Cramér-Rao Lower Bound Optimization of Flush Atmospheric Data System Sensor Placement

Soumyo Dutta,[*] and Robert D. Braun[†]

Georgia Institute of Technology, Atlanta, GA, 30332-1510, USA

Flush atmospheric data systems take measurements of the pressure distribution on the forebodies of vehicles and improve the estimate of freestream parameters during reconstruction. These systems have been present on many past entry vehicles, but design of the pressure transducer suites and the placement of the sensors on the vehicle forebody have largely relied on engineering judgment and heuristic techniques. This paper develops a flush atmospheric data system design methodology using Cramér-Rao lower bound optimization to define the smallest theoretical variance possible from the estimation process. Application of this methodology yields Pareto frontiers of possible optimal configurations and identifies the number of ports which serve as the point of diminishing returns. The methodology is tested with a simulated Mars entry, descent, and landing trajectory.

Nomenclature

$\mathbf{f}$	Objective function
H	Measurement sensitivity matrix
k	Number of sampled points for the objective function
n	Number of ports
P	State covariance matrix
$\mathbf{p}$	Parameter vector for the objective function
R	Measurement noise covariance matrix
q_∞	Freestream pressure, Pa.
α	Angle of attack, rad
β	Sideslip angle, rad
σ	Standard deviation

Subscript and Superscript

i, j	Indices
y	y-component
z	z-component
$\hat{}$	Non-normalized value
$^-$	Mean value

I.   Introduction

FLUSH atmospheric data systems (FADS) provide measurements of pressure on a vehicle's forebody during flight and in conjunction with inertial measurement unit (IMU) data enable the reconstruction of the vehicle's freestream conditions, angle of attack, and sideslip angle.[1,2] FADS provides critical information for post-flight reconstruction of planetary entry, descent, and landing (EDL) vehicles, which typically suffer from a dearth of on-board measurements to make parameters of interest observable. FADS sensors have

---

[*]Graduate Research Assistant, Daniel Guggenheim School of Aerospace Engineering, AIAA Student Member.
[†]David & Andrew Lewis Professor of Space Technology, Daniel Guggenheim School of Aerospace Engineering, AIAA Fellow.

American Institute of Aeronautics and Astronautics

flown in many missions, such as the Space Shuttle orbiters and the 2012 Mars Science Laboratory (MSL)[2,3] and an example of a FADS sensor integrated into the forebody of a re-entry vehicle can be seen in Fig 1.[4]
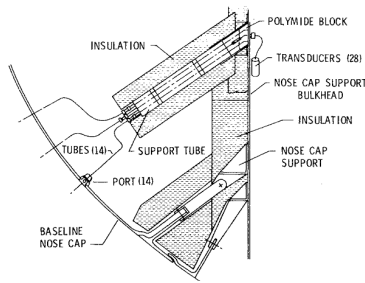

Figure 1. FADS integrated with the Shuttle Orbiter.[4]

Despite increasingly prevalent use, the methods for FADS design and sensor arrangement remain rudimentary. In spite of observations that different port configurations can vastly affect the effectiveness of the estimation,[5] past FADS sensors have always been placed in symmetrical annular or cruciform patterns based on engineering judgment rather than computationally-based rationale. FADS configurations are also often designed for fixed points in the trajectory, e.g. the sensor configuration is designed for Mach 5 and angle of attack of 2 deg., even though variations from the nominal condition leave these configurations suboptimal for the inverse estimation of parameters.

A thorough review of literature has shown two past studies[6,7] that have considered computational optimization for FADS sensor placement for EDL applications. Although both methods have advantages, they contain computationally-intensive objective functions that use the estimation method for the actual flight data reduction as their objective function. This makes the objective function values dependent on the estimation method used. The Cramér-Rao lower bound (CRLB) is the lowest bound on the estimated variance[8] irrespective of the estimation method chosen. Although a given estimator may not be able to reach the lower bound suggested by this quantity, the CRLB can serve as a good metric in producing the configuration of sensors that are most sensitive to a given set of parameters.[9] Reference 7 also considered a similar observability-based objective function in their work, but this work will be the first application of CRLB for FADS optimization.

This study demonstrates a methodology to optimize the design of a FADS sensor suite using CRLB as a design metric, improving the estimation of EDL parameters of interest including freestream conditions that are not typically observable using non-FADS re-entry sensors. Using CRLB instead of an inverse estimation process speeds-up the optimization, making the method more useful to designers while also grounding the objective function in estimation theory rather than in an empirically-based objective function. The methodology uses multi-objective optimization to allow the designer to explore the entire design space with one optimization process and produce Pareto frontiers of optimal configurations. A simulated Mars EDL trajectory, similar to the one flown by the Mars Science Laboratory, is used to test the optimization methodology.

II.   Past Use of Flush Atmospheric Data Systems

One of the first use of FADS were on the Viking landers.[10,11] The FADS sensors were arranged in an annular fashion with one port at the predicted stagnation point, as seen in Fig. 2. However, the transducers were add-ons during the latter stages of the design and did not go through a thorough calibration effort. This reason has been largely attributed to large amount of noise in the data which has made it generally unintelligible.[10]

The Shuttle Entry Air Data System (SEADS) program used a flush-mounted air data system on the shuttle's nose[4] with the port configuration (shown in Fig. 2(a)) arranged in a cruciform shape. This configuration was derived using heuristic methods dependent on engineering judgment.[4,12] Designers used error analysis to determine the minimum number of pressure ports and the ports were arranged in a cruciform manner to capture changes in the pitch and yaw plane. However, the cruciform configuration is only optimum if the trajectory has either angle of attack-only motion or sideslip angle-only motion. This configuration is non-optimum in terms of observability if both sideslip angle and angle of attack are non-zero at the same time. Since the SEADS configuration was optimized for a point in the trajectory, it is also not robust to variations from the nominal trajectory.

High-Angle-of-Attack Flush AirData Sensing (HI-FADS) systems have been used for aerodynamic test vehicles and conceptual studies for munitions guidance. The configurations were derived by adding annular arrays of pressure ports across the forebody of the vehicle, as seen in Fig. 2(b). Similar to SEADS, these

American Institute of Aeronautics and Astronautics

applications did not use physics-based optimization routines to select the transducer locations; instead, it was hoped that adding more ports at different radial and angular directions would capture the entire pressure distribution and allow for the estimation of the freestream condition.[13]

The air data system for the proposed Aeroassist Flight Experiment (AFE) (shown in Fig. 2(d)) was based on physics-based optimization. Deshpande et al.[6] used a gradient-based estimator and a genetic algorithm (GA) to optimize the distribution of the sensors in order to decrease the effect of normally distributed random noise from the pressure transducers. The residuals between the estimated parameters and their true values were then used in a single-objective function for the optimization routines. However, the study only considered reconstruction of a single trajectory point. As such, the reconstruction process that serves as the objective function for the optimization problem is expected to converge to a single trajectory state, similar to a situation in wind tunnel testing, but unlike the case of EDL reconstruction where the trajectory states are variable.

MSL carried a set of FADS transducers known as the Mars Entry Atmospheric Data System (MEADS). The MEADS science objective is to reconstruct dynamic pressure to within 2% and angle of attack and sideslip angle to within 0.5 deg. when the dynamic pressure is greater than 850 Pa.[3] To accomplish this, the MEADS sensors were arranged in a cruciform configuration around the forebody of the aeroshell (see Fig. 2(e)). The locations were based on the predicted pressure distribution on the aeroshell for a point in the trajectory where sideslip angle is small; however, no quantitative optimization procedure was conducted in the selection of the transducer locations. Based on the nominal trajectory, stagnation pressure is around P1 and P2, while P6 and P7 help reconstruct the sideslip angle. All ports besides P6 and P7 help reconstruct the angle of attack history. However, in reality with off-nominal trajectory conditions, the FADS configuration is non-optimum.[7]

(a) SEADS[4]

(b) HI-FADS[13]

(c) Viking[11]
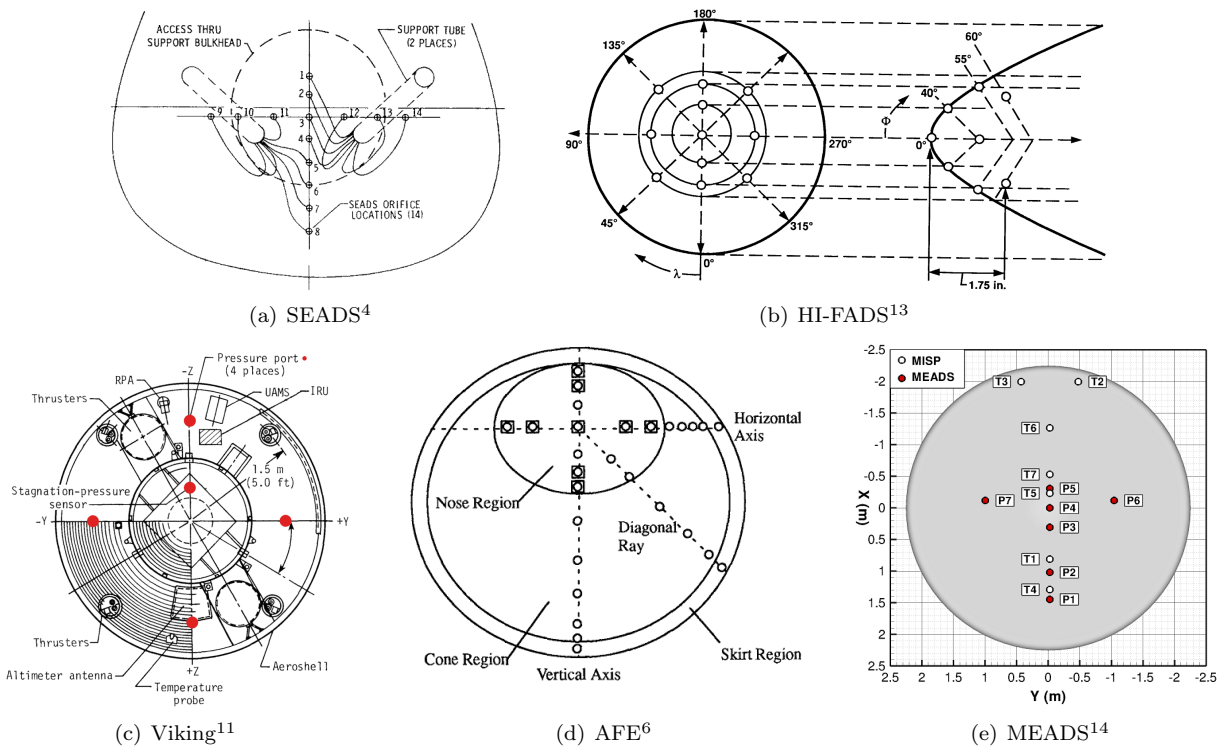
(d) AFE[6]

(e) MEADS[14]

Figure 2. Layouts of various FADS configurations.

As seen from these examples, traditional design of FADS configurations have used engineering judgment instead of physics-based modeling. The designers generally arrange the sensors for a fixed trajectory point rather than designing a system that will serve a time-varying trajectory. Bandwidth limits on on-board sensors for planetary entry missions make it crucial to make FADS configurations as efficient and optimized as possible in order to capture important pressure measurements under a range of conditions.

American Institute of Aeronautics and Astronautics

III. Multi-Objective Sensor Placement Optimization

The FADS sensor configuration problem can be cast as a multi-objective optimization problem with the objective of placing the sensors to accurately reconstruct parameters of interest of an inverse estimation process. Deshpande et al.[6] simplified this process by creating a single-objective optimization problem where they combined the multiple objectives of optimizing the reconstruction of dynamic pressure, angle-of-attack, and sideslip angle using weighting parameters. The use of these weighting parameters introduced subjectivity into the optimization process.

Dutta et al.[7] incorporated the concept of Pareto dominance in the optimization process. Pareto dominance allows one to find a set of optimal points that are an improvement over all other points in the design space.[15] The problem involves finding solutions that represent trade-offs among conflicting objective functions when multiple objectives are involved. The concept of domination, as described in Eq. (1), occurs when a objective function parameter vector $\mathbf{p}_1$ is better than another point $\mathbf{p}_2$ since the n-dimensional objective function vector $\mathbf{f}$ of $\mathbf{p}_1$ is no worse than the objective function vector of $\mathbf{p}_2$ and the function value $\mathbf{p}_1$ is strictly better that the function value of $\mathbf{p}_2$ along at least one dimension of the objective function.[15] All points that are non-dominated by any other point in the design space are members of the Pareto frontier.

$$
\begin{aligned}
\min \mathbf{f} &= \mathbf{f}(\mathbf{p}) \\
\forall i \in \{1, \cdots, n\} &: \mathbf{f}(\mathbf{p}_1)_i \leq \mathbf{f}(\mathbf{p}_2)_i \\
\exists j \in \{1, \cdots, n\} &: \mathbf{f}(\mathbf{p}_1)_j < \mathbf{f}(\mathbf{p}_2)_j
\end{aligned}
\tag{1}
$$

The optimization technique used for this paper is the same as that used in Ref. 7 —Non-dominated Sorting Genetic Algorithm II (NSGA-II).[15,16] NSGA-II is an evolutionary algorithm that can solve multi-modal problems such as the FADS sensor placement problem[7] better than traditional gradient-based methods, which are often stuck in local minima. NSGA-II uses Pareto dominance to find the best representation of the Pareto frontier,[15] and is considered a baseline technique in the field of multi-objective optimization.[17]

The optimization algorithm consists of three basic steps: initialization, sorting, and reproduction (see Fig. 3). A randomly generated population of feasible port locations is initially generated and then the members are assigned different fronts based on the objective function values. A crowding distance in the objective function space between different design parameters is also calculated to ensure diversity in the frontiers. In successive iterations or generations, the optimizer uses an elitist technique to pair the most Pareto dominant parents to produce children using tournament selection, crossover, and mutation operations and the process is repeated for a user-defined period of time. Additionally, a continuous search space is implemented using modifications to the traditional mutation and crossover operators.[18,19] It is expected that the final generation will be close to the theoretical Pareto frontier.[16]
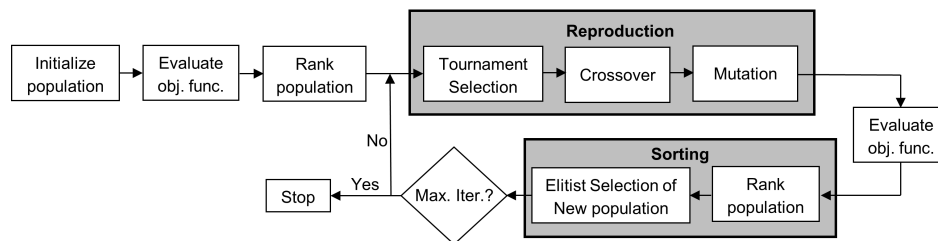

Figure 3. Flow diagram of Non-dominated Sorting Genetic Algorithm II.

IV. Observability and Cramér-Rao Lower Bound

Deshpande et al.[6] and Dutta et al.[7] used the residual between the estimated parameters and the true parameters as the objective function for FADS optimization. Although a proper metric for sensor locations sensitivity during parameter estimation, the residual-based function is time intensive since an inverse parameter estimation process has to be conducted every time the function is evaluated. This also makes the objective function values dependent on the estimation method chosen.

An optimized sensor suite should be independent of what specific estimation method is used to process the observations, leading to the concept of observability as a measure of how well the state vector can be

American Institute of Aeronautics and Astronautics

deduced from the outputs.[20] Observability metrics can be found for linear, time-invariant systems using the observability Gramian, but are hard to calculate for nonlinear, time-varying systems,[21] such as the FADS sensors on a re-entry vehicle. The Cramér-Rao Lower Bound provides a useful substitute by defining the theoretical lower bound of the expected uncertainty for an estimation process. CRLB is independent of the estimation method used and is defined as the inverse of the Fisher information matrix, which due to the Gauss-Markov Theorem results in a simple inequality as shown in Eq. (2). In this expression P is the state covariance matrix, H is the measurement sensitivity matrix, and R is the measurement noise covariance.

$$P \geq \left(H^T R^{-1} H\right)^{-1} \tag{2}$$

If the measurement uncertainties are uncorrelated and can be represented by an identity matrix, the CRLB simplifies further to the expression in Eq. (3), which is only a function of the Jacobian of the measurement equation with respect to the state vector (or H). For FADS, the measurement equation is the pressure measured by a transducer and this value is a function of the transducer location, trajectory states, and atmospheric conditions.[22]

$$P_{\text{lowest}} = \left(H^T H\right)^{-1} \tag{3}$$

V. Test Problem

The CRLB is calculated using the true state vector. These true states are provided here by the Program to Optimize Simulated Trajectories II (POST2),[23] which was used to generate a nominal EDL trajectory that is presented in Fig. 4. The trajectory is for a 4.5 m, 70-deg sphere-cone with the same geometry and specifications as MSL. This POST2-generated trajectory represents the truth data and the CRLB is calculated with this truth data. MSL's FADS objectives were defined while the dynamic pressure was 850 Pa. or above; so in this investigation, the FADS sensor optimization is also limited to this time frame.

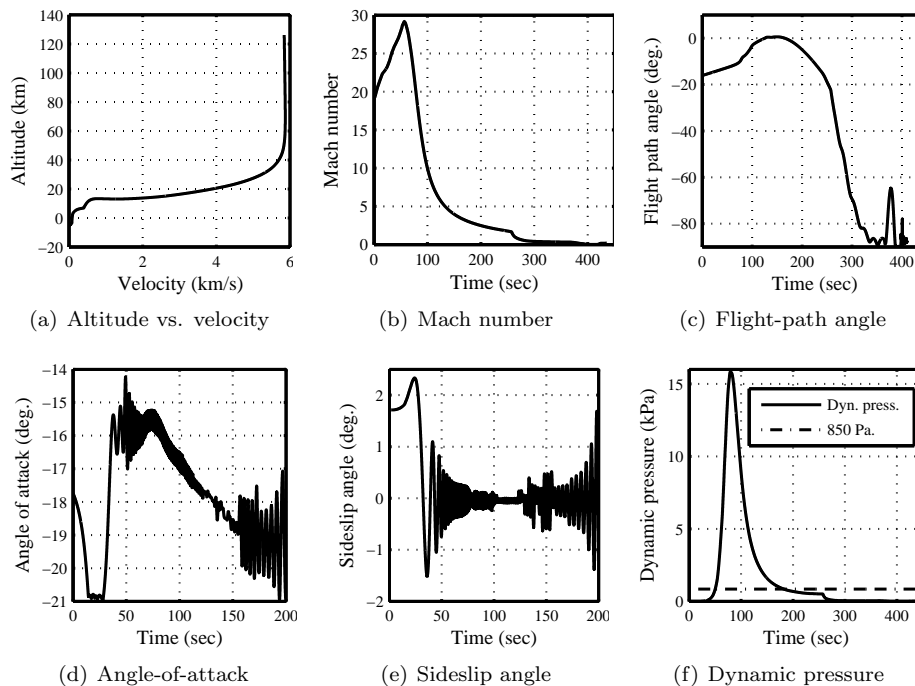| (a) Altitude vs. velocity | (b) Mach number | (c) Flight-path angle |
| (d) Angle-of-attack | (e) Sideslip angle | (f) Dynamic pressure |

Figure 4. Simulated Mars EDL trajectory used for creating the dataset for sensor location optimization.

VI. CRLB Sensitivity to Trajectory

Dynamic pressure (q_∞), angle of attack (α), and sideslip angle (β) are often the parameters of interest in FADS applications[4, 6, 12, 13] and were also the quantities for which MSL's science objectives were specified.[3] These terms serve as the state vector here, which means that H and the CRLB-calculated P will be calculated

American Institute of Aeronautics and Astronautics

with respect to these parameters. If only the diagonal of P is considered, then one gets uncorrelated variance of each of the parameters of interest as seen in Eq. (4).

$$\left(H^T H\right)^{-1} = P_{\text{lowest}} = \begin{bmatrix} \hat{\sigma}_\alpha^2 & & \\ & \hat{\sigma}_\beta^2 & \\ & & \hat{\sigma}_{q_\infty}^2 \end{bmatrix} \tag{4}$$

FADS optimization will locate a sensor configuration that minimizes the non-normalized, standard deviation ($\hat{\sigma}$) - square-root of the variance - for each parameter of interest over the length of the trajectory. Due to multiple parameters of interest, the function is multi-objective and the optimal configurations will be part of Pareto frontiers. However, since the CRLB is calculated at a given trajectory condition, there will be a CRLB for every trajectory point. The CRLB values throughout the trajectory need to be combined into one objective function vector that describes a metric of observability for a given FADS configuration.

To decide how to combine the various CRLBs into one metric, a sensitivity study of the CRLB-based, parameter standard deviations was conducted. These standard deviation values were normalized by the MSL science objective values, i.e. 0.5 deg. for angle of attack and sideslip angle and 2% of the actual dynamic pressure, to yield a normalized standard deviation (σ). NSGA-II was used to find Pareto frontiers at a few selected trajectory points for a seven-port configuration and the fronts' sensitivities to variations in trajectory were studied.

Figure 5 captures the Pareto frontier variation with Mach number. The range of the Pareto frontier decreases at certain Mach numbers - signifying an improvement in observability - but the trend is not based on an increase or decrease in Mach number. This is somewhat surprising since the flow regime actually changes from high, hypersonic speeds to low, supersonic speeds and one would assume that observability would change with the flow regime.
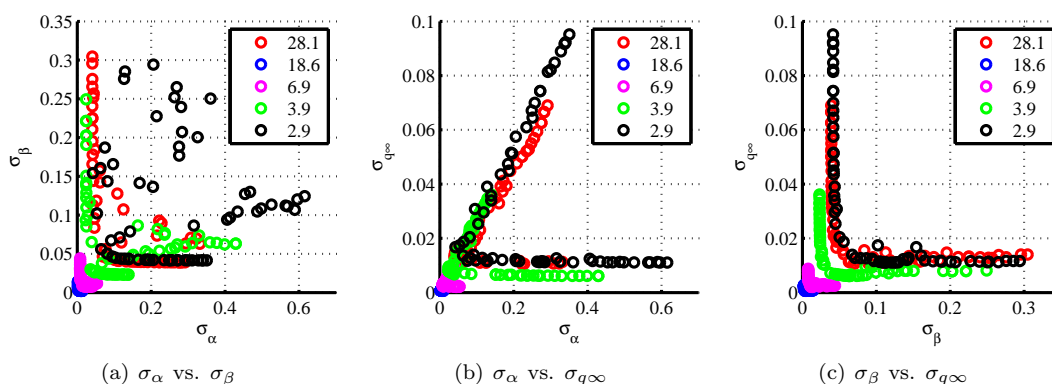

Figure 5. Sensitivity of CRLB to Mach number for the seven-port optimized configuration.

Instead, a stronger correlation is seen with dynamic pressure (Fig. 6), where higher dynamic pressure values lead to lower standard deviations and better observability. Due to CRLB's strong sensitivity to dynamic pressure, the metric used to aggregate CRLBs from various points in the trajectory should sample different dynamic pressure values to have a diversity of q_∞ conditions.

Angle of attack and sideslip angle also have an effect on the CRLB. Figure 7 displays the variation in Pareto frontiers at various sideslip angles for low q_∞ and high q_∞ conditions. Varying the sideslip angle slightly leads to a discernible change in the Pareto frontier, suggesting that variations in this quantity should be captured in the aggregate metric. Variations in angle of attack displayed similar trends.
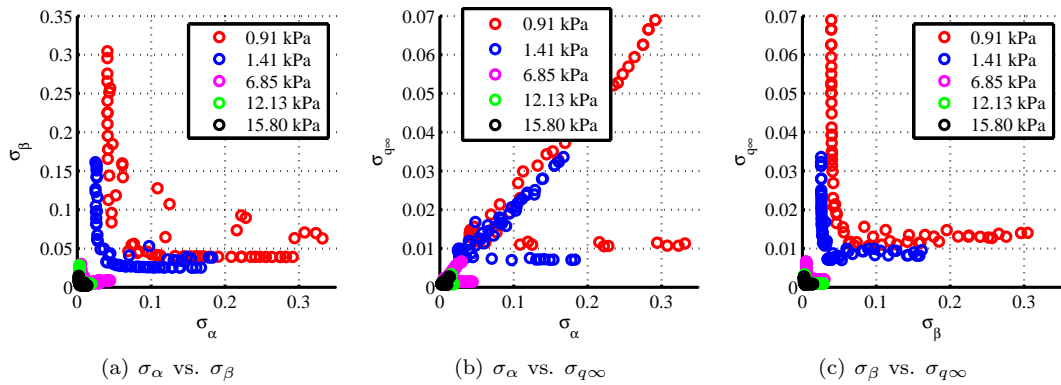
American Institute of Aeronautics and Astronautics

(a) σ_α vs. σ_β (b) σ_α vs. $\sigma_{q\infty}$ (c) σ_β vs. $\sigma_{q\infty}$

Figure 6. Sensitivity of CRLB to dynamic pressure for the seven-port optimized configuration.

(a) σ_α vs. σ_β (b) σ_α vs. $\sigma_{q\infty}$ (c) σ_β vs. $\sigma_{q\infty}$
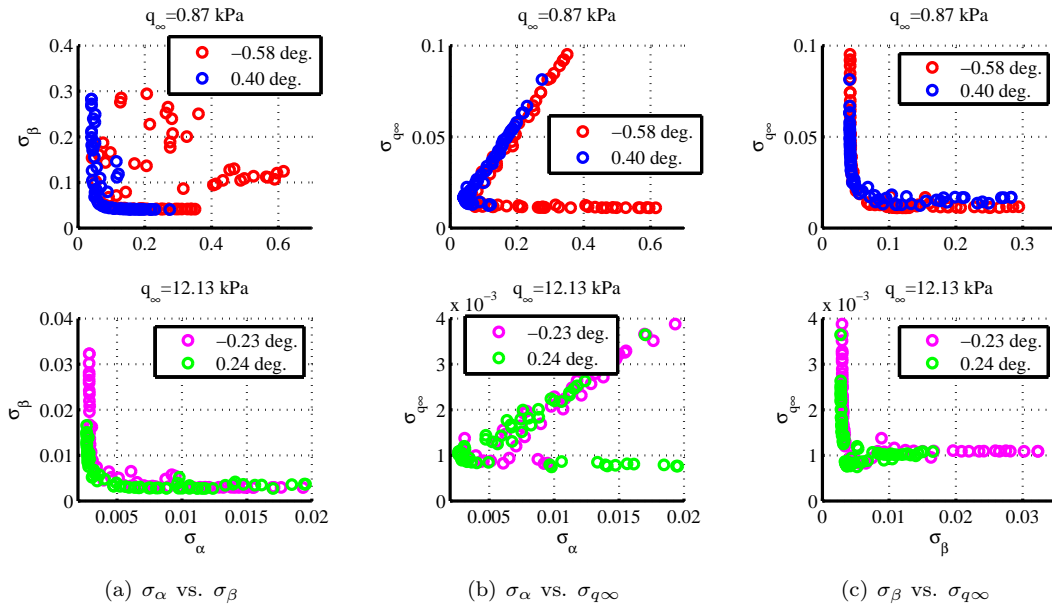
Figure 7. Sensitivity of CRLB to sideslip angle for the seven-port optimized configuration.

VII. Objective Function Formulation

In this investigation, 20 sample locations from the test problem trajectory were selected for the aggregate objective function. These discrete trajectory states are shown in Fig. 8 overlaid on the continuous trajectory. One can see that there has been equal distribution given to low and high dynamic pressure values and a variety of angle of attack and sideslip angles.
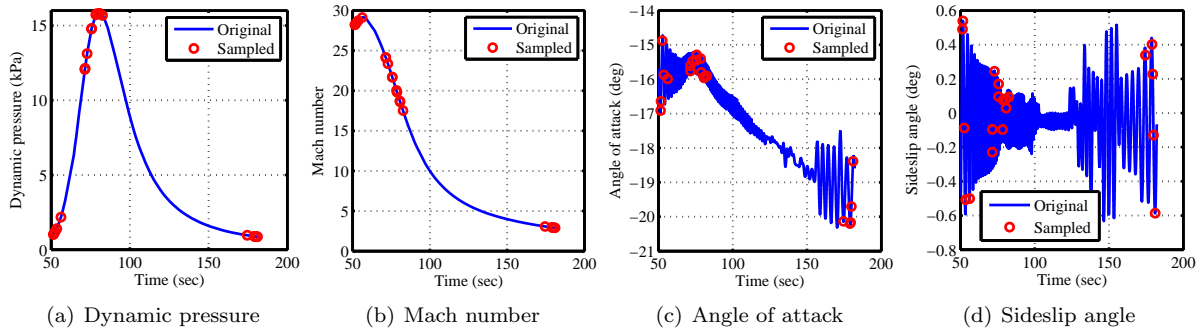

(a) Dynamic pressure   (b) Mach number   (c) Angle of attack   (d) Sideslip angle

Figure 8. Sampled trajectory points for the aggregate objective function.

American Institute of Aeronautics and Astronautics

Using these sampled points, the objective function $\mathbf{f}$ is defined in Eq. (5), while the optimization problem is defined in Eq. (6) where $\bar{}$ is the component-wise, arithmetic average of the objective functions. The inequality constraint in Eq. (6) is used to maintain a minimum spacing between the n port locations ($\mathbf{p}$), and this minimum distance $d_{\min}$ is chosen as 5 inches as it was done in Ref. 7.

$$\mathbf{f}(\mathbf{p})_i = [\sigma_{\alpha,i}, \quad \sigma_{\beta,i}, \quad \sigma_{q\infty,i}]^T \quad \forall \quad i = \{1, \cdots, k\} \tag{5}$$

$$\begin{aligned} \min \quad & \bar{\mathbf{f}}(\mathbf{p}) \\ \text{s.t.} \quad & |\mathbf{p}_i - \mathbf{p}_j| \le d_{\min} \quad \forall \quad i, j \in \{1, \cdots, n\} \end{aligned} \tag{6}$$

VIII.   Optimization Results

A.   Implementation and Computational Effort

The NSGA-II method is an evolutionary algorithm so finding the true Pareto frontier is not guaranteed. The CRLB-based optimization was conducted with a population of 128 candidate configurations over 500 generations and this process was repeated 10 times. The final Pareto frontiers were found using the combined results. Experimentation showed that a population size of 128 provided a good distribution across the design space to capture the near-optimal Pareto frontier and that 500 generations were enough to reach a stable set of non-dominated points. This process required in 640,000 function calls. These calls were made regardless of the type of configuration being optimized, i.e. a three-port or a seven-port situation had the same number of function calls. Additionally, the number of generations, population size, and number of repetitions were selected with conservatism to ensure that the optimization converged. It is possible that similar results could be achieved with far fewer function calls.

On the other hand, a brute-force search, where each possible configuration was checked, would demand a significantly larger number of function calls, especially as one increased the number of pressure transducers being optimized. This is described by Fig. 9 where the number of function calls using a Monte Carlo-like process on a representative grid of possible transducer locations and the CRLB-based, NSGA-II optimization process are compared. A smart culling process can reduce the grid size and the number of possible combinations to evaluate, but as shown in Fig. 9 even a coarser search grid size uses significantly more function calls compared to the NSGA-II optimization that uses a continuous search space.

(a) Possible ports (with 1585 points on grid)

(b) Comparison of number of function calls
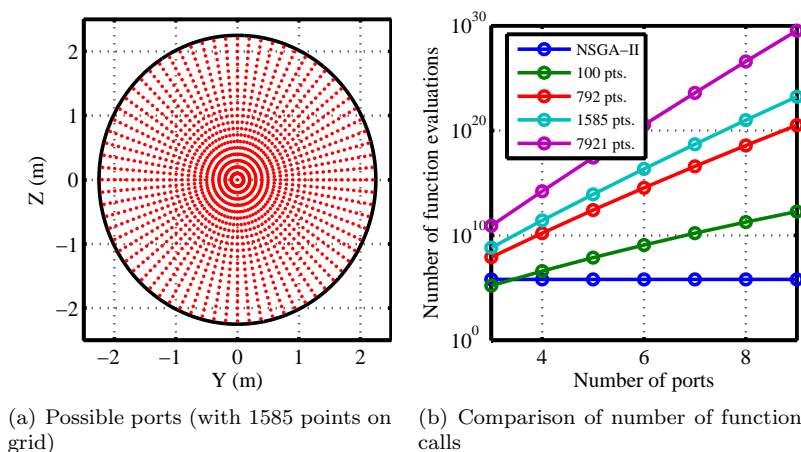
Figure 9. Comparison of function evaluations between a Monte Carlo-based method with various number of points on the grid and the CRLB-based optimization method.

In terms of computational speed, the CRLB-based objective function evaluation took around 10^{-3} s using a 3.4 GHz Intel i7 processor, with a slight increase in run time as the number of ports increased. On the other hand, Dutta et al.'s residual-based objective function took close to 25 s per run using the same hardware,[7] underscoring the improvement in speed if the CRLB-based function is used. Deshpande et al.[6] did not provide any computational data for comparison of their method.

American Institute of Aeronautics and Astronautics

B. Multi-objective Optimization Pareto Frontiers

The results of the CRLB-based FADS optimization are summarized in Fig. 10, which shows different views of the Pareto surface formed by the three-objective optimization. Pareto frontiers for three-port through nine-port configurations are shown. The two-port configuration did not provide a converged Pareto frontier in 500 generations and was excluded in this analysis.
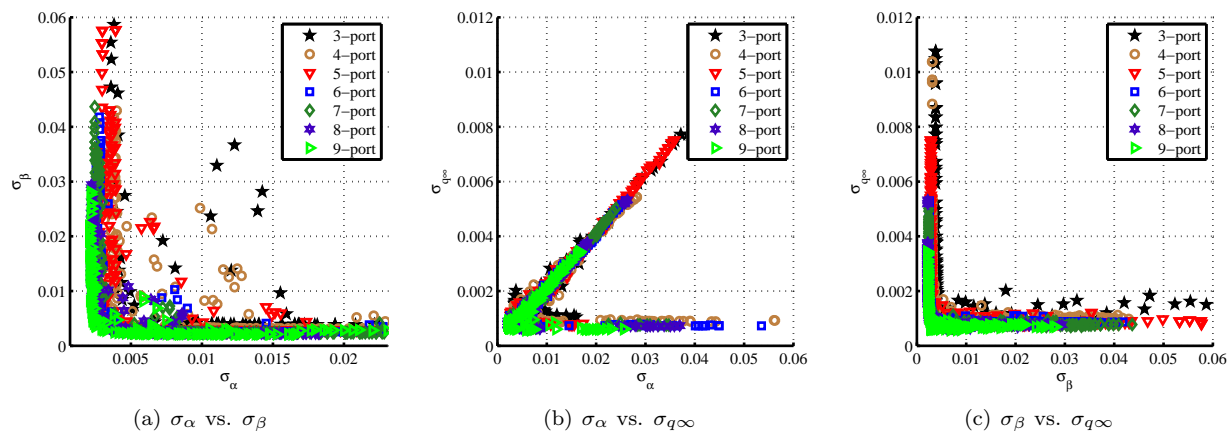

(a) σ_α vs. σ_β · (b) σ_α vs. $\sigma_{q\infty}$ · (c) σ_β vs. $\sigma_{q\infty}$

Figure 10. Pareto frontiers from multi-objective optimization for various port numbers.

Figure 10 shows that the Pareto frontiers come closer and closer to the origin as the number of ports increase. This is not surprising, since empirical evidence suggests that increasing the number of ports improves observability and leads to a lower objective function value. It also appears that the frontiers coalesce upon each other and not much is gained in observability after the six-port configuration. A six-port configuration thus appears to be the point of diminishing returns. The identification of the point of diminishing returns is investigated further later in this section.

Some representative configurations from the Pareto frontier are shown in Fig. 11. Although there is a structure to the port configurations, there was no constraint for symmetry and thus the optimized configurations are non-symmetrical. The representative configurations chosen in Fig. 11 are for either minimum σ_α, σ_β, or $\sigma_{q\infty}$ and the layouts exhibit these qualities. Dynamic pressure observability is achieved by placing ports near the stagnation point, which for this trajectory was around $y = 0$ and $z = -1$ m. Angle of attack observability is achieved by placing ports in the pitch plane on either side of the origin, while sideslip angle observability is maintained in a similar way except in the yaw plane. Numerical effects of the optimization are apparent in Fig. 11(b), where intuition would suggest that all of the ports would have $z = 0$, and Fig. 11(c), where all of ports don't have $y = 0$. Due to numerical noise, the optimization may not capture these nuances very well.

(a) Min. $\sigma_{q\infty}$ 3-port · (b) Min. σ_β 4-port · (c) Min. σ_α 7-port
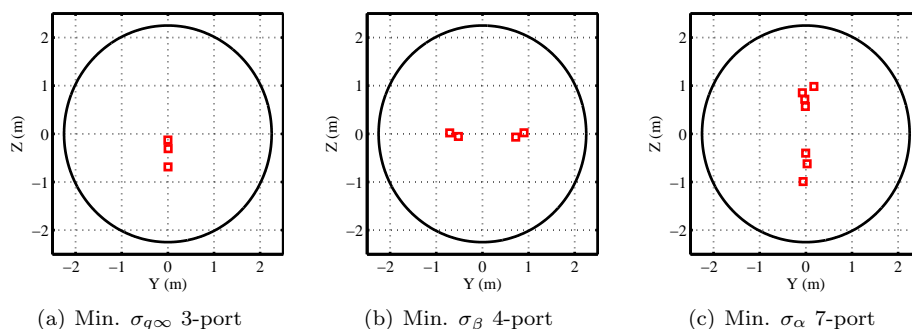
Figure 11. Port configurations from some representative points of the Pareto frontiers.

Although single objective optimal results are interesting to study, designers are more interested in configurations that can achieve good performance in all of the objective functions. Every point of the Pareto frontier is a non-dominated solution and hence it is hard to pick one point over another; however, one can define an equally-weighted compromise point which is closest to the ideal solution. The definition of this compromise

American Institute of Aeronautics and Astronautics

point may differ due to the type of weighting applied; however, the simplest such compromise point would come from a linear weighting scheme. Figure 12 explains the meaning of this linearized, equally-weighted compromise point on a nominal Pareto frontier.
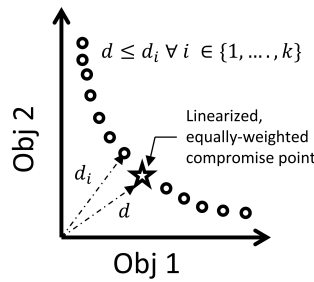

Figure 12. Definition of the linearized, equally-weighted compromise point of a Pareto frontier.

Using the linearized, equally-weighted compromise point as a benchmark of a good design, Fig. 13 shows some representative optimal configurations for various number of ports. Some broad design ideas can be gleaned from these configurations. It appears that annular like layouts - where ports are laid out in rings - are more preferred using this benchmark than cruciform layouts that were seen in some past configurations (Fig. 2). Additionally, many of these ports are concentrated near a ring of radius 0.5 m, which is near the area of a change in curvature as the aeroshell shape transitions from a spherical segment to the sharp cone. A change in curvature or geometry would make a port located in that region very sensitive to changes in the trajectory. Finally, all of the configurations have a port or two located near the stagnation point, suggesting that measuring pressure in this region improves observability of all of the parameters.

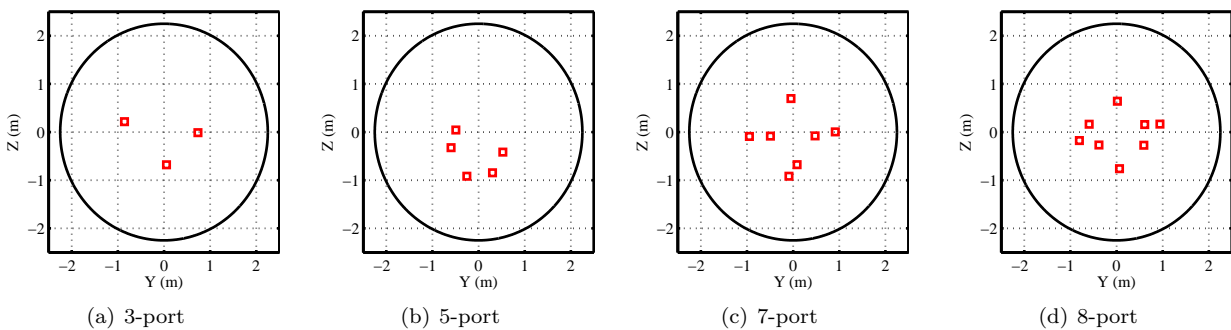(a) 3-port   (b) 5-port   (c) 7-port   (d) 8-port

Figure 13. Port configurations of some of the linearized, equally-weighted compromise points of the Pareto frontiers.

Interestingly, some of the configurations shown could be simplified further. For example, the 7-port configuration in Fig. 13(c) shows two ports near $z = -1$. If these two ports are combined to create a 6-port configuration, the objective function values does not degrade significantly from the 7-port values. Thus, although numerical optimization can quickly narrow down the design space to a list of good designs, it still leaves room for intuitive improvements by the designer.

The linearized, equally-weighted compromise point also allows one to visualize the point of diminishing returns. The diminishing point is apparent in Fig. 14, where the objective function values of the representative point of the Pareto frontier are plotted for different port configurations. If one is interested in only dynamic pressure reconstruction, a 5-port configuration seems to suffice as the point of diminishing returns. The Pareto contours in Figs. 10(b) and 10(c) that contain dynamic pressure dependency also support this assertion. However, when all of the parameters are considered together, one needs at least 6-ports to reach the point of diminishing returns, since the marginal return point is not reached for angle of attack and sideslip angle until this port configuration as seen in Fig. 14 and the Pareto frontier in Fig. 10(a).

American Institute of Aeronautics and Astronautics

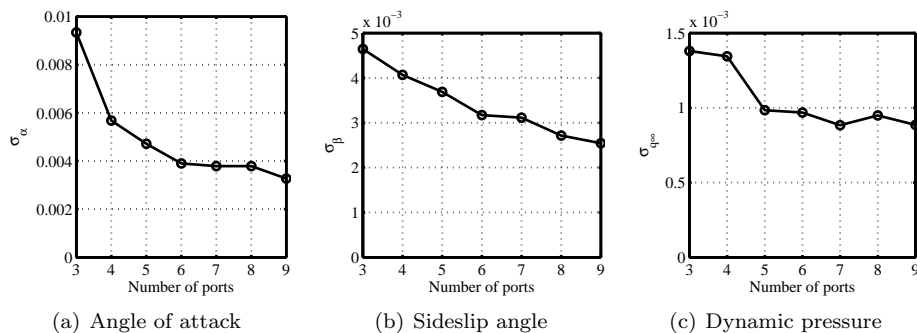(a) Angle of attack   (b) Sideslip angle   (c) Dynamic pressure

Figure 14. Identification of the point of diminishing return for non-symmetric configurations using objective values of the linearized, equally-weighted compromise points.

IX. Sensitivity to Pressure Models

The pressure model used to evaluate the objective function has a sensible effect on the optimization results. In Sec. VIII, a Computational Fluid Dynamics (CFD) derived pressure distribution was used in the function evaluations. However, CFD results have some uncertainties associated with them. One can use the classical Newtonian model to represent the pressure distribution, as was done by Deshpande et al.[6] Figure 15 captures the effect of using various pressure distributions by showing the Pareto frontiers for a 6-port configuration with the nominal CFD distribution, a CFD-based distribution perturbed randomly by 5%, and a Newtonian distribution.
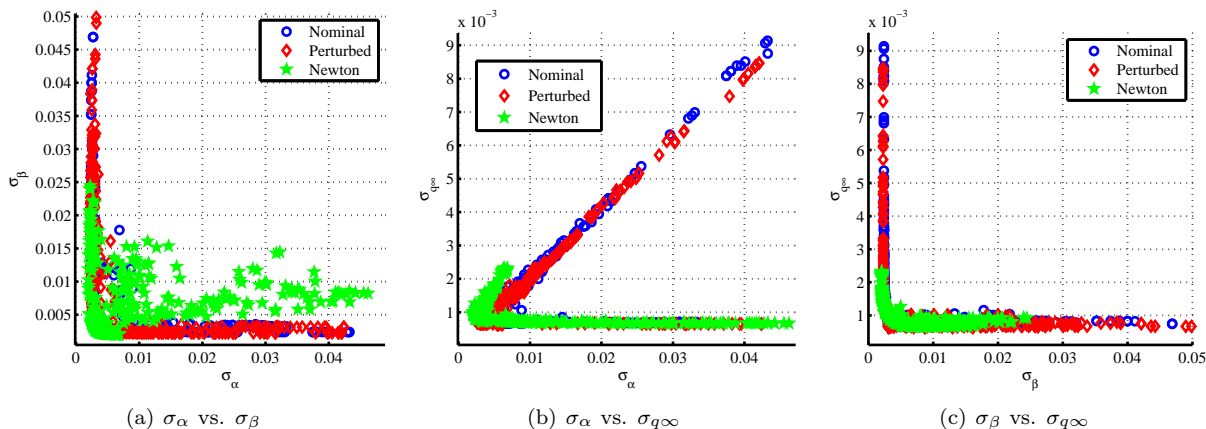

(a) σ_α vs. σ_β (b) σ_α vs. $\sigma_{q\infty}$ (c) σ_β vs. $\sigma_{q\infty}$

Figure 15. Comparison of Pareto frontiers for the 6-port configurations using various pressure models.

One does not see a major difference between the results of the two CFD-based optimizations, but the Newtonian distribution's Pareto frontier in the α-β slice appears less structured. Since the Newtonian distribution is based on a smooth function $2\sin\theta^2$ there are multiple port configurations that have similar objective function values and that makes the objective function space multi-modal.

Similar conclusions can be drawn when looking at the configurations described by the Pareto frontiers. Figure 16 shows the minimum σ_α configurations for the 6-port case using different pressure models. As expected for a suite making angle of attack more observable, all three configurations have transducers that are located on the pitch plane and have sets of ports that are on either side of the origin to increase the sensitivity to changes in the angle of attack. Due to the accumulation of ports in two locations, it seems that if one was only interested in angle of attack reconstruction a 2-port solution could suffice. In reality, designers are interested in reconstructing more than one parameter and hence would not be interested in an optimal configuration for only one parameter.

The nominal and perturbed CFD-based configurations yield extremely similar results, while the Newtonian configuration is different. As the CFD-distribution is not as smooth as the Newtonian pressure distribution, the objective function space is less multi-modal and the configurations shown in Figs. 16(a) and 16(b) represent samples from a basin of attraction. The Newtonian distribution-based objective space

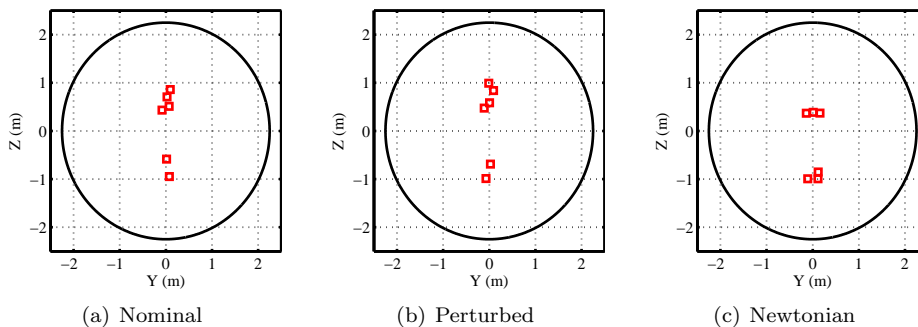American Institute of Aeronautics and Astronautics

Figure 16. Optimal σ_α 6-port configuration using various pressure models.

is more multi-modal and vastly different looking configurations are represented in the Pareto frontiers.

This exercise underscores the need to use computational methods to optimize a FADS suite and to tailor it for the proper conditions. Simply relying on engineering judgment and pressure distribution predictions from one set of tools - the modus operandi of designing FADS configurations in the past - is not enough to design a robust sensor suite. The variations caused by using different pressure distributions can be significant.

X. Sensitivity to Trajectory Perturbations

One of the main assertions of this comprehensive FADS placement optimization procedure is to make the chosen configuration robust and optimal over the entire trajectory. The effect of trajectory variations is clearly visible in Fig. 17 which shows the Pareto frontiers of a 6-port configurations using the nominal trajectory defined in Sec. VII and another trajectory perturbed by 5% from the nominal. Even though the perturbation is small, the Pareto frontiers show a visible difference.
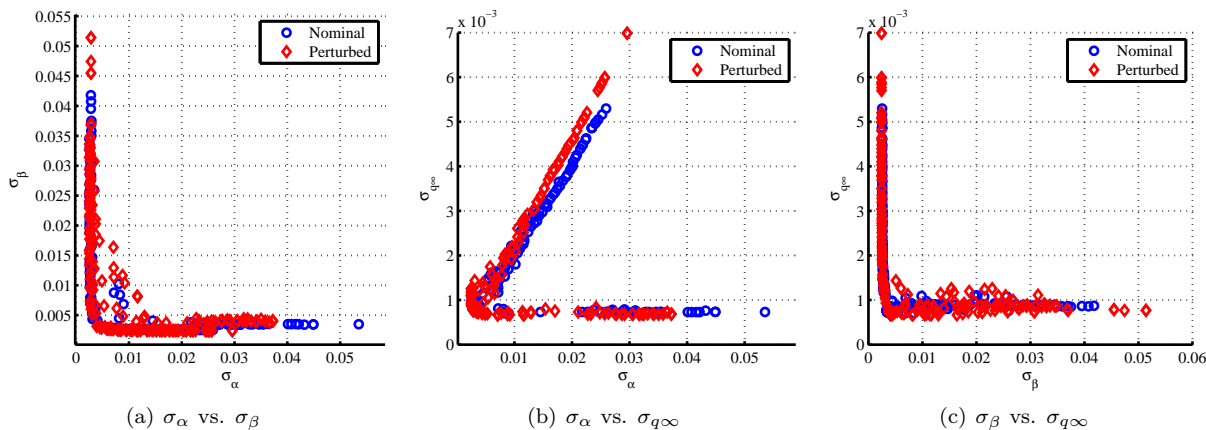

Figure 17. Comparison of Pareto frontiers for the 6-port configurations using various trajectories.

The effect of trajectory variation is also apparent in the optimized FADS configurations for minimum $\sigma_{q\infty}$ shown in Fig. 18. As expected, the ports optimizing the reconstruction of dynamic pressure are centered around the stagnation point. However, the slight difference between the optimal $\sigma_{q\infty}$ nominal (Fig. 18(a)) and perturbed (Fig. 18(b)) trajectory leads to a different looking port configuration. On the other hand, Fig. 18(c) shows a very different looking configuration that does not have the best $\sigma_{q\infty}$ value for either case but is still robust to the two different trajectories. This emphasizes the effect of trajectory perturbation and why FADS optimization should be performed across the entire trajectory and not at a single point of the trajectory. This way solutions that are robust to such perturbations can be found instead of optima based on point designs.

American Institute of Aeronautics and Astronautics
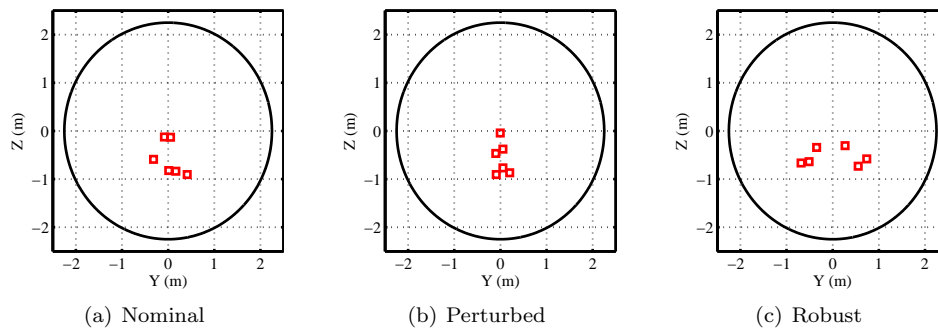
Figure 18. Optimal $\sigma_{q\infty}$ 6-port configuration using different trajectories.

XI.   Optimization with Symmetry Constraints

Past FADS sensors have had symmetric configurations (Fig. 2) and the FADS optimization study conducted by Deshpande et al.[6] explicitly set symmetry as a constraint. Due to the preference of symmetry in these past configurations, the optimization was also conducted with a symmetric constraint to look at how this affected the optimal configurations. The new objective function is shown in Eq. 7 and this optimization was repeated with various numbers of even-numbered pressure ports.

$$
\begin{aligned}
\min \quad & \bar{\mathbf{f}}(\mathbf{p}) \\
\text{s.t.} \quad & |\mathbf{p}_i - \mathbf{p}_j| \le \mathrm{d}_{\min} \quad \forall \quad i,j \in \{1, \cdots, n\} \\
& \mathbf{p}_{i,y} = -\mathbf{p}_{j,y} \quad \forall \quad i \in \{1, \cdots, n/2\} \quad \text{and} \quad j = i + n/2 \\
& \mathbf{p}_{i,z} = \mathbf{p}_{j,z}
\end{aligned}
\tag{7}
$$

The Pareto frontiers of the design space are shown in Fig. 19. Once again, 2-port configurations were excluded due to their poor convergence in the optimization. The number of ports that serves as the point of diminishing returns may be determined using Fig. 20, which shows the objective function of the linearized, equally-weighted compromise points. For certain objectives, like sideslip angle, there seems to be little difference in objective value by increasing the number of ports and the point of diminishing returns appears to be at 4-port configurations. If one is interested in only sideslip angle reconstruction, a 4-port configuration could suffice. But overall, considering all of the objectives at once, it appears that 6-port configurations are the points of diminishing returns as the Pareto frontiers coalesce upon each other as the number of ports increase and only marginal improvement in the uncertainty is gained by increasing the number of ports.
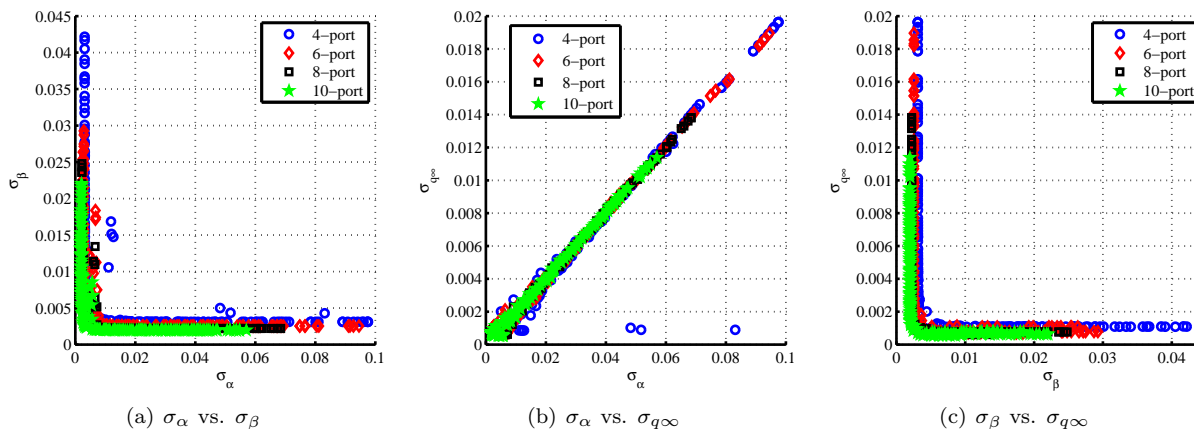

Figure 19. Pareto frontiers from multi-objective optimization for various port numbers with symmetry constraints.

The port configurations related to the linearized, equally-weighted compromise point of the Pareto frontiers are shown in Fig. 21.

There are some generalizing trends that can be observed when comparing the representative symmetric, linearized, equally-weighted compromise point configurations with their non-symmetric counterparts in

American Institute of Aeronautics and Astronautics

(a) Angle of attack   (b) Sideslip angle   (c) Dynamic pressure
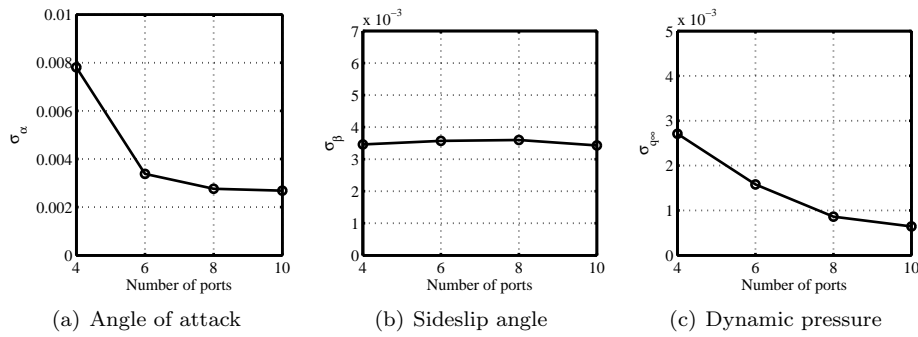
Figure 20. The point of diminishing return of symmetric configurations found using objective function values of the linearized, equally-weighted compromise point.
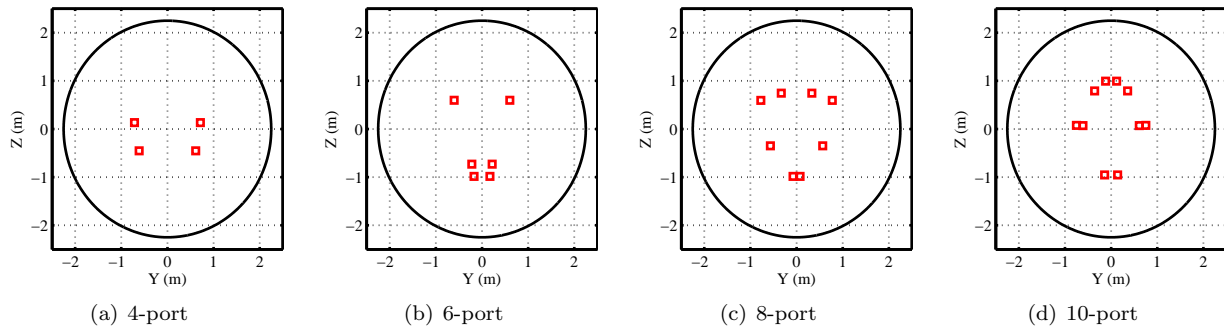

(a) 4-port   (b) 6-port   (c) 8-port   (d) 10-port

Figure 21. Port configurations of the linearized, equally-weighted compromise points of the symmetric optimization's Pareto frontiers.

Fig. 13. Similar to the situation with the non-symmetric cases, the optimal configurations appear to be annular rather than cruciform shaped. The ring of ports are in the region where the aeroshell shape transitions from a spherical segment to a cone. These design guides seem to reinforce lessons learned from the non-symmetric optimization. However, upon comparing their respective objective function values, as shown in Table 1, the effect of the slight differences between the two optimizations are apparent. The table shows that although the symmetric constraint leads to slight improvements in some objective function values over the non-symmetric cases, there is always one objective function value where the symmetric case performs very poorly compared to its non-symmetric counterpart. It can be inferred then that symmetric constraints may hinder the observability of the sensor suite in some fashion over the non-symmetric constrained results.

Table 1. Comparison between non-symmetric and symmetric configurations using the linearized, equally-weighted compromise points.

Ports	Percent difference from non-symmetric values		
	$\%\sigma_\alpha^*$	$\%\sigma_\beta^*$	$\%\sigma_{q\infty}^*$
4	57.30	-14.92	11.37
6	-13.26	37.22	-5.31
8	-27.12	32.55	-11.95

XII. Optimizing for Low Dynamic Pressure and Wind Speed Reconstruction

The observability of angle of attack, sideslip angle, and dynamic pressure were optimized by the objective function chosen in this study. However, wind speeds are also often important parameters of interest, and there are techniques that leverage FADS measurements and on-board IMU data to estimate these quantities.[24] Thus, the observability of wind speeds can also be a quantity that is added to the objective function. However, past studies have shown that the wind speed estimation is more a function of the IMU-based

American Institute of Aeronautics and Astronautics

velocity reconstruction,[2] so other changes to the objective function have to be also made to reflect this situation.

Another potential modification is to capture the effect of the measurement or sensor uncertainty in the objective function. Recall that for this objective function formulation the measurement noise covariance, R, was assumed to be an identity matrix. In actual sensors, the measurement uncertainty varies based on the flight regime or the dynamic pressure value and this can be reflected by varying R with the trajectory. In fact, FADS transducers are usually classified as either high dynamic pressure or low dynamic pressure sensors and engineers often design a port configuration for only one of these situations. For example, the MEADS suite that flew on MSL was only optimally calibrated for dynamic pressures above 850 Pa. although the transducers continued to take pressure measurements well below that limit. So one can optimize port configurations by including a varying R in the objective function and obtain results where one set of ports are optimized for high dynamic pressure regimes and another set is optimized for the low dynamic pressure regime.

XIII. Conclusions

Flush atmospheric data systems often serve as critical sensors of orientation angles and freestream conditions in atmospheric entry and flight dynamics applications. However, the design of most past flush atmospheric data system configurations was based on engineering judgment rather than a physics-based model. Additionally, the few studies that looked into the optimization of flush atmospheric data systems either did not consider the entire trajectory or had computationally expensive objective functions.

The current investigation uses the concept of observability to calculate objective functions that improve the ability to estimate angle of attack, sideslip angle, and dynamic pressure. Specifically, the Cramér-Rao Lower Bound is used to define the lowest possible standard deviation of the parameters of interest. The effect of trajectory is considered in creating the objective function value and it is found that dynamic pressure plays the most important role in the value of the Cramér-Rao-based uncertainties. An evolutionary optimization technique based on Genetic Algorithms is used to conduct the multi-objective optimization and Pareto frontiers are found for various port configurations. The optimization is conducted at first without any symmetrical constraints and then with symmetry enforced as a constraint. If one is interested in a single objective optimization, such as angle of attack or dynamic pressure, a 4 or 5-port configuration may be the point of diminishing returns depending on the parameter of interest. But when all parameters are considered, both symmetric and non-symmetric cases showed that a 6-port configuration is the point of diminishing return for the test problem which is based on the Mars Science Laboratory trajectory. Hence, adding an additional port after the 6th port has minimal gain.

The port configurations from some representative Pareto frontier points suggest design trends for improving observability. For instance, putting ports in annular fashion especially in the zone where a sphere-cone transitions from the spherical segment to the conical frustum seems to improve observability in all parameters of interest. Moreover, putting one or two ports near the stagnation point improves estimation of dynamic pressure greatly and this improvement also benefits other parameters of interest. The sensitivities of the optimization to pressure models and trajectory perturbations were also considered and underscored the need of flush atmospheric data system designs that are robust to a variety of conditions.

This study demonstrated a low computational cost, multi-objective optimization method that can be used to address the gap of physics-based flush atmospheric data system design methods. This method quickly allows a designer to determine design trends for robust configurations that maximize the observability of parameters of interest and the methods are easily amenable to include other parameters of interests.

Acknowledgments

American Institute of Aeronautics and Astronautics

References

[1]Dutta, S., Braun, R., Russell, R., Striepe, S., and Clark, I., "Comparison of Statistical Estimation Techniques for Mars Entry, Descent, and Landing Reconstruction," *Journal of Spacecraft and Rockets,* accessed May, 24, 2013. doi: 10.2514/1.A32459.

[2]Karlgaard, C. D., Beck, R. E., Keefe, S. A., Siemers, P. M., White, B. A., Engelund, W. C., and Munk, M. M., "Mars Entry Atmospheric Data System Modeling and Algorithm Development," AIAA 2009-3916, *AIAA Thermophysics Conference,* San Antonio, TX, 2009.

[3]Gazarik, M. J., Wright, M. J., Little, A., Cheatwood, F. M., Herath, J. A., Munk, M. M., Novak, F. J., and Martinez, E. R., "Overview of the MEDLI Project," IEEEAC 1510, *IEEE Aerospace Conference,* Big Sky, MT, 2008.

[4]Siemers, P. M. and Larson, T. J., "Space Shuttle Orbiter and Aerodynamic Testing," *Journal of Spacecraft*, Vol. 16, No. 4, 1979, pp. 223–231.

[5]Cobleight, B., Whitmore, S., Haering, J. E., Borrer, J., and Roback, V., "Flush Airdata Sensing System Calibration Procedures and Results for Blunt Forebodies," Tech. rep., NASA TP-1999-209012, 1999.

[6]Deshpande, S., Kumar, R., Seywald, H., and Siemers, P., "Air data system optimization using a genetic algorithm," AIAA 1992-4466, *AIAA Guidance, Navigation, and Control Conference,* Hilton Head, SC, 1992.

[7]Dutta, S., Braun, R., and Karlgaard, C., "Atmospheric Data System Sensor Placement Optimization for Mars Entry, Descent, and Landing," *Journal of Spacecraft and Rockets,* accessed July 12, 2013. doi: 10.2514/1.A32515.

[8]Bar-Shalom, Y., Li, X., and Kirubarajan, T., *Estimation with Applications to Tracking and Navigation*, John Wiley & Sons, Inc., New York, NY, 2001.

[9]Nehoria, A. and Hawkes, M., "Performance Bounds for Estimating Vector Systems," *IEEE Transactions in Signal Processing*, Vol. 48, No. 6, 2000, pp. 1737–1749.

[10]Ingoldby, R., Michel, F., Flaherty, T., Doty, M., Preston, B., Villyard, K., and Steele, R., "Entry Data Analysis for Viking Landers 1 and 2," Tech. rep., NASA CR 159388, 1976.

[11]Blanchard, R. C. and Walberg, G. D., "Determination of the Hypersonic-Continnum/Rarefied-Flow Drag Coefficient of the Viking Lander Capsule 1 Aeroshell from Flight Data," Tech. rep., NASA TM 1793, 1980.

[12]Pruett, C. D., Wolf, H., Heck, M. L., and Siemers, P. M., "Innovative Air Data System for the Space Shuttle Orbiter," *Journal of Spacecraft and Rockets*, Vol. 20, No. 1, 1983, pp. 61–69.

[13]Whitmore, S., Moes, T., and Larson, T., "High Angle-of-Attack Flush Airdata Sensing System," *Journal of Aircraft*, Vol. 29, No. 5, 1992, pp. 915–919.

[14]Edquist, K. T., Dyakonov, A. A., Wright, M. J., and Tang, C. Y., "Aerothermodynamic Design of the Mars Science Laboratory Heatshield," AIAA 2009-4075, *AIAA Thermophysics Conference,* San Antonio, TX, 2009.

[15]Deb, K., *Multi-Objective Optimization using Evolutionary Algorithms*, John Wiley & Sons, Ltd., 2001.

[16]Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, 2002, pp. 182–197.

[17]Reyes-Sierra, M. and Coello Coello, C. A., "Multi-Objective Particle Swarm Optimizers: A Survey of the State-of-the-Art," *International Journal of Computational Intelligence Research*, Vol. 2, No. 3, 2006, pp. 287–308.

[18]Deb, K. and Agarwal, R., "Simulated Binary Crossover for Continuous Search Space," *Complex Systems*, Vol. 9, 1995, pp. 115–148.

[19]Beyer, H.-G. and Deb, K., "On Self-Adaptive Features in Real-Parameter Evolutionary Algorithms," *IEEE Transactions in Evolutionary Computation*, Vol. 5, 2001, pp. 250–270.

[20]Tapley, B. D., Schutz, B., and Born, G., *Statistical Orbit Determination*, Elsevier Academic Press, Burlington, MA, 2004.

[21]Crassidis, J. and Junkins, J., *Optimal Estimation of Dynamic Systems*, Chapman & Hall/CRC, 2004.

[22]Dutta, S. and Braun, R., "Mars Entry, Descent, and Landing Trajectory and Atmosphere Reconstruction," AIAA 2010-1210, *AIAA Aerospace Sciences Meeting,* Orlando, FL, 2010.

[23]Striepe, S., Powell, R., Desai, P., Queen, E., Brauer, G., Cornick, D., Olson, D., and Peterson, F., *Program To Optimize Simulated Trajectories (POST II), Vol. II: Utilization Manual,* Version 1.1.6.G, 2004.

[24]Kelly, G. M., Findlay, J. T., and Compton, H. R., "Shuttle Subsonic Horizontal Wind Estimation," *Journal of Spacecraft*, Vol. 20, No. 4, 1983, pp. 390–397.

American Institute of Aeronautics and Astronautics