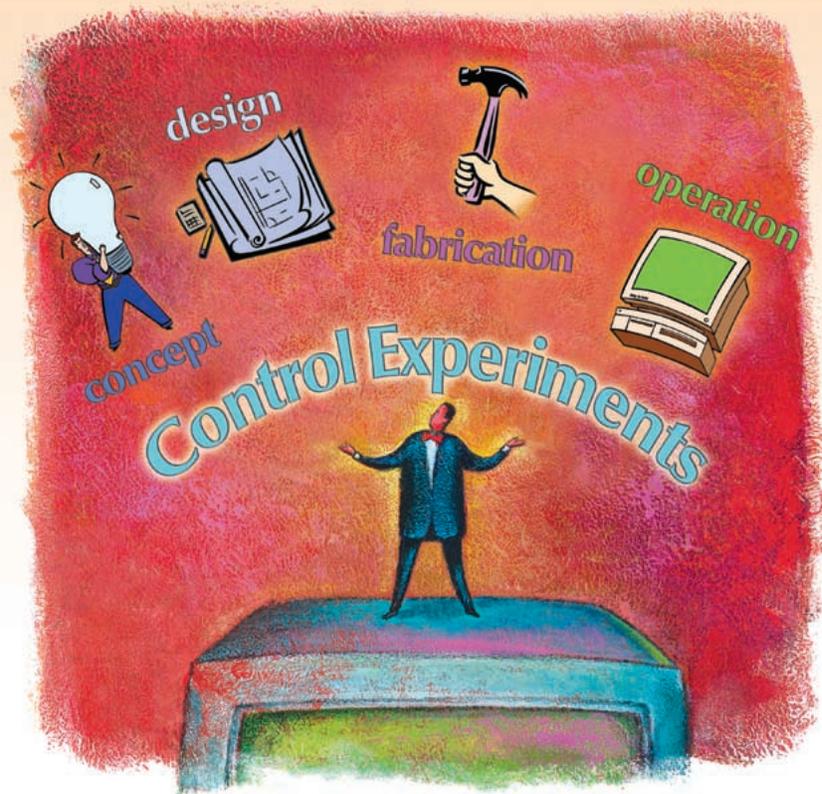


A Formation Flight Experiment

Constructing a testbed for research in control of interconnected systems.



©MASTERSERIES & EYEWIRE

By Jeffrey M. Fowler and Raffaello D'Andrea

This article discusses the design process for a formation flight experiment constructed at Cornell University. The primary motivation for the testbed is to explore control design and analysis for interconnected systems. These systems are comprised of many similar units that interact directly with their nearest neighbors and that have sensing and actuating capabilities at every unit. The resulting interconnected systems often display rich and complex behavior, even when the units have tractable models and interact with their neighbors in a

simple and predictable manner. In addition to formation flight, there are many examples of such engineered systems, including automated highway systems [1], satellite constellations [2], cross-directional control in paper processing applications [3], and microcantilever array control for massively parallel data storage [4]. In this class one can also consider lumped approximations of partial differential equations, including the deflection of beams, plates, and membranes and the temperature distribution of thermally conductive materials [5] embedded with sensors and actuators.

Among the challenges associated with controlling these types of systems, the two most important are

- Dealing with a large state space and a large number of inputs and outputs. This combination typically leads to a computationally demanding, if not intractable, control design problem.
- Issues of implementation. It is typically not feasible to control these types of systems with centralized algorithms, the result of most optimal control design techniques, as these require high levels of connectivity, impose a substantial computational burden, and are typically more sensitive to failures and modeling errors than decentralized or distributed schemes.

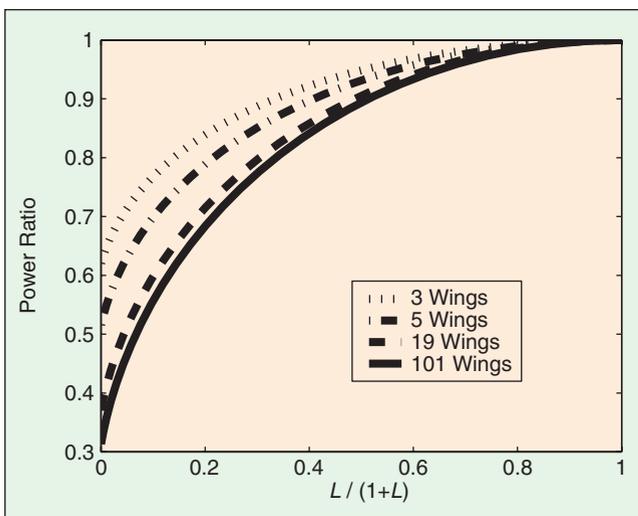


Figure 1. The drag reduction ratio for formation flight. The ratio of the total induced drag of a formation to the total induced drag of vehicles flying independently is a function of the lateral geometry of the formation [12]. Note that near the left edge, where induced drag is reduced the most, the induced drag is very sensitive to the lateral separation of the wings. L is the wing-tip separation in units of span, so the left edge of the graph corresponds to the port wing-tip of one wing aligned with the starboard wing-tip of the next. An elliptical lift distribution is assumed.

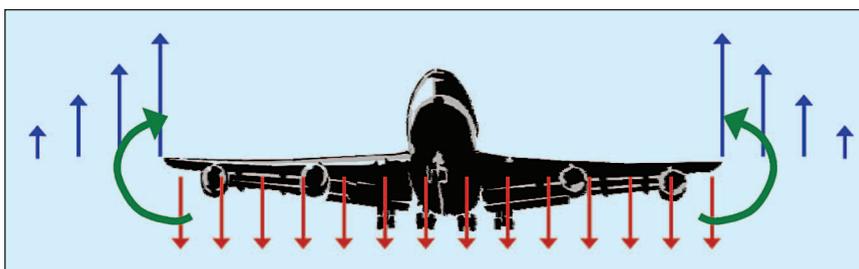


Figure 2. A representation of down-wash and up-wash created by an aircraft and the associated shed vortices. The red lines denote the down-wash, the blue lines the up-wash, and the green arcs the shed vortices.

There is a substantial body of literature devoted to this class of problems [6]–[11], and it is an active area of research.

We are investigating these issues in the context of an experimental platform representative of airplanes flying in a vee formation. Airplanes flying in such a formation would enjoy a substantial reduction in induced drag, quantified in Figure 1, extending the range of a squadron or allowing solar powered airplanes to stay aloft indefinitely [13], [14]. This reduction in drag is thought to be the reason geese fly in formation [15].

The essential mechanism is captured in Figure 2. An aircraft generates both up-wash and down-wash when it is flying. In particular, in steady flight the down-wash is concentrated within the span of the aircraft, while the up-wash is concentrated near the wingtips and quickly decays away from this location. If a downstream aircraft can fly close enough to an upstream aircraft, it can use the generated up-wash to reduce the amount of power required to stay aloft. Associated with down-wash and up-wash are wing-tip vortices, which can be described as weak “tornadoes” that can be extremely energetic and persist for long periods of time. The presence of these vortices is one reason for large spacings between aircraft landing or taking off consecutively at airports [16].

In actual experimental results with two F/A 18 aircraft, Hummel demonstrated a 15% fuel savings for the downstream aircraft [17]. No fuel savings were reported for the upstream aircraft; this is consistent with potential flow models of flying in formation [18], which predict that the benefits of flying in formation are mainly realized by downstream aircraft if the aircraft separation, in the stream-wise direction, exceeds several wingspans. The same aerodynamic coupling that provides the drag benefit can also introduce spatial instability into the system [13], [14]. In addition, Figure 1 demonstrates that minor changes in the lateral spacing of the formation can incur substantial costs in induced drag.

Purely decentralized control, in which the formation-keeping autopilots of the various airplanes do not communicate, may fail to counteract this instability where control

based on more information about the structure of the system would succeed. A weak disturbance on a craft near the front of the formation can result in gross oscillations a few vehicles back, preventing the formation from enjoying much of the potential gain in efficiency. A centralized controller could theoretically deliver stability and performance but would impose high costs in computation and communication. Furthermore, a centralized controller would have to

be tailored to each formation: if the number of wings in a formation were to change, the airplanes would have to transfer to a completely different controller. By applying a distributed controller that is interconnected in the same way as the plant, we seek to address these liabilities while maintaining performance close to that of the centralized controller, as summarized in Figure 3.

We have chosen airplane formation flight as a problem mainly because of the intrinsic importance of this application and also because it is inherently a difficult and challenging problem. In particular, due to the complexity of vortex generation and shedding and the resulting interaction between a shed vortex and a downstream lifting surface, it is difficult to obtain a first principles *quantitative* model that accurately captures the dynamics of the overall system (see [20], for example).

Successfully controlling airplanes that fly in formation will entail a complete control design cycle: first principles modeling to understand the underlying physical phenomena, system identification to fit the experimental data to a physically motivated model, model-based control design, control system implementation, and finally, assessing the performance of the controller. Note that in practice this design cycle is iterative rather than sequential in nature. Thus, even though our main objective is to develop and test new algorithms for controlling interconnected systems, we feel that the best way to achieve this objective is to embed control algorithm development within a complete design cycle. This embedding must be tempered with the requirement that the resulting experiment be easy to work with to maximize the time spent on control design and analysis and minimize the time spent on operating and maintaining the experiment.

Design of the Experiment

The goals for the design of a control experiment are considerably different from product-oriented design objectives. In a control experiment, the physical system may be made purposely more difficult to control to examine more sensitively the performance of a control scheme under study. Phrased another way, design decisions for creating a high-performance, closed-loop system tend toward an open-loop system that is easier to control, while design decisions for studying a high-performance control scheme tend toward an open-loop system that is harder to control. As discussed in [21], a system attribute that often leads to difficult control design problems is simultane-

ity of multiple impediments. Designing a high performance *and* structured control system, distributed or decentralized, in the presence of system uncertainty, unmodeled dynamics, and noise falls into this class of problems.

It should be stressed, however, that the mechanisms that render the control design problem difficult cannot be arbitrary in nature; they must either conform to the control design methodology being investigated or be representative of the real system which is motivating the experimental abstraction. In an ideal situation, they should satisfy both of these constraints.

Configuration Design

For a formation flight experiment, we wish to control the total induced drag. We plan to study a control scheme in which the controller has the same spatial structure as the aerodynamic formation and to compare the performance of that control scheme with the performance of more standard types of control. This plan guides two fundamental design choices: the qualitative structure of each aerodynamic unit and the number of units.

The general concept of the problem calls for each aerodynamic unit to be an airfoil that moves relative to the others. The pertinent design choice is thus: how should the airfoil move? Of the six standard degrees of freedom (DOF), sway (lateral motion) is the one that most impacts the induced drag of the formation as a whole. Sway is most effectively controlled via roll. Roll is most effectively controlled with ailerons. Yaw, pitch, and heave (vertical motion) have second-order effects on the induced drag of the formation; we consider them decoupled from sway and roll so that in a practical system they would be controlled independently. Surge (streamwise motion) affects the induced drag locally, but moderate surge does not affect

Centralized Control	Distributed Control	Decentralized Control
Best Performance		Worst Performance
Depends on Formation	Independent of Formation	Independent of Formation
Heavy Communication	Light Communication	No Communication
Demanding Calculation	Light Calculation	Light Calculation
Difficult to Synthesize	Easy to Synthesize	Easy to Synthesize

Figure 3. A comparison of centralized, distributed, and decentralized control design and implementation. The comparisons are based on H -infinity optimization for the centralized control design, on H -infinity optimization with the coupling dynamics captured as noise for the decentralized control design, and on the results in [10], [19], and [11], for distributed control design.

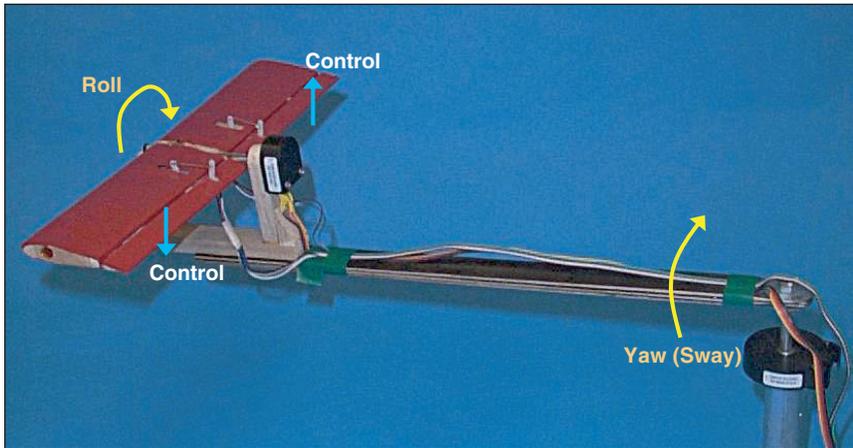


Figure 4. First-generation apparatus. The apparatus was a single wing with two degrees of freedom, roll and sway (lateral motion). Sway was implemented as yaw about an axis behind the wing. Both degrees of freedom were measured with low-friction optical encoders. Control was effected with servomotor-actuated ailerons.

the drag of the formation as a whole. Simulation of a half-vee formation of wings, which included a potential-flow model of the fluid dynamical coupling (similar to the approach described in [13]), with these two DOF confirmed that such a system does exhibit the necessary qualities: noise generates sway, and the system is controllable to the nominal position. Thus the physical system should have roll and sway DOF. The first apparatus built, pictured in Figure 4, was a single wing with these two DOF. For reasons of manufacturability, sway was implemented as yaw about an axis some distance behind the wing.

A minimum of two wings must be built to implement the spatial coupling integral to the control problem and the drag reduction. A larger formation can potentially present a more challenging control problem because disturbances have more opportunity to propagate in the spatial dimension. This propagation allows for the empirical observation of controller performance in a system having significant spatial structure. We chose to build a formation of four wings on the basis of the size of the available wind tunnel and readily available components for the wings such as servomotors and optical encoders. We chose a half-vee formation for the wings because it captures the essential dynamics of a full vee and thus makes the best use of the four wings.

Parameter Design

Some design decisions were directed by circumstantial constraints. The smallest available servomotors were selected to move the ailerons. To fit these servomotors inside the airfoil, a chord of about 9 cm was required. The test section of the wind tunnel is 1.2 m wide. This space must accommodate all of the wings placed tip to tip, while allowing for the motion of the outer wings

away from their nominal positions and avoiding the boundary layers at the edges of the test section. With four wings, a wingspan of about 24 cm is appropriate, resulting in an aspect ratio of about 2.7. For comparison purposes, to fit five wings in the wind tunnel each span could be no more than about 19 cm. The chord would remain 9 cm because of the servomotor, so the aspect ratio would be only 2.1. The assumptions used for Figure 1 are less applicable at the smaller aspect ratio. Additionally, the aerodynamic disturbances caused by the support structure in the center of the wing become more dominant for a smaller wingspan.

The airfoil cross section, NACA 0018, was chosen based on its ability to house the servomotors, the linearity of its lift curve for moderate angles of attack at low Reynolds numbers, and its significant lift under such conditions [22]. A rectangular platform was used because it generates stable and persistent trailing vortices [23], is actually quite inefficient and thus generates strong vortices [16], and may be manufactured with little variation from wing to wing. These qualities are helpful in keeping the aerodynamic coupling among wings consistent and substantial.

The length of the cantilevered arm between the yaw axis and the wing, for effecting the sway degree of freedom, was set at about 50 cm. In Figure 1, the transition from low, desirable power ratios to power ratios near one occurs at around $L/(1+L) \approx 0.2$, where the wingtip separation L is about 0.25 span, or 6 cm. An arm of 50 cm length yaws by 0.12 rad (about 7°) at a sway position of 6 cm. A 6-cm sway on a 25-cm arm would be about 0.24 rad. Such a large yaw angle combined with a restoring roll angle can combine to diminish the effective angle of attack, eliminating the lateral lift necessary to move the wing back toward the nominal position. The relationship is

$$\sin \alpha_E = \sin \alpha_0 \cos \phi + \cos \alpha_0 \sin \theta \sin \phi$$

where α_E is the effective angle of attack, α_0 is the nominal angle of attack, ϕ is the yaw angle about the nadir, and θ is the roll angle about the forward-pointing roll axis. This relationship may be approximated by $\alpha_E - \alpha_0 = \phi \sin \theta$. With $\phi = \alpha_0 = 7^\circ$ and $\theta = 30^\circ$, α_E is about half of α_0 . An arm half as long would double ϕ so that α_E would be near zero and the wing would generate no lift in that position. An arm too long could result in heaving oscillations. While heave has a second-order effect on the dynamics of interest, such motion would constitute a continual disturbance

interfering with intentionally applied disturbances. The second generation apparatus, pictured in Figure 5, comprised four wings designed under these considerations.

Myriad decisions were made on more detailed levels of the design, outside the scope of this article. In some cases, these decisions were based on practical issues arising with the first- and second-generation apparatuses. Many decisions were guided by considerations of manufacturability, repeatability, robustness, and ease of maintenance. For example, the yaw arm was mounted using L-brackets because they make a stiff joint and are simple to machine; the airfoil was rapid prototyped to provide uniformity among wings with a structure too complex to fabricate by hand; the shafts were mounted with bearings, spacers, and shaft collars to keep nonlinear friction low, resulting both in a harder, more realistic control problem and in a system less apt to change over time; and the airfoil was made to allow nondestructive access to its interior for service. In general, roll inertia and yaw inertia were kept low: put simply, inertia may be added to the system with ease, but it is difficult to remove inertia.

Real-Time Control Environment

The main requirements of the real-time control environment are:

- easy to use and operate
- the interface can handle a modest number of inputs and outputs
- can implement large order controllers or computationally demanding controllers

We settled on a real-time control system from dSPACE, mainly because of its ability to handle all of the encoder inputs and pulse-width modulation outputs and because we had one available from a previous experiment. The centralized nature of the real-time control environment has obvious benefits for comparing various control strategies: one can readily implement centralized controllers, in addition to distributed and decentralized controllers. All control design, analysis, and implementation were performed with MATLAB and Simulink.

The Three Generations

The design of the experiment was implemented in three stages. The first was simply to build, test, and control one wing, as depicted in Figure 4. We were able to stabilize the wing with a simple observer-based controller. The model used was a so-called white box model in which all parameters were measured directly or estimated, and no experimental data was needed.

Once this was successfully achieved, we moved on to what we thought would be the final design of the formation, or what we now refer to as the second generation, depicted in Figure 5. Unfortunately, we were not able to obtain a good model of the coupling between aircraft. We



Figure 5. Second-generation apparatus. The apparatus was essentially four copies of the first generation apparatus, mounted in a half-vee formation. The wingspan was reduced to fit all four wings into the wind tunnel.

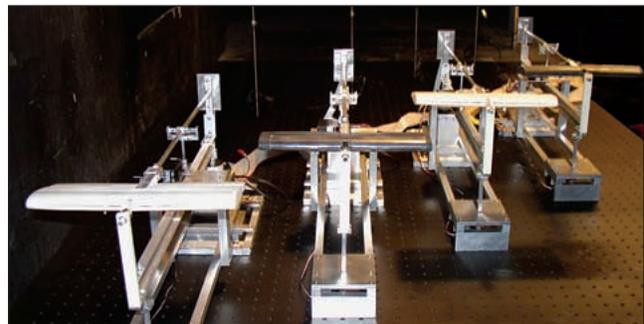


Figure 6. Third-generation apparatus. The apparatus allowed for more repeatable experiments than the second generation, among other improvements. This apparatus is viewed from upstream inside the wind tunnel.

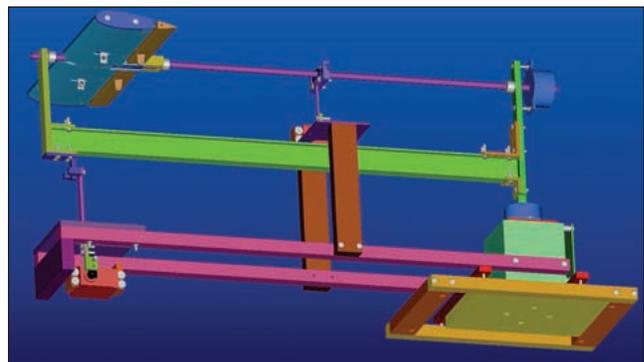


Figure 7. Pro/ENGINEER drawing of the third-generation design. The CAD model was used for rapid prototyping the aerodynamic parts and for machining the aluminum parts.

discovered that there was substantial Coulomb friction, or stiction, in the various bearings that introduced large uncertainty in the coupling dynamics.

From the lessons learned in the first two generations, we proceeded to the third generation, shown in Figures 6 and 7. All the observable deficiencies of the first two generations were addressed in the third incarnation of the experiment. Noticeable changes included rapid prototype wings, 20,000 count encoders for the yaw axis, new low friction bearings, new methods for routing the cables with minimal interference to the motion, and servo-controlled roll and yaw locks to facilitate system identification and system initialization.

Each module was made to bolt to an optical table to allow for different formation configurations. The wings

can be mounted with different streamwise spacing, with different lateral spacing, or possibly in a more general formation. A formation other than a variation on a vee formation is less applicable to the practical and control problems we are studying but may be useful for investigating alternative control problems abstracting different practical scenarios.

Wing Dynamics

Physical considerations motivate a general form for the dynamics of each wing (except for the lead wing) given by

$$\begin{aligned} \dot{X} &= F(X, U, X_u), \\ (\theta, \phi) &= CX, \end{aligned} \quad (1)$$

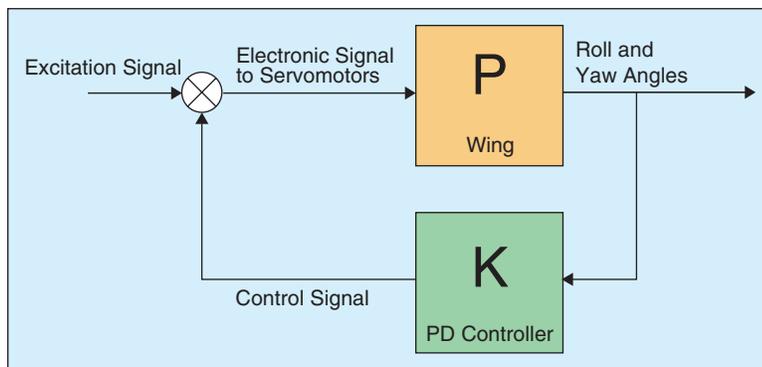


Figure 8. Configuration for local, linear system identification. A wide-band disturbance was injected into the control signal. The control signal and sensor output signals were analyzed using MATLAB's System Identification Toolbox.

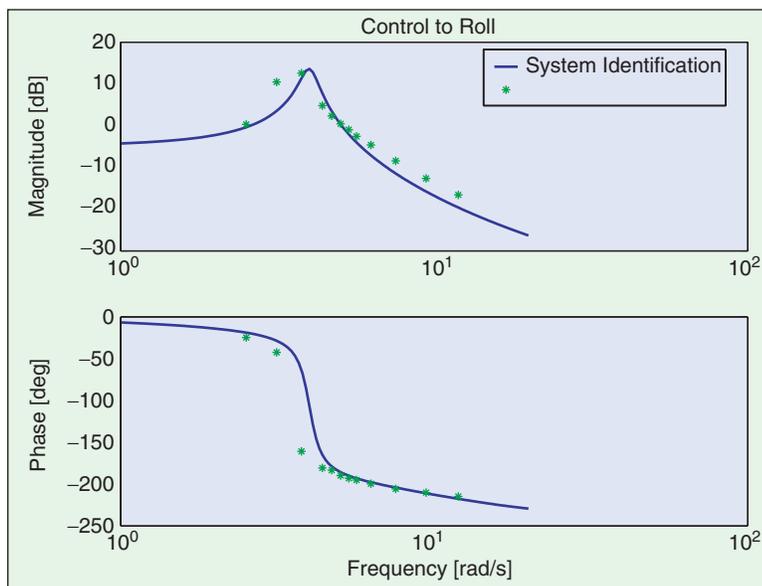


Figure 9. Frequency response of roll angle to control input. The identified transfer function from control input to roll angle conforms to the observed response of the system to sinusoidal excitations.

where X is the local state, U is the local control input, X_u is the upstream state, θ is the roll angle, ϕ is the yaw angle (sway), F is a static nonlinearity, and C is a matrix. One immediate feature of this model is the assumption that the downstream aircraft does not affect the local dynamics; we shall see experimental evidence for this assumption shortly. This assumption is also supported by classical potential flow analysis [18], which states that for airfoil separations beyond several wing-spans, the coupling from downstream to upstream becomes negligible. The dynamics of the lead wing are slightly modified, since there is no upstream wing. For the purposes of this discussion, we concentrate on the wings embedded in the formation, or wings 2, 3, and 4. For all the experiments, the lateral wing-tip to wing-tip separation of the wings was zero, and the streamwise separation was 20 cm.

Local, Linear System Identification

By locking the yaw and roll axes of wings 1, 2, and 4, where wing 1 is the lead wing, and wing 4 is the tail wing, we can identify the local, linear, dynamics of wing 3. A complication, of course, is that the free wing is unstable and must be actively controlled to remain near a trim position. Control was achieved by hand-tuning a proportional-derivative controller, as depicted in Figure 8. In particular, the system was trimmed and stabilized about an operating point $(\theta, \phi, U) = (\bar{\theta}, 0, \bar{U})$. A nonzero roll angle and nonzero control effort were required to trim the yaw position to zero, corresponding to the solution of the dynamic equations given by

$$0 = F(\bar{X}, \bar{U}, 0)$$

$$(\bar{\theta}, 0) = C\bar{X}.$$

A broadband excitation signal injected into the control signal was used for system identification. The System Identification Toolbox in MATLAB was used to obtain a fourth-order, state-space model from control input to yaw and roll angles. The results are depicted in Figures 9 and 10. Included in these plots is data obtained by exciting the closed-loop system with sinusoidal signals. The fit is quite good; the identified model was in fact used to design an observer-based controller and an H-infinity controller, both of which stabilized the system without iteration.

Nonlinear Coupling

In the second-generation design, we were unable to identify the wing coupling. The difficulty was due to large stiction, or Coulomb friction, that resulted in data with virtually no correlation. One of the attributes of the third-generation design is readily identifiable nonlinear coupling. The nonlinear coupling is a crucial aspect of the experiment: the ambient noise in the windtunnel is large enough to excite the system to a point where nonlinear effects dominate the fluid-dynamic coupling. A full characterization of the nonlinear coupling is beyond the scope of this article, but several prominent features can be presented.

The first experiment consisted of locking the roll and yaw axes of wings 4 and 1 and the yaw axes of wings 2 and 3. In other words, the only free DOF were the roll angles of wings 2 and 3. Wing 2, the upstream wing, was then cycled by injecting a slowly varying excitation signal into its local PD controller. Wing 3, the downstream wing, was commanded to a zero roll angle through its local PD controller, which was augmented with a small integral term on the roll angle. The required control effort for wing 3, as a function of the wing 2 roll angle, is depicted in Figure 11. For comparison purposes, a control effort of one saturates the control ailerons at an angular displacement of approximately $\pm 60^\circ$. Note that this identification corresponds to a very limited characterization of the nonlinearity F in (1).

The process was repeated with the roles of wings 2 and 3 interchanged. In other words, we wanted to observe the coupling from the downstream wing to the upstream wing. The results are found in Figure 12, which shows that the coupling from downstream wing to upstream wing is negligible when compared to the coupling from upstream to downstream.

These results can be explained by geometric arguments based on the location of the shed

vortices and the locations of the wings. When the wings are maximally coupled, any deviation from the trim condition will result in a decrease in coupling. Locally, this can be approximated by a parabola, which can then be used for nonlinear control design and analysis.

PD Control of the Formation

The hand-tuned PD controllers can be modified to control the formation, with the objective of minimizing the lateral separation of the wings. In particular, the local control signals were generated as

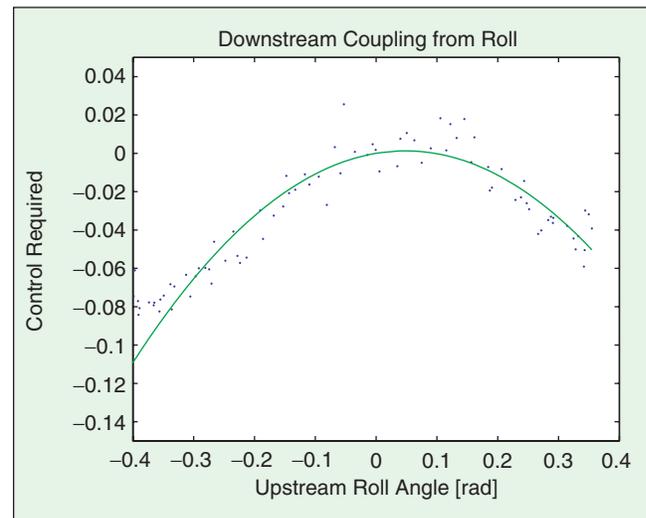


Figure 11. Roll coupling from upstream wing to downstream wing. A parabola is used to fit the data.

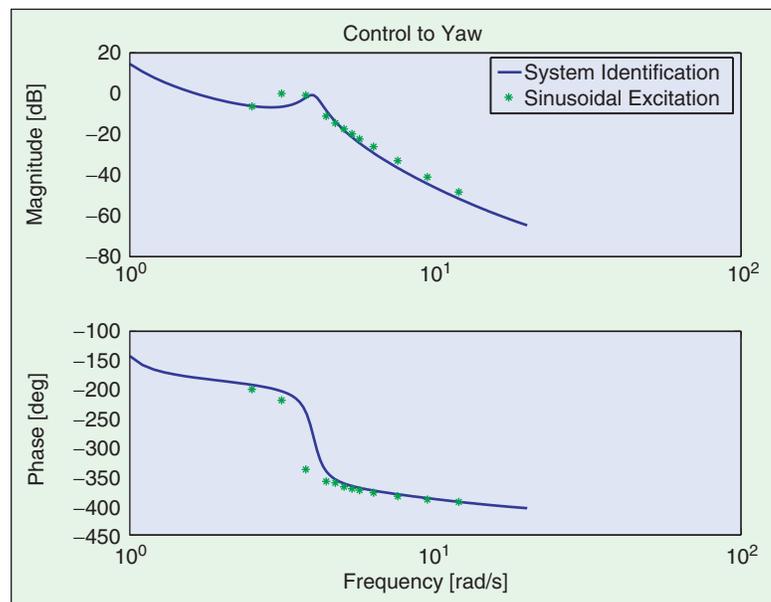


Figure 10. Frequency response of yaw angle to control input. The identified transfer function from control input to yaw angle conforms to the observed response of the system to sinusoidal excitations.

$$u(s) = k_1\theta(s) + k_2\dot{\theta}(s) + k_3(\phi(s) - \phi(s-1)) + k_4\dot{\phi}(s),$$

$$1 \leq s \leq 4, \quad \phi(0) := 0.$$

It should be noted that including a $\dot{\phi}(s-1)$ term in the feedback law promptly destabilized the formation. Two data sets are presented in Figures 13 and 14. The only external forcing of the system was the ambient noise in the wind tunnel. As can be seen, there is clear coupling in the motion of the wings. This coupling is to be expected, given the coupling in the control law. Also of interest is that the overall motion of the wings increases as we move further down the formation; this is not surprising given the simple control law being used and is an indication of string instability. What may be surprising, however, is how different the two data sets are. In the first data set, the yaw angle offsets decrease as we move down the formation and as time increases. In the second data set, the yaw angle offsets increase. This trend is a clear indication that the non-linear coupling, which is essentially quadratic, is playing a predominant role in the dynamics of the formation.

Concluding Remarks

The design of the experiment has been a long, but rewarding, experience. In addition to the two authors, the project has involved five undergraduate and master of engineering students in the design and construction aspects.

Prior to the third-generation version of the experiment, the results were mixed. The first generation, which consisted of only one wing, was successful when considered in isolation. The experiment displayed the right types of dynamics, consistent with what was predicted by simple scaling and physical arguments, and was controllable with a relatively simple controller. The experiment displayed

some limit-cycle behavior in the closed loop, due to the Coulomb friction nonlinearities.

With some minor modifications and improvements, the second generation was not as successful. Coulomb friction and other undesirable phenomena had a large impact on the coupling and identification of the wing dynamics. It should be stressed that these nonidealities led to some interesting control problems, including control of highly uncertain, highly nonlinear, interconnected systems. Nonetheless, the second generation did not meet either of our two basic requirements: 1) it was not a useful testbed for exploring control of interconnected systems as the theory is not ready

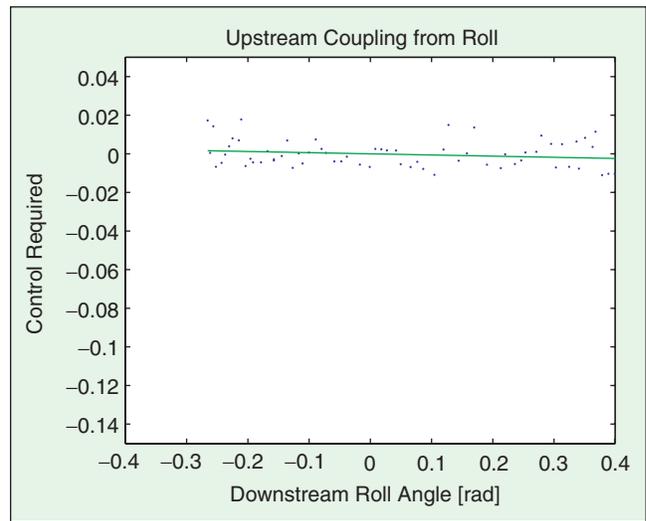


Figure 12. Roll coupling from downstream wing to upstream wing. The upstream coupling shown here is negligible compared to the downstream coupling in Figure 11.

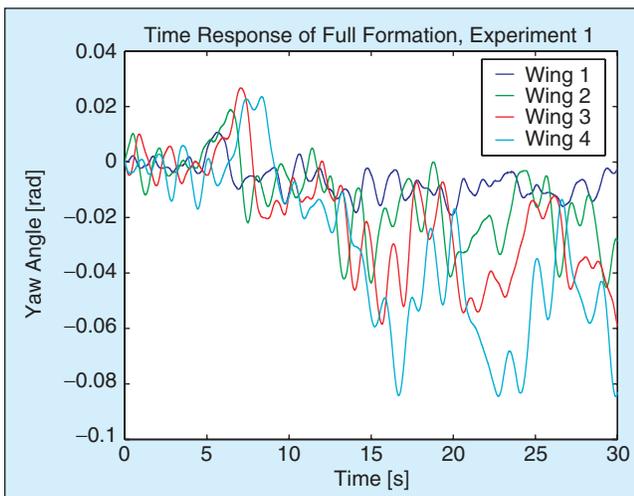


Figure 13. PD control of formation, first data set. The coupling among the wings is apparent. In this trial, yaw angle offsets decreased with time and from the first wing to the fourth wing.

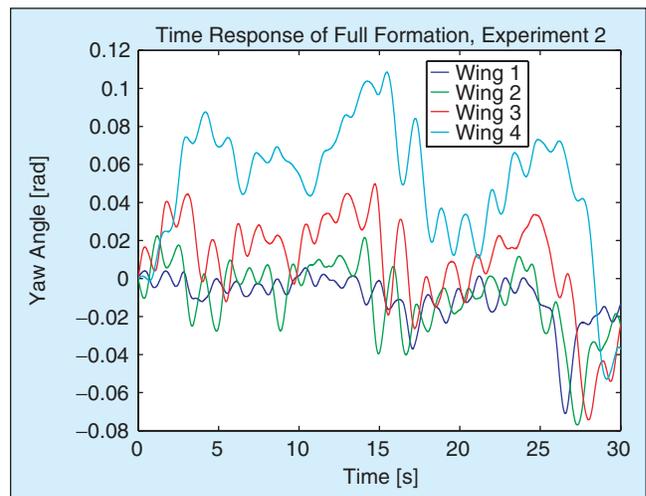


Figure 14. PD control of formation, second data set. The coupling among the wings is apparent. In this trial, yaw angle offsets increased with time and from the first wing to the fourth wing.

for both highly uncertain dynamics *and* strong nonlinearities, and 2) it was not a suitable abstraction of airplanes flying in formation, where Coulomb friction is not an issue.

The third generation has met the original objectives. In particular, the dynamics are identifiable, the nonlinearities are mild (they are continuous, for example), and the dynamics of the experiment are a reasonable abstraction of airplanes flying in formation.

Acknowledgments

The contributions of the following people are gratefully acknowledged: Ryusei Matsumoto and Hoon Kong designed, built, and analyzed the first generation version of the apparatus and helped start the multiwing testbed. Brett Lee and Steven Chang helped design and build the second-generation testbed. In the third-generation version of the apparatus Jin Woo Lee made great improvements to the electronics. Jeremy Miller helped to make similar major improvements to the design and executed much of the construction of the new system. We also wish to thank Hod Lipson and Evan Malone for their generous assistance in the fabrication of the new airfoils. Finally, we would like to thank the anonymous reviewers for their thoughtful and careful comments.

References

- [1] H. Raza and P. Ioannou, "Vehicle following control design for automated highway systems," *IEEE Contr. Syst. Mag.*, vol. 16, no. 6, pp. 43–60, 1996
- [2] G.B. Shaw, D.W. Miller, and D.E. Hastings, "The generalized information network analysis methodology for distributed satellite systems." Ph.D. dissertation, Massachusetts Institute of Technology, 1998.
- [3] G.E. Stewart, "Two dimensional loop shaping controller design for paper machine cross-directional processes," Ph.D. dissertation, Univ. of British Columbia, 2000.
- [4] M. Napoli and B. Bamieh, "Modeling and observed design for an array of electrostatically actuated microcantilevers," in *Proc. IEEE Conf. Decision and Control*, Orlando, FL, 2001, pp. 4274–4279.
- [5] M.E. Taylor. *Differential Equation I: Basic Theory*. New York: Springer-Verlag, 1996.
- [6] A. Pant, P. Seiler, and J.K. Hedrick, "Mesh stability of look-ahead interconnected systems," *IEEE Trans. Automat. Contr.*, vol. 47, no. 2, pp. 403–407, 2002.
- [7] B. Bamieh, F. Paganini, and M. Dahleh, "Distributed control of spatially invariant systems," *IEEE Trans. Automat. Contr.*, vol. 47, no. 7, pp. 1091–1118, 2002.
- [8] D. Swaroop and J.K. Hedrick, "String stability for a class of nonlinear systems," *IEEE Trans. Automat. Contr.*, vol. 41, no. 3, pp. 349–357, 1996.
- [9] D.D. Siljak. *Decentralized Control of Complex Systems*. San Diego, CA: Academic, 1991.
- [10] R. D'Andrea and G.E. Dullerud, "Distributed control of spatially interconnected systems," *IEEE Trans. Automat. Contr.*, to be published.
- [11] R. D'Andrea, C. Langbort, and R. Chandra, "A state space approach to control of interconnected systems," in *Mathematical Systems Theory in Biology, Communication, Computation and Finance*. J. Rosenthal and D.S. Gilliam, Eds. New York: Springer, IMA Book Series, 2003, pp. 157–182.
- [12] R.S. Shevell. *Fundamentals of Flight*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [13] J.D. Wolfe, D.F. Chichka, and J.L. Speyer, "Decentralized controller for unmanned aerial vehicle formation flight," in *Proc. AIAA Guidance, Navigation, and Control Conf.*, San Diego, CA, 1996, AIAA paper 96-3833.
- [14] D.F. Chichka and J.L. Speyer, "Solar-powered, formation-enhanced aerial vehicle system for sustained endurance," in *Proc. American Control Conf.*, Philadelphia, PA, 1998, pp. 684–688.
- [15] P.B.S. Lissaman and C.A. Shollenberger, "Formation flight of birds," *Science*, vol. 68, pp. 1003–1005, 1970.
- [16] J.D. Anderson. *Fundamentals of Aerodynamics*. New York: McGraw Hill, 1991.
- [17] D. Hummel, "The use of aircraft wakes to achieve power reductions in formation flight," in *AGARD Conf. Proc.*, Dresden, Germany, 1996, pp. 1–13.
- [18] L. Prandtl and O.G. Tietjens, *Applied Hydro and Aeromechanics*. Mineola, NY: Dover, 1934.
- [19] R. D'Andrea, "A linear matrix inequality approach to decentralized control of distributed parameter systems," in *Proc. American Control Conf.*, Philadelphia, PA, 1998, pp. 1350–1354.
- [20] K.S. Wittmer, W.J. Devenport, M.C. Rife, and S.A.L. Glegg, "Perpendicular blade-vortex interactions," *AIAA J.*, vol. 33, no. 9, pp. 1667–1674, 1995.
- [21] D.S. Bernstein. "What makes some control problems hard?" *IEEE Contr. Syst. Mag.*, vol. 22, no. 4, pp. 8–19, 2002.
- [22] S.J. Miley, *A Catalog of Low Reynolds Number Airfoil Data for Wind Turbine Applications*. College Station, TX: Department of Aerospace Engineering, 1982.
- [23] C.H.K. Williamson, personal communication, August 1999.

Jeffrey M. Fowler is a Ph.D. student of mechanical engineering at Cornell. He received a bachelor of mechanical engineering degree from the Georgia Institute of Technology in 1999 and a master of science degree in mechanical engineering from Cornell. He is a fellow in Cornell's interdisciplinary graduate training program in nonlinear systems. He also received an NSF Graduate Research Fellowship.

Raffaello D'Andrea received the B.A.Sc. degree in engineering science from the University of Toronto in 1991 and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology in 1992 and 1997, respectively. Since then, he has been with the Department of Mechanical and Aerospace Engineering at Cornell University, where he is an associate professor. He received the University of Toronto W.S. Wilson Medal in Engineering Science, an IEEE Conference on Decision and Control best student paper award, an American Control Council O.Hugo Schuck Best Paper award, an NSF CAREER Award, and a DOD sponsored Presidential Early Career Award for Scientists and Engineers. He was the system architect and faculty advisor of the RoboCup world champion Cornell Autonomous Robot Soccer team in 1999, 2000, 2002, and 2003 and third-place winner in 2001. His recent collaboration with Canadian artist Max Dean, "The Table," an interactive installation, appeared in the Biennale di Venezia in 2001.

