

Visualization of Growth Curve Data from Phenotype Microarray Experiments

Janet S. Jacobsen¹, Dominique C. Joyner², Sharon E. Borglin²,
Terry C. Hazen^{2,4}, Adam P. Arkin^{3,4}, E. Wes Bethel¹

¹Computational Research Division, Lawrence Berkeley National Laboratory,

²Earth Sciences Division, Lawrence Berkeley National Laboratory,

³Physical Biosciences Division, Lawrence Berkeley National Laboratory,

⁴Virtual Institute for Microbial Stress and Survival

{JSJacobsen, DCJoyner, SEBorglin, TEHazen, APArkin, EWBethel}@lbl.gov

Abstract

Phenotype microarrays provide a technology to simultaneously survey the response of an organism to nearly 2,000 substrates, including carbon, nitrogen and potassium sources; varying pH; varying salt concentrations; and antibiotics. In order to more quickly and easily view and compare the large number of growth curves produced by phenotype microarray experiments, we have developed software to produce and display color images, each of which corresponds to a set of 96 growth curves. Using color images to represent growth curves data has proven to be a valuable way to assess experiment quality, compare replicates, facilitate comparison of the responses of different organisms, and identify significant phenotypes. The color images are linked to traditional plots of growth versus time, as well as to information about the experiment, organism, and substrate. In order to share and view information and data project-wide, all information, plots, and data are accessible using only a Web browser.

Keywords--- *Desulfovibrio vulgaris*, growth curve, phenotype microarray, space efficient color display, information visualization

1. Introduction

Phenotype microarrays (PM) provide a high-throughput technique for screening the response of an organism to thousands of chemical treatments. We have developed a suite of software tools that display growth curve data generated by PM experiments on the researcher's computer screen without having to install any software other than a Web browser. The software presents the growth curve data at different levels of detail and using different display and visualization techniques so that the researcher may easily and quickly assess experiment quality, compare replicates, compare the responses of different organisms to chemical treatments,

and determine which growth curves represent interesting and significant phenotypes. The tools were developed at the request of and with design guidance from experimentalists generating PM data.

With current PM technology, it is possible to screen the response of an organism to as many as 1,920 chemical treatments in a matter of days. The chemical treatments are provided as substrates in microwells on 96-well plates that are inoculated with the organism(s) being studied. The Omnilog[®] instrument [1], which provides a digital readout of the growth in each well, can scan fifty 96-well plates at one time, so that replicate data can be produced, or more than one organism can be studied. The result is that in one to four days, data for 4,800 growth curves can be generated. By representing sets of growth curves as color images, we are able to display growth curves corresponding to thousands of treatments on a single Web page with a window size of 8½ by 12 inches. This was done at the request of the experimentalists in order to match data analysis efficiency with data generation capability.

The traditional way of looking at cell growth is to plot cell density (or a surrogate for cell density) vs. time. Under normal conditions in a microwell (*i.e.*, limited source of nutrients), cell growth shows a sharp increase (this is the log phase), reaches a maximum, and then levels off as cell increase is offset by cell death. When an organism is subjected to a stressor, its growth is affected, resulting in poor or no growth, a delay in the onset of growth, or a combination. When using PM technology, analysis of growth is complicated by factors such as noise in the digital signal, under or over inoculation of the microwells, or chemical reactions between the media the cells were grown in (prior to inoculation) and the chemical treatment in the microwell.

Currently, researchers performing PM experiments have three options for viewing data: use the vendor-supplied software for looking at the data [2], import the data into another plotting program (*e.g.*, spreadsheet), or look at digital images of the plates [3]. The third option

may be appropriate for studies limited to growth/no growth results as only the endpoint data can be assessed. In general, however, the onset of growth and the rate of growth are as important as the maximum cell density reached, so that it is essential to be able to look at kinetic plots of growth versus time. Our software tools were developed with that need in mind.

None of the three options above is ideal for sharing large amounts of data with project participants at different institutions. Our approach is to provide remote users, using only a Web browser, with the ability to display and access growth curve data stored in a project repository. This approach saves the researcher from having to import the PM data into a plotting program and generate thousands of plots, does not require the researcher to export images as files and email them to colleagues, and does not require the installation of any software other than a Web browser.

In addition, we have implemented a number of color maps (transfer functions) that map the magnitude of growth to a color in order to generate color images representing growth curves. Though this is not a new visualization technique, the application to visualizing growth curve data, in particular, growth curve data from phenotype microarray experiments, is new. Our approach combines the advantage of the 'small multiple' [4] representation of data with the use of color to highlight anomalies and facilitate comparisons. The vendor-supplied software also uses small multiples, but relies on curve shape or 'difference' plots to compare experiments. Because the eye is very good at discriminating differences in color, using color to represent of growth has a significant advantage over using shape when trying to scan or compare hundreds of growth curves at one time. As described later, the color images are linked to growth curves plotted in the traditional manner, so the user has a choice of how the data are displayed.

The next section discusses the microbe that we are studying, followed by more details about PM experiments. The software tools that we have developed are described in detail, including examples of how color images are used to detect anomalous results, compare replicates, and compare the responses of different organisms to chemical treatments. We conclude with a brief discussion of other techniques for mining data from PM experiments.

2. Use of Phenotype Microarrays

The United States Department of Energy's (DOE) Genomics:GTL program [5] is investigating the potential of using microbes for bioremediation, energy production, and carbon sequestration. Microbes are ubiquitous and have adapted to practically every environmental niche on earth. Some microbes live and even thrive in conditions generally thought to be inhospitable to life.

Funded by the Genomics:GTL program, the Virtual Institute for Microbial Stress and Survival (VIMSS) [6] has been studying the response of *Desulfovibrio vulgaris*

Hildenborough (DvH), an anaerobic, sulfate-reducing, bacterium found in sediments, to environmental stressors such as oxygen, sodium chloride, nitrate, nitrite, and low and high pH levels. These environmental stressors are typical of environmental conditions found in DOE waste sites contaminated with metals and radionuclides. The VIMSS project uses not only phenotype microarray experiments to study response characteristics, but also gene expression, metabolomic, and proteomic experiments [7].

Phenotype microarray experiments primarily are being used to survey the phenotypic response of DvH and other microbes of interest to a larger number of stressors than could be studied using other more time and labor-intensive analysis methods. PM experiments also are being used to compare the response of wild-type DvH to constructs created by knocking out selected genes [8].

3. Phenotype Microarray Experiments

In a PM experiment, each microwell of a 96-well plate is inoculated with a culture containing cells that have been harvested at mid-log phase. (The log phase is the stage of cell growth where the cells are dividing at a constant rate.) The plate is then placed in the Omnilog® instrument that records a digital image of the change in color or opacity of each well on the plate.

Each PM plate consists of a 96 microwells arranged in eight rows, labeled A to H, and twelve columns, labeled 1 to 12. Generally, each PM plate contains substrates from one or two mode of action (MOA) categories (Table 1). MOA refers to the expected effect that the chemical in the well will have on the organism.

Plate Type	Mode of Action Category
PM01-PM02	carbon source
PM03	nitrogen source
PM04	sulfur, phosphorus sources
PM05	nutritional supplement
PM06-PM08	nitrogen utilization
PM09	osmotic sensitivity, toxicity
PM10	pH
PM11-PM20	inhibitors, chemical sensitivity

Table 1. Mode of action categories for phenotype microarray plates PM01 to PM20.

Because DvH is an anaerobe (its growth is inhibited by oxygen), studying DvH using phenotype microarrays presents a particular challenge [9]. DvH cells must be grown and handled in anaerobic conditions. After the PM plates are inoculated with the DvH cell culture, the PM plates are sealed in oxygen-impermeable bags before being placed in the Omnilog®. Run times for wild-type DvH and DvH mutant PM experiments vary from 24 to 96 hours.

DvH is grown in a defined minimal medium that reacts with the hydrogen sulfide produced by the microbe as it grows. The product is a black iron sulfide

precipitate, which as it accumulates, results in a darkening (increase in turbidity) of the microwell. Cell growth also contributes to turbidity.

4. Visualization of Growth Curve Data

For the purpose of this discussion, a PM dataset consists of growth curve data from a phenotype microarray experiment involving a single organism and a set of distinct PM plates run on the same day with the same inoculum. We have developed two categories of interfaces that display growth curves in one of two formats: either as traditional growth curves (measure of growth vs. time) or as color images. Because the color image format is linked to the traditional format, we describe the PMRowView interface first.

4.1. PMRowView Interface

Figure 1 shows an example of the typical way in which the PMRowView interface is used. The figure was constructed using screenshots of the interface rather than a text-block diagram in order to give a better feel for how the interface operates. (The screenshots may be enlarged to see additional detail.)

The first page of the PMRowView interface (Fig. 1a) allows the user to select one or more PM datasets, a single plate type, and a row (12 wells) on that plate. The user has the option to pre-select the colors to use for plotting the growth curves for each dataset, and in addition, may choose to plot log values (the underlying plotting program does not support log scales), or normalize the curves by the maximum value.

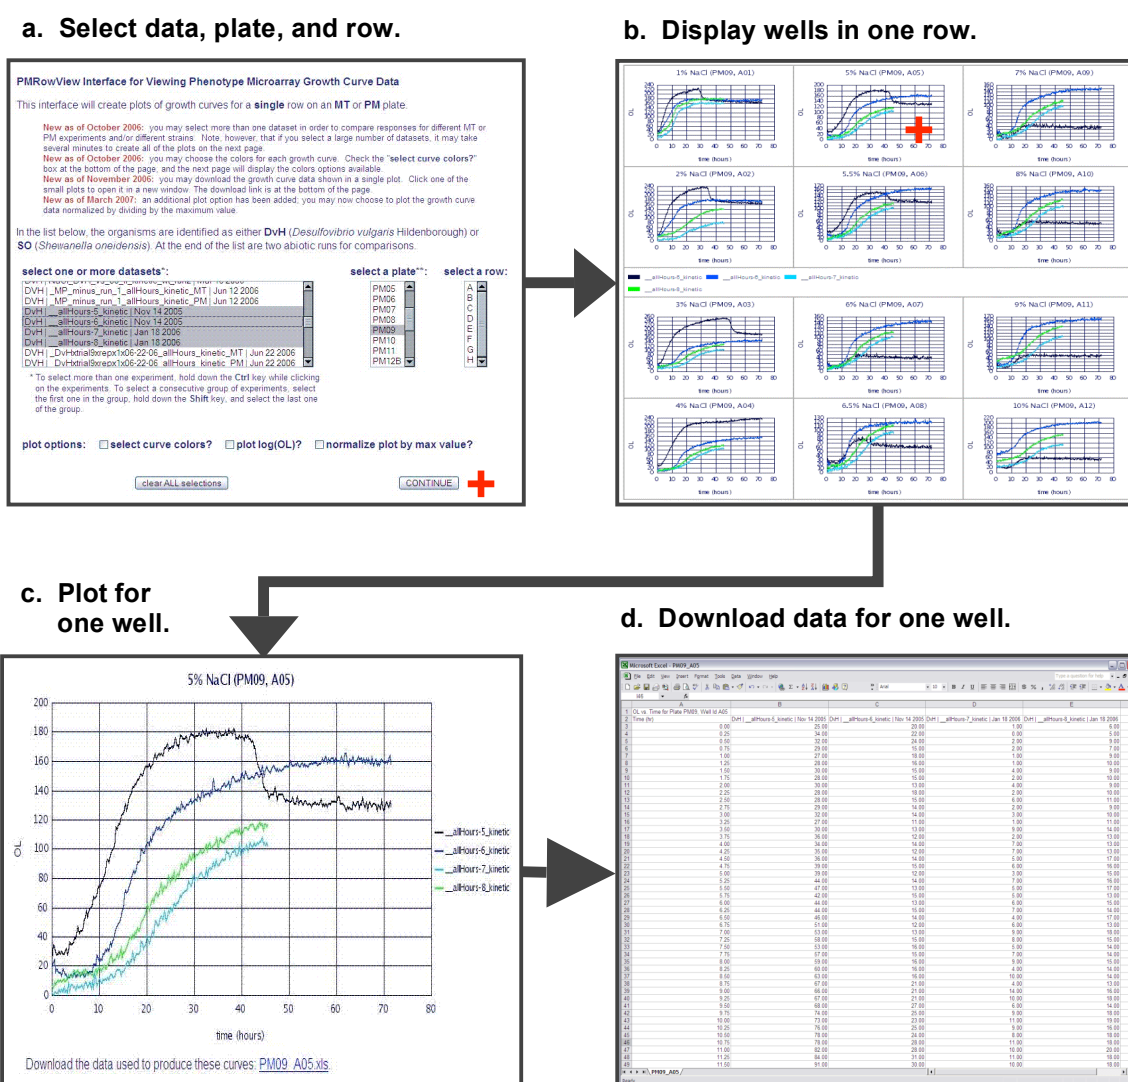


Figure 1. Typical use of the **PMRowView** interface. (a) On the home page for the interface, the user selects datasets, plate type, and row on plate, and clicks the Continue button (red cross). (b) Screenshot shows the results of selecting four datasets, plate type PM09 and row A. (c) Growth curve for well selected by clicking on the plot for well A05 (red cross in b). (d) Spreadsheet, which can be downloaded by clicking on the download link (red cross in c), of growth curve data plotted in (c).

The interface will flag any datasets not containing the selected plate type (some PM experiments are run with fewer than 20 plates) and will plot the growth curves for the datasets that contain the selected plate type. The PMRowView interface plots the growth curves twelve to a page in a three-column, four-row format (Fig. 1b). A view of one plot (at the approximate size that it would be viewed on the Web page), showing how each small plot is labeled with the chemical name, plate type, and well id, of one of the small plots in Fig. 1b, is shown in Fig. 2.

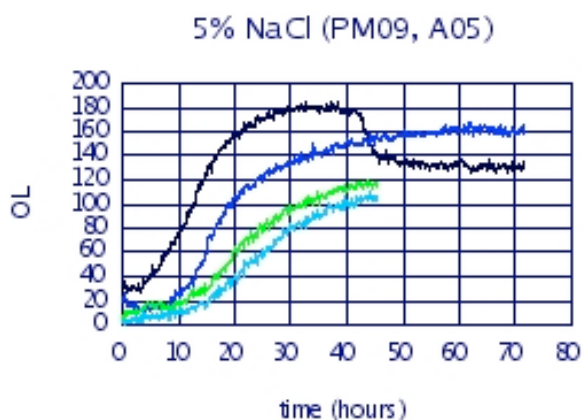


Figure 2. Example of one of twelve plots that make up the PMRowView display. The plot shows two pairs of technical replicates. The blue and black curves are biological replicates of the green and turquoise curves. The label on the y-axis, OL, stands for Omnilog units referring to the measurement unit of the Omnilog® instrument.

Note that the vertical scales on each plot in Fig. 1b vary. The reason for this is to see the maximum amount of detail in each growth curve. We plan to implement a plot option in the future so that all growth curves can be plotted on the same scale.

Because some chemicals are used to treat four consecutive wells (using different doses) on a plate, the plots shown in Fig. 1b are placed on the page so that when a chemical is repeated, growth curves for all four doses appear in one column in order to facilitate comparing dose responses. A link on the Web page (not shown) allows the user to advance to the growth curves for the next row on the plate (for the selected dataset(s)) without having to return to the selection menu page.

If the user clicks on one of the twelve small plots, a new window will open showing a larger version of the plot (Fig. 1c). In the new window, the user is given the option to download the growth curve data for the curve(s) shown in the plot (Fig. 1d) in spreadsheet format.

The plots produced by the PMRowView interface are generated on the fly. The growth curve data are read

from files; information about each selected dataset (organism and experiment date) and the chemical treatments are read from a relational database.

The PMRowView interface is useful for single row (12 wells) or single well comparisons of technical replicates (same inoculum), biological replicates (same organism, but different inoculum), and the responses of different strains of an organism.

An extension of the PMRowView interface (not shown) displays pre-averaged replicates (either biological or technical) and error bars (\pm one standard deviation).

4.2. PMColorMap Interface

The PMColorMap interface displays a growth curve as a thin horizontal line in which the magnitude of turbidity is represented by color. Using this visualization technique, and by vertically stacking the horizontal lines representing the growth curves, it is possible to display growth curves for all 96 wells on a single PM plate in a relatively small space. In a Web browser, each color image in the preview mode is approximately 1 ½ by 2 inches. An example of such a color image is shown in Fig. 3, which also includes annotation that explains the direction of increasing time and which region and which line of the color image corresponds to each row and well on the PM plate.

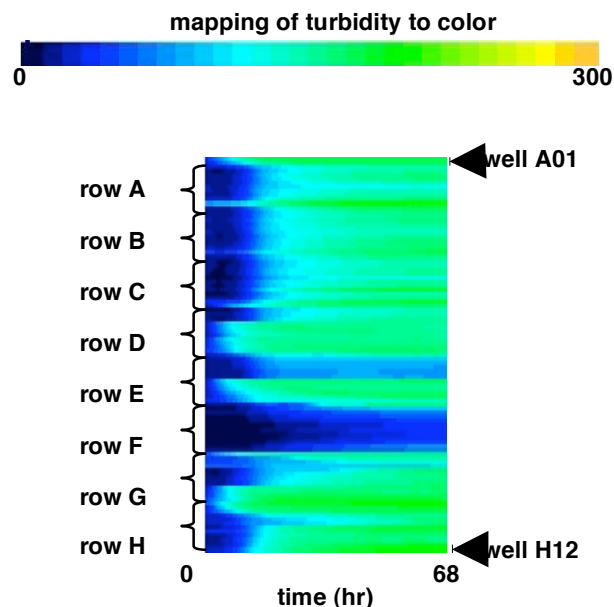


Figure 3. Color image for plate PM09 showing color map (top), time scale (bottom), and row and well locations (sides).

The growth curve data are smoothed using a moving window average algorithm [10] prior to generating the corresponding color image. The reason for smoothing

the data prior to generating the color images is that PM data are noisy. Without smoothing the data, the amount of noise in the data results in color images with rapid changes in color (*vs.* time) that make it difficult to see trends in the data. Our goal is to be able to use the color images to show general characteristics of and trends in the data (see Section 4.3), which requires displaying the growth curve data as smoothly varying and thus precludes other coloring methods (*e.g.*, two-tone pseudo coloring [11]).

We experimented with different methods of averaging the data, as well as different color maps, until we achieved result that satisfied the experimentalists. Because two people working on the VIMSS project have some degree of colorblindness, we implemented three color maps to color code the growth curve. Figure 4 illustrates how the data for the same 96-well PM plate looks when colored using each of the three available color maps. The color map whose use is illustrated in Fig. 4b, uses a 'color safe' palette [12], that excludes colors that a person with colorblindness cannot easily distinguish.




Figure 4. Examples of a single 96-well dataset colored coded using the three PMColorMap options: (a) standard, (b) color safe, (c) gray scale. Images are shown reduced in width. The colors for color map b were chosen from a color safe palette [12] suitable for people who are colorblind.

Figure 5 shows the typical use of the PMColorMap interface for comparing data from a single PM plate used in multiple PM experiments. As with Fig 1, this figure was constructed using screenshots instead of text blocks. Larger versions of the color images shown in Figs. 5b and 5c are presented in Figs. 3 and 6, respectively.

The use of the PMColorMap interface shown in Fig. 5 illustrates Shneiderman's "Visual Information Seeking Mantra" [13]: "overview" (Fig. 5b), "zoom" (Fig. 5c), "filter" (Fig. 5d), "details" (Fig. 5e, 5f). The first page of the interface (Fig. 5a) allows the user to select the datasets, plate type, and color map. The selected data are displayed as small color images in the 'preview' mode shown in Fig. 5b.

The four small color images in the preview mode show the equivalent of 384 growth curves in a single view. Even at the scale of the reduced-size screenshot, this example shows a couple of interesting features that would be very time consuming to discover looking at hundreds of growth curve plots. The four plots are two pairs of technical replicates. The two technical replicates on the right hand side show good agreement with each other, but not with the two biological replicates (two images on the left) that they correspond to. In addition, the image on the far left shows a region (yellow circle) where the wells were not properly inoculated.

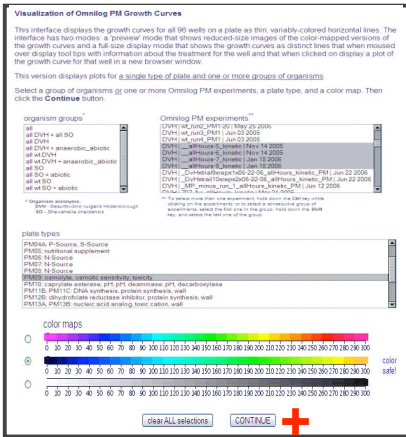
Clicking on one of the small color images in the preview mode (Fig. 5b) opens a separate window that shows the data in the 'full-size' mode (Fig. 5c). The full-size color image is about 6 by 5½ inches (see Fig. 6).

In the full-size version, the thin horizontal lines that represent the growth curves are slightly separated (see Fig. 6) so that when the cursor is moved slowly over individual lines, a 'tool tip' appears that displays the plate type, well id, chemical name, and the mode of action of the chemical (Table 1). If the user clicks on a colored line, a separate window opens that displays a plot of growth *vs.* time for that well (Fig. 5e). The separation between the colored lines (white spaces between the colored lines in Fig. 6) is necessary to help the user to be able to distinguish between individual colored lines so that the user is able to select (by clicking on the colored line) the growth curve for the well of interest.

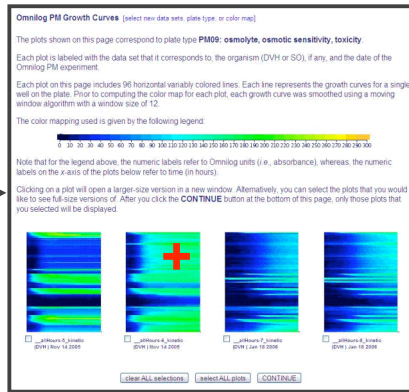
If the user clicks on one of the row labels to the left of the color image (*e.g.*, the red cross in Fig. 5c or "row A" in Fig. 6), a window opens that displays the growth curves for all twelve wells in that row using the PMRowView interface (Fig. 5d). In this way, the user can easily navigate between the 96-well, plate-level view, and the more detailed single-row (12 wells) view or single-well view. Figs. 5e and 5f are analogous to Figs. 1c and 1d.

Though not shown, there is a second version of the PMColorMap interface that displays all of the data for all of the PM plates run for a single experiment. This interface provides a quick check of the overall quality of an experiment.

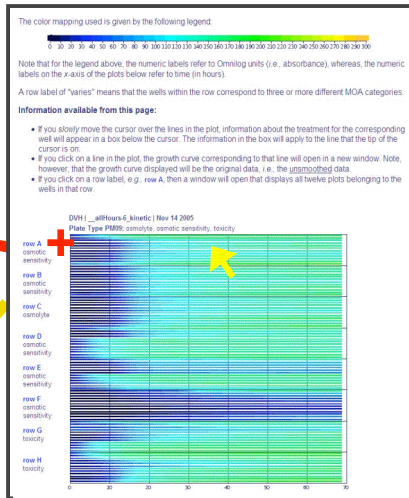
a. Select data, plate, color map.



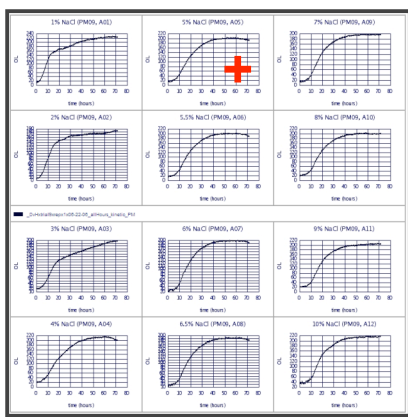
b. Preview data for 4 x 96 wells.



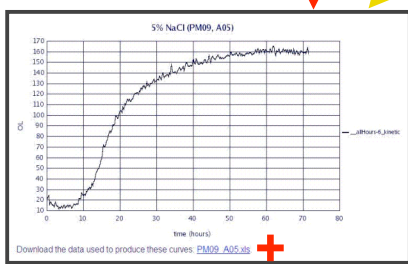
c. Full-size view of data for 96 wells.



d. Display wells in one row.



e. Plot for one well.



f. Download data for one well.

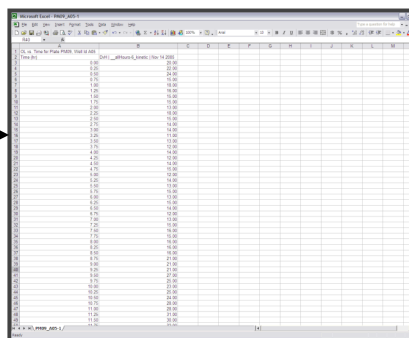


Figure 5. Typical use of the **PMColorMap** interface. (a) On the home page for the interface, the user selects datasets, plate type, and color map, and clicks the Continue button (red cross). (b) Preview page shows the results of selecting plate type PM09 and datasets from four PM experiments using DVH (4 x 96 wells). Yellow circle shows a group of wells that were not properly inoculated. (c) Results of the user selecting (red cross in b) a single dataset (96 wells) from the preview mode. Next user can opt to display data for all of the wells in a single row (red cross) or from a single well (yellow arrow). (d) Screenshot showing the result of clicking on the "row A" label in c, *i.e.*, growth curves for the twelve wells in row A. (e) Growth curve for single well selected either from c (red cross) or d (yellow arrow). (f) Spreadsheet, which can be downloaded by clicking on the download link (red cross in e), of growth curve data plotted in (e).

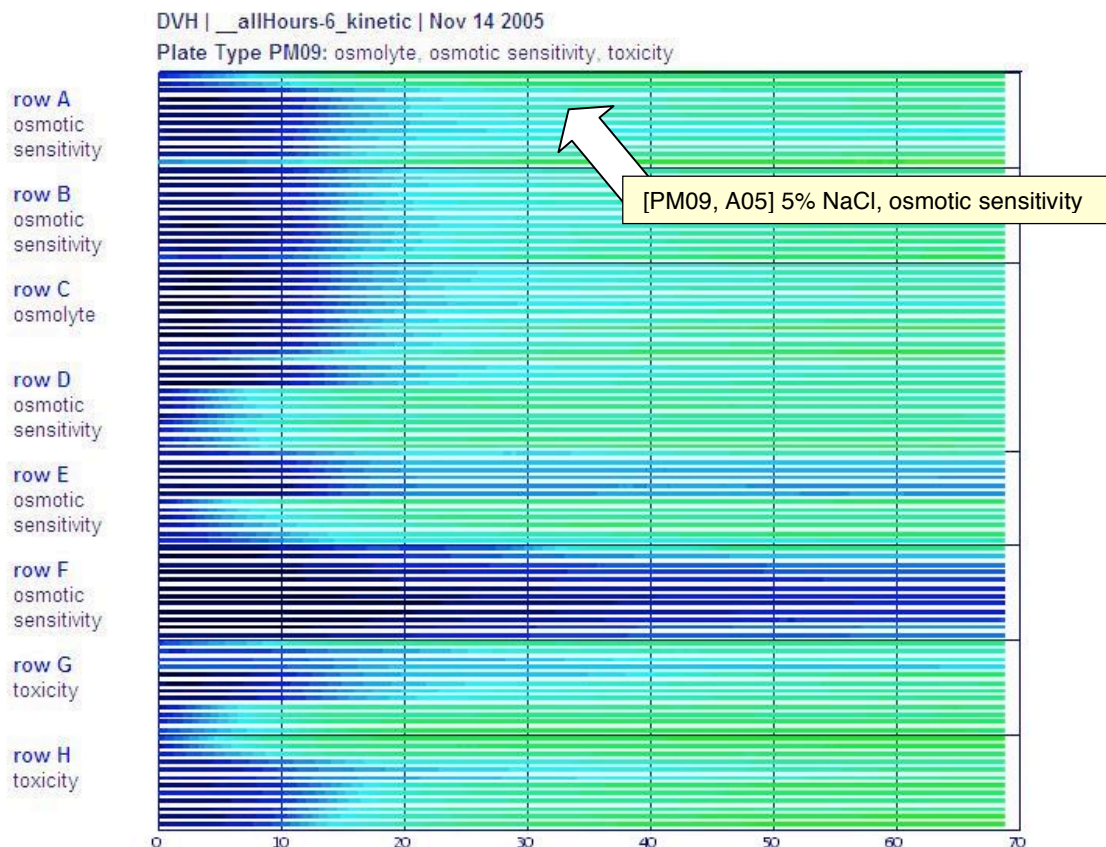


Figure 6. Full-size color image. To provide detailed information about the well (chemical and mode of action), the yellow label appears when the cursor is positioned over (in this example) the 5th color line. The row labels to the left of the figure are linked to the PMRowView interface (see text).

4.3. Application of the PMColorMap Interface

Because of its compact display of growth curves, the preview mode of the PMColorMap interface shows at a glance whether there are anomalies resulting from equipment malfunctions, chemical reactions between the medium that the cells were grown in and the substrate in the microwell, or poor cell growth. Figures 7 to 9 show examples of each of these.

The sharp spikes shown in the growth curve in Fig. 7a are the result of instrument malfunction or operator interrupt, *i.e.*, they are not biological in nature. Note that due to the smoothing of the growth curve data (described in section 4.2), the spikes shown in Fig. 7a are not mapped to the color orange (per the color map shown in Fig. 4) in Fig. 7b. The smoothing method is not used to eliminate spikes in the data because in fact, it is useful to be able to quickly identify which plates for a given PM experiment show spikes.

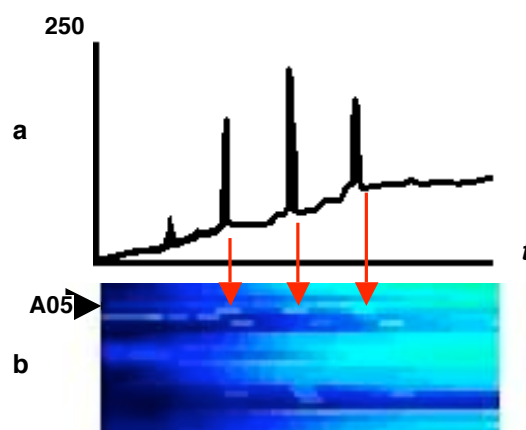


Figure 7. Example of using color image representation of growth curves to identify wells where camera readings spiked. (a) Plot of turbidity vs. time for well A05 showing three large spikes in signal. (b) Color image of rows A and B. Black arrow points to well A05 whose growth curve is shown in (a). The vertical scale for the color image in (b) has been increased to show more detail.

Figure 8b shows part of a color image for plate type PM20. The orange band at the bottom of the color image corresponds to wells B09-B12. These wells are pre-treated with tetrazolium violet, which reacts with the growth medium that the DvH cells were grown in. The result is that when the DvH inoculum is added to these wells, the wells darken immediately resulting in high readings of opacity. Wells such as these are excluded from further analysis because the high values are a result of a chemical reaction, not a biological process. Using the PMColorMap interface makes it easy to detect chemical reactions occurring between growth media and well pre-treatments.




Figure 8. Example of using color images to identify wells in which the growth medium reacted with the substrate resulting in high values. (a) Plot of turbidity vs. time for well B12. (b) Color image of rows A and B. Arrow points to well B12 whose growth curve is shown in (a).

The two color images shown in Fig. 9 are technical replicates (same inoculum) of data from the same plate type. Given the number of phenotype microarray experiments that have been run as part of the VIMSS project, our experience tells us that when comparing images such as those in Fig. 9, quite a number of wells on plate corresponding to the image on the right were not properly inoculated with cells so that the wells show little growth (indicated by a right brace). The PMColorMap interface thus provides a quick way to screen experiments to decide which ones may need to be re-run, and it also helps the experimentalist decide on which datasets and wells to focus.




Figure 9. Two technical replicates of the same PM plate. The color image on the left shows normal growth. The color image on the right is a technical replicate of that on the left. Comparison of the two color images suggests that rows A09 to C12, G03 to G08, and rows H08 to H12 were not properly inoculated.

Figure 10 provides an example of using the PMColorMap interface to compare the phenotypic responses of two different organisms. The figure shows four images for a single plate (PM15; mode of action - chemical inhibitors, chemical sensitivity). Two of the images (Fig. 10a and 10b) are technical replicates of wild-type DvH, and the other two (Fig. 10c and 10d) are technical replicates of a mutant of DvH constructed by knocking out the gene DVU3335, which is part of the *kdp* operon [8]. These pairs of images show that for this phenotype microarray experiment, the technical replicates reproduce the same behavior, *i.e.*, Fig. 10a is similar to Fig. 10b, and Fig. 10c is similar to Fig. 10d. Similarity of results from replicates, *i.e.*, reproducibility, gives us confidence when comparing the results of phenotypic responses.

Comparing the images for the wild-type DvH (Figs. 10a and 10b) with those for the mutant (Figs. 10c and 10d) shows that the growth of the mutant was more robust (showed more vigorous growth, *i.e.*, higher values) than that of the wild-type DvH suggesting that knocking out the DVU3335 gene conferred an advantage on the mutant in the presence of the chemicals on plate PM15.

In comparing the phenotypic responses of the wild-type DvH and the knockout mutant to a range of PM substrate, the PMColorMap interface played a key role in helping the experimentalists to decide which phenotype microarray experiments were of good quality, which experiments needed to be re-run, and which PM plate types showed interesting and significant differences between the responses of the wild-type DvH and the knockout mutant [8].

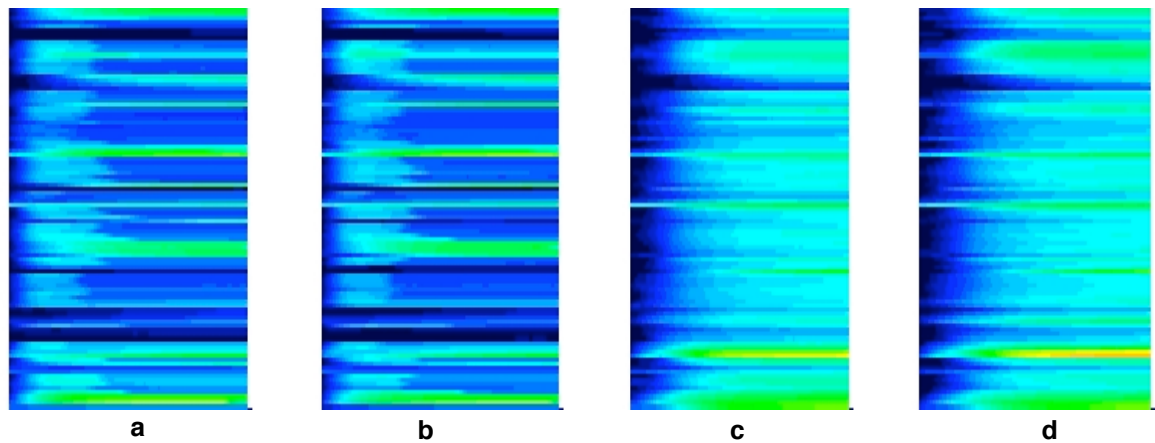


Figure 10. Using color images to compare the response of the wild-type DvH with a knockout mutant of DvH. Images a and b are technical replicates of wild-type DvH. Images c and d are technical replicates of the knockout mutant of DvH. Images a and b show good agreement in growth characteristics, as do c and d. For this PM plate, the growth of the knockout mutant (c and d) was more vigorous than that of the wild-type.

4.2. Technical Details

The PMRowView and PMColorMap interfaces are Perl-CGI scripts. The Perl module GD [14] is used both for drawing the growth curves on the fly and for creating the color images displayed by the PMColorMap interface. The color images are created when the PM datasets are uploaded to the VIMSS file server. In addition, information about the PM datasets (name, organism, date that the data were generated) is stored in a relational database. The information is accessed from the database when the color images are displayed or the growth curves drawn.

Though the PMRowView and PMColorMap interfaces are Web-based, downloading VIMSS phenotype microarray data is restricted to VIMSS project participants at this time.

Future Directions

The PMRowView and PMColorMap interfaces provide convenient and flexible ways to view and compare phenotype microarray data; however, due to the amount of data being generated, we plan to use additional computational methods to assess the quality of and to facilitate comparing datasets. In general, two approaches are used to characterize growth curves: growth curve models [14, 16] and calculation of growth indices [17]. We will evaluate the latter method first and will link the display of the growth curves with the calculated growth indices in order to be able to easily evaluate how well the growth indices reflect the actual growth.

Summary

Because of the large amount of growth data generated by phenotype microarray experiments, current ways of viewing these data are insufficient for displaying hundreds to thousands of growth curves at one time, and are far from ideal for sharing data and results between researchers working at different institutions.

We have solved the first problem by creating color images of all 96 growth curves corresponding to a single PM plate. This novel and space-efficient way of displaying growth curves has a significant advantage over traditional growth curve plots in that the overall quality of a phenotype microarray experiment and the reproducibility of technical or biological replicates can be quickly assessed by looking at small multiples of color images on a single Web page. The responses of different strains to the substrates on a single PM plate can be easily compared using the color image approach, and because the color images are linked to more detailed views of the growth curve data, the user can quickly confirm his/her assessments.

The PMRowView and PMColorMap interfaces solve the second problem of data access by providing easy and convenient access to thousands of growth curves and the raw data via a Web browser. The interfaces display the data differently, and at different levels of detail, in order to enable assessment of data quality, to facilitate different types of comparisons (biological replicate, technical replicates, different strains), and to highlight sought for, or unexpected, phenotypic responses.

Our tools, developed entirely using open source software, provide both new and traditional ways of looking at phenotype microarray data, use a relational database to store information about the data, and by providing data displays and raw data via the Internet, make it possible to easily share large amounts of information and data with project collaborators anywhere. The software development was guided by the needs and requests of the experimentalist generating the PM data. The experimentalists are using the software interfaces to assess data quality, compare replicates, and determine which datasets to focus their time and attention on.

Acknowledgements

We would like to thank the anonymous reviewers for their suggestions on how to improve this paper.

This work was supported by the Genomics:GTL Program, Office of Biological and Environmental Research, Office of Science, U.S. Department of Energy, through contract DE-AC02-05CH11231 with Lawrence Berkeley National Laboratory, Berkeley, California.

References

- [1] Biolog, Inc. URL: <http://www.biolog.com>
- [2] L. Zhou, X.-H. Lei, B.R. Bochner, and B.L. Wanner. Phenotype microarray analysis of *Esherichia coli* K-12 mutants with deletions of all two-component systems. *J. of Bacteriology*, **185**(16):4956-4972. August 2003.
- [3] K.D. Loh, P. Gyaneshwar, E.M. Papadimitriou, R. Fong, K.-S.Kim, R. Parales, Z. Zhou, W. Inwood, and S. Kustu. A previously undescribed pathway for pyrimidine catabolism. *PNAS*, **103**(13):5114-5119, March 28, 2006.
- [4] E.R. Tufte. *Envisioning Information*. Cheshire, CT. Graphics Press, 1990.
- [5] DOE Genomics:GTL Program URL: <http://genomicsgtl.energy.gov>
- [6] Virtual Institute for Microbial Stress and Survival (VIMSS) URL: <http://vimss.lbl.gov>
- [7] A. Mukhopadhyay, Z. He, E.J. Alm, AP. Arkin, E. Baidoo, S.C. Borglin, W. Chen, T.C. Hazen, Q. He, H.-Y. Holman, K. Huang, R. Huang, D.J. Joyner, N. Katz, M. Keller, P. Oeller, A. Redding, J. Sun, J. Wall, J. Wei, A. Yang, H.-C. Yen, J. Zhou, J., and J.D. Keasling. Salt stress in *Desulfovibrio vulgaris* Hildenborough: An integrated genomics approach. *J. of Bacteriology*, **188**(11):4068-4078. June 2006.
- [8] A. Mukhopadhyay, D.C. Joyner, J.S. Jacobsen, T.C. Hazen. Phenotypic characterization of a *Desulfovibrio vulgaris* high affinity K⁺ uptake mutant. *In preparation*.
- [9] S.E. Borglin, D.C. Joyner, J.S. Jacobsen, and T.C. Hazen. Anaerobic growth of *Desulfovibrio vulgaris* Hildenborough in microwell plates. *In preparation*.
- [10] Wolfram MathWorld Moving Average URL: <http://mathworld.wolfram.com/MovingAverage.html>
- [11] T. Saito, H.N. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, T. Kaseda. Two-tone pseudo coloring: compact visualization for one-dimensional data. INFOVIS 2005. *IEEE Symposium on Information Visualization*. October 23-25, 2005.
- [12] Colorblind Web Page filter URL: <http://colorfilter.wickline.org>
- [13] B. Shneiderman. The next generation of graphical user interfaces: information visualization and better window management. *Displays*, **17**(3,4):125-129. May 1997.
- [14] Comprehensive Perl Network URL: <http://www.cpan.org>
- [15] S. Lopez, M. Prieto, J. Dijkstra, M.S. Dhanoa, and J. France. Statistical evaluation of mathematical models for microbial growth. *Int. J. of Food Microbiol.*, **96**(3):289-300. November 15, 2004.
- [16] J.M. Ponciano, F.P.J. Vandecasteele, T.F. Hess, L.J. Forney, R.L. Crawford, and P. Joyce. Use of stochastic models to assess the effect of environmental factors on microbial growth. *Applied and Environmental Microbiology*, **71**(5): 2355-2364. May 2005.
- [17] J. Warringer, E. Ericson, L. Fernandez, O. Nerman, and A. Blomberg. High-resolution yeast phenomics resolves different physiological features in the saline response. *PNAS*, **100**(26):15724-15729. December 23, 2003.