1

Introduction

The general theme of the book is to reproduce the research and insights that led the authors through their seminal studies into airplane control allocation. There is much research remaining to be done in the field of control allocation, and by following the thinking that preceded the fruitful directions taken by other researchers, new areas of inquiry will be opened.

The authors defend their geometrical approach to visualizing the problem as one that provides greater insight into the mechanisms of the methods of solution that exist or may be contemplated. This is particularly true when considering the processes of reconfiguring the controls following the identification of a failure.

It is emphasized that the primary interest of the authors and the focus of the book is airplanes. Thus, we stick to a relatively small number objectives in the allocation problem, corresponding primarily to the three rotational degrees-of-freedom of airplanes and secondarily to the linear degrees-of-freedom. We acknowledge that there are many other fields that have similar problems, and believe our research lays a sound basis for other researchers to modify our results to apply them to their particular interests.

With respect to rigorous mathematical proofs, none will be found here. The authors are not mathematicians, as will be readily confirmed by any real mathematician who picks up this book. We certainly never thought before embarking on this research that ‘null space’ and ‘airplane’ would ever be used in the same sentence. We typically began by sketching a two-dimensional figure on the blackboard, something that seemed ‘intuitively obvious’ to us, and wondering if that figure generalized to higher dimensions.

Most important results have been proved in other sources: the many technical papers, theses, and dissertations that arose from our research, or in textbooks, particularly books that deal with linear algebra. Many of these publications are presented in Appendix C. Here we will just make claims that we are pretty sure are true. For instance, rather than prove that convexity is preserved under the mappings we describe, we will just assert it and perhaps give compelling evidence of its truth.

1.1 Redundant Control Effectors

The origins of our research into airplane control allocation lay in earlier research into model-following and dynamic-inversion control laws. The nature of model-following and

---

Aircraft Control Allocation, First Edition. Wayne Durham, Kenneth A. Bordignon and Roger Beck.

Companion website: www.wiley.com/go/durham/aircraft_control_allocation

dynamic-inversion algorithms is such that one is required to find a vector of control effector deflections that yield a desired moment, force, or acceleration. With three moments and three controls, the answer for a linear problem is a trivial matrix inversion. The physical limits of the control effectors does not affect the solution since the solution is unique. That is, for an airplane with ganged ailerons, a rudder, and an elevator, the combination of these effectors that will generate a specific moment vector is unique; if one or more saturates then the problem is not in the math but in the hardware.

Early problems arose when considering an airplane whose horizontal tails were not ganged to generate pitching moments only, but that could be displaced differentially as well to generate rolling moments and, unintentionally, yawing moments. By considering the left and right horizontal tails as independent we now have four control effectors for the three components of the moment vector to be generated. The linearized control effectiveness matrix (to be defined in Eq. (2.20)) is no longer square, but has three rows and four columns.

Figure 1.1 depicts a variety of control effector types. The airplane is USAF S/N 71-0290, the F-15 ACTIVE (Advanced Control Technology for Integrated Vehicles). The canards and horizontal tails are all-moving surfaces. The two vertical stabilizers have hinged rudders at their trailing edges. The wings have trailing-edge ailerons and flaps. Finally, both engines have axisymmetric thrust vectoring capabilities. Each of these various control effectors is capable of independent action.

Redundant control effectors are employed to extend an airplane's performance envelope, typically in the low-speed regime. Thrust vectoring generates moments long after conventional

Figure 1.1 F-15 ACTIVE

flapping control surfaces have lost effectiveness at low dynamic pressure. Thrust vectoring enhances the dog-fighting potential of the F-22 Raptor, and permits maneuvers such as Pugachev's Cobra to be performed in other aircraft.

Clever control allocation is not needed in high-speed flight, where more than ample forces and moments can be generated with small effector deflections. In low-speed flight aerodynamic effectors lose effectiveness and must be combined with other effectors (aerodynamic or propulsive). However, if there are more control effectors than moments or accelerations to be generated, methods of allocating these controls are needed.

We are now faced with a 'wide' control effectiveness matrix: more columns than rows. As we will see there is a simple mathematical way to 'invert' such matrices. The real problem arises when the physical limitations of the control effectors are considered. In other words, control effectors have hard deflection limits that cannot be exceeded. When simple mathematical solutions to the problem are used, it is possible for one or more effectors to be *unnecessarily* commanded past its limits, meaning that it will saturate. When a control effector is saturated, the assumptions on which the flight control system was based are no longer valid.

1.2 Overview

We will begin by discussing aircraft flight dynamics and control. This will consist of a very brief overview of flight dynamics and its nomenclature, offered to provide the reader with explanations for some of the terms used subsequently.

Next we will spend some time describing dynamic inversion control. This form of control lends itself naturally to the control allocation problem as we have posed it. We will briefly discuss 'conventional' control, and even more briefly mention model-following control. All three forms of control law determination have some need of control allocation.

After formulating the problem, we address the geometry of control allocation. We do this first considering two-moment problems. Two-moment problems have application in aircraft control, since often the lateral-directional problem (rolling and yawing) is treated separately from the longitudinal problem (pitching). Moreover, it is much easier to make figures on the page of two-dimensional objects than of three-dimensional ones.

The geometry of the three-moment problem is a natural extension of that of the two-moment problem, and it is discussed in detail. For each of the two- and three-moment problems a metric is offered that permits comparison of different control allocation methods for their effectiveness in solving the problem. This gives rise to the idea of a 'maximum set' of moments that can be generated using different control allocation schemes, and its importance is discussed.

A large section on solution methods follows. We explore all the allocation methods of which the authors have first-hand experience, and most are accompanied by numerical illustrations. One of the control allocation methods—linear programming—is briefly discussed. Because of the current interest among researchers in the subject of linear programming solutions, there is a separate section (Appendix A) that further explores linear programming in greater detail.

All the preceding has been based on a global problem: the total set of control deflections that yield the whole moment vector. Now we turn our attention to a local problem. Digital flight control computers solve the allocation problem scores or even hundreds of times a second. We look within one frame of the computer's operation and consider not just how far an effector can move, but how fast. This permits us to incorporate rate limits into the problem. This framewise allocation comes with a serious drawback sometimes called 'windup'. The remedy to windup is not hard and comes with some beneficial side effects.

Next we briefly explore control allocation and flight control system design. Example designs are given for a roll-rate command, pitch-rate command, and a sideslip controller. Finally, the consequences of using a non-optimal control allocation method are graphically displayed.

At the end of the text there is a chapter on some of the real-life applications of the previously described research. Lessons learned from the design of the X-35 control system are presented.

Throughout the book we occasionally make reference to simulation code. The MATLAB®/Simulink® based code is available at a companion website to this book, and it is fully explained in Appendix B. The simulation code offers different modules that implement the various control allocation methods described in the book. Readers are free to adapt and use this code to further explore the concepts of control allocation.

We feel that a common simulation source is essential to creating reproducible results. Many technical papers present simulation results with insufficient information about that simulation for the reader or reviewer to reproduce them. There are many assumptions inherent in one researcher's simulation code that cannot be conveyed in a brief paper but that may greatly affect the results one obtains.[1] The simulation code we provide came to us courtesy of the Swedish Defence Research Agency with their permissions. MATLAB®/Simulink® code is available in student and academic editions and should be very widely accessible.

The final appendix is an annotated bibliography, in which we clean out our files of control allocation and dynamic inversion papers and list them, along with their abstracts or other descriptive material. This appendix is a good place to look to see if anyone is pursuing interests related to yours. It is inevitable that some have been overlooked, either through our inattention or the fault of some search engine or other. Our loosely-enforced cut-off criterion for inclusion in this list was that the source be refereed. Conference papers were generally not included unless the material was unique and relevant.

Finally, we wish to emphasize that the content of this book reflects only topics with which the authors are personally familiar. It is not a survey of the control allocation literature. There are many very good and sound areas of control allocation research that we have not addressed, except as indicated in the annotated bibliography (Appendix C). There one will find works by Marc Bodson, Jim Buffington, Dave Doman, Dale Enns, Tony Page, and many others with whom the authors have had collegial exchanges.

In particular we appreciate a group familiarly known as 'Bull studs': John Bolling, Josh Durham, Michelle Glaze, Bob Grogan, Matt Hederstrom, Jeff Leedy, Bruce Munro, Mark Nelson, Tony Page, Kevin Scalera, and the others who toiled in the 'Sim Lab' to help make sense of all this. And lastly Fred Lutze, who pondered the various stages of our progress and occasionally said, 'You *could* do that, but it would be wrong.'

---

[1] For example, whether one begins the calculations in a given frame assuming that the last commanded controls have been achieved, or using the actual deflections that resulted. It can make a big difference. Most of what this statement means will be made clear in Chapter 7. See Bodson and Pohlchuck (1998) for some more insight.

References

Aerodata Model in Research Environment (ADMIRE), Ver. 3.4h, Swedish Defence Research Agency (FOI), Stockholm, Sweden, 2003.

Bodson, M and Pohlchuck, E 1998 ‘Command Limiting in Reconfigurable Flight Control’ *AIAA J. Guidance, Control, and Dynamics*, **21**(4), 639–646.