

NWSLite: A Non-Parametric Prediction Utility for Resource-Restricted Devices

Selim Gurun Chandra Krintz Rich Wolski
Computer Science Department
University of California, Santa Barbara
{gurun,ckrintz,rich}@cs.ucsb.edu

Abstract

Computation off-loading, i.e., remote execution, has been shown to be effective for extending the computational power and battery life of resource-restricted devices, e.g., hand-held, wearable, and pervasive computers. Remote execution systems must predict the cost of executing both locally and remotely to determine when off-loading will be most beneficial. These costs however, are dependent upon the execution behavior of the task being considered and the highly-variable performance of the underlying resources, e.g., CPU (local and remote), bandwidth, and network latency. As such, remote execution systems must employ sophisticated, prediction techniques that accurately guide computation off-loading. Moreover, these techniques must be efficient, i.e., they cannot consume significant resources, e.g., energy, execution time, etc., since they are performed on the mobile device.

In this paper, we present NWSLite, a computationally efficient, highly accurate prediction utility for mobile devices. NWSLite is an extension to the Network Weather Service (NWS), a dynamic forecasting toolkit for adaptive scheduling of high-performance Computational Grid applications. We significantly scaled down the NWS to reduce its resource consumption yet still achieve accuracy that exceeds that of extant remote execution prediction methods. We empirically analyze and compare both the prediction accuracy and the cost of NWSLite and a number of different forecasting methods from existing remote execution systems. We evaluate the efficacy of the different methods using a wide range of mobile applications and resources.

1 Introduction

Remote execution is emerging as a promising technique to extend the computational power and battery life of resource-restricted environments, e.g., hand-held, wearable, and pervasive computers. Using remote execution,

parts of program execution are off-loaded from battery-powered mobile devices to wall-powered, higher performance platforms. As such, remote execution can significantly extend the usefulness of devices by enabling execution of a wide-range of resource-intensive applications, e.g., augmented reality, natural language translation, feature recognition, collaborative computing, etc., in a mobile environment.

Key to the successful implementation of remote execution, is *prediction*. That is, we must predict when the cost of performing remote execution will not outweigh its benefits. For example, if a task will take significantly less time if executed remotely due to superior resource performance, the system should off-load the task. Alternately, if the cost of remote execution will exceed that of local execution, e.g., if significant amounts of data must be transferred to perform the task or if the remote host is busy, then the system should perform the operation locally.

The cost of a remotely executed operation consists of the time to transfer data (and possibly code) from the device to the target, the time to transfer result information, e.g., data, status, rendered graphics, etc., from the target back to the device, and the time to execute the operation at the target. The cost of a locally executed operation is the local execution time. These costs can be decomposed or translated to consider other metrics, e.g., battery consumption, response time, application fidelity. However, regardless of their form, these values must reflect what the costs *will be* when the operation is eventually performed.

To complicate matters, each of these costs is dependent upon the highly variable performance available from the underlying resources as well as the application input data. Network bandwidth and latency dictate the time required for communication and CPU availability on both the device and the target impacts local and remote execution time, respectively. Moreover, execution cost is dependent upon the length of time the application tasks will execute, which is commonly dependent upon program inputs.

To predict these costs, extant remote execution systems

employ statistical techniques that use past behavior to predict the future [21, 8, 10, 2]. The goal of these systems is to enable high prediction accuracy. That is, they attempt to reduce prediction error – the difference between predicted values and the measurement values that they predict. Techniques that result in large prediction errors can cause incorrect decisions to be made about the best execution choice for the device. As such, prediction accuracy plays an important role in the performance of a remote execution system.

However, another characteristic of prediction techniques that is often not considered, is computational cost – performing the prediction itself consumes device resources. This cost can be high (particularly in terms of power consumption) since statistical techniques commonly use floating-point operations and most mobile devices do not implement a floating-point co-processor. Instead, they rely on software emulation of floating-point instructions making them highly resource-intensive operations.

In this paper, we consider both the accuracy and cost of commonly used forecasting technologies in remote execution settings. We evaluate techniques that are currently used for remote execution and present a novel resource performance prediction utility for resource-restricted devices, called *NWSLite*. *NWSLite* is a low-cost, yet, highly accurate prediction service that is an extension of the Network Weather Service (NWS), a resource performance measurement and prediction toolkit originally developed for scheduling high-performance, scientific applications in Computational Grid [11] environments [27, 30, 4, 26, 25]. We modified the NWS forecasting model to reduce its resource consumption footprint to enable its use in a mobile setting.

NWSLite can be incorporated into any mobile framework that uses prediction. It makes non-parametric forecasts of any resource for which measurement values can be supplied. As such, we can use it for prediction of CPU load, memory availability, and network bandwidth and latency, as well as file I/O and execution time of an application’s operations (tasks).

In this study, we empirically compare both the accuracy and cost of *NWSLite* to the original NWS as well as to two prediction algorithms (described in [21] and [19]) currently used for prediction in remote execution frameworks. We analyze these performance characteristics for a wide range of applications and resources: application execution time, availability, wired-network bandwidth and latency, and wireless bandwidth. Our results show that *NWSLite* enables prediction accuracy that in many cases significantly exceeds that of the predictors to which we compare. In addition, it consumes significantly less computational resources than its predecessor (the original NWS).

In summary, the contributions we make with this paper are:

- A light-weight, highly-accurate, prediction utility for resource-restricted environments,
- A predictive framework that uses algorithms that are *non-parametric*, i.e., they do not require externally determined, resource-specific parameters or manual fine-tuning,
- A general utility for mobile devices that can be used to forecast efficiently the performance of a wide range of resources,
- A comparison of the prediction accuracy and execution cost of the *NWSLite* to popular predictive techniques, and
- An empirical evaluation that shows that *NWSLite* can significantly improve prediction accuracy for remote execution.

In the following sections, we overview the use of resource performance prediction to facilitate remote execution. In Section 3, we describe the prediction methodologies for two popular remote execution systems to which we compare our work. We then detail the design and implementation of *NWSLite* in Section 4 and present an empirical evaluation of its efficacy both in terms of prediction accuracy and resource consumption (Section 5). Finally, we conclude in Section 6.

2 Background

Remote execution is a popular technique that is used to extend the computational capability of mobile, resource-restricted, devices [8, 22, 21, 15, 32]. Figure 1 depicts the general design of a remote execution system. Using remote execution, application tasks are off-loaded from battery-powered mobile devices to wall-powered, higher-performance servers.

To decide whether a particular task should be off-loaded, a remote execution system must first compute the *demand* of the application task. In this work, we define demand as task execution time but additional constraints can be applied to this value, e.g., response time, computational fidelity, power consumption, etc., according to the overall goals of the system.

To determine how best to accommodate demand, a remote execution system must evaluate how best to employ its *supply* – the set of resources, local and remote, that it has available to it for task execution. The system computes whether computation off-loading will be beneficial, according to its set of constraints, using a cost model. When cost of local execution exceeds that of remote execution (including all necessary communication), the system off-loads work to the server. The cost model must




Figure 1: Components of a typical remote execution system. The decision process includes forecasting the available resource supply both at the client and server and application resource demand.

consider both the task execution characteristics as well as the highly-variable performance of the underlying resources that dictate computation and communication performance. A general cost evaluation function can be specified as:

$$F_L(T_{ET_M} * CPU_M) \leq F_R\left(\frac{T_{input}}{BW} + (T_{ET_R} * CPU_S) + \frac{T_{output}}{BW} + (HS * LAT)\right)$$

where the local cost is F_L and the remote cost is F_R . The local cost is some function (computational fidelity, execution time, response time, battery power) of the time to execute the task on the mobile device (T_{ET_M}) given the available fraction of the CPU on the device (CPU_M). The remote cost is some function of the sum of the time required for four constituent operations:

1. The time required for transfer of the input data, and any needed program code (T_{input}) given the available bandwidth (BW) between the device and server;
2. The execution time at the server given the fraction of CPU available at the server ($T_{ET_R} * CPU_S$);
3. The time for transfer of results, e.g., data, status, rendered graphics, etc., back to the device given the available bandwidth between the server and device ($\frac{T_{output}}{BW}$); and;
4. The time for any handshake protocol (HS) to establish remote execution.

Since (4) commonly consists of very short packet communication between the device and the server, the handshake operation is impacted by the latency in the network link between the client and server ($HS * LAT$).

Remote execution systems must compute what the cost of remote and local execution *will be* when the task is

eventually executed. That is, these systems must *predict* this cost *on the device* to determine when to off-load. As such, remote execution systems must employ sophisticated forecasting techniques to enable accurate prediction of the cost of remote (and local) execution. We describe extant prediction technologies used for remote execution and other mobile device optimizations in the next section.

3 Prediction Algorithms for Remote Execution

Two advanced, prediction-based, remote execution systems to which we compare our work, are Spectra [8] and the remote processing framework (RPF) by Rudenko et al [21]. We briefly describe these systems in this section and discuss the predictors that each employs.

3.1 Spectra/Odyssey

Spectra[8, 9] is the remote execution component of Aura, a pervasive computing environment. Spectra allows the user to dictate goals, such as minimizing energy use, and then tries to achieve this goal by means of local, remote or hybrid, i.e., partially remote, execution decisions. To estimate the behavior of an application, Spectra uses the algorithms developed for Odyssey, another component of Aura.

Odyssey uses *multi-computational fidelity* to dynamically trade-off resource demand for application quality. The fidelity adjustments are performed at the granularity of an *operation*, which is the smallest unit of execution that can be perceived by the user as a response to a request [18].

Since the performance of off-loading is dependent upon the underlying resource availability, Odyssey couples task demand with resource supply. That is, Odyssey makes forecasts of the underlying resource performance, of CPU, memory, and network bandwidth and latency. It uses these predictions with those from the fidelity resource functions to identify opportunities for dynamic off-loading.

Odyssey estimates CPU availability, by first assuming that CPU cycles are evenly distributed among all processes. The CPU model is:

$$S_{cpu} = \frac{P}{N + 1}$$

where P is the processor clock speed, N is the number of runnable processes and S_{cpu} is the available CPU cycles. It uses a smoothing filter to estimate the number of processes in the next interval:

$$N_{t+1} = \alpha N_t + (1 - \alpha)n(p) \quad (1)$$

In this equation n is a function of observed number of processes in the current interval and defined as:

$$n(i) = \begin{cases} n_r - 1 & \text{If } p \text{ is runnable} \\ n_r & \text{Otherwise} \end{cases}$$

where n_r is the number of runnable processes at time t .

To predict network bandwidth and latency, Odyssey uses a different prediction model based on exponential smoothing:

$$new = \gamma(measured) + (1 - \gamma)old \quad (2)$$

Odyssey uses $\gamma = 0.75$ for round trip time and $\gamma = 0.875$ for throughput.

An exponential smoothing predictor with a constant gain factor does not adapt dynamically to the changes in the system. To address this issue Kim et.al. [14] developed a flip-flop filter. The flip-flop filter employs two exponential smoothing predictors: one with $\gamma = 0.1$ and the other with $\gamma = 0.9$. Using a small gain factor, the former can respond to changes more quickly than the latter one. However, a small gain is more susceptible to a transient noise. As such, by using both gain factors, the filter is able to achieve the benefits from both while avoiding the limitations of each. Our prediction methodology, described in the next section, takes a similar but more general approach.

Finally, to estimate the resource consumption of applications, Odyssey and Spectra use an online-learning predictor. The predictor maintains data-specific coefficients that are used to model the cost behavior of the application for a particular dataset. However, computing the initial values of coefficients requires off-line training. The online-learning algorithm updates coefficients using recursive least squares regression with exponential decay. Due to exponential decay characteristics, more weight is given to the recent observations.

This technique works well for applications that render pictures as the user or camera scans a scene, e.g., as is done in augmented reality programs, due to the slow, smooth transitions between scenes. That is, the resource consumption behavior for the generation of a scene will be similar to that of a neighboring scene. A limitation of this type of statistical technique is that numerical computation errors can accumulate after each recursion causing algorithm to become unstable and diverge [5].

The recursive least squares method is an efficient way to predict the value y when it is dependent on a set of parameters x such that $y = Ax + w$, and w is the measurement error or noise. The general formula is given by:

$$A_k = A_{k-1} - P_k \{ x_k x_k^T A_{k-1} - x_k y_k \}$$

$$P_k = \{ P_{k-1} - P_{k-1} x_k [\alpha + x_k^T P_{k-1} x_k]^{-1} x_k^T P_{k-1} \} / \alpha$$

where α is the decay factor and y_k is the measurement at time k . In the equation above, y_{k+1} is predicted by $A_{k+1} x_k$. The P_k matrix is commonly referred to as the *history* or *filtering factor* [31].

The value of the exponential decay factor determines the agility of the method. A smaller value increases the responsiveness, but decreases the capability of filtering out any noise. Odyssey uses a decay factor of 0.5, which makes it very agile. Our observations show that, with such a small value the method can become unstable and diverge if the observed data is noisy or moderately variable. Based on our experimentation using several different parameters on a large dataset, we found that a decay factor of 0.8 is more appropriate. We refer to this technique using this decay factor as *LSQ* in our experimental results.

3.2 Remote Processing Framework (RPF)

The remote processing framework (RPF) uses a different execution model than that of Spectra and other remote execution systems [21]. RPF models a single metric – power consumption – of the executing task. That is, it collects history data on the power consumption of previous tasks and uses it to predict the consumption of future tasks. The comparison between local and remote task execution is performed on the device itself. All code and data are shared and kept synchronized between the device and server. The server simply processes the tasks transmitted by the device and returns the results.

The device and server interact via a simple handshake protocol. During this protocol, the device verifies that the server has enough resources and the appropriate software packages. Following this, the device predicts whether off-loading will reduce power consumption. The RPF system uses a *smoothing filter* to make its forecasts given prior history. In particular, the RPF cost model is the *parameterized function*:

$$f_{n+1} = (1 - \alpha) * \frac{\sum_{i=n-k}^n v_i}{k} + \alpha * f_n \quad (3)$$

where f_i is the forecasted value, v_i is the measured value and i is the measurement index. α and k determine how conservative the forecaster is: A small k combined with a large α will result in higher responsiveness to recent changes.

RPF is parameterized with two parameters, k and α . Unfortunately, the authors in [21] do not specify default values for these parameters, nor do they discuss the values that they used. To enable our evaluation and comparison to this technique, we empirically evaluated several different combinations of parameters for a large set of data. We then selected best-performing set across all experiments. For the data we detail in this study, the best overall parameterization is $k = 20$ and $\alpha = 0.0$. With this par-

parameterization, RPF becomes a “sliding window average” over a fixed window (size k) of previous data history. The data for different parameterizations can be found in the appendix of this paper.

Note that the RPF smoothing filter (Equation 3) is the same as the equation for CPU prediction in Odyssey (Equation 1) when $k = 1$. In addition, the smoothing filter is the same as the bandwidth and latency prediction function in Odyssey (Equation 2) when $k = 1$ and $\alpha = 1 - \gamma$.

4 NWSLite

Odyssey and RPF employ parameterized prediction methodologies to forecast the cost of local and remote execution. The parameters are identified through empirical evaluation and are specific not only to the executing application but also to the individual tasks. In addition, different types of resource data (task execution time and characteristics, and network and CPU performance) require different predictors or different parameterization of the same predictors to be effective. Moreover, prediction of task execution behavior in Odyssey requires extensive off-line training. Our approach to the problem of remote execution cost forecasting takes a different approach. Specifically, it is one that is *non-parametric, automatic, and completely general*.

To enable this, we developed NWSLite, an accurate and efficient prediction utility for mobile devices. NWSLite is an extension of the Network Weather Service [29], a freely available toolkit [20], originally developed for Computational Grid computing [11, 3]. The Computational Grid is a computing paradigm for the development of software systems that enables dynamic acquisition of resources from a heterogeneous and non-dedicated resource pool. To extract performance from these systems, application schedulers must use predictions of future resource behavior to determine how the application can best use the available resources.

The NWS operates a distributed set of performance sensors, from which it periodically, and unobtrusively, collects performance measurements, applies a set of statistical forecasting techniques to individual performance histories, and generates forecast reports for the resources being monitored, which it disseminates via a number of different APIs in near-real-time [30]. Currently, the NWS provides sensors for end-to-end TCP/IP bandwidth and latency, available CPU and memory, battery power, and disk storage, and is used in a large number of different of Grid technologies.

NWS prediction uses a mixture-of-experts approach to prediction instead of relying on a single model. It implements a large set of models, each having its own parameterization. Given a performance history of observed mea-

	Name	Average Cost
1	Last Value	0
2	Running Mean	3
3	5% Exp Smooth	3
4	10% Exp Smooth	3
5	15% Exp Smooth	3
6	20% Exp Smooth	3
7	30% Exp Smooth	3
8	40% Exp Smooth	3
9	50% Exp Smooth	3
10	75% Exp Smooth	3
11	90% Exp Smooth	3
12	5% Exp Smooth, with 0.1% trend	10
13	10% Exp Smooth, with 0.1% trend	10
14	15% Exp Smooth, with 0.1% trend	10
15	20% Exp Smooth, with 0.1% trend	10
16	30% Exp Smooth, with 0.1% trend	10
17	Median Window 31	88
18	Median Window 5	16
19	Sliding Median Window 31	124
20	Sliding Median Window 5	26
21	30% Trimmed Median Window 31	106
22	30% Trimmed Median Window 51	169
23	Adaptive Median Window 5-21	171
24	Adaptive Median Window 21-51	455

Table 1: NWS Forecasters and the Approximate Costs of Each. We show cost in column three as the number of floating point operations performed.

surement values, it generates a forecast for each measurement. NWS ranks each predictor by computing the prediction errors (the difference between measured and forecasted values). Each time a forecast is requested, NWS recalculates the ranking across all predictors using the most recent history and chooses the most-accurate model. The implementation of NWS that we extended uses the 24 prediction models shown in Table 1.

This mixture-of-experts method achieves its accuracy by employing wide range of statistical models, each of which may be most appropriate at a given time, for a given resource. This method also has other important advantages. First, even though the individual NWS models may be parametric, the overall system is not. The only input to the system is the measurement history. Second, NWS can easily adjust itself to changes in the characteristics of the data series by switching to another model. Third, it can be used on any type of data for which measurements can be made. There is no distinction between CPU availability and network bandwidth, for example.

Because the NWS was originally designed to support performance applications in wired settings, its designers put a premium on speed and extensibility. As such, it consumes significant resources to perform a single prediction since many models are evaluated at once. The Av-

erage *Cost* column shows the number of floating point instructions executed for each predictor (all are computed for each forecast made) on average. To enable its use in resource-restricted environments, we have significantly reduced this consumption without sacrificing appreciable accuracy. To this end, we first evaluate the cost of NWS prediction in terms of dynamic floating point instructions.

Given a history of measurements and their predicted values, prediction error can be defined using the square of the errors:

$$E = \sum_{i=1}^n (f_i - v_i)^2 \quad (4)$$

where f_i is the output of the predictor, v_i is the measurement and n is the length of history.

Since the NWS uses a mixture-of-experts approach, all forecasters are invoked (logically in parallel) and a single winner is selected and used for the next estimation. We use zero-one integer variables $s_{i,j}$ to denote the winning forecaster:

$$s_{i,j} = \left\{ \begin{array}{ll} 1 & \text{If model } j \text{ is used to predict} \\ & \text{measurement } i \\ 0 & \text{Otherwise} \end{array} \right\} \quad (5)$$

Specifically, if $s_{i,j}$ is 1, the i^{th} forecast is made using predictor j . If $s_{i,j}$ is 0, the predictor is not the winner for the i^{th} forecast. If we set k to be the number of models in NWS, using (4), we can formulate prediction error of NWS as:

$$E = \sum_{i=1}^n \sum_{j=1}^k (f_i - v_i)^2 s_{i,j} \quad (6)$$

Similarly, we can compute the cost of using the winning forecasters (in terms of floating point instructions, c) as:

$$C = \sum_{i=1}^n \sum_{j=1}^k c_j s_{i,j} \quad (7)$$

Theoretically, it is possible to optimize NWS by running it with different combination of internal models on a set of representative data and then removing the least efficient ones. However, the search space is prohibitive: There are a total of 2^{24} combinations. To reduce the search space, we used a heuristic that evaluates how much the total computation cost and error would change if a forecaster u is substituted with another forecaster v throughout the series.

Formally, this process can be expressed as:

$$s'_{i,j} = \left\{ \begin{array}{ll} 1 & \text{If model } j \text{ is winner forecaster} \\ & \text{for measurement } i \text{ and } j \neq u \\ 1 & \text{if model } j \text{ is not winner forecaster} \\ & \text{for measurement } i \text{ and } j = v \\ 0 & \text{Otherwise} \end{array} \right\} \quad (8)$$

where $E_{u,v}$ and $C_{u,v}$ are defined same as (6) and (7) using $s'_{i,j}$ instead of $s_{i,j}$

We computed $E_{u,v}$ and $C_{u,v}$ for every pair of u and v on a set of six representative traces and represented it as matrix with u as rows and v as columns. This representation provides a very compact form with which we can evaluate the efficiency of each model: Every column of the matrix shows how much the error rate would change if v had been used instead of u . For example, $E_{2,1}$ shows the new error if *last value* is used instead of *running mean*. If the $E_{2,1}$ is smaller than original NWS's error rate for *all* the trace files, then we consider *last value* to be a better predictor than *running mean*. Similarly, if in an extreme case, all the values of column 2 are smaller than original NWS's error rate, then *running mean* outperforms the original NWS. Even though, it is theoretically possible, we did not come across an example of such a case.

Our methodology is similar to off-line, profile-based optimization research in which a set of representative program inputs are used to collect profile information that is used to guide optimization [17, 24, 16]. Empirical experiments then use a different set of inputs to evaluate the efficacy of the technique. Here, we used six traces to identify NWS forecasters that enable high accuracy at low cost. We then evaluated NWSLite using over 300 different traces. Given the evaluation matrix, we used a set of empirical rules with which we eliminated forecasters. We removed any model

- that had more than 1% error rate across all traces,
- for which there is another model with significantly lower cost that can replace it, with a slight but acceptable (less than 5%) increase in error rate, and
- for which there is a combination of other models that enable a similar error rate.

We ruled out many models directly using the first criteria. For example, replacing *30% trimmed median window 31* with *running mean* generated an increase in error rate of at most 0.2%. On the other hand, for *median window 31* the *running mean* was only 0.2% higher in 5 of the 6 traces, except the last for which it had an error rate that was 32% higher. In the remaining trace, *median window 5* had almost the same error rate. As such, we included *median window 5* and omitted *30% trimmed median window 31*. Eventually, we identified five predictors (shown

Name	Trace Size	Description
Application	20 traces 17870 predictions	Interactive, 3-D rendering application CPU demand. Measurements are CPU time from user request to program response.
Network Bandwidth [20]	132 traces 750476 predictions	Observations of 64Kb-1Mbyte TCP data transfers. 3 configurations: UIUC LAN (inter-cluster), UIUC campus-wide network (intra-cluster), and cross-country Internet (UIUC-UCSD)
CPU load [20]	59 traces 6000697 predictions	Fraction of CPU occupancy time a standard user process can obtain Observations are in 10 seconds intervals.
Network Latency [20]	134 traces 750305 predictions	Round trip time of TCP. Transferring 4 bytes and measuring acknowledge time. Granularity levels same as network bandwidth.
Wireless Bandwidth [23]	1 trace 3028 predictions	4 access points on same subnet. Traces include 195 users, 300000 flows and 4.6 GB of network traffic. Bandwidth measured in 1 minute intervals

Table 2: Datasets Used for Evaluation

Input Scene	Applications	
	GLVU	Radiosity
castle	Yes	
cessna	Yes	Yes
chevy	Yes	
cloister	Yes	
cup		Yes
dragon		Yes
ground-table-land	Yes	Yes
ground-riverain-valley		Yes
shuttle	Yes	Yes
venus		Yes

Table 3: Applications and Inputs Used for Evaluation. We collected 10 trace files per application (3-D scene rendering programs). We used some programs multiple times with different input sets, i.e., different navigation paths. Empty entries mean that the application failed to process the particular scene.

in bold in Table 1) that trade off cost and prediction error most effectively.

5 Evaluation

To empirically evaluate the efficacy of NWSLite, we performed experiments using a wide range of datasets, applications, and metrics. We compared the performance of NWSLite against the NWS, LSQ, and RPF (the prediction systems described in Section 3). We implemented all of forecasters as efficiently as possible using the C language; we compiled each using gcc and -O2 optimization. Unlike NWSLite and the NWS, the LSQ and RPF methods are parametric models and hence, require parameterization. For each model, we created a pool of parameter settings, that included the published values [18, 9, 19] as well as our own values, resulting in 18 different forecasters. We selected the best performing parameterization for

each over all of the datasets we considered.

In the next section, we describe the experimental methodology (datasets and applications). We then detail the metrics we use in Section 5.2 results using these metrics in Section 5.3.

5.1 Experimental Methodology

We evaluated the efficacy of the prediction techniques using traces from a number of different types of data. In total, we performed experiments on 346 traces which produced more than 7 million predictions. We considered prediction of the execution time of application tasks, wired and wireless network bandwidth, wired network latency, and CPU load availability. All of the data traces, except for application execution times, were made freely available to us via web-sites of research groups around the country [20, 1, 13]. These groups collect, analyze, and post resource performance measurements from a wide range of resources for others to use. We provide the details on the different datasets in Table 2 and we refer to each of the different types of data sets (CPU, bandwidth, applications, etc.) as “groups”. In this section, we provide an overview of these datasets.

We generated traces of task execution times from augmented reality applications that are commonly used in similar mobile studies and highly suitable for remote execution: Independent tasks used in 3-D scene rendering and image fidelity adjustment that are computationally-intensive and easily divided into components for off-loading. The applications and inputs that we considered are shown in Table 3.

GLVU [12] allows navigating inside a 3-D scene by rendering the scene from any viewpoint of user. From an augmented reality view, Radiator [28] complements GLVU by computing the lighting effects for a given scene. Both applications can easily be divided into *operations* [18], which are a suitable unit for remote execution and fidelity adjustment. An operation (which we also

refer to as a task) is the smallest user-visible execution unit, such as viewpoint change in a rendering operation. For each application we rendered a set of 10 scenes which produced a total of 17870 operations. We consider the prediction performance for applications to be the accuracy with which the prediction system forecasts the CPU demand of each operation.

The bandwidth, CPU availability, and latency data were collected as a part of the NWS project [20]. NWS network sensors use active network probes to collect TCP/IP latency and bandwidth data on a group of geographically distributed hosts connected via local, wide area, and Internet networks. Each probe establishes a TCP connection, transmits a fixed amount of data, and tears down the connection. Network sensors measure network bandwidth using a 64 KByte data transfer and network latency using a 4 byte data transfer.

The NWS CPU sensors combine the information from Unix system utilities *vmstat* and *uptime* with periodic active CPU occupancy tests to provide measurements of CPU availability. The *uptime* utility reports the average number of processes in the run queue over the last one, five and fifteen minutes. The sensor uses the average load over the one minute period and computes the CPU availability by using the idle, user, and system time output from *vmstat* utility. The CPU availability is measured as the fraction of CPU occupancy time a standard user process can obtain.

The wireless bandwidth traces we used were collected during the SIGCOMM'01 conference [23]. The conference building was covered with four 802.11b access points. The traces span a 3 day period capturing 300000 flows generated by 195 users consuming a total of 4.6 GB of bandwidth.

5.2 Evaluation Metrics

We present our empirical evaluation of the different prediction systems in terms of both accuracy and computational cost. We use three metrics, described in this section, to evaluate predictor accuracy. We use instruction count (both total and floating point) as the metric for predictor cost.

The first of the three metrics we use to evaluate predictor accuracy is **error deviation**. We define error deviation as:

$$MSE = \frac{\sum_{i=1}^n (x_i - y_i)^2}{n}$$

$$\text{Error deviation} = \sqrt{MSE} \quad (9)$$

where x is the set of n predictions and y is the set of n corresponding observations. The mean square error (MSE) is the average square prediction error over the n pairs, (x, y) . The error deviation is the square root of the mean square

error. Error deviation describes the error in absolute terms and represents (in analogy) the *standard deviation* of the errors with respect to the *expectation* constituted by the forecast. Error deviation accounts for outliers and is more sensitive to incorrect predictions than is *absolute error* in which the absolute value of the error is used.

However, the error deviation is most meaningful when comparing the performance of predictors on the same time series. To provide a comparison across different series, we use a second metric that is the ratio of error deviation over the average observed value, i.e., the **relative error rate**:

$$\text{Relative error rate} = \frac{\sqrt{MSE}}{\text{observed_mean}} \quad (10)$$

This metric provides insight into how severe the error is in terms of the magnitude of the average measured value. For example, an error of 2Mb/s is large in a 10Mb/s link, but may not be significant in a 100Mb/s link.

The third metric we use for reporting prediction error is similar to relative error rate, however, instead of using the mean as the expected value, we use the absolute value of the forecast. This metric, called **predictability**, indicates how predictable the series relative to the forecasts it generates. It differs from the *relative error* in that it treats each forecast as a *conditional expectation* that it uses to normalize the error, instead of using the overall measurement mean. We compute predictability as:

$$\frac{\sum_{i=1}^n \frac{|x_i - y_i|}{|x_i|}}{n} \quad (11)$$

5.3 Results

We next present the results from our empirical comparison between NWSLite and competing prediction systems: The Network Weather Service (NWS), Odyssey (LSQ), and the Remote Processing Framework (RPF). We first present results for prediction accuracy and then evaluate the computational cost of each system.

5.3.1 Predictor Accuracy

Table 4 compares the **error deviation** (Equation 9) of the predictors using three representative traces, for brevity. In the application (APP) and CPU availability (CPU) datasets, we sorted the traces with respect to the *error deviation / average* of NWSLite and selected the best, worst, and median, which we report in the table. For the wired network data (bandwidth (BW) and latency (LAT)), we instead report data for three different types of links: intra-cluster, inter-cluster, and inter-campus (across country). For wireless (WBW), we only have a single trace and thus show data only for it.

Description	Units	Avg	NWSLite	NWS	LSQ	RPF
APP1 - best	msecs	148845.000	5287.856	5358.179	8180.561	22013.694
APP2 - median		9179.390	1322.139	1329.372	2385.072	5702.085
APP3 - worst		169753.000	135125.056	138064.335	145384.166	186430.176
BW1 - within cluster	Mbits/sec	65.801	17.161	16.958	52.112	17.191
BW2 - cross-cluster		76.522	13.308	13.329	59.279	13.507
BW3 - cross-country		4.536	0.878	0.859	78.063	1.164
CPU1 - best	CPU fraction	1.992	0.016	0.016	13.905	0.029
CPU2 - median		0.543	0.017	0.017	14.451	0.049
CPU3 - worst		1.391	2.672	2.684	3.113	2.661
LAT1 - within cluster	msecs	13.936	16.873	16.890	41.121	17.048
LAT2 - cross-cluster		2.345	8.309	8.319	46.829	8.337
LAT3 - cross-country		77.217	14.295	12.753	81.820	13.149
WBW	Kbits/sec	206.674	193.782	194.498	255.254	261.744

Table 4: Error Deviation for a Set of Representative Traces. The third column is the average of the measured values, the next four columns show the error deviation for each of the prediction systems. The APP and CPU datasets are sorted with respect to *error deviation / average* and best, median and worst cases are shown. For the BW and LAT datasets, the average error deviation within cluster, across cluster and across country are reported.

The first three columns of the table shows the description, trace name, and value units for each trace. The third column, Avg, shows the average observed value. The final four columns show the error deviation for each of the four predictors: NWSLite, NWS, LSQ, and RPF. LSQ and RPF are parameterized as described in Section 5, and identify the best-performing, converging parameterizations of each technique.

The NWS and NWSLite have almost identical error deviations in every case. LSQ performs well for applications (as was shown in prior work [18]), but it is the worst-performing predictor for all other types of data. NWSLite performs better than the NWS for *lat3*, however, the difference is less than 2%. NWSLite performs better than LSQ and RPF in almost every case, and is significantly better than both LSQ and RPF in most cases. For example, in the application group, for both *shuttle* and *cloister* NWSLite performs 3 times better than RPF. The wireless dataset is especially challenging. All the forecasters show a high error rate.

Figure 2 shows the **relative error rate** of the predictors across all of the traces in each group. The information in the graph confirms the results of Table 4. NWSLite performance is very similar to that of the NWS; in all groups it enables the best prediction error. LSQ is ineffective for the bandwidth, CPU, and latency groups. RPF performs quite well for the CPU and bandwidth groups; and exceeds NWSLite performance for network latency by 1.5%. RPF is the worst predictor however, for the application and wireless groups. For the application group, the average error rate of RPF is 86% higher than that of NWSLite. As a special case, we also compared the performance of predictors with Odyssey’s specialized smooth-




Figure 2: Relative Error Rate (Equation 10). This metric shows how severe the error is with respect to the average measured value. The LAT has the highest relative error rate among all forecasters, however, as most latency observations are very small (around 1 msecs), the absolute error is small.

ing filters for bandwidth and latency, which we refer to as ODY-BW and ODY-LAT (not shown in figure). ODY-BW performed as good as the rest of the group in bandwidth prediction, however, ODY-LAT is the second worst performer in latency dataset. The appendix section of this paper contains more detail on the results.

Figure 3 shows the **predictability** (Equation 11) of the series given each predictor. This metric assumes that predictor is a valid conditional expectation that can be used to normalize the error at each point of the trace. The lower the value the more accurate the forecaster. Since the variance of the results is high, we normalized the results to NWSLite for each group.

The predictability results support our findings in Figure 2. NWSLite is as accurate as NWS in all cases, and it performed significantly better than the parameterized forecasters in most cases. The single exception is the latency dataset, in which RPF is the winner. However, the difference between RPF and NWSLite is very small. In contrast, the accuracy of RPF is significantly worse in all other cases, emphasizing the difficulty of finding a good parameterization for the general case. These results also show that, except the application dataset, LSQ is always worse than smoothing-filter-based predictors. In application dataset, LSQ is approximately 40% more accurate than RPF, however, it is still significantly worse than NWSLite. The predictability of NWSLite is considerably higher even than the highly tuned predictors ODY-LAT and ODY-BW (not shown in figure). For the latency dataset, ODY-LAT is 13% less predictable than NWSLite; whereas in bandwidth dataset, NWSLite does 21% better than ODY-BW.

An interesting case is the behavior of RPF in Figures 2 and 3; even though the relative error rate of RPF is small, its predictability is not. This is due to the characteristics of CPU dataset - the CPU availability values are in the range $(0, 1)$, or $(0, n)$ if there are n processors. As such, most of the time the values are a fraction of 1. This results in a small value for the sum of square errors even though the errors are high relative to the expected value.

5.3.2 Computational Cost of Prediction

In addition to studying prediction error, we also considered the cost of performing prediction on a resource-restricted device. To our knowledge, no prior studies that use prediction on mobile devices consider the resource consumption of the predictors themselves.

We first compare the predictors in terms of instructions required for one prediction. We extracted this information by using the SimpleScalar [6] simulator. Figure 4 shows the average cost of each predictor. NWSLite uses 55 floating point instructions per forecast. Even though this is more than the cost of RPF and LSQ, which use 8




Figure 3: Predictor Predictability (Equation 11). Due to high variation among forecasters, the values are normalized to NWSLite for each group. The lower the value, the more accurate the forecaster.

and 42, respectively, the accuracy of NWSLite exceeds both of these predictors significantly.

As most resource-restricted devices lack a floating point co-processor, floating point instructions are very expensive. We break down the instruction counts into floating-point and non-floating-point instructions in the first two columns of Table 5.

We also executed the predictors on a real resource-restricted device: An iPAQ H3800 hand-held computer from Compaq [7]. The iPAQ has a 206 MHz Intel StrongArm CPU and runs Familiar Linux version 0.5.3. The execution times (in microseconds) are shown in the final column of the table. These times include the cost of IO to read the trace file from flash memory and to print the results.

The execution time of NWSLite is approximately 4% that of NWS but enables prediction accuracy that is nearly equivalent. Given that it requires only 381 microseconds to execute a prediction, including the IO, NWSLite is

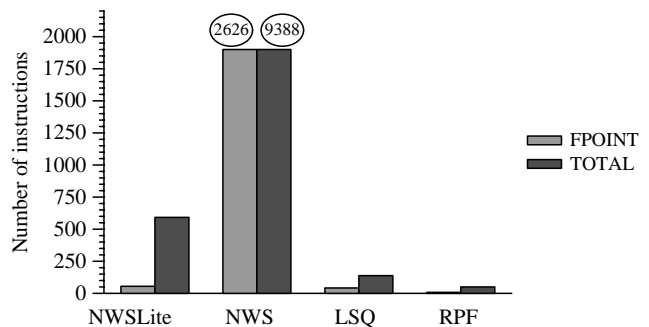


Figure 4: Forecaster Cost as Number of Instructions Executed (floating-point (FPOINT) and TOTAL) per Prediction

Prediction System	Floating Point	Total Instructions	Execution time (microsecs)
NWSLite	55	592	381.34
NWS	2626	9388	10231.31
LSQ	42	138	295.27
RPF	8	50	154.9

Table 5: Execution Cost Comparison per Prediction

a more attractive solution for on-line forecasting using resource-restricted devices, than the parametric and less accurate models of Odyssey and the RPF.

5.4 Summary of the Results

We summarize the result of our findings in Table 6. To make our results comparable to previous studies [18], we report summary performance in terms of percentile error. We define the X percentile error, E_X , as the maximum *absolute* error for X% of the experiments. For example, for the bandwidth dataset, E_{95} of NWSLite is 25.6 meaning that 95% of the time the prediction error of NWSLite is within 25.6 kilobits/second. The reason we use absolute rather than relative error is to avoid skewed data in CPU and latency datasets. We report the results for NWS, NWSLite, LSQ, RPF as well as for the two other smoothing filters that we studied, ODY-LAT (the Odyssey network latency predictor) and ODY-BW (the Odyssey network bandwidth predictor).

The results show that NWS and NWSLite are general enough that they perform well in all datasets. Even though parameterized forecasters can match NWSLite in some datasets, they fail in others. As an example, the performance of ODY-BW is close to NWSLite in APP dataset, but it is significantly higher in BW, CPU and LAT datasets. The same pattern also exists for ODY-LAT and RPF. RPF matches NWSLite in BW and LAT, but it is significantly worse in APP and CPU datasets.

Another pattern in the results is that both NWS and NWSLite perform better than all others when a higher percentage of predictions considered. This suggests that, NWS and NWSLite can better adjust themselves to sudden changes in performance patterns by switching to another model; the other models must simply rely on their static parameters.

The wireless bandwidth dataset is significantly different than other datasets. The error rates are very high, i.e., E_{90} is around 200Kbits/sec on a 11Mbits/sec link, hence none of the forecasters performed at a satisfactory level. This emphasizes the need for additional study of and novel forecasters for wireless network bandwidth data.

The success of NWSLite results from its capability

to dynamically switch between a carefully chosen set of competing models based on previously observed accuracy. If the dynamics of the observed dataset changes over time, NWSLite can adapt to the new conditions; the prediction systems of Odyssey and RPF cannot and as such are data (input) dependent. For example, exponential smoothing with a gain of 0.05 can be the most accurate predictor at some point, however, a transient or permanent change can occur so that the running mean can become the most accurate. In this case, NWSLite will respond by switching to running mean if the change is persistent enough to cause the aggregate error ranking to change. Odyssey and RPF are statically configured by a set of predetermined parameters. Thus, even though there are individual cases that other predictors can match the accuracy of NWSLite, they are unable to do well across dynamically changing series and to different types resource performance data.

The flip-flop filter extension to Odyssey [14], described in Section 2, incorporates some adaptivity by using two different parameter settings in its exponential smoothing predictor. However, exponential smoothing cannot always produce the best prediction accuracy (given any gain parameters). NWSLite incorporates exponential smoothing using two different gain factors but is more general and adaptive than this filter since it considers a wide range of other prediction techniques that can enable significant improvements in accuracy at low computational cost.

6 Conclusions

By off-loading tasks from the resource-restricted devices to wall-powered, high-performance servers, remote-execution can significantly extend the capability and battery life of mobile and pervasive devices. To determine when to offload, these devices must make forecasts about the efficacy of doing so. A device must estimate the cost of both remote and local execution given the highly-variable underlying resource performance as well as the characteristics of the task to be executed. As such, it must employ sophisticated prediction techniques that are both accurate and light-weight, i.e., they do not consume significant device resources.

We present a light-weight, computationally efficient, prediction utility for mobile devices called NWSLite. The system is an extension of the Network Weather Service (NWS), a dynamic measurement and forecasting toolkit designed and developed for adaptive application scheduling in Computational Grid environments (performance-oriented distributed systems). We identify 5 of the 24 NWS forecasters for NWSLite implementation, that trade-off computational cost for predictor accuracy most effectively.

	APP		BW		CPU		LAT		WBW	
	E90	E95	E90	E95	E90	E95	E90	E95	E90	E95
NWSLite	3319.000	7336.000	10.271	25.699	0.019	0.043	15.772	24.566	198.130	351.090
NWS	3343.000	7459.000	9.601	25.580	0.018	0.038	15.801	24.502	202.771	358.798
LSQ	5866.552	13305.338	14.105	28.459	0.058	0.115	16.415	26.867	230.591	422.977
RPF	17147.400	38696.700	10.596	25.561	0.080	0.209	16.187	24.915	326.340	533.047
ODY-LAT	3759.839	8806.945	9.923	39.717	0.025	0.094	16.318	29.848	197.429	335.172
ODY-BW	3458.320	7894.141	7.384	42.541	0.021	0.079	16.883	31.494	192.992	354.560

Table 6: Results in Summary: Percentile Error. We define the X percentile error, E_X , as the maximum absolute error for X% of the experiments. The table compares the E_{90} and E_{95} of all forecasters for all 5 datasets and prediction systems studied.

We evaluate NWSLite using over 300 different traces of application execution times, CPU availability, wired network bandwidth and latency, and wireless bandwidth. In addition, we compare NWSLite to the NWS and to two other extant remote execution prediction systems. We find that NWSLite consistently outperforms the latter and achieves prediction accuracy similar to that of the NWS. However, NWSLite achieves this level of accuracy at a significantly lower execution cost than the NWS.

References

- [1] A. Balachandran, G. Voelker, P. Bahl, and P. Rangan. Characterizing user behavior and network performance in a public wireless lan. In *ACM International Conference on Measurement and Modeling of Computer Systems*, 2002.
- [2] R. Balan, M. Satyanarayanan, S. Park, and T. Okoshi. Tactics-based remote execution for mobile computing. In *International Conference on Mobile Systems, Applications, and Services*, 2003.
- [3] F. Berman, G. Fox, and T. Hey. *Grid Computing: Making the Global Infrastructure a Reality*. Wiley and Sons, 2003.
- [4] F. Berman, R. Wolski, S. Figueira, J. Schopf, and G Shao. Application level scheduling on distributed heterogeneous networks. In *Supercomputing 1996*, 1996.
- [5] G. Bottomley and S. Alexander. A novel approach for stabilizing recursive least squares filters. *IEEE Transactions on Signal Processing*, 39, 1991.
- [6] D. Burger and T. Austin. The simplescalar tool set, version 2.0. Technical Report 1342, UW Madison Computer Sciences, June 1997.
- [7] Compaq Computer Corporation. *iPAQ Pocket PC*. <http://www.compaq.com/products/handhelds/pocketpc/>.
- [8] J. Flinn, D. Narayanan, and M. Satyanarayanan. Self-tuned remote execution for pervasive computing. In *Hot Topics in Operating Systems(HotOS-VIII)*, pages 61–66, Germany, 2001.
- [9] J. Flinn, S. Park, and M. Satyanarayanan. Balancing performance, energy, and quality in pervasive computing. In *International Conference on Distributed Computing Systems (ICDCS '02)*, pages 217–226, 2002.
- [10] J. Flinn and M. Satyanarayanan. Energy-aware adaptation for mobile applications. In *Symposium on Operating Systems Principles*, pages 48–63, 1999.
- [11] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, Inc., 1998.
- [12] GLVU source code and documentation, Feb 2002. <http://www.cs.unc.edu/~walk/software/glvu/>.
- [13] The grid application development software project (GrADS). <http://hipersoft.cs.rice.edu/grads/>.
- [14] Minkyong Kim and Brian Noble. Mobile network estimation. In *Mobile Computing and Networking*, pages 298–309, 2001.
- [15] U. Kremer, J.Hicks, and J.M.Rehg. A compilation framework for power and energy management on mobile computers. In *International Workshop on Languages and Compilers for Parallel Computing (LCPC'01)*, August 2001.
- [16] C. Krintz. Coupling On-Line and Off-Line Profile Information to Improve Program Performance. In *International Symposium on Code Generation and Optimization (CGO)*, March 2003.
- [17] C. Krintz and B. Calder. Using Annotation to Reduce Dynamic Optimization Time. In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 156–167, June 2001.
- [18] Dushyanth Narayanan and M. Satyanarayanan. Predictive resource management for wearable computing. In *International Conference on Mobile Systems, Applications, and Services*, 2003.
- [19] B. Noble, M. Satyanarayanan, D. Narayanan, J. Tilton, J. Flinn, and K. Walker. Agile application-aware adaptation for mobility. In *sixteenth ACM symposium on Operating systems principles*, pages 276–287. ACM Press, 1997.
- [20] The Network Weather Service Home page – <http://nws.cs.ucsb.edu>.
- [21] A. Rudenko, P. Reiher, G.Popek, and G.Kuenning. The remote processing framework for portable computer power saving. In *ACM Symp. Appl. Comp.*, San Antonio,TX, February 1999.

- [22] A. Rudenko, P. Reiher, G. Popek, and G. Kuenning. Saving portable computer battery power through remote process execution. *Mobile Computing and Communications Review*, 2(1):19–26, January 1998.
- [23] Wireless LAN Traces from ACM SIGCOMM'01. <http://ramp.ucsd.edu/pawn/sigcomm-trace/>.
- [24] M. Smith. Overcoming the challenges to feedback-directed optimization. In *ACM SIGPLAN Workshop on Dynamic and Adaptive Compilation and Optimization (Dynamo00)*, January 2000.
- [25] N. Spring and R. Wolski. Application level scheduling: Gene sequence library comparison. In *Proceedings of ACM International Conference on Supercomputing*, July 1998.
- [26] S. Sucu and C. Krintz. ACE: A Resource-Aware Adaptive Compression Environment. In *International Conference on Information Technology: Coding and Computing (ITCC)*, April 2003.
- [27] Martin Swamy and Rich Wolski. Representing Dynamic Performance Information in Grid Environments with the Network Weather Service. In *2nd IEEE International Symposium on Cluster Computing and the Grid*, May 2002.
- [28] A. J. Willmott. Radiator source code and online documentation, Oct 1999. <http://www.cs.cmu.edu/ajw/software/>.
- [29] R. Wolski. Dynamically Forecasting Network Performance Using the Network Weather Service. *Journal of Cluster Computing*, 1:119–132, January 1998.
- [30] R. Wolski, N. Spring, and J. Hayes. The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing. *Future Generation Computer Systems*, 1999.
- [31] P. Young. *Recursive Estimation and Time-Series Analysis*. Springer-Verlag, 1984.
- [32] Z.Li, C.Wang, and R.Xu. Computation offloading to save energy on handheld devices:a partition scheme. In *Proc. of International Conference on Compilers,Architectures and Synthesis for Embedded Systems (CASES)*, pages 238–246, 2001.

Appendix

In this section, we provide more detail on the performance of forecasters. First, we evaluate the effectiveness of all forecasters using *relative error rate* and *predictability*. These metrics are described in section 5.2. Then we discuss the effect of parameterization on the linear regression and smoothing filter based forecasters.

First we compare six forecasters - NWS, NWSLite, LSQ, RPF, ODY-LAT and ODY-BW- on network bandwidth, application CPU demand and wireless network bandwidth datasets. We further divide the network bandwidth traces into three groups: In-cluster, across-cluster and cross-country.

A typical computation cluster is a few high performance UNIX and Linux workstations connected through a high speed local area network. The local area network is a 100Mb/s or 1000Mb/s (Gigabit) Ethernet. The clusters are connected through other LAN's or university backbone networks.

In-cluster and across-cluster datasets consist of measurements inside a cluster and across clusters. Cross-country dataset includes probes among clusters that are scattered throughout the country, including UCSB, UCSD, UTK and University of Illinois at Urbana-Champaign.

We show the results in Figures 5 and 6. NWS and NWSLite are invariably among the best forecasters. The LSQ is most appropriate for application dataset, but it is almost always the worst forecaster in other cases. On the other hand, the exponential smoothing based filters do better than LSQ, however, they still can not match the performance of NWSLite. All smoothing filters, in particular ODY-LAT and ODY-BW, fail when predictability metric is considered. As a side note here, the relative error rate considers the *outliers*, as it adds the square of prediction error. On the other hand, the predictability metric considers the absolute error. Thus, our conclusion is that smoothing filters leave less outliers but they are still worse than NWS and NWSLite in terms of precision.

When individual datasets considered, the network bandwidth data presents interesting results. The predictability evaluation shows that cross-country network bandwidth is more predictable than high-speed local area networks. That could be due to the characteristics of the applications that dominate the clusters.

As the Figures 5 and 6 show, the parametric models sometimes do as well as non-parametric models for a *particular* subset of data. However, if they are poorly configured, the results could be disastrous. To prove this point, in table 7, we compare the relative error rate of two different LSQ forecasters: one with a decay factor 0.5 and the other with 0.8. The former one is the original parameter Narayanan used to estimate application resource consumption. A factor as low as 0.5 provides a very agile forecaster that adapts quickly to any change in the system. However, noise can easily degrade its performance to unacceptable levels. As the table shows, the agile predictor has a very high error rate in noisy environments such as network bandwidth and latency. Furthermore, it diverges (error rate goes to infinity) on a subset of CPU availability traces.

We explain the divergence this way: Some of the workstations that we observed had long idle periods with sporadic high activity periods. As the CPU availability is a fraction of 1, in data traces this corresponds to a long sequence of values that are very close to 0, and then sudden jumps. This, when combined with a long observation period, causes the numerical errors to aggregate. Eventually, due to the recursive nature of the algorithm, it becomes unstable and diverges. More information about divergence in least squares linear regression is given in [5].

On the other hand, a higher exponential decay parameter creates a predictor that gives more weight to past. Such a model is less agile but more immune to any transient noise. After our empirical evaluations with many parameters, we found that $\alpha = 0.8$ provides a good trade-off between agility and precision. We found that this predictor has an error rate two or three order of magnitude better than the agile one.

Table 8 shows the parameterizations we evaluated for RPF forecaster. The forecasters to the left of the table use a smaller history window than the ones to the right, making them more agile. We found that the forecasters with a larger history window are slightly better. Figure 7 compares relative error rate of all RPF parameterizations on the overall dataset. On the other hand, the more agile predictors do better in application dataset but do worse as figure 8 shows. Overall, R10 provides the best parameterization for RPF.




Figure 5: **Comparison of forecasters:** The figures (a) through (e) compare the performance of forecasters on individual datasets. The evaluation metrics is relative error. NWS and NWSLite are winners in all cases. ODY-LAT and ODY-BW forecasters are just behind them. LSQ is the worst performer in almost all cases except the application dataset. The figures at the left side (a,c,e) belong to fixed network bandwidth dataset divided into smaller groups. (a) describes relative error rate in a computation cluster, which is generally connected via a high speed local area network. (c) and (e) are similar to (a) except that (c) is across computation clusters (such as a university campus network) and (e) is cross country. (such as a wide area network). (b) and (d) belong to wireless network bandwidth and application datasets.




Figure 6: **Comparison of forecasters:** The figures (a) through (e) compare the performance of forecasters on individual datasets. The evaluation metrics is predictability. NWS and NWSLite are winners in all cases. Surprisingly, the ODY-BW, which is specially developed to estimate network bandwidth is the worst performer in predicting in-cluster and across-cluster network bandwidth. The descriptions of datasets for figures (a to e) is given in figure 5 caption.

Dataset	APPL	WBW	LAT-1	LAT-2	LAT-3	CPU	RTT-1	RTT-2	RTT-3
$\alpha = 0.5$	0.20	4021.30	4850.43	5005.36	15167.1	Div	84292.88	98750.82	6968.61
$\alpha = 0.8$	0.23	1.23	0.78	0.78	23.19	14.12	12.77	14.12	2.55

Table 7: Effect of exponential decay on LSQ: The exponential decay factor determines the weight of history on predicting the future values. A value as low as $\alpha = 0.5$ makes the predictor very agile. That is it can respond to the changes quickly. However, any transient change or noise can adversely effect the outcome. This table compares the relative error rate of least squares method with two different decay factors $\alpha = 0.5$ and $\alpha = 0.8$. Except the application dataset, which is virtually noise free, a low agility factor always results in a **much** higher error rate. The datasets used in the evaluation are: Application, wireless bandwidth, network bandwidth in 3 different resolutions (in cluster, across cluster, cross-country), CPU availability, network latency in 3 different resolutions (in cluster, across cluster, cross-country). *Div stands for divergence for at least some traces*

Name :	R0	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14
α	0.0	0.1	0.2	0.5	0.9	0.0	0.1	0.2	0.5	0.9	0.0	0.1	0.2	0.5	0.9
k	5					10					15				

Table 8: **Different parameterizations of RPF:** To find the best parameterization, we tried 15 different combinations of k and α (see formula 3). The effect of history on new forecast decreases with a higher α . The parameter k determines how large the history window is; the higher it is, the larger the window.






Figure 7: **Overall comparison of RPF forecasters:** This figure compares the effectiveness of RPF forecasters with different parameters in terms of relative error rate (see formula 10). As the figure shows, the difference between forecasters is very small suggesting the formula 10 is not sensitive to parameters. However, this is not true. In some datasets the difference is very visible. (see figure 8). Overall, the best parameterization is R10.



(a)



(b)

Figure 8: **Parameterizations of RPF:** The effectiveness of RPF forecasters depends on the selected parameters. The figures above compare the effect of parameters in terms of relative error rate. The application dataset is highly sensitive to the chosen parameter (shown in Figure (a)), but the CPU dataset (Figure (b)) is not. In application dataset, the R0 parameterization is slightly **better** than the rest. This suggests that for our applications (GLVU and Radiosity), the resource demand is dependent on most recent history. On the other hand, in the CPU dataset, the parameterizations with a shorter history are slightly **worse** than the rest.