Tuning Database Configuration Parameters with iTuned

Songyun Duan, Vamsidhar Thummala, Shivnath Babu
Duke University

ABSTRACT

Database systems have a large number of configuration parameters that control memory distribution, I/O optimization, costing of query plans, parallelism, many aspects of logging, recovery, and other behavior. Regular users and even expert database administrators struggle to tune these parameters for good performance. The wave of research on improving database manageability has largely overlooked this problem which turns out to be hard to solve. We describe iTuned, a tool that automates the task of identifying good settings for database configuration parameters. iTuned has three novel features: (i) a technique called Adaptive Sampling that proactively brings in appropriate data through planned experiments to find high-impact parameters and high-performance parameter settings, (ii) an executor that supports online experiments in production database environments through a cycle-stealing paradigm that places near-zero overhead on the production workload; and (iii) portability across different database systems. We show the effectiveness of iTuned through an extensive evaluation based on different types of workloads, database systems, and usage scenarios.

1. INTRODUCTION

Consider the following scenario from a small to medium business (SMB) enterprise. Peter, a Web-server administrator by training, maintains the Web-site of a ticket brokering company that employs eight people. Over the past few days, the Web-site has been sluggish. Peter collects monitoring data, and tracks the problem down to poor performance of queries issued by the Web server to a backend database.

Realizing that the database needs tuning, Peter runs the database tuning advisor. (SMBs often lack the financial resources to hire full-time database administrators, or DBAs.) Peter uses system logs to identify the workload W of queries and updates to the database. With W as input, the advisor recommends a database design (e.g., which indexes to build, which materialized views to maintain, how to partition the data). However, this recommendation does not solve the current problem: Peter has already designed the database this way based on a previous invocation of the advisor.

Peter recalls that the database has *configuration parameters*. For lack of better understanding, he had set them to default values during installation. Maybe the parameters need tuning, so Peter pulls out the 1000+ page database tuning manual. He finds many dozens of configuration parameters like buffer pool sizes, number of concurrent I/O daemons, parameters to tune the query optimizer's cost model, and others. Being unfamiliar with most of these parameters, Peter has no choice but to follow the tuning guidelines given. One of the guidelines look promising: if the I/O rate is high, then increase the database buffer pool size. However, on following this advice, the database performance drops even further. (We momentarily show an example of such behavior.) Peter is puzzled, frustrated, and undoubtedly displeased with the database vendor.

Most of us would have faced similar situations before. Tuning database configuration parameters is hard but critical: bad settings can be orders of magnitude worse in performance than good ones. Changes to some parameters cause local and incremental effects on resource usage, while others cause drastic effects like changing query plans or shifting bottlenecks from one resource to another. These effects vary depending on hardware platforms, workload, and data properties. Groups of parameters can have nonindependent effects, e.g., the performance impact of changing one parameter may vary based on different settings of another parameter.

iTuned: Our central contribution is a tool, called *iTuned*, that automates parameter tuning. iTuned can provide a very different experience to Peter. He starts iTuned in the background with the database workload W as input, and resumes his other work. He checks back after half an hour, but iTuned has nothing to report yet. When Peter checks back thirty minutes later, iTuned shows him an intuitive visualization of the performance impact each database configuration parameter has on W. iTuned also reports a setting of parameters that is 18% better than the current one. Another hour later, iTuned has a 35% better configuration, but Peter wants more improvement. Three hours into its invocation, iTuned reports a 52% better configuration. Now, Peter asks for the configuration to be applied to the database. Within minutes, the actual database performance improves by 52%; and Peter is very happy.

To understand the technical innovations in iTuned, let us now consider a simple, but real, example. Figure 1 is a *response surface* that shows how the performance of a complex TPC-H query [18] in a PostgreSQL database depends on the *shared_buffers* and *effective_cache_size* parameters. *shared_buffers* is the size of PostgreSQL's main buffer pool for caching disk blocks. The value of *effective_cache_size* is used to determine the chances of an I/O hitting in the OS file cache; so its recommended setting is the size of the OS file cache. Some observations from Figure 1:

- The surface is complex and nonmonotonic.
- Performance drops sharply as *shared_buffers* is increased beyond 20% (200MB) of available memory; causing a "increase
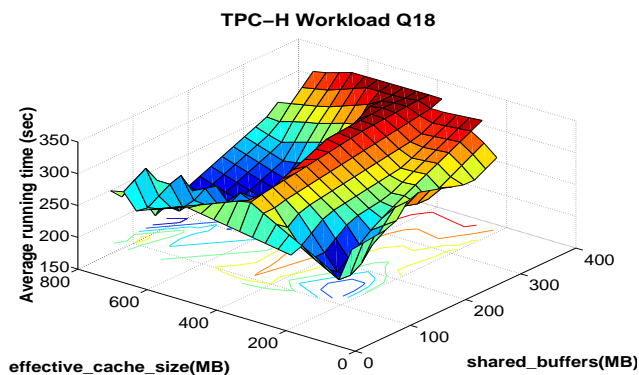
Figure 1: 2D projection of a response surface for TPC-H Query 18; Total database size = 4GB, Physical memory = 1GB

buffer pool size" rule of thumb to degrade performance.
- The effect of changing *effective_cache_size* is different for different settings of *shared_buffers*. Surprisingly, the best performance comes when both parameters are set low.

Typical database systems contain few tens of parameters whose settings can significantly impact workload performance.[1] What automated tools do users have today for holistic tuning of these parameters? Perhaps shockingly, the answer would be "very few or none".

The majority of tuning tools focus on the logical or physical design of the database. For example, index tuning tools are relatively mature (e.g., [4]). These tools use the query optimizer's cost model to answer *what-if questions* of the form: how will performance change if index I were to be created? Unfortunately, such tools do not apply to parameter tuning because the settings of many high-impact parameters are not accounted for by these models.

Many tools (e.g., [16, 19]) are limited to specific classes of parameters like buffer pool sizes. IBM DB2's Configuration Advisor recommends default parameter settings based on answers provided by users to some high-level questions (e.g., is the environment OLTP or OLAP?) [10]. These tools are based on predefined models of how parameter settings affect performance. Developing such models is nontrivial [20] or downright impossible because response surfaces can differ markedly across database systems (e.g., DB2 Vs. PostgreSQL), platforms (e.g., Linux Vs. Solaris; databases that are run on virtual machines), workloads, and data properties.[2] Furthermore, DB2's Configuration Advisor is helpless if the recommended defaults are still unsatisfactory.

Users are forced to rely on trial-and-error or rules-of-thumb from manuals and experts. The following tuning rule from an authoritative PostgreSQL source [12] highlights their predicament (*work_mem* is memory used by sort and hash operators):

> Adjust *work_mem* upwards for: large databases, complex queries, lots of available RAM. Adjust it downwards for: low available RAM or many concurrent users. Finding the right balance spot can be hard.

How do expert DBAs overcome these hurdles? They often run *experiments* to perform what-if analysis during parameter tuning. A typical experiment would consist of:
- Create a replica of the production database on a test system.
- Initialize database parameters on the test system to a chosen setting. Run the workload that needs tuning, and observe the resulting performance.

---

[1]The total number of parameters may be more than a hundred, but most have reasonable defaults.

[2]Section 7 provides empirical evidence.

iTuned takes a leaf from the book of expert DBAs. Each experiment gives a point on the response surface. Since reliable techniques for parameter tuning have to be aware of the underlying response surface, a series of carefully-planned experiments is a natural approach to parameter tuning. iTuned is not the first to advocate an experiment-driven approach for parameter tuning. Reference [17] applied such an approach to tune four parameters in BerkeleyDB. The tuned settings were impressive, however, 37 days were spent in running experiments in parallel on five machines.

Users don't always expect instantaneous results from parameter tuning; they would rather get recommendations that work as described. (Reference [10] estimates that configuring large database systems takes on the order of 1-2 weeks.) Nevertheless, to be practical, an automated parameter tuning tool has to produce good results within few hours. In addition, several questions need to be answered like: which experiments to run? where to run experiments? what-if the SMB does not have a test database platform?

1.1 Our Contributions

To our knowledge, iTuned is the first practical tool that uses planned experiments to tune database configuration parameters. We make the following contributions.

Planner: iTuned's experiment planner uses a novel and methodical technique, called *Adaptive Sampling*, to select which experiments to conduct. Adaptive Sampling uses the information from experiments done so far to estimate the utility of new candidate experiments. No assumptions are made about the shape of the underlying response surface, so it can deal with simple to complex surfaces.

Executor: iTuned's experiment executor uses a novel approach to conduct online experiments in a production environment while ensuring near-zero overhead on the production workload. The executor is controlled through high-level policies. It hunts proactively for idle capacity on the production database, hot-standby databases, as well as databases for testing and staging of software updates. The executor's design is particularly attractive for databases that run in cloud computing environments providing pay-as-you-go resources.

Representation of uncertain response surfaces: iTuned introduces *GRS*, for *Gaussian process Representation of a response Surface (GRS)*, to represent an approximate response surface derived from a set of experiments. GRS enables: (i) visualization of response surfaces with confidence intervals on estimated performance; (ii) visualization and ranking of parameter effects and inter-parameter interactions; and (iii) recommendation of good parameter settings.

Scalability: iTuned incorporates a number of features to reduce tuning time and to scale to many parameters: (i) a sensitivity-analysis algorithm that quickly eliminates parameters with insignificant effect; (ii) planning and conducting parallel experiments; (iii) aborting low-utility experiments early, and (iv) workload compression.

Evaluation: We establish the advantages of iTuned comprehensively through an empirical evaluation along a number of dimensions: multiple workload types, data sizes, database systems (PostgreSQL and MySQL), and number of parameters. We compare iTuned with recent techniques proposed for parameter tuning both in the database [5] as well as other literature [17, 22]. We consider how good the results are and the time it takes to produce them.

2. ABSTRACTION OF THE PROBLEM

Response Surfaces: Consider a database system with workload W and d parameters $x_1, \ldots, x_d$ that a user wants to tune. (The notation used throughout this paper is summarized in Table 1.) The values of parameter x_i, $1 \le i \le d$, come from a known domain $dom(x_i)$. Let *DOM*, where $DOM \subseteq \Pi_{i=1}^{d} dom(x_i)$, represent the space of possible settings of $x_1, \ldots, x_d$ that the database can have.

Notation	Description
$x_1, \ldots, x_d$	Parameters for tuning
$dom(x_i)$	Domain of feasible settings for x_i
X	A setting of $x_1, \ldots, x_d$ from the respective domains
y	Performance metric of interest for tuning
W	Workload of interest for tuning
X^t	Transpose of X
$\hat{y}(X)$	Mean of the estimation of y at setting X
$v^2(X)$	Variance of the estimation of y at setting X
$Y(X)$	Probability density function of the estimation of y
$\langle X^{(i)}, y^{(i)} \rangle$	Samples collected so far through experiments
$\vec{f}(X)$	Vector of basis functions
$\vec{\beta}$	Vector of regression coefficients
GRS	Gaussian process representation of response surface
$corr(X, X')$	Correlation function used in GRS
$Z(X)$	Zero-mean Gaussian process used in GRS
$EIP(X)$	Expected improvement when next expt. is done at X

Table 1: Notation used in the paper

Let y denote the performance metric of interest. Then, there exists a response surface, denoted S_W, that determines the value of y for workload W for each setting of $x_1, \ldots, x_d$ in *DOM*. That is, $y = S_W(x_1, \ldots, x_d)$. S_W is unknown to iTuned to begin with. The core task of iTuned is to find settings of $x_1, \ldots, x_d$ in *DOM* that give close-to-optimal values of y. In iTuned:

- Because iTuned runs experiments, it is very flexible in how the database workload W can be specified. iTuned supports the whole spectrum from the conventional format where W is a set of queries with individual query frequencies [4], to mixes of concurrent queries at some multi-programming level, as well as real-time workload generation by an application.
- y is any performance metric of interest, e.g., y in Figure 1 is the time to completion of the workload. In OLTP settings, y could be, e.g., average transaction response time or throughput.
- Parameter x_i can be one of three types: (i) database or system configuration parameters (e.g., buffer pool size); (ii) knobs for physical resource allocation (e.g., % of CPU); or (iii) knobs for workload admission control.

Experiments and Samples: Parameter tuning is performed through experiments planned by iTuned's planner, which are conducted by iTuned's executor. An experiment involves the following actions that leverage mechanisms provided by the executor (Section 5):
1. Setting each x_i in the database to a chosen setting $v_i \in dom(x_i)$.
2. Running the database workload W.
3. Measuring the performance metric $y = p$ for the run.

The above experiment is represented by the setting $\langle X \rangle = \langle x_1 = v_1, \ldots, x_d = v_d \rangle$. The outcome of this experiment is a *sample* from the response surface $y = S_W(x_1, \ldots, x_d)$. The sample in the above experiment is $\langle X, y \rangle = \langle x_1 = v_1, \ldots, x_d = v_d, y = p \rangle$.

As iTuned collects such samples through experiments, it learns more about the underlying response surface. However, experiments cost time and resources. Thus, iTuned aims to minimize the number of experiments required to find good parameter settings.

3. OVERVIEW OF ITUNED

Gridding: *Gridding* is a straightforward technique to decide which experiments to conduct. Gridding works as follows. The domain $dom(x_i)$ of each parameter x_i is discretized into k values $l_{i1}, \ldots, l_{ik}$. (A different value of k could be used per x_i.) Thus, the space of possible experiments, $DOM \subseteq \Pi_{i=1}^{d} dom(x_i)$, is discretized into a grid of size k^d. Gridding conducts experiments at each of these k^d settings. Gridding is reasonable for a small number of parameters. This technique was used in [17] while tuning four parameters in the Berkeley DB database. However, the exponential complexity

makes gridding infeasible (curse of dimensionality) as the number of parameters increase. For example, it takes 22 days to run experiments via gridding for $d = 5$ parameters, $k = 5$ distinct settings per parameter, and average run-time of 10 minutes per experiment.

SARD: The authors of [5] proposed *SARD* (Statistical Approach for Ranking Database Parameters) to address a subset of the parameter tuning problem, namely, ranking $x_1, \ldots, x_d$ in order of their effect on y. SARD decides which experiments to conduct using a technique known in Statistics as the *Plackett Burmann (PB) Design* [9]. This technique considers only two settings per parameter—giving a 2^d grid of possible experiments—and picks a predefined $2d$ number of experiments from this grid. Typically, the two settings considered for x_i are the lowest and highest values in $dom(x_i)$. Since SARD only considers a linear number of corner points of the response surface, it can be inaccurate for surfaces where parameters have nonmonotonic effects (Figure 1). The corner points alone can paint a misleading picture of the shape of the full surface.[3]

Adaptive Sampling: The problem of choosing which experiments to conduct is related to the sampling problem in databases. We can consider the information about the full response surface S_W to be stored as records in a (large) table T_W with attributes $x_1, \ldots, x_d, y$. An example record $\langle x_1 = v_1, \ldots, x_d = v_d, y = p \rangle$ in T_W says that the performance at the setting $\langle x_1 = v_1, \ldots, x_d = v_d \rangle$ is p for the workload W under consideration. Experiment selection is the problem of sampling from this table. However, the difference with respect to conventional sampling is that the table T_W is never fully available. Instead, we have to pay a cost—namely, the cost of running an experiment—in order to sample a record from T_W.

The gridding and SARD approaches collect a predetermined set of samples from T_W. A major deficiency of these techniques is that they are not *feedback-driven*. That is, these techniques do not use the information in the samples collected so far in order to determine which samples to collect next. (Note that conventional random sampling in databases is also not feedback-driven.) Consequently, these techniques either bring into too many samples or too few samples to address the parameter tuning problem.

iTuned uses a novel feedback-driven algorithm, called *Adaptive Sampling*, for experiment selection in parameter tuning. Adaptive Sampling analyzes the samples collected so far to understand how the surface looks like, and where the good settings are likely to be. Based on this analysis, more experiments are done to collect new samples that add maximum utility to the current samples.

Suppose n experiments have been run at settings $X^{(i)}$, $1 \leq i \leq n$, so far. Let the corresponding performance values observed be $y^{(i)} = y(X^{(i)})$. Thus, the samples collected so far are $\langle X^{(i)}, y^{(i)} \rangle$. Let $X^\star$ denote the best-performing setting found so far. Without loss of generality, we assume that the tuning goal is to minimize y.

$$X^\star = \arg \min_{1 \leq i \leq n} y(X^{(i)})$$

Which sample should Adaptive Sampling collect next? Suppose the next experiment is done at setting X, and the performance observed is $y(X)$. Then, the improvement *IP(X)* achieved by the new experiment X over the current best-performing setting $X^\star$ is:

$$IP(X) = \begin{cases} y(X^\star) - y(X) & \text{if } y(X) < y(X^\star) \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

---

[3]The authors of SARD mentioned this problem [5]. They recommended that, before invoking SARD, the DBA should split each parameter x_i with nonmonotonic effect into distinct artificial parameters corresponding to each monotonic range of x_i. This task is nontrivial since the true surface is unknown to begin with. Ideally, the DBA, who may be a naive user, should not face this burden.

Adaptive Sampling: Algorithm run by iTuned's Planner

1. Initialization: Conduct experiments based on Latin Hypercube Sampling, and initialize GRS and $X^\star = \arg\min_i y(X^{(i)})$ with collected samples;
2. Until the stopping condition is reached, do
3.   Find $X_{next} = \arg\max_{X \in DOM} EIP(X)$;
4.   Executor conducts the next experiment at X_{next} to get a new sample;
5.   Update the GRS and $X^\star$ with the new sample; Go to Line 2;

Figure 2: Steps in iTuned's Adaptive Sampling algorithm

Ideally, we would like to pick the next experiment X so that the improvement *IP(X)* is maximized. However, a proverbial chicken-and-egg problem arises here since the improvement depends on the value of $y(X)$ which will be known only after the experiment is done. We can instead compute *EIP(X)*, the *expected improvement* when the next experiment is done at setting X. Then, the experiment that gives the maximum expected improvement is selected.

$$X_{next} = \arg\max_{X \in DOM} EIP(X) \qquad (2)$$

$$EIP(X) = \int_{p=-\infty}^{p=+\infty} IP(X)\texttt{pdf}(Y(X) = p)dp \qquad (3)$$

Here, $\texttt{pdf}(Y(X) = p)$ is the probability density function of the predicted performance $y(X)$ at X. *IP(X)* is defined by Equation 1. Recall that *DOM* is the set of all feasible parameter settings.

iTuned's Workflow: The challenge in Adaptive Sampling is to compute *EIP(X)* based on the $\langle X^{(i)}, y^{(i)}\rangle$ samples collected so far. The crux of this challenge turns out to be the generation of the probability density function of the predicted performance at X.

Figure 2 shows iTuned's workflow for parameter tuning. Once invoked, iTuned starts with an initialization phase where some experiments are conducted for bootstrapping. Adaptive Sampling starts with the initial set of samples, and continues to bring in new samples through experiments selected based on *EIP(X)*. Experiments are conducted in a seamless fashion in the production environment using mechanisms provided by the executor.

Roadmap: Section 4 describes Adaptive Sampling. Details of the executor are presented in Section 5. iTuned's scalability-oriented features are described in Section 6.

4. ADAPTIVE SAMPLING
4.1 Initialization

As the name suggests, this phase bootstraps Adaptive Sampling by bringing in samples from an initial set of experiments. A straightforward technique is random sampling which will pick the initial experiments randomly from the space of possible experiments. However, random sampling is often ineffective when only a few samples are collected from a fairly high-dimensional space. More effective sampling techniques come from the family of *space-filling designs* [13]. iTuned uses one such sampling technique, called *Latin Hypercube Sampling (LHS)* [9], for initialization.

LHS selects m experiments from a space of dimension d (i.e., parameters $x_1, \ldots, x_d$) as follows: (1) the domain $dom(x_i)$ of each parameter is partitioned into m equal subdomains; and (2) m experiments are chosen from the space such that each subdomain of any parameter has one and only one sample in it. LHS has two important advantages:

- LHS samples are very efficient to generate because of their similarity to *permutation matrices* from matrix theory. Generating m LHS samples involves generating d independent permutations of $1, \ldots, m$, and joining the permutations on a position-by-position basis.

- In general, experiments done through LHS give much better space coverage than through random sampling. LHS guarantees that the settings in the chosen experiments are spread evenly over the ranges of each parameter.

However, LHS by itself does not rule out bad spreads (e.g., all samples spread along the diagonal). iTuned addresses by problem by generating multiple sets of LHS samples, and finally choosing the one which maximizes the minimum distance between any pair of samples. That is, suppose l different sets of LHS samples $L_1, \ldots, L_l$ were generated. iTuned will select the set $L^\star$ such that:

$$L^\star = \arg\max_{1 \le i \le l} \min_{X^{(j)}, X^{(k)} \in L_i, j \ne k} \texttt{dist}(X^{(j)}, X^{(k)})$$

Here, $\texttt{dist}$ is a common distance metric like the Euclidean distance. This technique avoids bad spreads.

4.2 Picking the Next Experiment

Let the samples collected so far be $\langle X^{(i)}, y^{(i)}\rangle$, $1 \le i \le n$. As discussed in Section 3, we need to compute the expected improvement that comes from doing the next experiment at a setting X. One approach is to derive a *regression model* [9] that can estimate $y(X)$ based on the $\langle X^{(i)}, y^{(i)}\rangle$ samples available so far. Such a regression model would have the form:

$$y = \vec{f}^{\,t}(X)\vec{\beta} + \varepsilon(X) \qquad (4)$$

Here, $\vec{f}(X) = [f_1(X), f_2(X), \ldots, f_h(X)]^t$ is a vector of basis functions, and $\vec{\beta}$ is the corresponding $h \times 1$ vector of regression coefficients. The t notation is used to represent the matrix transpose operation. $\varepsilon(X)$, given by $\varepsilon(X) = y(X) - \vec{f}^{\,t}(X)\vec{\beta}$, is called the *residual* because it represents the difference between the true value and the value estimated via regression. The residuals are assumed to follow identical and independent normal distributions.

For example, some response surface may be represented well by the regression model: $y = 0.1 + 3x_1 - 2x_1x_2 + x_2^2$. In this case, $\vec{f}(X) = [1, x_1, x_2, x_1x_2, x_1^2, x_2^2]^t$, and $\vec{\beta} = [0.1, 3, 0, -2, 0, 1]^t$.

Problems with conventional regression models, and iTuned's solution: Conventional regression models assume that the residuals ε_i and ε_j at any pair of settings $X^{(i)}$ and $X^{(j)}$ are independent. However, the response surface of performance with respect to parameter settings is predominantly continuous. Thus, the residuals at two nearby settings tend to be correlated, violating the assumption of independent errors in the model. A related, but bigger, problem with these models is that they do not capture the probability density function $\texttt{pdf}(Y(X))$ of the performance metric. Recall from Equation 3 that $\texttt{pdf}(Y(X))$ is required to compute the expected improvements from experiments that have not been done yet.

iTuned addresses both these problems by modeling the residual $\varepsilon(X)$ using a *Gaussian process* $Z(X)$. We first define Gaussian processes, and then describe how iTuned uses them to create the *Gaussian process Representation of a response Surface (GRS)*.

Definition 1. **Gaussian Process:** Let χ be a subspace of *DOM*. We say that $Z(X)$, for $X \in \chi$, is a Gaussian process provided that for any $l \ge 1$ and any choice of $X^{(1)}, \ldots, X^{(l)}$ in χ, the vector $[Z(X^{(1)}), \ldots, Z(X^{(l)})]$ has a multivariate normal distribution. $Z(X)$ is determined by its mean and covariance functions. $\square$

Intuitively, a Gaussian process is a stochastic process for which any finite linear combination of samples are normally distributed.

Definition 2. **Gaussian process Representation of a response Surface (GRS):** A GRS represents a response surface $y(X)$ as: $y = \vec{f}^{\,t}(X)\vec{\beta} + Z(X)$. Here, the residual in the regression is modeled by a Gaussian process $Z(X)$ with zero mean and covariance

function $\texttt{Cov}(Z(X^{(i)}), Z(X^{(j)})) = \alpha^2 corr(X^{(i)}, X^{(j)})$. $corr$ is a pairwise correlation function defined as $corr(X^{(i)}, X^{(j)}) = \Pi_{k=1}^{d} \exp(-\theta_k |x_k^{(i)} - x_k^{(j)}|^{\gamma_k})$. $\alpha, \theta_k \geq 0, \gamma_k > 0, 1 \leq k \leq d$ are constants. □

GRS's covariance function $\texttt{Cov}(Z(X^{(i)}), Z(X^{(j)}))$ represents the predominant phenomenon in response surfaces that if settings $X^{(i)}$ and $X^{(j)}$ are close to each other, then their respective residual values are correlated. As the distance between $X^{(i)}$ and $X^{(j)}$ increases, the correlation decreases. The parameter-specific constants θ_k and γ_k capture the fact that each parameter may have its own rate at which the residuals become uncorrelated. We will describe how these constants are set and give an example momentarily. GRS has the following attractive features:

- Unlike conventional regression models, GRS enables us to capture the probability density function $\texttt{pdf}(y(X))$ based on the samples collected through experiments conducted so far. We prove that GRS helps even further by enabling us to derive a closed form for *EIP(X)* from Equation 3.
- We will prove empirically using real and synthetic data that GRS is powerful enough to capture the response surfaces that arise in parameter tuning. (Gaussian processes have been used to great success on complex tasks like simulation of fire evolution and aircraft flight [13].)
- As we show momentarily, GRS enables us to naturally balance the twin tasks of *exploration* (understanding the surface) and *exploitation* (going after known high-performance regions) that arise in parameter tuning. It is nontrivial to achieve this balance, and many previous techniques [5, 17] lack it. Furthermore, GRS enables easy update as well as validation.

Lemma 1. **Prediction using GRS:** Suppose a GRS is generated from n collected samples $\langle X^{(i)}, y^{(i)} \rangle, 1 \leq i \leq n$. For any X, the GRS generates an estimate of $y(X)$ that is normally distributed with mean $\hat{y}(X)$ and variance $v^2(X)$ where:

$$\hat{y}(X) = \vec{f}^t(X)\vec{\beta} + \vec{c}^t(X)\mathbf{C}^{-1}(\vec{y} - \mathbf{F}\vec{\beta}) \qquad (5)$$

$$v^2(X) = \alpha^2[1 - \vec{c}^t(X)\mathbf{C}^{-1}\vec{c}(X)] \qquad (6)$$

$\vec{c}(X) = [corr(X, X^{(1)}), \ldots, corr(X, X^{(n)})]^t$, $\mathbf{C}$ is an $n \times n$ matrix with element i, j equal to $corr(X^{(i)}, X^{(j)}), 1 \leq i, j \leq n$, $\vec{y} = [y^{(1)}, \ldots, y^{(n)}]^t$, and $\mathbf{F}$ is an $n \times h$ matrix with the ith row composed of $\vec{f}^t(X^{(i)})$. □

Proof: Recall that the joint distribution of $y(X)$ and $Y^n = [y(X_1), y(X_2), \ldots, y(X_n)]^t$ is a $(1+n)$-dimensional Gaussian distribution

$$\binom{y(X)}{Y^n} \sim N_{1+n}\left[\binom{\vec{f}^t(X)}{F}\beta, \alpha^2 \begin{pmatrix} 1 & \vec{c}^t(X) \\ \vec{c}(X) & \mathbf{C} \end{pmatrix}\right]$$

The conditional distribution of $y(X)$ given Y^n is still a Gaussian distribution with mean and variance as expressed in Equation (5) and (6) [13]:

$$(y(X)|Y^n = [y_1, y_2, \ldots y_n]^t) \sim N\left[\hat{y}(X), v^2(X)\right]$$

□

Note that $\vec{f}^t(X)\vec{\beta}$ in Equation 5 is simply a plug in of X into the regression model from Definition 2. The second term in Equation 5 is an adjustment of the prediction based on the errors (residuals) seen at the sampled settings, i.e., $y^{(i)} - \vec{f}^t(X^{(i)})\vec{\beta}, 1 \leq i \leq n$. Intuitively, the prediction at X can be seen as a weighted sum of the values $y^{(i)}$ observed through experiments; where the weights are determined by the correlation function from Definition 2. Since
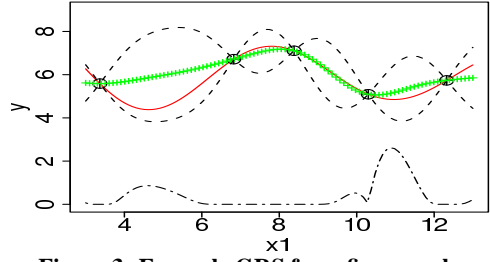

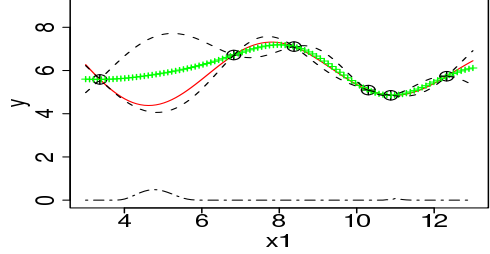**Figure 3: Example GRS from five samples**

Figure 4: Example of EIP computation

the correlation function weighs nearby settings more than distant settings, the prediction at X is affected more by y values observed at the nearby settings.

Also note that the variance at X—which is the *uncertainty* in the GRS's estimate $\hat{y}(X)$ at X—depends on the distance between X and the settings $X^{(i)}$ where experiments were done to collect samples. Intuitively, if X is close to one or more settings $X^{(i)}$ where we have collected samples, then we will have more confidence in the prediction than the case where X is far away from all settings where experiments were done. Thus, GRS captures the uncertainty in estimated values in an intuitive fashion.

Lemma 1 gives us the necessary building blocks to compute the expected improvements from experiments that have not been done yet. We first give an example to illustrate the basic ideas of GRS.

Example 1. The solid (red) line near the top of Figure 3 is a true one-dimensional response surface. Suppose five experiments are done, and the collected samples are shown as circles in Figure 3. iTuned creates a GRS from these samples. The (green) line marked with "+" symbols represents the predictions $\hat{y}(X)$ generated by the GRS as per Lemma 1. The two (black) dotted lines around this line denote the 97% confidence interval, namely, $[\hat{y}(X) - 2v(X), \hat{y}(X) + 2v(X)]$. For example, at $x_1 = 8$, the predicted value is 7.2 with confidence interval [6.4, 7.9]. Note that, at all points, the true value (solid line) is within the confidence interval; meaning that the GRS learned from the five samples is a good approximation of the true response surface. Also, note that at points close to the collected samples, the uncertainty in prediction is low. The uncertainty increases as we move further from the collected samples. □

Recall from Lemma 1 that the estimate of $y(X)$ based on the n collected samples $\langle X^{(i)}, y^{(i)} \rangle, 1 \leq i \leq n$, is normally distributed with mean $\hat{y}(X)$ and variance $v^2(X)$. Hence it follows that the probability density function of $y(X)$ is:

$$\texttt{pdf}(Y(X) = p) = \frac{1}{\sqrt{2\pi}v(X)} \exp(\frac{-(p - \hat{y}(X))^2}{2v^2(X)}) \qquad (7)$$

Theorem 1. The expected improvement from conducting an experiment at X is:

$$EIP(X) = \int_{p=-\infty}^{p=y(X^\star)} (y(X^\star) - p)\texttt{pdf}(Y(X) = p)dp \qquad (8)$$

EIP(X) has the following closed form:

$$EIP(X) = v(X)[\mu(X)\Phi(\mu(X)) + \phi(\mu(X))] \qquad (9)$$

Here, $\mu(X) = \frac{y(X^\star) - \hat{y}(X)}{v(X)}$. Φ and ϕ are $N(0, 1)$ normal cumulative distribution and density functions respectively.

Proof: Substituting Equation 1 into Equation 3, we have

$$
\begin{aligned}
EIP(X) &= \int_{p=-\infty}^{p=+\infty} IP(X)\texttt{pdf}(Y(X) = p)dp \\
&= \int_{p=-\infty}^{p=y(X^\star)} (y(X^\star) - p)\texttt{pdf}(Y(X) = p)dp \\
&= \int_{p=-\infty}^{p=y(X^\star)} [y(X^\star) - \hat{y}(X) \\
&\quad + \hat{y}(X) - p]\texttt{pdf}(Y(X) = p)dp
\end{aligned}
$$

Note that $\int_{p=-\infty}^{p=y(X^\star)} [y(X^\star) - \hat{y}(X)]\texttt{pdf}(Y(X) = p)dp$

$$
\begin{aligned}
&= [y(X^\star) - \hat{y}(X)]\Phi(\frac{y(X^\star) - \hat{y}(X)}{v(X)}) \\
&= v(X)\mu(X)\Phi(\mu(X))
\end{aligned}
$$

and $\int_{p=-\infty}^{p=y(X^\star)} [\hat{y}(X) - p]\texttt{pdf}(Y(X) = p)dp$

$$
\begin{aligned}
&= -\int_{t=-\infty}^{t=\frac{y(X^\star)-\hat{y}(X)}{v(X)}} t * v(X)\phi(t)dt \quad \{\text{let } t = \frac{p - \hat{y}(X)}{v(X)}\} \\
&= v(X)\phi(\mu(X))
\end{aligned}
$$

So

$$
\begin{aligned}
EIP(X) &= v(X)\mu(X)\Phi(\mu(X)) + v(X)\phi(\mu(X)) \\
&= v(X)[\mu(X)\Phi(\mu(X)) + \phi(\mu(X))]
\end{aligned}
$$

$\square$

Therefore, the next experiment should be run at the setting

$$
X_{next} = \arg \max_{X \in DOM} EIP(X)
$$

Recall that *DOM* is the set of all the feasible configuration settings. Intuitively, the next experiment to run should be picked from regions where there is high uncertainty, which is expressed as $v(X)$ in (9), or the predicted value can improve over the current best setting, which is expressed as $\mu(X)$ in (9). In regions where the current GRS from the observed samples is uncertain about its estimate, i.e., where $v(X)$ is high, exploration is preferred to reduce the model uncertainty. At the same time, in regions where it is possible to achieve better performance, i.e. $\mu(X)\Phi(\mu(X)) + \phi(\mu(X))$ is high, the current GRS is used to pick samples around the current good setting $X^\star$ for exploitation. There is a tradeoff between exploration (global search) and exploitation (local search).

Example 2. The dotted line at the bottom of Figure 3 shows EIP*(X) along the x_1 dimension. (All* EIP *values have been scaled by 40 to make the plot fit in this figure.) There are two peaks in the* EIP *plot. (I)* EIP *values are high around the current best sample (X^* with x_1=10.3), encouraging local search (exploitation) in this region. (II)* EIP *values are also high in the region between x_1=4 and x_1=6 because no samples have been collected near this region; the higher uncertainty motivates exploring this region. Adaptive Sampling with conduct the next experiment at the highest* EIP *point, namely, x_1=10.9. Figure 4 shows the new set of samples as well as the new* EIP*(X) after the GRS is updated with the new sample. As expected,* EIP *around x_1=10.9 has reduced.* EIP*(X) now has a maximum value at x_1=4.7 because the uncertainty in this region is*

still high. Adaptive Sampling will experiment here next, bringing in a sample close to the global optimum at x_1=4.4.

4.3 Overall Algorithm and Implementation

Figure 2 shows the overall structure of iTuned's Adaptive Sampling algorithm. So far we described how the initialization is done and how *EIP(X)* is derived. We now discuss how iTuned implements the other steps in Figure 2.

Finding the Setting that Maximizes EIP: Line 3 in Figure 2 requires us to find the setting $X \in DOM$ that has the maximum *EIP*. Since we have a closed form for EIP, it is efficient to evaluate EIP at a given point. In our implementation, we pick $k = 1000$ settings (using LHS sampling) from the space of feasible settings, compute their EIP values, and pick one that has the maximum value to run the next experiment.

Initializing the GRS and Updating it with New Samples: It follows from Definition 2 that initializing the GRS with a set of $\langle X^{(i)}, y^{(i)} \rangle$ samples, or updating the GRS with a newly collected sample, involves deriving the best values of the constants α, θ_k, and γ_k, for $1 \le k \le d$, based on the current samples. This step can be implemented in different ways. Our current implementation uses the well-known and efficient statistical technique of *maximum likelihood estimation* [21].

When to Stop: When does Adaptive Sampling stop (Line 2 in Figure 2)? The easy case is when the user issues an explicit stop command once they are satisfied with the tuned performance. iTuned incorporates a novel stopping condition that can handle the harder cases, namely, when iTuned is invoked (i) in the auto-tuning mode, and (ii) by a nonexpert user.

Intuitively, Adaptive Sampling can stop when the maximum expected improvement over all settings $X \in DOM$ falls below a threshold. However, there is a possible pitfall: if the current GRS does not represent the underlying response surface reasonably well, then the expected improvement values at some settings X may differ from the actual improvement that X gives. iTuned safeguards against this problem by leveraging the properties of a GRS and the statistical testing methodology of *cross validation* [21].

Let $\langle X^{(j)}, y^{(j)} \rangle$, $1 \le j \le n$ be the set of samples collected so far. iTuned performs the following test:

1. Remove the sample $\langle X^{(i)}, y^{(i)} \rangle$ from the set.
2. Use the remaining $n - 1$ samples to generate a GRS, and use it to predict the performance at $X^{(i)}$. Recall from Lemma 1 that this prediction has a normal distribution with some mean, denoted $\hat{y}_{-i}$, and variance, denoted v_{-i}^2. (The subscript $-i$ indicates that the sample $\langle X^{(i)}, y^{(i)} \rangle$ is not used.)
3. Based on the properties of standard normal distributions, a popular test is done to check whether $z_{-i} = \frac{\hat{y}_{-i} - y^{(i)}}{v_{-i}}$ lies within the 97% *confidence interval*. (The test succeeds if $-2 \le z_{-i} \le 2$.)

The above steps are repeated for each of the n samples by removing them one at a time. If the z_{-i} value in each case lies within the 97% confidence interval, then, with high probability, the GRS from the current n samples is a good representation of underlying true response surface.[4] If that is true, and the maximum expected improvement is below a threshold, then Adaptive Sampling can stop.

5. ITUNED'S EXECUTOR: A PLATFORM FOR RUNNING ONLINE EXPERIMENTS

We now consider where and when iTuned will run experiments. There are some simple answers. If parameter tuning is done before

---

[4]While we collect a fixed number of samples during initialization, the same test could be used to find the number of initial samples.

the database goes into production use, then the experiments can be done on the production platform itself. If the database is already in production use and serving real users and applications, then experiments could be done on an offline test platform. Previous work on parameter tuning (e.g., [5, 17]) assume that experiments are conducted in one of these settings.

While the two settings above—preproduction database and test database—are practical solutions, there are not sufficient because:

- The workload may change while the database is in production use, necessitating retuning.
- A test database platform may not exist (e.g., in an SMB).
- It can be nontrivial or downright infeasible to replicate the production resources, data, and workload on the test platform.

iTuned's executor provides a comprehensive solution that addresses concerns like these. The guiding principle behind the solution is: exploit underutilized resources in the production environment for experiments, but never harm the production workload. The two salient features of the solution are:

- **Designated resources:** iTuned provides an interface for users to designate *which* resources can be used for running experiments. Candidate resources include (i) the production database (the default for running experiments), (ii) standby (failover) databases backing up the production database, (iii) test database(s) used by DBAs and developers, and (iv) staging database(s) used for end-to-end testing of changes (e.g., bug fixes) before they are applied to the production database. Resources designated for experiments are collectively called the *workbench*.
- **Policies:** A policy is specified with each resource that dictates *when* the resource can be used for experiments. The default policy associated with each of the above resources is: "if the CPU, memory, and disk utilization of the resource for its *home use* is below 10% (threshold t_1) for the past 10 minutes (threshold t_2), then the resource can be used for experiments." Home use denotes the regular (i.e., nonexperimental) use of the resource. The two thresholds are customizable. Only the default policy is implemented currently, but we are exploring other policies.

iTuned's implementation consists of a front-end that interacts with users, and a back-end consisting of the planner which plans experiments using Adaptive Sampling, and the executor which schedules planned experiments on the workbench as per user-specified (or default) policies. Monitoring data needed to enforce policies is obtained through database monitoring tools.

The design of the workbench is based on splitting the functionality of each resource into two: (i) *home use*, where the resource is used directly or indirectly to support the production workload, and (ii) *garage use*, where the resource is used to run experiments. We will describe the home/garage design using the standby database as an example, and then generalize to other resources.

All database systems support one or more hot standby databases whose home use is to keep up to date with the (primary) production database by applying redo logs shipped from the primary. If the primary fails, a standby will quickly take over as the new primary. Hence, the standby databases run the same hardware and software as the production database. It has been observed that standby databases usually have very low utilization since they only have to apply redo log records. In fact, [7] mentions that enterprises that have 99.999% (five nines) availability typically have standby databases that are idle 99.999% of the time.

Thus, the standby databases are a valuable and underutilized asset that can be used for online experiments without impacting user-facing queries. However, their home use should not be affected, i.e., the recovery time on failure should not have any noticeable increase. iTuned achieves this property using two *resource con-*
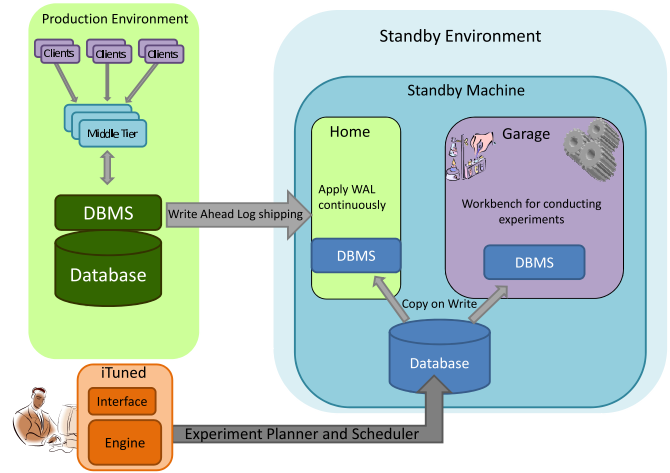

Figure 5: The executor in action for standby databases

tainers: the home container for home use, and the garage container for running experiments. iTuned's current implementation of resource containers using the *zones* feature in the Solaris OS [14]. CPU, memory, and disk resources can be allocated dynamically to a zone, and the OS provides isolation between resources allocated to different zones. Resource containers can also be implemented using virtual machine technology which is becoming popular [15].

The home container on the standby machine is responsible for applying the redo log records. When the standby machine is not running experiments, the home container runs on it using all available resources; the garage lies idle. The garage container is *booted*—similar to a machine booting, but much faster—only when a policy fires and allows experiments to be scheduled on the standby machine. During an experiment, both the home and the garage containers will be active, with a partitioning of resources as determined by the executor. Figure 5 provides an illustration. For example, as per the default policy stated earlier, home and garage will get 10% and 90%, repectively, of the resources on the machine.

Both the home and the garage containers run a full and exactly the same copy of the database software. However, on booting, the garage is given a *snapshot* of the current data (including physical design) in the database. The garage's snapshot is logically separate from the snapshot used by the home container, but it is physically the same except for *copy-on-write* semantics. Thus, both home and garage have logically-separate copies of the data, but only a single physical copy of the data exists on the standby system when the garage boots. When either container makes an update to the data, a separate copy of the changed part is made that is visible to the updating container only (hence the term copy-on-write). The redos applied by the home container do not affect the garage's snapshot. iTuned's implementation of snapshots and copy-on-write semantics leverages the Zettabyte File System [14], and is extremely efficient (as we will show in the empirical evaluation).

The garage is *halted* immediately under three conditions: when experiments are completed or the primary fails or there is a policy violation. All resources are then released to the home container which will continue functioning as a pure standby or take over as the primary as needed. Setting up the garage (including snapshots and resource allocation) takes less than a minute, and tear-down takes even less time. The whole process is so efficient that recovery time is not increased by more than a few seconds.

While the above description focused on the standby resource, iTuned applies the same home/garage design to all other resources in the workbench (including the production database). The only difference is that each resource has its own distinct type of home

Feature	Description and Use
Sensitivity analysis	Identify and eliminate low-effect parameters
Parallel experiments	Use multiple resources to run parallel expts
Early abort	Identify and stop low-utility expts quickly
Workload compression	Reduce per-experiment running time without reducing overall tuning quality
Semantic knowledge	Exploit *advisory parameters* in database systems
Incremental tuning	Cluster parameters to ensure independent effects across clusters; tune one cluster at a time
Interactive tuning	Get user feedback from intermediate results

Table 2: Features that improve iTuned's efficiency

use which is encapsulated cleanly into the corresponding home container. *Thus, iTuned works even in settings where there are no standby or test databases.*

6. IMPROVING ITUNED'S EFFICIENCY

Experiments take time to run. This section describes features that can reduce the time iTuned takes to return good results as well as make iTuned scale to large numbers of parameters. Table 2 gives a short summary. The first three features are fully integrated into iTuned, workload compression is currently a simple standalone tool, and the last three features will be implemented in future.

6.1 Eliminating Unimportant Parameters Using Sensitivity Analysis

Suppose we have generated a GRS using n samples $\langle X^{(i)}, y^{(i)} \rangle$. Recall that, given any setting X, the GRS can produce a prediction with mean $\hat{y}(X)$ and variance $v^2(X)$. Using the GRS, we can compute the expected value of y when $x_1=v$ as:

$$E(y|x_1=v)=\int_{dom(x_2)} \cdots \int_{dom(x_d)} \hat{y}(v, x_2, \ldots, x_d) dx_2 \cdots dx_d \quad (10)$$

Intuitively, Equation 10 averages out the effects of all parameters other than x_1, and $E(y|x_1)$ is a function of x_1 measuring its effect on y. If we consider l equally-spaced values $v_i \in dom(x_1)$, $1 \leq i \leq l$, then we can use Equation 10 to compute the expected value of y at each of these l points. A plot of these values, e.g., as shown in Figure 3, gives a visual feel of the overall effect of parameter x_1 on y. We term such plots *effect plots*. In addition, we can consider the variance of these values, denoted $V_1 = \text{Var}(E(y|x_1))$. Intuitively, if V_1 is low, then y does not vary much as x_1 is changed; hence, the effect of x_1 on y is low. On the other hand, large V_1 means that y is sensitive to x_1's setting.

Therefore, we define the *main effect* of x_1 as $\frac{V_1}{\text{Var}(y)}$ which represents the fraction of the overall variance in y that is explained by the variance seen in $E(y|x_1)$. The main effect of the other parameters $x_2, \ldots, x_d$ is defined in a similar fashion. Any parameter with low main effect can be set to its default value with little negative impact on performance, and need not be considered for tuning.

6.2 Running Multiple Experiments in Parallel

If the executor can find enough resources on the workbench, then iTuned can run $k > 1$ experiments in parallel. (Section 9 discusses how cloud computing is making resources cheaper to acquire.) The batch of experiments from LHS during initialization can be run in parallel. Running k experiments from Adaptive Sampling in parallel is nontrivial because of its sequential nature. A naive approach is to pick the top-k settings that maximize EIP. However, the pitfall is that these k samples may be from the same region (around the current minimum or with high uncertainty), and hence redundant.

We set two criteria for selecting k parallel experiments: (I) Each experiment should improve the current best value (in expectation);

(II) The selected experiments should complement each other in improving the GRS's quality. iTuned determines the next k experiments to run in parallel as follows:

1. Select the experiment $X^{(i)}$ that maximizes the current EIP.
2. An important feature of GRS is that the uncertainty in prediction (Equation 6) depends only on the X values of collected samples. Thus, after $X^{(i)}$ is selected, we update the uncertainty estimate at each remaining candidate setting. (The predicted value, from Equation 5, at each candidate remains unchanged.)
3. We compute the new EIP values with the updated uncertainty term $v(X)$, and pick the next sample $X^{(i+1)}$ that maximizes EIP. The nice property is that $X^{(i+1)}$ will not be clustered with $X^{(i)}$: after $X^{(i)}$ is picked, the uncertainty in the region around $X^{(i)}$ will reduce, therefore EIP will decrease in that region.
4. The above steps are repeated until k experiments are selected.

6.3 Early Abort of Low-Utility Experiments

While the exploration aspect of Adaptive Sampling has its advantages, it can cause experiments to be run at poorly-performing settings. Such experiments take a long time to run, and contribute little towards finding good parameter settings. To address this problem, we added a feature to iTuned where an experiment at $X^{(i)}$ is aborted after $\Delta \times t_{min}$ time if the workload running time at $X^{(i)}$ is greater than $\Delta \times t_{min}$. Here, t_{min} is the workload running time at the best setting found so far. Be default, $\Delta = 2$.

6.4 Workload Compression

Work on physical design tuning has shown that there is a lot of redundancy in real workloads which can be exploited through workload compression to give 1-2 orders of magnitude reduction in tuning time [3]. Reference [3] proposed an approach where the given workload is partitioned based on distinct query templates, and a representative subset is picked per partition via clustering. To demonstrate the utility of workload compression in iTuned, we came up with a modified approach. We treat a workload as a series of execution of query mixes, where a query mix is a set of queries that run concurrently. An example could be $\langle 3Q_1, 6Q_{18} \rangle$ which denotes three instances of TPC-H query Q_1 running concurrently with six instances of Q_{18}. We partition the given workload into distinct query mixes, and pick the top-k mixes based on the overall time for which each mix ran in the workload.

6.5 Using Database-specific Knowledge

It is common to have database parameters whose settings affect the query execution plan chosen by the optimizer, but do not affect anything else including resource allocation and database configuration. We term such parameters *advisory parameters*. PostgreSQL's *effective_cache_size parameter* (recall Section 1) is an example. More common examples include parameters used as inputs to the optimizer's cost model, e.g., the cost of a sequential I/O.

Consider two settings $X^{(i)}$ and $X^{(j)}$ that differ in the settings of advisory parameters only. Despite this difference, suppose the optimizer picks the same set of execution plans for $X^{(i)}$ and $X^{(j)}$. If iTuned "knows" about advisory parameters, then it can avoid running an experiment at $X^{(j)}$ if an experiment has already been done at $X^{(i)}$ (since the same plans would run in the same environment). This optimization is important and frequently applicable because typical databases have a number of advisory parameters, most of which are high-impact because they can change execution plans.

6.6 Other Techniques

One approach for scalability is analyze interactions among the effects of different parameters. Recall that the main effect of parameter x_1 is defined as $\frac{V_1}{\text{Var}(y)}$. Similarly, an *interaction effect* be-

tween x_1 and x_2 can be defined as $\frac{\mathrm{Var}(E(y|x_1,x_2))-V_1-V_2}{\mathrm{Var}(y)}$, where:

$$E(y|x_1=v_1, x_2=v_2) = \int_{dom(x_3)} \cdots \int_{dom(x_d)} \hat{y}(v_1, v_2, x_3, \ldots, x_d) dx_3 \cdots dx_d$$

Intuitively, the interaction effect between x_1 and x_2 is high if the effect of x_1 on y is very sensitive to x_2's setting. That is, different settings of x_2 cause different effects from x_1. We can identify important interaction effects using the above equation, and then partition the parameters in disjoint groups such that no cross-group interactions exist. iTuned could then take a *divide-and-conquer* approach to parameter tuning, i.e., tuning one group of parameters at a time, probably ranking the groups in some order.

7. EMPIRICAL EVALUATION

Our experimental setup involves a local cluster of machines, each with four 2GHz processors and 3GB memory, running PostgreSQL 8.2 on Solaris 10. One machine runs the production database. The other machines are used as hot standbys, test platforms, or workload generators. Recall from Section 5 that iTuned's policy-based executor can run experiments on the production database, standbys, and test platforms. By default, we use a standby database for experiments.

7.1 Methodology and Summary

We first summarize the different types of empirical evaluation conducted and the results obtained.

- Section 7.2 breaks down the overhead of various operations in the API provided by iTuned's executor, and shows that the executor is noninvasive and efficient.
- Section 7.3 shows real response surfaces that highlight the issues motivating our work, e.g., (i) why database parameter tuning is not easy for the average user; (ii) how parameter effects are highly sensitive to workloads, data properties, and resource allocations; and (iii) why optimizer cost models are insufficient for effective parameter tuning, but it is important to keep the optimizer in the tuning loop.
- Section 7.4 presents tuning results for OLAP and OLTP workloads of increasing complexity that show iTuned's ease of use and up to 10x improvements in performance compared to default parameter settings, rule-based tuning based on popular heuristics, and a state-of-the-art automated parameter tuning technique. We show how iTuned can leverage parallelism, early aborts, and workload compression to cut down tuning times drastically with negligible degradation in tuning quality.
- iTuned's performance is consistently good with both PostgreSQL and MySQL databases, demonstrating iTuned's portability.
- Section 7.5 shows how iTuned can be useful in other ways apart from recommending good parameter settings, namely, visualizing parameter impact as well as approximate response surfaces. This information can guide further manual tuning.

The tuning tasks in our empirical evaluation consider up to 26 database configuration parameters. By default, we consider the following 11 parameters for OLAP workloads in PostgreSQL: (P1) shared_buffers, (P2) effective_cache_size, (P3) work_mem, (P4) maintenance_work_mem, (P5) default_statistics_target, (P6) random_page_cost, (P7) cpu_tuple_cost, (P8) cpu_index_tuple_cost, (P9) cpu_operator_cost, (P10) memory allocation, and (P11) CPU allocation. Table 3 gives the exhaustive list of all the parameters.

7.2 Performance of iTuned's Executor

We first analyze the overhead of the executor for running experiments. Recall its implementation from Section 5. Table 4 shows the various operations in the interface provided by the executor, and the overhead of each operation. The Create Container operation is done once to set up the OS environment for a particular

Operation by Executor	Time (sec)	Description
Create Container	610	Create a new garage (one time process)
Clone Container	17	Clone a garage from already existing one
Boot Container	19	Boot garage from halt state
Halt Container	2	Stop garage and release resources
Reboot Container	2	Reboot the garage (required for adding additional resources to a container)
Snapshot-R DB	7	Create read-only snapshot of database
Snapshot-RW DB	29	Create read-write snapshot of database

Table 4: Overheads of operations in iTuned's executor

tuning task; so its 10-minute cost is amortized over an entire tuning session. This overhead can be cut down to 17 seconds if the required type of container has already been created for some previous tuning task. Note that all the other operations take on the order of a few seconds. For starting a new experiment, the cost is at most 48 seconds to boot the container and to create a read-write snapshot of the database (for workloads with updates). A container can be halted within 2 seconds, which adds no noticeable overhead if, say, the standby has to take over on a failure of the primary database.

7.3 Why Parameter Tuning is Nontrivial

The OLAP (Business Intelligence) workloads used in our evaluation were derived from TPC-H running at scale factors (*SF*) of 1 and 10 on PostgreSQL [18]. The physical design of the databases are well tuned, with indexes approximately tripling and doubling the database sizes for SF=1 and SF=10 respectively. Statistics are always up to date. The heavyweight TPC-H queries in our setting include Q1, Q7, Q9, Q13, and Q18.

Figure 1 shows a 2D projection of a response surface that we generated by running Q18 on a TPC-H SF=1 database for a number of different settings of the eleven parameters from Section 7.1. The database size with indexes is around 4GB. The physical memory (RAM) given to the database is 1GB to create a realistic scenario where the database is 4x the amount of RAM. This complex response surface is the net effect of a number of individual effects:

- Q18 (Large Volume Customer Query) is a complex query that joins the Lineitem, Customer, and Order tables. It also has a subquery over Lineitem (which gets rewritten as a join), so Q18 accesses Lineitem—the biggest table in TPC-H—twice.
- Different execution plans get picked for Q18 in different regions of the response surface because changes in parameter settings lead to changes in estimated plan costs. These plans differ in operators used, join order, and whether the same or different access paths are used for the two accesses to the Lineitem table.
- Operator behavior can change as we move through the surface. For example, hybrid hash joins in PostgreSQL change from one pass to two passes if the *work_mem* parameter is lower than the memory required for the hash join's build phase.
- Resource interference can happen. For example, if a hybrid hash join in PostgreSQL starts to create temporary files on disk, the accesses go through the OS file cache which competes for RAM with *shared_buffers*. Thus, increasing *shared_buffers* can degrade performance if hybrid hash joins are spilling to disk.

It took us several days of effort, more than a hundred experiments with PostgreSQL, as well as email conversations with PostgreSQL developers to understand the unexpected nature of Figure 1. (We point the interested reader to a commentary at [6].) It is unlikely that a non-expert who wants to use a database for some application—say, Peter in Section 1—will have the knowledge (or patience) to tune the database like we did. Surfaces like Figure 1 show how critical experiments are to understand which of many different effects dominate in a particular setting.

Table 3: Parameters considered

SNo	Parameter	Description
p1	shared_buffers	Shared buffers defines a block of memory that PostgreSQL will use to hold requests that are awaiting attention from the kernel buffer and CPU
p2	effective_cache_size	Effective cache size allows PostgreSQL to make best possible use of RAM available the server. It tells PostgreSQL the size of OS data cache. So that PostgreSQL can draw different execution plan based on that data.
p3	work_mem	Work mem sets maximum limit on memory that a database connection can use to perform sorts.
p4	default_statistics_target	Sets the default statistics target for table columns
p5	random_page_cost	Sets the planner's estimate of the cost of a nonsequentially fetched disk page
p6	cpu_tuple_cost	Sets the planner's estimate of the cost of processing each row during a query
p7	cpu_index_tuple_cost	Sets the planner's estimate of the cost of processing each index row during an index scan
p8	cpu_operator_cost	Sets the planner's estimate of the cost of processing each operator in a WHERE clause
p9	maintenance_work_mem	Used for maintenance operations like CREATE INDEX, VACUUM and ALTER TABLE ADD FOREIGN KEY
p10	checkpoint_segments	Maximum distance between automatic WAL checkpoints, in log file segments (each segment is normally 16 megabytes)
p11	checkpoint_timeout	Maximum time between automatic WAL checkpoints, in seconds
p12	wal_buffers	Number of disk-page buffers allocated in shared memory for WAL data
p13	max_prepared_connections	Sets the maximum number of transactions that can be in the "prepared" state simultaneously
p14	autovacuum	If on, automates the execution of VACUUM and ANALYZE
p15	fsync	If fsync is on, then PostgreSQL make sures that updates are physically written to disk
p16	ebs	Number of emulated browsers for simulating TPC-W workload
p17	workloadtype	Workload type for TPC-W. It can be of three types: Browsing mix, Ordering mix and Shopping mix
p18	mysqld_table_cache	MySQL shared_buffers per table
p19	mysqld_sort_buffer_size	MySQL work_mem
p20	mysqld_key_buffer_size	MySQL shared_buffers
p21	buy	RUBiS parameter indicating the number of buyers in an auction
p22	browse	RUBiS parameter indicating the number of browsing connections
p23	sell	RUBiS parameter indicating the number of sellers in an auction
p24	aboutMe	RUBiS parameter indicating the number of connections checking aboutMe information
p25	memory	Amount of memory available
p26	CPU	Amount of CPU available

The average running time of a query can change drastically depending on whether it is running alone in the database or it is running in a concurrent mix of queries of the same or different types. For example, consider Q18 running alone or in a mix of six concurrent instances of Q18 (each instance has distinct parameter values). At the default parameter setting of PostgreSQL for TPC-H SF=1, we have observed the average running time of Q18 to change from 46 seconds (when running alone) to 1443 seconds (when running in the mix). For TPC-H SF=10, there was a change from 158 seconds (when running alone) to 578 seconds (when running in the mix).

Two insights come out from the results presented so far. (More such results are in the technical report [6].) First, query optimizers compute the cost of a plan independent of other plans running concurrently. Thus, optimizer cost models cannot capture the true performance of real workloads which consist of query mixes. Second, it is important to keep the optimizer in the loop while tuning parameter settings because the optimizer can change the plan for a query when we change parameter settings. While keeping the optimizer in the loop is accepted practice for physical design tuning (e.g., [4]), to our knowledge, we are the first to bring out its importance and enable its use in configuration parameter tuning.

Figure 6 shows a 2D projection of the response surface for Q18 when run in the 6-way mix in PostgreSQL for TPC-H SF=10. The key difference between Figures 1 (Q18 alone, TPC-H SF=1) and 6
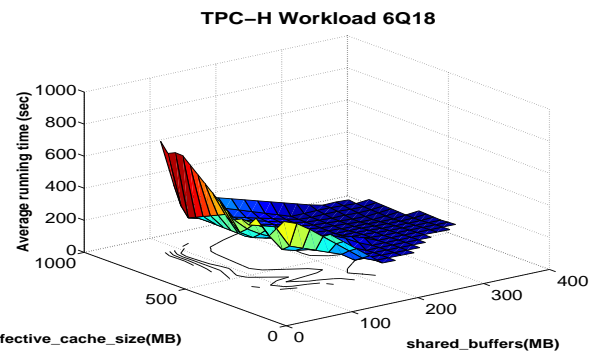
Figure 6: Impact of shared_buffers Vs. effective_cache_size for workload W4 (TPC-H SF=10)

(Q18 in 6-way mix, TPC-H SF=10) is that increasing *shared_buffers* has an overall negative effect in the former case, while the overall effect is positive in the latter. We attribute the marked effect of *shared_buffers* in Figure 6 to the increased cache hits across concurrent Q18 instances. Figures 7 and 8 show the response surface for a workload where *shared_buffers* has limited impact. The highest impact parameter is *work_mem*. This workload has three instances of Q7 and 3 instances of Q13 running in a 6-way mix in PostgreSQL for TPC-H SF=10. All these results show why users can have a hard time setting database parameters, and why experiments that
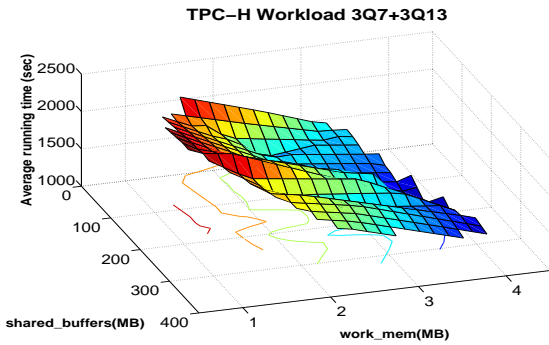
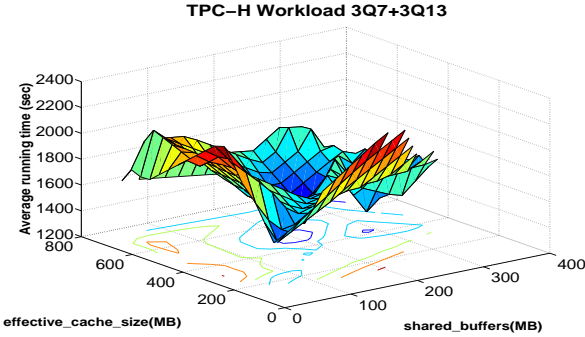**Figure 7: Impact of shared_buffers Vs. work_mem for workload W5 (TPC-H SF=10)**

Figure 8: Impact of shared_buffers Vs. effective_cache_size for workload W5 (TPC-H SF=10)

can bring out the underlying response surfaces are inevitable.

7.4 Tuning Results

We now present an evaluation of iTuned's effectiveness on different workloads and environments. iTuned should be judged both on its *quality*—how good are the recommended parameter settings?—and *efficiency*—how soon can iTuned generate good recommendations? Our evaluation compares iTuned against:

- **Default** parameter settings that come with the database.
- **Manual rule-based** tuning based on heuristics from database administrators and performance tuning experts. We use an authoritative source for PostgreSQL tuning [12].
- **Smart Hill Climbing (SHC)** is a state-of-art automated parameter tuning technique [22]. It belongs to the hill-climbing family of optimization techniques for complex response surfaces. Like iTuned, SHC plans experiments while balancing exploration and exploitation (Section 4.2). But, SHC lacks key features of iTuned like GRS representation of response surfaces, executor, and efficiency-oriented features like parallelism, early aborts, sensitivity analysis, and workload compression.
- **Approximation to the optimal setting:** Since we do not know the optimal performance in any tuning scenario, we run a large number of experiments offline for each tuning task. We have done at least 100 (often, 1000+) experiments per tuning task over the course of six months; the detailed numbers are in [6]. The best performance found is used as an approximation of the optimal. This technique is labeled *Brute Force*.

iTuned and SHC do 20 experiments each by default. iTuned uses the first 10 experiments for initialization. Strictly for the purposes of evaluation, by default iTuned uses only early abort among the efficiency-oriented techniques from Section 6.

Figure 9 compares the tuning quality of iTuned (I) with Default (D), manual rule-based (M), SHC (S), and Brute Force (B) on a range of TPC-H workloads at SF=1 and SF=10. The performance metric of interest is workload running time; lower is better. The

workload running time for D is always shown as 100%, and the times for others are relative. (The absolute numbers are in [6].) To further judge tuning quality, these figures show the rank of the performance value that each technique finds. Ranks are reported with the prefix R, and are based on the range of performance values observed by Brute Force; lower rank is always better. Figures 9 also shows (above I's bar) the total time that iTuned took since invocation to give the recommended setting. Detailed analysis of tuning times is done later in this section.

11 distinct workloads are used in Figure 9, all of which are nontrivial to tune. Workloads W1, W2, and W3 consist of individual TPC-H queries Q1, Q9, and Q18 respectively running at a *Multi-Programming Level (MPL)* of 1. MPL is the maximum number of concurrent queries. TPC-H queries have input parameters. Throughout our evaluation, we generate each query instance randomly using the TPC-H query generator *qgen*. Different instances of the same query are distinct with high probability.

Workloads W4, W5, and W6 go one step higher in tuning complexity because they consist of mixes of concurrent queries. W4 (MPL=6) consists of six concurrent (and distinct) instances of Q18. W5 (MPL=6) consists of three concurrent instances of Q7 and three concurrent instances of Q13. W6 (MPL=10) consists of five concurrent instances of Q5 and five concurrent instances of Q9.

Workloads W7 and higher in Figure 9 go the final step in tuning complexity by bringing in many more complex query types, much larger numbers of query instances, and different MPLs. W7 (MPL=9) contains 200 query instances comprising queries Q1 and Q18, in the ratio 1:2. W8 (MPL=24) contains 200 query instances comprising TPC-H queries Q2, Q3, Q4, and Q5, in the ratio 3:1:1:1. W9 (MPL=10), W10 (MPL=20), and W11 (MPL=5) contain 100 query instances each with 10, 10, and 15 distinct TPC-H query types respectively in equal ratios. The results for W7-N shown in Figure 9 are from tuning 30 parameters.

Figure 9 shows that the parameter settings recommended by iTuned consistently outperform the default settings, and is usually significantly better than the settings found by SHC and common tuning rules. iTuned gives 2x-5x improvement in performance in many cases. In fact, iTuned's recommendation is usually close in performance to the approximate optimal setting found (exhaustively) by Brute Force. It is interesting to note that expert tuning rules are more geared towards complex workloads (compare the M bars between the top and bottom halves of of Figure 9).

As an example, consider the workload W7-SF10 in Figure 9. The default settings give a workload running time of 1085 seconds. Settings based on tuning rules and SHC give running times of 386 and 421 seconds respectively. In comparison, iTuned's best setting after initialization gave a performance of 318 seconds, which was improved to 246 seconds by Adaptive Sampling (77% improvement over default). iTuned's sensitivity analysis found the shared_buffers parameter to have the most impact on performance. The default setting of 32 MB for shared_buffers is poor. The rule-based setting of 200 MB is better, but iTuned found a setting close to 400 MB where the performance is far better.

Figure 9 shows that iTuned takes on the order of tens of hours to find good settings for complex workloads. Figure 10 gives the absolute tuning values by executing a single instance of workload in seconds. Reference [10] estimates that configuring large database management systems takes on the order of one to two weeks, so one to two days of time spent parameter tuning is acceptable; especially considering that iTuned gives 2x-5x improvement in performance in many cases. More importantly, Figure 11 shows that iTuned's tuning time can be reduced by orders of magnitude using the techniques we proposed in Section 6. Early Abort uses $\Delta = 2$ and
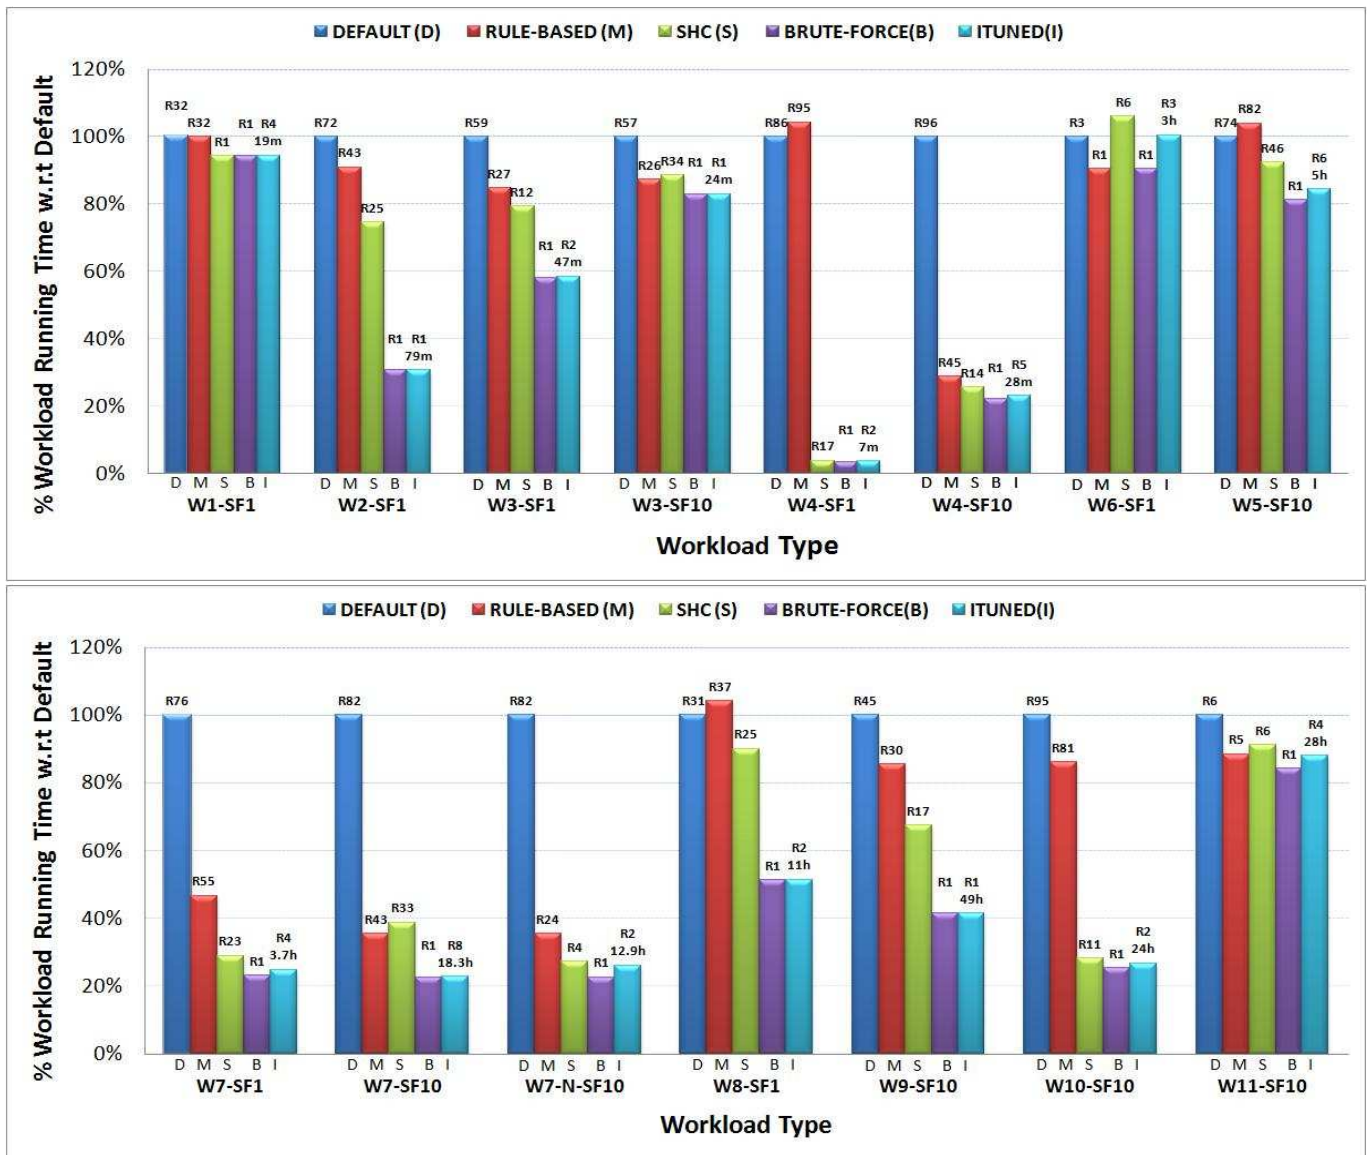
Figure 9: Comparison of tuning quality. iTuned's tuning times are shown in minutes (m) or hours (h). Ri denotes Rank i

workload compression picks the top mix in the workload.

For each of the complex workloads from Figure 9, we show iTuned's tuning time with and without different techniques. It is clear that these techniques can reduce iTuned's tuning time to at most a few hours. The drop in tuning quality across all these scenarios was never more than 1%. In general, we have found workload compression to be even more effective in parameter tuning than in physical design tuning. Intuitively, parameter settings are less sensitive to which queries get picked in the compressed workload compared to, say, index selection.

Because of space constraints, we have only given representative results in this paper. A number of other empirical results—including OLTP workloads, MySQL, and different performance metrics—are given in [6]. Table 5 gives a brief summary that shows iTuned's consistent good performance. TPC-W is an e-Commerce benchmark that simulates the activities of a retail website. Our experiments with TPC-W are based on a 48000-transaction workload

Workload	Perf. Metric	#Params	Quality (Rank)	Tuning time (Hours)
TPC-W (MySQL)	Response time	7	R1	3.2
TPC-W (MySQL)	Throughput	7	R4	7.6
TPC-W (PostgreSQL)	Response time	20	R23	2.5
TPC-W (PostgreSQL)	Throughput	20	R8	2.5
RUBiS (MySQL)	Response time	6	R1	6.1
RUBiS (MySQL)	Throughput	6	R2	6.6

Table 5: Sample of iTuned's results on TPC-W and RUBiS dataset

on a 6GB database. RUBiS [1] is Web service benchmark that implements the core functionality of an auction site like eBay.

WORKLOAD TYPE	QUERY MIX	DEFUALT	RULE-BASED	SHC	OPTIMAL	ITUNED
W1-SF1	Q1	107	107	101	101	101
W2-SF1	Q9	310	282	231	95	95
W3-SF1	Q18	315	267	250	183	184
W3-SF10	Q18_10	158	138	140	131	131
W4-SF1	6Q18	1443	1505	55	51	52
W4-SF10	6Q18_10	578	167	148	128	133
W5-SF10	3Q7+3Q13_10	1167	1057	1237	1057	1173
W6-SF1	5Q1+5Q9	907	943	838	738	765
WORKLOAD TYPE	QUERY MIX	DEFUALT	RULE-BASED	SHC	OPTIMAL	ITUNED
W7-SF1	Q1Q18	338	158	98	78	84
W7-SF10	Q1Q18	1085	386	421	246	248
W7-N-SF10	Q1Q18	1085	386	295	246	283
W8-SF1	Q2Q3Q4Q5	406	423	366	208	209
W9-SF10	CM1	1451	1240	979	601	601
W10-SF10	CM1_1800	5910	5089	1674	1505	1583
W11-SF10	CM2	1152	1020	1052	971	1014

Figure 10: Comparison of tuning quality for single instance of workload. iTuned's tuning times shown in seconds
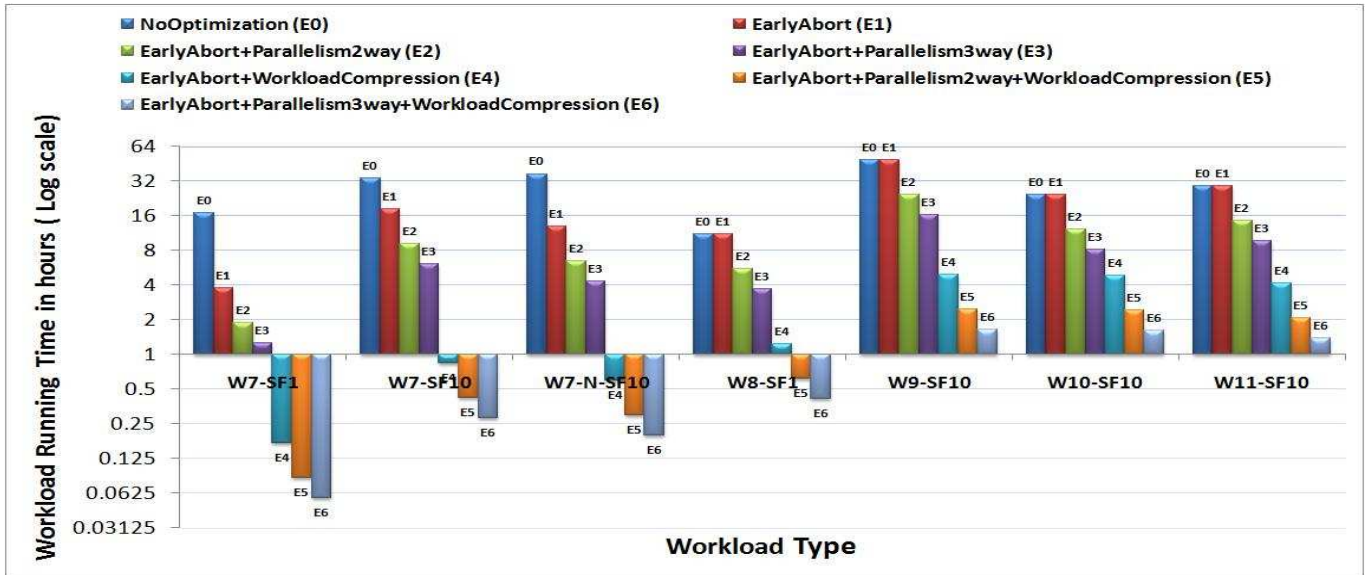

Figure 11: Comparison of iTuned's tuning times in the presence of various efficiency-oriented features

7.5 Sensitivity Analysis

This section evaluates two important features of iTuned: sensitivity analysis of database parameters and effects plots for visualization; see Section 6.1. We use both real workloads and complex synthetic response surfaces in our evaluation. We compare iTuned's performance against SARD [5] which is described in Section 3. Recall that, unlike iTuned, SARD is not an end-to-end tuning tool, and can be misled by nonmonotonic effects of parameters.

Our concerns about SARD were validated by a simple evaluation. We chose three popular and hard benchmark functions from the optimization literature: Griewank, Rastrigin, and Rosenbrock [22]. All three functions have a global optimum of 0. We used the functions to generate response surfaces with 20 parameters each.

Of these 20 parameters, 5 are important—i.e., they impact the shape of the surface significantly—while the remaining 15 are unimportant. On the Griewank and Rastrigin surfaces—which have significant nonmonotonic behavior—SARD completely failed to identify the unimportant parameters. As iTuned did experiments progressively, it never classified any important parameter as unimportant. By the time fifty experiments were done, iTuned was able to clearly separate the five important parameters from the unimportant ones.

Tables 6 gives end-to-end tuning results for three techniques: (i) SARD+AS, where SARD is used to identify the important parameters, and then Adaptive Sampling is started with the samples collected by SARD used for initialization; (ii) SHC (does not do sensitivity analysis), and (iii) iTuned. Note that lower numbers are better in all cases. iTuned clearly outperforms the alternatives.

Workload	Optimal	SARD+AS	SHC	iTuned
Griewank	0	28.6	28.7	2.0
Rastrigin	0	200.8	209.1	26.1
Rosenbrock	0	40.2	160.5	7.9
W2-SF1	95	240 (R29)	231 (R24)	95 (R1)
W3-SF1	11	43 (R20)	67 (R24)	12 (R4)
W6-SF1	390	450 (R63)	417 (R20)	403 (R5)
W8-SF1	208	208 (R1)	289 (R4)	208 (R1)

Table 6: Sensitivity analysis. For W2, W3, W6, W8, rank and performance of best setting (secs) are shown. Lower is better
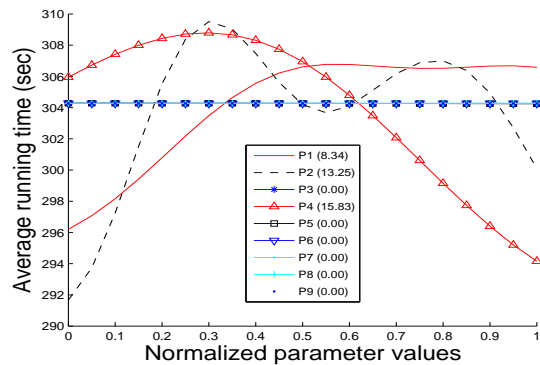
Figure 12: Effect plot for workload W4 (TPC-H Scale Factor 1)

Figure 13: Effect plot for workload W4 (TPC-H Scale Factor 10)

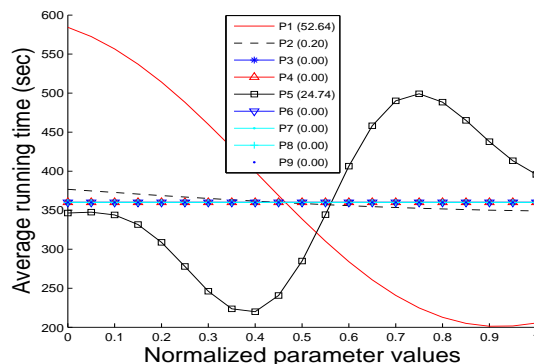**Figure 14: Effect plot for workload W5 (TPC-H SF=10)**

A very useful feature of iTuned is that it can provide intuitive visualizations of its current results. Figure 14 shows an effect plot (recall Section 6.1) generated by iTuned based on 10 experiments for the workload whose surface is shown in Figures 7 and 8. Figure 12 and 13 shows the effect plot for workload W4 for SF=1 and SF=10. The parameters P1-P9 correspond to the first nine PostgreSQL parameters listed in Section 7.1. Without knowing the actual response surface, a user can quickly grasp the main trends in parameter impact based on the effect plot. Note how the plot mirrors the trends in Figures 7 and 8. Effect plots of other workloads are in [6].

In summary, as few as twenty experiments chosen smartly by iTuned can produce a wealth of information in a reasonable amount of time to aid both naive users and expert DBAs in tuning database configuration parameters.

8. RELATED WORK

Databases have fairly mature tools for physical design tuning (e.g., index selection [4]). However, these tools do not address configuration parameter tuning. Furthermore, these tools depend on the cost models in the query optimizer so are limited in that these models do not capture the effects of many parameters.

Surprisingly, very little work has been done on tools for holistic tuning of the many configuration parameters in modern database systems. Most work in this area has either focused on specific classes of parameters (e.g., [16]) or on restricted subproblems of the overall parameter tuning problem (e.g., [5]). IBM DB2 provides an advisor for setting default values for a large number of parameters [10]. DB2's advisor does not generate response surfaces, instead it relies on built-in models of how various parameters affect performance [5]. As we show this paper, predetermined models may not be accurate in a given setting. SARD (discussed in Section 3) and [17] are also related to iTuned. SARD focuses on ranking parameters in order of impact, and is not an end-to-end tuning tool. Reference [17] proposed techniques to learn a probabilistic model using samples generated from gridding, which was then applied to tune four parameters in Berkeley DB. Gridding becomes very inefficient
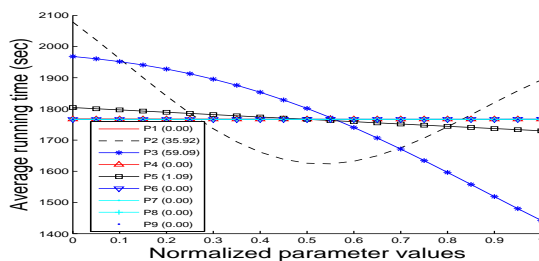
as the number of parameters increase. Section 7 also compared iTuned with a technique based on hill climbing (e.g., [22]) that has been applied to parameter tuning. None of the above techniques have an equivalent of iTuned's executor or the efficiency-oriented features from Section 6.

Techniques for tuning specific classes of parameters include solving analytical models [19], using simulations of database performance (e.g., in Oracle database), and control-theoretic approaches for online tuning [16]. These techniques are all based on predefined models of how changes in parameter settings affect performance. Reference [15] proposed techniques to tune the CPU and memory allocations to databases running inside virtual machines. However, the focus was not on planning experiments to learn the underlying response surfaces. All the above techniques can benefit from the Adaptive Sampling and executor ideas in iTuned.

Traditional database sampling deals with the problem of sampling from a large dataset, while our approach of Adaptive Sampling is about drawing samples from a response surface that is never materialized fully. Adaptive Sampling shares goals, but not techniques, with conventional database problems like online aggregation [8], acquisitional query processing [11], and sampling for statistics estimation [2]. For example, [2] gives a two-phase adaptive method in which the sample size required to reach a desired accuracy is decided based on a first phase of sampling. In contrast, Adaptive Sampling can adapt after each sample is brought in.

Oracle 11g introduced the SQL Performance Analyzer (SPA) to help DBAs measure the impact of database changes like upgrades, parameter changes, schema changes, and gathering optimizer statistics [23]. (Quoting from [23], "it is almost impossible to predict the impact of such changes on SQL performance before actually trying them.") SPA conducts experiments where SQL statements in the workload are executed with and without applying a change. However, Oracle 11g does not provide an experiment planner that can automatically handle complex tuning tasks like parameter tuning. Finally, experiments are used to collect data in many domains like chemical and mechanical engineering, so-

cial science, and computer simulation. While iTuned shares overall guiding principles with experiment planning in these domains, the requirements and algorithms differ.

9. CONCLUSION

We described iTuned, a tool that automates the task of recommending good settings for database configuration parameters. iTuned has three novel features: (i) Adaptive Sampling to proactively bring in appropriate data through planned experiments to find high-impact parameters and high-performance parameter settings, (ii) executor to supports online experiments in production database environments through a cycle-stealing paradigm that places near-zero overhead on the production workload; and (iii) portability across different database systems. We showed the effectiveness of iTuned through an extensive evaluation based on different types of workloads, database systems, and usage scenarios.

Cloud computing can make resources available at extremely cheap rates for experiments. In fact, for almost all the tuning tasks from Section 7.4, all the required experiments can be done on Amazon Web Services in a budget less than fifteen U.S. dollars! This cost includes the cost of using Amazon EC2 instances for the CPU and memory resources required by each experiment, and Amazon Elastic Block Storage for storing and accessing TPC-H data in a PostgreSQL database.

10. REFERENCES

[1] E. Cecchet, J. Marguerite, and W. Zwaenepoel. Performance and Scalability of EJB Applications. In *OOPSLA*, November 2002.

[2] S. Chaudhuri, G. Das, and U. Srivastava. Effective use of block-level sampling in statistics estimation. In *SIGMOD Conference*, 2004.

[3] S. Chaudhuri, A. K. Gupta, and V. R. Narasayya. Compressing sql workloads. In *SIGMOD Conference*, 2002.

[4] S. Chaudhuri and V. R. Narasayya. Autoadmin 'what-if' index analysis utility. In *SIGMOD Conference*, 1998.

[5] B. K. Debnath, D. J. Lilja, and M. F. Mokbel. SARD: A statistical approach for ranking database tuning parameters. In *SMDB*, 2008.

[6] S. Duan, V. Thummala, and S. Babu. Tuning database configuration parameters with ituned. Technical report, Mar. 2009. Available at http://www.cs.duke.edu/˜shivnath/ituned.

[7] R. G. Freeman and A. Nanda. *Oracle Database 11g New Features*. McGraw-Hill Osborne Media, 2007.

[8] J. Hellerstein, P. Haas, and H. Wang. Online aggregation. In *Proc. of the 1997 ACM SIGMOD Intl. Conf. on Management of Data*, May 1997.

[9] C. R. Hicks and K. V. Turner. *Fundamental Concepts in the Design of Experiments*. Oxford University Press, 1999.

[10] S. Kwan, S. Lightstone, et al. Automatic Configuration of IBM DB2 Universal Database. In *IBM Perf. Technical Report*, Jan. 2002.

[11] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. The design of an acquisitional query processor for sensor networks. In *Proc. of the 2003 ACM SIGMOD Intl. Conf. on Management of Data*, pages 491 – 502, June 2003.

[12] Postgresql performance optimization. wiki.postgresql.org/wiki/Performance_Optimization.

[13] T. J. Santner, B. J. Williams, and W. Notz. *The Design and Analysis of Computer Experiments*. Springer, first edition, July 2003.

[14] *Dtrace, ZFS, and Zones in Solaris*. www.sun.com/software/solaris.

[15] A. A. Soror, U. F. Minhas, A. Aboulnaga, K. Salem, P. Kokosielis, and S. Kamath. Automatic virtual machine configuration for database workloads. In *SIGMOD Conference*, 2008.

[16] A. J. Storm, C. Garcia-Arellano, S. Lightstone, Y. Diao, and M. Surendra. Adaptive self-tuning memory in db2. In *VLDB*, 2006.

[17] D. G. Sullivan, M. I. Seltzer, and A. Pfeffer. Using probabilistic reasoning to automate software tuning. In *SIGMETRICS*, 2004.

[18] TPC-H and TPC-W Benchmarks from the Transaction Processing Council. http://www.tpc.org.

[19] D. N. Tran, P. C. Huynh, Y. C. Tay, and A. K. H. Tung. A new approach to dynamic self-tuning of database buffers. *ACM Transactions on Storage*, 4(1), 2008.

[20] G. Weikum, A. Moenkeberg, C. Hasse, and P. Zabback. Self-tuning database technology and information services: from wishful thinking to viable engineering. In *VLDB*, 2002.

[21] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, second edition, June 2005.

[22] B. Xi, Z. Liu, M. Raghavachari, C. H. Xia, and L. Zhang. A smart hill-climbing algorithm for application server configuration. In *WWW*, pages 287–296, 2004.

[23] K. Yagoub, P. Belknap, et al. Oracle's sql performance analyzer. *IEEE Data Engineering Bulletin*, 31(1), 2008.