

# Phase-Independent Rhythmic Analysis of Genome-Wide Expression Patterns

CHRISTOPHER JAMES LANGMEAD,<sup>1</sup> ANTHONY K. YAN,<sup>1</sup>  
C. ROBERTSON MCCLUNG,<sup>2</sup> and BRUCE RANDALL DONALD<sup>1,2,3,4</sup>

## ABSTRACT

We introduce a model-based analysis technique for extracting and characterizing rhythmic expression profiles from genome-wide DNA microarray hybridization data. These patterns are clues to discovering rhythmic genes implicated in cell-cycle, circadian, or other biological processes. The algorithm, implemented in a program called RAGE (Rhythmic Analysis of Gene Expression), decouples the problems of estimating a pattern's wavelength and phase. Our algorithm is linear-time in frequency and phase resolution, an improvement over previous quadratic-time approaches. Unlike previous approaches, RAGE uses a true distance metric for measuring expression profile similarity, based on the Hausdorff distance. This results in better clustering of expression profiles for rhythmic analysis. The confidence of each frequency estimate is computed using Z-scores. We demonstrate that RAGE is superior to other techniques on synthetic and actual DNA microarray hybridization data. We also show how to replace the discretized phase search in our method with an exact (combinatorially precise) phase search, resulting in a faster algorithm with no complexity dependence on phase resolution.

**Key words:** functional genomics, DNA microarray analysis, mRNA expression time-series, rhythmic genes, circadian genes, Hausdorff distance, autocorrelation, cell-cycle, expression profile similarity, computational genomics.

## 1. INTRODUCTION

THE EXPRESSION PATTERNS OF MANY GENES associated with circannual (yearly), circadian (daily), cell-cycle and other periodic biological processes are known to be rhythmic. Conversely, the expression profiles of genes associated with aperiodic biological processes (e.g., tissue repair) are not rhythmic. The functional significance of previously uncharacterized genes, therefore, may be inferred if they exhibit rhythmic patterns of expression synchronized to some ongoing biological process.

DNA microarray experiments are an effective tool for identifying rhythmic genes when a time-series of expression levels are collected. Unlike northern blots and real-time PCR, which study one gene at a time,

---

<sup>1</sup>Dartmouth Computer Science Department, Hanover, NH 03755.

<sup>2</sup>Dartmouth Department of Biological Sciences, Hanover, NH 03755.

<sup>3</sup>Dartmouth Chemistry Department, Hanover, NH 03755.

<sup>4</sup>Dartmouth Center for Structural Biology and Computational Chemistry, Hanover, NH 03755.

DNA microarray hybridization time-series experiments can reveal the expression patterns of entire genomes. Chronobiologists are therefore able to assign putative functional properties to large numbers of genes based on the results of a single experiment. However, the large volume of data generated by hybridization experiments makes manual inspection of individual expression profiles impractical. Separating the subset of genes whose expression profiles are rhythmic from the thousands or tens of thousands that are not requires computer assistance. Ideally, the algorithms for providing such assistance should be efficient and have well-understood performance guarantees.

We have designed and implemented an algorithm to identify and characterize the properties of rhythmic genes from DNA microarray hybridization time-series data. Our approach specifically addresses issues of computational complexity, statistical significance, and morphological similarity.

The identification of rhythmic genes may be viewed as a pattern-recognition problem—the goal is to identify sinusoidal RNA expression patterns in massively parallel gene expression data. Each expression profile may be viewed as a scalar function of time. A stored set of “model” functions may be compared with an unknown function (expression profile) that has been obtained by experiment. The model may be either (a) an “ideal” synthetic sinusoid or (b) another gene expression profile. In (a), one tries to fit a family of different ideal sinusoids to the data, to determine if the data is periodic, and if so, to determine the best fit. The sinusoids may differ in frequency and phase and may be damped. In (b), the model may be a known rhythmic gene, in which case one attempts to find genes with similar profiles.

In either case, the difference between each model shape and unknown shape is computed, and the model that is closest to the unknown shape is reported as the best match. Strong arguments from the machine vision and pattern recognition literature argue that for such applications, the function used to measure the difference between model and data should be a metric (Arkin *et al.*, 1991; Mumford, 1987). This means that for a class of expression profiles the difference function  $d$  should obey the following properties, for any three profiles  $X$ ,  $Y$ , and  $Z$ :

$$d(X, Y) \geq 0 \text{ for all } X \text{ and } Y. \quad (1)$$

$$d(X, Y) = 0 \quad \text{if and only if } X = Y \quad \text{(Identity).} \quad (2)$$

$$d(X, Y) = d(Y, X) \quad \text{for all } X \text{ and } Y \quad \text{(Symmetry).} \quad (3)$$

$$d(X, Y) + d(Y, Z) \geq d(X, Z) \quad \text{for all } X, Y \text{ and } Z \quad \text{(Triangle Inequality).} \quad (4)$$

As argued in Donald *et al.* (1992), the triangle inequality is of particular importance, because it guarantees that if several model expression profiles are similar to a given data expression profile, then these profiles also must be similar to one another. Thus, for example, it is not possible for two highly dissimilar model profiles to be similar to the data profile for the same gene. Current microarray analysis methods generally compare profiles using functions that are *not* metrics and thus may report that several dissimilar models match the same data, which is highly counterintuitive. In addition to obeying metric properties, an expression profile comparison method should also be easy to compute in order for it to be of practical use. The method we describe can be computed efficiently both in theory and in practice. Previous algorithms that use hierarchical clustering (e.g., Eisen *et al.*, 1998) run in time  $O(n^2l)$ , where  $n$  is the number of genes represented in the microarray data and  $l$  is the number of time-series points. Other algorithms (e.g., Harmer *et al.*, 2000) that estimate both the frequency and phase of gene expression profiles using pattern recognition run in time  $O(nmp \log l)$ , where  $m$  is the frequency resolution, and  $p$  is the phase resolution. These methods can take up to a week of wall-clock CPU time to analyze data from a single gene chip experiment and suffer from the use of nonmetric similarity measurements. We describe an algorithm that runs in time  $O(n(m+p)f(l))$ , where  $f(l)$  is the time to compute the Hausdorff distance ( $f(l) = O(l^2)$  deterministic and  $O(l)$  probabilistic time). Next, we replace our discretized phase search with an exact

TABLE 1. CDC15, CDC28, FIBROBLAST AND *Drosophila* DATASETS<sup>a</sup>

<i>Experiment</i>	<i>Organism</i>	$\Delta t$ (minutes)	# samples	# periods	# genes
CDC15 <sup>b</sup>	<i>S. cerevisiae</i>	10/20	24	3.2	6,178
CDC28 <sup>c</sup>	<i>S. cerevisiae</i>	10	17	1.8	6,220
Fibroblast <sup>d</sup>	<i>H. sapiens</i>	15/30/60/120/240	12	1.5	9,712
<i>Drosophila</i> <sup>e</sup>	<i>D. melanogaster</i>	240	6	0.83	14,010

<sup>a</sup>The column labeled “ $\Delta t$ ” gives the time period between successive time points. If there is more than one  $\Delta t$  listed, then the data was nonlinearly sampled using a combination of the specified times. The column labeled “# periods” indicates the number of cell-cycle periods that fit within the duration of the sample interval.

<sup>b</sup>Spellman *et al.*, 1998.

<sup>c</sup>Cho *et al.*, 1998.

<sup>d</sup>Iyer *et al.*, 1999.

<sup>e</sup>McDonald and Rosbash, 2001.

(combinatorially precise) phase search. This eliminates the factor of  $p$  entirely, resulting in an overall complexity of  $O(nmf(l) + nl^3\alpha(l)\log l)$ , where  $\alpha$  is the extremely slow-growing inverse of Ackerman’s function. In all cases,  $l$  may be treated as a small constant, since in today’s technology,  $l$  is never more than a small constant  $l_{max} \ll n$  (for example, typically,  $l \leq 24$ , and  $n \approx 6500$ —see Table 1). This simplification obtains a complexity bound of  $O(nmp)$  for previous algorithms versus  $O(n(m+p))$  and  $O(nm)$  for ours. Our algorithm processes 32,000 12-sample mRNA time-series in two hours on a single processor Pentium-class workstation.

Our chief contributions are as follows:

1. The use of autocorrelation to define a phase-independent search over frequency- and phase-space. This allows us to perform two linear-time searches, one in frequency- and one in phase-space, as opposed to a quadratic-time search over frequency-phase space.
2. The use of the undirected Hausdorff (UH) distance to compare similarity of expression profiles. UH satisfies the axioms of a distance metric on the space of expression profiles, unlike previous measures, resulting in a robust and rigorous basis for clustering.
3. Testing our methods on publicly available gene expression data and a comparison of the results to previous analyses.
4. The application of our methods to find circadian genes in a microarray data set that has previously been searched only for cell-cycle genes.

### 1.1. Organization of paper

We begin, in Section 2, with a review of the relevant biology and a summary of three publicly available DNA microarray hybridization time-series datasets. Section 3 categorizes existing techniques for extracting rhythmic profiles from microarray data, including a discussion of their limitations and computational complexity. In Section 4, we detail our method and analyze its computational complexity. Section 5 presents the results of the application of RAGE to simulated and real biological data. Finally, Section 6 discusses these results.

## 2. BACKGROUND

There are many examples of DNA microarray time-series experiments in the literature (e.g., Cho *et al.*, 1998; Claridge-Chang *et al.*, 2001; Grundschober *et al.*, 2001; McDonald and Rosbash, 2001; Schaffer *et al.*, 2001; Harmer *et al.*, 2000; Spellman *et al.*, 1998; Iyer *et al.*, 1999; Yagita *et al.*, 2001). In many of these experiments, the primarily goal was to identify genes whose expression patterns were periodic over the length of the experiment. For example, cell-cycle regulated (e.g., Cho *et al.*, 1998; Spellman *et al.*, 1998) and circadian (e.g., Harmer *et al.*, 2000; Yagita *et al.*, 2001) genes have been identified from their expression profiles in hybridization experiments.

Several research labs have made their raw data available to the public<sup>1</sup> facilitating the development of improved techniques. The Davis lab at Stanford has released the yeast data presented by Cho *et al.* (1998) on the CDC28 mutant of yeast. The Botstein lab has released the data from their yeast experiment on the CDC15 mutant of yeast presented by Spellman *et al.* (1998). The Brown lab at Stanford has released the human fibroblast data presented by Iyer *et al.* (1999). The Rosbash lab at Brandeis has recently released the data from their circadian experiment on *Drosophila* presented by McDonald and Rosbash (2001). In this section, we briefly summarize the biological background relevant to these datasets.

### 2.1. Yeast, fibroblast, and drosophila datasets

The CDC15 and CDC28 experiments were designed to identify cell-cycle-regulated genes in yeast (*Saccharomyces cerevisiae*). The eukaryotic cell-cycle is the 4-stage process by which a single cell replicates into two daughter cells. The four stages, named G1, S, G2, and M, have distinct roles. The chromosomes are prepared for replication in G1. The DNA and centrioles are replicated in S. The cell is prepared for separation in G2. Finally, the cell divides in M (mitosis), and the process begins again. This process takes about 90 minutes in yeast and 16 hours in fibroblasts. Thus, the authors of the CDC15, CDC28, and fibroblast experiments were looking for uncharacterized genes whose expression profiles were periodic with those frequencies.

The fibroblast experiment was also designed to identify cell-cycle-regulated genes but in human fibroblast cells, instead of yeast. Unlike yeast, the cell-cycle for human fibroblasts is approximately 16 hours.

The *Drosophila* data set was generated as part of an experiment to identify circadian and circadian-regulated genes. Circadian rhythms are biological processes that are synchronized to the diurnal cycling of light and dark. The goal of circadian microarray experiments is to find the genes associated with the circadian clock. Table 1 details the content of the CDC15, CDC28, fibroblast, and *Drosophila* datasets.

## 3. PRIOR WORK

A variety of techniques have been developed to extract the rhythmic genes from these datasets. The various techniques fall into two categories: *spectral* and *cluster-based* analyses. In this section, we discuss each type, citing specific examples.

### 3.1. Spectral techniques

The Fourier transform is a standard tool for detecting periodic fluctuations in discretized signals. Spellman *et al.* (1998) used the Fourier transform as one component of a hybrid technique for determining the frequency and phase of gene expression profiles in the CDC15 and CDC28 datasets. The limitations of the Fourier transform are well understood. The range of detectable frequencies within a signal and the resolution to which they can be resolved are particularly relevant to DNA microarray data. The frequency resolution obtainable on short time series, such as those generated in typical microarray experiments, is often not adequate for resolving wavelengths of interest. The interested reader is directed to Appendix A.1 for a longer discussion of these limitations.

The size limitations on the datasets are typically not biological but rather financial. For example, for *Arabidopsis* studies, each Affymetrix chip costs \$400. That cost will increase to \$500 when the full genome chip for *Arabidopsis* becomes available. There is also a \$400 per chip processing fee. Thus, a 24-data-point time-series with a replication factor of 3 costs \$57,600. When these costs drop and the number of points per experiment rises, the Fourier transform will become a more effective tool. Until that time, nonspectral techniques will likely dominate. We discuss these techniques in the next section.

Skiena and coworkers (Filkov *et al.*, 2001) give an algorithm for estimating frequencies of periodic genes, using the correlation coefficient and an unusual time-division strategy. While complexity bounds are not given, the algorithm appears to be efficient. The phase search of Filkov *et al.* (2001) is unique: they look for *aggregate* shifts. They find, for example, that the CDC15 dataset is phase-shifted relative to the CDC28 set. Our method estimates the phase offset of individual genes relative to the start of the experiment.

---

<sup>1</sup>Stanford Microarray Database: <http://genome-www4.stanford.edu/MicroArray/SMD/>

### 3.2. Cluster-based analysis


A clustering algorithm takes as input a set of items and a method for comparing the similarity between pairs of items. The outputs of the algorithm are subsets/clusters of the input set where the average similarity between pairs within a cluster is higher than the average similarity between items from different clusters. The second class of time-series analysis techniques clusters gene expression profiles to model profiles. Unknown genes are attributed the properties of the model to which they are most similar. The two most important distinctions among clustering methods are 1) how the models are generated and 2) which similarity measurements are used. The choice of models and similarity measurements affects both the complexity and accuracy of the resulting algorithm.

**3.2.1. Model generation.** There are two primary means for generating models for a clustering algorithm. In the first approach the models are generated from the data itself. The clustering technique developed by Eisen *et al.* (1998), which was used for the original analysis of the CDC15, CDC28, and fibroblast datasets, is an example of a data-generated clustering technique. It is often the case that among the microarray data there are a number of genes of known function. Such genes can be used as models. Data-derived models have a potential advantage in that they implicitly include noise models. There are nontrivial variations found among microarrays and in the steps leading up to hybridization. Consequently, the actual expression profiles of rhythmic genes sometimes deviate from an ideal model. When these variations are systematic across all genes, the use of data-derived models is especially useful. Whether or not such noise is in fact systematic is a valid concern.

The disadvantage of data-derived methods is that one has no control over the models. This places a particular burden on the similarity measurement used to compare profiles. The similarity measurement must be somewhat forgiving of small variations between otherwise similar shapes. Consider the following example (Figs. 1A and 1B), comparing the shapes of two sinusoids differing in phase by 90 degrees using the correlation coefficient:

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\left(\sum_{i=1}^n (x_i - \bar{x})^2\right) \left(\sum_{i=1}^n (y_i - \bar{y})^2\right)}} \quad (5)$$

where  $x$  and  $y$  are the signals being compared and  $\bar{x}$  and  $\bar{y}$  are the respective means of those signals.



**FIG. 1.** The sinusoid in A is 90 degrees out of phase with the sinusoid in B. The correlation coefficient (Equation 5) of A and B is 0, or no correlation. The signal in C is the autocorrelation of A. The signal in D is the autocorrelation of B. Note they are nearly identical due to the phase-independent nature of the autocorrelation. Furthermore, unlike the Fourier transform, the autocorrelation is not affected by nonlinearly sampled data. Therefore, it is more suitable for a larger class of microarray experiments.

The correlation coefficient is used by quite a few clustering algorithms (e.g., Eisen *et al.*, 1998; Filkov *et al.*, 2001). The correlation coefficient of the two signals in Figs. 1A and 1B is 0, indicating that they are not similar. Clearly, the two shapes have a lot in common. Furthermore, the correlation coefficient violates three of the four criteria for being a metric outlined in Section 1. Perhaps the most important violation is that the correlation coefficient does not satisfy the triangle inequality. A related statistical measurement, the cross-correlation, computes correlation coefficients over all relative shifts of the two input signals. Consequently, the cross-correlation is also a nonmetric.

Another approach to model-based clustering is implemented in the program CorrCos (Harmer *et al.*, 2000). CorrCos clusters gene expression data to *synthetic* sinusoidal models using cross correlations. The advantage of synthetic models is that the program has complete control over the models. CorrCos generates one thousand sinusoidal models of differing frequencies. For each frequency model, 101 phase variations are also generated. Each gene expression profile is then compared, using the cross-correlation, to each of the 101,000 synthetic models. The frequency and phase of the model most closely matching the expression profile are assigned to that gene.

The time complexity of CorrCos is  $O(nmp l \log l)$ , where  $n$  is the number of genes,  $m$  is the number of frequency models,  $p$  is the number of phase variants generated for each frequency model, and  $l$  is the length of the time series. The time is  $O(l \log l)$  to compute a cross correlation. Hence, this brute-force search over frequency and phase grows quadratically with the frequency and phase resolution ( $m$  and  $p$ ) of the search. In practice, CorrCos (compiled Fortran code) takes about one week to run on microarray data of size 32,000 genes. In contrast, our method only takes about two hours (interpreted MATLAB code) on the same data set.

In summary, there are a number of problems with the existing approaches for detecting and characterizing rhythmic genes in microarray time-series data. Spectral methods are not appropriate because the typical microarray experiment generates relatively short time-series. Consequently, cluster-based techniques are required. Of the cluster-based approaches, data-derived models are problematic because they lack flexibility and, when using a phase-sensitive similarity measurement such as the correlation coefficient, can lead to nonintuitive results. Synthetic model-based clustering addresses some of these problems at the expense of computational complexity. Any generate-and-test-based technique using a phase-dependent similarity measurement will have similar properties.

In contrast, we introduce the *autocorrelation* to render our search algorithm phase-independent (Section 4.1). The (hypothetical) use of a *nonmetric* similarity measurement (such as the correlation coefficient) on an autocorrelated phase-independent representation would still suffer from the drawbacks presented in Section 1. Therefore, we employ the *Hausdorff distance* (a true metric) on the autocorrelated signals in order to obtain superior matching and clustering performance (Section 4.2). By attacking the problem of time-series analysis using these new methods, we hope to strengthen the computational armamentarium of the chronobiologist.

## 4. RAGE (RHYTHMIC ANALYSIS OF GENE EXPRESSION)

We have developed an algorithm named RAGE to address the problems above. Like CorrCos, RAGE is a synthetic model-based clustering technique. However, RAGE is more efficient than CorrCos, running in time  $O(n(m+p)f(l))$ , where  $f(l)$  is the time to compute the undirected Hausdorff distance between an  $l$ -point model function and an  $l$ -point gene expression profile (Section 4.2). Hence, the time complexity of RAGE grows linearly with increased frequency ( $m$ ) and phase ( $p$ ) resolution. RAGE achieves the better bound by using phase-independent transformations of the data and models. RAGE also utilizes a true mathematical metric, the undirected Hausdorff distance metric, for computing expression profile similarities. A summary of the RAGE algorithm is given in Fig. 2.

### 4.1. Phase-independence

The computational complexity of the CorrCos algorithm stems from a brute-force search over the space of frequency and phase. We claim it is not necessary to search this entire space. It is possible to decouple the phase and frequency searches by transforming both the raw data and the models into a phase-

```

Let  $G \leftarrow$  Gene Data
Let  $A \leftarrow$  AutoCorr( $G$ )
for  $\omega = \omega_{lo}$  to  $\omega_{hi}$ 
     $M_\omega \leftarrow$  AutoCorr(GenerateModel( $\omega$ , 0))
for  $i = 1$  to num-genes
     $\Omega_i \leftarrow$  ClusterFreq( $A_i, M$ )
for  $i = 1$  to num-genes
    for  $\phi = \phi_{lo}$  to  $\phi_{hi}$ 
         $P_\phi \leftarrow$  GenerateModel( $\Omega_i, \phi$ )
         $\Phi_i \leftarrow$  ClusterPhase( $G_i, \Omega_i, P$ )
for  $i = 1$  to num-genes
     $Z_i \leftarrow$  Z-Score( $G_i, \Omega_i, G$ )

```

**FIG. 2.** RAGE algorithm. The function AutoCorr takes a set of gene expression profiles as input and returns a set of autocorrelated gene expression profiles. GenerateModel takes a frequency ( $\omega$ ) and phase ( $\phi$ ) as input and returns a sinusoid of the specified frequency and phase. The length of the model is assumed to be the same size as the length of the expression profiles in  $G$ . The set  $M$  is the set of autocorrelated (i.e., phase-independent) models, and  $M_\omega$  is the model of frequency  $\omega$ . ClusterFreq takes as an autocorrelated expression profile and a set of models and returns the frequency of the model that has the lowest Hausdorff distance to the input expression profile. The set  $\Omega$  is the set of frequency assignments, as determined by ClusterFreq, and  $\Omega_i$  is the frequency assigned for each gene  $G_i$  in  $G$ . The set  $P$  is the set of phase-sensitive models, and  $P_\phi$  is the model of phase  $\phi$  and frequency  $\Omega_i$ . ClusterPhase takes as input a single gene expression profile, the frequency that has been assigned to that gene, and a set of phase-sensitive models and returns the phase of the model that has the lowest Hausdorff distance to the input expression profile. Note that  $P$ , the set of models for a given frequency ( $\Omega_i$ ) can be cached for better performance. The set  $\Phi$  is the set of phase-assignments, as determined by ClusterPhase, for each of the genes in  $G$ , and  $\Phi_i$  is the phase for gene  $G_i$ . Z-Score takes as input an expression profile, the frequency assigned to that gene, and  $G$ . Z-score randomly chooses a constant-sized set  $R$  of genes from  $G$ . The distribution of Hausdorff distances between  $G_i$  and the expression profiles in  $R$  is used to compute the Z-score of the Hausdorff distance between the model of frequency  $\Omega_i$  and  $G_i$ .

independent form. The autocorrelation  $\Psi_x$  of a signal  $x$  is a frequency-sensitive but phase-independent representation of  $x$ :

$$\Psi_x(l) = \int_{-\infty}^{+\infty} x(t)x(t-l)dt \quad (6)$$

where  $x$  is the function,  $x(t)$  is the function evaluated at  $t$ , and  $x(t-l)$  is a translated version the function. Figures 1C and 1D show the autocorrelation of the two signals from Figs. 1A and 1B, respectively. Note that the two autocorrelations are nearly identical. The interested reader is directed to Appendix A.2 for more information on the autocorrelation. In our implementation, each signal is zero-meaned and normalized prior to computing the autocorrelation.

RAGE estimates the frequency of expression profiles using the autocorrelations of both the model and gene data. After the frequency has been estimated, phase-variations of the winning frequency model are generated, and a second clustering is performed using these phase models and the original, nonautocorrelated data. Note that the phase search is conducted over a single frequency, and not all frequencies (unlike CorrCos).

#### 4.2. Undirected Hausdorff distance

RAGE uses the undirected Hausdorff distance metric (UH) instead of the correlation coefficient to cluster expression profiles. The undirected Hausdorff distance ( $H$ , Equation (7)) calls the directed Hausdorff

distance ( $h$ , Equation (8)) as a subroutine (Donald *et al.*, 1992; Huttenlocher and Kedem, 1990):

$$H(A, B) = \max(h(A, B), h(B, A)) \quad (7)$$

$$h(A, B) = \max_{a \in A} \min_{b \in B} (\rho(a, b)), \quad (8)$$

where  $A$  and  $B$  are two point sets in the plane and  $\rho$  is a distance metric. In our implementation,  $\rho$  is the Euclidean distance in the plane between two points. In fact,  $\rho$  can be any metric. The UH distance essentially measures the maximum outlier between two sets of points. It is commonly used in the computer vision community for object recognition tasks. We treat the model and data expression profiles as scalar functions; the Hausdorff distance is used to compare the similarity of their graphs, treated as point sets, in the plane. The deterministic time-complexity  $f(l)$  for computing the Hausdorff distance is  $O(l^2)$ . UH distance can be computed on curves (such as our functions) in expected time  $O(l)$  (Belogay *et al.*, 1997). One potential limitation of the UH is that it can be sensitive to outliers. However, this can be easily solved using *quantiles* (Donald *et al.*, 1992). The basic strategy is to compute the UH on a subset of the complete set of distances measured using  $\rho$ . For example,  $h$  might be computed using only the 75% quantile of maximum distances. However, by using quantiles, some metric properties of the UH are lost. We did not use quantiles in our implementation.

#### 4.3. Statistical significance

Any cluster-based approach computes the similarity between each gene and model. The gene is assigned to the model with the highest similarity. In algorithms such as CorrCos and RAGE, all of the models are rhythmic. Thus, aperiodic genes are initially clustered to rhythmic models. A postprocessing step is required to filter such genes from each cluster. There are number of ways to approach this problem. RAGE employs a statistical approach.

The  $Z$ -score ( $z = \frac{x-\mu}{\sigma}$  where  $x$  is a member of the distribution,  $\mu$  is the mean of the distribution, and  $\sigma$  is the standard deviation) is a common statistical technique for estimating the significance of a given score. Given a normal distribution of similarity scores, the  $Z$ -score of any individual similarity measurement is the number of standard deviations that score lies from the mean. The  $Z$ -score can in turn be transformed into a probability ( $P(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}$ ). Hence, the confidence of a given estimate is  $1 - P(z)$ .

For each gene, RAGE precomputes a distribution of UH distances between the gene and a constant number of randomly selected genes from the data set.<sup>2</sup> The confidence associated with the UH distance of the winning frequency model is computed using this distribution. One-sided probabilities are used so that only statistically small UH distances are considered. The biologist may then sort the genes by these confidence scores.

#### 4.4. Algorithmic complexity

The complexity analysis of RAGE is as follows. The conversion of the raw gene data to its autocorrelated form can be computed in time  $O(nl \log l)$  where  $n$  is the number of genes and  $l$  is the length of the time-series for each gene. The generation of the set of models,  $M$ , is computed in time  $O(ml \log l)$  where  $m$  is the number of models (determined by the desired frequency resolution,  $m$ ). The clustering of each gene to its nearest model is done by computing the Hausdorff distance between each of the  $n$  genes and each of the  $m$  models. This can be done in time  $O(nmf(l))$ , where  $f(l)$  is the time to compute the Hausdorff distance (Section 4.2). The creation of the set  $P$  of phase models takes  $O(npl)$  time where  $p$  is the number of phase models (determined by the desired phase resolution,  $p$ ). The clustering of each (frequency assigned) gene to the nearest phase model takes  $O(npf(l))$ . Finally, computing the  $Z$ -scores takes time  $O(cnf(l))$  where  $c$  is a constant corresponding to the number of genes randomly selected to compute the distribution required for calculating  $Z$ -scores. Note that  $c$  is determined by the need for statistical significance. It does not grow as the size of  $n$  increases. This gives an overall complexity of  $O(nl \log l + ml \log l + nmf(l) + npl + npf(l) + cnf(l))$ . Since  $l$  is  $O(1)$  in today's experiments, the expression

<sup>2</sup>In our experiments, 500 other gene expression profiles were randomly selected.



reduces to  $O(n(m+p))$ , asymptotically faster than CorrCos' complexity of  $O(nmp)$ . In the future, if gene-chip experiments become cheaper,  $l$  will not be a constant, although we can still expect  $l \ll n$ , and  $l \ll m$  is likely for the foreseeable future. Hence, our algorithm would run in time  $O(n(m+p)l^2)$  (deterministic) or  $O(n(m+p)l)$  (probabilistic) time, as opposed to  $O(nmpl \log l)$  for CorrCos.

#### 4.5. Exact phase search and improved algorithm

We can improve on this algorithm as follows. For each of the  $n$  genes, the phase search as presented above takes time  $O(pf(l))$  per gene. We can replace this discretized phase search with an exact (combinatorially precise) phase search, with running time  $O(l^3\alpha(l) \log l)$  per gene, where  $\alpha$  is the inverse of Ackerman's function. This eliminates the factor of  $p$  entirely, resulting in an overall complexity of  $O(nmf(l) + nl^3\alpha(l) \log l)$ . We achieve this as follows. Given two point sets  $A$  and  $B$  in the plane, Donald *et al.* (1992) give an algorithm for computing

$$\operatorname{argmin}_t H(A_t, B), \quad (9)$$

where  $A_t$  is a rigid translation of set  $A$  by vector  $t \in \mathbb{R}^2$ . If  $A$  and  $B$  have  $p$  and  $q$  points, respectively, the algorithm runs in time  $O(pq(p+q)\alpha(pq) \log(pq))$ .

**Definition 1 (Voronoi Surface (Donald *et al.*, 1992)).** *Given a set  $S$  of points in  $\mathbb{R}^2$  and a metric  $\rho$  on  $\mathbb{R}^2$ , consider the function  $d(x) = \min_{s \in S} \rho(x, s)$ . The Voronoi surface of  $S$  is the graph of  $d$ ,  $\{(x, d(x)) \mid x \in \mathbb{R}^2\} \subset \mathbb{R}^3$ .*

The Voronoi surface is a two-dimensional surface in a three-dimensional space. Intuitively, the Voronoi surface looks like an "irregular egg carton." Following Donald *et al.* (1992), to determine the translation  $t$  that minimizes Equation (9), we must identify the value of  $t$  that minimizes the upper envelope  $U$  of all the Voronoi surfaces defined by the sets  $S_i = a_i \ominus B$  and  $S'_j = A \ominus b_j$ , for all  $a_i \in A$  and all  $b_j \in B$ . Here,  $\ominus$  denotes the *Minkowski difference*, so that  $S_i = \{a_i - b \mid b \in B\}$  and  $S'_j = \{a - b_j \mid a \in A\}$ . The Voronoi surface has  $O(pq(p+q)\alpha(pq))$  local minima. Hence, Claim 2 follows.

**Claim 2 (Donald *et al.*, 1992).** *The minimum undirected Hausdorff distance under translation of two sets of points in the plane (and the translation that achieves this minimum) can be computed in time  $O(pq(p+q)\alpha(pq) \log(pq))$ .*

We use this algorithm as follows:  $B$  is the expression profile for a gene  $G_i$ ,  $A = \Omega_i$  is a frequency model, and  $A_t$  is the frequency model translated by  $t$ . We restrict  $t$  to be a phase shift (pure one-dimensional  $x$ -translation), as follows. The algorithm finds the minimum of the upper envelope  $U$ . We are interested only in a one-dimensional translation (the phase) along the  $x$ -axis. We can use the algorithm of Donald *et al.* (1992), by intersecting the upper envelope  $U$  of the Voronoi Surfaces, with the plane  $\pi_y^{-1}(0)$ , where  $\pi_y : \mathbb{R}^3 \rightarrow \mathbb{R}^1$ ,  $\pi_y(x, y, z) = y$ , and then finding minima in  $U \cap \pi_y^{-1}(0)$ . It is also possible to construct the Voronoi surface of a collection of one-dimensional distance functions directly, as a 1D surface in a 2D space.

## 5. RESULTS

RAGE has been applied to both synthetic and real microarray data. In Section 4, we demonstrated that RAGE is computationally more efficient than CorrCos due to the use of a phase-independent search. Our second claim is that the Hausdorff distance metric is superior to the correlation coefficient as a distance measure. To assess the benefit of using the Hausdorff distance metric, we implemented two versions of RAGE. The first one is as described in this paper. The second implementation (RAGE-CC) is exactly the same as RAGE except that it uses the correlation coefficient to compare function shapes instead of the Hausdorff distance metric.

TABLE 2. SYNTHETIC DATA<sup>a</sup>

(A)				(B)		
	<i>Models</i>	<i>Number</i>	<i>% noise</i>		<i>RAGE</i>	<i>RAGE-CC</i>
	Random	5,000	N/A	Number	1,498	2,199
	10 ± 2-hr periodic	500	5%	Mean $\Delta$ frequency (hours)	2.1	8.5
	15 ± 2-hr periodic	500	5%	Mean $\Delta$ phase (hours)	2.2	15.7
	24 ± 2-hr periodic	500	5%	False-Positives	6%	33%
				False-Negatives	7%	1%

<sup>a</sup>(A) Synthetic gene expression profiles. A set of 6,500 signals were generated; 5,000 of these were random signals and the remaining 1,500 were sinusoids. The sinusoids were grouped into three categories: 10-, 15-, and 24-hr periodic of 500 genes each. The actual frequency of each sinusoid was randomly chosen at  $\pm 2$  hrs from the category. All genes were assigned random phases. (B) Simulation results. After running RAGE and RAGE-CC on the simulated gene data, the accuracy of the results were evaluated. Given a sorted list of the confidence scores assigned to each gene, the top 50% of the nonzero scorers were investigated. The number of genes in that set is reported in the first row. Since the actual phase and frequency of the simulated genes were known, it is possible to compare the estimated frequency and phases against their true values. The second and third rows give the mean deviation of the estimated frequency and phase from their true values. False-positives indicates what percentage of the identified genes were from the set of random genes. False-negatives indicates the percentage of the 1,500 periodic genes missed. Smaller numbers for false positives and negatives are better.

### 5.1. Simulated data

Our first comparison of RAGE versus RAGE-CC was on a set of 6,500 simulated gene profiles. RAGE outperforms RAGE-CC in accuracy of phase and frequency estimation. That is, the average difference between the *actual* wavelength of the synthetic gene and the estimated frequency is smaller when using the Hausdorff metric. The interested reader is directed to Appendix A.3 for more details on this experiment.

### 5.2. *CDC15*, *CDC28*, *fibroblast* and *drosophila* data

RAGE's performance was also examined on the microarray data described in Section 2.1. The results of those analyses are in Tables 3A–3C. RAGE is more accurate than RAGE-CC at estimating the frequency on the real microarray data. Figures 3A–3D show some selected clusters of genes from the four datasets that RAGE estimated as highly periodic.

TABLE 3. ACCURACY<sup>a</sup>

(A)				(C)			
	<i>CDC15</i>	<i>RAGE</i>	<i>RAGE-CC</i>		<i>Fibroblast</i>	<i>RAGE</i>	<i>RAGE-CC</i>
	Mean $\Delta\omega$ (minutes)	4.39	7.99	Mean $\Delta\omega$ (hours)	2.68	3.05	
	St. dev. $\Delta\omega$ (minutes)	15.05	24.41	St. dev. $\Delta\omega$ (hours)	2.89	3.12	

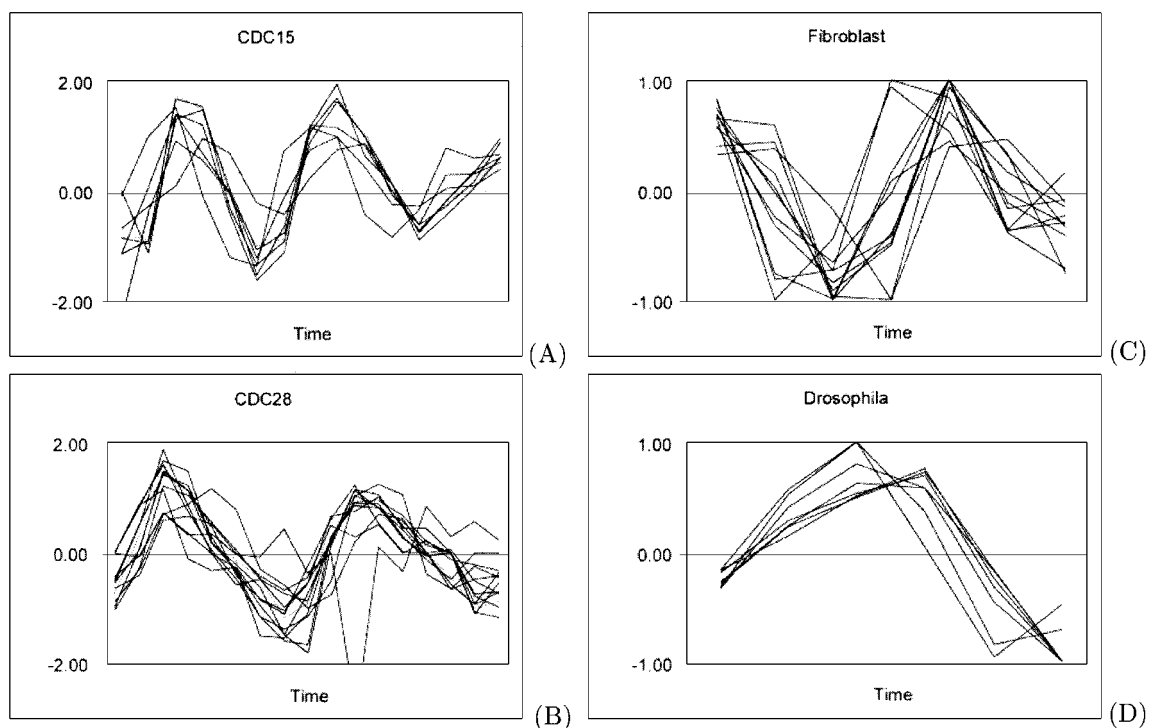
(B)	<i>CDC28</i>	<i>RAGE</i>	<i>RAGE-CC</i>
	Mean $\Delta\omega$ (minutes)	8.57	11.62
	St. dev. $\Delta\omega$ (minutes)	15.8	35.05

<sup>a</sup>Results on actual microarray data. RAGE is more accurate than RAGE-CC at estimating the wavelengths of cell-cycle regulated genes. Tables A and B compare RAGE and RAGE-CC's accuracy at estimating the wavelengths of 104 known cell-cycle-regulated yeast genes identified by White *et al.* (1987), Chapman and Johnston (1989), Siede *et al.* (1989), Johnston *et al.* (1990), Araki *et al.* (1991), Fitch *et al.* (1992), Wan *et al.* (1992), Igual *et al.* (1996), and Caro *et al.* (1998). The wavelengths of those genes is assumed to be 90 minutes. The value  $\Delta\omega$  is the average difference (in minutes) between each program's estimate and 90 minutes. (C) Compares RAGE and RAGE-CC's accuracy at estimating the wavelength of the 479 fibroblast genes identified by Iyer *et al.* (1999). The wavelength of those genes is assumed to be 16 hours. The value  $\Delta\omega$  is the average difference (in hours) between the program's estimate and 16 hours.

TABLE 4. CIRCADIAN GENES<sup>a</sup>

(A)	Gene	Wavelength (hours)	(B)	Gene	Wavelength (hours)
	Cry1	21		<i>Period</i>	18
	BMA	24		<i>Timeless</i>	20
	Per2	24		<i>Vrille</i>	18
				<i>Clock</i>	20
				<i>Cryptochrome</i>	24
				<i>Takeout</i>	24

<sup>a</sup>(A) Three of the seven genes identified as circadian by Yagita *et al.* (2001) are present in the fibroblast cell-cycle dataset. RAGE correctly estimates the frequency of those genes. (B) RAGE correctly estimates the frequency of six known circadian genes (*period*, *timeless*, *vrille*, *clock*, *cryptochrome*, *takeout*) present in the *Drosophila* dataset (McDonald and Rosbash, 2001).



**FIG. 3.** Representative clusters from results on actual microarray data. (A) RAGE identified 367 genes from the CDC15 dataset as having  $90 \pm 10$  minute wavelengths. Nine of these (YDL055C, YDL003W, YDR113C, YGL163C, YML021C, YBR010W, YNL031C, YBR009C, YDL155W) with similar frequencies and phases are shown. (B) RAGE identified 321 genes from the CDC28 dataset as having  $(90 \pm 10)$ -minute wavelengths. Fourteen of these (YDL127W, YMR307W, YNL289W, YJL187C, YOR074C, YLR103C, YOL090W, YDR113C, YNL102W, YOR368W, YNL312W, YJL173C, YKL042W, YKL096W) with similar frequencies and phases are shown. (C) RAGE identified 535 genes from the fibroblast dataset as having  $(16 \pm 2)$ -hour wavelengths. Eleven of these (T62968, H16592, N75026, R87731, H61274, H05133, AA057433, N52170, AA047641, AA025786, H62473) with similar frequencies and phases are shown. (D) RAGE identified 180 genes from the *Drosophila* dataset as having  $(18\text{--}20)$ -hour wavelengths. Seven of these (CG11594, CG10616, *Zw*, CG15240, CG11425, CG8286, CG14279) with similar frequencies and phases are shown. Note that the *Drosophila* dataset was recorded over 20 hours, not 24. Thus, the lowest frequency that can be accurately detected in the data has a wavelength of 20 hours.

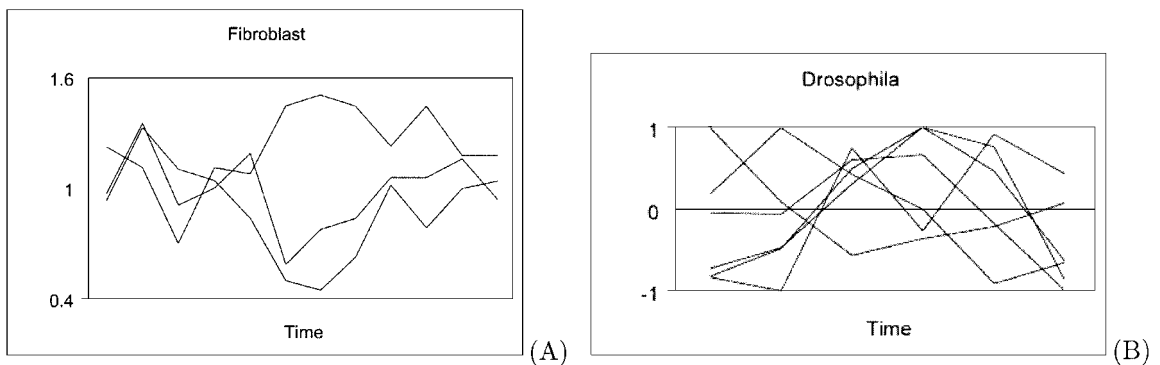
The task of any DNA microarray analysis technique is to find a subset of the genes with a specified property (e.g., cell-cycle regulated, circadian, etc.). Of course, unlike the synthetic data sets described above and in Appendix A.3, the notions of false positives and negatives are not well defined. In some biological systems, there exist genes whose expression profile properties have been well characterized in the literature. For example, there are 104 known cell cycle-regulated genes in yeast (White *et al.*, 1987; Chapman and Johnston, 1989; Siede *et al.*, 1989; Johnston *et al.*, 1990; Araki *et al.*, 1991; Fitch *et al.*, 1992; Wan *et al.*, 1992; Igual *et al.*, 1996; Caro *et al.*, 1998). Therefore, it is possible to evaluate a given method in terms of false negatives when such information is available. However, due to experimental conditions, it is possible for one or more of these known genes to have atypical expression profiles. It is always necessary to go back to the data and examine the profiles of genes that one expected to find.

False positives are much more difficult to evaluate. One can examine the data and look for gross errors. It is customary for biologists to use DNA microarrays as highly parallel screens. Genes with unknown function that exhibit rhythmicity are subsequently studied using slower, more traditional assays.

Direct comparison of the accuracy of multiple techniques on real data is, therefore, best done in terms of false negatives against known genes. Beyond the known genes, it is not possible to do quantitative comparisons other than examining which genes the two methods both identify versus the genes that were uniquely identified by each method. Qualitative assessments, such the number of genes returned as a fraction of the entire genome, can also be helpful.

RAGE identifies a total of 367 genes as cell-cycle regulated in the CDC15 dataset, including 91, or 88%, of the 104 known cell-cycle-regulated genes. In the CDC28 dataset, 321 genes are identified including 81, or 78%, of the 104 known cell-cycle-regulated genes. It is not uncommon to consider the results from multiple datasets on the same organism when examining genetic data (e.g., Birnbaum *et al.*, 2001). By combining the RAGE results on the two yeast data sets, 100, or 96%, of the known genes are found. In contrast, the analysis by Spellman *et al.* (1998) of the same two datasets yielded only 95, or 91%, of the 104 known cell-cycle-regulated genes. RAGE finds 396 rhythmic genes in both yeast datasets with wavelengths consistent with the length of the cell-cycle in yeast. Spellman *et al.* (1998) found 800. It is interesting that RAGE's more conservative estimate actually finds more of the 104 known cell-cycle-regulated genes.

In the human fibroblast dataset, RAGE identifies 535 genes as cell-cycle regulated, including 444, or 93%, of the 479 known cell-cycle-regulated genes. As noted in Section 2.1, the human fibroblast dataset was collected over a 24-hour time period. While this is not a true circadian experiment, it is known that the clock genes are expressed in peripheral tissues such as fibroblasts (Balsalobre *et al.*, 1998). Hence, it should be possible to look for circadian genes in the same dataset. This had not been done before; we ran RAGE to try it. Our results were compared to those of Yagita *et al.* (2001) who found circadian genes in cultured rat fibroblasts, using (different) microarray experiments and data. Three of the seven genes



**FIG. 4.** Circadian genes identified by RAGE. **(A)** The expression profile of three genes (*Cry1*, *Per2*, and *BMAL*) from the human fibroblast dataset. These genes were identified as circadian by RAGE. The genes *Cry1*, *Per2*, and *BMAL* are close human homologs of rat circadian genes identified by Yagita *et al.* (2001). **(B)** The expression profiles of six known circadian genes (*period*, *timeless*, *vrille*, *clock*, *cryptochrome*, *takeout*) from the *Drosophila* data set. These genes were identified as circadian by RAGE.

identified as circadian by Yagita *et al.* (2001) are present (as close human homologs) in the fibroblast cell-cycle dataset. RAGE correctly estimates the frequency of those genes: *Cry1*, *Per2*, and *BMAL*. In Table 3A, RAGE's estimate of wavelengths for *Cry1*, *Per2*, and *BMAL* is reported. RAGE successfully estimates the wavelengths of these three genes as being circadian. Figure 4D shows the expression profiles for *Cry1*, *Per2*, and *BMAL* from the human fibroblast dataset.

Finally, RAGE was used to analyze the *Drosophila* dataset. Of the 14,010 genes in the complete dataset, RAGE identifies 180 genes as circadian while McDonald and Rosbash (2001) identified 134. Both RAGE and McDonald and Rosbash (2001) identify 6 of the known circadian genes (*period*, *timeless*, *vriille*, *clock*, *cryptochrome*, *takeout*). In all, 131 genes were identified by both McDonald and Rosbash (2001) and RAGE. Figure 3D shows some of the circadian profiles discovered by RAGE. Figure 4B shows the profiles of *period*, *timeless*, *vriille*, *clock*, *cryptochrome*, and *takeout*.

## 6. CONCLUSION

Genome-wide mRNA expression time-series experiments are an important source of biological information. The discovery of periodic gene expression profiles is especially useful for the study of rhythmic processes such as the cell-cycle and the circadian clock. The sheer volume of data generated by microarray experiments prohibits manual inspection of all the data. Therefore, algorithms for identifying rhythmic genes are needed.

Fourier-based techniques are not yet appropriate for microarray data because the number of time-points in a typical experiment is too small to yield adequate frequency resolution. Model-based techniques are therefore needed. Of the existing model-based techniques, those using synthetic models, such as CorrCos, are very accurate, but are computationally inefficient. We have presented a novel technique that is computationally much more efficient. It gains its efficiency through the use of a phase-independent search of frequency-space. Furthermore, RAGE uses a true mathematical metric to compute the similarity between gene expression profiles. The Hausdorff distance is a more accurate measurement and, therefore, RAGE tends to give more accurate frequency and phase estimates on both simulated and actual microarray data. The Hausdorff distance may prove useful in other microarray data applications, such as clustering. We are presently exploring its performance in that application. Finally, the RAGE algorithm can be combined with maximum entropy spectral analysis, resulting in even greater improvements in accuracy and efficiency. For work in this direction, please see Langmead *et al.* (2002).

## 7. SOFTWARE

The RAGE software is available by contacting the authors.

## APPENDICES

### A.1. Fourier transform

The Nyquist frequency is the maximum frequency that can be detected within a sampled signal. It is one half the sampling rate. The minimum frequency that can be detected in a sampled signal is determined by the sampling interval. If the signal is collected over  $n$  time units, then the lowest detectable frequency is 1 cycle per  $n$  time units. The resolution of a spectrum is determined by the number of points collected. If there are  $m$  points in the signal, the resulting spectrum will have  $m$  equally spaced frequency intervals from the lowest frequency to the Nyquist frequency.

An important assumption made by the Fourier transform is that the data is linearly sampled. The CDC15 dataset, for example, is nonlinearly sampled, making the Fourier transform inappropriate. Of course, it is sometimes possible to discard data points in order to create a linearly sampled signal, but as noted above, eliminating data points reduces the resolution of the resulting spectrum.

For example, the CDC15 dataset was collected over 290 minutes. The sampling was nonlinear. It is possible to create a linearly sampled dataset of 15 time points spaced 1 every 20 minutes. The Nyquist frequency of this reduced set is one cycle per forty minutes, and the minimum frequency is one cycle per 290 minutes. Since there are 15 time points in the reduced set, the resolution of the spectrum is one cycle per  $(290 - 40)/(15 - 1) = 17.9$  minutes. In other words, two periodic functions must differ by one cycle per 17.9 minutes (or more) in order to be well resolved (distinguishable) by the Fourier transform.

Spellman *et al.* (1998) used a variant of the discrete Fourier transform (Equation (10)).

$$X(\omega) = \frac{1}{N} \sum_{n=0}^{N-1} x(n) e^{-\frac{2\pi i \omega n}{N}} \quad (10)$$

Spellman *et al.* (1998) noted “the magnitude of the Fourier transform was unstable for small variations of  $\omega$ ” and were forced to average the results of multiple transforms in order to come to a final frequency estimate. The issue was primarily one of resolution in the actual signal. Not surprisingly, the authors resorted to a hybrid approach including nonspectral methods to estimate frequencies. The final frequency and phase estimate was a combination of both spectral and nonspectral methods.

### A.2. Autocorrelation

The autocorrelation is truly phase-independent for periodic signals of infinite length. For finite-length signals, the numeric error associated with signals differing only in phase is proportional to the difference in phase and inversely proportional to the number of periods in the signal. In our experiments on simulated data modeled on the CDC15, CDC28, and fibroblast datasets, that error represents approximately a 1% error in frequency estimation. By contrast, the calculated frequency resolution of the real datasets used in this paper represents between a 15% and 25% error (see Section A.1). The noise in the real datasets is also likely to mask the error associated with the autocorrelation.

### A.3. Simulated data results

A set of 6,500 simulated gene expression profiles were generated. Of those, 1,500 were periodic and the remaining ones were aperiodic random signals. The elements of that set are detailed in Table 2A. The results of running RAGE and RAGE-CC on that dataset are reported in Table 2B. RAGE outperforms RAGE-CC in accuracy of phase and frequency estimation. That is, the average difference between the *actual* wavelength of the synthetic gene and the estimated frequency is smaller when using the Hausdorff metric. RAGE also has a lower rate of false-positives. We conclude that the Hausdorff metric is superior to the correlation coefficient.

## ACKNOWLEDGMENTS

We would like to thank Ryan Lilien and all members of Donald Lab for helpful discussions and suggestions. We thank Dr. M. Straume, author of CorrCos (Harmer *et al.*, 2000), for reading and commenting on a draft of this paper. Thanks to Drs. J.M. Groh and H. Farid for their suggestions on the signal processing and statistical aspects of RAGE. This work is supported by the following grants to B.R.D.: National Institutes of Health (R01-GM65982), National Science Foundation (IIS-9906790, EIA-0102710, EIA-0102712, EIA-9818299, EIA-0305444, and EIA-9802068), a Guggenheim Foundation Fellowship, and an equipment grant from Microsoft Research. C.R.M. was supported by National Science Foundation grants MCB-9723482 and MCB-0091008.

## REFERENCES

Araki, H., Hamatake, R.K., Morrison, A., Johnson, A.L., Johnston, L.H., and Sugino, A. 1991. Cloning DPB3, the gene encoding the third subunit of DNA polymerase II of *Saccharomyces cerevisiae*. *Nucl. Acids Res.* 19, 4867–4872.

- Arkin, E., Chew, L.P., Huttenlocher, D.P., Kedem, K., and Mitchell, J.S.B. 1991. An efficiently computable metric for comparing polygonal shapes. *IEEE Trans. Patt. Anal. Mach. Intell.* 13(3), 209–216.
- Balsalobre, A., Damiola, F., and Schibler, U. 1998. A serum shock induces circadian gene expression in mammalian tissue culture cells. *Cell* 93, 929–937.
- Belogay, E., Cabrelli, C., Molter, U., and Shonkwiler, R. 1997. Calculating the Hausdorff distance between curves. *Information Processing Letters* 64(1), 17–22.
- Birnbaum, K., Benfey, P.N., and Shasha, D.E. 2001. cis Element/transcription factor analysis (cis/TF): A method for discovering transcription factor/cis element relationships. *Genome Res.* 11, 1567–1573.
- Caro, L.H., Smits, G.J., van Egmond, P., Chapman, J.W., and Klis, F.M. 1998. Transcription of multiple cell wall protein-encoding genes in *Saccharomyces cerevisiae* is differentially regulated during the cell cycle. *FEMS Microbiol. Lett.* 161, 345–349.
- Chapman, J.W., and Johnston, L.H. 1989. The yeast gene, *DBF4*, essential for entry into S phase is cell cycle regulated. *Exp. Cell Res.* 180, 419–428.
- Cho, R., Campbell, M., Winzeler, E., Steinmetz, L., Conway, A., Wodicka, L., Wolfsberg, T., Gabrielian, A., Landsman, D., Lockhart, D., and Davis, R. 1998. A genome-wide transcriptional analysis of the mitotic cell cycle. *Mol. Cell* 2, 65–73.
- Claridge-Chang, A., Wijnen, H., Naef, F., Boothroyd, C., Rajewsky, N., and Young, M.W. 2001. Circadian regulation of gene expression systems in the drosophila head. *Neuron* 32, 657–671.
- Donald, B.R., Kapur, D., and Mundy, J. 1992. Huttenlocher, D., and Kadem, K. *Symbolic and Numerical Computation for Artificial Intelligence*, 201–219, Academic Press, Harcourt Jovanovich, London.
- Eisen, M., Spellman, P.T., Botstein, D., and Brown, P.O. 1998. Cluster analysis and display of genome-wide expression patterns. *Proc. Natl. Acad. Sci.* 95(25), 14863–14868.
- Filkov, V., Skiena, S., and Zhi, J. 2001. Analysis techniques for microarray time-series data. *Proc. 5th Ann. Intl. Conf. on Computational Biology*, 124–131.
- Fitch, I., Dahmann, C., Surana, U., Amon, A., Nasmyth, K., Goetsch, L., Byers, B., and Futcher, B. 1992. Characterization of four B-type cyclin genes of the budding yeast *Saccharomyces cerevisiae*. *Mol. Biol. Cell* 3, 805–818.
- Grundschober, C., Delaunay, F., Phlhofer, A., Triqueneaux, G., Laudet, V., Bartfai, T., and Nef, P. 2001. Circadian regulation of diverse gene products revealed by mRNA expression profiling of synchronized fibroblasts. *J. Biol. Chem.* 276, 46751–46758.
- Harmer, S., Hogenesch, J.B., Straume, M., Chang, H.S., Han, B., Zhu, T., Wang, X., Kreps, J.A., and Kay, S.A. 2000. Orchestrated transcription of key pathways in *arabidopsis* by the circadian clock. *Science* 290, 2110–2113.
- Huttenlocher, D.P., and Kedem, K. 1990. Computing the minimum Hausdorff distance for point sets under translation. *Proc. 6th ACM Symp. Computational Geometry*, 340–349.
- Igual, J.C., Johnson, A.L., and Johnston, L.H. 1996. Coordinated regulation of gene expression by the cell cycle transcription factor Swi4 and the protein kinase C MAP kinase pathway for yeast cell integrity. *EMBO J.* 15, 5001–5013.
- Iyer, V.R., Eisen, M.B., Ross, D.T., Schuler, G., Moore, T., Lee, J.C.F., Trent, J.M., Staudt, L.M., Hudson, J. Jr., Boguski, M.S., Lashkari, D., Shalon, D., Botstein, D., and Brown, P.O. 1999. The transcriptional program in the response of human fibroblasts to serum. *Science* 283, 83–87.
- Johnston, L.H., White, J.H., Johnson, A.L., Lucchini, G., and Plevani, P. 1990. Expression of the yeast DNA primase gene, *PRI1*, is regulated within the mitotic cell cycle and in meiosis. *Mol. Gen. Genet.* 221, 44–48.
- Langmead, C.J., McClung, C.R., and Donald, B.R. 2002. A maximum entropy algorithm for rhythmic analysis of genome-wide expression patterns. *Proc. IEEE Computer Society Bioinformatics Conference (CSB), Stanford University*, 237–245.
- McDonald, M.J., and Rosbash, M. 2001. Microarray analysis and organization of circadian gene expression in *Drosophila*. *Cell* 107, 567–578.
- Mumford, D. 1987. The problem of robust shape descriptors. *Proc. 1st Int. Conf. Comput. Vision*, 602–606.
- Schaffer, R., Landgraf, J., Accerbi, M., Simon, V., Larson, M., and Wisman, E. 2001. Microarray analysis of diurnal and circadian-regulated genes in *Arabidopsis*. *Plant Cell* 13, 113–123.
- Siede, W., Robinson, G.W., Kalainov, D., Malley, T., and Friedberg, E.C. 1989. Regulation of the *RAD2* gene of *Saccharomyces cerevisiae*. *Mol. Microbiol.* 3, 1697–1707.
- Spellman, P., Sherlock, G., Zhang, M.Q., Iyer, R.I., Anders, K., Eisen, M.B., Brown, P.O., Botstein, D., and Futcher, B. 1998. Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces cerevisiae* by microarray hybridization. *Mol. Biol. Cell* 9(12), 3273–3297.
- Wan, J., Xu, H., and Grunstein, M. 1992. *CDC14* of *Saccharomyces cerevisiae*. Cloning, sequence analysis, and transcription during the cell cycle. *J. Biol. Chem.* 267, 11274–11280.
- White, J.H., Green, S.R., Barker, D.G., Dumas, L.B., and Johnston, L.H. 1987. The *CDC8* transcript is cell cycle regulated in yeast and is expressed coordinately with *CDC9* and *CDC21* at a point preceding histone transcription. *Exp. Cell Res.* 171, 223–231.

Yagita, K., Tamanini, F., van der Horst, G.T.J., and Okamura, H. 2001. Molecular mechanisms of the biological clock in cultured fibroblasts. *Science* 292, 278–281.

Address correspondence to:  
*Bruce Randall Donald*  
*6211 Sudikoff Laboratory*  
*Dartmouth Computer Science Department*  
*Hanover, NH 03755*

*E-mail:* brd@cs.dartmouth.edu