data require changes to the equations. The above process can be time-consuming when we desire to predict for a range of parameters. Some analytical models [13; 14; 21; 22; 34] require the conversion of source code into a control flow chart for the ease of framing equations. Our model predicts the performance of an application on a range of input parameters without requiring a new set of equations, as FiM uses a machine learning model to learn hardware parameters. Ad hoc analytical models and structured analytical models have been developed to predict the network communication behavior of an application [11]. They use predefined models such as timeline diagrams showing various network overheads. This type of models cannot efficiently capture different patterns like scatter, broadcast and gather. The concept of using queuing theory to model such communication and predict the expected communication time is a novel idea described in our paper.

7 CONCLUSIONS

We present a novel performance modeling technique (FiM) to predict the execution time of iterative multi-stage data processing applications running on parallel computing paradigm. We combine different modeling techniques, such as stochastic Markov modeling, machine learning techniques and queuing theory, to predict the end-to-end execution time. FiM estimates the time required for both data calculation and data communication across a range of input datasets, application configurations and parallel hardware parameters such as number of processes. We demonstrate that FiM helps system designers and application programmers choose optimal hardware parameters and application parameters. More importantly, our prediction models are parameterized using small datasets but can predict accurately for large datasets. We evaluated our approach using NAS Parallel Benchmarks and real iterative data processing applications. We rank the number of processes according to the actual and predicted results from FiM and calculate the correlation between the actual and predicted rankings. Our results show that FiM obtains high correlation in the range of 0.80 to 0.99, which indicates considerable accuracy of our technique. In the future, we plan to expand the scope of our prediction model to investigate other communication patterns like all-to-all communication. We will also capture scenarios where there is an overlap between the communication and computation phases. Large-scale computing platforms, including GPUs, will further be considered as a target environment.

REFERENCES

- [1] Information technology services - research computing. <https://www.northeastern.edu/rc/>.
- [2] Nasa advanced supercomputing division, nas parallel benchmarks. <http://www.nas.nasa.gov/publications/npb.html>.
- [3] AUSTIN, LARSON, AND ERNST. SimpleScalar: An infrastructure for computer system modeling. *Computer* 35, 2 (2002), 59–67.
- [4] BAILEY, D. H., AND SNAVELY, A. Performance modeling: Understanding the past and predicting the future. In *European Conference on Parallel Processing* (2005), Springer, pp. 185–195.
- [5] BARNES, B. J., ROUNTREE, B., LOWENTHAL, D., REEVES, J., DE SUPINSKI, B., AND SCHULZ, M. A regression-based approach to scalability prediction. In *Proceedings of the 22nd annual international conference on Supercomputing* (2008), ACM, pp. 368–377.
- [6] BHIMANI, J., LEESER, M., AND MI, N. Accelerating K-Means clustering with parallel implementations and GPU computing. In *High Performance Extreme Computing Conference (HPEC), 2015 IEEE* (2015), IEEE, pp. 1–6.
- [7] BHIMANI, J., MI, N., AND LEESER, M. Performance prediction techniques for scalable large data processing in distributed MPI systems. In *Performance Computing and Communications Conference (IPCCC), 2016 IEEE 35th International* (2016), IEEE, pp. 1–2.
- [8] BHIMANI, J., MI, N., LEESER, M., AND YANG, Z. FiM: performance prediction for parallel computation in iterative data processing applications. In *Cloud Computing (CLOUD), 2017 IEEE 10th International Conference on* (2017), IEEE, pp. 359–366.
- [9] BINKERT, N., BECKMANN, B., BLACK, G., REINHARDT, S. K., SAIDI, A., BASU, A., HESTNESS, J., HOWER, D. R., KRISHNA, T., SARDASHTI, S., ET AL. The gem5 simulator. *ACM SIGARCH Computer Architecture News* 39, 2 (2011), 1–7.
- [10] CASANOVA, H. Simgrid: A toolkit for the simulation of application scheduling. In *Cluster computing and the grid, 2001. proceedings of first ieee/acm international symposium on* (2001), IEEE, pp. 430–437.

- 1145 [11] CHEN, W., ZHAI, J., ZHANG, J., AND ZHENG, W. LogGPO: An accurate communication model for performance prediction of MPI programs. *Science in*
1146 *China Series F: Information Sciences* 52, 10 (2009), 1785–1791.
- 1147 [12] CHEN, X. E., AND AAMODT, T. M. A first-order fine-grained multithreaded throughput model. In *High Performance Computer Architecture, 2009.*
1148 *HPCA 2009. IEEE 15th International Symposium on* (2009), IEEE, pp. 329–340.
- 1149 [13] CHENNUPATI, G., SANTHI, N., BIRD, R., THULASIDASAN, S., BADAWY, A. A., MISRA, S., AND EIDENBENZ, S. A scalable analytical memory model for
1150 CPU performance prediction. In *Proceedings of the 8th International Workshop on High Performance Computing Systems. Performance Modeling,*
1151 *Benchmarking, and Simulation, PMBS (2017), PMBS@SC '17, Springer*, pp. 114–135.
- 1152 [14] CHENNUPATI, G., SANTHI, N., EIDENBENZ, S., AND THULASIDASAN, S. An analytical memory hierarchy model for performance prediction. In *2017*
1153 *Winter Simulation Conference (WSC) (2017)*, pp. 908–919.
- 1154 [15] CHENNUPATI, G., SANTHI, N., EIDENBENZ, S., ZERR, R. J., ROSA, M., ZAMORA, R. J., PARK, E. J., NADIGA, B. T., LIU, J., AHMED, K., AND OBAIDA, M. A.
1155 *Performance Prediction Toolkit (PPT)*. Los Alamos National Laboratory (LANL), 2017. <https://github.com/lanl/PPT>.
- 1156 [16] DE MELO, A. C. The new linux perf tools. In *Slides from Linux Kongress* (2010).
- 1157 [17] FORTUNE, S., AND WYLLIE, J. Parallelism in random access machines. In *Proceedings of the tenth annual ACM symposium on Theory of computing*
1158 (1978), ACM, pp. 114–118.
- 1159 [18] GIBBONS, P., MATIAS, Y., AND RAMACHANDRAN, V. The queue-read queue-write PRAM model: Accounting for contention in parallel algorithms.
1160 *SIAM Journal on Computing* 28, 2 (1998), 733–769.
- 1161 [19] GROPP, W., AND LUSK, E. User guide for mpe: Extensions for mpi programs. *Argonne National Laboratory* (1998).
- 1162 [20] HOEFLER, T., SCHNEIDER, T., AND LUMSDAINE, A. Loggopsim: simulating large-scale applications in the loggops model. In *Proceedings of the 19th*
1163 *ACM International Symposium on High Performance Distributed Computing* (2010), ACM, pp. 597–604.
- 1164 [21] ISLAM, T. Z., THIAGARAJAN, J. J., BHATELE, A., SCHULZ, M., AND GAMBLIN, T. A machine learning framework for performance coverage analysis of
1165 proxy applications. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2016), IEEE,
1166 pp. 46:1–46:12.
- 1167 [22] JAIN, N., BHATELE, A., ROBSON, M. P., GAMBLIN, T., AND KALE, L. V. Predicting application performance using supervised learning on communication
1168 features. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (2013), ACM, pp. 95:1–95:12.
- 1169 [23] JANSSEN, C. L., ADALSTEINSSON, H., CRANFORD, S., KENNY, J. P., PINAR, A., EVENSKY, D. A., AND MAYO, J. A simulator for large-scale parallel computer
1170 architectures. *Technology Integration Advancements in Distributed Systems and Computing* 179 (2012).
- 1171 [24] KERBYSON, D., ALME, H., HOISIE, A., PETRINI, F., WASSERMAN, H., AND GITTINGS, M. Predictive performance and scalability modeling of a large-scale
1172 application. In *Proceedings of the 2001 ACM/IEEE conference on Supercomputing* (2001), ACM, pp. 37–37.
- 1173 [25] LEE, S., MEREDITH, J., AND VETTER, J. COMPASS: A framework for automated performance modeling and prediction. In *Proceedings of the 29th ACM*
1174 *on International Conference on Supercomputing* (2015), ACM, pp. 405–414.
- 1175 [26] MAGNUSSON, P. S., CHRISTENSSON, M., ESKILSON, J., FORSGREN, D., HALLBERG, G., HOGBERG, J., LARSSON, F., MOESTEDT, A., AND WERNER, B. Simics:
1176 A full system simulation platform. *Computer* 35, 2 (2002), 50–58.
- 1177 [27] MITRA, R., JOSHI, B. S., RAVINDRAN, A., MUKHERJEE, A., AND ADAMS, R. Performance modeling of shared memory multiple issue multicore machines.
1178 In *2012 41st International Conference on Parallel Processing Workshops* (2012), IEEE, pp. 464–473.
- 1179 [28] PATEL, A., AFRAM, F., CHEN, S., AND GHOSE, K. Marss: a full system simulator for multicore x86 cpus. In *Proceedings of the 48th Design Automation*
1180 *Conference* (2011), ACM, pp. 1050–1055.
- 1181 [29] PLACKETT, R. L. Karl pearson and the chi-squared test. *International Statistical Review/Revue Internationale de Statistique* (1983), 59–72.
- 1182 [30] RICO-GALLEGO, J.-A., DÍAZ-MARTÍN, J.-C., AND LASTOVETSKY, A. L. Extending τ -Lop to model concurrent MPI communications in multicore clusters.
1183 *Future Generation Computer Systems* 61 (2016), 66–82.
- 1184 [31] RODRIGUES, A. F., HEMMERT, K. S., BARRETT, B. W., KERSEY, C., OLDFIELD, R., WESTON, M., RISEN, R., COOK, J., ROSENFELD, P., COOPERBALLS, E., AND
1185 JACOB, B. The structural simulation toolkit. *SIGMETRICS Perform. Eval. Rev.* 38, 4 (2011), 37–42.
- 1186 [32] SAAVEDRA-BARRERA, R. H., AND CULLER, D. E. An analytical solution for a markov chain modeling multithreaded. Tech. rep., Citeseer, Berkeley, CA,
1187 USA, 1991.
- 1188 [33] SHAO, Y. S., REAGEN, B., WEI, G.-Y., AND BROOKS, D. Aladdin: A pre-rtl, power-performance accelerator simulator enabling large design space
1189 exploration of customized architectures. In *ACM SIGARCH Computer Architecture News* (2014), vol. 42, IEEE Press, pp. 97–108.
- 1190 [34] SNAVELY, A., CARRINGTON, L., WOLTER, N., LABARTA, J., BADIA, R., AND PURKAYASTHA, A. A framework for performance modeling and prediction. In
1191 *Proceedings of the 2002 ACM/IEEE Conference on Supercomputing* (2002), SC '02, IEEE, pp. 1–17.
- 1192 [35] SNIR, M. *MPI—the Complete Reference: The MPI core*, vol. 1. MIT press, 1998.
- 1193 [36] SPAFFORD, K., AND VETTER, J. Aspen: a domain specific language for performance modeling. In *Proceedings of the International Conference on High*
1194 *Performance Computing, Networking, Storage and Analysis* (2012), IEEE Computer Society Press, p. 84.
- 1195 [37] WILLIAMS, S., WATERMAN, A., AND PATTERSON, D. Roofline: an insightful visual performance model for multicore architectures. *Communications of*
1196 *the ACM* 52, 4 (2009), 65–76.
- 1197 [38] WILTON, S., AND JOUPLI, N. CACTI: An enhanced cache access and cycle time model. *Solid-State Circuits, IEEE Journal of* 31, 5 (1996), 677–688.
- 1198 [39] YUAN, M., AND LIN, Y. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical*
1199 *Methodology)* 68, 1 (2006), 49–67.

- 1197 [40] ZHENG, G., KAKULAPATI, G., AND KALE, L. V. Bigsim: a parallel simulator for performance prediction of extremely large parallel machines. In *18th*
1198 *International Parallel and Distributed Processing Symposium, 2004. Proceedings.* (2004).
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248