Design Variables		Code	Selected baseline
Eccentricity	[-]	A	0.02
Semi major Axis	[km]	B	9400
Inclination	[deg]	C	19.7
Payload Aperture Diameter	[m]	D	0.817
Satellite Transmitters Output Power (RF power)	[W]	E	1
Satellite Aperture-Antenna Diameter	[m]	F	0.1
Telemetry Data Rate	[Mbps]	G	1.25
Type of Solar Cells	[-]	H	8
Type of Batteries	[-]	I	9
Payload TRL	[-]	J	3
Launcher	[-]	K	2

Table 4.8   Design-variable settings for the selected design baseline.

4.3 Summary

In this chapter we presented global methods for supporting the engineering team during the design of complex space systems. Global MOO was used to support the design of complex space-mission architectures, starting from the mission objective and the major top-level requirements, by automatically generating and evaluating different scenarios.

As an example, a study on a hypothetical future space station to be implemented as a safe heaven for lunar missions and as a re-usable spaceship for deep-space missions is presented. The method generates different architectures with the definition of the systems and the phases in which they will be used. The systems taken into account comprehend modules for human support in space and propulsion modules. The mission architectures are then quantitatively and qualitatively evaluated. A comparative scenario cost-analysis is performed taking into account mission architecture, masses, technical complexity and system design. Finally, a mission success qualitative figure is estimated through mission duration, complexity and systems technology readiness level. The results demonstrated that a Skylab-like space station, thus with a single module, performs better than a multi-module space station in most of the cases. A dual-volume space station would allow for an increased level of the performance, albeit at a much higher cost.

Previous studies demonstrated that not all the mathematically-optimal solutions are also optimal from an engineering point of view. Robustness should also be taken into account. Therefore, in the second part of this chapter we presented a methodology for post-optimality studies to assess the robustness of the Pareto-optimal solutions, computed with an MOO algorithm. The proposed Pareto-Robust Optimization Approach (PROA) is based on AMH for sampling the design region in the neighborhood of the Pareto-optimal solutions. It allows for estimating a metric for the Pareto-Robustness and contributes to improving convergence of the known Pareto-front towards the true Pareto-front. The results demonstrated that PROA is able to compute the Pareto-robustness of the mathematically-optimal solutions correctly, steering the attention of the design team to the most Pareto-robust regions of the design space. The Pareto-Robustness metric discussed in this chapter is independent from the specific method used to compute the Pareto front, since it is solely based on the Pareto-dominance principle and it is applied after the computation of the Pareto front itself. The analysis of the test case of an Earth-observation mission has shown that the combination of a global MOO method with PROA effectively steered the design process, limiting the effort of the engineering team in the search for a single optimal and robust (Pareto-robust) design region. Indeed, the design region selected by PROA provided solutions far from three of the constraints, while still remaining close to the initially known Pareto front. The graphical information on sensitivity analysis and

contour plots can be considered a valuable aid for the engineering team, providing much more insight in the problem than any other single-point design methodology. In this case, the solution identified as the most Pareto-robust one can be considered quite uncommon, but still very reasonable due to the assumptions, presenting a satellite in a non sun-synchronous Medium Earth Orbit.

<div align="right">

Chapter *5*

Robust Optimization

</div>

As shown in the previous chapter, the Pareto Robust Optimization Algorithm is able to provide robustness information to variables dispersion of the system configurations (the solution to the optimization problem) that are considered Pareto-optimal. This approach only focusses on the solutions that are on the PF_{known}. PROA is fast, relatively easy to implement and computationally cheap, but there is the risk of missing out solutions that may be much more robust than these on the Pareto front: there may be robust solutions in regions of the objective space that are slightly away from PF_{known}. In this chapter we propose a robust optimization approach that takes robustness into account already from the beginning of the optimization process. We use the robustness concept to determine the effect of uncertainties on the design of the system.

Uncertainties in design variables and environmental factors are common in many engineering problems, and they must be taken into account when searching for robust optimal solutions. In robust MOO it is common practice to optimize the average performance instead of the nominal objective functions. To compute average performance, and to determine the compliance of the solutions to the constraints, sampling is needed in a neighborhood around each individual and the performance of each sample point must be evaluated. This increases the computational cost of robust optimization. In this chapter we present a repository-based approach that reduces the number of evaluations needed during robust optimization. Unlike most of the approaches available to date, we introduce methods to keep the joint probability density function of the input variables intact, when pre-existing points from the repository shall be used. This allows for computationally cheaper robust optimization also in the presence of non-uniform uncertain-variable distributions. The robust optimization of unmanned entry capsules, considering continuous shape-variation models, aerothermodynamics, flight mechanics, and thermal-protection system models at the same time is a valuable test-bed for the method presented here. In this chapter we discuss the results of minimizing the mass of the capsule(s) while maximizing the internal volume as well as the re-usability. In Section 5.1 we introduce the robust-optimization problem and relate it to the *atmospheric entry vehicle* test-case that is described in Appendix D. In Section 5.2 we refer to work in literature that is related to our original contribution to the field of robust-optimization. In Section 5.3 the proposed Double-Repository Archive Maintenance Scheme is presented, while in Section 5.4 we briefly describe the atmospheric entry vehicle problem that we consider. In Section 5.5 we provide the results of the analysis and in Section 5.6, finally, conclusions are drawn.

5.1 Introduction

The conceptual design of entry vehicles is commonly done in a number of sequential steps. One usually begins with a generic shape to get a first estimate of the aerodynamic properties

and uses a mass-point model for the initial trajectory design. Gradually, more detail is added and the outer shape is changed to accommodate specific mission and/or trajectory requirements.

This shape will largely define the aerothermodynamic characteristics of the vehicle. Since aerothermodynamic challenges, such as vehicle heating, remain one of the most difficult problems in atmospheric re-entry, an exploration of the possible shapes for a vehicle early in the design is advisable. It is advantageous to use a continuous model for the analysis, so that one is not limited to the analysis and comparison of a limited number of shapes, but is instead free to analyze any shape in the design space (Sudmeijer and Mooij, 2002). With only 5 geometric parameters it is possible to already model the geometry of Apollo-like shapes, as demonstrated later in Section 5.4.

The internal layout of the subsystems is usually only addressed at a later stage and the designers have to make sure that the mass properties (total mass, location of the center of mass and inertia tensor) meet the requirements. Deviations from these requirements can jeopardize the entire mission, because the loads on the vehicle may change, or the stability and control properties cannot be handled by the Guidance, Navigation, and Control (GNC) system any more. Further, uncertainties related to the entry conditions, environment, the characteristics of the thermal-protection system, and the design and allocation of the equipment on board, pose the multidisciplinary problem to be particularly cumbersome.

In this chapter we propose a multidisciplinary, robust optimization approach. This approach is then used for the design of unmanned entry capsules in support of the activities of the International Space Station (ISS). The problem is handled by minimizing the total mass of the capsules, while maximizing the internal available volume for carrying payload. As a third objective, we propose the maximization of the re-usability of the capsules, which can be seen as an attempt to push towards cheaper and more efficient solutions.

The shape, aerothermodynamic, and dynamic mathematical models are adapted from the work of Dirkx and Mooij (2011). It was demonstrated that the proposed simplified aerodynamic model can predict the aerodynamic forces and moments for ballistic shapes sufficiently well for use at a conceptual design stage. The multidisciplinary design framework is now enriched with a Thermal-Protection System (TPS) model, encompassing re-usable and ablative materials, as well as active cooling mechanisms. This allows for a complete conceptual design of an entry capsule.

Uncertainties in the design variables and environmental factors are integrated into the optimization process to handle *probabilistic* constraints. A probabilistic constraint is a constraint in the design or objective space that shall be satisfied with a pre-defined confidence level. The optimizer thus drives the search of optimal capsules towards those solutions that have the best *expected* performance under uncertain conditions, and that also meet the constraints with a given confidence level, pre-selected by the designer/decision-maker. A sampling-based approach is used to estimate the expected performance of the capsules and to determine the compliance with the probabilistic constraints. For each design point to be evaluated by the optimizer, a set of additional design points is generated around it, according to the joint Probability Density Function (PDF) of the uncertain variables and uncertain environmental factors, and evaluated. To limit the computational effort of the robust optimization, we adopt a double-repository archive maintenance scheme to save all the design-variable combinations computed during the process such that previous design points can be reused at future steps. The double-repository scheme allows to preserve the joint PDF of the uncertain input variables, therefore it is generally applicable with any type of multivariate distribution as input.

5.2 Related work

Although literature on robust (multi-objective) optimization is available, it seems to be still at a pioneering stage. In this section we summarize the main ideas and concepts regarding robust optimization that literature provides, and briefly discuss the main contribution of the proposed approach.

5.2.1 A survey on robust optimization

Jin and Branke (2005a) recognize that explicit averaging is the most common approach to robust optimization. This means that the expected value of the performance is taken into account instead of the value computed with a single point in the design space. The expected value is usually computed with Monte-Carlo integration (Branke, 1998):

$$f_{exp}(X) = \sum_{i=1}^{n} \frac{1}{n} f(X_i) \tag{5.1}$$

The model is indicated by f, while X represents a given point in the design space. In a problem with k design factors, $X = [x_1, x_2, \ldots, x_k]$. In this case $X_i = X + \delta_{\mathbf{i}}$, where the vector X_i is a set of n sample points in the proximity of X generated by the uncertainties in (some of) the design variables distributed according to the joint PDF ($\delta_{\mathbf{i}}$ is a realization from the joint PDF, $\delta_{\mathbf{i}} = [\delta_1, \delta_2, \ldots, \delta_k]$).

The expression in Eq. (5.1) is also called multi-objective robust solution of type I (Deb and Gupta, 2005). Deb and Gupta (2005) propose also an alternative way of approaching a multi-objective robust optimization problem introducing the so-called solutions of type II. In this case, the problem is to optimize the nominal objectives $f(X)$ given the following constraint:

$$\frac{\|f_{exp}(X) - f(X)\|}{f(X)} < \eta \tag{5.2}$$

The optimizer tends towards these solutions that are optimal, given a normalized change in the perturbed objective vector with respect to the nominal objective vector, restricted to a specified limit η.

It is also common that multi-objective algorithms are developed for robust design, where one of the objectives is the expected value and the other is the variance of the performance (Hassan and Crossley, 2008; Luo and Zheng, 2008b). In these cases the underlying problem is not multi-objective, and the interest of the designers is to capture the trade-off between robustness and optimality. With true multi-objective problems the approach of optimizing for both the expected value and the variance of each objective will become computationally expensive even with only few objectives, and the final Pareto front would be difficult to interpret.

Constraints in MOO engineering problems cannot be neglected, as, for instance, designing re-entry vehicles, as presented here. With robust-optimization problems the constraints become probabilistic. The reason is that they have to be verified considering a certain envelope of the joint PDF of the design variables and uncertain factors in the design space, and a certain envelope of the joint PDF of the performance on the objective space. Robust multi-objective optimization, with constraints, is even less explored in literature.

The robust solutions of type II could be considered suitable for dealing with probabilistic constraints. This is due to the fact that η gives rise to an envelope around the nominal objective vector. The constraints are then checked considering the envelope, not only the nominal objective vector. The drawback is that η is not necessarily linked to the probability distribution of the design variables, nor to the probability distribution of the objectives. Deb *et al.* (2009b) propose an integration of reliability-based concepts into evolutionary multi-objective algorithms
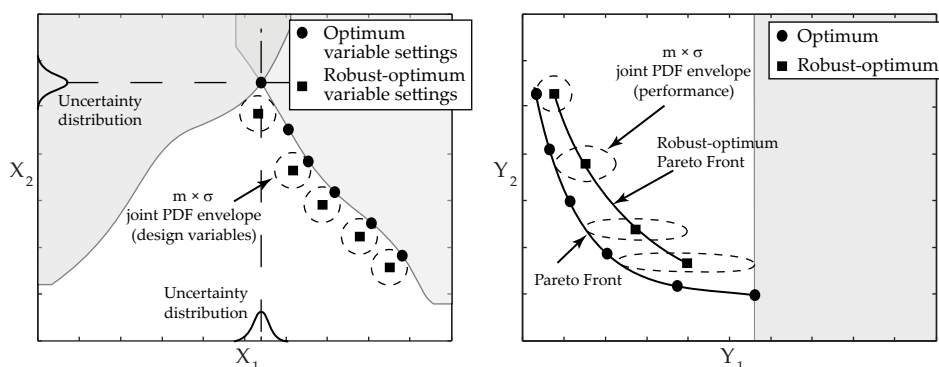
Figure 5.1   Schematic example of the effect of the probabilistic constraints in the design space (left) and objective space (right). The shaded areas represent the constraints. The dashed lines are the $m \times \sigma$ envelopes of the joint PDF around each specific design point in both the design and objective space.

to deal with this challenging problem. They consider the utilization of analytical methods to compute the envelope of the joint PDF in the design space that provides a pre-defined confidence level. Then, the constraints are checked given that particular envelope rather than a single point in the design space.

This means that the constraints have to be satisfied with that specific pre-defined confidence level, $m \times \sigma$. The parameter σ is the standard deviation of the joint PDF, while m is a factor specified by the designers. The analytical approach is feasible when the constraints are in the design space only. When the constraints are also defined in the objective space it is not generally valid anymore. The joint PDF of the performance given the joint PDF of the input may not be trivial to compute analytically, especially in the case of complex non-linear problems, see Figure 5.1, for instance. Sampling-based approaches may be useful: given a sample in the design space computed with the input joint PDF, the joint PDF of the performance can be estimated using mixture models, for instance, and thus the $m \times \sigma$ envelopes can be estimated as well (McLachlan and Peel, 2000; Marin *et al.*, 2005).

Sampling-based approaches can provide a way of checking probabilistic constraints in the objective space, but it is often argued that this option may come at the expense of a larger computational cost. This comment is true in general, also for unconstrained robust-optimization only. Indeed, Eq. (5.1) needs multiple sample points to be computed, and its accuracy increases proportionally to the sample size. Possible solutions proposed to reduce the computational cost of robust optimization include the use of variance-reduction techniques for sampling (Loughlin and Ranjithan, 1999), for instance, or also using *neighborhood* solutions previously computed during the optimizations process (Branke, 1998). Some other authors also propose repository-based approaches to save these *neighborhood* solutions in such a way to maintain a well-spread distribution in the repository (Kruisselbrink *et al.*, 2010). Paenke *et al.* (2006) propose, instead, to use metamodels computed with *neighborhood* solutions and to perform the uncertainty analysis using the metamodel instead of the real model, to save computational time.

The main drawback that we see in the direct utilization of *neighborhood* solutions from a repository, is that this approach cannot be used in the presence of specific joint PDF of the input variables. It will work only when the input variables are uniformly distributed (and that is not necessarily always the case) and when there is already a *sufficient* number of sample points in the design space. Without a mechanism for keeping the joint PDF of the uncertain input variables, repository-based robust-optimization cannot be used in general and the advantage obtained by re-using old solutions cannot be exploited.

5.2.2 Main contributions

In this chapter we propose a double-repository archive maintenance scheme, to allow for re-utilization of old solutions in the design space and at the same time to preserve the joint PDF of the input factors (**?**). The scheme allows for the storage of two separate repositories. One repository contains the values of the design variables distributed according to the joint PDF in the design space, this is the *real repository*. The other repository contains the values of the same design variables in a non-dimensional, uniformly-distributed space, this is the *fictitious repository*. Specific rules for transforming the design points from the real to the fictitious repository, and vice versa, are discussed later in this paper. Whether a certain design point in the real repository is suitable to be reused or not, is determined on the basis of its Euclidean distance from the actual design point. The threshold Euclidean distance that allows to determine if two design points are *close enough* is dependent on the joint PDF. A large value of the joint PDF means a large value for the threshold Euclidean distance, and vice versa.

5.3 Double-repository archive maintenance scheme

The approach we propose to efficiently incorporate robustness into an optimization process is based on the Unified Sampling Method (USM) presented in Chapter 3 of this thesis. The main purpose of the USM is to allow for taking all types of uncertainty distributions into account, also epistemic (*i.e.,* , where the designer only estimates the probability intervals), exploiting the properties of a *low-discrepancy* sampling technique like the one developed by Sobol' (Sobol', 1979; Antonov and Saleev, 1979). The basic idea is to uniformly sample in the probability space (between 0 and 1) and then to transform the sample to the actual design space using the Cumulative Distribution Function (CDF). In Figure 5.2 we show an example of this procedure considering two independent, uncertain design variables that are distributed according to a Gaussian distribution.

First, the *real* design space (horizontal axis of Figure 5.2(a)) is subdivided into uniform intervals that are transformed to the probability space using the inverse Gaussian CDF (vertical axis of Figure 5.2(a)), see the dashed lines of Figure 5.2(a). This procedure allows us to subdivide the probability space into *bins* with a width proportional to the value of the PDF. We call these *bins* Basic Probability Assignments (BPAs). Then, a uniform sampling is performed on the probability space, as shown in Figure 5.2(b), obtained with 300 sample points from a Sobol' sequence. The choice of using a Sobol' sequence is due to its characteristic of being a space-filling sampling method that creates a well-distributed set of points. Any other sampling method can be used for this step.

At this point, the CDF is used to transform the design points in the actual design space, as shown in Figure 5.2(c), see also the solid lines of Figure 5.2(a). As can be seen from Figure 5.2(c) the sample points are distributed according to a bi-variate Gaussian distribution, as expected. The number of BPAs considered in the double transformation will affect the quality of the final distribution, the more the better. A trade-off exists, however, with the speed of computation of the final joint PDF: using many BPAs will require a higher computational time, in proportion to the computational cost of each single simulation. According to our experience 30 to 50 BPAs, for each variable, provide already a good balance between the two.

The double-repository archive maintenance scheme is based on the implementation of the principle explained in Figure 5.2 during the whole optimization process. To obtain a robust estimation of performance and constraints, by preserving the input joint PDF, a sample as shown in Figure 5.2(c) needs to be generated and evaluated for every single design point that the optimizer analyzes. Doing so, Eq. (5.1) can be evaluated and the probabilistic constraints can be checked as previously discussed. Suppose to begin an optimization process evaluating
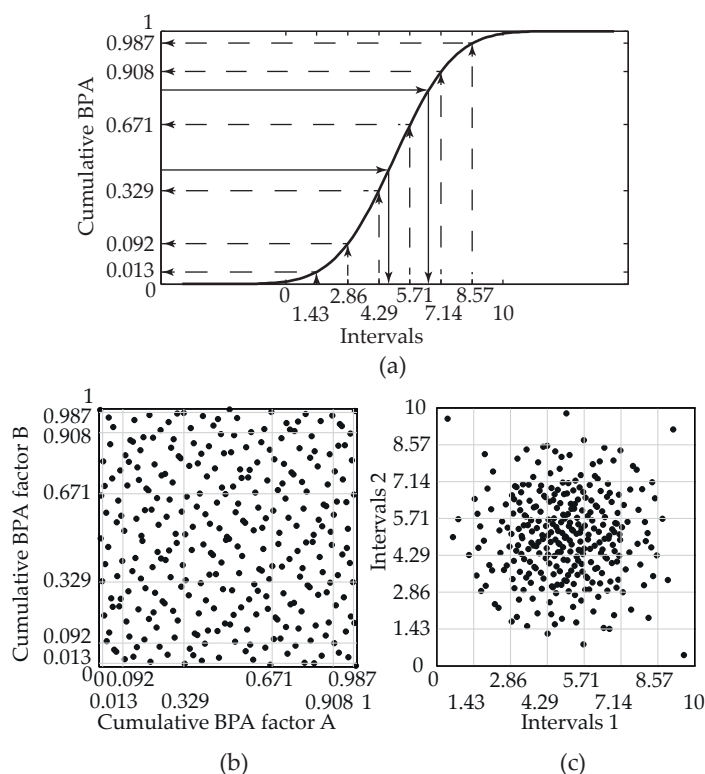
Figure 5.2   Unified Sampling Method. (a)Cumulative Density Function of a Gaussian distribution of one single variable. (b)Uniform sampling in the probability space of two independent variables. (c)Resulting sampling in the design space.

the first design point. A sample is generated from the joint PDF, then Eq. (5.1) is computed, constraints are checked for the desired confidence level, and all these points are saved in the repositories, that were empty prior to the beginning of the optimization. The sample points are saved in two distinct repositories. The *real repository* contains the true values of the design points, the points computed as shown in Figure 5.2(c). The second repository, the *fictitious repository*, contains the design points in a non-dimensional space instead. They are the points computed as shown in Figure 5.2(b).

When the next design point needs to be evaluated, another sample is generated. At this stage, there are already design points in both repositories. To determine whether some of these pre-existing points can be re-used or not we use the approach schematically shown in Figure 5.3. Using the non-dimensional version of the new sample, the crowding distances for each point are computed. The crowding distance is a measure used to understand if a pre-existing sample point is *close enough* to the sample point, which has yet to be evaluated. If so, the new sample point is substituted with the pre-existing one, thus saving computational time. In Figure 5.3(a), the black dots represent the sample points in the non-dimensional space, the radius of each black circle represents the crowding distance of that specific point instead. The crowding distance depends on the BPA to which the sample point belongs. The smaller the BPA, the smaller the crowding distance. In particular, for each BPA the crowding distance is determined as half the minimum distance between two sample points belonging to the same BPA. If there is only one sample point in one BPA, instead the crowding distance is determined as half of the smallest dimension of the BPA itself. This is done because, for events with smaller probability (*i.e.,* towards the edges of Figure 5.3(a)) the value of the PDF gets more sensitive to small offsets. Once the crowding distance is computed, the sample is transformed and checked against the points in the real repository.
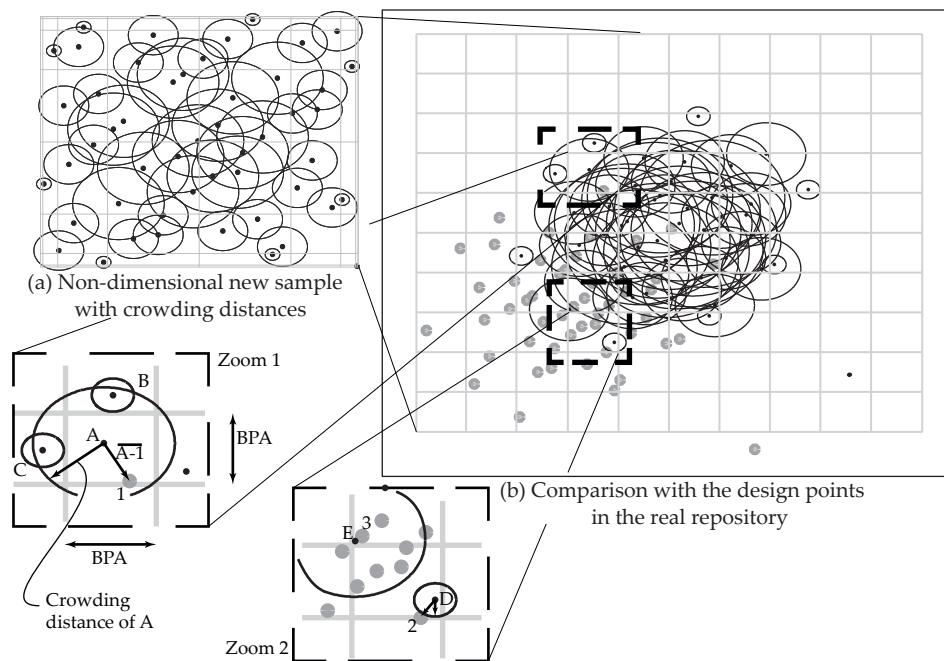
Figure 5.3   Double Repository Archive Maintenance Scheme, example. (a) The black dots are the sample points in the non-dimensional space, the radius of each black circle represent the crowding distance of that specific point. (b) The gray points are these belonging to the real repository already, coming from previous evaluations, the black dots are the point belonging to the new sample.

In Figure 5.3(b), the gray points are those belonging to the real repository already, coming from previous evaluations, and the black dots are the points belonging to the new sample instead. To preserve the joint PDF, the sample points are checked for each single BPA of the new sample. It is therefore guaranteed that, also in case of replacement of the design points with points already in the repository, in each BPA the number of sample points remains equal to that of the original sample. For instance, in Figure 5.3(b) Zoom 1, the *design point A* is substituted by the pre-existing *design point 1*, because they are in the same BPA and because the distance *A-1* is less than the crowding distance of *A*. In Figure 5.3(b) Zoom 2, the *design point D* cannot be substituted by the pre-existing *design point 2*. Even though they are in the same BPA, their relative distance is larger than the crowding distance of *D*. Further, even though there are more pre-existing points in that BPA, they are not taken into account because the original BPA only presents one design point, namely *point D*.

With the same reasoning one can conclude that *point E* is substituted by *point 3* and that no more points are added to that specific BPA. This is a clear example of the fact that a joint-PDF preservation mechanism is necessary. Considering all the sample points in this BPA will disrupt the distribution, leading to inaccurate estimation of the average performance and probabilistic constraints. The final result of the substitution is presented in Figure 5.4. This represents the true repository with both the samples from the first (gray dots) and the second (black dots) iteration of the optimization process. We would like to mention again that the gray dots have already been evaluated while the black dots have not. The squares in Figure 5.4 are the pre-existing points that substitute some of the points in the new sample, indicated by the cross symbols. In this case it means that 11 simulations are not executed and that Eq. (5.1) and the constraints are evaluated with the values from the simulations coming from the remaining black dots plus the values from the old simulations (the squares). This approach for sampling the design space using the double-repository is called *recursive sampling*.
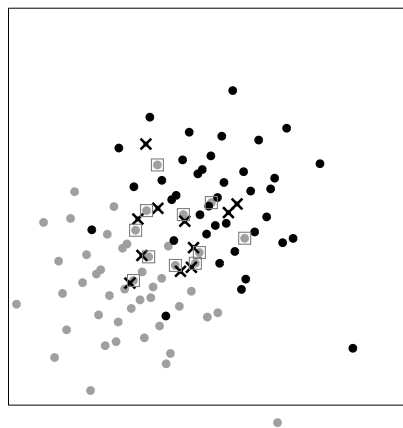
Figure 5.4   Double Repository Archive Maintenance Scheme, final sample. ● Sample points already in the real repository. ● New sample points. ✖ points from the new sample that are not simulated. ☐ points from the repository that are re-used.

5.3.1   Validation of the double-repository archive maintenance scheme

The double-repository archive maintenance scheme is validated with some test problems found in literature. The first set of problems is derived from the work of Deb and Gupta (2005), see Eqs. (5.3) to (5.5).

The test problem 1 can be expressed as follows:

$$
\begin{aligned}
\text{Minimize} \quad & (f_1(X), f_2(X)) = (x_1, h(x_1) + g(X)\,S(x_1)) \\
\text{Subject to} \quad & 0 \le x_1 \le 1 ,\ -1 \le x_i \le 1 \quad \forall i = 2, 3, \ldots, k , \\
\text{where} \quad & h(x_1) = 1 - x_1^2 , \\
& g(X) = \sum_{i=2}^{k} \left(10 + x_i^2 - 10\cos(4\pi x_i)\right) , \\
& S(x_1) = \frac{\alpha}{0.2 + x_1} + \beta x_1^2
\end{aligned}
\tag{5.3}
$$

Here, $\alpha = 1$ and $\beta = 1$ are used. The test problem 2 is identical to test problem 1, except that here $\alpha = 1$ and $\beta = 10$. The test problem 3 can be expressed as follows:

$$
\begin{aligned}
\text{Minimize} \quad & (f_1(X), f_2(X)) = (x_1, h(x_2) + g(X)\,S(x_1)) \\
\text{Subject to} \quad & 0 \le x_1, x_2 \le 1 ,\ -1 \le x_i \le 1 \quad \forall i = 3, 4, \ldots, k , \\
\text{where} \quad & h(x_2) = 2 - 0.8\exp\left(-\left(\frac{x_2 - 0.35}{0.25}\right)^2\right) - \exp\left(-\left(\frac{x_2 - 0.85}{0.03}\right)^2\right) , \\
& g(X) = \sum_{i=3}^{k} 50 x_i^2 , \\
& S(x_1) = 1 - \sqrt{x_1}
\end{aligned}
\tag{5.4}
$$

The test problem 4, finally, can be expressed as follows:

$$
\begin{aligned}
\text{Minimize} \quad & (f_1(X), f_2(X)) = (x_1, h(x_1, x_2) + g(X)\,S(x_1)) \\
\text{Subject to} \quad & 0 \le x_1 \le 1 ,\ -0.15 \le x_2 \le 1 ,\ -1 \le x_i \le 1 \\
& \forall i = 3, 4, \ldots, k , \\
\text{where} \quad & h(x_1, x_2) = 2 - x_1 - 0.8\exp\left(-\left(\frac{x_1 + x_2 - 0.35}{0.25}\right)^2\right) - \\
& \exp\left(-\left(\frac{x_1 + x_2 - 0.85}{0.03}\right)^2\right) , \\
& g(X) = \sum_{i=3}^{k} 50 x_i^2 , \\
& S(x_1) = 1 - \sqrt{x_1}
\end{aligned}
\tag{5.5}
$$

The results of the validation are shown in Figures 5.5 and 5.6. In particular, in Figure 5.5 we present the solutions to the test problems 1 to 4, according to the definition of robustness in Eq. (5.1), *i.e.,* multi-objective robust solutions of type I. The results are presented with
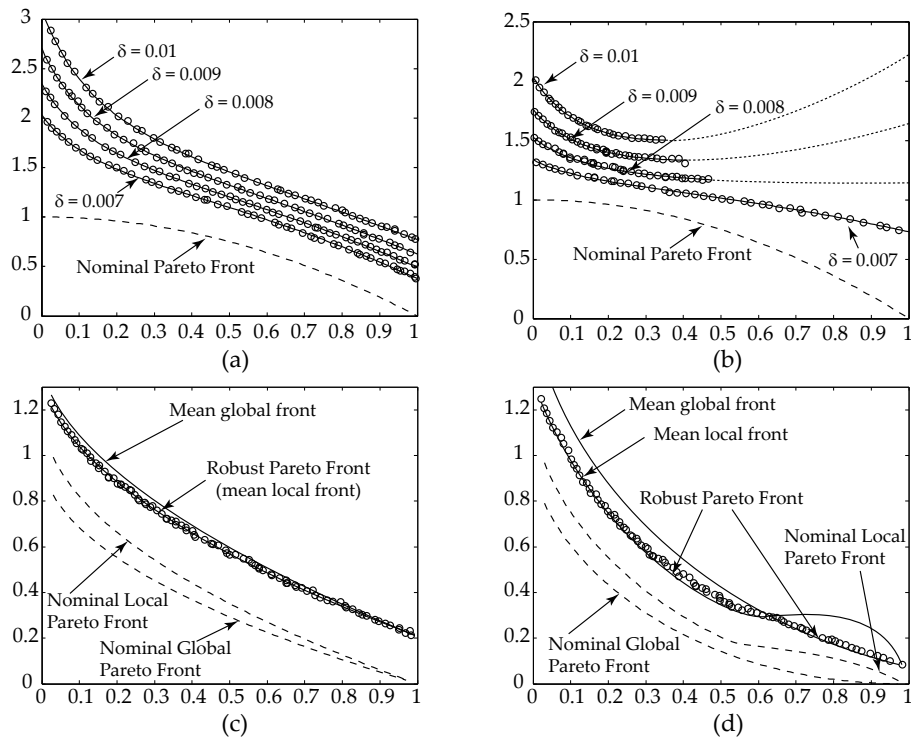
Figure 5.5   Robust Pareto-front (type I robustness) of the test problems 1 to 4 with uniform distribution of the uncertain design variables. On the axis there are the objectives f_1 and f_2. The symbol δ indicates the width of the uncertain boundaries for the design variables. Graphs (a) and (b) were obtained with the values of δ indicated in the figure. Graphs (c) and (d) were obtained with $\delta = 0.03$.

an increasing value of δ, that represents the extent of the uncertain boundaries of the design variables. Therefore, $\delta = 0.001$ corresponds to an uncertainty of $\pm 0.1\%$ around the design variables, for instance. With increasing values of δ the robust solutions gets more and more distant from the nominal optimal Pareto front, see Figure 5.5(a) and (b). The mean effective fronts are generated as a consequence. This effect is typical when considering robust solutions of type I. All solutions that lie on the mean effective fronts are identical to those lying on the nominal Pareto-optimal front in terms of design-variable settings, in the case of Figure 5.5(a). The shape, instead, is quite different and presents an interesting stimulus for discussion. The mean effective fronts are non-convex, compared to the nominal Pareto front that is convex, instead. This characteristic can potentially induce a robust-optimization problem to be more complex to be solved by an optimizer. Non-convex problems are harder to solve by an optimizer. In Figure 5.5(b) not all the solutions corresponding to the nominal Pareto front are then also considered robust solutions. The introduction of robustness, makes part of the Pareto front being dominated in the robust Pareto front (see dashed lines). This means that for $\delta > 0.007$ there are original solutions that are so sensitive to variables perturbation that they do not lie on the robust Pareto-front. In Figure 5.5(c) and (d) the robust Pareto front was obtained with $\delta = 0.03$. These two problems present a global and local nominal Pareto front each. The solid lines represent the mean effective local and global fronts. In Figure 5.5(c), we can observe that the local and global Pareto fronts switch role when robustness is considered. The mean effective local Pareto front is the robust Pareto front. This means that the nominal global Pareto front presents solutions that are less robust, *i.e.,* more sensitive to variable perturbations, than the solutions on the local Pareto front. In case of Figure 5.5(d), the robust Pareto front is a mix between the mean effective global Pareto front and the mean effective local Pareto front.
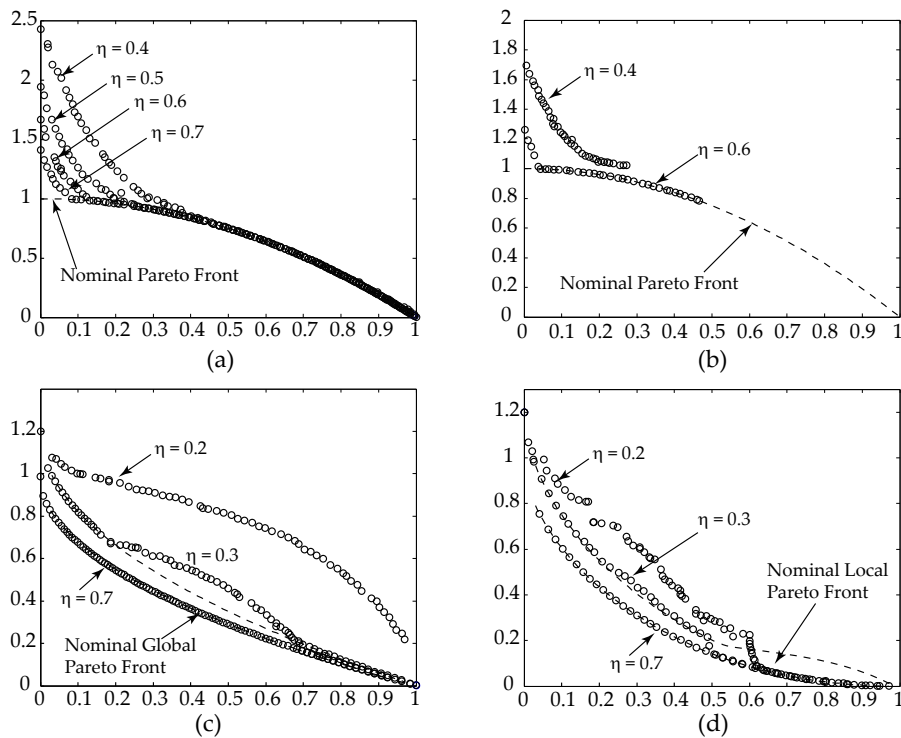
Figure 5.6   Robust Pareto-front (type II robustness) of the test problems 1 to 4 with uniform
                distribution of the uncertain design variables. On the axis there are the objectives f_1
                and f_2. The symbol δ indicates the width of the uncertain boundaries for the design
                variables. Graph (a) was obtained with $\delta = 0.007$, graph (b) with $\delta = 0.006$, and graphs
                (c) and (d) were obtained with $\delta = 0.03$.

In Figure 5.6 we show the results of the robust multi-objective type II solutions of Problems
1 to 4. The same problems defined earlier are now solved using MOO using the constraints
indicated in Eq. (5.2), using different values of η, with $\delta = 0.007$ for Problem 1, $\delta = 0.006$
for Problem 2, and $\delta = 0.03$ for Problems 3 and 4. The results demonstrate that the definition
of robustness affects the robust Pareto front that is obtained. Indeed, the differences between
Figures 5.5 and 5.6 are only due to the different definition of robustness, that is type II for the
results in Figure 5.6.

In Figure 5.7 it is demonstrated how the recursive sampling allows for a substantial reduc-
tion of the number of model evaluations for obtaining the robust Pareto front, if compared
to an approach that does not make use of a repository. The lines in Figure 5.7 represent the
number of model evaluations as a function of the number of generations of the optimization
algorithm, used for the problem in Figure 5.5(a). For an approach without repository, as the ap-
proach used by Deb and Gupta (2005), the trend is linear, as shown by the single-dashed line.
In this case, the number of model evaluations to obtain the results is equal to 7.5 million. The
continuous lines, instead, represent the trend of the number of model evaluations for varying
values of δ, using the double repository archive maintenance scheme, *i.e.,* recursive sampling.
A smaller δ makes the intersection of the design-point sets generated for each single evalua-
tion less likely to happen, thus requiring in general more evaluations (less substitutions) than
the cases when δ is larger. In general, there is a significant reduction of the number of model
evaluations, reducing them to 2 million, circa.

To reduce the computational load even more, also an adaptive selection of the sample
points could be implemented, following the advice of Luo and Zheng (2008a). Adaptive selec-
tion means that the sample points are considered one-by-one, and that Eq. (5.1) is computed
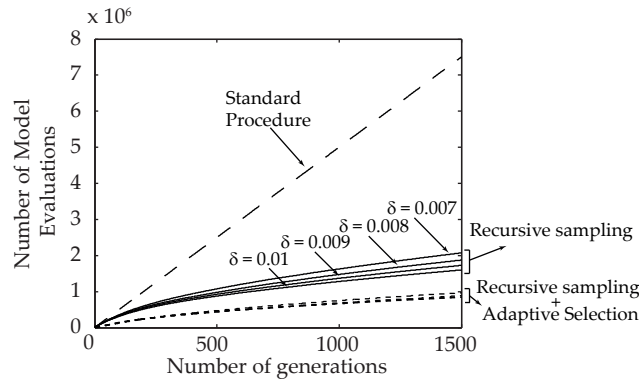
Figure 5.7   Evolution of the number of model evaluations in the optimization process of the test problem shown in Figure 5.5(a). The symbol δ indicates the width of the uncertain boundaries.

progressively with one sample point at the time. When the value of Eq. (5.1) does not change for subsequent executions of the model, the simulations are stopped for the current sample. Suppose that one decides that the sample for computing the average performance should contain 100 sample points. If, after evaluating 60 sample points the average performance does not change *substantially*, then the remaining 40 sample points are not simulated, saving computational time. The effect of a combined utilization of the recursive sampling and the adaptive selection is also shown in Figure 5.7. Referring to the results presented in Figure 5.7, the recursive sampling allows to save approximately 70% in terms of number of model evaluations. For moderately computationally expensive models, this difference can turn an infeasible analysis into a feasible one. The Sobol' sampling technique allows to exploit the adaptive selection at its best, reducing a further 10 % the computational cost of the analysis.

As discussed in the previous section, the type of robustness considered in this thesis, for robust MOO, is close to the type II robustness introduced by Deb and Gupta (2005). However, we consider the envelope around each design point in the design space, and the corresponding objective points in the objective space, determined by a certain probability level of the joint-PDF, selected by the engineering team.

To demonstrate the principle, we test it on a modified version of a test problem proposed by Deb *et al.* (2005), namely *DTLZ9*. The original problem is modified by adding a linear constraint in the design space, see the gray area in Figure 5.8 (right), and an additional constraint in the objective space. The problem is thus formulated as follows:

$$\text{Minimize } f_j = \sum_{i=j-1}^{j} x_i^{0.25}, j = 1, 2$$

$$\text{Subject to } g_1(x) = f_2^2(x) + f_1^2(x) - 1 \geq 0 \tag{5.6}$$

$$g_2(x) = f_1 - 0.1 \geq 0$$

$$g_3(x) = x_1 + x_2 - 0.7 \geq 0$$

$$0 \leq x_i \leq 1, \text{ for } i = 1, 2$$

The original problem does not include the constraints g_2 and g_3, and the nominal Pareto front is shown as the continuous line in Figures 5.8, 5.9, and 5.10 (right). The corresponding variable sets are represented by the continuous line in Figures 5.8, 5.9, and 5.10 (left). When adding the constraints g_2 and g_3 the design and objective space are modified, and the resulting Pareto front and relative variable settings change as well. The diamond symbols in Figures 5.8, 5.9, and 5.10 represent the new Pareto front (right), and the relative variable settings (left) when the constraints g_2 and g_3 are added to the original problem.

Figure 5.8   Robust optimization results, problem Eq. (5.6). The gray areas represent the constraints. Robust solutions with confidence level equal to 0.9.
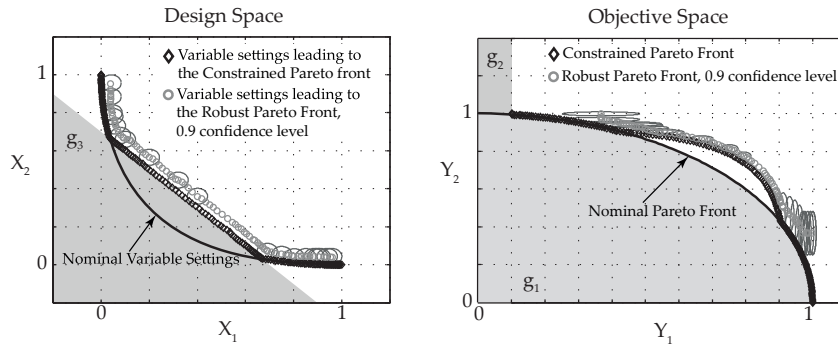

Figure 5.9   Robust optimization results, problem Eq. (5.6). The gray areas represent the constraints. Robust solutions with confidence level equal to 0.97725.

At this point we add uncertainty to the design variables. In particular, we use two normal distributions for the two variables with $\sigma = 0.03$. In Figures 5.8, 5.9, and 5.10 we show the effect of the confidence level on the robust Pareto front. In these figures the large circles/ellipses represent the envelopes of the joint PDF on the input variables (left) and on the objectives (right) that yield a probability level of 0.9, 0.97725, *i.e.*, 2σ, and 0.99875, *i.e.*, 3σ. The robust-optimization process pushes the Pareto front away from the non-robust Pareto front in such a way to meet the constraints with a confidence level of 90%. The same phenomenon is also visible in the design space.

The approach to robust optimization, using reliability concepts, was also explored by Deb *et al.* (2009b) using a double-loop optimization approach. The first loop is used to drive the f_{eff} towards the Pareto front, and the second loop is used to check for the probabilistic constraints. As a last test case for our approach we use a test problem proposed by Deb *et al.* (2009b) in the framework of reliability optimization:

$$\text{Minimize } f_1 = x_1, \, f_2 = \frac{1 + x_2}{x_1}$$
$$\text{Subject to } x_2 + 9x_1 - 6 \geq 0,$$
$$-x_2 + 9x_1 - 1 \geq 0, \tag{5.7}$$
$$0.1 \leq x_1 \leq 1, 0 \leq x_2 \leq 5$$

The problem has 2 uncertain variables with $\sigma = 0.03$, normally distributed. The results are presented in Figure 5.11 for increasing value of required confidence level. As in the previous cases we use values for the confidence level that are equal to 0.9, 0.97725, and 0.99875. Further,
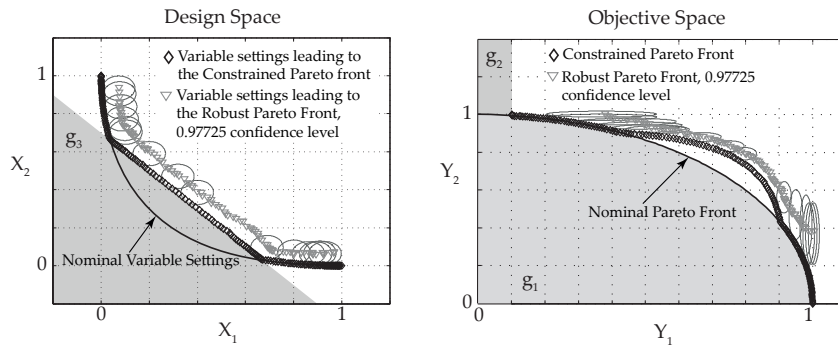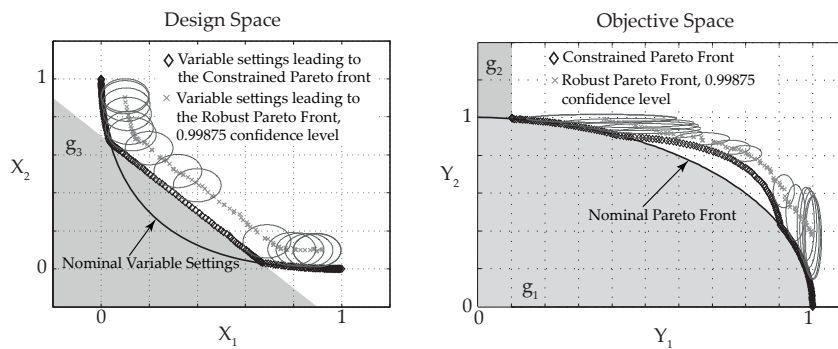
Figure 5.10   Robust optimization results, problem Eq. (5.6). The gray areas represent the constraints. Robust solutions with confidence level equal to 0.99875.
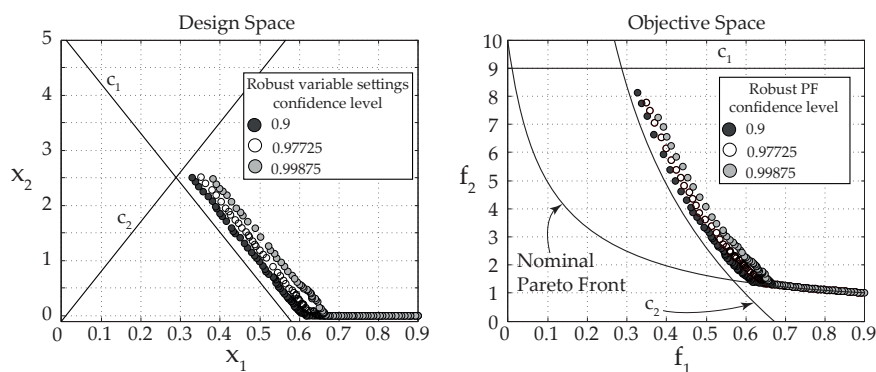

Figure 5.11   Robust optimization results, problem Eq. (5.7).

as in the previous cases, the resulting Pareto fronts are pushed back towards the feasibility region of the design and objective space.

5.4   Atmospheric entry vehicle model

In this section we provide a brief overview of the mathematical model used for the analysis presented in this paper. In particular we describe the parametrization of the capsule and the Thermal Protection System. Capsule-shape generation, aerothermodynamics, flight mechanics and guidance algorithms are adapted from earlier work of Dirkx and Mooij (2011), and Dirkx (2011). For a detailed discussion on these aspects we refer the reader to the original literature.

5.4.1   Capsule parametrization

The capsules that will be studied here belong to the family of the Apollo-like capsules. These are axial-symmetric capsules that can be defined by 5 parameters: nose radius R_N, side radius R_S, rear part half angle θ_C, mid radius R_m, and rear part length L_C, see Figure 5.12. The parametrization consists of four matched analytical geometries, namely a sphere segment, a torus segment, a conical frustum and again a spherical segment at the back. Since the shape is axial-symmetric, the entire surface geometry is defined by the cross-section shown in Figure 5.12. In the discussion regarding the TPS, we will refer to nose and rear part (or conical part) for the capsule. The gray area of Figure 5.12 represents the nose of the capsule, the white area, instead, represents the real part (or conical part).
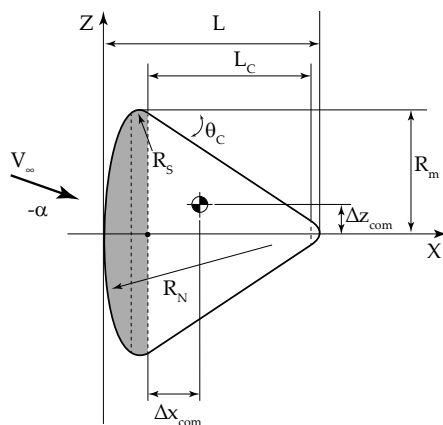
Figure 5.12   Schematic representation of the geometrical parameters of the capsules. The gray area is the nose of the capsule, the white area is the rear part. Adapted from Hirschel and Weiland (2009).

The shift of the center of mass in the vertical and longitudinal direction are additional variables that may be taken into account in the optimization process. However, the shift in the longitudinal direction does not have a significant effect on the static-stability properties and the dynamic behavior of the capsules (Dirkx, 2011). Therefore ΔX_{com} will not be used as a design parameter. The following two constraints must hold for the capsule to have a feasible shape:

$$
\begin{aligned}
R_m &< R_N \\
L_c &< \frac{R_m - R_s\left(1 - \cos\theta_C\right)}{\tan\theta_C}
\end{aligned}
\tag{5.8}
$$

5.4.2   Flight mechanics and capsule GNC

The aerodynamic and gravitational forces are the only external forces acting on the vehicle. The effect of wind is neglected, resulting in an atmosphere that is rotating with the Earth. The atmospheric properties are determined from the 1976 Standard Atmosphere (NOAA/NASA, 1976). There is no thrust force to take into account since the re-entry is considered to be un-powered. Other perturbing forces are neglected, like solar radiation and third-body attraction, since the errors induced by neglecting them are small if compared to errors due to inaccuracies in the aerodynamic coefficients (Dirkx and Mooij, 2011). The gravity field model includes the central gravity and the J2 term. The equations of motion used to propagate the trajectory of the vehicle are expressed in spherical coordinates. Their derivation is presented in detail in separate literature by Mooij (1994) and Dirkx (2011). The state variables that are taken into account are the following:

- Radial position r. This variable denotes the scalar distance from the center of the reference frame, in this case the distance from the center of the Earth.

- Longitude τ. This angle is measured from the Greenwich meridian, positive in East direction.

- Latitude δ. This angle is measured from the equator, positive in North direction.

- Speed V. The scalar velocity is measured with respect to the ground, which in this case equals the airspeed.

- Flight path angle γ. The flight path angle is the angle between the ground velocity vector and the local horizontal plane.

- Heading angle χ. The heading angle is the angle between the North direction in the local horizontal plane and the ground velocity vector projected onto this plane. It is measured positive clockwise.

The capsules that will be generated during the optimization process will not be targeted to a landing site or a Terminal Area Energy Management (TAEM) interface. Instead, entry conditions will be specified, while the end conditions will be kept free. It was decided to focus the attention on the effect of including the TPS in the robust optimization process rather than obtaining a set of optimal capsules with optimal guidance algorithms. Therefore the guidance algorithms considered here are very simple and leave room for a follow-up research. The pitch guidance consists in imposing a certain angle of attack to have trimmed conditions during re-entry. The control system, which is not modeled, will be assumed to be capable of steering the capsule to the desired attitude. The trimmed angle of attack follows from:

$$\alpha_{tr} = \alpha \left. \right|_{C_m=0} \tag{5.9}$$

The angle of attack is therefore chosen such that the pitch moment coefficient equals zero. This will of course influence the aerodynamic forces on the vehicles, but the resultant forse is not taken into account when determining the trim angle of attack. If such an angle of attack cannot be found within certain bounds, the vehicle will be labeled as untrimmable. Since the trajectory propagation is 3DOF, the time-dependent process by which the capsule changes attitude is not analyzed, but is assumed to be instantaneous. Attitude stability of the vehicle is important since instability will make it unlikely for the vehicle to retain its trimmed conditions throughout the flight. For this reason, the following condition will be imposed on the capsule:

$$C_{m_\alpha} \left. \right|_{\alpha=\alpha_{tr}} < 0 \tag{5.10}$$

For the definition of the lateral guidance profile, simplified equations for the time derivative of the velocity and flight path angle are used. A spherical Earth with only a central gravity term is assumed. The centrifugal term due to the Earth's rotation is neglected, but the Coriolis term is included, as it has an appreciable influence in the hypersonic phase. Although its magnitude becomes lower than that of the centrifugal term for low velocities, such velocities will not be encountered for the trajectories generated here, since the Mach number used as lower limit is equal to 3. The bank angle of the capsule is determined by imposing $\dot{\gamma} < 0$. This condition, together with the assumptions just mentioned, will lead to the following relationship for the bank angle, (Dirkx and Mooij, 2011):

$$\cos\sigma = \frac{m}{L} \left(g \left(1 - \frac{V^2}{V_c^2} \right) \cos\gamma - 2\omega_P V \cos\delta \sin\chi \right) \tag{5.11}$$

The bank angle is modulated according to Eq. (5.11), until $\cos\sigma > 0$, at which point the vehicle no longer has sufficient lift to be able to fly at constant flight path angle. When this occurs, the bank angle is set at 0 degrees. This guidance approach will force the vehicle to fly with the maximum possible flight path angle in the central part of the re-entry trajectory, where most of the constraints (*e.g.*, g-load constraint) are more stringent due to large values of the dynamic pressure. The bank angle modulated according to Eq. (5.11) will allow the capsules to fly at a quasi-equilibrium glide condition. This condition will allow, in most cases, to obtain trajectories that do not violate the load constraints (Shen and Lu, 2003).

5.4.3   Aerothermodynamics

The velocity of the capsule with respect to the atmosphere during the entry phase generates aerodynamic forces and moments. In addition to that, the capsule needs to release energy while re-entering to be able to land with accelerations that the structure can withstand. Due to the presence of the atmosphere, this energy is converted in heat, and part of this heat is transferred to the capsule.

The role of the aerothermodynamic analysis is to compute the aerodynamic coefficients of the forces and moments (pressure coefficients) and to compute the heat flux impinging on the capsule. In the model used for the analyses presented in this thesis, these coefficients are computed using local inclination methods. The capsule 3D geometry, generated according to the capsule parametrization previously described, is discretized using a paneled surface mesh. Then, the angle θ of each panel with respect to the free stream flow is the only parameter that determines the pressure coefficient on that panel. Panel methods are simplified when compared, for instance, to Computational Fluid Dynamics (CFD) tools. However, reasonable results can be obtained by using them. In fact, these methods are widely used for conceptual design and optimization, see for instance the publications from Maughmer *et al.* (1990), Theisinger and Braun (2009), Kinney (2004), and Cruz and White (1989). Recently Dirkx and Mooij (2011) demonstrated that an accuracy of better than 15% can be obtained when comparing the results with wind-tunnel data for the Space Shuttle orbiter and Apollo capsule.

In the study presented here, several methods are taken into account to describe the pressure distribution on each of the panels: Newtonian method, Modified Newtonian method, and Tangent Wedge and Tangent Cone methods. For each panel in which the capsule is discretized, a method for estimating the pressure coefficient is automatically selected, depending on the angle of incidence of the impinging flow. The selection method is extensively discussed and tested by Dirkx (2011).

5.4.4   Thermal Protection System

The Thermal Protection System (TPS) of a re-entry vehicle is usually very hard to model. CFD software allows to approach the problem considering a three-dimensional discretization of the vehicle and the surrounding fluid along the trajectory. In general the temperature profiles obtained using this approach are reliable but usually require a heavy computational load for each single vehicle, each single trajectory, and each single attitude. When also non-steady-state thermo-chemical reactions are taken into account, with material decomposition, charring, pyrolysis, and mass injection effects, the problem becomes even more complex.

In literature, examples can be found to reduce the complexity of modeling re-entry problems to two-dimensional and one-dimensional cases. Chen and Milos (2001), for instance, present a two-dimensional approach to the thermal transfer problem. They conclude that such a simplification with respect to the three-dimensional settings could be beneficial for speeding-up the design process of entry vehicles in general, maintaining a certain quality of the solutions. One-dimensional approaches have been investigated more often, especially during the early stages of the Apollo programme. For instance, a complete one-dimensional finite difference model of a charring ablative thermal protection system is provided by Swann (Swann and Pittmann, 1962; Swann *et al.*, 1965).

More recently, general approximated one-dimensional formulations of the heat transfer problem were proposed, demonstrating that at a preliminary design stage these methods can be used for fast analysis while still providing satisfactory results (Martinelli and Braun, 2011; Ferraiuolo and Manca, 2011). The materials' thermo-physical properties are of crucial importance for the determination of the correct temperatures of the TPS layers, and material consumption in the case of ablation or evaporation/sublimation phenomena. These properties
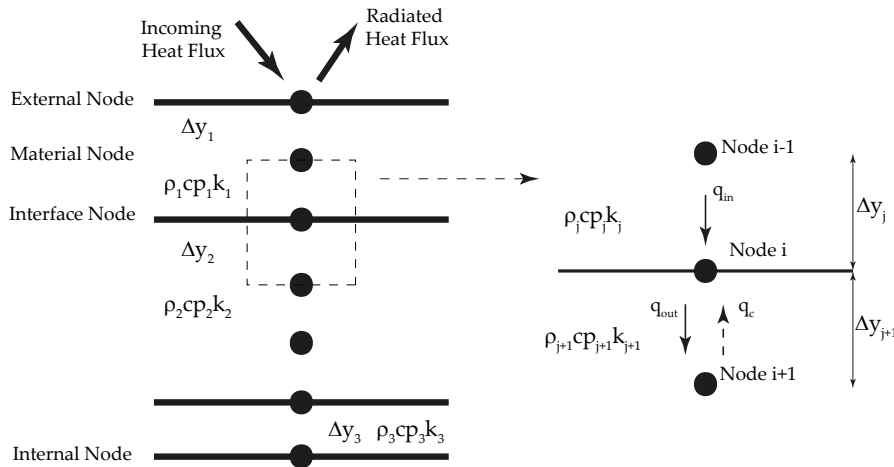
Figure 5.13   Nodes description of a general finite difference model, adapted from Liu and Zhang (2011).

are determined by means of wind-tunnel experiments, but most of the time they are kept classified. Only some data on TPS materials thermo-physical properties can be found (Williams and Curry, 1992). If the thermo-physical properties of the materials are not known, or known with large uncertainties, the analyses could have a limited validity, even in the case of a detailed three-dimensional model of the heat-transfer problem.

The main purpose of the TPS model presented in this study is to combine the transient thermal analysis in a large optimization problem with mission, trajectory, and shape optimization of an entry vehicle. This is done in such a way to substitute the classical heat-flux limits with temperature limits, therefore directly linking the trajectory to the design choices related to the TPS. Due to their reduced execution time, considering that a robust-optimization analysis usually requires a larger number of model evaluations when compared to traditional optimization, a general one-dimensional finite-difference formulation that simulates the stagnation-point conditions and determines the in-depth temperature state along the thickness of the TPS was considered for this analysis. The stagnation point will experience the largest heat-flux and heat-load during the atmospheric entry, thus providing a conservative estimate of the performance of the designed TPS. The material properties and the specific issues related to each material and each solution considered here are described in Appendix D. In general, one-dimensional models provide less accurate results than three-dimensional CFD models. However, this is beneficial for the computational load of the analysis. Robust optimization, as will be shown later in this chapter, may be used also to make conclusions on the assumptions made when choosing a certain simplified model or a more detailed one, and when uncertainty is present also with respect to the material properties. The one-dimensional finite-difference model used in this paper is schematically shown in Figure 5.13. It is adapted from the formulation presented by Liu and Zhang (2011). The TPS is divided into different layers along its axial dimension. Each layer is discretized in isothermal nodes, the distance of which may vary from layer to layer. The thermo-physical properties of the layers are assumed to be isotropic.

The nodes are classified in external, material, interface, and internal nodes. The thermal behavior of the external node can be described as follows:

$$\rho c_p A \frac{\partial T}{\partial t} \partial y = \frac{\partial}{\partial y}\left(kA\frac{\partial T}{\partial y}\right) + \dot{q}_{in} - \dot{q}_{out} \tag{5.12}$$

where ρ indicates the density of the material, c_p the specific heat, A the area, and k the thermal

conductivity. The term $\dot{q}_{in}$ is the incoming heat flux, while the term $\dot{q}_{out}$ represents in general the radiated heat plus the heat absorbed by transpiration gases in case of ablative material. A finite-difference formulation of Eq. (5.12) is the following:

$$T_1^{t+1} = [\dot{q}_{in} - \dot{q}_{out}] \frac{2\Delta t}{\rho_1 c_{p,1} \Delta y_1} + T_1^t \left(1 - \frac{2k_1 \Delta t}{\rho_1 c_{p,1} \Delta y_1^2}\right) + T_2^t \frac{2k_1 \Delta t}{\rho_1 c_{p,1} \Delta y_1^2} \tag{5.13}$$

where the symbol t represents the time, and the subscripts 1 and 2 represent the nodes numbers. The symbols Δt and Δy are the temporal and spatial step-size respectively. In the material nodes only conduction and storage of heat are modeled:

$$\rho c_p A \frac{\partial T}{\partial t} \partial y = \frac{\partial}{\partial y}\left(kA\frac{\partial T}{\partial y}\right) \tag{5.14}$$

Its finite-difference formulation can be expressed as follows:

$$T_i^{t+1} = \left(1 - \frac{2k_j \Delta t}{\rho_j c_{p,j} \Delta y_j^2}\right) T_j^t + \frac{k_j \Delta t}{\rho_j c_{p,j} \Delta y_j} \left(T_{i+1}^t + T_{i-1}^t\right) \tag{5.15}$$

Here, j indicates the layer to which the nodes belong. The model of the interface node takes into account an extra flux, q_c, that is present in the case of pyrolysis phenomena for ablative materials or heat absorption due to water evaporation, for instance:

$$\frac{\partial}{\partial y}\left(-k_j A \frac{\partial T}{\partial y_j}\right) = q_c - \frac{\partial}{\partial y}\left(k_{j+1} A \frac{\partial T}{\partial y_{j+1}}\right) \tag{5.16}$$

The assumption is that the pyrolysis manifests itself on a line rather than on a finite thickness region of the material. For interfaces where no additional heat flux is present the flux q_c is zero. Its finite-difference formulation can be expressed as follows:

$$\begin{aligned}
T_i^{t+1} &= \frac{2\Delta t}{\rho_j c_{p,j} \Delta y_j + \rho_{j+1} c_{p,j+1} \Delta y_{j+1}} \left(\frac{k_{j+1}}{\Delta y_{j+1}} T_{i+1}^t + \frac{k_j}{\Delta y_j} T_{i-1}^t\right) + \\
&+ \left[1 - \frac{2\Delta t}{\rho_j c_{p,j} \Delta y_j + \rho_{j+1} c_{p,j+1} \Delta y_{j+1}} \left(\frac{k_{j+1}}{\Delta y_{j+1}} + \frac{k_j}{\Delta y_j}\right)\right] T_i^t - \\
&- \frac{\dot{q}_c}{A} \frac{2\Delta t}{\rho_j c_{p,j} \Delta y_j + \rho_{j+1} c_{p,j+1} \Delta y_{j+1}}
\end{aligned} \tag{5.17}$$

At the inner structure nodes, similarly to the material nodes only the heat conduction and heat storage is modeled. A finite-difference formulation can be expressed as follows:

$$T_i^{t+1} = \left(1 - \frac{2k_i \Delta t}{\rho_i c_{p,i} \Delta y_i^2}\right) T_i^t + \frac{2k_i \Delta t}{\rho_i c_{p,i} \Delta y_i^2} T_{i-1}^t \tag{5.18}$$

The finite-difference formulation presented here allows for solving the problem with a number of nodes that may vary case by case, and with an explicit forward integration technique. To obtain a faster execution of the thermal transient analysis, the results presented in this chapter were obtained with an implicit formulation of the problem. It was decided to describe here the explicit formulation, because it is more direct and it allows for linking it directly to the involved physical phenomena. Having an implicit formulation means that the state at time t depends on previous states but also on the current one. In that way even when the time step is large, the integration error can be kept relatively low. There are references in literature for solving this type of problems with an implicit formulation (Press *et al.*, 2007). For the analyses presented here, an implicit formulation was used.

For the analysis presented in this thesis six TPS concepts were developed and implemented. A detailed description of the TPS concepts is provided in Appendix D. As a metallic TPS solution we consider *PM2000*, a highly oxidation resistant alloy. It presents an elevated content of Aluminium and Chromium. Therefore, PM2000 is able to form an Al_2O_3 layer when used at high temperatures under oxidizing atmospheres, protecting it from oxidation (Plansee, 2012). The metallic TPS solutions may also present water as cooling mechanism. In case of a direct water cooling TPS solution, when the temperature of the water reaches the boiling limit, the heat from the skin is absorbed by the water for its evaporation leaving the temperature of the skin at a constant value. The assumption is that the water is cooling through nucleate boiling. Therefore we consider the temperature of the innermost skin, which is in contact with the water, being at the same temperature as the water itself. Enhanced-radiation cooling is another TPS configuration that can be selected for the analysis. It is a concept developed at Delft University of Technology (Buursink, 2005). When enhanced-radiation cooling is adopted, the cooling principle is still based on evaporation but the engineering solution is substantially different. Water is contained in a porous material, ZAL-15 for instance (a mixture of Alumina 85 % and Silica 15%) that is detached from the external skin. At the interface the heat is tranferred through radiation, therefore, the equilibrium between the innermost layer of the skin and the ZAL-15 shall be taken into acount in the transient model. Ceramic materials were also considered as TPS engineering solution. In particular, Ultra High Temperature Ceramics (UHTC) were taken into account, using ZrB$_2$-SiC as one possible representative. The last category of TPS solutions considered for the analysis presented foresees the utilization of ablative materials. In particular, phenolyc nylon, carbon phenolic, and phenolic-impregnated carbon ablator (PICA) are considered.

These TPS concepts have been classified in fully re-usable (re-usability index equal to 3), when the nose and rear part of the capsule are made of re-usable materials. Partially re-usable (re-usability index is equal to 2) indicates the situation where only the nose or the rear part of the capsule are made of re-usable materials and the other is made of ablative material. The capsule is considered to be non-reusable (re-usability index is equal to 1) when both the nose and the rear part present ablative solutions. The fully re-usable engineering solutions are metallic uncooled and ZrB$_2$-SiC, for either the nose or rear part, or both. Further, active cooling mechanisms can be selected by the optimizer for nose and cone. In particular, for the nose direct water-cooling is considered, while enhanced-radiation cooling can be applied to the rear part. The assumptions, and the specific challenges related to each of these solutions, and a validation of the mathematical models are discussed in Appendix D.

The re-usability index is also related to the maximum temperature reached during the re-entry of the capsule. After a certain temperature threshold the capsule is considered to be less reusable if the external temperature gets closer to the limit temperature of the material, see Appendix D for the temperature limits of each material.

5.4.5 Initial conditions and trajectory simulation

In this section we present the results obtained using an Apollo-like capsule with Carbon Phenolic TPS for the nose and Phenolic Nylon TPS for the rear part. This is done to show the type of results that will be analysed during the optimization process. Figure 5.14 shows the entry trajectory of the capsule, and the evolution in time of other important parameters using the following initial conditions:

$$h = 120km, \tau = 225.5°, \delta = -23.75°, V = 10.668km/s, \gamma = -6.68°, \chi = 49.6°$$

In Figure 5.14(a) we show the trend of the altitude and the velocity of the capsule in time. The capsule dives into the atmosphere until $t = 100$ seconds where it experiences an inversion of the trend, presenting a small skip. This means that the capsule at that point has a large

lift, due to a combination of angle-of-attack, lift coefficient, velocity, and air density. To force the capsule to continue the re-entry into the Earth's atmosphere, the control system acts on the angle of attack (decreasing it) and on the bank angle (increasing it). This combined effect makes sure that the vertical component of the lift-force is reduced, so that the capsule can continue its descent towards the Earth's surface. The same effect happens on a reduced scale around $t = 350$ seconds. There is again a small tendency of the capsule to skip but the variation of the bank angle performed by the control system, and the fact that the capsule has less energy than before, do not allow the skip to happen. The G-load follows the altitude profile. In correspondence to the two points in the trajectory where the skip is initiated, the G-load has its peaks. The heat-flux profile is reflected on the temperature profile: the nose temperature goes up following the heat-flux. Each line in Figure 5.14(e) represents the temperature trend of one layer of the TPS, for the nose and for the rear part. At a certain epoch of the simulation, around $t = 60$ seconds, the pyrolysis of the Carbon Phenolic material begins. This is evident from the temperature plateau experienced by the third layer of the material, that keeps the temperature constant by changing phase from solid to gas. Finally, in Figure 5.14(f) we present the three-dimensional geometry of the capsule used for the simulation.
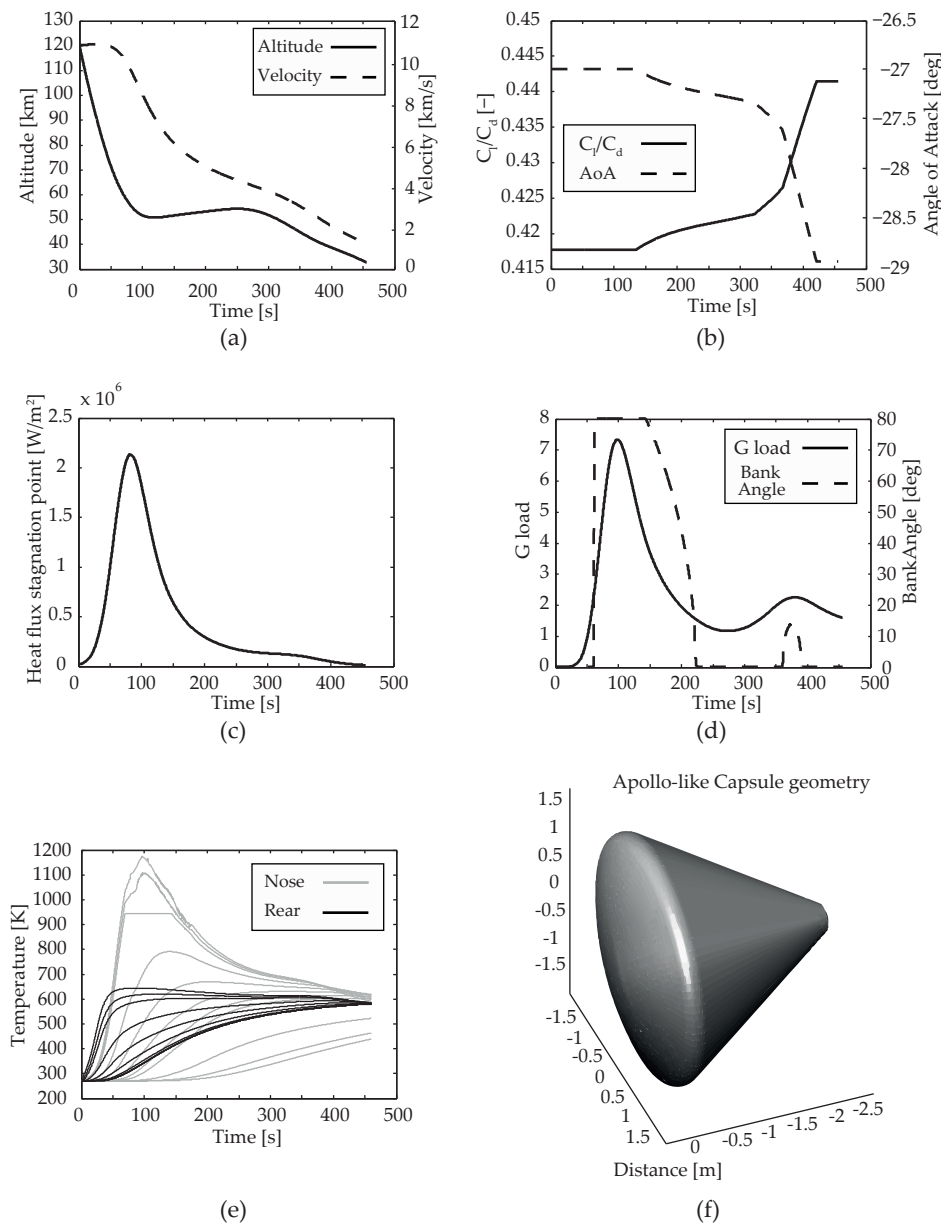
Figure 5.14   Example of the simulated entry of an Apollo-like capsule.

5.5   Unmanned entry capsules servicing the ISS

The main focus of this section is to describe the results obtained from the robust-optimization of unmanned entry capsules used as cargo service to the ISS. The analysis is carried out using the model described in the previous section and relevant literature, with the robust-optimization method introduced in Section 5.3. The objective is to find trade-offs between minimum mass and maximum volume of the capsules, while maximizing their re-usability at the same time.

The analysis is performed by introducing 9 constraints. The external-skin temperature of nose and rear part shall not exceed the limits indicated in Appendix D for each type of material. The TPS innermost-layer temperature shall not exceed 600 K. The capsules shall not exceed a longitudinal load factor of 8 g during the entry phase. The capsules shall be stable and controllable within certain ranges of the angle of attack (Dirkx and Mooij, 2011). Finally, in

Design Variables		Intervals	
		Min	Max
Nose radius	[m]	1	5
Corner radius	[m]	0.001	0.4
Rear part angle[a]	[deg]	−60	30
Capsule radius[a]	[m]	0.5	5
Capsule length	[m]	0.2	3
Z-position com	[m]	0	0.1
Type of TPS nose[b]	[-]	1	6
Thickness Nose	[m]	0.0015	0.1
Type of TPS rear part[c]	[-]	1	6
Thickness rear part	[m]	0.0015	0.1
Initial water-mass nose	[kg]	5	1000
Initial water-mass cone	[kg]	5	1000
Initial flight-path angle	[deg]	−6	−2
Initial bank angle	[deg]	0	90

Table 5.1   Reusable capsules design parameters settings.[a]Constraints of equation 5.8 must be considered. [b]1-ZiB2-SiC, 2-Metallic uncooled (heat sink), 3-Phenolic Nylon, 4-Carbon Phenolic, 5-Metallic direct water cooling, 6-PICA. [c]1-ZiB2-SiC, 2-Metallic uncooled (heat sink), 3-Phenolic Nylon, 4-Carbon Phenolic, 5-Metallic enhanced radiation cooling, 6-PICA.

the presence of ablative material or water as a cooling mechanism, consumed material should be less than what is actually available, and this applies for TPS of both the nose and rear part of the capsule.

The design parameters and the ranges used for the optimization are presented in Table 5.1.

The analysis is performed in two steps. First, in Section 5.5.1 the results of the optimization are discussed. Then, the robust-optimization is done in Section 5.5.2 and results are compared.

5.5.1   Optimization

In Figure 5.15 the Pareto front obtained as a result of the non-robust optimization process is presented. Three branches can be clearly distinguished. These branches correspond to fully-reusable capsules, partially reusable capsules, and non-reusable capsules. The optimizer is pushing the solutions towards minimum mass, maximum re-usability and maximum volume, as expected.

The projections of the Pareto-optima solutions on the mass-volume plane show that the performance of the various capsules with different materials are comparable. This means that from the mass-volume trade-off point of view, entry conditions and design parameter settings exist for which, for instance, a capsule with the nose made of ZrB2-SiC and the rear part with metallic enhanced-cooling behaves as a capsule with fully ablative nose and rear part. The capsules in the Pareto front range from very small and light to heavy and large. Some of the large capsules may simply not be delivered into LEO orbit with current launchers. However, it was decided to set wide limits to the geometric variables to see the evolution of the shapes and the engineering solutions on the efficient and robust-efficient Pareto fronts.

In Figure 5.16 the Pareto front is presented with a focus on the fully-reusable capsules branch of Figure 5.15.

Almost half of the re-usable capsules have both nose and rear part made of ZrB2-SiC material. However, in the same front there are also capsules with a nose made of ZrB2-SiC and
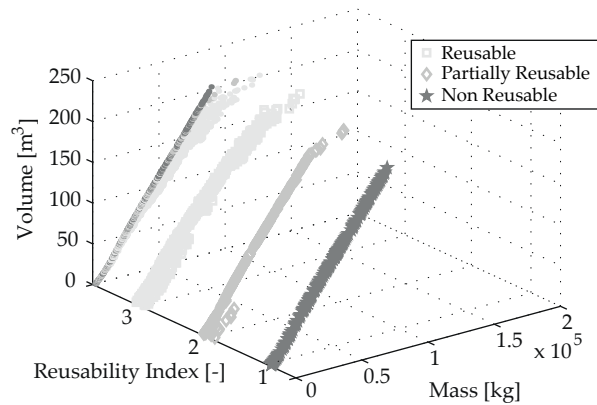
Figure 5.15   Complete Pareto front, robotic capsules servicing the ISS.

the rear part made of metallic material. Amongst these capsules, the metallic-uncooled and the metallic enhanced-cooling concepts are present. Metallic water cooling for the nose is not so common amongst the optimal solutions in the Pareto front. This may be explained from the fact that the tendency is to converge towards large values of nose radius, to maximize the available volume of the capsule. This makes the nose surface very large, thus eventually a larger amount of water is needed for a given heat-load to keep the skin cool. A large nose-radius has two effects. First, it reduces the peak heat-flux on the capsule. At the same time, the total heat load on the nose increases, with the consequent increase of the amount of water required to cool the metallic nose. The increased required water (increased volume and mass) makes the nose metallic-cooled solution less efficient in terms of mass and volume. In some cases, however, there are conditions for which the peak heat flux and the heat load are such that a thin metallic external surface cooled by water is more efficient than a thicker metallic skin protecting the capsule with the heat-sink principle, therefore without water cooling from the inside.

As explained earlier, the re-usability index is determined such that it gets worse as the maximum temperature reached by the material gets closer to the limit temperature. Amongst the solutions on the Pareto front of Figure 5.16, an improvement in the re-usability index causes the capsules to be less efficient in terms of mass and volume. This is what is indicated by the arrow *Increasing Re-usability* on top of the projection of the Pareto-optima solutions on the mass-volume plane. This effect is more evident as the capsules get larger and heavier. Indeed, a large thickness of nose and rear-part skin, and a large amount of water when applicable, cause the maximum temperature reached at the surface to decrease (so getting farther apart from the limit) but also the mass to increase and the available volume to decrease.

In Figure 5.16 some of the capsules obtained during the optimization process are illustrated. An expert eye would already recognize the fact that these capsules get worse in terms of aerodynamics and flight-mechanics characteristics as their size increases. However, more quantitative results are summarized in Tables 5.2 and 5.3. The most relevant trends of the five capsules are also shown in Figures 5.17 and 5.18.

The maximum L/D ratio shows a general decreasing trend going from small to larger capsules. This means that the behavior of the capsules gets closer to that of a ballistic entry as their size increases. The larger capsules look much like the Russian *Foton* capsules. For the few capsules with metallic nose, in general the g-load is taken to the limit by the optimizer. This is obtained by having steep trajectories that in turn cause the peak heat flux to increase and heat load to decrease. The main driver for the amount of water needed to cool the capsule nose is the heat load. Therefore a low heat load will also have an effect on reducing the amount of water to be stored on board and thus reducing the mass of the capsule.
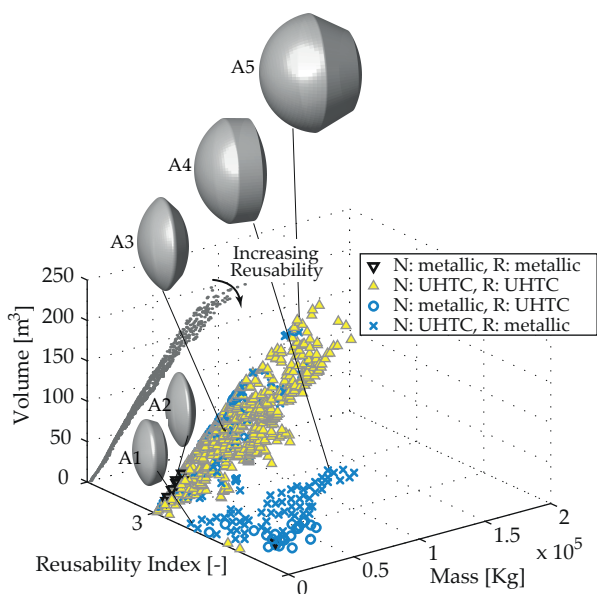
Figure 5.16   Pareto front with fully re-usable capsules, robotic capsules servicing the ISS.

		A1	A2	A3	A4	A5
Nose radius	[m]	3.64	4.99	4.89	5	5
Corner radius	[m]	0.34	0.25	0.33	0.4	0.3
Rear part angle	[rad]	−0.34	−1.05	−0.79	−0.06	−0.47
Capsule radius	[m]	2.27	3.07	3.88	4.35	4.95
Capsule length	[m]	0.6	0.2	0.97	1.95	3
Z-position CoM	[m]	0.1	0.065	0.07	0.09	0.08
Type of TPS nose	[-]	ZrB2-SiC	Water cooled	ZrB2-SiC	ZrB2-SiC	ZrB2-SiC
Thickness Nose	[m]	0.0015	0.0045	0.0015	0.002	0.011
Type of TPS rear part	[-]	Metallic Uncooled	Enhanced radiation cooling	Enhanced radiation cooling	Enhanced radiation cooling	ZrB2-SiC
Thickness rear part	[m]	0.0087	0.0015	0.0025	0.0015	0.0015
Initial mass water nose	[kg]	-	730	-	-	-
Initial mass water cone	[kg]	-	19.5	175	200	-
Initial flight-path angle	[deg]	−2.75	−3.95	−2.89	−2	−2.2
Initial bank angle	[deg]	84.9	60.73	32.14	20	40

Table 5.2   Reusable capsules design parameters settings.

The range that the capsules fly in the atmosphere, from the initial 120 km altitude until Mach 3 is reached, is proportional to the L/D ratio and the flight-path angle. The value of Mach 3 is selected because below this Mach number, the aerodynamic coefficients which are calculated can no longer be assumed to be valid (Dirkx and Mooij, 2011).

As expected, the capsules show an increased flight range as the trajectory gets shallower (flight-path angle closer to zero) and the L/D ratio increases. The range also gives an indication for the heat load that the capsules experience. The larger the range, the larger the time the capsules fly in the atmosphere. This results, in most of the cases, in a larger heat load on the TPS, see for instance Figure 5.17(c).

In the case of capsule A2, even if the heat load is the lowest (the heat load is determined as the total area under the heat-flux curve) the amount of water needed to keep the nose relatively