

GT2009-60148

IMPROVED COMPRESSOR MAPS USING APPROXIMATE SOLUTIONS TO THE MOORE GREITZER MODEL

Chris Drummond

Institute for Information Technology
National Research Council Canada
Ottawa, Ontario, Canada, K1A 0R6
Email: Chris.Drummond@nrc-cnrc.gc.ca

Craig R. Davison

Institute for Aeronautical Research
National Research Council Canada
Ottawa, Ontario, Canada, K1A 0R6
Email: Craig.Davison@nrc-cnrc.gc.ca

ABSTRACT

Producing compressor maps is time consuming, costly and error prone and many data samples must be collected to give sufficient accuracy. Even then, expert input is typically required to fine tune the map to the appropriate shape. In this paper, we take some of that expertise and incorporate it in the smoothing process. The main piece of knowledge used is the cubic approximation for speed lines derived from the Moore Greitzer model. This well accepted approximation captures much of the general performance properties of compressors. But it is also widely recognized as only being very roughly true of real compressors. Nevertheless, we show that embedding this approximation, however limited, in the smoothing process results in accurate interpolation and extrapolation. The aim of this work is to substantially reduce the need for human input in the fitting process. We also anticipate a number of other benefits: less data is needed, with the commensurate time and money saved; the data collection process can be monitored for possible problems; changes in the map can be quantified and, when sufficiently small, data collection can be terminated.

1 Introduction

Producing an accurate compressor performance map is a crucial part of gas turbine engine design. A map is needed to characterize engine performance and work out different engine parameters such as specific fuel consumption. It also forms a vital part of the cycle deck, to provide an accurate simulation of

the complete engine. It is critical, therefore, to gather sufficient data to accurately represent compressor performance over the full range of shaft speeds and mass flows. Collecting data samples, and using these to produce a compressor map, is time consuming, costly and has inherent measurement uncertainty. However much data is collected, it cannot possibly cover the complete operating range. Inevitably some form of interpolation, and even extrapolation, is needed to give full coverage. Even then, expert input is typically required. Points must be tweaked, changing their values to fine tune the map to the appropriate shape.

Once the samples have been collected, there are many ways to carry out interpolation. One simple, and very popular, method is to connect adjacent points by a straight line. Unfortunately, this method may need a great many data points to obtain acceptable accuracy. Additionally, this method does nothing to remove noise created by measurement uncertainty. Different functions, such as polynomials, can be fitted to the sample points. Polynomial fitting is notoriously unstable, small perturbations can cause large changes in fit. A popular method that combines some of the desired smoothness of polynomials with the flexibility of linear interpolation uses different types of splines. Here, we investigate the use of splines with different sorts of penalty functions that control their smoothness.

In this paper, we take some technical knowledge and incorporate it in the smoothing process. The main piece of knowledge used is the cubic approximation for speed lines derived from the Moore Greitzer model. This well accepted approximation captures much of the general performance properties of compres-

sors. It has been used extensively, not only to explain the normal operation of an engine, but also to model the effects of compressor surge. Although it seems to capture much of compressor performance, it does so only at an abstract level. It is widely recognized as only being very roughly true of compressors that are used in real engines. Nevertheless, we show that embedding these approximations, however limited, in the smoothing process results in accurate interpolation and extrapolation. We also propose a corollary to provide efficiency contours at a similar level of accuracy.

The aim of this work is to substantially reduce the need for human input in the fitting process. We also anticipate a number of other benefits. Less data is needed, with the commensurate time and money saved. The data collection process can be monitored for possible problems, as reasonable approximations to the final map can be made with very little data and updated incrementally. An adequate number of data points can be determined. Changes in the map can be quantified due to each additional data point obtained; when the changes are within the acceptable error range data collection can be terminated.

2 Interpolation, Smoothing and Extrapolation

The aim of this section is to show how to incorporate expert insight into the interpolation-extrapolation process. There are many standard techniques that do not make use of such knowledge, each with their own advantages and disadvantages. Let us begin by looking at some of the most commonly used ones. We will continue by discussing the abstract knowledge about compressors used by our approach. We will then show how to incorporate this knowledge into the penalty functions of different types of spline.

2.1 Commonly used Interpolation Methods

The black dots in Figure 1 are samples along a speed line, representing a constant shaft speed. The x-axis is the corrected mass flow and the y-axis the pressure ratio. In all our figures for compressor maps, pressure ratio and mass flow are relative to the design point i.e. $PR_{rel} = (PR - 1)/(PR_{dp} - 1)$ and $W_{rel} = W_c/W_{c,dp}$, respectively.

One very simple, and very popular, method to interpolate between the points is piecewise linear interpolation. Here, adjacent points are connected by straight lines, see legend in Figure 1. With this type of interpolation, the result is far from smooth. There is a sudden change in direction at every data point. Yet, abrupt changes in the gradient of the line do not make much physical sense. So, particularly when the spacing of points is large, we would expect the accuracy of this type of interpolation to be poor.

We could use polynomial interpolation, which would result in a much smoother curve. Unfortunately, it has been known




Figure 1. Fitting a Speed Line

for more than a century [1] that it often gives wildly oscillating solutions, particularly when there are a large number of points. Figure 1 bears this out for this set of speed line sample points. Although passing perfectly through each data point, the resulting curve is not plausible and interpolates very poorly. The piecewise linear scheme is clearly better. If we could improve on the smoothness of the curve then it should not only be more theoretically plausible but also a better interpolator. Instead of connecting the points with straight lines, higher degree polynomials can be used. The dashed curve in Figure 1 has cubic polynomials connecting the points, i.e. a cubic spline, with knots at the points. This is produced by summing the basis functions ϕ , in this case b-splines [2], weighted by the w values, as shown in Equation 1. The c values indicate the center of the ϕ basis functions, the sample points in this example. Although this has eliminated the discontinuity in the gradient of the curve, it has done little to change the accuracy of interpolation.

$$f(x) = \sum_{k=1}^K w_k \phi(x - c_k) \quad (1)$$




Figure 2. Smoothing the Speed Line

2.2 Smoothing out Noise

Interpolation requires the curve to go through every point. Yet, we certainly would not be surprised if the readings that gave us some of the points were not exact. Errors can arise from many sources and our fitting procedure should, at least, be resistant to them. If we allow the curve to come close (in the least squares sense) but not exactly through the points, we should get a smoother fit.

Complete low degree polynomials, as opposed to piecewise ones, are very smooth functions which should remove much of the noise. Figure 2 shows a series of polynomials of lower degree than used for interpolation, so they no longer go through every point. Degree 6 has a much reduced oscillation, even though it is very close to the points, the black dots. For degrees 4 and 3, the oscillation is less, but at a sacrifice of some distance to the samples. In the area indicated by the dashed circle, the two polynomials of lower degree produce a more realistic curve than simply connecting the two points by a straight line. The 4th degree polynomial, in particular, goes very close to the points. Unfortunately, it still introduces this unintuitive oscillatory behavior which can become very exaggerated when extrapolating, as seen on the left of the figure.




Figure 3. Cubic approximation to Moore Greitzer model

2.3 The Technical Knowledge

In this paper, our aim is to trade-off flexibility and smoothness, using technical knowledge to define what is meant by smooth. This should give us physically plausible speed lines with good interpolation properties. The shape of a compressor map is constrained by the physics of the compressor. Moore and Greitzer produced a mathematical model of the compressor [3] which captures many of its important characteristics. Equation 2 is derived from this model and defines the approximate shape of any speed line on the compressor map. The basic equation defines a cubic as shown in Figure 3. By normalizing the speed lines (the pressure ratio is divided by N^2 , the mass flow by N where N is the rate of compressor rotation), a single cubic representing all speed lines is produced. This relationship is very approximate but when combined even with a small amount of data can produce accurate speed lines.

$$\frac{PR}{N^2} = C + H\left(1 + \frac{3}{2}\left(\frac{W_c}{WN} - 1\right) - \frac{1}{2}\left(\frac{W_c}{WN} - 1\right)^3\right) \quad (2)$$

It is important to note that there are three degrees of freedom in this equation: H the (semi) height, W the (semi) width and C

the intersection point with the y axis. A cubic polynomial has four degrees of freedom, the extra degree of freedom is nullified by requiring the gradient to be zero when the mass flow is zero.

Removing the normalization factors results in a set of speed lines, as would be seen in one commonly used visualization of a compressor map. Figure 4 shows speed lines for a range of normalized shaft speeds varying from 0.1 to 1.0 in 0.1 steps. Due to the dependence of the pressure ratio on the square of shaft speed, we get a quadratic, the dashed line in the figure, connecting the points of maximum pressure ratio on each speed line. Equation 2 can also be used to generate a reasonable approximation to the efficiency contours usually seen in the compressor map. We simply input the pressure ratio generated by the various cubics into Equation 4.

Equation 4 was developed from Equation 3, which is the standard definition for isentropic efficiency. The numerator is replaced with the well known ideal gas relationship for an isentropic process, to put the isentropic enthalpy rise in terms of the pressure ratio. To obtain a result for the denominator, we make the following assumptions:

1. All stages are similar
2. Inlet velocity to each stage is axial
3. Axial velocity from each stage is proportional to rotational speed

The validity of assumption 1 will depend on the design of the compressor; the validity of assumptions 2 and 3 will change with the operating point. However, the objective is to characterize the general shape of the efficiency contours not to obtain an accurate representation of any particular map. Applying these assumptions, along with Euler's pump equation summed across all the stages, results in Equation 4. This provides the isentropic efficiency in terms of pressure ratio, from Equation 2, and corrected shaft speed, provided as an input. The constant c , under the assumptions above, is a function of machine geometry and specific heat of the working fluid. It is chosen arbitrarily to produce efficiency contours that appear reasonable. The ratio of specific heats γ is set for dry air at standard conditions.

$$\eta_s = \frac{h_{02s} - h_{01}}{h_{02} - h_{01}} \quad (3)$$




Figure 4. Generating Multiple Speed Lines

$$\begin{aligned} \eta_s &= \frac{T_{01} c_p \left(\left(\frac{p_{02}}{p_{01}} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right)}{h_{02} - h_{01}} \\ &= \frac{\left(\frac{p_{02}}{p_{01}} \right)^{\frac{\gamma-1}{\gamma}} - 1}{c \left(\frac{N}{\sqrt{T_{01}}} \right)^2} \end{aligned} \quad (4)$$

Combining these equations produces a simple model which captures many of the abstract properties of a compressor, and thereby the abstract shape of the map. It captures the rollover of the speed lines, where pressure and mass flow decrease together, which leads to compressor surge. It captures the close connection between the speed line maxima and a line drawn along the ridge representing maximum efficiency. This line, often called the backbone of the compressor map, has been used before in the scaling of compressor maps [4, 5]. We do not expect, however, that the figure precisely captures any real compressor. The original cubic form is known to be useful but approximate. The equation for efficiency contains unrealistic assumptions, so is certainly not a precise model of all turbo-compressors.

2.4 Physics-based Smoothing

Interestingly, in our attempts to fit polynomials to the sample points along a speed line, a cubic gave a reasonable fit. From the last section, we now have a good idea why. Still, the cubic fit was far from perfect. We want to bias the fit towards a cubic polynomial but include some additional flexibility to allow for variations between speed lines and between compressors. A thin plate spline [6] biases the fit towards a straight line, but generalizes to other polynomials, including cubics. A further generalization, L-splines [7], uses different penalty functions based on linear differential operators of which one choice of operator gives rise to the thin plate spline.

There are many ways to realize this type of spline. One is to use the b-splines demonstrated earlier, but instead of simply minimizing the fit to the sample points, a penalty functional P , as shown in Equation 5, is included in the minimization. There is now a trade-off between the squared error of fit, how close the curve is to the points, and a penalty functional, measuring the deviation from our desired shape. The trade-off is controlled by the factor λ .

$$E(f) = \sum_i^n (y_i - f(x_i))^2 + \lambda P(f) \quad (5)$$

We use a two dimensional version of this type of spline to fit the efficiency surface from the samples. We do not have a closed form equation for the efficiency surface given by Equation 4 and shown in Figure 4 as contours. However, slices through the shape suggest that each dimension is approximately a quadratic. To allow greater flexibility in the backbone of the map, we do not penalize the cross terms. So, the integral in Equation 6 contains only 2 terms, the two cross terms, $\frac{\partial^3 f(\cdot)}{\partial x^2 \partial y}$ and $\frac{\partial^3 f(\cdot)}{\partial x \partial y^2}$, are missing.

$$E(f) = \sum_i^n (z_i - f(\cdot))^2 + \lambda \int \left(\left(\frac{\partial^3 f(\cdot)}{\partial x^3} \right)^2 + \left(\frac{\partial^3 f(\cdot)}{\partial y^3} \right)^2 \right) dx dy \quad (6)$$

For the speed lines, we use an alternative type of spline, based on a truncated power series. The two types splines are closely related, in fact we can transform from one to another easily by matrix multiplication [8]. This is useful to produce numerically stable fitting algorithms. The advantage of the truncated power series representation, as shown in Equation 7, is that it makes the constraints more explicit. The first summation is a polynomial of degree p , the second summation is over a set of one sided functions of the monomial of the same degree. These functions are zero for negative values and begin at the equally

spaced c_k values. We can penalize any coefficient in this equation. If we strongly penalize only the β values, the result will be very close to a polynomial, a cubic in our case. If we only weakly penalize the β values, we will get a piecewise polynomial function as we saw earlier in the paper.

$$f(x) = \sum_{k=0}^p \alpha_k x^k + \sum_{k=p}^K \beta_k (x - c_k)_+^p \quad (7)$$

However, we do not simply want to fit to a general cubic because Equation 2 is somewhat constrained, having only three coefficients. In addition, these three coefficients are shared by all speed lines. To restrain the cubic form, we penalize α_1 encouraging a gradient of zero when the mass flow is zero. To make all lines similar, we penalize deviations between the α_p and β_k of adjacent speed lines. Thus ensuring that the curves are based on similar cubics and with similar local smoothing. We end up fitting multiple splines to samples of multiple speed lines at the same time. This requires as many versions of Equation 7 as there are speed lines in the map.

One of the reasons that our fit for the cubic polynomial was not exact is that the simple model does not allow for compressor choke, which makes the right hand side of the speed line become close to vertical. Generally, fitting any function of x where the gradient is close to infinity is problematic. To overcome this limitation, we represent the speed lines in a parametric form. Instead of the pressure ratio being a function of mass flow, both pressure ratio and mass flow are functions of another value $(PR(n), Wc(n))$. This idea is very similar to other schemes for smoothing maps. Kurzke [9] has proposed many different ways of matching points on different speed lines. We use the one where n is the distance normal to the backbone. Where we differ is that we use a scaling based on the shaft speed, which is the effect of removing normalization of the x -axis. Another advantage of this approach is that it gives us a more flexible fit as we can allow this backbone to vary, capturing an important difference between compressor maps [4].

Fitting these splines now forms part of an overall procedure as follows:

1. Normalize pressure ratio and mass flow
2. Fit a two dimensional spline to the efficiency values
3. Locate the backbone of the efficiency surface
4. Find the Euclidean distance of each point to backbone
5. Scale distance according to normalized speed
6. Fit multiple one dimensional splines to $(PR(n), Wc(n))$ samples for all speed lines

All stages of the fitting procedure are carried using quadratic programming. Quadratic programming [10] allows the minimization of a quadratic function, in this case squared error in




Figure 5. High Pressure Compressor

fit, while enforcing inequality constraints. These constraints are used on the two dimensional splines to enforce a downward gradient at the edges of the map. They are used on the speed lines to enforce a strict order on the values for mass flow and to minimize the possibility of the crossing of the speed lines.

3 Investigating the Fit

We now apply the procedure of the previous section to two maps for two different types of compressor. Figure 5 is for a high power compressor [11], Figure 6 is for a fan, the core section [12]. The black dots are the data samples. The contour lines represent the efficiency of the compressor, according to the fit using a two dimensional spline with the appropriate penalty function. The gap between contour lines represents a difference of 0.02 in efficiency. The dashed line is the backbone of the map which passes through the sharpest parts of the contour lines. In the top left hand corner of each figure is a histogram of the absolute residual error in efficiency after the fit.

Each data point also has a black solid line through it, which is the same length as the error in efficiency at that point. Most of the errors are small and may well be due to noise in the measurements. They are less than the gaps between contour lines, so are hard to see except in the lower left hand corner of the map. Here, errors are due to inaccuracy of the fit. The errors,

however, tend to occur only for the lower shaft speeds, and far from the backbone. So, they do not effect the important central region of the map. The efficiency surface produced is smooth as indicated by the smooth contours. This makes more physical sense than the use of linear interpolation between corresponding points. The latter process would give abrupt changes in the surface, and therefore in the contour lines.

Based on distance to the backbone of individual points, the speed lines have been formed. One can see there is some degree of smoothing in terms of mass flow and pressure ratio, particularly for the lower values of speed. The irregularity of the points suggests that the smoothing is reasonable, only removing measurement uncertainty. The speed lines are also extrapolated. At the top of each line, this produces a natural looking curve, consistent with the general understanding of the shape of these lines. Still, a simple interpolation procedure, lines just connecting the sample points, would not have produced excessive errors. As we show below, it is when there are even fewer points that the technique presented in this paper is of most value.

For the fan core map of Figure 6, the errors are even smaller. None are greater than 0.015 and most are smaller than 0.01, half the distance between contour lines. The efficiency surface is also smooth, as indicated by the contour lines. What is of particular note, here, is that some of the speed lines (those labeled 0.875, 1.05, 1.1 in the figure) have only a single point. At speeds where there are more points, the lines exhibit the normal curved shape. As the coefficients are shared between the lines, even those with a single point follow this general shape.

In Figures 7, 8 and 9, we show the effect of reducing the number of sample points even further for the fan core map. Figure 7 shows the fit after 10% of the samples have been removed. There is a very tiny change in the map. The shape of the efficiency surface contours look very similar but between speed lines 0.5 and 0.7, the area surrounded by the innermost contour line has grown. However, there is little or no change in the position and shape of the speed lines.

If we remove a further 10% of the samples, the result is shown in Figure 8, there is now an appreciable change, albeit rather small. The efficiency surface has dropped at the top right hand corner, but little else is different. The speed lines are very similar in shape to those of Figure 7. Of particular note is the speed line for a value of 1.1, at the top right of the figure. The single point anchoring this line has been removed, yet the position and shape is only slightly changed. This can be seen more clearly in Figure 10 where the top two speed lines, 1.05 and 1.1, for the different number of samples are compared. Thus, extrapolation outside the given range of points has achieved a reasonably accurate speed line.

If we remove a further 10% of the points as shown in Figure 9 there is little further change in the efficiency contours. There is more of a change though in the speed lines. Their position is very similar to that in Figures 7, 8 but some of the curvature has been




Figure 6. The Fan Core




Figure 8. Removing 20% of Samples for Fan Core




Figure 7. Removing 10% of Samples for Fan Core




Figure 9. Removing 30% of Samples for Fan Core




Figure 10. The Top Two Lines for Different Samples

lost. There are no longer sufficient points on any single speed line to dictate the shape of other lines.

In this section, we have shown that introducing a bias based on technical knowledge about compressors has led to the effective smoothing out of measurement uncertainty. In addition, the shape of the efficiency contours and the position and shape of the speed lines was shown to be very robust to the number of samples. This should lead to a number benefits: less data should be needed, with the commensurate time and money saved; the data collection process can be monitored for possible problems; changes in the map can be quantified and when sufficiently small data collection can be terminated.

4 Limitations and Future Work

The main limitation of this work is that the factors, trading-off shape preferences and fit to data points, have been arrived at by manual tuning. This has typically only required working with a couple of maps, preferably quite different types, say a fan and a high pressure compressor. Once these values were selected, however, informal experiments showed they were effective for a large number of other maps. Nevertheless, there are ways of automatically assigning values in smoothing that could be very useful here. One commonly used approach is generalized cross-validation [13] which would be well worth exploring.

It might also be argued that the exact trade-off is best decided by an expert, so any system incorporating this approach should allow some input from the user.

The two dimensional spline generally produces a smooth and accurate fit to samples of compressor efficiency. It was less accurate, however, for low shaft speeds. The efficiency surface of the simple model shown in Figure 4 tends to become sharper at lower speeds. The present method of smoothing does not allow for this. This may produce physically questionable values close to the origin. Decreasing the amount of smoothing in this region should go a long way towards addressing this problem.

We have used a very simple model that ignores many additional factors for which there exist mathematical models. For example, Gravdahl and Egeland [14] give a model for compressor choke. It would be interesting to see how much more complex models might, if indeed at all, improve the fit. This might have particular benefit for high pressure compressors that choke at high speeds. At present, as long as a single point remains, we get good interpolation and extrapolation. But if there was no remaining point, the position of the line may be considerably in error. Using a model for choke might substantially reduce this.

We chose the present representation of compressor performance as it is the most commonly used one. However, we might also consider other non-dimensional representations of compressor performance. Experimental comparison would tell us whether or not they such representations would interpolate and extrapolate more effectively.

5 Conclusions

In this paper, we took technical expertise and incorporated it in the compressor map smoothing process. The main piece of knowledge used was the cubic approximation for speed lines derived from the Moore Greitzer model. This well accepted approximation captures much of the general performance properties of compressors, but is also widely recognized as being only a very rough approximation of real compressor performance. The important point to note is that even though this expertise was of a very general nature, when combined with a small number of samples from a real compressor, it was able to produce an accurate performance map.

Acknowledgments

We would like to thank our colleagues in the Institute of Aerospace Research and the Institute of Information Technology for many useful discussions and suggestions. We would particularly like to thank Jeff Bird, Sylvain Létourneau and Chusheng Yang.

REFERENCES

- [1] Runge, C., 1901. “Über empirische funktionen und die interpolation zwischen äquidistanten ordinaten”. *Zeitschrift für Mathematik und Physik*, pp. 224–243.
- [2] de Boor, C., 1978. *A Practical Guide to Splines*. Springer-Verlag, New York.
- [3] Moore, F., and Greitzer, E., 1986. “A theory of post-stall transients in axial compressor systems: Part i development of equations”. *Journal of Engineering for Gas Turbines and Power*, **108**, pp. 68–76.
- [4] Kurzke, J., and Riegler, C., 2000. “A new compressor map scaling procedure for preliminary conceptional design of gas turbines”. In Proceedings of the ASME International Gas Turbine And Aeroengine Congress And Exposition.
- [5] Sellers, J. F., and Daniele, C. J., 1975. DYNGEN-A program for calculating steady-state and transient performance of turbojet and turbofan engines. Tech. Rep. NASA-TN D-7901, NASA.
- [6] Duchon, J., 1977. “Splines minimizing rotation-invariant seminorms in Sobolev spaces”. In *Constructive Theory of Functions of Several Variables*, Vol. 571 of *Lecture Notes in Mathematics*. Springer, Berlin, pp. 85–100.
- [7] Ramsay, J. O., and Silverman, B. W., 1997. *Functional Data Analysis*. Springer-Verlag, New York.
- [8] Ruppert, D., Wand, M., and Carroll, R., 2003. *Semi-parametric regression*. Cambridge University Press, Cambridge, New York.
- [9] Kurzke, J., 2003. “Model based gas turbine parameter corrections”. In Proceedings of the ASME International Gas Turbine And Aeroengine Congress And Exposition.
- [10] Nocedal, J., and Wright, S. J., 2006. *Numerical Optimization*, 2nd ed. Springer-Verlag, Berlin, New York.
- [11] Tesch, W., and Pase, R., 1983. “Design and performance of a low aspect ratio, high tip speed, multi-stage axial compressor”. In Proceedings of the Nineteenth AIAA Joint Propulsion Conference, AIAA-83-1161.
- [12] Cornell, W., 1975. Health experimental quite engine program - summary report. Tech. Rep. NASA CR-2519, NASA.
- [13] Craven, P., and Wahba, G., 1979. “Smoothing noisy data with spline functions”. *Numerical Mathematics*, **31**, pp. 377–403.
- [14] Gravdahl, J. T., and Egeland, O., 1999. *Compressor Surge and Rotating Stall: Modeling and Control*. Springer Verlag, London.