

An Example of Dynamic Graphics for Manufacturing Process Data *

William F. Eddy

Audris Mockus

Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Revised September 8, 1992

Abstract

Dynamic graphical methods can be usefully applied to the analysis of continuous manufacturing process data. This paper describes some data collected during an ongoing project which addresses a particular manufacturing problem and describes a method developed to display the data. The nature of the manufacturing problem and the computational requirements of a dynamic graphical display prevented the more common interactive approach. We describe a graphical display we developed for these data and discuss the impact our display has had on the manufacturing process. In an appendix we briefly describe a system we have developed to display such data animations.

*Partially supported by ONR Contract N00014-91-J-1024, by Bureau of the Census under Joint Statistical Agreement 91-25, and by an anonymous industrial client. The equipment used to generate the videotape was acquired with partial support from NSF Grant DMS 88-05406. The authors thank Vijay Nair and two anonymous reviewers for their very helpful comments on earlier versions of this paper.

1 Introduction and Summary

An industrial client was interested in the spatial and temporal distribution of temperature across a metal plate. Process engineers believe that maintaining a constant temperature across this plate is critical to the quality of the product. In order to maintain a constant temperature the engineers have installed a feedback control system which regulates the temperature of the metal plate by regulating the flow of an electric current through the plate.

To assess the uniformity of the temperature distribution, to assess the adequacy of the control system, and to consider additional control systems and strategies the engineers designed a specially instrumented version of the metal plate. Data were collected from this special plate under various test conditions over time intervals that corresponded to a natural cycle in the production process. The data were analyzed by a variety of traditional statistical methods. These analyses showed that under the test conditions the temporal variation in temperature at some points on the plate was nearly forty times the variation at the control points. The spatial variation in temperature under the test conditions across the whole plate during the natural cycle of the process was nearly four hundred times the variation at the control points. In order to more clearly understand the spatial-temporal variation in temperature a graphical data animation method was developed to display the spatial and temporal distribution of temperature over the entire surface of the metal plate. We have created a videotape which displays the spatial-temporal distribution of temperature over the plate under several different test conditions.

In Section 2 we provide a brief overview of existing dynamic graphical methods. Some general characteristics of data gathered from processes such as the one we study together with the associated graphical needs are discussed in Section 3. In Section 4 we describe in some detail the particular manufacturing process we are working with and the means by which data were gathered from the process. We describe the method we used to fit a smooth function to the data and describe the dynamic graphical display we created in Section 5. In Section 6 we describe the resulting display for two specific example data sets. We describe the effect on the engineers and the process in Section 7. Finally, in an appendix we describe the system used to create the display.

2 Dynamic Graphical Methods

Graphical methods are widely utilized for the analysis of data and are invaluable for the presentation of results. Univariate and bivariate data are easily displayed. Effective display of higher dimensional data presents a challenge. Usually displays of higher dimensional data are presented by projecting into (or conditioning on) a two-dimensional subspace. Often, a sequence of subspaces are displayed creating a dynamic display. Generally, dynamic graphical displays provide an enormous increase in the amount of information conveyed to the viewer compared to static displays.

A basic problem with either dynamic or static displays of multidimensional data is the selection of which two-dimensional subspaces to use. One general technique which has been fairly widely used is the “Grand Tour,” see [1]; this method and other similar methods provides a smooth sequence of projections for viewing and thus lend themselves naturally to dynamic displays. A more recent study of projection methods which specifically considers the interactive approach to dynamic graphics is [6]. A basic drawback of projection methods relates to dimensionality; as the number of dimensions in the data increases the number of projections which need to be viewed to maintain the same minimum angle between projections increases exponentially.

An alternative approach to high-dimensional graphics based on parallel coordinate plots has been pioneered by [7]; see also, [8]. The basic idea is quite simple: instead of picturing coordinate axes as being orthogonal, display them as parallel. This has the great advantage that a quite large number of coordinates can be displayed simultaneously and the great disadvantage that the resulting pictures depend strongly on the ordering of the coordinates. Furthermore, since few users have knowledge of projective geometry, interpretation of observed structure in the data is often difficult.

A different approach is based on the scatterplot matrix. A collection of methods is considered in [2] – identification, deletion, linking, brushing, scaling, and dynamic parameter control – methods by which the user may interact with an array of two dimensional scatterplots. These methods, unlike the projection methods, do not change the original coordinate system. Since the number of scatterplots increases with the square of the dimension, this method becomes infeasible beyond a very small number of dimensions.

Techniques for the dynamic display of multidimensional data are of partic-

ular interest in statistical applications. The collection [4] edited by Cleveland and McGill contains descriptions of several interactive dynamic graphical methods. Unfortunately, all of those methods are only suitable for relatively low dimensional problems (say, up to 9). With many tens of variables and hundreds of observations (or even larger data sets), a new approach will be needed to avoid throwing away the bulk of the information for the sake of a low dimensional display.

3 Continuous Manufacturing Processes

Continuous manufacturing processes are typically controlled by a large number (hundreds) of automatic control systems. Consequently data is automatically collected on a large number of variables at regular intervals. To date very little effort has been made to simultaneously examine all (or even a reasonable fraction) of these variables for their effect on the typical production process.

Continuous production control is vital in the aluminum, steel, glass, and chemical industries, among others. For these applications, process control typically involves 20 to 200 control variables, and perhaps 10 to 100 quality (response) variables are recorded. The quality variables often form either individual time series or a set of repeated measures, and these features frequently apply to the control variables as well. The objective of the control is to monitor and adjust the process to ensure that the quality measures of the products meet the specifications of the customers. A basic difficulty in such control problems is that the available information relating control and quality variables is typically obtained during product manufacture, and thus is observational in character. Designed experiments are extremely difficult, due to the very large number of control variables and the practical difficulties of perturbing a large scale industrial process.

The large number of control variables means that any sort of statistical analysis is working at the very edge of what is possible. A typical analysis would be based on a nonparametric model (simply because the structure of such large data sets is unlikely to be clearly described by a parametric model). Nonparametric modeling consists of

1. *Data Smoothing.* Multivariate smoothing procedures for serially correlated data; and

2. *Data Modeling.* Fitting and evaluating competing adaptive (and usually additive) nonparametric or semiparametric high-dimensional models that incorporate variance stabilizing transformations, general low-order interaction terms, categorical data and prior knowledge on functional relationships.

Both of these steps are going to require the use of dynamic graphics. Here we utilize dynamic graphics for the first of the steps. Because of the huge computational requirements of the smoothing that is required in the particular problem we are considering we are led to the use of dynamic graphics that are not interactive. This point will be discussed further below.

4 Our Data

Because the industrial client has asked that the details of their production process remain confidential, that process will not be described in any detail here. The process generates product on a continuous basis and the product is collected into units for further processing and shipment to the customer. A unit is “filled” during approximately thirty minutes of production time. The crucial aspect of the process that we are focusing on here concerns the spatial and temporal distribution of temperature across a thin rectangular metal plate. The plate is oriented with its faces horizontal and the product flows vertically through a large number (thousands) of small holes in the plate and enters its final physical state. A diagram of the plate as seen from above is given in Figure 1. The dots in the diagram indicate the locations of thermocouples used for gathering test data and will be described further below.

Process engineers believe that it is essential to have the temperature maintained at a constant (in time) uniformly (in space) over the surface of the plate. The mass flow rate of the product is proportional to temperature (over a narrow range of operating temperatures). Consequently uniformity of the temperature (in time) increases uniformity of the product. Furthermore, uniformity (in space) of the temperature improves process efficiency. Of course, there is a temperature gradient over the plate because of the detailed physical structure that is a necessary aspect of the production process and because of boundary (of the plate) conditions. These boundary conditions are not only at the four edges but also the two faces of the plate. Additionally, the




Figure 1: Schematic of the Metal Plate Showing Locations of Thermocouples and Electrical Conductors

replacement of a production unit when full with a new empty unit introduces transient temperature effects on the metal plate. For temperatures outside the operating range, the process simply fails to produce the desired product.

The temperature of the plate is affected by three main factors:

1. the flow of product down through holes in the plate;
2. the (uncontrolled) flow of air across the external (bottom) surface of the plate from front to back;
3. the (controlled) flow of electric current horizontally from right to left through the plate.

The project described in this paper is a portion of one of the earlier parts of a larger long term effort to understand the contributions of each of these three factors to the spatial-temporal behavior of the temperature of the plate. Ultimately, one might hope to describe the spatial distribution of temperature across the plate by use of the heat equation (a partial differential equation) subject to the boundary conditions on the four edges and two faces of the plate. We are a long way from doing that successfully and are rather in a descriptive and mildly exploratory mode of study of the plate.

To achieve an approximately constant mass flow rate (and hence uniformity of the product) the present control system consists of two thermocouples measuring the temperature at points near the front edge of the plate. The two locations are indicated by large dots in Figure 1. The use of two

thermocouples is an improvement over the previous use of only one control thermocouple. The reason these locations are on the front of the plate is partly historical (that is where the single thermocouple had been located) and partly for convenience. Also, since the air flow is from front to back, changes in the air flow will be noticed more rapidly by the control system if the thermocouples are near the front. The left to right locations were chosen for geometric symmetry. The outputs of the two thermocouples are added by a process controller which raises the amperage of a 6 volt electric current through the plate (and, because of electrical resistance, the temperature) when the measured temperature falls (and lowers the amperage when the temperature rises). In Figure 1 the vertical bars at the right end of the plate indicate where the electric current enters the plate (identified with the letter I); the vertical bars at the left end indicate where it leaves the plate (identified with the letter O). The gap between the two bars at each end was an earlier attempt to make the temperature more uniform across the plate. When the bar had extended across the entire end of the plate there was a considerably hotter region in the middle of each end of the plate.

The process engineers believed that the control system was able to maintain the control points within plus or minus one-half degree of the desired temperature. Examination of the temperature measurements from those two thermocouples confirm this.

Process engineers were interested in developing a more detailed understanding of the temporal and spatial distribution of the temperature for several reasons. If the temperature distribution were better understood it should be possible:

1. to develop flow control fins to direct air at the “hot spots” on the plate;
2. to change the detailed structure of the plate to reduce variations in the resistivity (and hence temperature) of the plate; and
3. to develop a more complex electric control system to maintain a more uniform temperature.

To gather that information the process engineers designed a special metal plate with 72 thermocouples in addition to the control thermocouples (bringing the total to 74) at various locations on the plate. The particular locations were dictated by the structure of the plate and the conditions of the production process. Subject to these binding constraints, the locations were chosen

so that the distance from any point on the plate to the nearest thermocouple was small. The result was, unfortunately, not a regular pattern. Figure 1 uses small dots to indicate the locations on the plate of the 72 additional thermocouples.

The part of the project described in this paper was focusing on the effect of changing the current flow through portions of the metal plate. One interest was in determining the time period during which the temperature of the plate would complete its change in response to a change in the current (the *time constant* of the plate in engineering jargon). A related interest was in determining the behavior of the control system in response to these changes. The specific objective of the whole project was to develop a new control system which would provide more precise spatial and temporal uniformity of the temperature of the plate.

There are five short dark bars on the edges of the plate in Figure 1, two in the back identified with the letters A and B and three in the front identified with the letters C, D, and E. These bars indicate the locations of auxiliary current taps which were installed on this special plate for the specific purpose of altering the current flow through the plate. Various test conditions were created by connecting two of these taps together. This provided an alternate flow path for the electric current. Approximately 1% of the maximum available control current would follow one of these alternate paths when available. This small change in the current was sufficient to cause profound changes in the heat distribution across the plate.

Data were gathered automatically during continuous operation of the process under various test conditions from the 74 thermocouples (which included the two which were used to control the process). The temperature at each thermocouple was determined by averaging the output of a digital voltmeter over 100 powerline cycles (60 cycles per second). The result was stored in two bytes of binary data. The time interval between observations was determined by the overall capabilities of the data gathering equipment and remained constant at approximately 6 seconds throughout. Data were gathered for 300 time intervals. This time period corresponds to the production of a single unit of product. Note that $30 \text{ minutes} \times 60 \text{ seconds/minute} \doteq 300 \text{ intervals} \times 6 \text{ seconds/interval}$. This data gathering process was repeated a number of times under varying conditions resulting in several data sets.

There are $74 \times 300 = 22,000$ observations in each data set and at least as many ways to think about, visualize, and model these data. Although in




Figure 2: Time Series Plot of Output from Typical Thermocouple

one sense we have 74 variables, in another sense we have only three variables (two giving a position on the plate and the third the temperature at that location). Because of this special spatial aspect of these data we are not necessarily faced with the usual problems of very high-dimensional data (74 dimensions in this case).

The data sets were first examined using standard techniques. Time series plots were made for the data from each thermocouple. A typical such plot (from the A-D test described in Section 6) is given in Figure 2. The vertical scale is given (for confidentiality reasons) as an arbitrary linear transformation of temperature and the horizontal scale is given as the number of 6 second time intervals.

A two-way analysis of variance was calculated with the two factors being time and thermocouple number with a single observation in each cell. It quickly became apparent that the temperature on the plate remained constant neither in time nor in space, although, at the two locations where the controller thermocouples measured the temperature it did remain approximately constant. Several other standard statistical analyses were performed and graphical displays were generated which showed, roughly speaking, that the temperature at some locations on the plate under these test conditions varied over time by as much as plus or minus *twenty degrees* from a typical value at that location (this should be compared with the plus or minus one-half degree variation at the control points). These studies also showed

that over time within a data set the temperature across the plate could vary by as much as plus or minus *two hundred degrees* from nominal under these test conditions. The spatial variation was typically one order of magnitude greater than the temporal variation. This had an important influence on the particular dynamic display we created. These studies also showed that over a long period of time (corresponding to a number of separate data sets) the output of certain of the individual thermocouples became noisier and noisier. Eventually there was no information at all from certain thermocouples because the recorded values for the particular thermocouple were so erratic (or because the thermocouple ceased conducting current). When the output of an individual thermocouple became sufficiently degraded we ignored it in all future data sets.

5 Designing a Display

The process engineers were interested in viewing the temperature distribution across the plate as a function in time, corresponding exactly to the real physical situation. Consequently, we are interested in fitting a function $f(x, y, t)$ of the two spatial coordinates (x, y) and the one time coordinate (t) to the 74×300 observations and making a visual display of the function. There are three obvious approaches to fitting a smooth function of (x, y, t) to these data:

- Smoothing a piecewise constant function where each piece is some contiguous (x, y, t) region around a particular thermocouple at a particular time and the constant is the observed temperature;
- Estimating a smooth function from a nonparametric statistical model based on the three variables by, e.g., Loess [3] or MARS [5] (We did not seriously consider any parametric models.);
- Interpolating a smooth function to the data by kriging or some more direct technique.

Because of the irregular spacing of the thermocouples in space each of these types of methods will be computationally intensive. Additionally, because we wish to display the resulting function there will be substantial additional computation. Because the signal to noise ratio in these data is relatively high,

we believe that any reasonable fitting method will produce qualitatively the same display. We were led then to choose a fitting method which seemed computationally simplest. We chose a variant of kriging to be described below.

We decided to display the function of (x, y) at a grid of 120 pixels high by 460 pixels wide and then vary t . This size was roughly proportional to the actual size of the metal plate. (A pixel is a *picture element* and is the smallest individual component that can be displayed.) One reviewer pointed out that it might be very interesting to display the function of (x, t) and vary y or (y, t) and vary x ; we agree and intend to do so when the resources are available. We decided to record the dynamic display on videotape and we decided to use color to indicate the value of the function. The standard television signal (in the U.S.) has 30 frames per second. Consequently, if we wish to have a videotape that lasts two and one-half minutes (150 seconds) we will have to determine the color of $150 \times 30 \times 120 \times 460$ pixels from the 74×300 observations. The amount of calculation required for even a linear function from the observations to the pixels is astounding. The selection of 150 seconds as a target for the length of the dynamic display is based on our experience that:

- Anything less than about two minutes would have more than two time intervals per second and would complete much too rapidly for thoughtful viewing;
- Anything more than about four minutes would move too slowly and the viewer would tend to lose interest.

In some of our data sets there was a great deal of “action” and we later decided to slow them down a bit. In other of our data sets there was very little action and we decided to speed those up a bit.

5.1 Our Fitting Method

For computational simplicity we chose a variant of kriging. The basic technique of kriging can be illustrated using the following example. Let s stand for a generic point (x, y, t) and assume $f(s)$ to be a zero-mean stationary Gaussian process with covariance function $\gamma(h)$. The available data consist of the $f(s_i)$, $i = 1, \dots, n$. The object of interest is $f(s_0)$ for several values of

s_0 . The best linear unbiased predictor for $f(s)$ is $\hat{f}(s) = \sum_i a_i f(s_i)$. It can be obtained by minimizing

$$\text{var}(\hat{f}(s_0) - f(s_0)) = \mathbf{a}\Gamma\mathbf{a}^T - \gamma_0\mathbf{a}^T + \gamma(0)$$

where

$$\begin{aligned}\Gamma &= (\Gamma_{ij}), i = 1, \dots, n, j = 1, \dots, n, \\ \Gamma_{ij} &= \text{Cov}(f(s_i), f(s_j)) = \gamma(s_i - s_j), \\ \gamma_0 &= (\gamma(s_0 - s_1), \dots, \gamma(s_0 - s_n)),\end{aligned}$$

and

$$\mathbf{a} = (a_1, \dots, a_n),$$

subject to the constraint $\sum_{i=1}^n a_i = 1$. That is equivalent to solving a set of linear equations

$$\mathbf{a}\Gamma = \gamma_0 \tag{1}$$

under the same constraint. If a trend is present in the process it can be estimated first and then the above procedure could be applied to the residuals.

In our application we need to estimate the temperature over the surface of the plate and in time. Taking the size of the image to be $120 \times 460 = 55,200$ pixels, and having $n = 74 \times 300 = 22,200$ (74 points in space and 300 time points) observations, we have to solve the 22200×22200 -dimensional Equation (1) $55,200 \times 4500 = 248,400,000$ times (150 seconds \times 30 frames/second = 4500). This is an unrealistic amount of computation. However, substantial simplifications are possible.

- Because the human eye does a certain amount of smoothing in time (due to its psychophysical properties) we decided to fit a function $f(x, y)$ at each time $t = 1, \dots, 300$ and then interpolate each pixel linearly in time to the 14 frames intermediate between each consecutive pair of data frames (150 seconds \times 30 frames/second \div 300 intervals = 15 frames/interval). This reduces the dimension of Equation 1 to 74×74 and the number of solutions to 55,200. It adds the computational cost of the linear interpolation of intermediate frames. We believe that this is a reasonable way to proceed because the eye does not detect variation in time as easily as variation in space.

- We wished to avoid trying to estimate and invert a 74 by 74 covariance matrix for each pixel. Because the signal to noise ratio seemed quite high, we felt that many reasonable linear combinations of the observations would produce a good fit. Thus, we simply chose weights $\{a_i\}$ of the form $\exp(-\gamma \cdot d(T_i, P_j))$ where $d(\cdot, \cdot)$ is Euclidean distance in the plane, T_i represents the i th thermocouple, P_j represents the j th pixel, and γ is a single “smoothing” parameter. If we don’t count the cost of computing, storing, and recovering each of the $74 \times 55, 200 \approx 4$ million weights, then the amount of computation is about 8 million floating point operations per frame or roughly 2.5 billion floating point operations for the 300 “data” frames.
- To reduce this computational burden further we decided to fit the linear combination to a regular grid of “pseudothermocouples” and then interpolate (using a product of linear splines) from these pseudothermocouples to the pixels within each rectangle of the regular grid. We chose a regular $7 \times 25 (= 175)$ subgrid of the 120×460 pixel grid for the pseudothermocouples.

Even with all these simplifications there is still a substantial amount of calculation. For each of the 300 data frames, a weighted least squares estimator of the temperature was calculated at each of the 175 pseudothermocouples (from the temperatures of the 74 irregularly placed thermocouples). After some limited experimentation we chose the value of $\gamma = .040$. This strongly limits the effect (on the estimated value at a particular pseudothermocouple) of the observed values at any thermocouples much further away than the closest one. Next a product of linear splines was interpolated to the pixels within each rectangle of the regular grid of pseudothermocouples. Finally we interpolated linearly in time for each of the 14 intermediate video frames between each consecutive pair of “data” frames.

Recall that the spatial variation in the data is ten times as great as the temporal variation at any fixed point in space. This means that any dynamic display which does not somehow remove the spatial variation would not reveal very much (if any) of the temporal variation. We chose to combine two views of the plate: one static view showing a typical state and one dynamic view showing the deviations from that typical state.

An example of the actual graphical display is given in Figure 3. The details of the corresponding test conditions will be described in Section 6.




Figure 3: Time Interval 255 for C-O and A-D Test

There are three regions in the display. The lower of the two rectangular regions is static. It displays the median (across time) temperature. We applied our kriging variant to these medians. The locations of the thermocouples used for fitting are indicated in white and the locations of the electrodes and current taps are indicated in black. The locations of the thermocouples in Figure 3 can be compared to Figure 1 to observe which thermocouples have failed. The upper rectangular region displays the deviation from that median. We applied our kriging variant to each of the 300 sets of residuals. Thus, in the dynamic display the lower region remains constant and the upper region changes with time. The middle region represents a time series plot of the median (of the thermocouples at each time point). A tick mark indicates each 50 time intervals (roughly 5 minutes). A vertical bar moves horizontally through the time series plot. The position of the bar indicates the time location of the particular frame. A reference color scale is given at the very bottom of the display. The numbers below the scale refer to the lower region of the display. To preserve confidentiality they have been subjected to a linear transformation. The numbers above the scale refer to the residuals displayed in the upper region and are given in degrees.

6 Two Example Data Sets

We now consider two example data sets. Each data set contained several outliers, and a few of the thermocouples had much larger variability than the others even after we removed the data corresponding to failed thermocouples. To remove the obvious “spikes” in the data we ran the output of each thermocouple through a “running median of five” filter. The increased variability was an early indication of thermocouple failure. We used our judgment to decide when a thermocouple was sufficiently degraded to call it failed and removed it from all future data sets.

6.1 First Example

Our first example considers a data set collected during a test when a shunt was connected between auxiliary current taps A and D. At time interval number 50 auxiliary current taps A and D were connected together (see Figure 1) providing an alternate path for the electric current and at time

interval 200 the shunt was removed. Figure 4 displays the temperature at the moment when the shunt was connected. Thus, the sum of the upper and lower regions of the display can be considered a typical pattern for this metal plate. Remember, of course, that the lower region is the median over a time period including the test conditions so that neither of the two individual regions may actually display a typical pattern.

There are clearly visible small cooler regions in the lower static part of the display near the current taps labeled A, B, and C (and somewhat less visible regions around current taps D and E). This cooling is presumed to be caused by the heat conducted away by the current tap even when there is no electric current flowing through the tap. Also a typical hot region between the two legs of the input conductor and down the middle of the plate can be easily observed.

In the upper dynamic part of the figure the hotter region roughly between A and D is an artifact of our decomposition of the data. Notice that in the static part of the display this region is slightly cooler (because this is the region that was cooled by shunting the current). Because the overall (time) temperature was cooler, at this time the residuals appear hotter.

In the time series plot in the middle one can observe typical “ringing” behavior when the current is initially shunted via the auxiliary tap at time interval 50. This is followed by a fairly steady change to the new state until the shunt is removed at time interval 200. Again there is a slight ringing. The test ended before the entire plate had returned to standard conditions.

Figure 5 displays the temperature at the moment when the shunt was disconnected. The static part of the display is naturally the same as Figure 4. The dynamic part of the display clearly shows the cooler region caused by the shunt. The cooler region between the current taps is not exactly rectangular. This corresponds (we presume) to the region where the current has a reduced flow.

6.2 Second Example

Our second example consists of two figures in addition to Figure 3 from the C-O and A-D test. In this particular test a conductor was connected between the auxiliary current tap labeled C and the output conductor labeled O during the period from time interval 115 to the end of the test. Additionally a conductor was connected between tap D and tap C during the period from




Figure 4: Time Interval 50 for A-D Test




Figure 5: Time Interval 200 for A-D Test

time intervals 115 and 260. Figure 6 displays the temperature at the moment when both the shunts were connected. Figure 4 displays the temperature when the A-D shunt was disconnected. Figure 7 displays the temperature at the end of the test.

Remembering that the control system is able to hold the average temperature of the two control points approximately constant, from the time series plot in the middle region, one can see that initially the control system is able to respond adequately to the drop in temperature caused by the shunts. Then, however, the median temperature over the whole plate drops over a period of about five minutes to a new level.

Again, there are clearly visible small cooler regions in the static part of the display near the current taps labeled A, B, C, and E (and somewhat less visible region around tap D).

In the upper (dynamic) regions of the three figures, one can see the (dramatic) effect of the shunts. In Figure 6 there is no particularly cool region. In Figure 3, one can see two apparently hotter regions on the left near the locations of the current taps A and C. This is caused by the fact that over the whole test the median temperature in those regions is cooler than nearby areas (as can be seen in the lower static region of the figure). The electric shunts clearly explain why the region between O and D is cooler than the rest. In Figure 7 the region between A and D has recovered to the standard temperature but the region between C and O remains cool. These effects are visually striking on the videotape.

7 Discussion

Unfortunately, attempts to display dynamic graphics on the static page are doomed to failure, so we have not even tried to convey the real sense of what the displays associated with these particular data look like. Nevertheless, we are convinced that the main message in these data sets is conveyed by the dynamic displays. Furthermore, it seems unlikely that this message can be transmitted to the process engineers (or senior management) in any other way that would be as persuasive.

The effect of the dynamic graphics on the process engineers was immediate. First, they began many tests of various methods to control the air flow across the metal plate in order to reduce the spatial variation of the




Figure 6: Time Interval 115 for C-O and A-D Test




Figure 7: Time Interval for C-O and A-D Test




Figure 8: New Locations of Thermocouples for Process Control

temperature. Variations of these tests included the introduction of “fins” to redirect the air flow across the plate. These tests are continuing.

Second, they ran a number of tests to find the best locations for the small number of thermocouples that can be used in the actual production process. These tests resulted in a decision to utilize a control system based on six thermocouples. This number was based in part on cost considerations and also on the difficulty of installing thermocouples anywhere other than the front edge of the plate. Given the choice of six thermocouples, the locations were determined by (a) historical considerations, (b) geometric symmetry, and (c) the difficulty of locations other than the front edge mentioned above. Figure 8 shows the locations of the six thermocouples. The weights in the linear combination of the six thermocouples used to control the electric current have not yet been finally determined. By the geometric symmetry, for overall control the three in the front should receive equal weights as should the three in the center. The relative weightings of the center three to the edge three is the subject of further study. There are presently eight in-plant versions of the new control system under evaluation. The engineers have built a complex system to provide local control of the electric current across the plate. This allows independent control of the temperature by automatic shunting of current across three roughly equal disjoint regions of the plate. Tests are currently under way to evaluate various control strategies.

We have been working with the process engineers on this project for about three years. We expect these tests (and our visualization of their results) to continue for some time (years) into the future and we expect that within a

The frame buffer generates an RGB (Red, Green, Blue) television signal with NTSC compatible timing. NTSC composite video is the standard television signal in the United States, Canada, and Japan. The RGB video signal is sent to a Bosch NTSC color encoder (labeled NTSC in Figure 9); the encoder converts the RGB video into composite video.

Another essential part of the TV system is a Sony VO-5850 Umatic video tape recorder (VTR) which uses 3/4 in. videotapes (labeled UMATIC in Figure 9), and has special editing capabilities. A prerecorded time code on the videotape allows the Umatic VTR to locate individual video frames and is used to control the editing of individual frames onto the videotape in place of the black images. Animation (the dynamic graphical display) is created by editing successive video frames onto the tape, which are replayed at the rate of thirty frames per second. The control signals for editing are sent to the VTR by a Diaquest controller (labeled CTRL in Figure 9) which resides in an IBM personal computer (labeled PC in Figure 9). The PC also contains a vide sync generator (labeled SYNC in Figure 9). This signal is used to maintain synchronization of the video signal in the various separate components of the system.

The Diaquest controller receives instructions from a program running on the VAX workstation via a serial line to the PC. Such a program creates a digitized image, loads it into the video frame buffer that generates the RGB video signal, and issues an editing instruction to the control unit in the PC. The control unit translates this instruction into control signals that are sent to the Umatic VTR, which in turn edits a frame onto the videotape at the correct point. This cycle is repeated until all the frames comprising the animation sequence have been successively created and recorded on the videotape.

The video equipment is completed by a VHS video cassette recorder (labeled VHS in Figure 9), which is used to transfer animations created on Umatic videotapes onto conventional VHS videotapes, and a Trinitron color monitor (labeled TV MONITOR in Figure 9) that can display both RGB component video signal and NTSC composite video signal. Additionally, there is a 35mm film recorder (labeled SLIDE CAMERA) which allows accurate photographs of the VAX monitor.

The actual sequence of events during the recording begins when the program running on the VAX loads an image into the Peritek frame buffer. The program then sends an editing command to the Diaquest controller in the

PC which in turns sends the instructions to the Umatic VTR. The Umatic VTR starts to roll the tape. The Umatic VTR typically discovers that its current location is further down the tape than the requested location. It stops rolling and rewinds the tape to a point at least 90 frames prior to the requested location and then starts to roll the tape again. When the tape reaches the correct location the image is inserted onto the tape in place of the prerecorded (black) image and the tape stops. This entire cycle takes approximately twenty to thirty seconds to execute, assuming the image is delivered to the frame buffer instantly. Consequently, just the mechanics of recording one minute of dynamic graphics (1800 frames) takes approximately 15 hours. The software for rendering the images and controlling the VTR is designed so that the calculations associated with the next image are being executed on the VAX while the PC is controlling the recording of the current image. This overlap means that if an individual image can be rendered in less than thirty seconds there is essentially no delay due to the computation.

We have recently replaced much of this equipment with a new device, a Sony LVR-5000A laser video disk system. This equipment replaces the Bosch NTSC color encoder, the Diaquest controller, and the Umatic videotape recorder. The main benefit of the new equipment is that the time needed to find and record a single frame on a disk is negligible compared to the approximately thirty seconds required to find and record a single frame on tape. An additional benefit during playback is the ability to instantaneously jump between arbitrary frames (for comparison), an impossible task with tape. Also we can play back the recorded video at many different speeds (which is impossible with tape). We have also moved our calculations from the VAXstation to a DECstation 5000/240. This reduces the computation time by roughly a factor of thirty.

References

- [1] Asimov, D. (1985). The Grand Tour for Viewing Multidimensional Data, *SIAM J. Sci. Stat. Comput.*, **6**, 128-143.
- [2] Becker, R.A., Cleveland, W.S., and Wilks, A.R. (1987). Dynamic Graphics for Data Analysis (with discussion), *Statistical Science*, **2**, 355-395.

- [3] Cleveland, W. S. and Devlin, S. J. (1988). Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting, *J. Amer. Statist. Assoc.*, **83**, 596-610.
- [4] Cleveland, W. S. and McGill, M. E. (1988). *Dynamic Graphics for Statistics*, Wadsworth, Inc., Belmont, CA.
- [5] Friedman, J.H. (1991). Multivariate Adaptive Regression Splines, *The Annals of Statistics*, **19**, 1-141.
- [6] Hurley, C. and Buja, A. (1990). Analyzing High-dimensional Data with Motion Graphics, *SIAM J. Sci. Stat. Comput.*, **11**, 1193-1211.
- [7] Inselberg, A. (1985). The Plane with Parallel Coordinates, *The Visual Computer*, Special Issue on Computational Geometry, **1**, 69-91.
- [8] Wegman, E.J. (1990). Hyperdimensional Data Analysis Using Parallel Coordinates. *J. Amer. Statist. Assoc.*, **85**, 664-675.