ICAS-94-0.2

CFD AND THE CHANGING WORLD OF AIRPLANE DESIGN

Paul E. Rubbert
Chief, Aerodynamics Research
The Boeing Company
Seattle, Washington

Abstract

This paper talks about the contributions that computational fluid dynamics is making to the *processes* used to design airplanes. This process-focused way of looking at what we do, why we do it, and how we do it is far different than the traditional focus of just a few years ago. The early parts of this paper provide a description of the workings of the commercial marketplace and the characteristics of processes that are used to design airplanes. Those descriptions serve to identify and define what constitutes "goodness" in the processes used to define, design and build airplanes. What emerges is an understanding of "goodness" that is quite different from the conventional traditions of the airplane business. A key parameter that emerges on the axes of many of the figures contained herein is *time*! Finally, the role of CFD in contributing to goodness is described, together with examples of CFD computations that are poised to revolutionize the process of aerodynamic design.

Introduction

It has been my great privilege to listen to a number of the Wright Brothers Lectures, and to read a number of the earlier ones given to me recently by George S. Schairer, a Wright Brothers lecturer of a previous generation. In reading through them I came upon a comment (1) attributed to B. Melville Jones, the first Wright Brothers lecturer, in which he stated "- - - I am instructed that the Wright Brothers Lecture should deal with subjects upon which the lecturer is engaged at the time, rather than with a general survey of some wide branch of aeronautical knowledge."

That is what I intend to do. As a manager of CFD research I have been spending a lot of time trying to figure out the directions in which to lead research, directions that will maximize its contributions to the traveling public, to the competitiveness and profitability of the company that employs me, and to the economic well-being of the airlines that buy our airplanes. This navel gazing and associated study has led me to a very different perspective and understanding about what is important and what is not. This is what I share with you in this paper. I hope that you enjoy it and that it stimulates you also to think about and to question why you do what you do. It is not

Copyright © 1994 by ICAS and AIAA. All rights reserved.

enough to be able to do things right. One must also excel at doing the right things!

Our Heritage

The early foundations of the airplane business were clearly built upon a strategy of innovation and technology, accompanied by rapid processes for design, fabrication and testing. That was the strategy employed by the Wright Brothers that enabled them to reduce to practice, the dream of powered flight. That strategy was further propelled by two world wars, during which innovation, technology and rapid design and fabrication processes led to rapid advances in the capabilities of airplanes.

The mindset of depending on innovation and technology as a competitive strategy continued to flourish, not only in the airplane business but widely throughout U.S. industry. As reported by Amai, (2) "During the two decades preceding the oil crises, the world economy enjoyed unprecedented growth and experienced insatiable demand for new technologies and new products. It was a period in which innovation strategy paid off handsomely. Innovation strategy is technology driven and thrives on fast growth and high profit margins. It flourishes in a climate featuring

- Rapidly expanding markets
- Consumers oriented more toward quantity rather than quality
- Abundant and low-cost resources
- A belief that success with innovative products could offset sluggish performance in traditional operations
- Management more concerned with increasing sales than with reducing costs."

Unfortunately, over time we forgot about the importance of accompanying our appetite for technology and innovation with rapid processes for design and fabrication. We let the seductiveness of advanced technology delude us into calmly accepting whatever added time and cost were required to develop and exploit new technologies and new design tools.

So it was during the early days of CFD, the decade of the 1960's and into the 1970's. CFD was born in an era that measured goodness in terms of technology, technology advancement, and innovation.

The rationale for supporting CFD research in those days was the pursuit of mission performance. More speed, more range, and greater payloads were the goals. The standard questions asked by management during reviews of CFD advancements were along the lines of:

- Functionality ~ Can you now do what you couldn't do before? If so, that is good and noble.
- Can your new wing fly at higher Mach numbers?
- Have your new technology or CFD tools allowed the design of thicker wings, with lighter structure and more fuel volume?
- Can you design a shock-free wing?

Almost nobody ever asked how long it took to do something, or how much it cost. Worshipers of a technology-driven innovation strategy don't ask those kinds of questions, because they operate under the principle that "new and wonderful things take as long as they take and cost whatever they cost." And with an almost unbroken track record of successful technology-driven advances dating back to the Wright Brothers, the value of doing new things was rarely questioned, no matter how long they took or how much they cost.

A Newer World

Those days are gone! (Or almost gone. I still encounter some of those same questions during reviews, and I still encounter people who entertain visions of burning me at the stake as a heretic.). The old technology-driven strategy is being replaced by one based on being "market driven" and "customer driven". This does not imply that advanced technology and innovation are no longer worthy of pursuit. What it does mean is that the criteria by which "goodness" is judged has changed. It is no longer good enough to develop technology that accomplishes something that had not been done before. Rather, the new measure of goodness is whether or not it adds value as seen from the eyes of a customer!

That realization has come none too soon, for we were beginning to choke on cost, indecisiveness and lack of commitment. In the early days of CFD, the wonderful vision of being able to actually compute some flow field was usually enough to justify and fund an R&D activity. But over time, as more people began to play in the CFD sandbox, as computers became more capable and more costly, and as we increased in knowledge of how to compute, our abilities to develop computational capabilities began to exceed the resources available to fund such work, even though that resource stream grew rapidly also. We had to face the fact that being able to do something technically was no longer an adequate justification for doing it. We had to come up with a better way of identifying what made sense to do.

That better way is found in the paradigm shift that the aerospace industry and many other industries are experiencing in their endeavors to remain competitive. They are moving beyond a technology-driven, innovation-based strategy to one that focuses on customers, both internal and external. They still value advanced technology and innovation, but the deciding factor that tells if it is worth doing is the question "does it add value for the customer?" In the airline and commercial airplane business, value to a customer is defined in terms of safety, reliability, comfort, convenience, timeliness and low ticket price. A product that is central to providing that value just happens to have wings, tail and engines and looks like an airplane.

This has had a profound effect on the subsonic commercial transport aircraft business. Figure 1 illustrates the fundamental change that has taken place. In the historic model of Figure 1a, advancing technology drove the desire to design new airplanes, airplanes whose operational performance in terms of speed, range, and fuel efficiency significantly exceeded that of its predecessors. The cost of the airplane was dictated by the cost of the technology embodied in its design. Its price in the marketplace was dictated by its cost to design and produce. And the final question was, will it sell? Some did and some didn't, as witnessed by the wreckage of numerous airplane companies or divisions that at one time or another ventured into the market place with commercial transport aircraft.

The new model, Figure 1b, is nothing less than the reverse of the old. The difference is that *the arrows all point the other way*! The birth of a new transport airplane today is driven not by technology, but by market requirements. The market, and not technology, dictates the kinds and numbers of new

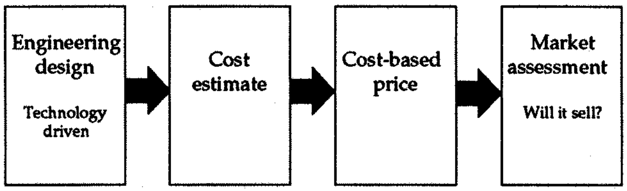

(a) The Historic Model

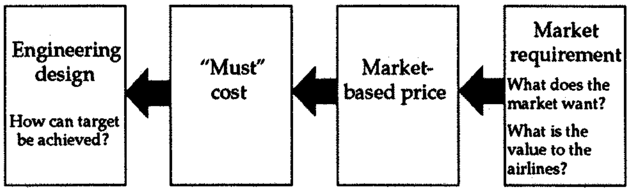

(b) The New Model

Figure 1. Product Decision Process Models

airplanes that are needed. The customers not only determine overall airplane parameters such as range, speed and size, but their desires are flowing down into the detailed component level. We are finding that wing span is no longer determined solely by balancing aerodynamic efficiency against structural weight, but is becoming defined by the customer airlines and their associated ground handling requirements. Choice of span is also affected by manufacturing costs, a longer wing being more expensive because of the longer machines and the larger number of fasteners required to assemble it. On the newest Boeing 737-700 design the airlines have requested the use of aluminium rather than composites for the nacelles, a choice that would not have been made if the design were technology driven or mission performance driven. Rather, it was driven by the fact that nacelles are vulnerable to being bumped by ground handling equipment, and aluminum is easier and less costly to inspect and repair.

The market also establishes the price of the airplane. If the price is such that the airplane adds value to their operations, the airlines will buy it, and if it doesn't, they won't. That market-based price thus establishes a lid on what the manufacturer's cost to design and build it must be.

And so the name of the game in technology supporting the airplane business is to be able to selectively exploit advanced technologies in ways that improve our cycle time for responding rapidly to market developments, and which add the most value in the eyes of the customer. In other words, *we not only have to be good at doing things right, but we have to be very good at figuring out the right things to do!*

Figure 2 shows a typical breakdown of total airplane-related operating cost (TAROC) for a long range subsonic jet transport. The feature that stands out in this figure is the dominant effect of cost-of-ownership, the price of the airplane, on total operating cost. This component of cost is 2 1/2 times bigger than the cost of fuel! In simple terms, this means that if some clever aerodynamicist can figure out an airplane shape that reduces drag and fuel consumption by 2 1/2% for a given airplane mission, but that entails a 1% increase in the purchase price of the airplane, the airline customer would experience no added value. For shorter range airplanes the cost of fuel is even less significant. This is the principal reason that laminar flow control or riblet

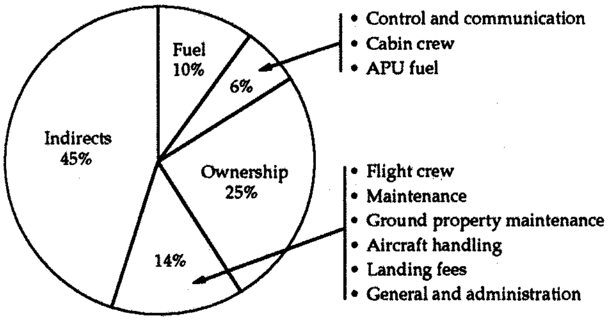

Figure 2. Total Airplane-Related Operating Cost (TAROC) Breakdown; Typical for U.S. International Rules, 6,000-nmi Mission

technology have not been offered to date on transport aircraft. Both of those technologies reduce fuel consumption, but currently the added cost of design, fabrication, installation and inservice maintenance cancel the value of reduced fuel consumption.

So, the name of the game has really changed. We now know how to design commercial transport airplanes that fly at the speeds that the airlines desire. We know how to design them to fly as far as customers want to go, and we can design them to carry as many people as desired. Mission performance is no longer the driving factor that it once was*. The driving factor today is economic performance, the ability of the airplane to do its job at less overall cost to the airline and to the traveling public, with the utmost in safety and reliability.

Does this mean that there is little to be gained by further investment in CFD? If we continue to ask the same old questions concerning payoff in terms of higher drag rise Mach number, thicker and lighter wings, etc., the answer may well be yes, at least for subsonic transport aircraft. We have been attacking those challenges for decades, and we are probably reaching the point of diminishing returns.

What we have to learn to do is to ask the right questions. The name of the game today is to design and build airplanes that deliver the best in economic performance, and to be able to design and deliver them at the time when the market wants and needs them. The way that that is accomplished is through processes, the processes that are used to understand the market and the customer, and the processes that are used to design and build airplanes.

So, the right question to ask is *"What is the payoff from investing in CFD in terms of improving our ability to deliver airplanes at the right time, when the market needs and wants them, and which deliver the utmost in economic performance in the hands of the airline customers?" The answer to that can be found by examining the role of CFD in the processes that are used to design airplanes, because that is where the leverage is.* The remainder of this paper is focused on understanding and clarifying the role of CFD in improving the processes by which airplanes are designed. The indicators of goodness that emerge are far different than the traditional way that we have been taught to think!

Fundamentals of a Competitive Market

Airplanes, particularly commercial transport aircraft, are sold competitively in the marketplace. The factors that mostly influence the success (or failure) of an airplane in its market niche are:

(i) how well it fits the mission requirements of its customer's market niche in terms of passenger capacity, range and speed
(ii) the economic performance and reliability that the airplane delivers in everyday airline service,
(iii) availability for delivery when the market wants and needs the airplane.

Being an engineer who is accustomed to thinking and communicating with figures

---

* The design of a supersonic transport remains more technology driven because the achievement of the required long ranges to serve the Pacific market at costs which are competitive with subsonic airplanes remains a challenge. Mission performance, and associated advanced technologies that may enable the achievement of that performance, remain of critical importance to the future economic viability of a supersonic transport.

and sketches, I have attempted to express these fundamental factors in graphical form as shown in Figure 3. The vertical axis in Figure 3, labeled "Goodness,"is a measure of economic performance of an airplane as operated by an airline. Goodness is a global measure that includes how well an airplane is sized to fit its market niche, its purchase price, its fuel efficiency, its maintenance costs and its dispatch reliability, because all of these influence its overall operational cost. The accountants within the airline industry are getting to be pretty good at quantifying goodness.

The horizontal axis is a measure of time, or date of market entry. There is a line that ranges from the lower left to the upper right of the figure which I term the "zero percent market share" line. If your airplane, when plotted on this figure, falls below or to the right of the zero percent market share line, you cannot sell any, the reason being that your competitor's product beats yours in terms of its economic performance and in terms of date of availability. If your product is second best in both of these factors, it is dead.

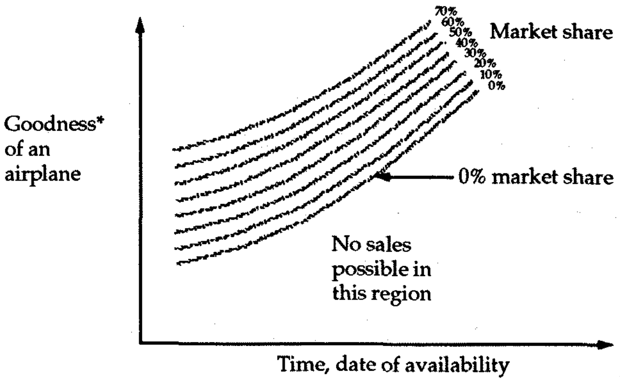

* Goodness is an integral measure that includes purchase price, fuel efficiency, maintenance costs, and dispatch reliability.

Figure 3. Variation of Market Share With Goodness and Date of Availability

In order to capture market share, your airplane must live in that portion of goodness-time space that lies above and to the left of the zero percent market share line. Figure 3 displays what must be the shape of increasingly greater percent market share lines, marching upward and to the left. Increased market share can only be captured by a product that delivers increasingly superior economic performance, over its competitors, and/or which is available sooner, rather than later.

We will now leave this discussion on the fundamentals of a competitive market and return to it later, after examining the fundamentals of processes.

Fundamentals of a Process

Airplanes, like most products, are designed using processes. When I speak of a process, I don't mean a function, a task, or an activity. Rather, a process is(3) "a collection of activities that takes one or more kinds of input and creates an output that is of value to a customer." An example that I often use to illustrate what is meant by a process is the example of a recipe in a cookbook. The typical recipe is a description of a process. It says things like "mix so much of ingredient A with ingredient B, stir for 2 minutes, bake for 20 minutes," etc. That is a process. Stirring, baking, etc. are tasks or activities or functions within the process. The entire collection of activities, in proper sequence, comprise the recipe or process. Execution of the process produces an output that is of value to a customer. In the case of a recipe, that output is a ready-to-eat product, which is of greater value to the eater (customer) than the raw ingredients which entered into the processes.

Processes have a number of common characteristics. One is that they should be designed to meet the needs of the customer. If the eater desires roast beef and the cook uses a recipe that produces tomato soup, the customer will not be satisfied. Or, if the eater desires to be served in ten minutes but doesn't receive his meal for an hour, he will not be satisfied.

Another characteristic that processes need (but sometimes don't have) is a clearly defined owner, someone who is responsible for the integrity of the process. The old saying of "too many cooks spoil the brew" tells us that if nobody is clearly in charge of the process, the product will suffer. In my observation, that is just as true in designing airplanes as in cooking.

Other characteristics of processes are that they cost money and take time to execute. Those are aspects that can be measured and quantified provided the process is well defined, stable and repeatable. The fast food restaurants know very well how long it takes and what it costs to produce a hamburger. Airplane companies badly need to know how long it takes and what it will cost to produce an airplane.

Another characteristic of processes is that they produce products that contain variation. In other words, the product deviates in one way or another from the norm. In the hamburger business, one expression of variation would be the amount by which the meat is undercooked or overcooked, which is only ascertained after the product has been bitten into by the customer. In the airplane business, one measure of variation in the design process is the amount by which the airplane's performance, as measured by actually flying it, deviates from the expectations of performance that were provided by the design process.

Yet another characteristic of processes is that they can be improved. The renown chefs of the world achieve their fame by doing a better job of satisfying their customers, the eaters. They generally achieve that by improving their recipes while using customer satisfaction as the measure of goodness, and by stabilizing the improvements so that the results are repeatable, time after time. And they are aware of the fact that the tastes of the market can change with time, requiring that they continuously monitor and interact with their customers and that they adapt their recipes over time to changing tastes. The top chefs of the world are frequently pictured chatting with their customers on the restaurant floor, and for good reason. The top airplane companies of the world are also doing so.

The smart chefs have also learned to separate process improvement activities from process execution. When trying a new recipe or a new ingredient for the first time, they execute it off-line, on a small scale, and sample it themselves. Only when they are satisfied that it produces, reliably, a product that will satisfy customers, do they add that product to the menu. In the world of airplane designing we have not been so careful about that. The consequences over the years have been overruns in schedule and budget, missed performance guarantees, etc.

One characteristic that seems common to all design processes is that the "functional goodness" of the product being designed increases with the amount of time that one is willing to spend on designing it. As an example, in a process called "wing design", there usually exists very early in the process a global definition of the wing in terms of span, wing area, sweep, overall thickness, etc. If one were to stop designing at that early point in the process and actually build that wing, its overall functional performance would be less good than if one spent more time in designing it. But, eventually one encounters the law of diminishing returns, which says that additional investments of time will yield increasingly smaller gains in goodness. Also, the cost of executing a process increases monotonically the longer a process is allowed to run. These ideas can be displayed graphically as shown in Figure 4a. The same features apply also to production processes, and indeed to the overall global process of designing and producing an airplane.

In a nonsubsidized market environment where the cost of executing a process must be passed on to the customer, the concept of "economic goodness" associated with a product of that process should accompany the concept of "functional goodness". I define economic goodness to be the result of subtracting the cost of executing the process, which must be passed on to the customer, from the functional goodness of the product as sketched in Figure 4b. Higher cost makes a product "less good" in the eyes of a customer who must pay the bill.

In the business of cooking, the cycle time and cost to execute a new recipe is short, and so the chef can well afford to continue to refine the recipe until no further improvements are possible. He can run his process far to the right, down the asymptote of Figure 4a. In the airplane business, the cycle time and cost of designing and producing something as complex as an airplane are so high that traveling down the asymptote is not a viable option, as witnessed by the common cry of wing designers to "give us one more cycle and we can get a better wing" as the date for freezing the design approaches, and their cry is rebuffed.

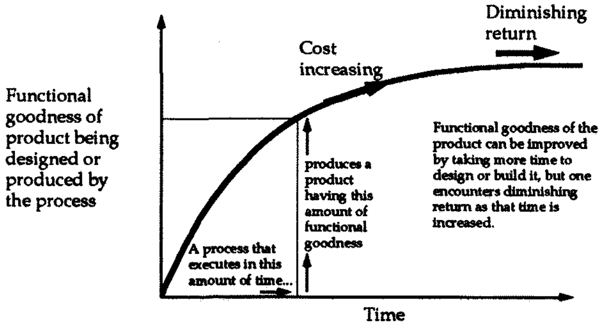

(a) Functional Goodness

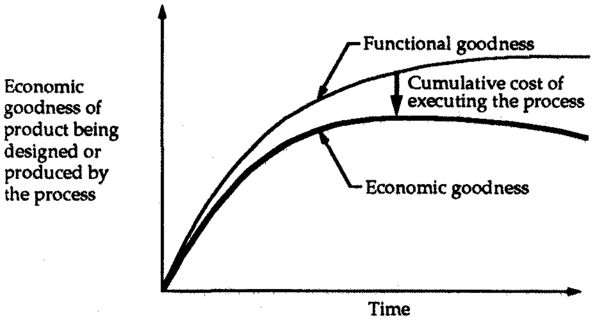

(b) Economic Goodness

Figure 4. Fundamental Characteristics of Processes

The Influence of Fast Processes on Market Share

Another feature of design processes that can be displayed graphically is the effect of process flow time or cycle time. Figure 5 shows the relationship between a faster process and a slower process. Both processes may asymptotically approach the same level of functional goodness of the product being designed, as sketched in Figure 5a, but one does it in less time than the other. Figure 5b displays faster and slower processes in terms of economic goodness, where the cost vs. time for executing both processes is assumed to be the same for both. It should be obvious how to draw the curves for other cost vs. time scenarios. However, it should be pointed out that in most industries, including aerospace, the largest contributors to cost are

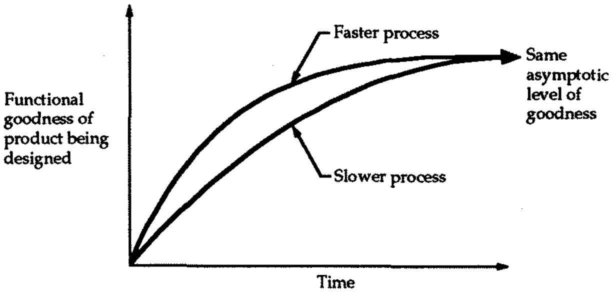

(a) Functional Goodness

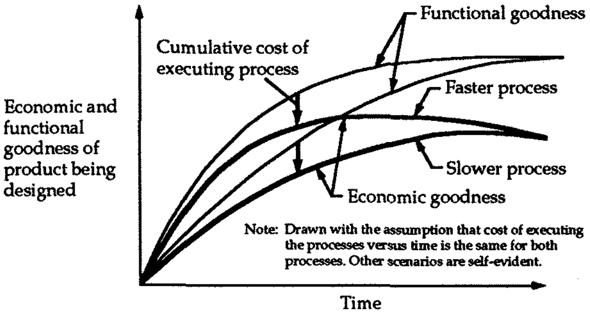

(b) Economic and Functional Goodness

Figure 5. Slower Versus Faster Process

items whose cost is proportional to time, such as computers, facilities, etc. Therefore the old rule of thumb of estimating cost based on labor hours is no longer a good indicator of cost. A better rule of thumb is to assume that cost will be proportional to the calendar time required to do the job.

Now let us resurrect the market share plot of Figure 3 and overlay on it the process curves from Figure 5. Here, the process curves relate to the entire overall process of defining requirements, designing, building and selling the product. We can do this overlay because both figures have a vertical axis of economic goodness and a horizontal axis of time.

Let us examine two scenarios. In the first scenario it is assumed that the cost of executing the process is inconsequential, such that economic goodness can be equated to functional goodness of the product. In the second scenario it is assumed that the cost of executing the process is significant, and that it is proportional to the time required to execute the process.

The resulting overlays for these two scenarios are shown in Figures 6a and 6b respectively. The major points that leap out from the figures are that:

1. *Fast processes* that compress the calendar time required to bring the product to market *have tremendous leverage on market share.*

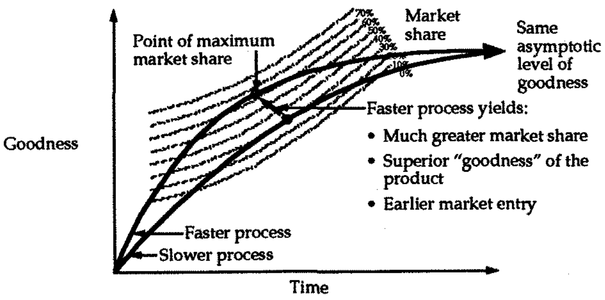

(a) Assuming Cost of Executing Processes Is Inconsequential, Such That Functional Goodness and Economic Goodness Are Equal

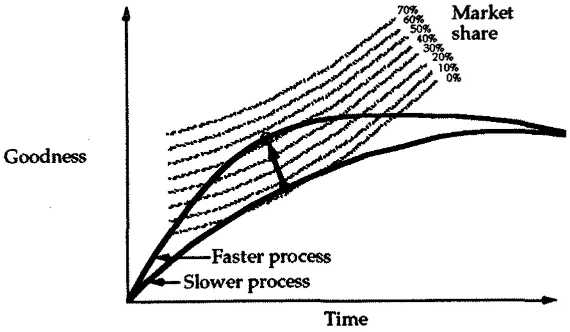

(b) Assuming Cost of Executing Each Process Is Proportional to the Time Spent in Executing Each Process

Figure 6. Influence of Fast Versus Slow Process on Market Share

2. *Faster processes allow one to achieve a greater level of functional goodness,* in the product than slower processes, even though both processes may be capable of achieving the same asymptotic level of functional goodness. In fact, the asymptotic level of goodness that a process may be capable of achieving can mean relatively little. What counts is the ability to get good quickly!

3. *Faster processes produce early entry* into the market.

4. *One should not allow a design or production process to run down the asymptote to perfection,* to maximize functional goodness, *even if the added cost of taking more time is insignificant.* The maximum market share is achieved with less than the asymptotically perfect level of functional goodness.

These observations are rather independent of the cost of executing the processes as indicated by the similarities of Figures 6a and 6b. A fifth observation is that increased market share can obviously be achieved by finding ways to reduce the cost of executing the process and thereby improving the economic goodness of the product produced.

One outstanding example in the airplane industry of the value of fast processes is the saga of the Boeing 747. That airplane was designed, produced and certified in the breathtakingly short time span of about 3 years rather than the customary 5. Yes, the design was far from perfect. The production was planned and executed in great haste, with the result that scrap and rework became a way of life in the factory. But the facts are that that airplane came to literally own the long haul, high capacity market, and did so early enough that no other manufacturer attempted to compete against it for 25 years. That airplane brought greater comfort and lower ticket prices to long haul travelers, which have been positive factors in spurring the growth of air travel. And the profits and market share garnered by its manufacturer have been handsome indeed!

The above observations represent a major change in thinking for those of us who have deep roots in an old culture where functional performance seemingly was the only parameter by which advances in technology were measured. As a developer of CFD capabilities, I have found that the key question in technical reviews over the years has seemingly always been "how much better aerodynamic performance has your new CFD tool produced?" We would respond by dutifully running out the asymptote of Figure 4a for the next year so that in the next annual review we could show an incremental improvement in asymptotic goodness. By hook or by crook the target was to raise the level of asymptotic goodness, with little regard to how much time or money it took to achieve it. And as the years passed, the incremental asymptotic improvements became less and less as we approached the asymptotic limit.

It is now becoming increasingly clear that the focus of CFD R&D activities must be changed. We must learn to equate goodness with speeding up the process of designing something, rather than measuring only the level of goodness that can be achieved by running a design process far down the asymptote. *The key question for the CFD developer must become "by how many days did you reduce the time required to execute the design process?"*" In other words, the target is to be able to arrive at a good design quickly, one that provides an acceptable level of functional performance but perhaps not the ultimate that could be achieved.

A good example of the value of this new philosophy arose last year when Boeing began to think seriously about upgrading its 737 family of airplanes in response to changing customer needs. We had to change the wing in order to provide the range and speed that the airlines wanted, and a key question was how best to do it. The three options were to (i) modify the old wing, (ii) replace the wing with the outboard section of the 757 wing, or (iii) design an entirely new wing. With the aid of fast, CFD-based design processes, wings representing all three options were quickly designed.

The subsequent verification wind tunnel test verified that all three options could provide the required functional performance desired by the airlines. Therefore, the choice of which wing to select could be based on other factors, such as cost. In essence, fast, reliable design processes enabled Boeing to make the design choices that optimize the <u>economic performance</u> of the airplane. The differences in aerodynamic performance between the three candidate wings were not the deciding factor, and all could provide the required functional performance of achieving the required speed and range.

Variation Associated With Processes

Another feature of processes is that they produce output that contains variation. As an example, the process of manufacturing a machined part produces parts that are all slightly different from one another, the part-to-part differences being a measure of variation. This can be described statistically. One can define the probability of the output of the process deviating from the expectation, and in that way quantify the variation.

And so it is with processes that design something as complex as an airplane, or even a wing. The design process produces well-defined expectations on how the designed product will perform (range, speed, weight, fuel consumption, etc.). It is generally found that the actual performance measured by building and flying the airplane is somewhat different. I define this difference between design expectation and actual performance to be a measure of the variation contained within the design process. On one airplane program the airplane may end up performing 2% better than expectation, while the next airplane program might produce an airplane that is 1% worse than expectation. Those are measures of the variation associated with the design process.

It will also be found that there is plane-to-plane variation in performance, even though all are built to the same design. Those plane-to-plane differences are measures of the variation contained within the manufacturing and assembly processes, whereas the difference in performance between design expectation and the average performance realized by sampling a number of airplanes of the same design is a measure of the variation contained within the design process.

A feature common to all design processes is that one's level of knowledge concerning how the designed product will actually perform increases with time as one proceeds with the work of designing. As an example, early in the design of a wing we have a certain expectation of how it will perform based on knowledge of its sweep, thickness, etc. But after we have refined the design in detail, conducted wind tunnel tests, etc., we know with greater precision what its performance actually will be. And so, the variation between expectation and actual performance becomes less as one proceeds farther downstream in a design process.

We can illustrate that graphically by adding a variation band about the process curve as shown in Figure 7. The expectation of the performance (e.g., "goodness) is the central line within the shaded band of variation. The actual performance that would be achieved at any point in time, which can only be measured by stopping the design process at that point in time and building and flying the design, would be different, lying somewhere within the upper and lower extremes of the shaded band.

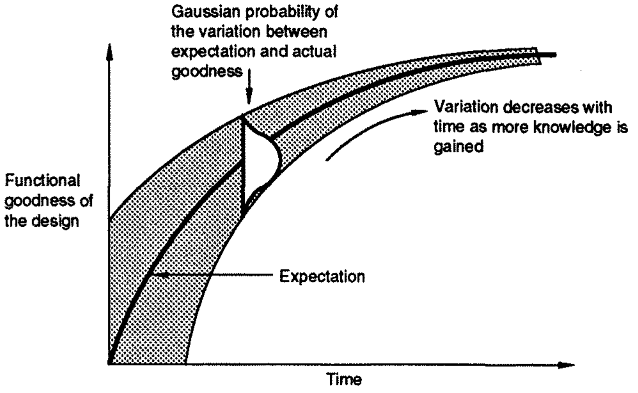

Figure 7. Variation Associated With a Design Process

The penalty for having excessive variation in the processes used for designing airplanes can be very high. The industry is replete with examples of missed guarantees of range, fuel consumption, and speed, and the consequences thereof in terms of market share and profit. These are the results of variation in the design processes. Cost overruns and late deliveries are other embodiments of variation in processes because they reflect deviation from expectation in terms of ability to predict the cost and cycle time of executing processes.

Conservatism in design practices does not solve the problem of variation. For example, if the performance of an airplane turns out during flight test to be better than expectation, this means that sales were probably lost by not claiming the true performance, and profits were lost by not crediting the value of the extra performance in the price of units that were already sold.

I find it interesting to observe how the industry tends to react to findings of variation in their processes. If the airplane performance turns out to be better than expectation, everybody thanks their lucky stars and goes on about their business. Nobody beats up the design team or on the design processes, and the companies don't rush to spend money to fix the variation that this reveals in their processes.

Alternately, if the airplane fails to meet its performance expectations and guarantees, the companies throw money at performance fixes in the form of weight and drag improvement programs, engine upgrades, etc. A side effect of this response is that researchers and process improvement staffs are led to believe that the most important things to work toward *ad infinitum* are incremental improvements in drag, fuel consumption, etc., by traveling further along the asymptote in Figure 4a. In fact, that is only the near term response required to fix the unforeseen problem. The long term response and focus should be to improve the underlying processes so as to reduce the variation contained within them. That is where the leverage lies in terms of increasing the competitiveness of the company!

Implications for CFD and CFD Research

It is becoming increasingly clear that we need to expand our focus and shift our ideas concerning what is important. The old paradigm has been to focus on exploiting CFD and CFD developments to achieve demonstrations of increased performance, measured in terms of higher drag rise Mach number, increased wing thickness, higher $C_{L\,max}$ or whatever. And those demonstrations were frequently conducted not within a true airplane design environment but in a research environment. In reality, what we were doing was usually a combination of (i) running further down the design process asymptote of Figure 4, (ii) raising the level of the asymptote somewhat, frequently accompanied by the necessity of introducing more operations into the design process and thereby increasing the time and cost required to achieve the design, and (iii) sometimes increasing the variation in the process by introducing more operations and by driving toward aerodynamic designs that were increasingly sensitive to small changes in geometry, Mach number, angle of attack, etc.

The new paradigm must be to focus on the idea of justifying and exploiting CFD to improve the processes by which airplanes are developed, with improvement defined in terms of

- reduced cycle time to execute the processes
- reduced cost to execute the processes
- reduced variation associated with the processes

This is where the competitive advantage resides.

I have found over the past couple of years that, while it takes some time to become accustomed to thinking of CFD research and CFD development in these terms, eventually it becomes completely natural and automatic to do so. I have also found that everything we do in CFD R&D can be readily viewed in these terms. And because the first two measures of improvement (cycle time and cost to execute a process) are particularly amenable to quantification and estimation, it would appear that the decision processes we use to determine how much money to invest in what kinds of CFD development should be somewhat more amenable to rational argument and justification than in the past (but only when the principal players in the organization have learned to equate value with improvements in process cycle time, cost, and variation. Until that happens, be prepared for frustration.).

What About Airplane Performance-Are We Now Abandoning It In Favor of Process Performance?

If we focus on improving cycle time, cost, and variation associated with airplane design processes, have we abandoned the quest for improved airplane performance? The answer to this is no. The truth is that the path to improved performance is by improving the processes by which we seek that improved performance. I remember the era of 1970 when Dr. Richard Whitcomb, of NASA Langley, was carrying out research to demonstrate the feasibility of achieving cruise Mach numbers in the neighborhood of $M_\infty = 0.98$ for transport aircraft. His principle tools were the wind tunnel, a file for altering the geometry of a wind tunnel model, some primitive CFD, and an outstandingly creative and analytical mind. Every time he changed a model geometry and retested it, he learned something, and eventually he acquired an in-depth understanding of the principal design features that enabled the goal to be accomplished, and he demonstrated that it could be achieved.

A breakthrough technology that he employed was the file, wielded personally, that enabled him to change and retest a model within hours or days, rather than the months required to build another model. In other words, he shortened the cycle time of the learning process! If he had relied upon the conventional practice of designing and building a new model after every test, he probably would not have been able to achieve his goal before patience, resources and peer support were used up. Today, with the aid of modern CFD, a Richard Whitcomb could learn in weeks what then took him months, and he would have been able to achieve his goal of $M_\infty = 0.98$ performance much quicker and at less cost.

The facts are that the methods, tools and processes used in quest of knowledge of how to achieve improved aerodynamic performance are much the same as the ones used to design airplanes, namely CFD, geometry systems, wind tunnels, test instrumentation, etc. Improvements in the cycle time, cost and variation of these tools and processes have just as much payoff in enabling the researcher to figure out how to achieve some advance in "aerodynamic technology" sooner, cheaper, and with greater certainty. *We need to recognize that in most cases, the best route to increased knowledge and understanding is to improve the processes by which we acquire knowledge and understanding.* This is the path by which most "advanced technology" is developed.

The Changing Character of CFD R&D

The old paradigm of CFD research was to be technology driven. As computers became more powerful, they allowed us to do more things, and so we did them. First came panel methods for linear flows, and 2-D boundary layer methods. Then 3-D boundary layer methods, transonic small disturbance and full potential methods, Euler, and various modelings of the Navier-Stokes equations. Then direct numerical simulation of turbulence and transition, and so on. Each, in their time, were leading edge challenges.

The primary contributions of those developments have been in two principle categories, (i) creation of knowledge and understanding, and (ii) procedures and tools for use in designing airplanes. Each time computer advances allowed us to advance to another level of physical modeling, the initial contributions and payoff were mostly in the category of knowledge and understanding. They enabled us to gain a much better understanding of aerodynamic flows and their behavior. Only later did the confidence, reliability, and automation of those codes develop to the point where they could or would be accepted as standard computational procedures for regular use in airplane design processes, and only a small percentage of codes ever achieved that status.

That scenario is portrayed in Figure 8. Spanning the figure is a band labeled "the leading edge of physics simulation." Each level of modeling, beginning with panel methods, then full potential, Euler, and so forth, experienced the same type of evolution with time when passing through the band from left to right. At points in time to the left of the band there were no capabili-

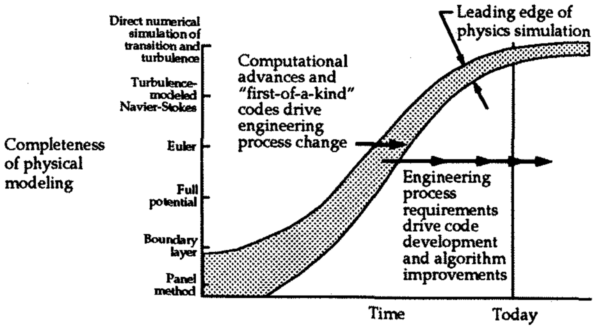

Figure 8. Crossing the Leading Edge

ties. Then, as computers became more powerful with time, it became possible at each level, at their respective points in time, to solve the governing equations on a computer. CFD researchers then proceeded to invent solution algorithms and build codes that solved those equations. They had entered the band of the leading edge.

Those types of activities, taking place within the leading edge band, are technology driven and properly so. They are justified within an overall vision that they will somehow contribute to our ability to design better airplanes, but the details of that vision are generally very unclear and imperfect. The primary impact of those technology-driven, leading edge codes on airplane design processes seems to be to build an awareness that those new types of computations can now be done. Relatively few of the early, leading edge codes actually end up in regular use for airplane design work.

But time marches on. With continued research, testing and demonstration, the vision of how an airplane design process could be made better eventually sharpens. At that point the design process itself begins to generate requirements concerning what the new code must be able to do, what geometry it must be capable of handling, what level of accuracy is required, what its cycle time must be, what other computing systems it must interface with, etc. At that point, CFD development has become process driven. It has emerged from the leading edge band of Figure 8 and entered the domain to the right of the band.

This transition, which takes place as CFD R&D traverses the leading edge band and experiences the change from being technology driven to process driven has strong implications for the conduct of computational R&D. A key parameter is the increased level of understanding required of the researcher/planner/code builder concerning the targeted airplane design processes. For endeavors within the leading edge band that are appropriate for employing a technology-driven strategy, a conceptual understanding or even a vision of a candidate airplane design process to be improved is generally sufficient, because the primary contribution of the research will probably be mainly knowledge and understanding. This is clearly a role that researchers in government laboratories and in academia are postured to carry out.

However, for endeavors that need to be process driven, such as the design and fabrication of a CFD code for use in an airplane design process, it is imperative that the planning of the code be done with a very clear understanding of the process which it is intended to improve. That is best done with active participation of the design engineers and the process owner who have hands-on experience with the process. It is carried out most effectively by the researcher/planner/code developers who reside in industry and can interact daily with the industrial designer community, and be privy to the proprietary processes of their respective companies.

Even then, one frequently ends up making only modest improvements to the targeted design process or processes because of two primary limitations. One is the natural tendency of people to think of process improvement in terms of doing the same things they did in the past, only faster. It is easiest to view the role of new technology or tools as a means for doing those same things faster. But in many cases this usually produces only modest benefits. It is much harder to visualize those fundamental changes to a process that frequently end up producing truly great advances.

The other limitation is risk. If a CFD development plan merely proposes to speed up some part of an existing airplane design process and then fails to meet its expectations, the customer knows that he can fall back on his old part of the process (albeit with a schedule slide). The impact on other parts of his process will be minimal. However, if the objective is to achieve fundamental and comprehensive changes that affect many parts of the process, then the penalty for failing to deliver as promised is much higher. This makes it much more difficult for the CFD developer to get customer ownership and acceptance of proposed code development that could yield some truly large gains. He frequently must resort instead to the technology driving mode by developing and demonstrating a pilot capability that prospective process-owning customers can "touch and feel", in order to get them to think in more visionary terms and to reduce their risk to acceptable levels.

Also, those larger changes to a process frequently extend beyond the confines of a single discipline and must be viewed cross-functionally, which increases the difficulty of achieving success.

Making the transition from technology-driven to process-driven R&D is not easy. It was not too many years ago that CFD visionaries were enamored with the vision (Figure 9) of a CFD engineer spending days or weeks in front of a Silicon Graphics workstation crafting complex grids about airplanes, sending the assembled grid off to the mainframe for crunching, and then spending days producing pretty colored pictures of various elements of the computation. This was appropriate technology for demonstrating what could be done with then-leading-edge flow solver technology. But it is not the type of process or tool that

Vision

Characteristics

- Long apprenticeship
- Master craftsman
- Low output, long time to fabricate product
- Pieces hand-fitted
- Tolerances not well-understood
- Validation consists of demonstrating that the final product functions about as expected

Observation

These characteristics also describe the way manufacturing processes functioned prior to the Industrial Revolution! Today they are as unacceptable for airplane design processes as they are for manufacturing.

Figure 9. An Earlier Vision of CFD Use

engineering process owners want or need for the task of designing airplanes. Its characteristic features were the following:

- long apprenticeship
- master craftsman, with quality strongly dependent on the skill of the craftsman
- low output, long time to fabricate product
- pieces hand fitted
- tolerances not well understood
- validation consists of demonstrating that the final product functions about as expected.

Those same characteristic features describe a way of doing things that disappeared in the world of manufacturing with the advent of the industrial revolution! The process owners whose job it is to produce airplane designs rapidly, reliably and at lower cost have mostly rejected this vision and for good reason. It was not customer driven and does not meet their process requirements, just as the 19th century craftsman and his tools does not meet the requirements of 20th century manufacturing processes.

The Search for Opportunities

As one acquires a more complete understanding of the processes by which airplanes are designed, and what is of importance with respect to improving those processes, a multitude of opportunities for CFD unfold. Some opportunities are cross functional in nature, such as incorporation of manufacturability constraints into CFD design codes to reduce the need to assess manufacturability after-the-fact. Another one within probable reach is to develop and demonstrate a degree of reliable accuracy with CFD along the perimeter of the flight envelope such that preliminary loads for structural design purposes can be produced concurrently with aerodynamic wing design. That would enable structural design to commence months earlier. And other opportunities for reducing cycle time, cost and variation exist just within the aerodynamic portions of design processes.

In searching out these opportunities, it soon becomes apparent that most design processes contain another feature in common, namely that a typical design process is comprised usually of a lengthy sequence or deployment of operations involving a sequenced array of rather specific, special purpose tools. There usually is no single tool or code that "does the design."

This same feature has been true of manufacturing processes for many decades now. Manufacture of a part may begin with a forging, then a lathe operation, drilling, then milling, proceeding to polishing and final inspection, with variation initially high in the early operations and becoming increasingly less as the part nears completion. There is no single tool that "does the whole job." And improving the cycle time, cost or variation of any of those tools can add value to the process.

The same is true of design processes. A typical sequence of operations involved in, say, aerodynamic wing design, entails the use of a number of codes for specific purposes The process might begin with a sequence of runs with a fast, cheap, wing-body code to examine the performance of a baseline wing over the entire flight envelope. That may be followed by a sequence of inverse design/analysis runs to improve the performance of the wing at selected design points and to achieve a better balance of performance across the envelope. That might be followed by more costly but more accurate Navier-Stokes solutions to "check" selected points in the flight envelope. Then, action may shift to a code such as TRANAIR(4, 5, 6, 7), capable of handling very complex geometries, to perform the work of nacelle integration and manufacturability integration. That in turn might be followed by multi-block Navier Stokes to check the aerodynamics at certain key points in the flight envelope. If the resulting wing was not too much different from one that had been wind tunnel tested previously, the aerodynamic lines could be frozen at that point, followed some months later by a wind tunnel test that played the role of "final inspection" and of filling the data base throughout the flight envelope at less cost and with greater accuracy than is possible with CFD today.

Other points of commonality between design processes and manufacturing processes are that:

- fabrication (of hardware or of designs) uses a rapid sequence of special purpose tools
- tolerances are a way of life
- uses the lowest cost and fastest tooling (or codes) that will produce the required thruput and tolerances
- processes are under control
  - preplanned sequence of steps
  - bounds on variation (tolerance) known and predictable
  - time-to-fabricate (time-to-design) planned and committed at the outset

In examining the specific role of a certain tool (code or wind tunnel) in a design process, it becomes apparent that each occupies a rather well-defined "range of usefulness" within the overall design process as shown in Figure 10. The upstream end of the range is generally established by cost and flow time. At this upstream end of its range the code or wind tunnel is competing with other tools that are faster and cheaper to use.

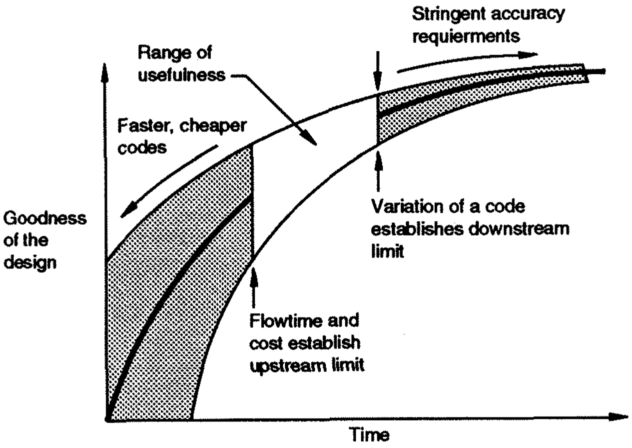

Figure 10. A Code's "Range of Usefulness" Within a Design Process

The downstream end of the range is generally restricted by limitations on variation (e.g. reliability of accuracy). The downstream portions of processes are increasingly demanding of low variation, and tools whose variation is higher cannot have a role in the downstream portions of a process. Perhaps the greatest single restriction to expanded use of CFD today is its frequent inability to produce the levels of reliable accuracy that are demanded by the downstream ends of aerodynamic design processes, a regime still dominated by the wind tunnel. Interestingly, the ability provided by CFD to carry out logic-based aerodynamic design is greatly increasing the demands for higher precision and reliability of accuracy in wind tunnel testing and testing practices, including exploitation of the highest Reynolds number capabilities that can be achieved.

It should be clear at this point that producing a CFD code with little detailed knowledge of the design process encompassing its intended use is at best a hit or miss proposition. It is not uncommon for such a code to have a negative range of usefulness, which is the principle reason that causes design engineers to ignore codes that outside developers toss over the fence.

The need for really understanding the intended use before designing a code was brought home to me recently when we received an inquiry from an outside organization that was contemplating the building of a 3-D code to predict aircraft icing, seeking our advice and concurrence. Our response was that we first had to understand the process by which icing issues were addressed in the development of a transport aircraft. We then had to look for opportunities for shortening that process, reducing its cost, and increasing the reliability of its findings, as well as understanding regulatory requirements. Only then could we hope to identify portions of that process amenable to improvement by CFD. This would also identify specific requirements that the CFD must meet in terms of the types of geometry it must handle, in terms of the flow time and cost to set up and run, and in terms of the reliability of accuracy that was needed. The availability of those specifications at the outset is crucial.

In seeking to further improve processes that are already CFD-based, one usually has benchmark data on the performance of that process's baseline CFD components. Improvement can then be rather easily defined in terms of incremental gains to flow time, cost, and accuracy over the baseline CFD components. It seems that continuous improvements of this type are always ongoing and always welcome, which is the path by which the large CFD systems in use by industry evolve to remain competitive and useful over an extended period of time. One example of that type of improvement opportunity arose as a result of the activities of our Boeing Technology Research Center in Moscow. We uncovered a Russian code that handles wing/body geometries and is based upon a full potential and boundary layer formulation, not unlike the standard workhorse Boeing code that is used extensively by us in the early parts of our wing design process. But, the Russian code, because of well-balanced and clever algorithm technology, runs in a few minutes on a workstation, which our code does not. And, it apparently can be pushed further into regimes of mildly separated flows. Those improved attributes may allow it to possibly play roles further upstream in the overall design process, perhaps becoming a key element in a new and improved conceptual or preliminary design process. And if it turns out that its level of accuracy (and, equally important, our ability to do enough work to truly understand and to stabilize the reliability of its accuracy) is comparable with that of our old workhorse wing/body code, then we can contemplate also replacing our workhorse. Since the Russian code is faster and cheaper to run, it appears to offer the potential of reducing the cycle time and cost of executing the wing design process.

And so we come to realize that there is another leading edge within CFD, one that spans the entire airplane design space across the spectrum from the fastest, least precise methods to the most expensive, time-consuming and accurate. That leading edge is sketched in Figure 11 as a relationship between accuracy or variation on one axis and cost/cycle time of carrying out a computational analysis on the other. Algorithm improvements that reduce variation, reduce cost, and reduce cycle time of CFD capabilities anywhere along the spectrum

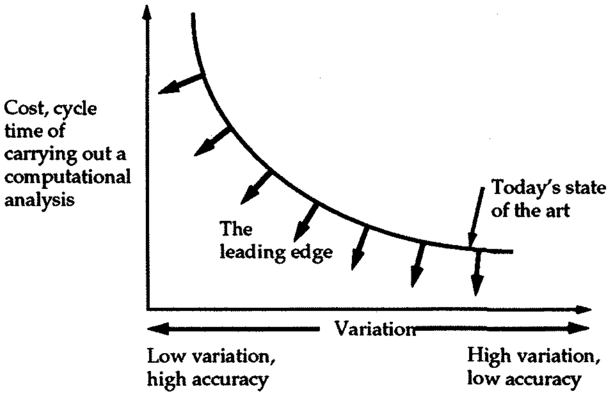

Figure 11. The Leading Edge of Airplane Design Space

are welcome and constitute a valuable improvement to the state of the art.

A Vision of the Future

In the early days of airplanes, the time spent in designing them was measured in weeks or months. But for the past 30 years or so, the process of designing airplanes has been held hostage partly by the aerodynamicist and his tools. Aerodynamic design of something as sophisticated as an airplane has come to be viewed as a process that must take months or even years, and the managements of airplane companies have been trained to think in those terms.

That trend has been aided and abetted by the contributions of CFD developers. They added the need to generate precisely-defined surface geometry lofts, to generate grids, to wait while the computer did its thing, to analyze and understand the entire flow field in great detail, and to acquire highly accurate and detailed surface pressure measurements in nearly all wind tunnel testing.

One serious consequence of that ponderously slow aerodynamic design cycle has been to limit our ability to make the proper cross-functional trades involving aerodynamics, loads, structures, systems and manufacturing. The time and resources needed by the aerodynamicist to arrive at a final aerodynamic loft had become so great that the ability to trade aerodynamic shape and performance for advantages in the other disciplines in rapid fashion during the design process has been quite restricted. As a consequence, the resulting airplane design process became mostly sequential, with the aerodynamicist taking up to a year or two to define the external lines, followed by loads determination and then structural design, with systems people struggling to find space for their systems. Last in line were the fabrication people who figured out how to build the design, and the cost ended up being whatever was necessary.

That is not the type of overall process that is well-suited to the optimization of an airplane's economic performance. It was perhaps better suited to the days when increasing the mission performance of the airplane was paramount, and the cost of achieving it was of secondary importance.

But we are now postured to change all of that. *My vision is to be able to carry out the detailed aerodynamic design of any portion of an airplane within a handful of days at most, and to do it in concert with the loads engineer, the structural designer, the systems person and the manufacturing expert sitting side by side in the same room, with computer systems that talk well with one another.*

In wing design, for example, the current process of balancing the high speed wing design across many points in the flight envelope will be handled by multipoint design optimization with design constraints that encompass the space requirements for systems, which ensure low-cost manufacturability of such things as shot-peenformed wing skins, and which display the sensitivities associated with the various constraints. The cycle time of a design iteration will be about 4 hours, at which time the aerodynamicist will review the constraint costs, in terms of aerodynamic performance, with the structures and systems and manufacturing engineers, and vice-versa. After 4 such cycles, taking about 3 days, one should have a wing that balances across the flight envelope and that is well-balanced in terms of manufacturability and manufacturing cost, and systems.

The next two days or so would be devoted to propulsion system integration. After a baseline run to uncover the aerodynamic problem areas, parallel runs involving design optimization technology would be made, asking CFD to show the designers how to cure the aerodynamic problems in several different ways, such as by recontouring the strut alone, the nacelle alone, the core cowl alone, or the wing undersurface alone. All disciplines, sitting together, would review the impact of those various aerodynamic solutions on their respective disciplines, and make decisions about which blend of options makes sense to pursue. Those design decisions would be input to the next cycle of the process by day's end. The results would be reviewed the next morning, with time for a couple of final iterations by mid afternoon. That evening the final design would be checked at a number of strategic points across the flight envelope, using probably the most expensive and most accurate CFD code. If all turned out well, the design would be frozen the next morning.

That scenario is not as far-fetched as it may seem. Many of the functions just described, such as multipoint design optimization, wing design optimization in the accompaniment of literally thousands of manufacturing constraints, and nacelle integration, accomplished independently with strut shaping, nacelle shaping, core cowl shaping or wing shaping, using constrained optimization methods, have been developed and demonstrated at Boeing. The following discussion provides some of a rapidly growing number of examples.

Figure 12 provides an example of the type of multipoint design optimization that will shorten the process of arriving at a wing design that is properly balanced across several points in the flight envelope. The first part of the figure displays drag rise curves of a baseline RAE 2822 airfoil and of three other airfoils derived from it. Those

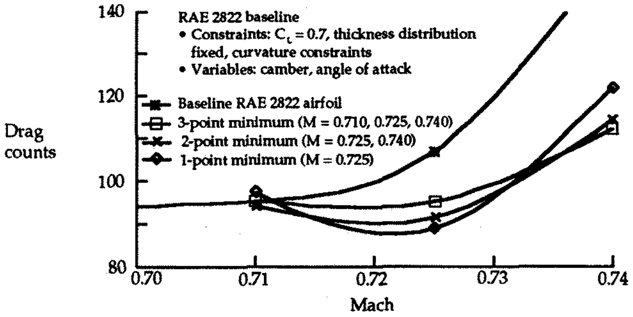

(a) Drag Rise Curves for Airfoils Optimized at One, Two, and Three Mach Numbers

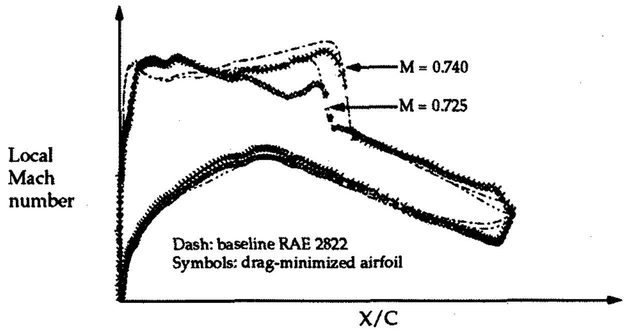

(b) Pressure Distributions for 2-Point Drag-Minimized Airfoil

Figure 12. Designs for Multipoint Weighted Drag-Minimized Airfoils

other three airfoils were derived by minimizing the weighted sum of the drag at one, two, and three points respectively in the flight envelope corresponding to three different free stream Mach numbers and a fixed value of lift. The thickness distribution of all three airfoils was constrained to be that of the RAE 2822 baseline, and there were additional constraints on surface curvature to ensure smoothness and manufacturability of geometry.

The second part of the figure displays pressure distributions of the two-point-optimized airfoil at its two design points. Shock wave strength and drag at both design points are lower than that of the baseline airfoil. These examples demonstrate the achievement of better performance than the baseline over broad regions of free stream Mach number, and provide insight into the trade between performance breadth and point design performance.

Figure 13 displays local Mach number distributions on a wing in the presence of a nacelle installation. The upper part of the figure shows the presence of an undesirable shock wave caused by mounting the nacelle and strut on a wing that originally had good aerodynamic characteristics. The lower part of the figure shows the aerodynamic result of redesigning the inboard wing to eliminate the undesirable aerodynamic characteristics. The computationally interesting aspect of this work is that it involved the use of 13,000 geometry constraints, several at each surface grid point. Those constraints, involving surface curvature, rate of change of curvature, and saddleback growth, were functions of local wing skin thickness. The purpose of those constraints was to arrive at wing skin shapes that can be easily and reliably manufactured with the shot-peening process that Boeing uses to manufacture wing skins. The significance of this is that the total elapsed time for redesigning and arriving at a wing which meets all aerodynamic and manufacturability requirements can now be envisioned to approach the order of one day!

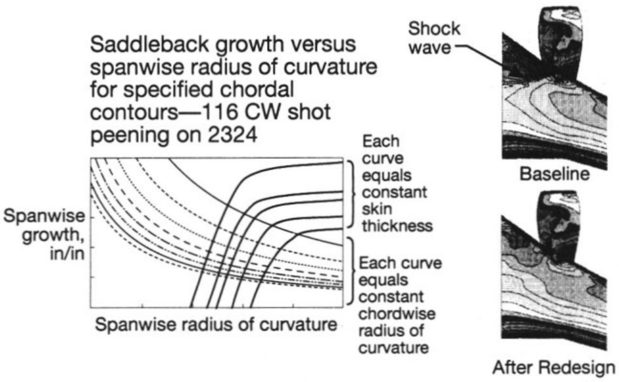

Figure 13. Wing Redesign in Presence of Nacelle With Manufacturing Constraints Involving Limits on Surface Curvature, Rate of Change of Curvature, and Saddleback Growth as a Function of Variable Wing Skin Thickness

Figure 14 shows two alternative design solutions to cure the problem of a shock wave standing between the primary core cowl of an engine and the undersurface of a wing. Part (a) of the figure displays the aerodynamic problem as it exists on one of the older Boeing airplanes. Part (b) of the figure shows the original strut on that airplane and the shape of a redesigned strut that eliminates the problem. The contours of the redesigned strut were constrained to wrap around the primary structure of the original strut. The redesign was carried out using constrained optimization methods, and the total redesign process from beginning to end can be carried out in a few hours.

Part (c) of the figure shows another way of solving the aerodynamic problem, namely by changing the shape of the primary core cowl of the engine that lies beneath the fan flow plume. This, too, can be done in a few hours.

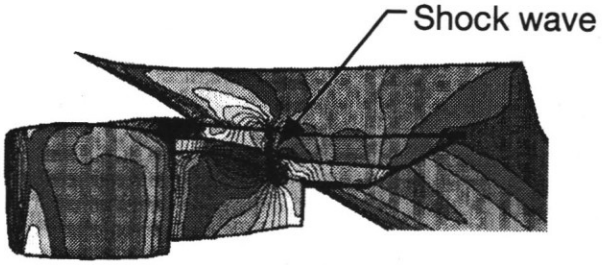

(a) Aerodynamic Problem With Baseline Engine Installation

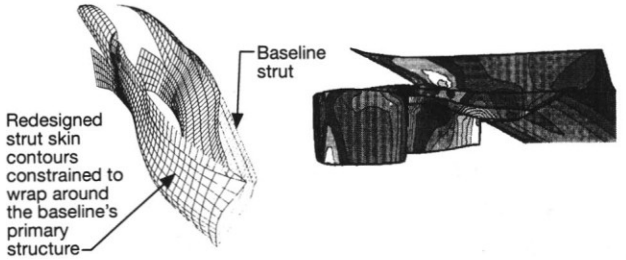

(b) Aerodynamic Problem Eliminated by Optimization-Based Strut Redesign Constrained to Accommodate the Structure of the Original Baseline Strut

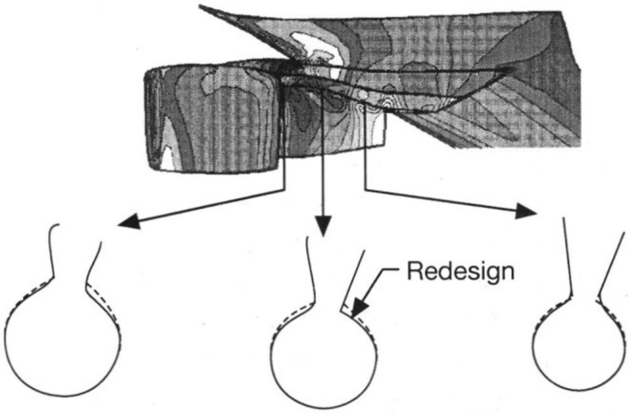

(c) An Alternative Solution–Aerodynamic Problem Eliminated by Optimization-Based Engine Core Cowl Redesign

Figure 14. Alternative Solutions to an Engine Installation Aerodynamic Problem

So, the logical approach to refining the design of a propulsion system integration is to first take a look at independently redesigning each separate piece of adjacent geometry to independently produce good aerodynamic behavior, using appropriate geometry constraints. For a configuration such as that of Figure 14, the candidate surfaces for accomplishing it would be the strut, the core cowl, the outer fan cowl or the undersurface of the wing. Those independent looks can all be done in one evening, so that the next morning the designer has at his/her disposal a number of design solutions. At that point the multidisciplinary team addresses the trades and compromises related to structural implications, volume for systems installation, access doors, manufacturing considerations, etc., and comes to an agreement on which combinations or blends of aerodynamic shaping from among the four independent design solutions seems to represent the best balance among their various disciplines. After the lunch break, the blended design solution candidate and probably some alternates are input and executed.

By the next morning, a primary plus alternate solutions involving blends of strut contouring, core cowl shaping, fan cowl and wing contouring will have been completed, and after final review and selection by the multidisciplinary team the job of designing the installation will have been completed. The entire exercise shouldn't take more than a few days.

Figure 15 shows the outcome of an interesting design exercise aimed at extending the payload/range capability of Boeing 747 aircraft. It is observed that the wing in-

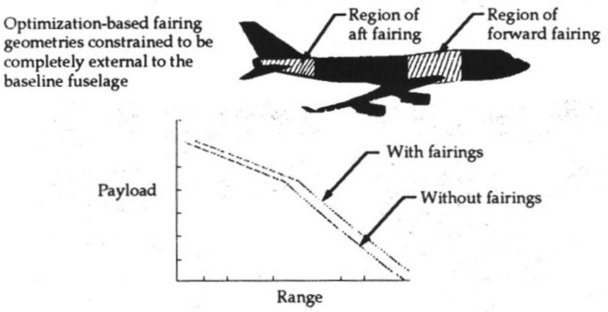

Figure 15. Performance Improvements Achievable With Exterior Fairings on Selected Regions of 747-400 Fuselage

duces supervelocities on the adjacent portion of the fuselage, and that the horizontal and vertical tail surfaces induce supervelocities on each other and on the fuselage. The design problem posed was to design fuselage fairings that are everywhere external to the fuselage of the unmodified airplane, and which are contoured so as to minimize velocity gradients and peak velocities. The resulting fairing installations, whose shapes were arrived at by means of design optimization CFD, were subsequently wind tunnel tested. The results demonstrated significant gains in both range and payload. But perhaps the most significant thing is that this type of design job, from beginning to end, can now be carried out in a couple of hours!

Figure 16 displays another job of this type, namely the reshaping of an overwing portion of a nacelle strut to eliminate a local region of high supervelocity.

Most of the design solutions shown in the previous figures actually took several days to complete, rather than the couple-of-hour times quoted in the vision. That is because we are still learning about the right things to do and how to do them, inventing as we go. Optimization technology is very good at exploiting weaknesses in formulation and at circumventing constraints that are not well thought out, and so a number of tries were usually required to learn how to do things right. But once we learn how to formulate and to properly constrain a particular class of design problem, we can automate the process and thereby reduce it to hours rather than days.

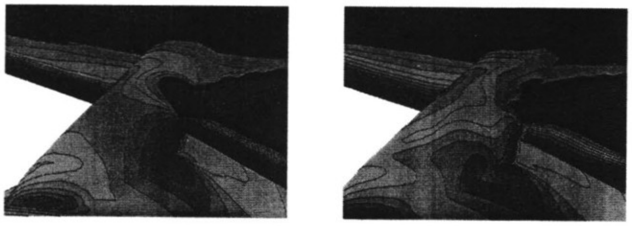

Baseline

Hot Spot Eliminated

Figure 16. An Example of Detailed Design Refinement Using Constrained Optimization

One example of the power of truly fast processes, in which we participated, took place a couple of years ago during the America's Cup boat races between Bill Koch and Dennis Conner. Prior to the races we had developed and validated a *process* for designing winged keels. We knew how long it took and we knew the levels of predictive accuracy, including drag, that the process would produce. The process was fast and streamlined. But we had not designed or optimized a keel.

The races commenced and Dennis started losing. At that point he called us for help and we unleased the process. We proceeded to design and optimize a keel for his boat, using the design and optimization methods that were shown in the previous figures. General Motors built and tested the keel design in their wind tunnel. A mold for the 9 ton keel was fabricated and shipped to San Diego, where the keel was cast and installed on the boat. One of our people assisted in installing and aligning the keel wings. Dennis took the boat out and raced, and won. Not once, but several times in a row, until Bill Koch responded by improving his boat. Bill eventually won, but then proceeded to pay us the ultimate compliment by replacing the keel on his boat with a new one before the final set of races against the Italians. When challenged by the race committee for adopting an untested keel, he reportedly responded that Dennis Conner had already tested it.

The significance of this story is that the total elapsed time from getting the call to winning the first race with the new keel was, I recall, about 25 days. That included design, test, fabrication and installation. The total elapsed time for design and optimization of the keel was about 48 hours, done over a weekend! Those are the kinds of results that are possible by focusing on the processes by which things are carried out, rather than focusing solely on the functionality of the product produced by a process, as has become the norm.

What better evidence is there of the need to change than the fact that most of the research papers presented in this conference show results only in terms of what one or another method ("process") can produce? Most do not even talk about the time required to accomplish what is shown!

We are not yet so good or so fast at designing for flow conditions that involve significant amounts of separated flow such as encountered at landing conditions with flaps deployed. But progress is being made and we understand what we want to be able to do. In due time that challenge will also be conquered.

A key enabling factor in accomplishing the vision will be the appropriate exploitation of information technology. We took a major step forward in the development of the new Boeing 777 transport by implementing digital definition and digital preassembly of the entire airplane (i.e., the traditional "mockup" was done entirely by computer). We also learned that doing aerodynamic design work within a geometry definition and lofting system that is different from the digital definition of the airplane structure, hardware and systems places too great a barrier between the aerodynamicist and the other disciplines. Next time we hope to have in place the geometry and digital bridges that will allow the aerodynamicist, the structures engineer, the systems engineer and the manufacturing engineer, sitting side by side, to make joint decisions in near-real time concerning the tradeoffs between their respective disciplines. That, I believe, is the proper path to approaching what has come to be called multidisciplinary design optimization, or MDO.

Another key element of this vision that is falling into place is that the accuracy, and the reliability of the accuracy, of CFD is becoming good enough to be increasingly accepted by design engineers as a basis for final design decisions. Figure 17 portrays the respective roles of CFD and the wind tunnel today and into the foreseeable future at Boeing. The development of candidate aerodynamic lines for all external parts of an airplane is done today with CFD, as are the trades and interplay with the other disciplines. The role of the wind tunnel is for after-the-fact design validation, and to measure and document the flight and performance characteristics throughout all regions of the flight envelope. Lines freeze of the Boeing 777 wing, including the modifications allowing the folding wing tip option that is featured on that airplane, was done based on CFD, with the wind tunnel

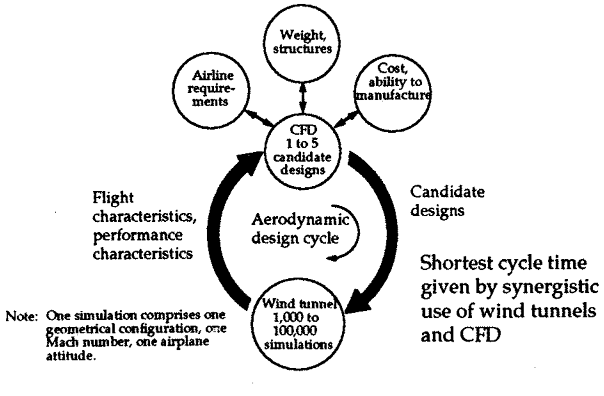

Note: One simulation comprises one geometrical configuration, one Mach number, one airplane attitude.

Shortest cycle time given by synergistic use of wind tunnels and CFD

Figure 17. The Aerodynamic Design Cycle

test of the final design taking place some months after lines freeze.

However, this vision does not imply in any way that wind tunnels will become obsolete. Rather, the roles of wind tunnels and CFD have evolved in ways that are complementary and which properly exploit their respective strengths. Figure 18 displays the characteristics of CFD and wind tunnels in terms of relative cost and flow time. It will be seen that their optimal usage is concentrated at opposite corners of the spectrum.

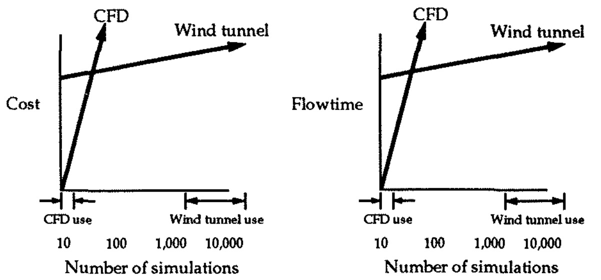

Note: One complete airplane development requires about 2.5 million aerodynamic simulations.

Figure 18. Cost and Flowtime Characteristics of Wind Tunnels and CFD

The strength of CFD is to provide an ability to rapidly and cheaply carry out a very small number of simulations. The strategy that has evolved for exploiting this strength is to concentrate CFD design on those small areas of the flight envelope (Figure 19) in the neighborhood of cruise, takeoff and landing. These are the regions that most influence the operational economics of the airplane, which is why so many of the design decisions involved in optimizing a commercial transport for economic performance can be carried out with CFD.

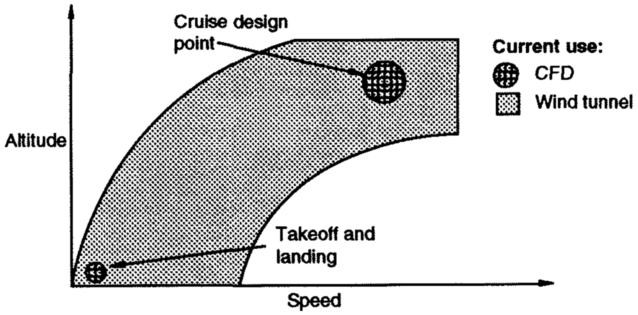

Figure 19. Role of CFD and Wind Tunnels for Simulating the Airplane Operating Envelope

But the airplane's handling characteristics, its ability to recover from upset conditions, and its performance must be accurately known throughout all regions of the flight envelope. That entails literally hundreds of thousands of simulations, with each combination of Mach number, angle of attack, angle of yaw, aeroelastic deflection and control surface positioning constituting a separate simulation. And the results of those simulations must be available early enough in the program to enable the development of flight simulators for handling characteristics evaluation, crew training, etc. The strength of the wind tunnel is to be able to carry out those hundreds of thousands of simulations within acceptable limits of cost and flow time, a task that is unthinkable with CFD.

A typical complete transport airplane development program today involves the conduct of approximately 2 1/2 million aerodynamic simulations. Only a tiny fraction of those are done with CFD. However, that is an extremely strategic and effective fraction because it is the key to returning the time scale of aerodynamic designing to that found acceptable by the Wright brothers and the other aviation pioneers who followed in their footsteps.

It seems that we are doing nothing less than returning to the strategy employed by the Wright brothers, namely to continue to depend on innovation and technology where it adds value to what we set out to do, accompanied once again by rapid processes for design and manufacturing. History is showing that the Wright brothers laid out a path from which we have strayed and must now return. They obviously knew a lot more about the airplane business than we give them credit.

REFERENCES

1. Smith, A. M. O., "High-Lift Aerodynamics, 37th Wright Brothers Lecture." AIAA Paper No. 74-939, 1974.

2. Imai, Masaaki, Kaizen. McGraw-Hill Publishing Company, 1986.

3. Hammer, M. and Champy, J. Reengineering the Corporation: A Manifesto for Business Revolution. Harper Collins Publishers, Inc., 1993.

4. Young, D. P., Melvin, R. G., Bieterman, M. B., Johnson, F. T., Samant, S. S., and Bussoletti, J. E. "A Locally Refined Rectangular Grid Finite Element Method: Application to Computational Fluid Dynamics and Computational Physics". Journal of Computational Physics, Vol. 92, pp. 1-66, 1991.

5. Bieterman, M. B., Bussoletti, J. E., Hilmes, C. L., Johnson, F. T., Melvin, R. G., and Young, D. P. "An Adaptive Grid Method for Analysis of 3D Aircraft Configurations". Comput. Meths. Appl. Mechs. and Eng., Vol. 101, pp. 225-249, 1992.

6. Huffman, W. P., Melvin, R. G., Young, D. P., Johnson, F. T., Bussoletti, J. E., Bieterman, M. B., and Hilmes, C. L. "Practical Design and Optimization in Computational Fluid Dynamics". AIAA 24th Fluid Dynamics Conference, Paper AIAA 93-3111, July, 1993.

7. Bussoletti, J. E., Johnson, F. T., Bieterman, M. B., Hilmes, C. L., Melvin, R. G., Young, D. P., and Drela, M. "TRANAIR: Solution Adaptive CFD Modeling for Complex 3D Configurations". Pp. 10.1-10.14 in Proceedings of the 1993 European Forum - Recent Developments and Applications in Aeronautical CFD, Royal Aeronautical Society, London, England, September, 1993.