

Electric versus Hydraulic Flight Controls: Assessing Power Consumption and Waste Heat Using Stochastic System Methods

William Schley
PC Krause & Associates

ABSTRACT

Of all aircraft power and thermal loads, flight controls can be the most challenging to quantify because they are highly variable. Unlike constant or impulsive loads, actuator power demands more closely resemble random processes. Some inherent nonlinearities complicate this even further. Actuation power consumption and waste heat generation are both sensitive to input history. But control activity varies considerably with mission segment, turbulence and vehicle state. Flight control is a major power consumer at times, so quantifying power demand and waste heat is important for sizing power and thermal management system components. However, many designers sidestep the stochastic aspects of the problem initially, leading to overly conservative system sizing. The overdesign becomes apparent only after detailed flight simulations become available. These considerations are particularly relevant in trade studies comparing electric versus hydraulic actuation. These two actuation types use power in fundamentally different ways. This paper provides methods to quantify power consumption and waste heat, by applying stochastic system methods. Both electric (electromechanical and electrohydrostatic) and conventional hydraulic actuation are discussed. Formulas are derived to quantify average and peak power demand. A complete set of waste heat mechanisms is also discussed, and methods are provided to quantify each one. For electric actuation, a method is provided to estimate regenerative power. Approximate methods are also addressed, to facilitate quick engineering estimates. In addition, the physical locations of waste heat generation are identified because these can impact thermal management system architecture.

CITATION: Schley, W., "Electric versus Hydraulic Flight Controls: Assessing Power Consumption and Waste Heat Using Stochastic System Methods," *SAE Int. J. Aerosp.* 10(1):2017, doi:10.4271/2017-01-2036.

INTRODUCTION

The loads on an air vehicle power system occur with a variety of duty cycles:

- Continuous-steady (e.g. avionics)
- Occasional-steady (e.g. landing gear retract)
- Impulsive (e.g. radar, electronic warfare)
- Continuous-variable (e.g. flight controls)

This mix adds to the design sizing challenges for power and thermal management systems. Many vehicle system components are large, specialized, and expensive, have long procurement lead times, and are impractical to resize once in hand. This puts pressure on designers to size components early in the program. In addition, the power available from the engines increases with engine speed but unfortunately, peak flight control power demand usually occurs at low speed during turbulence. Also, cost-cutting pressure favors the use of off-the-shelf components. This practice requires designers to accept some overdesign. To minimize the inevitable weight and size impacts, accurate estimates of peak and average power demand are essential. Variable loads and rates make flight control power demand particularly challenging to quantify. Flight control response to turbulence creates an even bigger challenge.

Current aircraft design practice usually follows three steps to quantify flight control power demand: In *Step One*, actuator sizes are estimated from aerodynamic estimates of control surface hinge moments, and from control surface rates and bandwidths needed to satisfy stability, maneuverability, and handling requirements.

Step Two entails making preliminary estimates of peak and average flight control power demand, and average waste heat. These are needed for sizing the related long-lead components such as central hydraulic pumps and thermal management components. This paper addresses conservatism often seen during this step. In hydraulic flight controls for example, designers may reason that the peak total hydraulic flow demand cannot exceed the sum of the individual actuator peak demands, so that sum serves as an estimate of the peak total demand. This is overly conservative because statistically, the actuators do not demand their individual peaks simultaneously. Then to estimate the average total flow demand, some designers arbitrarily use the Root-Sum-Square (RSS) of the individual actuator peak demands. This too, will be shown to be overly conservative, and skirts the true stochastic aspects of the problem.

Step Three involves doing six-degree-of-freedom (6DOF) simulation studies to more accurately quantify control surface load/rate envelopes, and enables a more accurate estimate of flight control

power demand. The approach is rigorous and thorough. It produces actuator input histories over the entire flight envelope for aggressive maneuvers, turbulence, and gusts. The input histories are then used to accurately quantify flight control peak and average power demand and average waste heat generated. These studies usually provide the first opportunity for designers to confront the true stochastic aspects of power consumption and waste heat, and it is during these studies that any overdesign first becomes apparent. Unfortunately, since the simulation requires detailed engineering models of aerodynamics, mass properties, propulsion, and flight control laws, the simulation configuration might undergo months of development before these results converge.

In the interim between Steps Two and Three, an issue may arise as the procurement of long-lead components begins. Designers must provide their still-conservative requirements to the suppliers, and hope that subsequent revisions can be produced soon enough to avoid significant cost and schedule impacts. The process is cumbersome and can be problematic.

The following discussion begins with a traditional approach for making detailed assessments of peak and average flow and power demands for hydraulic flight controls. We then propose a new method for making preliminary estimates that is more accurate than the methods currently used by many designers. We also locate and quantify the waste heat that is generated. We then direct the discussion to *electric* flight controls, and quantify electrical current and power demands, as well as the waste heat mechanisms. A second new method is proposed for preliminary power estimates for electric flight controls. By gathering together descriptions of these methods, the intent is to facilitate preliminary design, and enable more accurate trade studies of electric versus hydraulic flight controls.

QUANTIFYING POWER CONSUMPTION FOR HYDRAULIC ACTUATION

A conventional hydraulic servoactuator can be idealized as demanding variable fluid flow at a constant working pressure. In the simplest terms, the *instantaneous power draw is equal to the instantaneous flow demand multiplied by the working pressure*. Fluid is delivered by the central hydraulic system to the servoactuator as essentially unidirectional, high-pressure flow. Once used by the servoactuator, the fluid is returned with negligible pressure to the central hydraulic system where it is cooled, filtered, and pressurized back to the nominal supply pressure.

In actuality, the supplied fluid arrives at the servoactuator at a somewhat reduced pressure, due to viscous flow losses in the hydraulic system tubing. The fluid returned by the servoactuator has a pressure typically around 150 psi. The actual working pressure available to the servoactuator is the difference between the pressures at its supply and return ports. For example, a typical system having a 4,200 psi nominal supply pressure will provide the actuator with a working pressure averaging around 3,800 psi.

It is important to note that a *hydraulic servoactuator consumes power very nearly in proportion to flow, regardless of whether it performs any useful work*. This can be explained using [Figure 1](#). We define piston velocity, $\dot{x}_p$, as positive when extending. Velocity is modulated by Main Control Valve (MCV) position, x_v . We define signed “load flow”, Q_L , as positive when the actuator is extending. Fluid supplied to the servoactuator undergoes three pressure drops: The first one is across a metering orifice in the MCV, the second is across the piston head, and the third is across another metering orifice in the MCV. The pressure drop across the piston head reacts the load as shown. It is this pressure drop that performs useful work. Whereas, the two orifice pressure drops dissipate energy, but perform no useful work. Since flow exits the servoactuator with negligible pressure, the power demanded by the servoactuator is nearly completely spent by these three pressure drops. In the special case where the actuator moves under no load, power is consumed but no useful work is performed. In that case, all of the power is allocated to the MCV orifice pressure drops. Consequently, *if a hydraulic servoactuator is moving, it is consuming power*. (Note that a small amount of power is consumed even when stationary, due to internal leakage, as discussed in the coming section.)

In order to quantify hydraulic power demand, we first quantify load flow. We seek an approach that can accommodate arbitrary inputs because the results are known to be input sensitive. It is convenient to make use of a dynamic model such as the one shown in [Figure 2](#). This model follows the derivations and notation used in [Ref. 1]. Several simplifying assumptions are made in order to linearize the dynamics (small control surface deflections, matched rectangular flow metering orifices, lumped masses and mechanical compliances, and so forth). This model accounts for fluid compressibility and internal leakage. Here, we define *trim flight* as the condition in which moments on the vehicle are balanced. This condition generally requires non-zero control surface hinge moments. Aerodynamic hinge moments are modelled here as spring-like, and with a bias needed for trim. The spring-like portion is accounted for in the load stiffness, K , and the bias is applied as the trim hinge moment input, modelled as a unique constant for each flight segment or flight condition. A second input to the system is control surface deflection command, and the output is control surface deflection. Refer to [1] for the complete derivation based on first principles and guidance on computing the parameters. On a time-average basis, this model can quantify load flow with excellent accuracy for our purposes, because the linearizing assumptions are satisfied nearly all of the time.




Figure 1. Hydraulic fluid undergoes three pressure drops.




Figure 2. Basic linear dynamic model of a hydraulic servoactuator, showing how load flow is used to compute hydraulic flow and power demands.

This particular model is shown only to illustrate the approach. In general, the model may incorporate rigorous detail and nonlinearities. The approach only requires that the load flow signal is accessible. Load flow is defined in [1] as the average of the incoming and outgoing flows. In Figure 1, it can be seen that the flow demanded from the hydraulic supply is simply the absolute value of Q_L . Then to compute hydraulic power demand, we multiply the flow demand, Q_{DEM} by the working pressure. Note that both flow demand and power demand have a rectifying effect produced by the MCV. Experience shows that the assumption of constant working pressure yields good accuracy. Return pressure is neglected in this model, but can be taken into account by instead setting P_s equal to the difference between supply and return pressures. To summarize, this model-based approach accommodates any desired fidelity and arbitrary inputs, but requires some understanding of the first principles in order to compute the parameter values.

As an alternate approach, we may approximate load flow by multiplying piston velocity by piston area. This approximation of Q_L is quite accurate as a consequence of the high fluid bulk modulus and low internal leakage typical of modern hydraulic servoactuators. We then take the absolute value of this to get an approximation of flow demand, and multiply that by working pressure to obtain approximate power demand, as shown with broken lines in Figure 2.

Importantly, it is evident from this approach that *hydraulic power demand is very nearly proportional to output speed*. This approach accommodates arbitrary inputs. But in addition, it offers the advantage that it can be applied to any dynamic model in which the velocity signal is accessible. It does not require detailed understanding of the actuator design. *The approach only requires knowledge of the dominant dynamic characteristics of the actuator.* For example, in the approximate linear dynamic actuator model of

Figure 3, piston velocity can be readily accessed by expanding out the second order lag. Simple dynamic models such as this are frequently used in 6DOF simulations for their speed of execution.

Now, consider the driving inputs for the model. Power consumption and waste heat generation are input-sensitive. It therefore is important to identify realistic actuator inputs (i.e. control surface position command and trim hinge moment). The trim hinge moment is fairly constant during each flight segment. But the surface command, idealized for our purposes, is random in nature. It contains the combined responses to pilot inputs and vehicle disturbances. The vehicle and servoactuators respond as linear systems because the conditions for linearity are met nearly all of the time. This is particularly true for highly-augmented aircraft because designers go to lengths to prevent servoactuator stall and rate saturation. It is tempting to assume that hydraulic servoactuators consume the maximum input power when they are delivering the maximum output power. However as previously shown, they consume power in close proportion to speed. Experience shows that *for most manned aircraft, aggressive maneuvers typically require surface rates no greater than those seen during turbulent approach*. In addition, aggressive maneuvers allocate most of the control activity to only a subset of the servoactuators at a time. Whereas during turbulence, all the servoactuators show consistently high activity and operate very inefficiently. Moreover, maneuvers in general occur over a very small fraction of the flight time. Therefore, we can focus on just surface commands in response to turbulence.

Realistic surface command histories can be obtained from 6DOF simulation studies including turbulence, or from representative flight data captured for similar aircraft. It might also be convenient to characterize the random nature of the position command by using its Power Spectral Density (PSD). There is an effort currently underway, investigating the notion that surface command PSD for each flight segment might be accurately approximated using simple functions,

and further, that standard sets of PSD's might represent typical and worst-case control surface activity for entire classes of aircraft. Surface command PSD's might also yield more accurate estimates of power consumption when there is little or no turbulence. The usefulness of the input PSD can be seen in the following discussion.




Figure 3. Approximate linear dynamic model of a hydraulic actuator.

Signals in engineering systems often exhibit Gaussian behavior, which have normally distributed time samples. Gaussian models have long been accepted as adequate representations of flight turbulence. For example, the popular Dryden and von Karman models described in [2] and [3] provide Gaussian wind disturbances, scaled for various levels of intensity. A remarkable result from stochastic system theory is that *the output of a linear system driven by a Gaussian input, is also Gaussian* [4]. If we choose a linear model such as the one in Figure 2, it is striking that the only nonlinearity necessary for our purposes is the absolute value calculation. This system is linear up to point of load flow, Q_L . If Gaussian flight disturbances bring about Gaussian surface commands (treating the aircraft as a linear dynamic system), we expect Q_L to be Gaussian as well. Given the PSD of the position command input, $\Phi(\omega)$, the variance of the load flow can be computed [Ref. 5] using a numerical integration of,

$$\sigma_{Q_L}^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} |G(j\omega)|^2 \Phi(\omega) d\omega \quad (1)$$

where $G(s)|_{s=j\omega}$ is the transfer function from position command to load flow. *This integral captures how the unique input power in each frequency band is filtered by the system, and how the result contributes to the output variance.* Since the system is linear, $G(s)$ can be found easily, and only its magnitude is needed. The computation becomes straightforward if the form of $\Phi(\omega)$ is simple. Next, the *mean* of the load flow can be found using steady-state analysis of the mean inputs. With the mean and variance of the load flow in hand, all that remains is to find the mean and variance of the absolute value of the load flow.

For the absolute value function, $y = |x|$ in general, if the variable x is Gaussian with a mean, μ_x and variance, σ_x^2 , Appendix A derives mean and variance of y :

$$\mu_y = \sigma_x \sqrt{2/\pi} e^{-\mu_x^2/(2\sigma_x^2)} + \mu_x \operatorname{erf}\left(\frac{\mu_x}{\sigma_x \sqrt{2}}\right) \quad (2)$$

$$\begin{aligned} \sigma_y^2 = & \mu_x^2 + \mu_y^2 - 2\mu_x\mu_y \operatorname{erf}\left(\frac{\mu_x}{\sigma_x \sqrt{2}}\right) + \dots \\ & + \sigma_x \left[\sigma_x - 2\sqrt{2/\pi} \mu_y e^{-\frac{\mu_x^2}{2\sigma_x^2}} \right] \end{aligned} \quad (3)$$

Some examples are illustrated in Appendix A, showing probability densities for $|x|$, given various mean values of x . The broken lines show how the bell-shaped probability density curve for x folds back and adds to itself. The solid curve, representing the probability density of $|x|$, is the sum of the two curves with broken lines. (Note that the effect of the absolute value on the shape of the probability density curve is only significant for signals centered near zero. For progressively increasing bias, the distortion diminishes and we find that $\mu_y \rightarrow |\mu_x|$ and $\sigma_y \rightarrow \sigma_x$. In fact for positive signals having $\mu_x \geq 3\sigma_x$, we might choose to ignore the absolute value altogether.) The mean load flow is negligible for typical hydraulic servactuators, even for inputs with large non-zero means (though a small amount of flow compensates for internal leakage). As a result, the formulas simplify considerably:

$$\mu_y = \sigma_x \sqrt{2/\pi} \quad (4)$$

$$\sigma_y^2 = \left(1 - \frac{2}{\pi}\right) \sigma_x^2 \quad (5)$$

By applying these results to the load flow, we can quantify the mean and variance of its absolute value. These values enable us to quantify the actuator's mean and peak hydraulic flow demand for any flight segment for the given air vehicle. For example, we may choose to define peak flow as the mean value plus three sigma. Power demands can then be computed by multiplying the mean and peak flow demands by supply pressure.

Notice that although the load flow is Gaussian, the flow demand is not, as shown in Figure 4. If we rectify a zero-mean Gaussian process, the resulting process has a one-sided probability density function twice as tall, and with a non-zero mean. We shall refer to this as "rectified-Gaussian", for which the mean and variance are given by Eqs. (4) and (5). We now have a way of expressing statistics of the flow demand for a single servomotor responding to flight turbulence.




Figure 4. Comparison of probability densities for load flow (Gaussian) and flow demand ("rectified-Gaussian").

Next, we are interested in the statistics of the combined rectified-Gaussian flow demands from multiple actuators. To address this, we know that *the mean of a sum of random processes is the sum of the means* (assuming each actuator moves independently), and that *the variance of a sum is the sum of the variances*. It is also interesting that as the number of actuators increases, the probability density of their combined flow demand reverts back toward Gaussian; a consequence of the Central Limit Theorem. To illustrate with an example, the Figure 5 shows the density function resulting when we sum N rectified-Gaussian processes, each with $\sigma = 1$.




Figure 5. The probability density of the sum of N rectified-Gaussian processes reverts back to Gaussian as N increases.

The above method can be used to quantify the variance of the output of the linear portion of the system, if the PSD of each input is known. The effect of the absolute value is then addressed, and the individual servoactuator power demands are combined. However, input PSD's might not be available. We now propose a new engineering method for estimating hydraulic flow and power demands, without the PSD's. Although this method is approximate, we will show that it is more accurate than other methods commonly used for preliminary power estimates. We shall assume that the servoactuators have already been sized, and as a result, their individual maximum no-load flow demands are available. We are interested in quantifying total flow and power demand for the complete flight control system. Having already recognized servoactuator load flow responses as zero-mean and Gaussian, we need a way to scale the individual actuator flow statistics. In contrast to the previously-described PDS method, *this method simply equates 3σ to the servoactuator's no-load flow demand, Q_{NL}* . This is an approximation that treats attainment of Q_{NL} as a three-sigma event while flying through significant turbulence. The rationale for this is that a properly sized actuator should reach such a high flow only very rarely. Reaching Q_{NL} more frequently than this would mean that the actuator is rate-undersized, and probably rate limiting, which is dangerous. In addition, reaching Q_{NL} more rarely would mean that the actuator is oversized. In summary, for the i^{th} actuator, we set $\sigma_i = Q_{NLi}/3$.

Having scaled the σ_i , we can now compute the average of the combined flow demands from N servoactuators, as follows:

$$Q_{Avg} = \text{mean of a sum} = \text{sum of means} = \sum_{i=1}^N \text{mean}_i \quad (6)$$

where for a rectified-Gaussian distribution, we found earlier that,

$$\text{mean}_i = \sigma_i \sqrt{2/\pi} \quad (7)$$

Substituting leads to,

$$Q_{Avg} = \frac{1}{3} \sqrt{\frac{2}{\pi}} \sum_{i=1}^N Q_{NLi} = 0.2660 \sum_{i=1}^N Q_{NLi} \quad (8)$$

Note however that this average corresponds to significant turbulence. In calmer conditions, the total average flow demand may only be 5-10% of this value. (The PSD method may offer a more accurate estimate of average flow demand during *calm* conditions, and this is the primary motivation for assessing it for this application. But in contrast, the proposed method uses average flow as an intermediate step in computing peak flow.)

Next, we define the peak flow as the average plus three sigma:

$$Q_{3\sigma\text{Peak}} = Q_{Avg} + 3 \sqrt{\text{variance of the sum}} \quad (9)$$

where,

$$\begin{aligned} \text{variance of a sum} &= \text{sum of the variances} \\ &= \sum_{i=1}^N \text{variance}_i \end{aligned} \quad (10)$$

For a rectified-Gaussian distribution, we computed earlier,

$$\text{variance}_i = \left(1 - \frac{2}{\pi}\right) \sigma_i^2, \text{ where } \sigma_i = \frac{Q_{NLi}}{3} \quad (11)$$

Substituting gives,

$$\begin{aligned} Q_{3\sigma\text{Peak}} &= \frac{1}{3} \sqrt{\frac{2}{\pi}} \sum_{i=1}^N Q_{NLi} + \sqrt{\left(1 - \frac{2}{\pi}\right) \sum_{i=1}^N Q_{NLi}^2} \\ &= 0.2660 \sum_{i=1}^N Q_{NLi} + 0.6028 \sqrt{\sum_{i=1}^N Q_{NLi}^2} \end{aligned} \quad (12)$$

Finally, we compute average and peak power demand. We multiply flow by the full working pressure, even though some losses are present. This captures the full power demand, including losses. Note that some quiescent demands (such as for internal leakage and electrical power for the controller) have been neglected here, but can be added for completeness.

$$\text{Power}_{Avg} = Q_{Avg} P_S \quad (13)$$

$$\text{Power}_{3\sigma\text{Peak}} = Q_{3\sigma\text{Peak}} P_S \quad (14)$$

As an example, we have simulated the total flight control flow demand for a hypothetical aircraft having eight actuators of various sizes, moving independently. Their individual no-load flows are tabulated in [Table 1](#), and the time history of the total flow is shown in [Figure 6](#). The proposed method, though approximate, clearly provides more accurate estimates of peak and average total flow demand. This