2001-01-3004

Aerospace Manufacturing Cost Prediction from a Measure of Part Definition Information

Paul D. Collopy
DFM Consulting
Urbana, Illinois

and David J. H. Eames
Rolls-Royce
Indianapolis, Indiana

ABSTRACT

At Rolls-Royce, a study determined that manufacturing part cost can be estimated from dimensional data in the part definition. The method appears to be more accurate than mass-based parametric models or process simulation cost models. Additional advantages are: part definition information is causally related to part manufacturing cost; part definition information is easily available during the design phase; and mature part cost can be estimated even with new manufacturing processes.

The theoretical basis for the model uses Shannon's theory of information, work performed at MIT on calculating cost from information, and the results of the study at Rolls-Royce.

INTRODUCTION

Designers need an independent cost driver that is a true cause of cost and is available during the design phase. The driver should be causally related to cost, so that when the designer takes an action that increases cost, the action increases the driver and the driver increases cost.. Ideally, this cost driver would be easily deduced from the computer aided design part definition and would provide accurate cost estimation. In the best of all possible worlds, the estimate from this driver would be more accurate than manufacturing engineers' estimates, and therefore superior to manufacturing process simulator cost models.

This paper suggests that design information is a cost driver that satisfies all these criteria. Using an information-based cost model, a design engineer could receive from the computer aided design workstation instantaneous feedback on the cost impact of every design choice, every feature, every dimension, and every tolerance she enters. The paper also suggests potential follow-on work that could lead to a universal standard for capturing and displaying cost information to the designer.

Previous research [2] suggests that a simple linear estimate of process cost from quantity of information in design dimensions is more accurate than manufacturing engineering or shop cost estimation methods. The Rolls-Royce study suggests that a simple linear fit of complete part cost to quantity of information in the part definition is as accurate or more accurate than process-based models, even though the information-based model takes no account of the processes that will be used to manufacture the part other than a simple classification of part type.

A plausible reason for the success of the information metric as a predictor of cost is that information is a true link in the causal chain that leads to cost. We will see that, using the information metric:

- The addition of a feature increases the information metric and intuitively increases cost.
- Every facet, chamfer, counterbore, and so on that is added to a feature add dimensions to the definition and thereby increase the information metric. Each requires additional manufacturing operations that should normally increase cost.
- Every tightening of a tolerance increases the information metric and correspondingly increases cost.

Thus, in an intuitive way, the information metric quantifies the manufacturing cost burden caused by the part definition.

SIGNIFICANCE

Contemporary computer-aided design tools, such as CATIA, Unigraphics and Pro-Engineer, provide a virtual laboratory in which engineers can experiment with part definitions, leading to optimal design. However, this laboratory is critically flawed by its inability to assess the impact of design on manufacturing cost. Part definitions are represented as three dimensional solid bodies. Computer-aided design tools measure the volume of these bodies, thereby providing instantaneous mass estimates to the engineer. Finite element grids automatically generated within the body (see Figure 1) provide rapid assessment of strength, strain, and heat transfer properties. For parts that perform as aerodynamic flow boundaries, external grids are automatically generated and computational fluid dynamic models provide rapid assessment of flow capacities, drag, or aerodynamic efficiency. Three dimensional solid models can also be used for animated simulation of assembly operations. The results of all these analyses can be balanced in a formal optimization or in informal tradeoffs that effectively provide optimization. However, no design optimization can succeed without accounting for cost, and a correspondingly rapid analysis of part manufacturing cost is not available. Current costing techniques vary throughout the aerospace industry and include the use of both proprietary and nonproprietary methods. Most companies still retain a traditional cost estimating department that uses experienced individuals who access large proprietary databases.

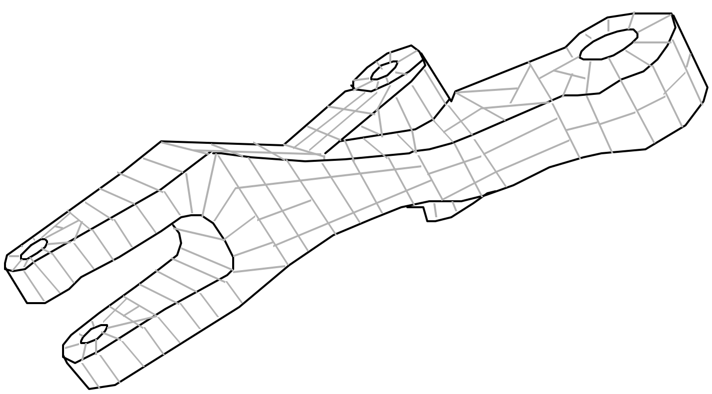

Figure 1: Automatically generated finite element grid

Cost estimates are generally available only after most tradeoffs are complete. Instead of being optimized in concert with other figures of merit, manufacturing cost is judged by a threshold: acceptable or not acceptable. Manufacturing costs are substantially higher under this approach than they would be if cost were optimized or balanced along with other figures of merit.

The work described in this paper is a first step in establishing whether part information correlates well with part cost. If so, a road map needs to be defined that leads us to a point where cost modules become standard features on all commercially available CAD systems.

STATE OF THE ART COSTING

Today's part cost estimating tools fall into two categories: manufacturing process planning tools and parametric cost estimators.

Manufacturing process planning tools generally require the design engineer to determine how each feature on the part will be machined and then, using parameters of the feature, estimate the cost of each manufacturing operation. In essence, they automate the analysis performed by manufacturing engineers when they plan the manufacturing processes for a part. The best of these models can quite accurately reproduce the cost estimates of manufacturing engineers. Today's state of the art, however, requires that the designer translate the design drawing into a parametric description of the part, requiring dozens to hundreds of parameters. The models cannot cope with new design processes—they are calibrated to the processes existing at the time the model was released. Also, manufacturing engineering cost estimates are generally far less accurate than estimates of other figures of merit, such as mass, flow, drag, or efficiency.

Parametric cost models statistically estimate part cost based on the correlation between historical cost data and part properties that are deemed to be cost drivers, such as mass and material. A positive statistical correlation between two properties, such as mass and cost, may exist because the first causes the second. Or the correlation may exist because the second property causes the first (see Figure 2). Or there may be some third factor that causes both mass and cost. In any case, a correlation is a valid basis for prediction. It is not, however, a sufficient argument for causation.

If A and B are correlated

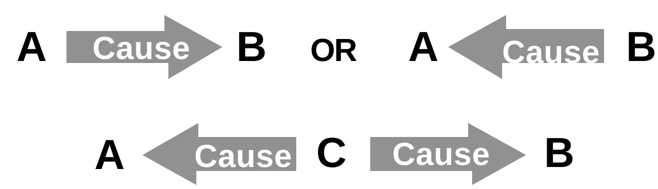

Figure 2: Correlation and causation

Price-H is an example of a popular aerospace cost model that uses mass as the primary cost driver and adjusts cost for material type. Mass is the only readily available measure of a wide variety of aerospace parts, so it is generally the principal cost driver by default. However, it is doubtful that mass actually *causes* cost to a major extent. Causation is unimportant when the model is used simply to predict cost. It is, however, a critical concern when engineers use the model to optimize part designs. The model will induce the engineer to reduce mass in order to reduce cost. Aerospace designers should immediately see the fallacy here: we normally incur *higher* manufacturing cost as we reduce mass by strategies such as more precise machining, scalloping out mechanically unloaded material, or employing composite parts. In design space excursions in the vicinity of an optimum design, reduced mass typically corresponds to higher cost, not lower cost. Thus, engineers who “drive” cost by turning the mass knob on a cost model are liable to bear off in a very wrong direction. Another serious deficiency of mass as a primary cost driver in engineering design is that, in most systems or products, mass is already a key figure of merit on which parts are assessed. When cost is predicted from mass, and cost and mass are input to an optimization algorithm, the effect is as if mass were counted twice and cost were ignored.

Material is also problematic as a cost driver. By *material* we mean the substance that constitutes a part, such as aluminum, stainless steel, Inconel 718, or graphite-reinforced polymeric composite. Our normal notion is that material is an important cost driver, because the cost of the material is a substantial portion of the cost of the part. However, for aerospace parts, this is not true. Consider the following:

> In 1996, according to The Boeing Company website, the list price of the Boeing 777-200 was $110 million. The empty weight of the aircraft is 314,300 lbs. [1]. The price per pound is therefore about $350. Titanium is one of the more expensive materials used in the aircraft. The cost of raw titanium (sponge) varies widely, but a typical value is $5 per pound [4]. Thus, if the whole aircraft were made of titanium, raw material cost would be less than 1.5% of price, easily less than 2% of manufacturing cost. Because titanium is more expensive than the most commonly used materials (aluminum and steel), this is an upper limit.

Thus, aerospace materials do **not** cause cost because the raw materials themselves are expensive (except in rare and unimportant cases). Materials cause cost because they impact processing costs. Therefore, to estimate the cost impact of a particular material selection, it is essential to know the processing steps for the part. For example, titanium is only somewhat more expensive to turn than aluminum, but it is much more expensive to cast: purity must be extremely high due to the severe consequences of microscopic inclusions, particularly oxidized titanium, and an alpha-phase casing usually forms which must be removed by another process, such as etching. Titanium is also expensive to weld because it rapidly oxidizes, and even burns. The cost impact of selecting titanium versus aluminum is therefore critically dependent on whether the part is turned, cast, or welded. Often the manufacture of a single part involves all these processes, and others, so the extent of each process would need to be known to properly assess the impact of material selection on cost. The point is that a single correction factor to adjust for material type can only provide a small improvement to part cost.

An even more difficult challenge is distinguishing materials at a level appropriate to estimating cost. A model may classify PH17-4 Stainless Steel, Inconel 718, Waspaloy and Udimet powdered metal alloys all as *steels*, although the manufacturing properties of these alloys are extremely different. Composites present an even graver problem. *Composite* is not a material type, it is a combination of materials. In many cases, the combining of the base materials is the most expensive element of a composite part’s manufacture. The number of different types of composites vastly outnumber the number of materials—if there are n materials, there can be $2^n - n - 1$ different combinations of two or more materials, each of which could yield a family of composites. It seems unrealistic to hope to classify the potentially astronomical number of composites finely enough to create a variable (composite type) that contributes in an important way to cost estimation while retaining a sufficiently large population in each class to perform a reliable cost regression.

In summary, existing cost models are deficient for use by design engineers. Parametric models use weight as a principal driver although weight is not an important cause of cost. Such models direct the design engineer that all efforts to reduce weight will be rewarded by lower cost, when generally the opposite is true. Secondarily, the models use material as a driver. Material is an important cause of cost, but the impact depends intimately on the manufacturing process steps, which are not available to the parametric cost model. Such information **is** available to cost models based on manufacturing planning, but such models are generally difficult to use, requiring a level of definition sufficient for manufacturing planning before any costs are available.

A BETTER WAY

This paper suggests that design information may be a more powerful, causal cost driver than weight or material and that cost models based on design information will be superior to contemporary models, particularly for use by design engineers. To understand the notion of design information as a cost driver will require some background in information theory and how to think about design definition as a communication channel

BACKGROUND

This section explores the design drawing or product definition as a communication channel from design to manufacturing. Shannon’s theory of information is used as a basis for quantifying the information content of messages in a communication channel. Prior work by Hoult and Muter is recapitulated, demonstrating the relationship between manufacturing cost and quantity of information on a design drawing.

COMMUNICATION DESIGN TO MANUFACTURING

Think about design as a planning process. Product designers plan the manufacture, operation and maintenance of a product. These plans are the essential output of design. The most complex plan is generally the plan for manufacturing the product.

For a machined part, the traditional representation of this plan is a design drawing, consisting of one or more sheets of paper. Three decades ago, these drawings were created by hand in pencil on a mylar master and duplicated onto paper using a diazo printing process. Fifteen years ago, the manual drawing process was imitated on computers where digital masters were stored and copies were printed or plotted on paper. Today, the plan for manufacturing a product is not a drawing, but a digital file, the *product definition*, which may never be committed to paper. Computer-aided design workstations display the part definition in three-dimensional solid representations.

For our purposes, all these representations are similar. Each consists of four classes of elements:

1. A broad visual representation of the part, whether plan, cross-section, or three-dimensional. Particularly with computerization, these elements have become simpler and convey relatively little information in modern drawings.
2. Standard features located on the visual representation, such as surfaces, holes and chamfers.
3. Feature parameters, primarily toleranced dimensions. For a hole, the parameters include the location of a hole (two dimensions in older drawing, up to five dimensions in a solid model), depth (one dimension), diameter (one dimension), and possible countersink and thread specifications and dimensions. A feature may also include a *count* such as “4 Places” indicating that a quantity of congruent features are required, differing only by location. Only one specification is made, but some pattern or algorithm must be specified (such as “equally spaced”) for determining the location of the remaining features.
4. Textual notes, such as descriptions of test processes. These processes may be parameterized. If so, the parameters along with tolerances on the parameters are generally included in the notes. For example, heat treatment of a part may be specified in a note. The specification should include the series of temperatures to hold at each temperature or to transition between temperatures. Each temperature and each time interval should properly be toleranced.

In our study we primarily relied on feature dimensions and their explicit or implicit tolerances. We also used feature counts and, to some extent, numerical data in textual notes.

QUANTITY OF INFORMATION

Quantity of Information as a property of communication is a concept classically presented by Shannon [3]. Shannon conceived of communication in its most fundamental form[1] as the conveyance from a source to a receiver of a string of symbols. The quantity of information in the string is measured by its *entropy*.

Entropy can broadly be defined as the quantity of microscopic possibilities that can exist given a macroscopic specification of a system. For example, the entropy of a vessel of neon gas is the number of possible microscopic states, specified by the velocity of every neon atom, for a given macroscopic state, specified by temperature and pressure. This number is hyperastronomical. For this and other reasons of convenience, entropy is represented as a value proportional to the logarithm of the number of microscopic states.

The entropy of a string of symbols (the quantity of information in a communication) is similarly defined. The macroscopic specification incorporates everything the receiver knows about the string before examining the symbols. Consider, for example, an engineering drawing communicated from a design organization to a manufacturing organization (a factory). The factory knows that the drawing will specify a part. The factory knows that the drawing will conform to conventions regarding engineering drawings or even to a particular standard that addresses the depiction of views and features and the labeling of dimensions and tolerances. The factory knows the language used for notes (English or German, for example) and the numeral set and conventions used to represent numbers (Arabic numerals, decimal notation). All these elements are part of the macroscopic specification of the class *drawing*. Given this macroscopic specification, the quantity of information in a drawing is broadly the number of possible distinct drawings that conform to the specification. As in the thermodynamics case, the quantity of information is represented as the logarithm of the number of possible drawings. For correspondence with the capacity of digital media, this logarithm is normally taken to the base two and the units of the quantity of information are *bits*, or binary digits. Shannon [3] proved, for simple cases, and asserted, for general cases, that the quantity of information in a message is the bit length of the shortest (most compressed) binary encoding of the message decodable by the receiver, using all of the receiver’s prior knowledge about the message.

PREVIOUS RESEARCH

Hoult and Muter [2] demonstrate that for a particular machining process, the time to process a part (and, by

[1] a discrete noiseless channel

inference, the cost) is proportional to the information content of the process definition on the part drawing.

The method that Hoult and Muter used for measuring information is

$$I = \sum_{k=1}^{N} Log_2\left(\frac{dimension_k}{tolerance_k}\right) \qquad (1)$$

where N is the number of dimensions that describe the process.

This means that, for every dimension used to describe the process, the dimension is divided by its respective tolerance. The result is a normalized dimension where the tolerance is the base. The logarithm determines the number of bits required to represent the normalized dimension. The summation yields the log of the product of all the normalized dimensions or, alternatively, the total number of bits necessary to describe all the normalized dimensions needed to define the process.

This relates to information theory and to entropy as follows. The macroscopic definition of the process is the process definition with all dimensions located but not stated. For each dimension, the tolerance is known and the dimension is expressed as a multiple of the tolerance. The microscopic definition includes the actual values of the tolerances. The entropy, I, is the logarithm of all the possible combinations of dimensions that could be in the definition. Information theory also views the definition, without the explicit dimensions, as a communications channel. The explicit values of the dimensions (normalized to their respective tolerances) combine to form a message. The information metric is the length of the message in bits, which is also equal to I.

Hoult and Muter gathered data on three processes: milling, turning, and laying up composite material. Across eleven data sets they observed an average r^2 of 0.8 for a root-mean-squared correlation coefficient of 0.9. In all cases where shop cost estimates by professional estimators were available for comparison, a simple estimation from the single design information metric had less error in predicting actual manufacturing time than the shop cost estimate.

To restate this in a way that may be more clear, Hoult and Muter demonstrated that, with no data about the part except one number, the quantity of design information, the cost of a single manufacturing process could be more accurately estimated than with all the design data and knowledge of the shop using traditional methods.

At Rolls-Royce we investigated whether this result could be extended to estimation of the total cost of a part (labor and material) where several manufacturing processes may be required.

PROCEDURE

We first separated parts into thirty part types. In the course of the study we learned that for some of these part types, such as electrical equipment, sensors and pumps, this method does not directly apply because the part drawing only describes the shape of the part. A specification is used to actually define the part. The exterior drawing has little to do with the part cost. (From the outside, an engine control and a voltage regulator may be very similar.) Thus, we ended with the twenty-two part types listed in Table 1.

For each part type we attempted to identify ten to twenty parts (part numbers) that were distinctly different and for which we had a drawing and a current actual manufacturing cost. For some types we were not able to find ten such parts. For others, parts were abundant and we included more than twenty. On average, we examined thirteen parts within each type.

For each part, we quantified the information on the drawing.

1. All dimensions were quantified according to equation (1). For dimensions without tolerances, such as reference dimensions, the drawing's default tolerance was used. (We debated whether to include reference dimensions, as they are not true dimensions. However, in many cases they were used in such a way as to direct processing operations, so that they caused cost like normal dimensions. Therefore we counted them.)
2. Indications of repetitions (such as "in 4 places") were counted as dimensions with a tolerance of one. That is, the bit string length of the count was used.
3. Numbers in textual notes were counted if they included a tolerance ("Heat treat for 4 hours plus or minus 30 minutes") or if they were obviously an integer, in which case a tolerance of one was used.
4. For drawings that included assembly instructions, one bit was added for every part number listed for assembly.

Due to the limited scope of this study, this algorithm was never tested against alternative approaches to measuring information, even as simple as different weightings of the measures listed above. It is very probable that there is an alternative assessment method that would provide better correlation and more accurate prediction of cost.

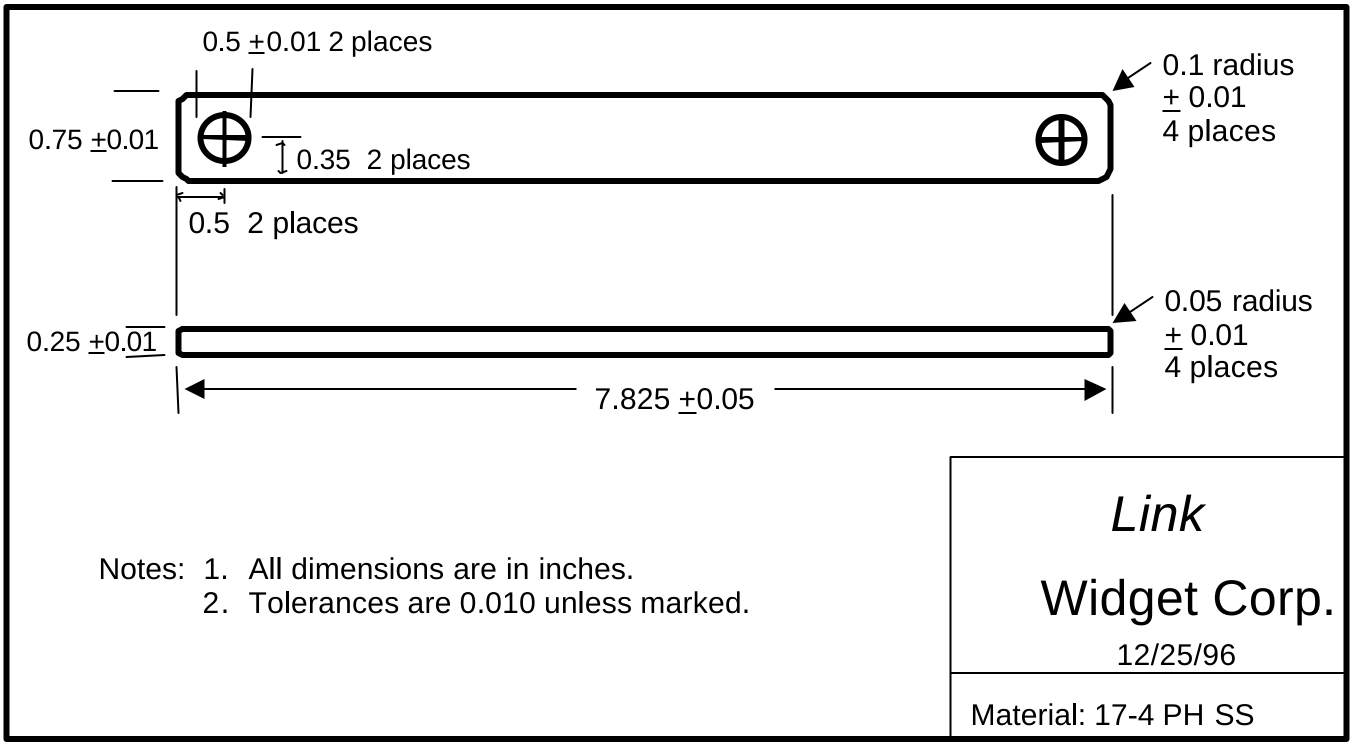

Figure 3: Example part definition

All the dimensions, tolerances, and so on, were entered into Excel spreadsheets, one sheet per part number, and the spreadsheets were stored in an Access database. Most drawings had a few hundred dimensions. Some had as many as one thousand.

EXAMPLE CALCULATION

To make this method more clear, an example is provided for a very simple part (Figure 3).

Dimension	Tolerance	Count	Bit String
0.5	0.01		5.65
		2	1
0.75	0.01		6.23
0.25	0.01		4.64
0.1	0.01	4	5.32
0.05	0.01	4	4.32
7.825	0.05		7.29
0.5	0.01	2	6.65
0.35	0.01	2	6.12
Quantity of Information			47.22

Table 1: Calculation of design information

There are eight dimensions on the drawing. Consider first the diameter of the round hole, which is 0.5". The tolerance is 0.01". The ratio of dimension to tolerance is 0.5 / 0.01 = 50. Thus, the dimension is 50, normalized to its tolerance. The base 2 logarithm of 50 is 5.65. That is, it takes 6 bits to represent the number 50 in binary, but the first bit is not really fully used. Encoded in a long string with other data, the number would only require 5.65 bits. This is summarized in the first line of Table 1.

This dimension applies to two places in the drawing (both of the holes). Two, being a count, does not need normalization. The bit string length associated with two is one (the base 2 logarithm of two). This is reflected in the second line of Table 1. Later lines combine the dimension and count into a single bit string length for convenience. When all the dimensions and counts have been assessed, the bit string lengths are totaled in the right hand column to about 47, which is the quantity of design information in the drawing.

DATA

Figure 4 displays all the data that were collected which, because of its proprietary nature, has been undimensionalized for this paper. Because the data cover a wide range, they are broken into three groups, each of which is plotted on a different range. Relative ranges are shown on the axes. The shadows of the smaller plots are indicated on the larger plots.

Statistics on each part type are summarized in Table 2. r^2 is the square of the correlation coefficient, and suggests the fraction of the cost which is caused by design information.

INTERPRETATION

WEAKNESSES OF THE DATA

The data compiled in Figure 4 and Table 2 should be viewed with the following considerations in mind:

1. As mentioned previously, the algorithm chosen for quantifying design information was a first guess, and it is reasonable to expect that it could be significantly improved with further research.
2. The part type distinctions were made prior to gathering the data. It seems likely that the data may suggest better groupings. For example, shafts and gears cluster around the same regression line and could be treated as a single group. (Notice however that some groups were effectively split, usually by material, after analyzing the data. For example, disks were divided into titanium disks and other disks.)
3. The blisks group and perhaps the tip tracks were too small to make the correlation data meaningful. In fact, none of the groups with less than twenty parts are very compelling.
4. There was no attempt to evaluate any aspect of the design definition other than numbers that appear on the drawing. A more advanced approach might be able to use more data in its evaluation.
5. Almost all the regression fits are linear. With more data there would be opportunity to consider some second order fits that may improve accuracy. Note though that Hoult and Muter [2] proposed an underlying theory that claims that fits should be linear.

ACCURACY OF INFORMATION BASED COST MODEL

The average error values shown in Table 2 suggest that a cost model based on design information would be quite competitive with other cost models in terms of predictive accuracy. Moreover, when the cost weighted mean average error is calculated, the result (22%) is really quite excellent compared to the state of the art. This means that if the cost of a whole system is being estimated, the result using this approach will be very accurate, possibly superior to any alternative.

It is intriguing that the highest cost parts work best with this model. These are the parts that are most distinctly aerospace parts (airfoils, disks, cases and frames). Possibly, the difference is because these parts are least likely to share manufacturing learning with other programs. It is much more likely for two aircraft engines to share a common bearing design than a common turbine disk design, particularly if the engines are made by separate companies. This suggests that the inaccuracy of the lower cost part types may be partly due to learning that occurs off-program and is not recognized. This is especially true for fasteners and seals, where cost is mostly determined by sales volume, not design.

Part Type	Number of Data Points	Correlation Coefficient	r^2	Average Error	Notes
Assembly	8	0.95	0.90	29%	
Bearing	10	0.65	0.43	48%	
Blisks	3	1.00	1.00	0%	
Bracket or Cover	10	0.72	0.52	56%	
Cases	12	0.99	0.98	13%	(1)
Compressor Blade	17	0.99	0.98	18%	
Compressor Vane	19	0.82	0.66	11%	(3)
Disk	20	0.89	0.80	15%	(2)
Fairing	16	0.96	0.92	38%	
Fastener	25	0.20	0.04	104%	
Gear	7	0.55	0.30	32%	
Housing or Frame	17	0.94	0.88	29%	(4)
Liner	3	0.98	0.96	9%	
Retainer	23	0.78	0.60	59%	
Ring	11	0.70	0.48	62%	
Seal	18	0.81	0.66	70%	
Shaft	22	0.85	0.72	41%	
Tip Track	5	0.21	0.05	14%	
Tube	13	0.72	0.52	67%	
Turbine Blade	6	0.98	0.95	13%	
Turbine Vane	8	0.31	0.10	15%	(3)
Wire	10	0.59	0.35	52%	
SUM	283				
MEAN	13	0.75	0.63	36%	
STD DEV	6.6	0.25	0.32	26%	
Mean weighed by cost		0.87	0.79	22%	

(1) Corrected for Composite vs. non-composite
(2) Corrected for Titanium vs. non-titanium
(3) Corrected for Number of Vanes
(4) Corrected for Nickel Alloy or Titanium vs. other material

Table 2: Summary statistics on cost estimation

IMPORTANCE OF PART TYPES

Notice that the collective data in the plots (Figure 4) show a strong positive correlation but much less than the average correlation of regressions within part types (Table 2). Thus, the model could not be so successful if the parts were not separated by part type. Why does this work so well?

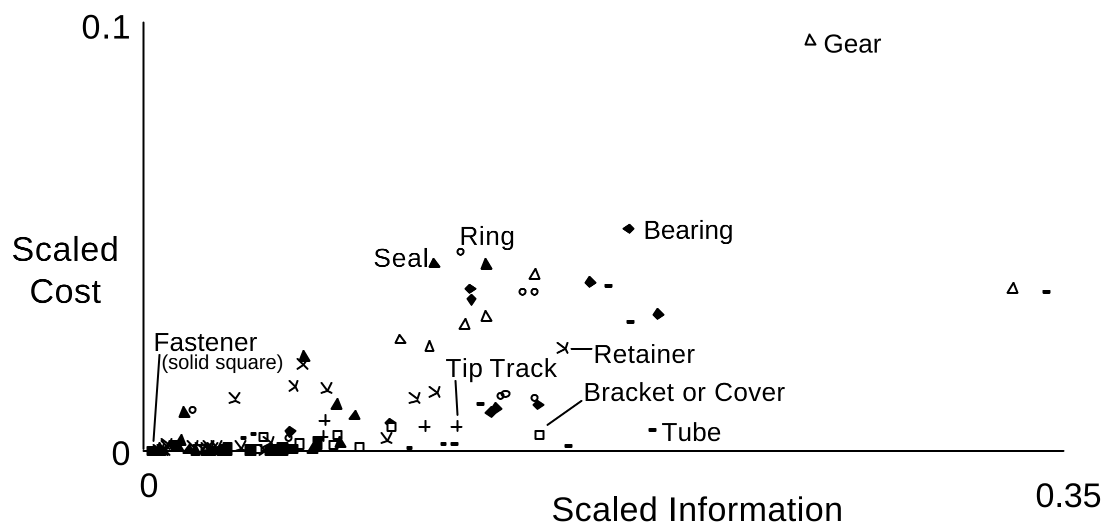

Figure 4a: Data for low cost parts

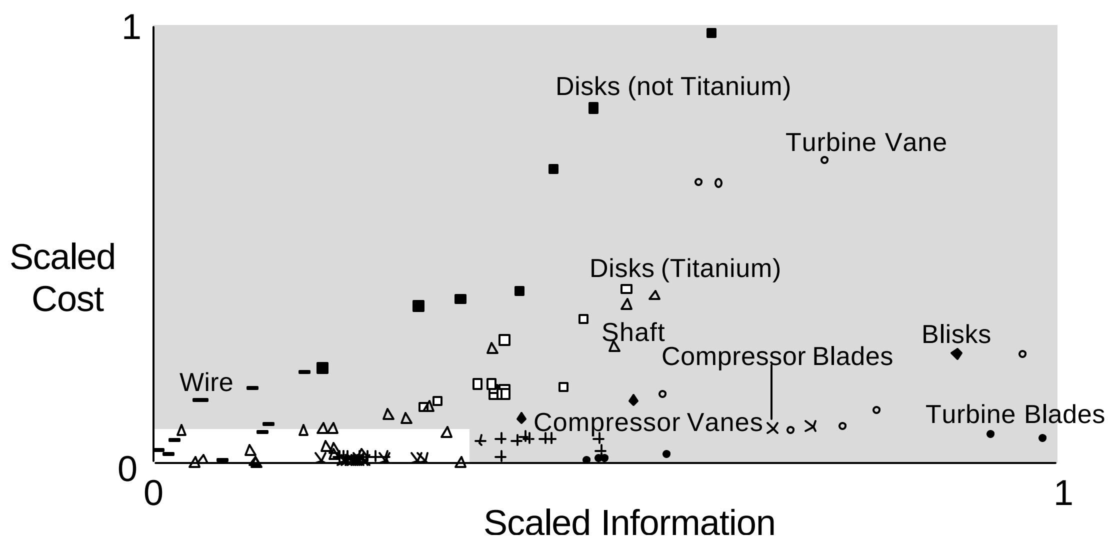

Figure 4b: Data for medium cost parts

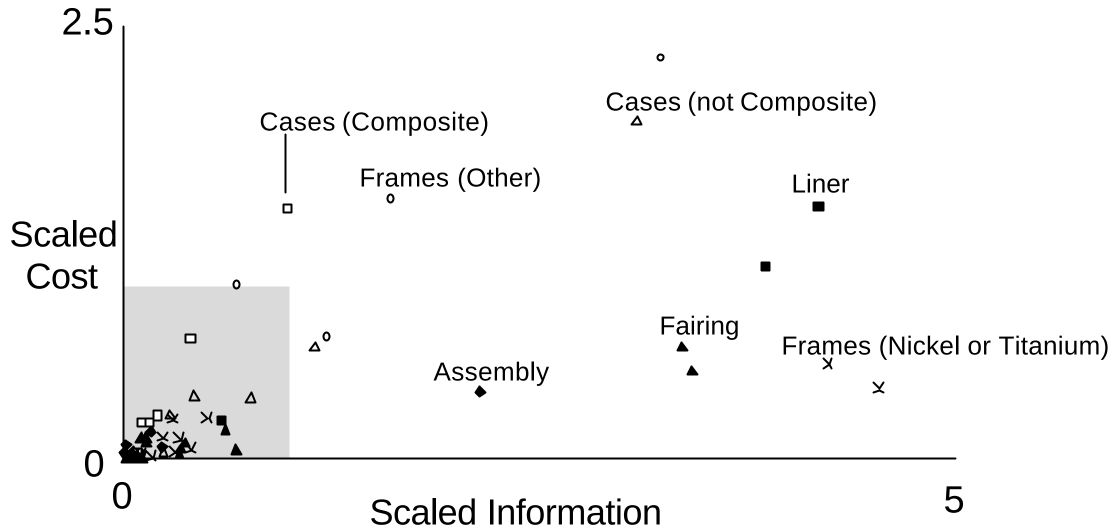

Figure 4c: Data for high cost parts

We can speculate on some possible explanations. Part types may correlate well with manufacturing processes (particularly in the high-cost, peculiar-to-aerospace parts like airfoils and disks). As Hoult and Muter [2] observed, different processes have different intrinsic relationships between design information and cost. For instance, on a lathe, every face of every feature must be turned separately. On the other hand, in an investment casting many features can be cast simultaneously.[2] Therefore, segregation by part types should improve correlations between design information and cost.

POTENTIAL NEXT STEPS

This investigation opened more questions than it answered. Following are some avenues for research and action:

1. What is the best way to quantify the design information in a drawing so as to provide the optimal guidance for design engineers who are balancing cost against performance?
2. Are there potential automated techniques for quantifying the design data in a computerized definition (CAD) system?
3. Since the potential benefits of an information-based CAD cost module would reach far beyond the aerospace industry, the authors recommend that a consortium be established to explore the feasibility of such a system. Such a consortium should be led by an independent organization trusted to handle proprietary information from a large number of organizations. These organizations, in turn, would be willing to share CAD descriptions and cost information in order to develop the cost-information database. The consortium should also include all of the major CAD vendors who recognize that a CAD cost module is a desirable product attribute. The first task of the consortium lead organization would be to outline the road map and timescale for the study.

CONCLUSIONS

We assert that the primary cause of the manufacturing cost of a part is the quantity of information that must be communicated from design to manufacturing in order to correctly make the part. Therefore, design environments that wish to provide the designer with manufacturing cost implications of his or her design decisions should measure the information content of the design as the basis of manufacturing cost estimates. This can be done simply on modern computer aided design workstations. Resulting estimates are superior to mass-based parametric cost models and even superior to process-based estimates traditionally performed manually by manufacturing engineers. Because information-based estimates need take no separate account of the manufacturing processes that will be used, mature manufacturing costs can be estimated for radical new parts and composite materials.

REFERENCES

1. *Aviation Week & Space Technology*, McGraw-Hill, New York. January 8, 1996.
2. Hoult, David P. and Muter Stuart W. "Cost Comparisons of Alternative Designs: an Information Based Method." The Manufacturing Institute at the Massachusetts Institute of Technology, 1993.
3. Shannon, Claude E. *The Mathematical Theory of Communication.* University of Illinois Press, Urbana, IL, 1963.
4. US Bureau of Mines. *Mineral Yearbook*. United States Government Printing Office, Washington, DC. 1990.

CONTACT

For further information, discussion, or comments please contact Paul Collopy at (888) 643-4393, e-mail paul@dfmconsulting.com or David Eames at (317) 230-2525, e-mail david.eames@rolls-royce.com.

[2] An extreme example is the photolithography / epitaxial deposition process used in the manufacture of integrated circuits. A fantastic amount of design information can be deposited into the product in parallel, yielding a very low partial derivative of cost to design information.