PERSPECTIVE

If we designed airplanes like we design drugs…

Walter S. Woltosz

Abstract In the early days, airplanes were put together with parts designed for other purposes (bicycles, farm equipment, textiles, automotive equipment, etc.). They were then flown by their brave designers to see if the design would work—often with disastrous results. Today, airplanes, helicopters, missiles, and rockets are designed in computers in a process that involves iterating through enormous numbers of designs before anything is made. Until very recently, novel drug-like molecules were nearly always made first like early airplanes, then tested to see if they were any good (although usually not on the brave scientists who created them!). The resulting extremely high failure rate is legendary. This article describes some of the evolution of computer-based design in the aerospace industry and compares it with the progress made to date in computer-aided drug design. Software development for pharmaceutical research has been largely entrepreneurial, with only relatively limited support from government and industry end-user organizations. The pharmaceutical industry is still about 30 years behind aerospace and other industries in fully recognizing the value of simulation and modeling and funding the development of the tools needed to catch up.

Forty years ago this year, fresh out of graduate school, I worked in the aerospace industry as a young engineer (actually, "rocket scientist"). My first aerospace job in 1971 was developing a computer program to simulate and optimize the ascent trajectory of the space shuttle to get the most

W. S. Woltosz (✉)
Simulations Plus, Inc., Lancaster, CA, USA
e-mail: walt@simulations-plus.com

payload into an orbit 150 nautical miles above the equator. The program consisted of six boxes of punch cards (2,000 per box) written in Fortran 4. The simulation included a rotating earth, modeled as an oblate spheroid (to calculate how gravity changes with both altitude and latitude), a NASA standard atmosphere model with the ability to simulate hot and cold days as well as wind profiles at different altitudes, and the numerical solution of all of the differential equations necessary to track the latitude, longitude, and altitude of the space shuttle over the rotating earth during flight.

Our primary computer was a Univac 1108 that filled a specially air-conditioned room occupying probably 1,000 square feet. A single run of the program simulated hundreds of trajectories as it worked to optimize the steering commands in pitch and yaw. Different ascent trajectories result in different complex interplays of thrust, drag, and gravity as the vehicle moved from the dense air at sea level to the near-vacuum of space, turning first southward toward the equator, then back to the east to end up in an equatorial orbit, all the while getting lighter as the fuel is expended, with a moving center of gravity, and quickly achieving supersonic flight where a complex pattern of shock waves forms around the orbiter, external tank, and booster rockets. The objective of optimizing the steering commands is to minimize the losses due to aerodynamic drag and gravity during the climb, so that fuel is used most efficiently to put as many pounds of payload into orbit as possible. Finding an optimized steering solution for a particular set of system weights and environmental conditions (hot day, cold day, winds, etc.) required 12–14 h of CPU time. My laptop today would execute the same equations and achieve an optimized solution in a fraction of a minute. But consider this—*NASA was funding this simulation and modeling effort 10 years before the first launch, and even before the vehicle design was finalized.*

By the way, the computer program for the space shuttle resulted in *accidentally* discovering the shuttle's signature roll maneuver—the program told us to roll it over on its back and fly upside down and we'd pick up about 8,000 pounds of free payload! A combination of more efficient aerodynamics and the required thrust vector angles for the main (liquid hydrogen/liquid oxygen) engines to balance the moments around the moving center-of-gravity provided the increase. Ask many astronauts today why it flies upside down and most can't tell you. *By developing this complex (for its day) simulation, we assembled many bits of known information and theory to discover something that was unknown prior to having the simulation capability, and that turned out to be extremely valuable.*

We used complex computer simulations in all of my aerospace jobs in the 1970s and 1980s, both to gain insight into our systems and to optimize their designs. They saved immeasurable time and money as we were able to try and fail an enormous variety of ideas quickly with no actual losses by iterating through countless virtual design options to find an optimal solution. That was 30–40 years ago!

Today, aerospace, automotive, electronics, and other industries routinely incorporate far greater detail into simulations and run them much faster to answer questions like: how much load a metal or composite part can take before bending or breaking, how aerodynamic surfaces might flutter at certain speeds and angles of attack, how the stresses in a very hot jet engine turbine blade will be affected as a fighter jet does a high-G maneuver, and many more complex behaviors. Numerical methods have evolved to speed up calculations along with the speed and memory improvements in computer hardware. In fact, today we even have "electronic wind tunnels" that can provide an accurate estimate of the aerodynamic forces on complex vehicles without the need to build scale models and test them in real wind tunnels. The solution of the Navier–Stokes equations for fluid flow was once considered so complex that even supercomputers were not expected to solve them with a fine enough grid to calculate such forces accurately.
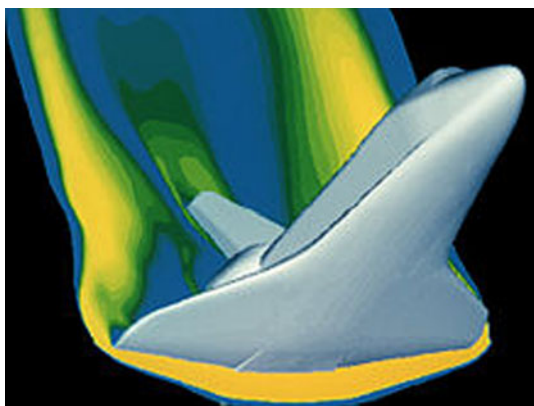
This Computational Fluid Dynamics (CFD) computer-generated image shows a model of the space shuttle. CFD has taken the place of wind tunnels for many evaluations of aircraft and, as computing power increases and computer models become more sophisticated, CFD will largely replace wind tunnels. Image credit: NASA.

Many major aerospace projects require a decade or more of R&D along with over a billion dollars in investment to get a single new product to market. Sound familiar? Most of these projects will involve *hundreds of millions of dollars spent on simulation and modeling*. That part does not sound familiar to pharmaceutical scientists. Yet the risks and costs associated with pharmaceutical R&D are on a level not so different from large development projects in aerospace, automotive, and other industries.

Why has pharmaceutical research and development lagged so far behind other industries in the development and application of simulation and modeling for research and development? I believe there are at least three main factors:

1. Aerospace simulation and modeling software involves complex physics and chemistry. *Pharmaceutical science adds biology to these, which increases variability and complexity.*

2. Aerospace simulation and modeling generally uses well-established inputs that are measured with relatively high accuracy and relatively small variance. So even though many inputs are required for a simulation as complex as the launch of a Space Shuttle to orbit, the inputs for the vehicle itself are well-known. The reason they are well-known is because the industry goes to the time and expense to perform experiments to get parameter values whose sole purpose is to be used as inputs to simulations. For example, the payoff from exhaustive (and expensive) wind tunnel experiments at a wide variety of altitudes (ambient pressures), angles of attack, and vehicle configurations (position of landing gear and flaps, etc.) is that they make possible accurate simulation of the entire flight envelope for a new vehicle. As opposed to the vehicle inputs, those for the environment (atmosphere) are less predictable, but like population virtual trials, those conditions are handled with Monte Carlo simulations that vary the conditions over the expected range of winds at various altitudes, temperatures, and air densities. The pharmaceutical industry largely does not yet recognize that *taking data with simulation and modeling in mind* could radically change the failure rate that plagues drug discovery and development. Pharmaceutical simulation and modeling typically has to use inputs that have been measured in a way that sacrifices accuracy for high-throughput go/no-go decisions.

3. Aerospace researchers are primarily engineers, who are trained to be generalists, integrating a variety of

disciplines as they approach problems (e.g., stress analysis, flight dynamics, heat transfer, thermodynamics, internal ballistics, solution of differential equations, computer programming, and other disciplines all enter into a typical rocket motor simulation of the 1970's). Pharmaceutical researchers are primarily scientists who are trained in narrow silos to become specialists rather than generalists. *It takes good generalists to get maximum benefit from system simulations that incorporate diverse areas of science including, for example, physiologies of various animals and human populations, formulation issues,* in vitro—in vivo *correlation methods, metabolism, transporters, solubility effects, fasted and fed state differences, numerical optimization, and machine learning methods.*

4. The knowledge base for many aspects of physics and chemistry provides a foundation upon which to build good mechanistic models of the systems under study. The knowledge base for pharmaceutical research is growing at an encouraging rate, yet today there remain many kinds of information that are not yet available for certain kinds of simulation and modeling. An example is the level of expression of various transporters in different tissues in different species, and their variances in populations. *Fortunately, progress has been made over recent decades and a very useful (if yet far from complete) knowledge base now exists upon which to build useful simulation and modeling tools.*

To say that pharmaceutical science is too difficult or lacks a sufficient knowledge base to build good mechanistic models was an argument 20–30 years ago—it is no longer a valid argument today. George Box is often quoted, "All models are wrong, some are useful". Current simulation and modeling tools for drug discovery and development are very useful, and improvements are coming rapidly. Yet their adoption has taken 10–15 years to reach current usage levels, which remain well below where they are in other industries.

There is no greater productivity tool than software. I'll say that again—*there is no greater productivity tool than software*! Don't believe it? Try working without your word processor, databases, spreadsheets, presentation software, e-mail, Internet, and so on. We use (and usually take for granted until we have a glitch) software in more ways than we realize. Managers have no reservations about purchasing software for such uses. Fortunately, these types of software are developed for millions of users, providing an economy of scale that supports sophisticated software at prices that even retirees can afford.

The development costs for sophisticated simulation and modeling software are very high—in the tens millions of dollars for the more sophisticated programs. Like simulation and modeling software used in aerospace and other industries, the number of users is relatively small, making the cost per user much higher than for commercial software sold in huge quantities. Specialized teams of scientist/programmers (or separate scientists and programmers—but that's a subject for another article) spend many person-years bringing such software to commercial standards and providing ongoing support and enhancements. Many senior pharmaceutical R&D managers grew up without exposure to the benefits of such tools, resulting in skepticism and difficulty recognizing the benefits of the insight they provide.

Many pharmaceutical scientists are skeptical of computer software that combines mathematical relationships in ways that are so complex that no human can grasp the full interplay of the equations and logic involved. Even the simplest mechanistic simulation of oral absorption and pharmacokinetics can involve hundreds of interacting differential equations. Aerospace researchers expect this and don't expect to be able to intuit results by examining a few key inputs. I have listened to numerous presentations at scientific meetings where speakers drew sweeping conclusions from a few parameters like polar surface area, logP, and molecular weight, with heads in the audience nodding in agreement with the speaker. Sorry, folks, but it's just not that simple! If it was, we'd be releasing new drugs every week!

Correlating molecular structures with a wide variety of activity and ADMET properties remains a particular challenge. If the often-heard number of potential drug-like molecules is truly on the order of 10^{62}, then clearly there is no hope for humankind to ever investigate all of them. In fact, the vast number of possible interactions among atoms within molecules and between molecules and their environments make it highly unlikely that humankind will derive purely mechanistic (and quantum) methods for predicting most properties from structure with experimental accuracy in the foreseeable future. "Activity cliffs" or "property cliffs" (minor changes in molecular structure that result in large changes in activity or other properties) are seen regularly that defy similarity (Tanimoto, nearest neighbor) rules and chemists' intuition. Changing a single atom in a molecule will affect every property of that molecule—not only affinity for the target, but all of its physicochemical and biopharmaceutical properties, metabolism, potential toxicities, etc. We simply don't have a large enough knowledge base from which to build ab initio models for every property of concern, so empirical methods are going to be around for a very, very long time in this part of pharmaceutical science. So what are we to do to get out of the "make and test" mode and go more into "design, then make" mode that other industries enjoy?

Fortunately, the availability of cheap computing power has enabled the evolution of powerful machine learning methodologies. Some of these methods are so complex that they defy direct human comprehension, but they have proved to provide the most accurate structure–property predictions available to date. Pharmaceutical scientists need to accept that there are tools that exceed the ability of humans to grasp in simple terms. In the aerospace industry, the use of computer numerical control (CNC) machining has enabled the manufacture of complex shapes from blocks of metal that would have been extremely difficult to machine by hand using the methods of a few decades ago. The use of composite materials has reduced weight while increasing strength, but the new science required to model such structures (which behave quite differently than metals) had to be developed.

In 1982, I developed a computer program to calculate the optimum filament winding angles for the first graphite/epoxy rocket motor. We knew enough from experiments to know how the composite material would behave in terms of both longitudinal (the length of the motor) and hoop (in the direction of the diameter of the motor) stresses. Making a rocket motor from a long filament of carbon fiber and a pot of epoxy had not been done before. The fiber has to be continuous from one end of the motor to the other to support the longitudinal stresses that try to stretch the motor when the pressure inside goes from zero to thousands of pounds per square inch in a fraction of a second. The fiber, wetted with epoxy, is wound over a mandrel, around each end and returns in barber-pole fashion, but at a relatively shallow angle. After the longitudinal winds are completed, a second process begins with a filament wound around the motor diameter in consecutive circles (hoops) to give it the strength it needs in the hoop direction. After all the windings are completed and after a curing process in an autoclave, the mandrel is removed and the case is ready for loading with propellant. We did exactly that, added the nozzle, and fired the motor—the world's first graphite/epoxy rocket motor, designed by a computer program, was a complete success.

We have to accept that we can learn things we don't know from things that we know—i.e., we may know a large number of individual facts and measurements (and molecular and atomic descriptors), but we don't know how they all interact. The pharmaceutical industry spends millions of dollars filling databases, spreadsheets, and reports with bits of diverse data, but with rare exceptions, no one can look through all of the data on a compound and tell you how it's going to behave in vivo. Only through highly sophisticated simulation and modeling do we have a chance at gaining insight into how all of the diverse properties play together.

The state-of-the-art today for predicting most properties from structure yields root-mean-square errors on the order of 0.3–0.5 log units, or about two-to three-fold. Aerospace designers would have a very tough time with such high uncertainty, since the safety factor added to some calculations is often as low as 1.4! Imagine if a part needs to carry a stress of 1,000 pounds, and is designed stronger to fail at an estimated 1,400 pounds in order to provide that safety factor of 1.4. If the actual error in the stress capability was 0.3 log units lower than expected, then the failure would occur at 700 pounds when it needed to handle 1,000, and disaster would follow. Fortunately, we don't have to deal with such tight safety factors for most drugs.

Every drug that fails in a clinical trial or after it reaches the market due to some adverse effect was "bad" from the day it was first drawn by the chemist. State-of-the-art *in silico* structure–property prediction tools are not yet able to predict every possible toxicity for new molecular structures, but they are able to predict many of them with good enough accuracy to eliminate many poor molecules *prior to synthesis*. This process can be done on large chemical libraries in very little time. Why would anyone design, synthesize, and test molecules that are clearly problematic, when so many others are available that can also hit the target? It would be like aerospace companies making and testing every possible rocket motor design rather than running the simulations that would have told them ahead of time that disaster or failure to meet performance specifications was inevitable for most of them.

The pharmaceutical industry is undergoing an awakening with respect to simulation and modeling tools. You can see it in the titles of presentations and posters at major scientific meetings, and in the number of smaller meetings with a strong focus in these areas. I predict that the day will come (probably not in my lifetime, but it will come) when pharmaceutical research and development will be so heavily driven by simulation and modeling tools that many fewer failures will occur in clinical and preclinical phases. Discovery efforts guided by de novo design tools available now offer the promise of more rapid discovery of good lead compounds and elimination of the majority of "losers" without the need to make and test them. Exploration of very large compound libraries automatically is already underway in a few organizations. I believe that many others will come to realize that simulation and modeling tools, properly applied, repay their costs many times over. A commitment to the investment in *developing skilled generalists* and *supporting the development of the tools* themselves has been made in relatively few organizations to date. Hopefully, the awakening we've seen in recent years will continue to grow and senior management will recognize that simulation and modeling does not cost money—it saves

considerable time and money. For an industry that requires a long-term view of research and development, recognizing the value of predictive tools would seem to be a no-brainer.