

# A Study of Dynamic Meta-Learning for Failure Prediction in Large-Scale Systems

Zhiling Lan, Jiexing Gu, Ziming Zheng

*Illinois Institute of Technology, Chicago, IL 60616*

Rajeev Thakur, Susan Coghlan

*Argonne National Laboratory, Argonne, IL, 60439*

---

## Abstract

Despite years of study on failure prediction, it remains an open problem, especially in large-scale systems composed of vast amount of components. In this paper, we present a dynamic meta-learning framework for failure prediction. It intends to not only provide reasonable prediction accuracy, but also be of practical use in realistic environments. Two key techniques are developed to address technical challenges of failure prediction. One is *meta-learning* to boost prediction accuracy by combining the benefits of multiple predictive techniques. The other is *a dynamic approach* to dynamically obtain failure patterns from a changing training set and to dynamically extract effective rules by actively monitoring prediction accuracy at runtime. We demonstrate the effectiveness and practical use of this framework by means of real system logs collected from the production Blue Gene/L systems at Argonne National Laboratory and San Diego Supercomputer Center. Our case studies indicate that the proposed mechanism can provide reasonable prediction accuracy by forecasting up to 82% of the failures, with a runtime overhead less than 1.0 minute.

*Key words:* failure prediction, meta-learning, dynamic techniques, large-scale systems, Blue Gene

---

---

*Email addresses:* {lan, jgu5, zzheng11}@iit.edu (Zhiling Lan, Jiexing Gu, Ziming Zheng), {thakur, smc}@mcs.anl.gov (Rajeev Thakur, Susan Coghlan).

# 1 Introduction

## 1.1 Motivations

To meet the insatiable demand in science and engineering, supercomputers continue to grow in size. Production systems with tens to hundreds of thousands of computing nodes are being designed and deployed [30]. Such a scale, combined with the ever-growing system complexity, is introducing a key challenge — reliability — in the field of high performance computing (HPC). Despite great efforts on the design of ultra-reliable components, the increase of system size and complexity has outpaced the improvement of component reliability. Recent studies have pointed out that the mean-time-between-failure (MTBF) of teraflop and soon-to-be-deployed petaflop machines are only on the order of 10 - 100 hours [4,41,22].

To address the reliability problem, considerable research has been done to improve fault resilience of computer systems and their applications through various technologies. Representative works include failure-aware resource management and scheduling [10,15,20], checkpointing [6,18,24,38], proactive or adaptive runtime resilience support [14,29]. The advance of these technologies, however, greatly depends on whether we can predict the occurrence of failure, i.e., failure prediction. For example, proactive fault tolerant methods, such as preemptive process migration, require failure forecasting to enable cost-effective failure avoidance. For reactive methods such as checkpointing, an efficient failure prediction could substantially reduce their operational cost by telling when and where to perform checkpoints, rather than blindly invoking actions periodically with an unwisely chosen frequency.

Despite years of study on failure prediction, it remains an unsolved problem, especially in large-scale systems composed of substantial amount of components. We summarize its key challenges from two aspects. First is *prediction accuracy*. Existing studies mainly concentrate on exploring one specific method to capture and discover failure patterns. As a matter of fact, in a large-scale system the sources of failures are numerous and complex, thus it is improper to expect a single method to capture all of failures alone. For example, many rule-based classifiers emphasize on discovering correlation relationships between warning messages and fatal events for failure prediction [39,23]. As we will show in our experiments, they are limited by the amount of fatal events occurring without any precursor warnings. Hence, relying on these methods alone is insufficient to provide an effective failure forecasting. Further, hardware and software upgrades are common at supercomputing centers, and system workloads tend to vary during system operation. These changes can drastically alter system behaviors [41]. As a result, static analysis that uses

a fixed set of historic data to learn failure patterns cannot adapt to system changes at runtime, thereby being incapable of providing accurate forecasting.

The next is with respect to *practical use*. While many predictive models have been presented so far, most of them merely focus on the algorithm-level improvement and are too complicated to be of practical use for online failure prediction [17,40]. In addition, to obtain sufficient failure patterns, many predictive methods require a long training phase (e.g., one or more years), thereby being unable to provide prediction service for a long period of time [8]. Given that most systems at supercomputing centers only have a couple of years in production, this requirement must be removed.

## 1.2 Main Contributions

In this study, we present a dynamic meta-learning framework for online failure prediction in large-scale systems. It intends to provide reasonable prediction accuracy, as well as be of practical use in realistic environments. Our framework consists of two parts to process and analyze system events: one is to preprocess system events by means of event categorization and filtering (i.e., *data preprocessing*), and the other is to examine the cleaned events for generating failure patterns and triggering failure warnings through continuous runtime event analysis (i.e., *failure prediction*). These two parts, along with their main components, are illustrated in Figure 1. The details of the main components will be described in Section 3 - 4.

Our method employs two key techniques to address the challenges listed above. First, *meta-learning* is explored to boost prediction accuracy by combining the benefits of multiple predictive methods. It enables us to discover a variety of failure patterns in large-scale systems, without constructing complex models of the underlying system. In this study, we integrate three widely-used predictive methods, namely association rule based learner [23,28], statistical rule based learner [16], and probability distribution [4], in the framework by applying a simple yet efficient ensemble learning method. Next, *a dynamic mechanism* is adopted to trigger relearning periodically and to adaptively extract effective rules of failure patterns by actively tracing prediction accuracy.

To demonstrate the effectiveness and practical use of our framework, we evaluate it with the real RAS (Reliability, Availability and Serviceability) logs collected from the production Blue Gene/L systems at Argonne National Laboratory (ANL) and San Diego Supercomputer Center (SDSC). The use of multiple RAS logs is to ensure our framework is not bias to any specific log and thus produces representative results expected in other systems as well. To comprehensively assess our framework, our experiments are structured to

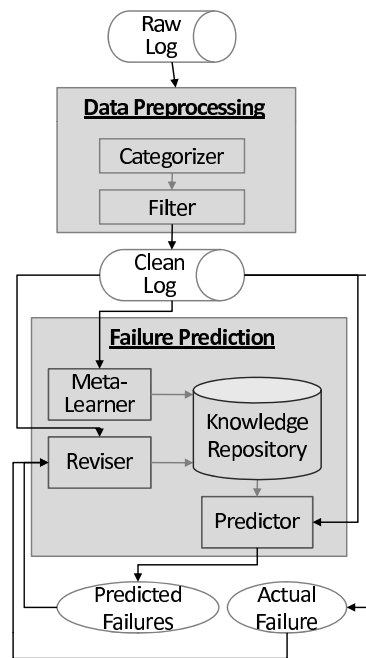


Fig. 1. Overview of our dynamic meta-learning framework for failure prediction

answer the following questions:

- Q1:** How much improvement is achieved by the meta-learning?
- Q2:** How much improvement is achieved by the dynamic approach?
- Q3:** How sensitive is prediction accuracy to prediction window size?
- Q4:** How much runtime overhead is introduced?

Our experiments demonstrate that meta-learning can effectively improve prediction accuracy by up to three times, and the dynamic approach is capable of adapting to system changes, even after a major system reconfiguration. For both systems, our method can provide reasonable prediction accuracy by predicting up to 82% of failures, with a runtime overhead less than 1.0 minute. Furthermore, prediction accuracy depends on how far away we are interested in forecasting failures. In general, the larger the window is, the higher the prediction coverage is, along with a higher false alarm rate. The rules of failure patterns change dramatically during system operation, which further proves that the dynamic approach is indispensable for better prediction. Finally, runtime overhead increases with the growing size of the training set. Overall speaking, we find that for both systems, the use of recent six-month training set can well balance between prediction accuracy and runtime overhead.

We note that three predictive methods, namely association rule based learning, statistical learning, and probability distribution, have been tested in our experiments. Rather than focusing on which predictive method is better, this study focuses on providing a general framework to dynamically combine multiple predictive methods for better failure prediction. We believe that other predictive methods like [17,37] can be easily integrated into our framework.

### *1.3 Paper Organization*

The rest of the paper is organized as follows. Section 2 gives the background information of Blue Gene systems and system logs. Section 3 describes the details of data preprocessing, followed by a detailed description of our dynamic meta-learning method in Section 4. The case studies with real failure logs are presented in Section 5. Section 6 discusses the related work and points out the key differences between this work and existing studies. Finally, Section 7 summarizes the paper.

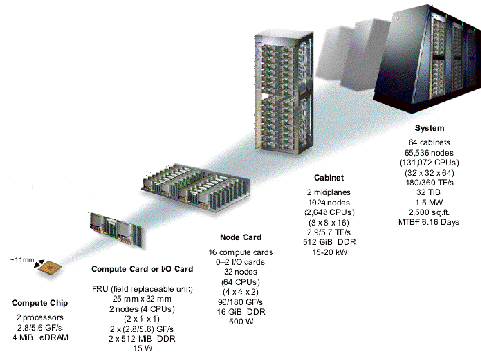


Fig. 2. Blue Gene/L System Overview

## 2 Background

### 2.1 Overview of Blue Gene/L

In this paper, we use RAS (Reliability, Availability and Serviceability) logs collected from the Blue Gene/L systems for case studies, thus in the below we give an overview of the systems and their RAS logging facilities. The proposed dynamic meta-learning framework can be easily extended for failure prediction of other large-scale systems.

System packaging is an integral aspect of Blue Gene/L systems (see Figure 2). As shown in the figure, the basic building block is called computer chip. Each computer chip consists of two PPC 440 cores, with a 32KB L1 cache and a 2KB L2 cache. The cores share a 4MB EDRAM L3 cache. A compute card contains two computer chips, a node card contains 16 compute cards, and a midplane holds 16 node cards with a total of 1,024 processors. In addition to compute nodes, a midplane is also populated with several I/O nodes which are configured to handle file I/O and host communication. Each midplane also has one service card that performs system management services like monitoring node heartbeat and checking errors. More details of the system architecture can be found in the literature [3].

In Blue Gene, the Cluster Monitoring and Control System (CMCS) service is implemented on the service nodes for the purpose of system monitoring and error checking. The service node, which is available in each midplane, acquires specific device information, such as RAS events, directly through the control network. Runtime information is collected from computer and I/O nodes by a polling agent, reported to the CMCS service, and finally stored in a centralized DB2 repository. This system event logging mechanism works in a granularity of less than 1 millisecond.

The entries in the RAS log include hard errors, soft errors, machine checks,

Table 1  
Attributes of RAS Events in Blue Gene

Attribute	Description
Record ID	An integer denoting event sequence number
Event Type	The mechanism through which the event is recorded
Event Time	Timestamp associated with the reported event
Job ID	Job that detects the event
Location	Place of the event (e.g., chip/node card/service card/link card)
Entry Data	A short description of the event
Facility	The service/hardware component experiencing the event
Severity	The level of severity of the reported event

and software problems. Information about scheduled maintenance, reboot, and repair is not included. Each record has eight attributes which are described in Table 1.

The SEVERITY attribute can be one of the following levels — INFO, WARNING, SEVERE, ERROR, FATAL, or FAILURE — which also denotes the increasing order of severity. INFO events are for the purpose of general information to administrators about the reliability of various hardware/services components in the system. WARNING events report unusual events in node cards, link cards, service cards or related services. SEVERE events provide more information about the reasons causing problems in node cards or service cards etc. ERROR events indicate problems that require further attention of administrators. An event with any of the above SEVERITY attributes is either informative in nature, or is related more to the initial configuration errors, and is thus relatively transparent to the applications/runtime environment. However, FATAL or FAILURE events (such as “uncorrectable torus error”, “communication failure socket closed”, “uncorrectable error detected in edram bank”, etc.) are more severe, and usually lead to system/application crashes. Our primary focus in this study is to predict FATAL and FAILURE events (denoted as *fatal events*<sup>1</sup>, while other events are denoted as *non-fatal events*). In [41], Oliner et al. have pointed out that some of the fatal events provided by the RAS log are not true fatal events. We have consulted with experienced system administrators at both ANL and SDSC, and removed these “fake” fatal events from the failure list.

<sup>1</sup> In the paper we use “failure” and “fatal event” interchangeably.

Table 2  
Log Description

Log	Period	Weeks	Event No.	Log Size
<b>ANL BGL</b>	Jan. 21, 2005 - Jun. 19,2007	112	5,887,771	2.27 GB
<b>SDSC BGL</b>	Dec. 6, 2004 - Jun. 11, 2007	132	517,247	463 MB

## 2.2 Test Logs

Two production Blue Gene/L systems are used in our experiments. One is at SDSC, which consists of three racks with 3,072 dual-core compute nodes and 384 I/O nodes. The configuration is chosen to support data-intensive computing. Each node consists of two PowerPC processors that run at 700 MHz and share 512 MB of memory, giving an aggregate peak speed of 17.2 teraflops and a total memory of 1.5 TB [31]. The other is at ANL, which has one rack with 1,024 dual-core compute nodes and 32 I/O nodes [32]. The aggregate peak performance is of 5.7 teraflops, with a total memory of 500 gigabytes. Both systems are mainly used for scientific computing. Table 2 summarizes the RAS logs used in our experiments.

The log from ANL has much more number of records, although the system has only one rack of nodes. This is due to a large quantity of error checking messages produced at the ANL site. For example, during the 50th week of the ANL’s log (between January 6 and January 13, 2006), there were over 1.15 million of machine checking information messages generated. System administrators at ANL ran diagnostics more frequently to cull out bad hardware faster, without applications seeing it.

## 3 Data Preprocessing

Raw logs generally contain many repeated or redundant information. This is because each computer chip runs a polling agent to collect the errors reported by the chip. As each job is assigned to multiple computer chips, any failure of the job will get reported multiple places — once from each of the assigned computer chips. Thus multiple components may report the same failure. Also, the logging mechanism records the events at a very fine granularity (e.g., in millisecond), but the recorded event time is generally in seconds or minutes, thus leading to multiple entries of an event with the same time stamp. Therefore, before a RAS log can be used for failure prediction, it needs to be processed to identify unique RAS events.

As shown in Figure 1, data preprocessing mainly consists of two components:



one is *event categorizer*, and the other is *event filter*. The categorizer aims at providing a precise list of RAS types, and the filter removes redundant data by conducting temporal compression at a single location and spatial compression across multiple locations. The goal of data preprocessing is to provide a list of unique events for failure prediction.

### 3.1 Categorizer

Event categorization is a time consuming process. It requires a deep understanding of system events, thus close collaboration with system administrators is essential for obtaining a list of meaningful event categories. Fortunately, for a specific system, the process only needs to be performed once. Once a standard categorizing of system events is constructed, we can use it for a long period of time, unless a drastic change occurs in the system (e.g., system reconfiguration). In case of minor changes during system operation, existing categorization technologies such as the one presented in [5] can be applied for dynamic tuning of event classifications.

We adopt a *hierarchical approach* for event categorization. We first divide system events into several high-level classifications, and then further group events into a number of subcategories based on their attributes. For the Blue Gene/L systems, ten high-level event categories are identified based on the *Facility* field, which are further divided into 219 low-level event types based on the *Severity* and *Entry Data* fields. Further, it is also necessary to distinguish these event categories into fatal or non-fatal groups for the purpose of data training. Non-fatal events indicate system warnings or information messages, while fatal events refer to those critical events that lead to system or application crashes. Although RAS logs from Blue Gene/L provide severity level for each event, it is not accurate since some fatal or failure events are not truly fatal at all [41]. By working with system administrators, we have identified and removed some of these events from the fatal list. Totally, there are 69 fatal events for the Blue Gene/L systems. Examples are shown in Table 3.

### 3.2 Filter

Event filtering is required to remove duplicated or unnecessary entries in the log. Common cleaning steps include removing duplicated entries, removing unnecessary entry attributes, correcting inaccurate attributes, preparing output files for corresponding learning methods, etc. In this study, we apply both temporal compression and spatial compression to remove duplicate entries by applying threshold based techniques. With temporal compression at a single

Table 3  
Event Categories in Blue Gene/L

<b>Main Category</b>	<b>Examples</b>	<b>No. of Fatal Categories</b>	<b>No. of Non-Fatal Categories</b>
APP	Load Program failure function call failure	10	7
BGLMASTER	segmentation failure BGLMaster restart info	2	2
CMCS	CMCS command info CMCS exit info	0	4
DISCOVERY	nodecard communication warning servicecard read error	0	24
HARDWARE	midplane service warning	1	12
KERNEL	broadcast failure cache failure cpu failure node map file error	46	90
LINKCARD	linkcard failure	1	0
MMCS	control network MMCS error	0	5
MONITOR	node card temperature error	9	5
SERV_NET	system operation error	0	1
TOTAL		69	150

location, events from the same location with identical values in the *Job ID* and *Location* fields are coalesced into a single entry, if reported within a predefined time duration. With spatial compression across multiple locations, we remove those entries that are close to each other within a predefined time duration, with the same *Entry Date* and *Job ID*, but from different locations.

How to decide an optimal threshold for filtering is still an open question. In this study, we adopt an *iterative approach* [42,43]. We first set the threshold to a very small number, and then gradually increase the number. The search stops when there is no significant change with respect to compression rate. Table 4 presents the numbers of events after applying different thresholds, where we separate the numbers according to the high-level event categories. The column where threshold is set to zero denotes the raw logs before any compression.

Table 4  
Number of Events with Different Filtering Thresholds (in seconds)

	Log	0s	10s	60s	120s	200s	300s	400s
APP	ANL	6758	1942	1827	1684	1566	1453	1378
	SDSC	26358	754	741	675	615	579	556
BGLMASTE	ANL	123	123	120	115	115	109	107
	SDSC	119	119	114	105	99	93	90
CMCS	ANL	302	295	292	286	284	283	280
	SDSC	437	433	421	404	384	362	356
DISCOVERY	ANL	18054	1727	1429	937	676	578	497
	SDSC	60748	3621	3356	1352	750	565	556
HARDWARE	ANL	1840	668	633	601	593	539	468
	SDSC	1648	422	349	316	296	283	278
KERNEL	ANL	5819166	59784	47998	40777	33847	26754	23823
	SDSC	426816	4238	4056	3940	3747	3595	3379
LINKCARD	ANL	64	30	18	15	13	11	10
	SDSC	188	120	107	95	92	88	82
MMCS	ANL	954	561	521	484	467	444	437
	SDSC	929	654	630	590	563	523	501
MONITOR	ANL	40509	19774	16120	15969	15834	15689	15421
	SDSC	0	0	0	0	0	0	0
SERV_NET	ANL	1	1	1	1	1	1	1
	SDSC	4	4	4	4	4	4	4

For both logs, the amount of compression of events achieved is not significant when the threshold greater than 300 seconds is used. Additionally, as RAS events are logged at a sub-second frequency, taking a higher threshold value will increase the chances of different events being clustered together. Hence, we choose 300 seconds as the threshold to coalesce events, which achieves above 98% compression rate for the logs.

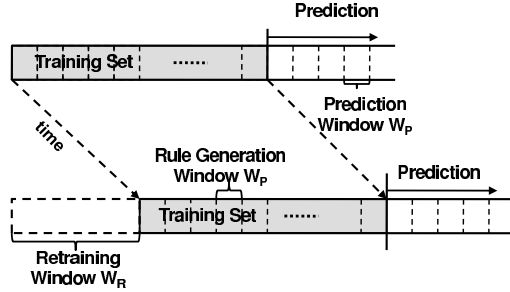


Fig. 3. Key Terms in Failure Prediction

## 4 Prediction Methodology

Our prediction method consists of three major components: the *meta-learner*, the *reviser*, and the *predictor* (see Figure 1). The *meta-learner* examines system events to discover various fault patterns by applying multiple predictive methods. The generated failure patterns or rules will be stored in a knowledge repository which encompasses all of the relevant information of failure patterns. It contains all the learned rules of failure patterns and corresponding ensemble rules for meta-learning. These rules of failure patterns are subjected to modifications made by the reviser at runtime. The *reviser* monitors prediction accuracy by comparing the predicted results and the actual failures, and then modifies the knowledge base. The training set used by the meta-learner and the reviser is periodically changed during system operation. The *predictor* actively examines system events. In case that the occurrence of an event triggers a matching pattern in the knowledge base, it will trigger a warning.

Distinguishing from existing studies like [23,16,17,12,37], our framework has two novel features. One is to exploit meta-learning(i.e., ensemble learning) to boost failure prediction, and the other is to dynamically learn failure patterns from a changing training set during system operation.

Before we go to the details of these components, we first present the main terms used in our framework (see Figure 3). The training set, which may be dynamically adjusted every  $W_R$  weeks (denoted as *retraining window*), is part of the log from which the meta-learner and the reviser use to generate the rules of failure patterns. In other words, the meta-learner and the reviser will be invoked every  $W_R$  weeks. The rules are generated with a fixed time window, generally in the order of a couple of minutes to hours (denoted as *rule generation window  $W_P$* ). The rules learned will be stored in the knowledge repository, which will be used by the predictor for failure prediction before the next retraining. The predictor actively monitors the events occurring during *prediction window*, whose size is the same as the rule generation window  $W_P$ , and in case of a matching rule, it will trigger a warning.

#### 4.1 Meta-learner

The meta-learner focuses on revealing and learning the cause-and-effect relations of system events by applying data mining techniques. Data mining, or knowledge discovery, is a computer-assisted process of searching and analyzing data sets for hidden patterns [11]. Meta-learning, also known as ensemble-learning, can be loosely defined as learning from learned knowledge [21]. It emphasizes on combining different individual models (denoted as *base learners*) to boost overall predictive effectiveness.

In this study, we choose three widely-used predictive methods, namely association rule based method [28,23], statistical rule based method [16], and probability distribution based method [4], as our base learners. The meta-learner intends to identify a preferable combination of these base learners. In the following, we first describe the base learners, followed by presenting our meta-learning method. Note that other base methods can be easily incorporated.

**Base Learners.** The first base learner is based on *association rules*. It examines *causal correlations* between *non-fatal* and *fatal* events by building association rules. In general, an association rule is in the form  $X \rightarrow Y$ , where the rule body  $X$  and  $Y$  are subsets of an event set. It states that a transaction that contains the items in  $X$  are likely to contain the items in  $Y$ . Association rules are characterized by two measures: *support* which measures the percentage of transactions that contain both items  $X$  and  $Y$ , and *confidence* which measures the percentage of item sets containing the items  $X$  that also contain the items  $Y$ . The problem of mining association rules consists of generating all the association rules from a set of items that have both *support* and *confidence* greater than the user-defined thresholds. Given that failure is rare event, low values of *support* and *confidence* are set for the purpose of capturing infrequent events.

On the training set, for each fatal event, we identify the set of non-fatal events preceding it within the rule generation window  $W_P$ . The set, including the fatal event and their precursor nonfatal events, is called an event set. We then apply the standard association rule algorithm to build rule models for event sets that are above the minimum *support* and *confidence*. The association rules will be in the form of  $\{e_1, e_2, \dots, e_k\} \rightarrow f, confidence$ , where  $e_i$  and  $e_j$  ( $1 \leq i, j \leq k$ ) are non-fatal events,  $f$  is a fatal event. For instance, two examples from the SDSC log are listed below:

*networkWarningInterrupt, networkError*  $\rightarrow$  *socketReadFailure*: 1.0  
*idoStartInfo, bglStartInfo*  $\rightarrow$  *fsFailure*: 0.79

Our second base learner emphasizes on discovering *statistical characteristics*,




Fig. 4. Temporal Correlations Among Fatal Events.

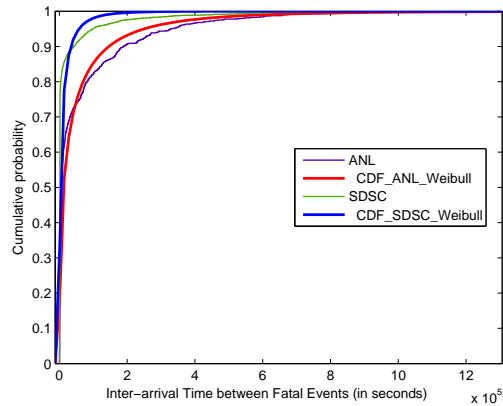


Fig. 5. Cumulative Distribution Functions (CDFs) of Fatal Events. The thin curves represent the actual fatal events, while the thick curves model the Weibull distributions of the events.

i.e., how often and with what probability will the occurrence of one failure influence subsequent failures, among fatal events and then using the obtained statistical rules for failure prediction. It is denoted as *statistical based method*. Studies have shown that temporal correlations among fatal events are common in large-scale systems [4,16,37]. Figure 4 plots fatal events per day occurred at ANL and SDSC. We can observe that a significant number of failures happen in close proximity, and our further analysis indicates that network and I/O stream related failures form a majority of such failures.

Specifically, on the training set, we calculate the probability of  $k$  failures occurred within the rule generation window  $W_P$ . If the probability is larger than a user-defined threshold, then a statistic rule is generated, along with its probability value. As an example, we have discovered that for both logs, if four failures occur within 300 seconds, then the probability of another failure is 99%.

The third base learner is called *probability based method*. It generates *probability distribution* of fatal events and stores it for failure prediction. Different from the above two methods which attempt to discover *short-term* (e.g., in the order of minutes) correlations among events for failure prediction, this method




Fig. 6. Meta-learning Method

recognizes that some failure events may not have any short-term precursor events and intends to utilize *long-term* failure behavior for failure prediction. Here, the long-term means the probability distribution of failure events, which is generally in the order of hours or even days.

Specifically, the method calculates inter-arrival times between adjacent fatal events and uses maximum likelihood estimation to fit a mathematical model to these data. Distributions like Weibull, exponential, and log-normal are examined for generating the cumulative distribution function (CDF) of fatal events. Figure 5 plots the CDFs of fatal events occurred at ANL and SDSC. By calculating the probability of possible failure based on the derived CDF function, this base method will trigger a warning if the probability is larger than a user-defined threshold, or equally saying, when the elapsed time since the last failure is longer than some threshold. For instance, on a training set from the SDSC log, the Weibull distribution of  $F(t) = 1 - e^{-(t/19984.8)^{0.507936}}$  is determined to be the best CDF to describe inter-arrival times between adjacent fatal events. Hence, if the threshold is set to 0.60, when the elapsed time since the last failure is 20000 seconds, a warning will be triggered because  $F(20000)(= 0.63)$  is larger than the threshold.

**Ensemble Learning.** There are many ways to combine base models, among which bagging, boosting, and stacking are well-known ensemble methods. In our study, we choose the *mixture-of-experts* model, which is a variation of stacking method [8,21]. Figure 6 illustrates our meta-learning process. The basic idea is simple: base learners are experts in some portion of the feature space, and the combination rule selects the most appropriate classifier for each instance. Based on verification on the training data, the meta-learner determines the ordering of rules used for prediction.

In our case studies, the order is association rule, followed by statistical rule, and finally probability distribution. Specifically, when an event occurs, if it is a non-fatal event, the meta-learner first checks whether it will trigger a matching of an association rule; if it is a fatal event, the meta-learner will check whether it will trigger a matching of a statistical rule. If no matching is found, the meta-learner will check the elapsed time since the last failure and apply the derived probability distribution of failures for failure forecasting.

## 4.2 Reviser

The reviser is responsible for modifying the candidate rules generated by the meta-learner via monitoring the actual observations and the predicted results. This is to ensure the effectiveness of the learned rules in the knowledge repository. As mentioned earlier, in order to capture infrequent items, the parameters used in the base learners may be adopted without much consideration regarding their effectiveness, thereby probably resulting in some bad rules. Thus, the reviser checks each rule in the knowledge repository by applying the ROC (Receiver Operating Characteristic) analysis [11]. It enables us to select optimal models and discard suboptimal ones independently from the class distribution. The reviser will examine each rule and only keep the rules which can provide satisfactory accuracy [9]. The detailed method is shown in Algorithm 1.

---

**Algorithm 1** The Reviser

---

For each rule  $r$  generated by the meta-learner:

- (1) Count its true positives  $T_P$ , false positives  $F_P$ , and false negatives  $F_N$  on the training data;
- (2) Calculate two metrics  $m_1(r)$  and  $m_2(r)$ :

$$m_1(r) = \frac{T_P}{T_P + F_P}, m_2(r) = \frac{T_P}{T_P + F_N};$$

- (3) Calculate  $ROC(r)$ :

$$ROC(r) = \sqrt{m_1(r)^2 + m_2(r)^2};$$

- (4) Keep the rule if its ROC value is larger than a predefined threshold  $MinROC$ ; otherwise, discard the rule.
- 

## 4.3 Predictor

The predictor actively monitors system events and triggers a warning when a rule is observed within the prediction window  $W_P$ . In order to be used for online forecasting, an *event-driven* method is adopted for its design [9]. That is, the predictor triggers a warning on the occurrence of events.

The detailed method is presented in Algorithm 2. The predictor maintains three event lists. One is called  $F - List$  which records a list of triggering events for each failure event. The second is called  $E - List$  which tracks a list of failure events that may be triggered by each event (fatal or non-fatal). The third is to keep the most recent events occurred within  $W_P$ . Upon an



occurrence of an event  $e$ , the predictor appends the event into the third list (Step 1), and then goes through its  $E - List$  to find out all possible failures that may be triggered by its occurrence (Step 2). For each possible failure  $f^i$ , the predictor checks its  $F - list$  to see whether a cause-and-effect rule is matched in the knowledge repository (Step 3 and 4).

---

**Algorithm 2** The Predictor

---

First, it creates two lists based on the learned rules:

$$F - List = \{f_i \Leftarrow \{e_{i1}, e_{i2}, \dots, e_{ik}\} : 1 \leq i \leq N_f\}$$

$$E - List = \{e_j \Rightarrow \{f_{j1}, f_{j2}, \dots, f_{jl}\} : 1 \leq j \leq N_e\}$$

where  $f_i$  is a fatal event and  $e_i$  is a fatal or non-fatal event,  $N_f$  is number of fatal events, and  $N_e$  is number of any events. During operating, when an event  $e$  occurs:

- (1) Append  $e$  into the monitoring event set  $E = \{e_1, e_2, \dots, e_n, e\}$  where the events are sorted in an increasing order of their occurrence times, remove  $e_i$  if its occurrence time is more than  $W_P$  before the occurrence time of  $e$ , i.e., keep the most recent events occurred within  $W_P$
  - (2) Go through the E-List of  $e$ , obtain the potential failures that may be triggered by  $e$ :  $\{f^1, f^2, \dots, f^k\}$
  - (3) For each potential failure  $f^i$ , go through its F-List:  $f^i \Leftarrow \{e_{i1}^i, e_{i2}^i, \dots, e_{ik}^i\}$
  - (4) If  $\{e_{i1}^i, e_{i2}^i, \dots, e_{ik}^i\} \subseteq E$ , then produce a warning that the failure  $f^i$  may occur within the time of  $W_P$ .
- 

## 5 Experiments

To evaluate the effectiveness of the proposed framework, we use the real RAS logs collected from the production systems at ANL and SDSC (see Table 2). Further, to comprehensively examine the framework, our experiments are structured to answer the key questions listed in Section 1.

### 5.1 Evaluation Metrics

Two metrics are used to measure prediction accuracy:

- (1) *Precision* is defined as the proportion of correct predictions to all the

predictions made.

$$precision = \frac{T_p}{T_p + F_p}$$

- (2) *Recall* is defined as the proportion of correct predictions to the number of failures.

$$recall = \frac{T_p}{T_p + F_n}$$

Here,  $T_p$  is number of correct predictions (i.e., true positives),  $F_p$  is number of false alarms (i.e., false positives), and  $F_n$  is number of missed failures (i.e., false negatives). Obviously, a good prediction engine should achieve a high value (closer to 1.0) for both metrics. We note that these metrics are also used by the reviser to determine whether a rule is effective or not (see Algorithm 1).

## 5.2 Results

In our experiments, the training set is initially set to six months, which will be dynamically adjusted during operation. The default retraining window  $W_R$  is four weeks, and the default prediction window  $W_P$ , also the rule generation window, is 300 seconds.

The minimum *support* and *confidence* values for association rules are set to 0.01 and 0.1 respectively. The low values are chosen for the purpose of capturing infrequent events. The rules that are not good will be removed by the reviser. There are three other parameters used by our framework, namely *MinROC* for the reviser, and the thresholds for statistical rule based learner and probability distribution based learner. In our experiments, *MinROC* is set to 0.7, and the thresholds for statistical rules and probability distribution based learner are set to 0.8 and 0.6 respectively. Choosing optimal values for these parameters is difficult, and often experimental determination might be the only viable option. We have tested different values, from a low value like 0.1 to a high value like 0.9, and found that these values can yield the best prediction accuracy for both logs. In general, low values for these parameters result in more failure rules and thus better failure coverage, at the expense of introducing more false alarms.

### 5.2.1 Q1: How Much Improvement Is Achieved by the Meta-learning?

In this set of experiments, we compare prediction results by using static meta-learner as against individual base learners (i.e., association rule, statistical

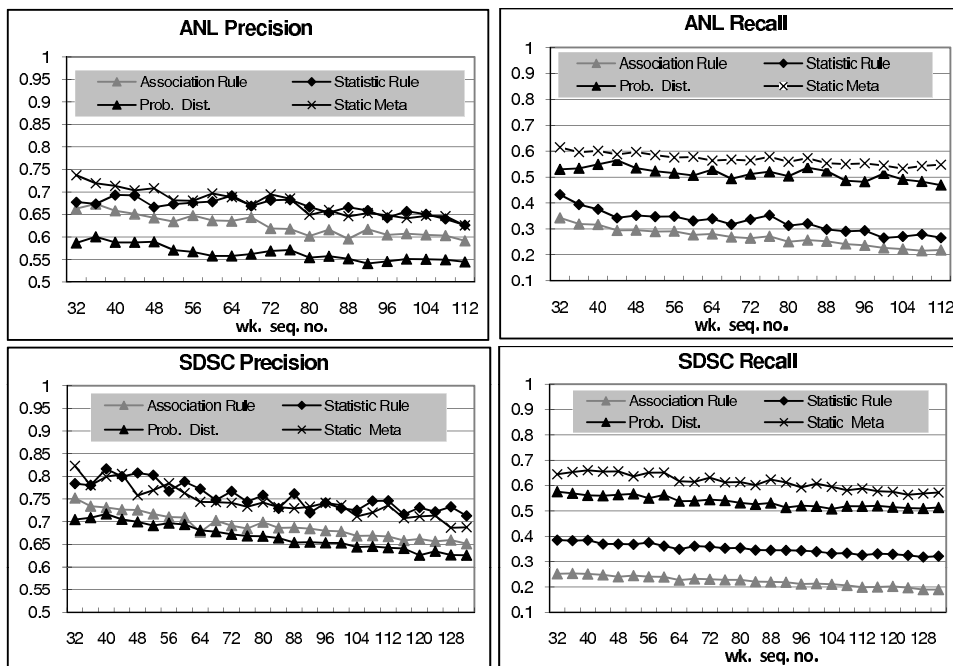


Fig. 7. Meta-learning versus base predictive methods. Each plot contains four curves, representing association rule based learner, statistical rule based learner, probability distribution based learner, and static meta-learner. Here, the “static” means that the meta-learner applies mixture-of-experts ensemble of the base methods without dynamic relearning. It is clear that meta-learning can substantially boost prediction accuracy in terms of both *precision* and *recall*.

rule, and probability distribution). Here, the “static” means that the meta-learner simply applies mixture-of-experts ensemble of the base methods, without dynamic training, retraining and revising. Hence, for both logs, the first six months are used as the training set, and the remaining parts are for testing. The results are shown in Figure 7, where the x-axis shows the sequence number of the week.

First, both *precision* and *recall* decrease as the time goes by, no matter which method is used. The reason is that all these methods are based on a static approach, meaning that they learn the rules on the six-month training set and then use these rules for prediction on the rest of the logs. The rules may well capture failure patterns at the beginning. However, system behavior is dynamically changing. As a result, the established rules become outdated, thereby resulting in lower prediction accuracy as the time goes by.

We make several interesting observations regarding each base method. First, statistical rule base method provides a reasonably good result for *precision*; however, it results in a low value for *recall*. It suggests that this method is only good at discovering certain types of failures which exhibit temporal correlations. Second, association rule base method has the worse results in terms of *recall*. This is mainly due to the fact that while this method well captures causal correlation between non-fatal and fatal events, it is limited by the proportion of fatal events without any precursor warnings (e.g., low *recall* values). Our analysis shows that for both logs, there are a large portion of fatal events (up to 75%) which are not preceded by any precursor non-fatal events. Third, the *recall* results provided by probability distribution based method are quite good (e.g., higher than 0.5 for both logs). Nevertheless, it can introduce many false alarms. The problem of probability distribution based method is that it cannot pinpoint the occurrence times of the failures, thereby giving many false alarms once the elapsed time since the last failure is large enough.

A Venn diagram of these base learners is presented in Figure 8. It shows the numbers of fatal events predicted by these base learners between the 44th and 48th week of the SDSC log. In total, there are 156 fatal events during this period, and 67 of them are captured by multiple base learners. The coverage of each base learner is as follows: association rule based learner 23.7% (37 fatal events), statistical rule based learner 37.2% (58 fatal events), and probability distribution based learner 56.4% (88 fatal events). The diagram clearly shows that it is improper to expect a single method to capture all of failures alone.

*Observation #1: In a large-scale system, there are numerous failure patterns in general; thus a single base learner is unlikely to capture all of them alone.*

Meta-learning can substantially improve *recall*, indicating that meta-learning




Fig. 8. A Venn diagram to show logical relations between association rule based learner (AR), statistical rule based learner (SR), and probability distribution based learner (PD) between the 44th and 48th week of the SDSC log. Each number represents the number of fatal events captured by one or multiple base learners. There are totally 156 fatal events occurred during this period of time. For example, the number six in the intersection of three circles indicates that six fatal events can be predicted by all these base learners.

can improve prediction coverage by capturing various fault patterns. The impact of using meta-learning on *precision* is not as significant as on *recall*, but still non-trivial, especially as compared to association rule based method and probability distribution method. Note that the meta-learner does not modify any of these base methods; instead, it dynamically chooses one base learner for failure forecasting upon each invocation. The benefit of using meta-learner is its ability to form a good integration of these base learners so as to improve both *precision* and *recall*.

*Observation #2: Meta-learning can substantially improve prediction accuracy by intelligently integrating multiple predictive methods without requiring complex system modeling.*

### 5.2.2 Q2: How Much Improvement Is Achieved by the Dynamic Approach?

In this set of experiments, we assess the benefits brought by the dynamic approach. Specifically, we analyze what is the appropriate size for the training set, how often to trigger relearning, whether it is necessary to perform dynamic revising, and how many rules are changed by applying dynamic relearning.

Figure 9 presents the answer to the first question, i.e., what is the appropriate size for the training set? In the figure, each plot consists of four curves: (1) *dynamic-whole* means to train the rules using all the historical data, e.g., in the 32nd week, the data in the previous 31 weeks is used for training; (2) *dynamic-6 mo* means to train the rules using the recent 6 months, e.g., in the 32nd week, the data in the previous 26 weeks is used for training; (3) *dynamic-3 mo* means to train the rules using the recent 3 months, e.g., in the 32nd week, the data in the previous 13 weeks is used for training; and (4) *static* means to use the initial 6-month data as the fixed training set. With the first method, as the time goes by, the training set is gradually increased every four weeks. With the second and third methods, the training set is sliding with the time every four weeks, with a fixed size of six months or three months respectively. With the fourth method, we always use the rules generated in the initial training set for failure prediction, i.e., without any retraining.

Clearly, *dynamic-whole* provides the best results in terms of both *precision* and *recall*, followed by *dynamic-6 mo*. Further, we can see that the accuracy difference between these two methods is generally less than 0.08. As we will show in Section 5.2.4, the overhead introduced by training on a large data set is not trivial. Therefore, in practice we suggest to make a tradeoff between better prediction accuracy and lower computation overhead. For these systems, we suggest to use the most recent 6-month data for dynamic training.

Next, it is shown that by using the *static method* without dynamically adjusting the training set, the prediction accuracy is monotonically decreasing. This

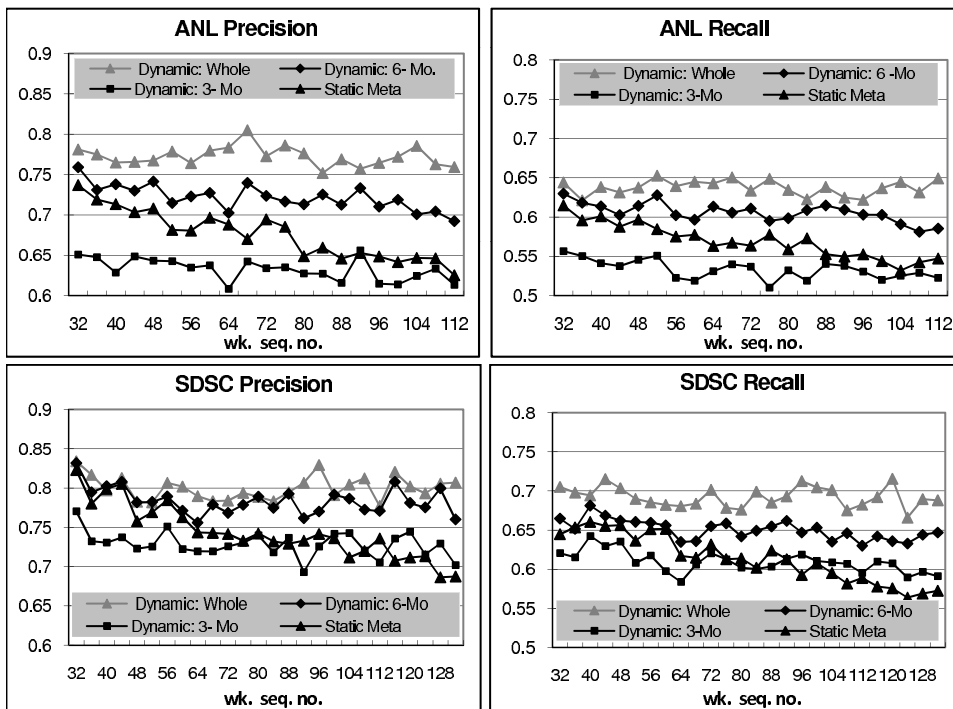


Fig. 9. What is the appropriate size for the training set? Each plot consists of four curves: (1) *dynamic-whole* means to train the rules using all the historical data; (2) *dynamic-6 mo* means to train the rules using the recent 6 months; (3) *dynamic-3 mo* means to train the rules using the recent 3 months; and (4) *static* means to use the first 6-mo training set. Clearly, *dynamic-whole* and *dynamic-6 mo* are the best. Combining the results shown in Table 5, we suggest that dynamic training on the most recent 6-month provides the best balance between prediction accuracy and runtime overhead.

is reasonable as the fixed rule set learned by static meta-learner is unable to adapt to new changes/reconfigurations occurring in the system.

Finally, the results produced by *dynamic-3 mo* are the worst among four different mechanisms. The reason is that this method relies on a limited amount of training data for rule generation and this could substantially limit its capability of discovering sufficient failure patterns for prediction. Nevertheless, as compared to *static*, the prediction results are more stable in the sense that they do not decrease dramatically with time. In summary, the plots indicate that dynamically adjusting the training set is needed; however it is not necessary to re-train the rules on the entire available data, generally the most recent few months are sufficient.

*Observation #3: Learned rules of fault patterns may not be applicable for very long, thus dynamically adjusting the training set is indispensable for good prediction accuracy. In general, using the most recent few months like six months makes a good tradeoff between accuracy and runtime overhead.*

Figure 10 answers the second question, i.e., how often to trigger relearning. It presents the results by using different retraining windows ( $W_R = 2, 4,$  or  $8$  weeks). While prediction accuracy generally remains similar, more frequent retraining can provide better accuracy by up to 0.06. Further, we notice that for the SDSC log, both *precision* and *recall* decrease more than 10% during the 64th week. This is due to the fact that the system went through a major system reconfiguration around this time. As a consequence, failure patterns were changed, thereby resulting in lower prediction accuracy during this period of time. Dynamic training is able to construct a new set of pattern rules. As we can see, both *precision* and *recall* are changed back to the normal range after a few retraining processes. Generally speaking, if the system is constantly changing or its workload is highly dynamic, frequent retraining is necessary, which can help to rapidly build up the effective rules for online prediction.

*Observation #4: The frequency to trigger relearning depends on system characteristics. If the system is highly dynamic, frequent retraining is necessary to maintain satisfactory prediction accuracy.*

The plots also indicate that our method can start to provide a good failure prediction service only after eight weeks of training. For the ANL log, *precision* is between 0.72 – 0.81 and *recall* is ranging between 0.56 – 0.66; for the SDSC log, *precision* is between 0.70 – 0.83 and *recall* is ranging between 0.59 – 0.70. In other words, our method does not need a long training phase to provide an acceptable prediction service. We shall also point out that even when the training set is two weeks (not shown in the figure), the predictor is still capable of capturing more than 43% of failures. In our previous study [14], we have found that runtime adaptive fault management is capable of providing positive



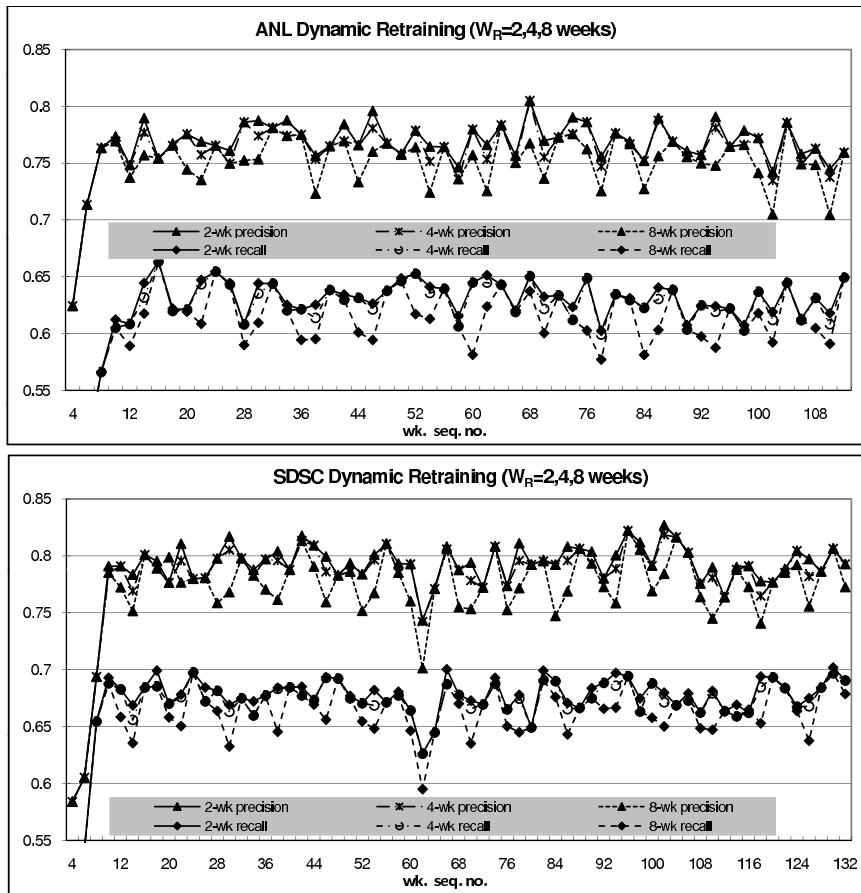


Fig. 10. How often to trigger relearning? The plots present the cases where the rules are re-trained every 2, 4, or 8 weeks, i.e.,  $W_R$  is set to 2,4, or 8. Obviously, more frequent retraining can boost prediction accuracy; however, the improvement is not drastic with the difference less than 0.06 in terms of both *precision* and *recall*.

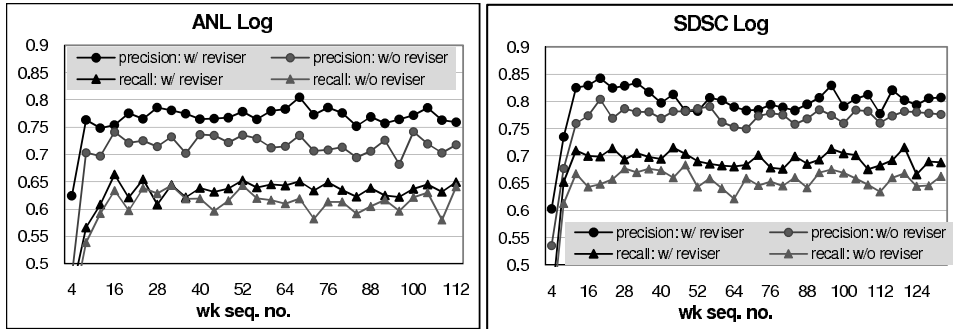


Fig. 11. Is it necessary to conduct dynamic revising? It is clear that dynamic revising can boost prediction accuracy by up to 6%.

performance gain as long as the underlying prediction mechanism can capture 30% of failures. Therefore, our dynamic meta-learning framework is able to serve runtime fault tolerant tool after as few as two weeks of training.

Figure 11 compares the prediction results with and without using the reviser. The plots show that dynamic revising can boost both *precision* and *recall*, and the improvement is up to 6%. As stated in Section 4, failures are rare events. In order to ensure these infrequent events to be analyzed, the parameters like *confidence* and *support* adopted in the association rules are typically chosen without much consideration to the effectiveness of the generated rules, thereby resulting in some rules that may mislead the prediction. The reviser acts like an additional learning process. It works on the candidate rules generated by the meta-learner, and filters out those rules that are not effective on the training set so as to improve prediction accuracy. The results shown in this figure demonstrate the necessity of using dynamic revising.

*Observation #5: Dynamic revising can help improve failure prediction by filtering out bad rules of fault patterns.*

Next, we examine the number of rules changed by using dynamic meta-learning, and the results are presented in Figure 12. Each plot contains four curves representing number of rules unchanged, number of rules added by the meta-learner, number of rules removed by the meta-learner, and number of rules removed by the reviser respectively.

It is clear that the numbers are dynamically changing (i.e., the rules are added or removed) during the operation. Initially when the training starts, there are only dozens of rules, which will be gradually popularized in the following retraining steps. For a period of one year, the knowledge repository will accumulate more than 100 rules for both systems. The number of unchanged rules starts to stabilize for the ANL log around the 70th week—about 140 – 160 rules. However, for the SDSC log, the number keeps increasing (up to 260 rules at the 120th week). In general, the number of rules used for runtime prediction

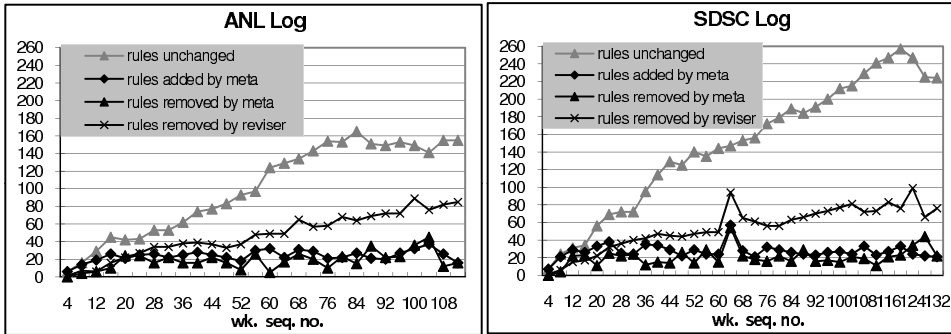


Fig. 12. Number of Rules Changed. We measure the numbers of rules that are unchanged, added by the meta-learner, removed by the meta-learner, and removed by the reviser. These numbers are constantly changing, indicating that the dynamic approach is essential for capturing varying pattern rules.

with the ANL log is between 60 – 115, and it is between 100 – 190 for the SDSC log. The difference between these logs is due to many factors, including system management, workload characteristics, etc. The change rate of rules, i.e., the ratio between changed and unchanged, ranges between 44% – 212%.

Further, we notice a substantial change occurs during the 64th week with the SDSC log, where 57 rules are added and 148 rules are removed. The change is significant since normally about 20 – 30 rules are added and 50 – 80 are removed per retraining. Between the 60th and 64th week, a system reconfiguration occurs. Our method retrains the rules every four weeks, meaning that it extracts a set of rules at the 60th week and then retrains the system at the 64th week. Due to the system change, these two sets of rules are quite different, thereby resulting in significant rule changes. This is consistent with the results shown in Figure 10.

The plots also show that the number of rules removed by the reviser is not trivial, by up to 80. This implies that the reviser can significantly remove non-trivial amount of rules. This result, combined with the information shown in Figure 11, proves that dynamic revising is indispensable for better prediction by removing bad rules.

*Observation #6: The rules of fault patterns are constantly changing during system operation. It further implies that dynamic relearning is essential for maintaining prediction accuracy.*

### 5.2.3 Q3: How Sensitive Is Prediction Accuracy to Prediction Window Size?

Figure 13 presents prediction results by using different prediction windows (5-min, 15-min, 30-min, 45-min, 1-hr, 1.5-hr, and 2-hr). The reason for choosing these durations is based on the results reported in [1,17] and our own experiments with different applications [14]. A time window smaller than 5 minutes

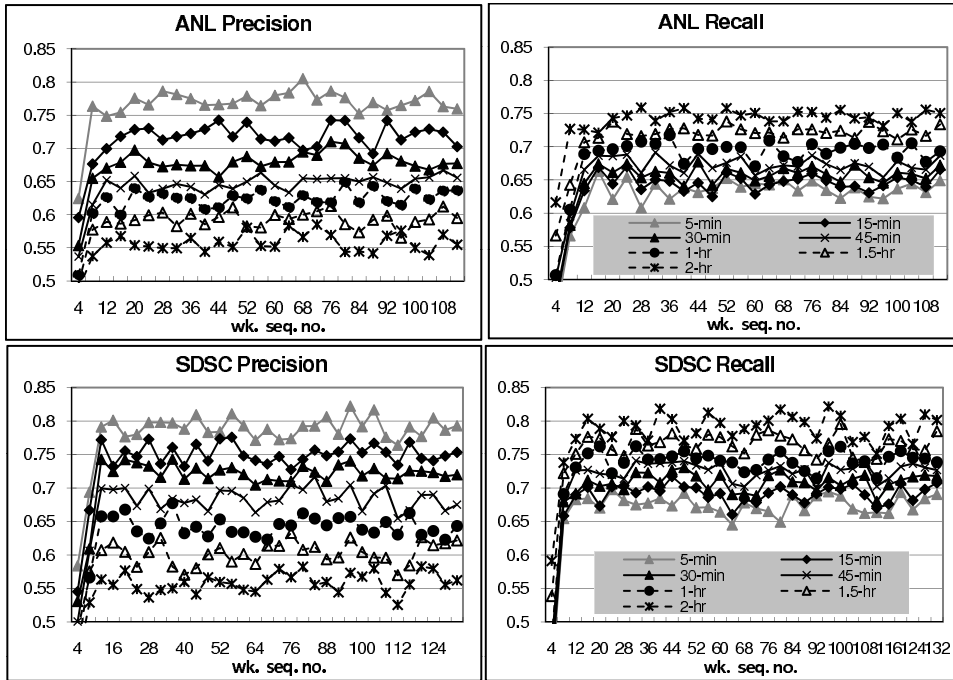


Fig. 13. Impact of Prediction Window. In general, the larger the window is, the higher the *recall* is and the lower the *precision* is.

may become too small for taking preventive action based on the prediction, whereas a time window larger than 2 hour will induce an increased overhead on the system as it will require maintaining the history of all the events for the duration of two hours. Also, the processing/analysis cost of these events for online failure prediction is not trivial.

The trend is obvious: the larger the prediction window is, the higher the *recall* is and the lower the *precision* is. When the prediction window is set to a larger number, it is more likely for the predictor to capture more events, thereby resulting in more chances to find a matching rule of failure pattern. This leads to a higher value for *recall*, meaning the predictor can capture more failures. As an example, when the prediction window is set to two hours, *recall* can be as high as 0.82. On the other hand, if the prediction window is large, it is also likely for the predictor to trigger false alarms due to the growing possibility of catching a misleading rule. For different prediction windows, the difference on *precision* is less than 0.25, and it is about 0.15 in terms of *recall*. Further, for all the cases, both *precision* and *recall* is generally above 0.55.

*Observation #7: The larger the prediction window is, the higher the recall is and the lower the precision is.*

Table 5  
 Operation Overhead (in minutes) as a Function of Training Size.

Training Size	Rule Generation				Rule Matching
	Stat_Rule	Asso_Rule	Prob_Dist	Ensemble & Revise	
3 mo	< 1	1	< 2	1	< 1
6 mo	< 1	2	< 2	1	< 1
12 mo	< 1	3	< 2	2	< 1
18 mo	< 1	4	< 2	2	< 1
24 mo	< 1	5	< 2	3	< 1
30 mo	< 1	6	< 2	4	< 1

#### 5.2.4 Q5: How Much Runtime Overhead Is Introduced?

Operation overhead depends on the size of training set. Table 5 summarizes the overheads as a function of training data, in which the overhead is classified into two parts, namely rule generation overhead and rule matching overhead. These times are measured on a local PC configured with a 1.6GHz Intel Pentium processor and 768 MB memory. Obviously, the overhead could be less when a more powerful PC is used.

The overhead mainly comes from rule generation, while the rule matching process (i.e., the event-driven predictor) is trivial, usually in dozens of seconds. As shown in the table, when the training set is set to 6 months (half of a year), the rule generation may take 6.0 minutes; and it can increase to 13 minutes when the training set is set to 30 months (two and a half year). Note that the rule generation process can be conducted in parallel when the production system is in operation, therefore this cost should not be counted into the actual runtime overhead for failure prediction. The actual runtime overhead introduced by the event-driven predictor is normally less than 1.0 minute. Thus, we believe the framework is feasible as a runtime prediction mechanism. Combining the results shown in this table and in Figure 9, we suggest that dynamic meta-learning on the recent 6 months is practical and time efficient.

*Observation #8: The runtime overhead is trivial (e.g., in dozens of seconds), while the major overhead introduced by rule generation can be conducted in parallel when the target machine is in operation.*

## 6 Related Work

Recognizing the importance of fault tolerance, the community has paid much attention to failure prediction. Existing predictive approaches can be broadly classified as *model-based* or *data-driven* methods. Model-based approach derives a probabilistic or analytical model of the system and triggers a warning when a deviation from the model is detected [7,13,25]. Examples include an adaptive statistical data fitting method called MSET presented in [26], Semi-Markov reward models described in [34], and a naive Bayesian based model to predict disk drive failures [12]. In large-scale systems, errors may propagate from one component to other component, which is commonly addressed by developing fault propagation models (FPM) [35]. *While model-based methods are effective for forecasting some failures, they seem too complicated to be practical for failure prediction in large-scale systems composed of tens of thousands of components.*

A data-driven method, such as using data mining techniques, attempts to learn failure patterns from historical data for failure prediction, without constructing an accurate model ahead of time. These methods extract fault patterns from system normal behaviors, and detect abnormal observations based on the learned knowledge without assuming a priori model ahead of time. For example, the group at the RAD laboratory has applied statistical learning techniques for failure diagnosis in Internet services [27,33]. The SLIC (Statistical Learning, Inference and Control) project at HP has explored similar techniques for automating fault management of IT systems [36]. Sahoo et al. have applied association rules to predict failure events in a 350-node IBM cluster [23]. In [16,17], Liang et al. have examined several data mining and machine learning techniques for failure forecasting in a Blue Gene/L system. Other representative works include system log analysis [4,41] and a prediction framework for networked systems [37].

*While this paper is built upon existing studies, it distinguishes from the above studies at several aspects.* First, unlike existing studies focusing on one specific predictive method, this paper presents a dynamic meta-learning framework to dynamically integrate existing predictive methods for better prediction. In this study, we have examined three widely-used predictive methods, namely association rule based learner [23,28], statistical rule based learner [16], and probability distribution based learner [4] in the framework. We believe other predictive methods can be easily incorporated into our framework. Second, this study emphasizes dynamic training and learning, which is rarely examined in the literature. By means of real system logs from production systems, we have demonstrated that dynamic relearning is essential to capture behavior changes during system operation. By examining our framework in various ways, we have shown that using the most recent few months like six months

makes a good tradeoff between accuracy and runtime overhead. Next, unlike offline log analysis studies, our prediction is event-driven, meaning that our framework triggers a warning on the occurrence of events during system operation. An event-driven approach is well suited for online failure prediction. Last but not the least, in addition to presenting the key techniques for boosting prediction accuracy, we have also systematically analyzed our framework and answered several key questions commonly raised in failure prediction. It provides a deep insight into failure prediction in large-scale systems. To the best of our knowledge, we are among the first to comprehensively evaluate the impact of different factors in failure forecasting.

## 7 Summary

In this paper, we have presented a dynamic meta-learning prediction engine for large-scale systems. Recognizing problems in failure prediction, our prediction mechanism relies on two key techniques to improve prediction accuracy in real systems. *Meta-learning* is applied to boost prediction accuracy by integrating multiple predictive methods, while a dynamic approach is employed to train the rules of failure patterns at runtime. Our prediction mechanism does not require a long training phase by dynamically adjusting the training set during system operation. Further, it can adapt to system changes, even after a major system reconfiguration. Our case studies with real system logs have demonstrated its effectiveness with a good accuracy, e.g., capturing up to 82% of failures. The studies have also shown that the proposed mechanism is practical and well suited for forecasting failures in real systems.

Our study has some limitations that remain as our future work. First, in the current design, the prediction window size is fixed. Our on-going work includes adaptively changing this window size such that the system can automatically tune its size to reduce the training cost, without sacrificing the prediction accuracy. Second, we plan to examine other data mining methods, such as decision tree and neural network, to popularize our base learners. We will also investigate other ensemble learning techniques to improve the meta-learner. Finally, more case studies with a variety of HPC systems will be conducted. Although our case studies focus on the Blue Gene/L systems, we believe the proposed mechanism is applicable to other systems. For the systems that do not have an error checking and logging facility, the first step is to develop a monitoring tool which is capable of gathering fault-related information from various system components and archive the information in a centralized repository. The proposed framework can be easily extended to these systems by linking to their event repositories.

## Acknowledgment

Zhiling Lan is supported in part by US National Science Foundation grants CNS-0834514, CNS-0720549, CCF-0702737, and a TeraGrid Compute Allocation. Susan Coghlan and Rajeev Thakur are supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357. We would like to thank John White at Revision<sup>3</sup> Company and Eva Hocks at San Diego Supercomputer Center for the discussion of the SDSC system log. Some preliminary results of this work were presented in [8,9].

## References

- [1] E. Elnozahy and J. Plank, Checkpointing for Peta-Scale Systems: A Look into the Future of Practical Rollback-Recovery, *IEEE Trans. on Dependable and Secure Computing*, vol. 1(2), pp. 97-108, 2004.
- [2] G. Hoffmann, F. Salfner, and M. Malek, Advanced Failure Prediction in Complex Software Systems, *Proc. of SRDS'04*, 2004.
- [3] A. Gara, M. Blumrich, D. Chen, G. Chiu, P. Coteus, M. Giampapa, R. Haring, P. Heidelberger, D. Hoenicke, G. Kopcsay, T. Liebsch, M. Ohmacht, B. Steinmacher-Burow, T. Takken, P. Vranas, Overview of the Blue Gene/L System Architecture, *IBM J. Res. & Dev.*, vol. 49(2/3), 2005.
- [4] B. Schroeder and G. Gibson, A Large-scale Study of Failures in High Performance Computing Systems, *Proc. of DSN'06*, 2006.
- [5] W. Peng, T. Li and S. Ma, Mining Logs Files for Computing System Management, *Proc. of ICAC'05*, 2005.
- [6] R. Gioiosa, J. Sancho, S. Jiang, F. Petrini, and K. Davis, Transparent Incremental Checkpointing at Kernel Level: A Foundation for Fault Tolerance for Parallel Computers, *Proc. of SC'05*, 2005.
- [7] A. Goyal, S. Lavenberg, and K. Trivedi, Probabilistic Modeling of Computer System Availability, *Annals of Operations Research*, 1987.
- [8] P. Gujrati, Y. Li, Z. Lan, R. Thakur, and J. White, A Meta-learning Failure Predictor for Blue Gene/L Systems, *Proc. of ICPP'07*, 2007.
- [9] J. Gu, Z. Zheng, Z. Lan, J. White, E. Hocks, and B. Park, Dynamic Meta-Learning for Failure Prediction in Large-Scale Systems: A Case Study, *Proc. of ICPP'08*, 2008.
- [10] Y. Zhang, M. Squillante, A. Sivasubramaniam and R. Sahoo, Performance Implications of Failures in Large-Scale Cluster Scheduling, *Proc. of Workshop on Job Scheduling Strategies for Parallel Processing*, 2004.



- [11] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*, 2nd Edition, Morgan Kaufmann, 2006.
- [12] G. Hamerly and C. Elkan, Bayesian Approaches to Failure Prediction for Disk Drives, *Proc. of ICML'01*, 2001.
- [13] J. Hellerstein, F. Zhang, and P. Shahabuddin, A Statistical Approach to Predictive Detection, *Computer Networks: The International Journal of Computer and Telecommunications Networking*, 2001.
- [14] Z. Lan and Y. Li, Adaptive Fault Management of Parallel Applications for High Performance Computing, *IEEE Trans. on Computers*, 57(12), pp. 1647-1660, 2008.
- [15] Y. Li, Z. Lan, P. Gujrati, and X. Sun, Fault-Aware Runtime Strategies for High Performance Computing, *IEEE Trans. on Parallel and Distributed Systems*, vol. 20(4), pp. 460-473, 2009.
- [16] Y. Liang, Y. Zhang, A. Sivasubramaniam, M. Jette, and R. Sahoo, Blue Gene/L Failure Analysis and Models, *Proc. of DSN'06*, 2006.
- [17] Y. Liang, Y. Zhang, H. Xiong and R. Shaoo, Failure Prediction in IBM BlueGene/L Event Logs, *Proc. of ICDM'07*, 2007.
- [18] A. Oliner, L. Rudolph, and R. Sahoo, Cooperative Checkpointing Theory, *Proc. of the International Parallel and Distributed Processing Symposium (IPDPS)*, 2006.
- [19] A. Oliner and J. Stearly, What Supercomputers Say: A Study of Five System Logs, *Proc. of DSN'07*, 2007.
- [20] A. Oliner, R. Sahoo, J. Moreira, M. Gupta, and A. Sivasubramaniam, Fault-Aware Job Scheduling for Blue Gene/L Systems, *Proc. of IPDPS'04*, 2004.
- [21] R. Polikar, Ensemble Based Systems in Decision Making, *IEEE Circuits and Systems Magazine*, vol.6(3), 2006.
- [22] D. Reed, C. Lu, and C. Mendes, Big Systems and Big Reliability Challenges, *Proc. of Parallel Computing*, 2003.
- [23] R. Sahoo, A. Oliner, I. Rish, M. Gupta, J. Moreira, and S. Ma, Critical Event Prediction for Proactive Management in Large-scale Computer Clusters, *Proc. of SIGKDD'03*, 2003.
- [24] M. Schulz, G. Bronevetsky, R. Fernandes, D. Marques, K. Pingali, and P. Stodghill, Implementation and Evaluation of a Scalable Application-level Checkpoint-Recovery Scheme for MPI Programs, *Proc. of SC'04*, 2004.
- [25] K. Trivedi and K. Vaidyanathan, A Measurement-based Model for Estimation of Resource Exhaustion in Operational Software Systems, *Proc. of ISSRE'99*, 1999.
- [26] K. Vaidyanathan and K. Gross, MSET Performance Optimization for Detection of Software Aging, *Proc. of ISSRE*, 2003.

- [27] M. Chen, A. Zheng, J. Lloyd, M. Jordan, and E. Brewer, Failure Diagnosis Using Decision Trees, *Proc. of ICAC'04*, 2004.
- [28] R. Vilalta and S. Ma, Predicting Rare Events in Temporal Domains, *Proc. of ICDM'02*, 2002.
- [29] C. Wang, F. Mueller, C. Engelmann and S. Scott, A Job Pause Service under LAM/MPI+BLCR for Transparent Fault Tolerance, *Proc. of IPDPS'07*, 2007.
- [30] The TOP500 Supercomputer Site, available at <http://www.top500.org/>
- [31] The Blue Gene System at SDSC, available at <http://www.sdsc.edu/us/resources/bluegene/>
- [32] The Blue Gene System at ANL, available at <http://www.bgl.mcs.anl.gov/>
- [33] P. Bodik, G. Friedman, L. Biewald, H. Levine, G. Candea, K. Patel, G. Tolle, J. Hui, A. Fox, M. Jordan, and D. Patterson, Combining Visualization and Statistical Analysis to Improve Operator Confidence and Efficiency for Failure Detection and Localization, *Proc. of The 2nd IEEE International Conference on Autonomic Computing (ICAC '05)*, 2005.
- [34] S. Garg, A. Puliafito, and K. Trivedi, Analysis of Software Rejuvenation Using Markov Regenerative Stochastic Petri Net, *Proc. of 6th International Symposium on Software Reliability Engineering*, 1995.
- [35] M. Steinder and A. Sethi, A Survey of Fault Localization Techniques in Computer Networks, *Science of Computer Programming*, vol. 53, 2004.
- [36] I. Cohen and J. Chase, Correlating Instrumentation Data to System States: A building Block for Automated Diagnosis and control, *Proc. of OSDI'04*, 2004.
- [37] S. Fu and C. Xu, Exploring Event Correlation for Failure Prediction in Coalitions of Clusters, *Proc. of SC'07*, 2007.
- [38] P. Hargrove and J. Duell, Berkeley Lab Checkpoint/Restart (BLCR) for Linux Clusters, *Proc. of SciDAC 2006 (publication LBNL-60520)*, 2006.
- [39] M. Joshi, R. Agarwal, and V. Kumar, Mining Needle in a Haystack: Classifying Rare Classes via Two-phase Rule Induction, *Proc. of SIGMOD'01*, 2001.
- [40] G. Weiss, Mining with Rarity: a Unifying Framework, *ACM SIGKDD Explorations*, vol. 6(1):719, 2004.
- [41] A. Oliner and J. Stearley, What Supercomputers Say: A Study of Five System Logs, *Proc. of the International Conference on Dependable Systems and Networks (DSN)*, 2007.
- [42] J. Hansen and D. Siewiorek, Models for Time Coalescence in Event Logs, *Proc. of Fault-Tolerant Computing (FTCS)*, 1992.
- [43] M. Buckley and D. Siewiorek, Comparative Analysis of Event Tupling Schemes, *Proc. of Fault-Tolerant Computing (FTCS)*, 1996.