Conflict Resolution in a Decentralized Air Traffic Concept of Operation

A Thesis
Presented to
The Academic Faculty

By

Antoine Genton

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
School of Aerospace Engineering

Georgia Institute of Technology
May 2015

Conflict Resolution in a Decentralized Air Traffic Concept of Operation

Approved by:

Dr. Amy Pritchett, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Eric Feron
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Jeff Shamma
School of Electrical and Computer Engineering
Georgia Institute of Technology

Date Approved: 01/06/2015

AKNOWLEDGEMENTS

I would like to express my gratitude to my advisor Dr. Amy Pritchett for her useful guidance through the learning process of this master thesis. Furthermore, I would like to thank Dr. Jeff Shamma for introducing me to some fundamental aspects of game theory, and Dr. Eric Feron for his constant support throughout this journey.

I also would like to thank my friends who helped me a lot, and were there during both good and bad days, suggesting me brilliant ideas and sometimes enjoying all-nighters with me. Adrien, Anil, Aparna, Arvind, Benjamin, Can, David, Frederic, Gabriel, Guliz, Jeremy, Raunak, Sebastien, and many other unlisted friends, those two years in your company were wonderful and I want to tell you that I would not have succeeded without you.

Finally, I would like to thank my mother, father, grandparents and loved ones who were always there for me despite the distance. I would like to thank them for having never forgotten their ‘petit américain’. I will be grateful forever for your love.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

The current air traffic concept of operations relies on a centralized process in which ground controllers are responsible for determining conflict-free trajectories. However, with new technologies such as ADS-B and GPS aircraft could directly interact together to resolve their own conflicts in a decentralized manner. The challenge is to guarantee aircraft separation while converging to reasonably fair resolutions for all aircraft. The difficulty is that aircraft have only limited information about how other aircraft evaluate the cost of conflict resolutions.

Thus, this thesis proposes to frame decentralized conflict resolution using game theory. A collaborative decentralized conflict resolution is developed as a sequential bargaining process between the different aircraft. At each step, aircraft propose personal trajectories to the other aircraft, corresponding to trajectories they would be ready to fly. Then they compute response trajectories, corresponding to trajectories they would have to fly to avoid the conflict if the other aircraft flew its personal trajectories. If any response trajectories are cheaper than the offered personal trajectories, an agreement is reached; otherwise, the aircraft need to compromise by offering more expensive personal trajectories at the next step.

Several pairwise conflict experiments, corresponding to different conflict geometries, were conducted to explore different ways of handling performance constraints and different ways of searching trajectories in the resolution space. Then, the bargaining process was demonstrated in a large scale simulation with more than a thousand aircraft flying over the Indianapolis Center, incurring more than five hundred conflicts. The traffic

sets were taken from real ETMS data over five hours, to represent ‘real’ conditions. 93% of the conflicts were successfully solved by the bargaining process.

CHAPTER 1: INTRODUCTION

1.1 Problem Statement

The current air traffic concept of operations relies on a centralized process in which ground controllers are responsible for determining conflict-free trajectories. However, the emergence of new technologies such as ADS-B and GPS affords more autonomous air traffic systems, where aircraft could directly interact together to resolve their own conflicts in a decentralized manner.

The real challenge here is to guarantee separation between aircraft while distributing the cost of resolutions fairly between the aircraft. However, the difficulty with decentralized conflict resolution is that aircraft have only limited information about their surrounding traffic, particularly with respect to how the other aircraft evaluate the cost of conflict resolutions. Several researchers have proposed decentralized conflict resolution algorithms using different approaches, such as prioritizing the conflicting agents [1], formulating and solving the conflicts as an optimal control problem [2], using AI techniques such as genetic algorithms [3], or a capacitance model representing aircraft as charged particles that repel each other [4].

This thesis proposes to frame decentralized conflict detection and resolution using game theory. A collaborative decentralized conflict resolution is developed as a sequential bargaining process between the different aircraft, i.e., the players. The goal of each aircraft is to minimize the cost associated with the trajectories that resolve the conflict. However, each aircraft doesn't exactly know the payoff, i.e., the cost function of the other involved

aircraft. This is actually the way it is in the real world: airlines prefer not to disclose their cost indices, by which they weigh the costs of delay and fuel burn.

The bargaining problem [5] is a game theory concept that aims at finding an equilibrium within a cooperative game. It represents situations in which there is a conflict of interest, since players want to follow their own objectives, but could still conclude with a mutually beneficial agreement. It is used when cooperation is required for a Pareto-efficient solution, i.e., one in which it is impossible to make anyone individual better off without making at least one other worse off.

Complicating the problem, performance constraints of the aircraft can limit the set of feasible conflict resolutions, and resolutions may be found in multiple dimensions. These two issues are framed in this thesis by building a bargaining process with two main design variables that explore methods of handling performance constraints and the multiple dimensions of the resolution space.

1.2 Objectives

- Develop a sequential bargaining process for a collaborative decentralized conflict resolution where both aircraft simultaneously propose resolutions until an agreement is reached, while each aircraft is able to protect its private information and account for performance constraints.
- Run a large scale experiment to assess the efficacy of this process using the multi-agent simulation engine WMC.

1.3 Overview of the Thesis

CHAPTER 2 presents the background of the problem of conflict resolution addressed in this thesis. The main notions involved in solving conflicts are detailed and the different inherent issues that come with it are presented. CHAPTER 3 describes the proposed approach and theoretical results used to build the collaborative decentralized conflict resolution algorithm. CHAPTER 4 gives an overview of the different computational tools and models used to build the simulation engine. CHAPTER 5 presents the results provided by the pairwise conflict experiments. CHAPTER 6 shows the output and conclusions drawn from the large scale experiment. CHAPTER 7 summarizes the contributions of this thesis work and gives some insights into what could be achieved next.

CHAPTER 2: BACKGROUND

This chapter defines the notion of conflict in an air traffic concept of operations. It characterizes decentralized conflict resolution processes and makes a clear distinction between decentralized processes and distributed algorithms. It also presents the main categories of decentralized processes and identifies relevant concepts in game theory.

2.1 Conflict Detection and Resolution

A conflict between two aircraft occurs when aircraft are closer than a given set of horizontal and vertical distances. Current en-route air traffic control rules often consider these distances to be five nautical miles horizontally and one thousand feet vertically [6] for aircraft operating in en-route airspace above 18,000 feet altitude.

Conflict detection consists in detecting aircraft trajectories in a given time frame that will lead to conflicts. Conflict resolution consists in performing changes to the flight plans of the aircraft so that conflict-free trajectories are ensured for a given period of time. For the conflict resolution to be effective, this period of time, the look-ahead time, must be greater than the time frame in which conflicts are detected. This guarantees that the conflict resolutions will clear the list of detected conflicts at a given point in time.

2.2 Resolution and Cost

Conflict resolution provides, in real-time, modified trajectories that maintain safety while minimizing cost as calculated for each aircraft according to its cost index. However aircraft don't generally share all the required information to optimize trajectories. Cost indices, for example, are kept secret by the airlines as part of their commercial strategies, even though they are central in the computation of the cost of a conflict resolution.

Further difficulties are inherently encountered in the conflict resolution problem due to the nonlinearities of the space of potential avoidance trajectories due to the performance limits of the aircraft. The common performance limits of aircraft are the range of speeds and altitudes in which they could operate, which are defined by the flight envelope of each aircraft. For example, the aircraft's maximum level cruise speed generally varies with aircraft weight, which decreases significantly through time as fuel is burnt. Of note, the performance limits of one aircraft at any point in time are also not known by neighboring aircraft.

Conflict solutions can lead to new downstream conflicts imposing their own control cost. Forecasting and controlling this future control cost, which could be represented by the number of maneuvers of future conflicts resolution trajectories and their magnitude, represents another challenge posed by the conflict resolution problem.

2.3 Centralized versus Decentralized Air Traffic Management Systems

In a centralized air traffic system, one agent, usually a ground controller, has the authority to determine the conflict resolution trajectories. This is the strategy adopted in the current air traffic system in which communications between aircraft are limited and not used actively by aircraft to adapt their trajectories. Ground controllers guarantee separation of the aircraft and establish smooth operations within the air traffic system by issuing clearances. Figure 1 shows the central role they play: when a conflict situation is detected between aircraft A & B, two distinct communication processes are triggered between each aircraft and the ground controller. Usually, aircraft just comply with the clearances given by the ground controllers and their role in the communication process is simply to read them back. The main weakness of such a system is that it relies on a decision maker who may have limited information about how each aircraft evaluates cost.

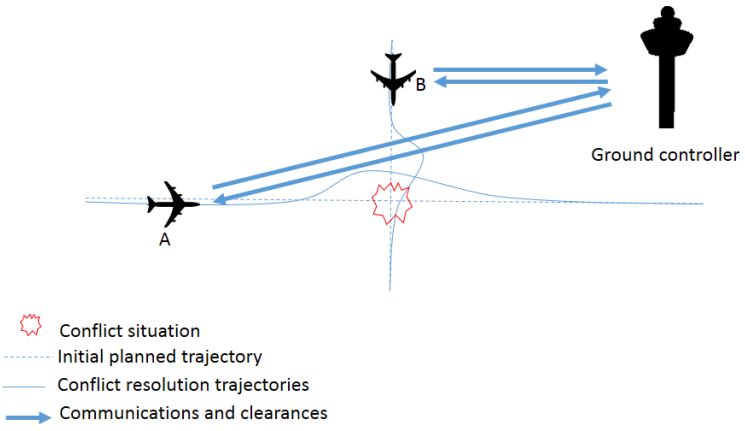

FIGURE 1: CENTRALIZED AIR TRAFFIC SYSTEM

In a decentralized air traffic system, aircraft have more autonomy and take part in the decision making process to, for example, decide their own strategies and resolutions

trajectories in case of conflict. These systems are difficult to implement but have the advantage of allowing the main stakeholders to actually be part of the decisions that concern themselves. Figure 2 shows a decentralized conflict resolution: aircraft A and B are here responsible for coming up with resolution trajectories, and deal with their own problem together.

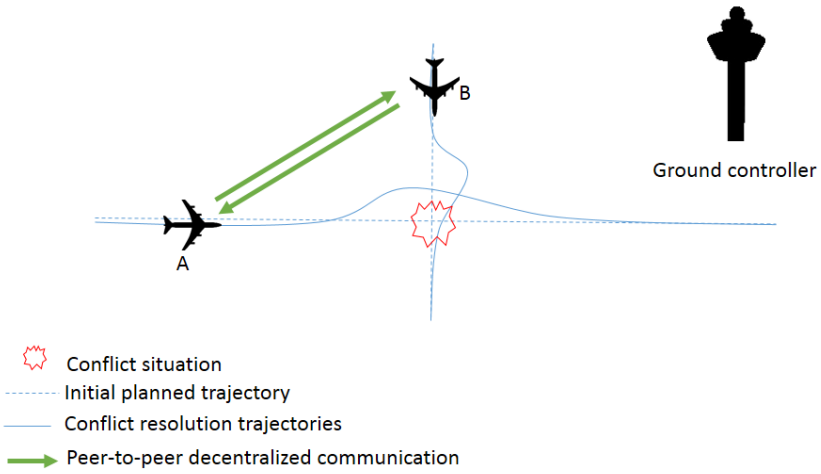

FIGURE 2: DECENTRALIZED AIR TRAFFIC SYSTEM

However, one difficulty to overcome when designing a decentralized system is the fact that each aircraft wants to follow its own objectives, which can be antagonistic to those of the other aircraft. Aircraft have therefore to make compromises to reach an agreement amongst themselves in the decision making process.

2.4 Distinction between Decentralized and Distributed

The difference between centralized and distributed is not important when one centralized agent with perfect knowledge of the whole system can get the same outcome

as distributed agents who share the same information and value cost the same way. On the other hand, truly decentralized systems are characterized by the fact that their outcomes are functions of information and cost functions known only by each individual agent, which means that centralized agents should not *a priori* be able to get the same outcome as decentralized agents.

It is interesting to note the potential value of decentralized operations in air traffic management. As mentioned above, cost indices are, for example, considered as strategic values by the airlines and are kept proprietary. Performance limits are also not known centrally. The interest of a decentralized conflict resolution system is then to take advantage of the private information that airlines are not willing to share to compute avoidance trajectories that are individually considered to be closer to optimal than trajectories computed by a centralized agent.

2.5 Main Categories of Decentralized Conflict Resolution Approaches

Different approaches of decentralized conflict resolution have already been discussed in the literature. Two main categories are those based on optimal control, and those based on capacitance and particles analogies.

Approaches based on optimal control formalize the conflict resolution problem as trying to minimize a global cost function [2]. The agent responsible for the actual computation of the avoidance trajectories would usually be an external agent, such as a ground controller, which means that this kind of systems is not truly decentralized.

The categories of decentralized conflict resolution algorithms based on capacitance and particles analogies consider the aircraft as charged particles with the same sign evolving towards their destinations, which are locations represented by fixed particles of the opposite sign [4]. Thus, aircraft are repelled by other aircraft and attracted by their destinations. This approach allows the use of well-established static charges and capacitance theory but has several issues. The first is the difficulty of ensuring that aircraft won't get pushed into the protected airspace of other aircraft. The second is preventing resolution trajectories outside of the flight envelopes of the aircraft.

In [7], Sislak & al. describe an iterative peer-to-peer collision avoidance algorithm, called IPPCA, very similar to what is developed in this thesis. It is based on high-level flight plan variations using evasion maneuvers. Its default behavior is to minimize the sum of the costs of the aircraft, therefore maximizing their social welfare, but it can also be configured to provide solutions for self-interested airplanes where airplanes optimize their own cost instead of the overall cost. Personal trajectories are successively offered by the aircraft without considering the other aircraft proposed personal trajectories, and the algorithm stops as soon as a pair of trajectories is found to be conflict-free. However, there is no guarantee that, for a fixed personal trajectory in the final conflict-free pair, the other trajectory is the cheapest that could have been computed by the other aircraft.

In [8], Wollkind & al. present an application to air traffic conflict resolution of the Monotonic Concession Protocol (MCP) developed by Zlotkin and Rosenschein [9]. This approach is also very similar to the one developed in this thesis. The MCP captures the incremental bargaining process that takes place between the two aircraft that simultaneously make proposals and counter proposals of progressively less value to

themselves. In their process, aircraft also pick the trajectories of the other aircraft in their successive proposals. These trajectories cannot therefore be guaranteed to minimize the other aircraft cost function, since it is not known.

2.6 Game Theory and Bargaining

Game theory examines strategic decision making in situations involving different agents. The objective of game theory is to model these situations to determine the optimal strategies for each agent, to predict equilibria of their collective games, and to examine whether decisions converge.

In this work, we are focusing on the development of a collaborative conflict resolution algorithm in a pairwise conflict. This is framed as a two-player game where both aircraft are trying to find optimal strategies to clear a conflict. The logical strategy that each aircraft would prefer would be to not maneuver, leaving the other aircraft to maneuver alone. This shows that the problem to solve is essentially a bargaining problem in which both aircraft have to compromise to reach an agreement on a conflict's resolution [5] [10] [11] [12].

The utility function of a player represents his preference among a set of different anticipations. An anticipation of an individual is considered here as a state of expectation which may involve the certainty of deterministic contingencies and various probabilities of stochastic contingencies [5]. Nash developed a few assumptions to develop the utility theory of an individual [5]:

- An individual offered two possible anticipations can decide which is preferable or that they are equally desirable.
- The ordering thus produced is transitive; i.e., if I is better than J and J is better than K then I is better than K.
- Any probability combination of equally desirable states is just as desirable as either.
- If I is better than J and J is better than K, then there is a probability combination of I and K which is just as desirable as J. This corresponds to an assumption of continuity.
- If $0 \leq p \leq 1$ and I and J are equally desirable, then $pI + (1-p)K$ and $pJ + (1-p)K$ are equally desirable, and I may be substituted for J in any desirability ordering relationship satisfied by J.

These assumptions are sufficient to show the existence of a satisfactory utility function that assigns a real number to each anticipation of an individual. This function is not unique: if u is such a function then so also is $ku + c$, where $k > 0$. Letting capital letters represent anticipations, such a utility function will verify the following properties:

- $u(A) > u(B)$ is equivalent to A is more desirable than B.
- If $0 \leq p \leq 1$ then $u(pA + (1-p)B) = pu(A) + (1-p)u(B)$.

  This corresponds to the important linearity property of a utility function.

The negative of a cost function f can be seen as a utility function u. For two aircraft denoted A and B, we can define $u_A(t_A) = -f_A(t_A)$ and $u_B(t_B) = -f_B(t_B)$. Hence it is

equivalent to reason with costs or utilities. The cost viewpoint is used in conflict resolutions.

Formally, a two-person bargaining problem consists of the set of all possible anticipations Σ, the feasible set of all attainable costs $F \subset \mathbb{R}^2$, and a disagreement point $d \in F$ which corresponds to the pair of costs when no agreement is reached. Solving a bargaining problem means finding an agreement $s \in F$ viewed as better than d for both players according to their personal preference, and considered as optimal with respect to certain criteria. In the case of conflict resolution, optimal is defined as lowest cost.

Assuming that players A and B have different preferences, represented by cost functions f_a and f_b, the set of all attainable costs could be written as:

$$F = \{(v_a, v_b) \in \mathbb{R}^2 | v_a = f_a(X), v_b = f_b(X), X \in \Sigma\} \quad (1)$$

The pair (F, d) represents a bargaining problem. Designing the set of all bargaining problems by $\mathcal{B}$, a bargaining solution is a function:

$$\varphi: \begin{cases} \mathcal{B} \to F \\ (F, d) \mapsto \varphi(F, d) = \big(\varphi_a(F, d), \varphi_b(F, d)\big) = s \end{cases} \quad (2)$$

Two different but complementary approaches have been developed so far to construct such bargaining functions. The first approach is strategic: the bargaining process is designed first and the properties of its outcome are studied, i.e., the bargaining solution is inherently defined by the chosen bargaining process. The second approach is axiomatic:

bargaining solutions are characterized by several axioms supposed to be reasonable and every function that satisfies these axioms is a bargaining function candidate.

This second approach led to a major result proved by Nash in 1950 [5]. He showed that, in the case where the set of all possible costs is convex and compact, there exists a unique bargaining solution that satisfies the following axioms:

1. **Pareto Optimality**

   As illustrated in Figure 3, a bargaining solution $\varphi(F, d)$ is Pareto optimal if, within a set of feasible solutions Σ, there is no other solution x that simultaneously provides lower costs than $\varphi(F, d)$ to both players A and B, i.e. for any $x \in \Sigma$, if $f_a(x) \leq \varphi_a(F, d)$ then $f_b(x) \geq \varphi_b(F, d)$, or if $f_b(x) \leq \varphi_b(F, d)$ then $f_a(x) \geq \varphi_a(F, d)$.

   This property simply exposes the fact that with a Pareto non-optimal outcome better solutions exist for either or both players. This axiom reduces the space of potential bargaining solutions, which should all be located on the Pareto frontier. However, it does not provide insights on how to order them.

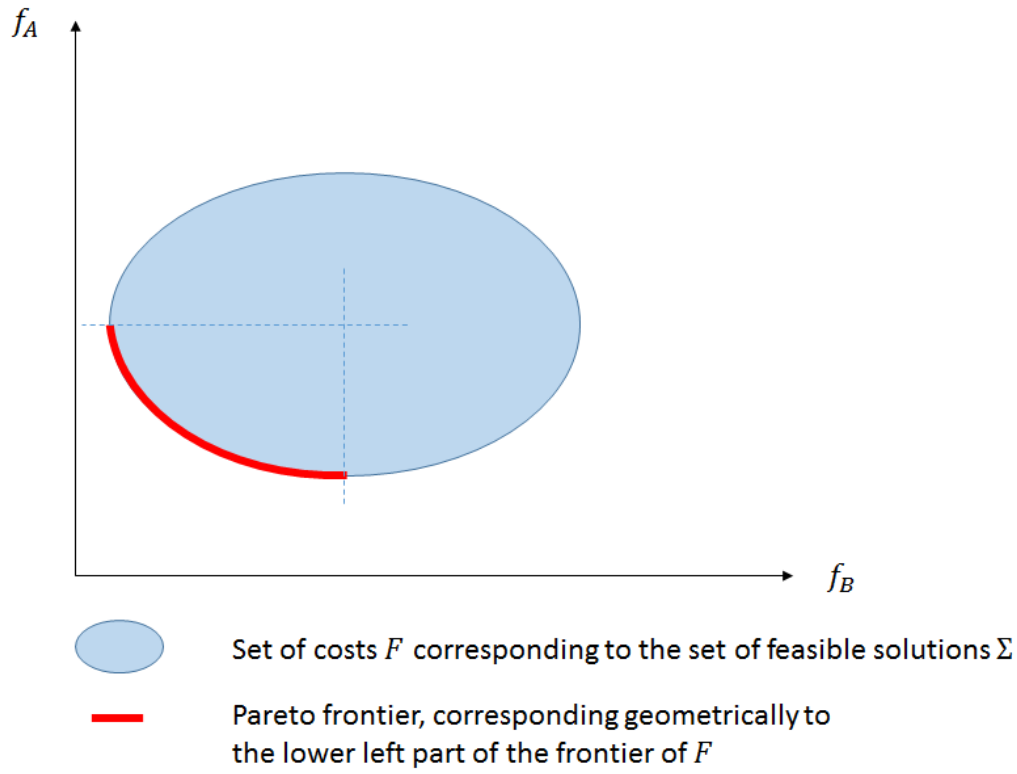

FIGURE 3: PARETO FRONTIER ILLUSTRATION

2. Symmetry

Let the function $T: \mathbb{R}^2 \to \mathbb{R}^2$ be defined by $T((x, y)) = (y, x)$. A bargaining solution $\varphi(F, d)$ is symmetrical if, for every bargaining problem $(F, d) \in \mathcal{B}$, $\varphi\big(T(F), T(d)\big) = T(\varphi(F, d))$. This basically means that the solution $\varphi(F, d)$ should not give preference to either player, as represented on Figure 4.

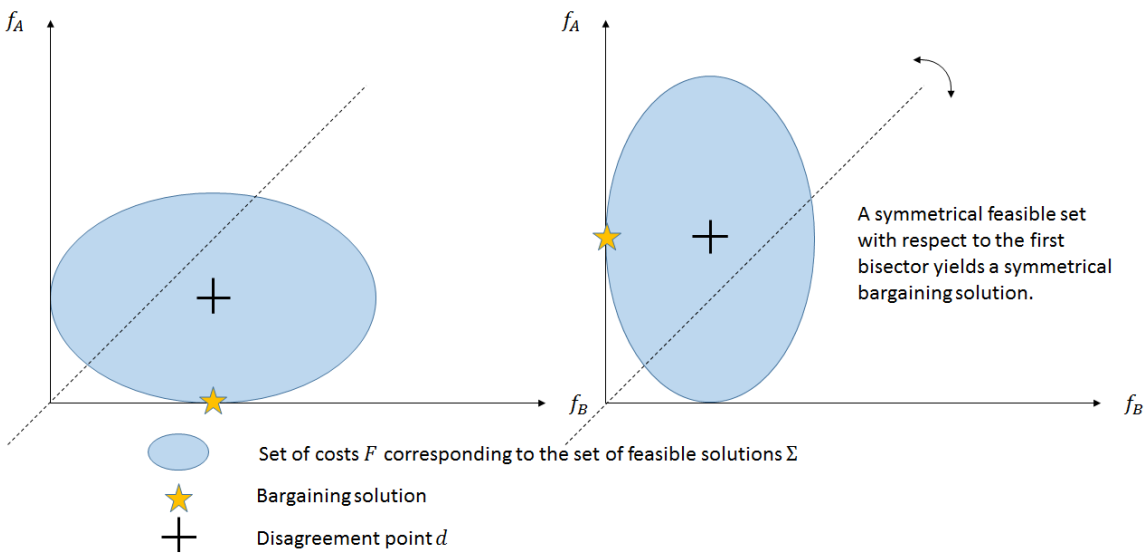

FIGURE 4: ILLUSTRATION OF THE AXIOM OF SYMMETRY

3. Invariance of Cost with Respect to Affine Transformations

This property states that an affine transformation of the cost functions maintaining the order over the preferences should not modify the bargaining solution. This is due to the non-uniqueness of cost functions, as explained earlier. If we consider an affine transformation of cost $G = (G_1, G_2)\colon \mathbb{R}^2 \to \mathbb{R}^2$ such that $G\big((x,y)\big) = (G_1(x), G_2(y))$, with $G_i(x) = c_i x + d_i$, $c_i > 0$, then for every bargaining problem $(F,d) \in \mathcal{B}$, $\varphi\big(G(F), G(d)\big) = G(\varphi(F,d))$.

4. Independence of Irrelevant Alternatives

A bargaining solution $\varphi(F,d)$ verifies this property if, for every (F,d) and (F',d) such that $F' \subset F$, with $\varphi(F,d) \in F'$, i.e. F' contains the solution of (F,d), then $\varphi(F',d) = \varphi(F,d)$, i.e. the solution of (F',d) is the same as the solution of (F,d). The corollary is that removing a subset of F not containing

the bargaining solution of (F, d) and therefore considered as irrelevant should leave the bargaining solution unchanged.

Nash showed that the unique bargaining solution satisfying the 4 properties presented above is defined by:

$$\varphi(F,d) = \max_{\substack{(v_a,v_b)\in F \\ v_a \leq d_a \\ v_b \leq d_b}} (d_a - v_a)(d_b - v_b) \tag{3}$$

Geometrically, the Nash bargaining solution corresponds to the point of the Pareto frontier that maximizes the area of the rectangle whose sides are parallel to the axes of the costs and opposite vertices are the considered solution and the disagreement point. This is depicted in Figure 5.

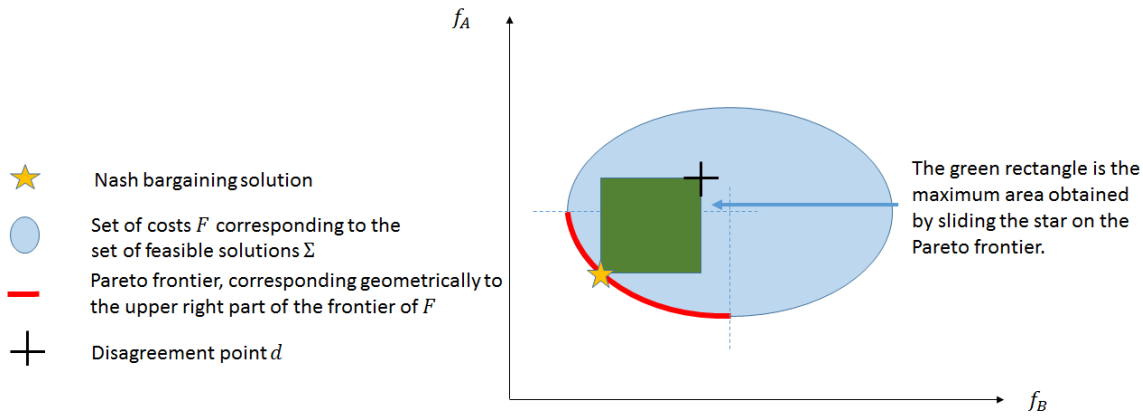

FIGURE 5: NASH BARGAINING SOLUTION

If the three first axioms have been well received by the scientific community, the fourth one has been discussed by several researchers such as Ehud Kalai and Meir Smorodinsky [13] who replaced it with another axiom, the axiom of monotonicity, which

states that, if (F, d) and (F', d) are two bargaining problems such that $F' \subset F$, and the minimum costs for both players are the same in both problems, then $\varphi(F', d) \geq \varphi(F, d)$. This new axiom, associated with the three previous ones, leads to another unique solution different from the Nash solution: the Kalai-Smorodinsky solution [13]. Geometrically, this solution corresponds to the point of the Pareto frontier located on the diagonal of the rectangle with opposite corners that are the disagreement point and the point whose x and y coordinates are respectively the minimum costs of each of players B and A. This solution is depicted in Figure 6.

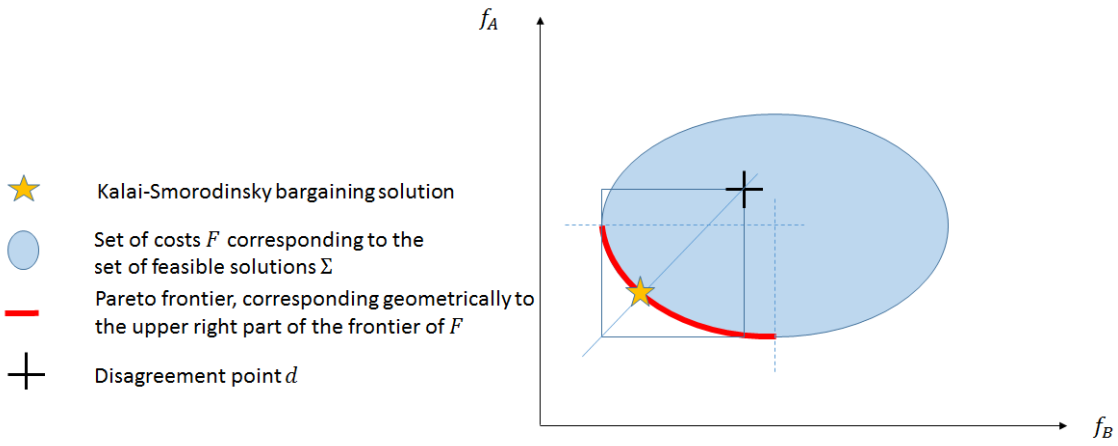

FIGURE 6: KALAI-SMORODINSKY BARGAINING SOLUTION

The difficulty is to provide a bargaining process that actually converges to the ideal solution through alternative offers. This shows the inherent link between the axiomatic and strategic approaches of bargaining. Some works have been developed in this discipline such as the Rubinstein bargaining model. Ariel Rubinstein provided a solution to a class of bargaining games that feature alternating offers through an infinite time horizon. One fundamental idea behind his solution is the fact that delays can themselves incur costs such

that payoffs converge to 0 with time and it is always better for the two players to reach an agreement in finite time [14].

CHAPTER 3: FORMAL DEFINITION OF A BARGAINING PROCESS FOR CONFLICT RESOLUTION

Consider a pairwise conflict of two aircraft A & B. If these aircraft remain on their initial respective trajectories, they will conflict. The goal is to define a process by which aircraft will find new trajectories that resolve the conflict and minimize cost. The aircraft also want these trajectories to be fair and they want to keep their confidential information private. As discussed in the previous chapter, being "fair" is hard to define. Here, the fairness of a conflict resolution is provided by the mutual agreement reached at the end of the bargaining process. The aircraft are never forced to accept a resolution, but act rationally to minimize their personal cost. This corresponds to a solution to bargaining given by a strategic approach, as opposed to an axiomatic approach. However, the fairness of the bargaining process can be assessed by examining whether the axioms linked to fairness listed in 2.6 are satisfied. In particular, an interesting metrics of fairness is given by the axiom of symmetry: the costs of the negotiated trajectories for the two aircraft should be relatively close when the conflict situation is symmetrical.

To interact properly, both aircraft in a conflict must apply the same form of a cost function to assess resolution trajectories. Their individual objectives are defined as the specific values of the parameters within the function, which can be held private. This allows for different evaluations of proposed conflict resolutions by each aircraft and represents a truly decentralized process whose outcomes couldn't have been produced by a centralized agent.

Formally denoting the space of private information by X, the space of public information by Y, and the space of feasible trajectories by T that actually clear the conflict situation and respect performance constraints, the evaluation function is:

$$f: X \times Y \times T \to \mathbb{R} \tag{4}$$

Assuming aircraft A & B are characterized by knowledge of their local sets of information $(x_A, y_A) \in X \times Y$ and $(x_B, y_B) \in X \times Y$, personal evaluation functions could be defined by $f_A: T \to \mathbb{R}$ and $f_B: T \to \mathbb{R}$ such that:

$$f_A(\cdot) = f(x_A, y_A, \cdot) \tag{5}$$

$$f_B(\cdot) = f(x_B, y_B, \cdot) \tag{6}$$

These functions represent the cost of personal resolution trajectories viewed by each aircraft, i.e., each aircraft measures its own trajectory cost. Since A does not know x_B and B does not know x_A, A has no way to actually compute the cost of trajectories as viewed by B and vice versa.

Starting from a pair of conflicting trajectories (c_A, c_B), a reasonable bargaining process must find a pair of resolution trajectories (t_A, t_B) mutually agreed upon by aircraft A and B. The bargaining solution will then be represented in the cost plan by the point of coordinates $\big(f_A(t_A), f_B(t_B)\big)$.

The disagreement point d of this bargaining problem will correspond to the situation in which both aircraft are maneuvering as if the other one was not wanting to

make any maneuvers, i.e., it would correspond to a safe but suboptimal situation. The feasibility set F corresponds in this problem to the set of all pairs of finite costs corresponding to pairs of trajectories that the aircraft can actually fly within their performance limits, that clear the conflict, and that are Pareto optimal.

This chapter defines a bargaining process in which both aircraft will agree upon a conflict resolution belonging to the feasibility set F. In this sequential bargaining process, players will be forced at each iteration to always offer alternative solutions they personally consider to be worse for themselves than their previous ones. Doing this will force the players to reach an agreement after enough negotiations. Delays may be costly but the bargaining process is assumed to happen so fast that delays won't be of any significant duration.

This process is designed so that aircraft can minimize their personal cost functions without getting information on the other aircraft's cost function. In addition, different methods for accounting for the performance limits of the aircraft and the multi-dimensionality of the resolution space are created by introducing two design variables.

3.1 Structure of the Bargaining Process

The goal is to reach a resolution mutually agreed by both aircraft. The difficulty is that both players have *a priori* no interest in deviating from the simple zero cost solution in which they don't maneuver and their opponent assumes all the cost. To avoid this, the proposed bargaining process will force the aircraft to deviate from these trivial but not-fair

solutions. A good-faith compliance with this process is assumed, enforced by regulation and perhaps built into the on-board automation negotiating for each aircraft.

The structure of this bargaining process is sequential, as represented in the flowchart in

Figure 7. Hence, each step of the bargaining process is divided into 3 sub-steps:

1. Each aircraft computes a set of "personal trajectories" in each dimension that it would agree to fly with a given cost requirement. These cost requirements are initially set to 0 and are increased with each iteration of the process to require compromises.
2. Each aircraft communicates the trajectories computed in 1 to the other aircraft.
3. Each aircraft computes "response trajectories" as the minimum cost resolution to the conflict if the other aircraft flew its personal trajectories.
4. Each aircraft compares the cost of the set of response trajectories it determined with the cost of its own personal trajectories computed in 1. If any of the response trajectories are cheaper than the personal trajectories, then an agreement can be reached; else no agreement is reached and both aircraft go one step further in the bargaining process and loop back to the first sub-step. Both aircraft have to make compromises: by increasing their cost requirements, they now have to offer personal trajectories that are slightly more expensive.

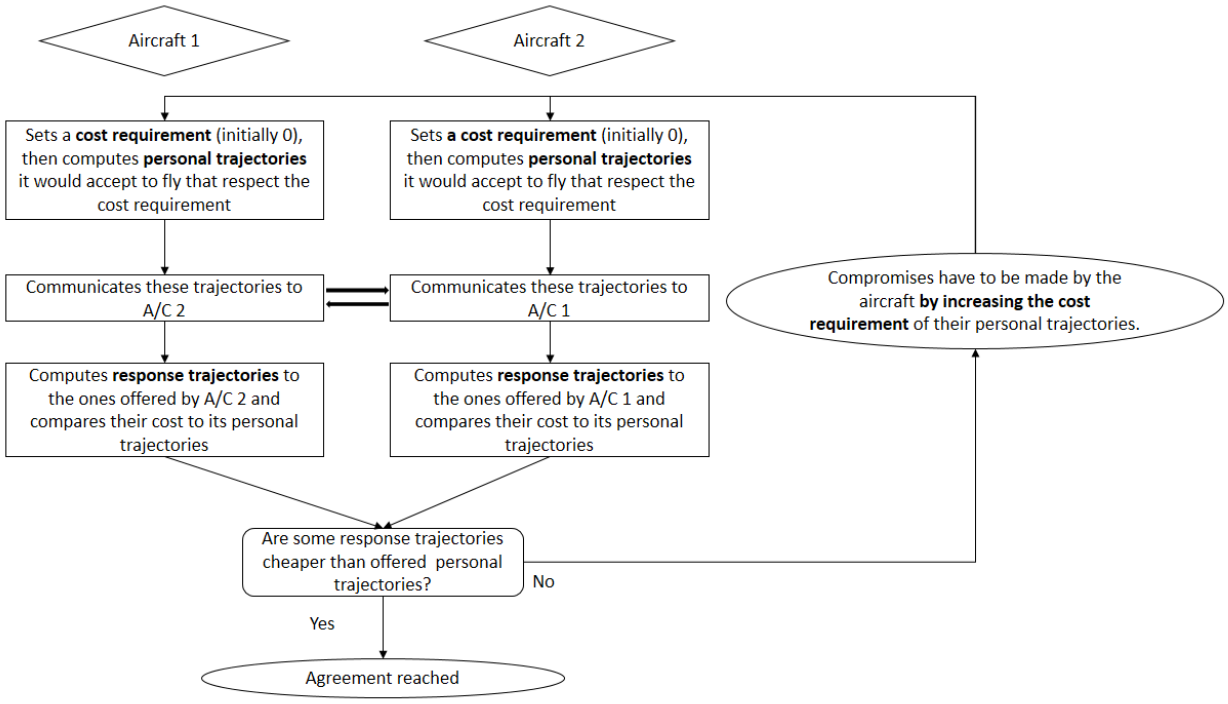

FIGURE 7: FLOWCHART OF THE SEQUENTIAL BARGAINING PROCESS

The trajectories computed at the sub-step 3 correspond to optimization problems subjected to constraints. Each aircraft tries to minimize its personal cost by providing trajectories that do not conflict with the trajectories offered by its opponent at the sub-step 1. Hence, the constraints on this optimization problem are the personal trajectories offered by the opponent at the sub-step 1 and the aircraft's performance constraints.

With this general architecture of the bargaining process, two important questions are raised: how to handle aircraft performance, and how to handle the multidimensionality of the resolution space. These questions each can be addressed by different approaches in the bargaining process, defined by the design variables described in the next two sections.

A third design variable is introduced to understand the impact of the maximum number number of iterations of the bargaining process on its convergence.

3.2 Design Variable 1: Representing Performance Constraints

The performance limits of the aircraft have a direct impact on the shape of the feasibility set. Consider, for example, the feasibility set of the bargaining problem with the following shape without performance limits, shown in Figure 8. The disagreement point here corresponds to the situation in which both aircraft maneuver as if the other one does not.

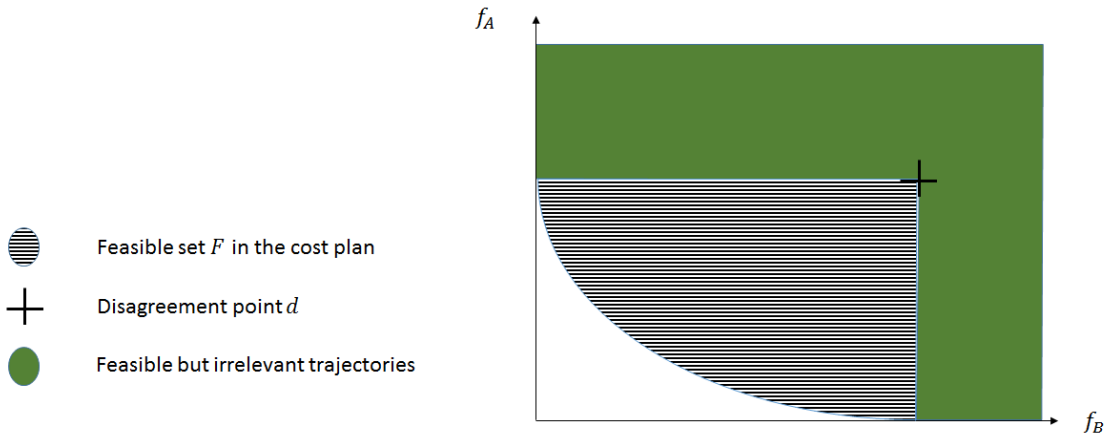

FIGURE 8: COST PLAN WITHOUT PERFORMANCE LIMITS

In this figure, two areas can be distinguished. The first one, in green, corresponds to feasible but irrelevant trajectories because they are Pareto suboptimal with respect to the disagreement point. The second area, in blue, corresponds to feasible trajectories which are at least better than the disagreement point.

In contrast, when aircraft are bounded by their performance limits, the feasible set will be reduced, as shown in Figure 9. In the worst case, the solutions that are Pareto-

optimal compared to the disagreement point can be completely cut out, as shown in Figure 10. Thus, aircraft performance and regulatory constraints can lead to situations in which the bargaining process might not converge because aircraft have little flexibility in their trajectories.

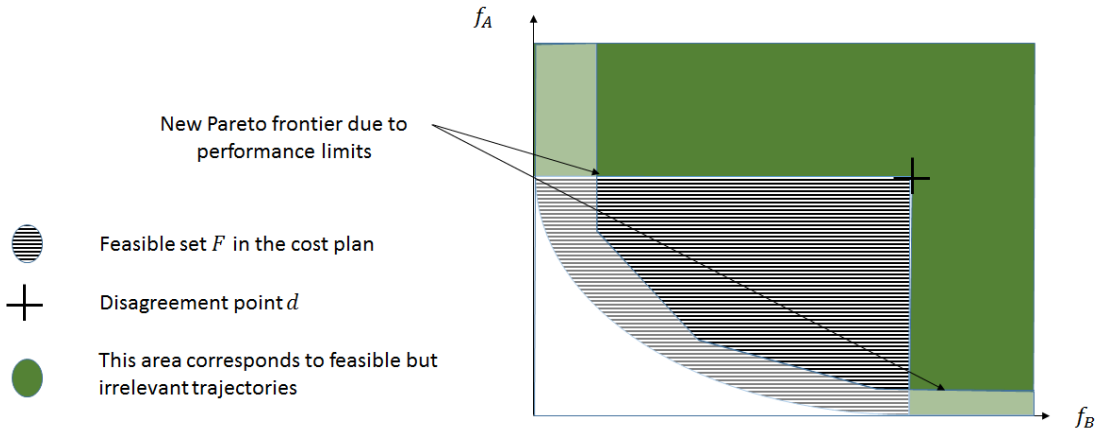

FIGURE 9: COST PLAN WITH RESTRICTIVE PERFORMANCE CONSTRAINTS REDUCING THE FEASIBLE SET

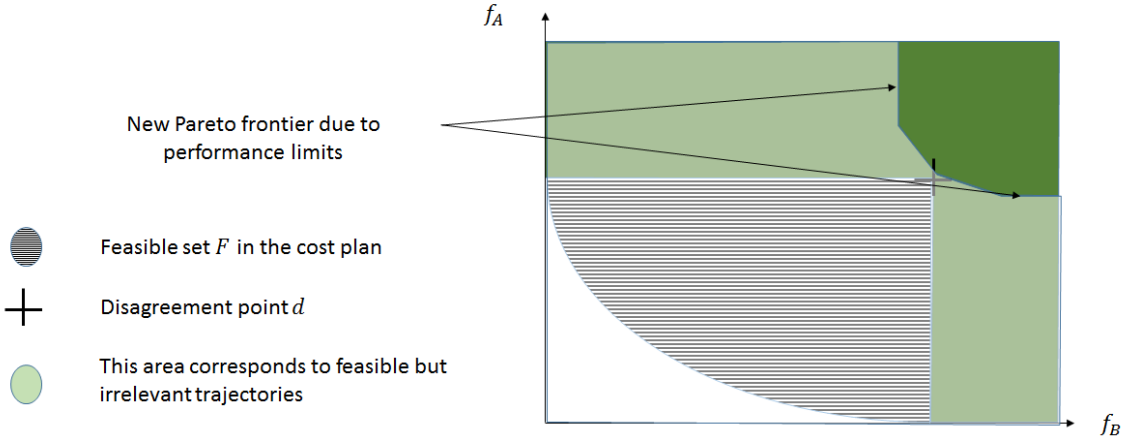

FIGURE 10: COST PLAN WITH MORE RESTRICTIVE PERFORMANCE CONSTRAINTS ELIMINATING THE FEASIBLE SET

The cost functions and definition of the feasible set can represent performance constraints in different ways addressed by the first design variable in this bargaining process, as shown in Figure 11. Specifically, the "clipped" method handles performance constraints by clipping the feasible set on allowable intervals in each dimension to directly constrain the resolution space. The advantage of this method is that it can represent precisely the real costs of the allowable trajectories. The cost functions can be computed at first, and then the feasible set clipped at the minimum and maximum possible resolutions. The potential problem with this approach is that there is no guarantee that solutions converge when the feasible set does not meet the axioms assumed in game theory.

The second "infinite cost" method assigns infinite cost to trajectories that exceed performance constraints. The advantage of this method is that good properties of the shape of the cost functions can be enforced, such as convexity. However, this method gives unrealistically-high costs for resolution trajectories close to the boundaries of the allowable part of the resolution space.

The third "finite cost" method builds a nonlinear mapping that provides real costs within the constraints then switches to very high costs when approaching the constraints. In this design, the performance constraints are accounted for with finite cost penalties instead of infinite cost barriers.

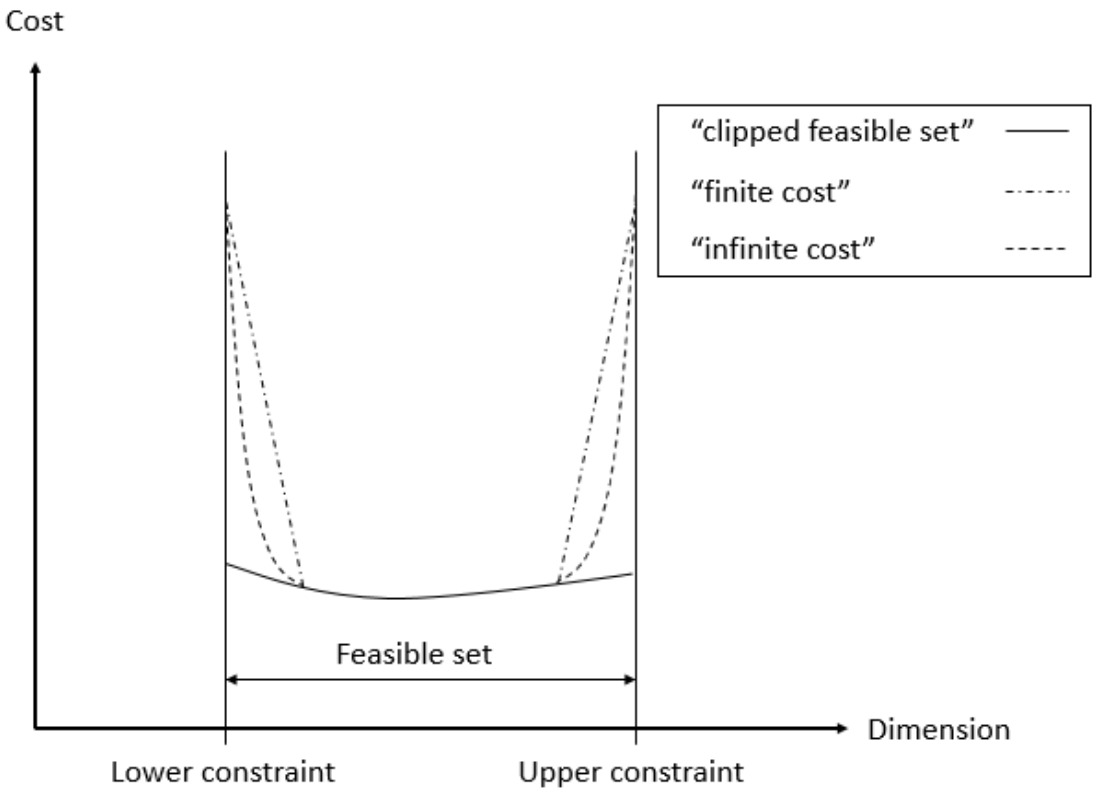

FIGURE 11: REPRESENTING PERFORMANCE CONSTRAINTS RELATIVE TO COST

3.3 Design Variable 2: Dimensionality of the Resolution Space

The space of resolutions represents the space of potential avoidance trajectories to clear the conflict. Aircraft have basically six levers of action:

1. Slow down
2. Speed up
3. Turn left
4. Turn right
5. Climb

6. Descend

These levers of action define a multi-dimensional space of resolutions constrained by the aircraft performance limits such as the maximum speed they can maintain at any given altitude and cruise ceiling. These different degrees of freedom yield a huge resolution space that cannot be entirely explored. The design choice made in this thesis is to define specific waypoints that correspond to a resolution in each of the different dimensions.

The second design variable within this bargaining process addresses how to handle this multidimensionality of the resolution space. For example, a parallel bargaining process might only examine personal and response trajectories that are within the same plane, or that are not. In this thesis, each personal or response trajectory will be only in one of the six possible dimensions, such that, for example, any aircraft won't climb and speed up in the same trajectory. However, this still provides a large number of possible maneuvers and allows for different aircraft in a conflict situation to maneuver in distinct dimensions. The number of trajectories computed has a direct impact on the computational effort and on the depth of the search of the resolution space.

The first way to incorporate multiple dimensions of the resolution space is, for each aircraft, to compute at the first sub-step six personal trajectories, one per dimension, and to communicate them to the other aircraft. Then, in the third sub-step, each aircraft computes six response trajectories, each in the same plane as the trajectory it is responding to. Hence, during the third sub-step, the costs of six response trajectories are compared with the costs of the personal trajectories offered by each aircraft. This corresponds to the situation in which every dimension is examined separately with potential pair of resolution trajectories each only in the same plane.

The second way to incorporate the multiple dimensions of the resolution space is for each aircraft to compute at the first sub-step six personal trajectories, one per dimension, and to communicate all of them to the other aircraft. Then, in the third sub-step, each aircraft computes six response trajectories for each trajectory received. Hence, during the third sub-step, thirty-six response trajectories are considered by each aircraft. The advantage of this method is that more potential solutions are explored than in 1, such that pair of resolution trajectories can be found in different dimensions. However, this requires more computational effort.

3.4 Design Variable 3: Convergence

The convergence mechanism ensures that the bargaining process converges in a finite number of steps. This mechanism must enforce compromises by forcing the aircraft to always propose personal trajectories that are more expensive than the previous ones. After a finite given number of steps, the pair of aircraft should simultaneously be in the situation in which they are offering personal trajectories that allows the other aircraft not to maneuver at all, if the problem is symmetric.

This is achieved using constants $\lambda_a > 0$ and $\lambda_b > 0$ such that, if a personal trajectory t_{K_i} has been proposed at step i by aircraft K, with a cost $f_K(t_{K_i})$, then aircraft K has to propose a personal trajectory $t_{K_{i+1}}$ at step $i+1$ such that $f_K(t_{K_{i+1}}) \geq f_K(t_{K_i}) + \lambda_K$. Intuitively, aircraft will progressively have to deviate from their initial zero-cost proposed personal trajectories to yield to each other.

Define the response trajectory for aircraft A to the zero-cost personal trajectory initially offered by B as having a cost c_{max}^{A}, corresponding to the situation in which A is the only aircraft to maneuver, and similarly, the response trajectory for aircraft B to the zero-cost personal trajectory initially offered by A as having a cost c_{max}^{B}, corresponding to the situation in which B is the only aircraft to maneuver.

With $\lambda_A = \frac{c_{max}^{A}}{n}$ and $\lambda_B = \frac{c_{max}^{B}}{n}$, after n steps, aircraft A would have to propose a personal trajectory with a cost of at least $n\lambda_A$=c_{max}^{A} and aircraft B would have to propose a personal trajectory with a cost of at least $n\lambda_B = c_{max}^{B}$. Doing so, an agreement should be reached no later than after n steps, because each aircraft has to offer personal trajectories whose costs are more expensive than the cost of the optimal trajectory it should perform if the other aircraft was not maneuvering. n corresponds here to the maximum number of iterations of the bargaining process and constitutes a third design variable.

The convergence might break down in the situation in which performance constraints are hit in every dimension before the end of the bargaining process. The aircraft would be unable to make any further compromises because they could not provide trajectories with higher personal costs. In this situation, the aircraft would be stuck at a given step and could not proceed to the next one in the bargaining process.

3.5 Example

To illustrate this bargaining process, consider a very simplified symmetrical situation in which two aircraft A and B are heading towards each other at 400 knots; the cost function is the same for both aircraft and corresponds to the extra length flown in the resolution trajectories in nautical miles. In this simple example the aircraft are only looking

for resolution trajectories in the horizontal plane, and performance constraints are not applied. Figure 12 depicts the initial situation.

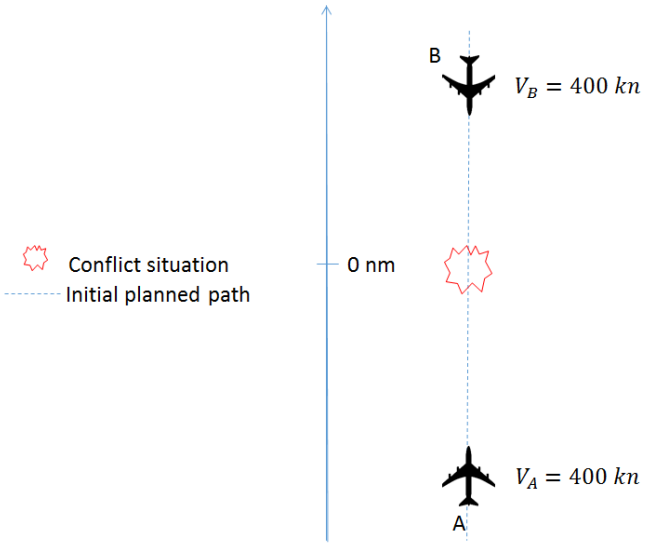

FIGURE 12: INITIAL CONFLICTING TRAJECTORIES IN A SIMPLIFIED EXAMPLE THAT ILLUSTRATES THE BARGAINING PROCESS

Step 0.1:

Each aircraft first proposes to not perform any maneuver, letting the other aircraft incur all the cost. This means that they propose personal trajectories (PT) of cost 0 to themselves.

Step 0.2:

Aircraft communicate the personal trajectories they computed in 0.1.

Step 0.3:

B receives the PT proposed by A, with its cost, and B's response trajectory (RT) solves the corresponding optimization problem of computing the trajectory with the lowest

possible cost that resolve the conflict with the trajectory proposed by A. The resulting RT has a cost of 0.7583 nmi. Likewise, A receives the PT proposed by B, with its cost, and computes its own RT as well. The resulting RT has a cost of 0.7574 nmi. (These two values should be equal due to the symmetry of the problem; the difference comes from numerical errors).

The two potential resolutions correspond to the following points on the cost plan, as shown in Figure 13.

$$\begin{cases} (0.7574,0) & \text{Unconditionally accepted by B} \\ (0,0.7583) & \text{Unconditionally accepted by A} \end{cases} \tag{8}$$

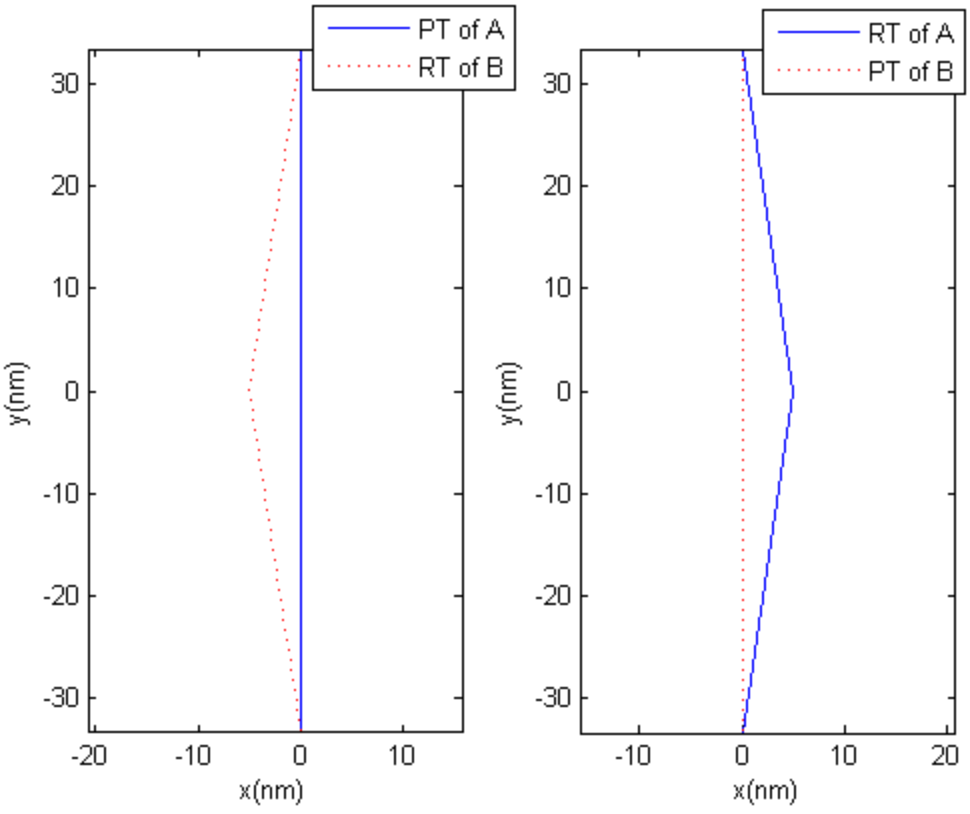

FIGURE 13: RESOLUTIONS EXPLORED AFTER STEP 0

Since $0.7574 > 0$, A does not accept (0.7574,0), and since $0.7583 > 0$, B does not accept (0,0.7583). No agreement is reached. At this point, it is known that $c_{max}^{A} = 0.7574$ nmi and that $c_{max}^{B} = 0.7583$ nmi. If we want the process to converge in less than $n = 12$ steps, we can compute the increased cost requirement for the next PT: $\lambda_A = 0.0631$nmi and $\lambda_B = 0.0632$ nmi.

Step 1:

B receives the PT proposed by A with cost to A of 0.0631, and computes its own RT, which has a cost of 0.3808. A receives the PT proposed by B, with cost to B of 0.0632, and computes its own RT, which has a cost of 0.3806. Hence the two potential resolutions correspond to the following points on the cost plan, as shown in Figure 14.

$$\begin{cases} (0.3806,0.0632) & \text{Unconditionally accepted by B} \\ (0.0631,0.3808) & \text{Unconditionally accepted by A} \end{cases} \quad (9)$$

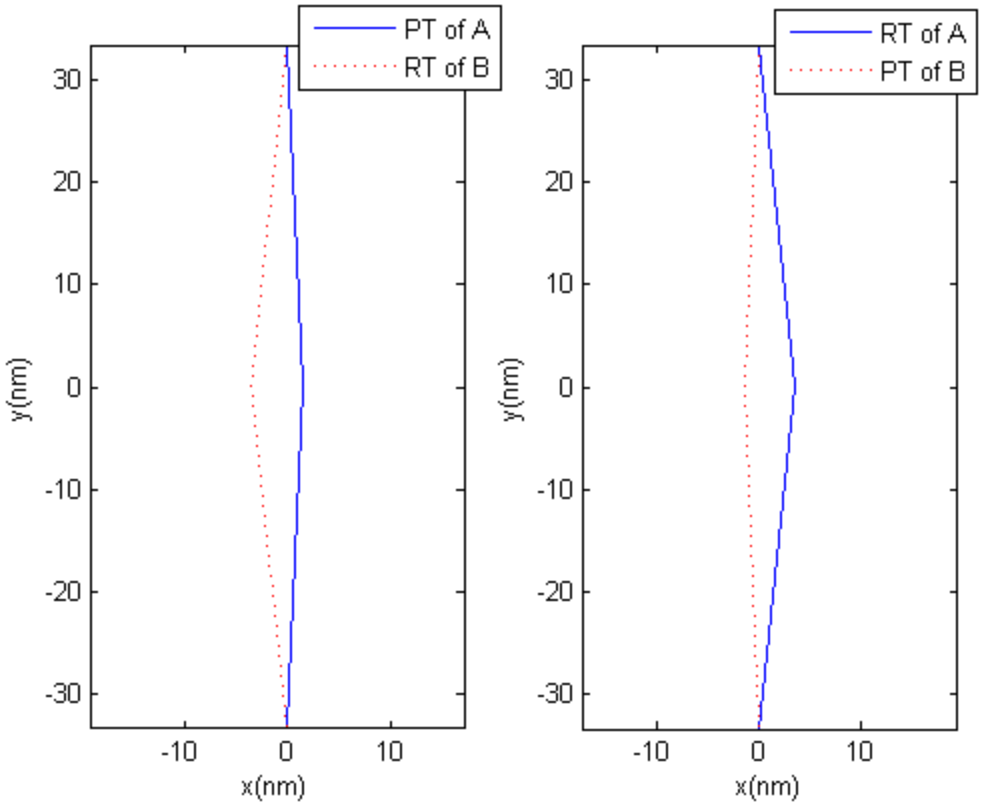

FIGURE 14: RESOLUTIONS EXPLORED AFTER STEP 1

Step 2:

Proceeding the same way as in the previous steps, B receives the PT proposed by A with cost to A of 0.1262, and computes its own RT, which has a cost of 0.2628. A receives the PT proposed by B with cost to B of 0.1264, and computes its own RT, which has a cost of 0.2626. Hence the two potential resolutions correspond to the following points on the cost plan, as shown in Figure 15:

$$\begin{cases} (0.2626,0.1264) & \text{Unconditionally accepted by B} \\ (0.1262,0.2628) & \text{Unconditionally accepted by A} \end{cases} \tag{10}$$

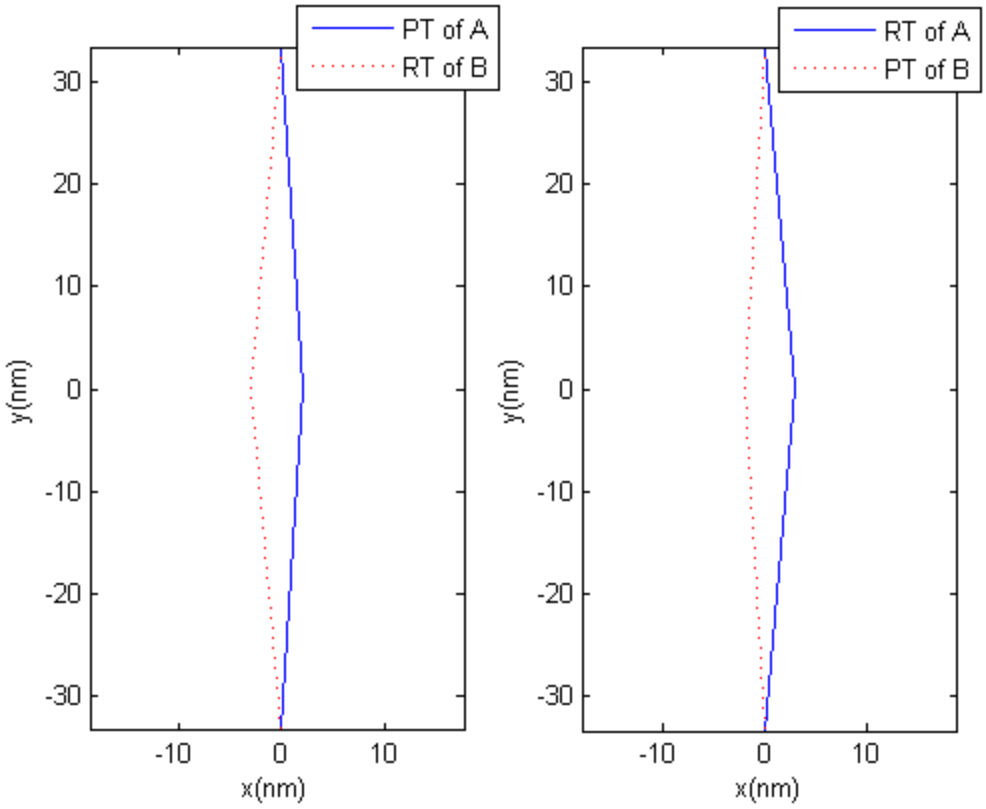

FIGURE 15: RESOLUTIONS EXPLORED AFTER STEP 2

Again, no agreement is reached here. The bargaining process finally stops at the 3rd step and gives two almost identical resolutions. The pair of aircraft can choose randomly between the two resolutions shown in Figure 16.

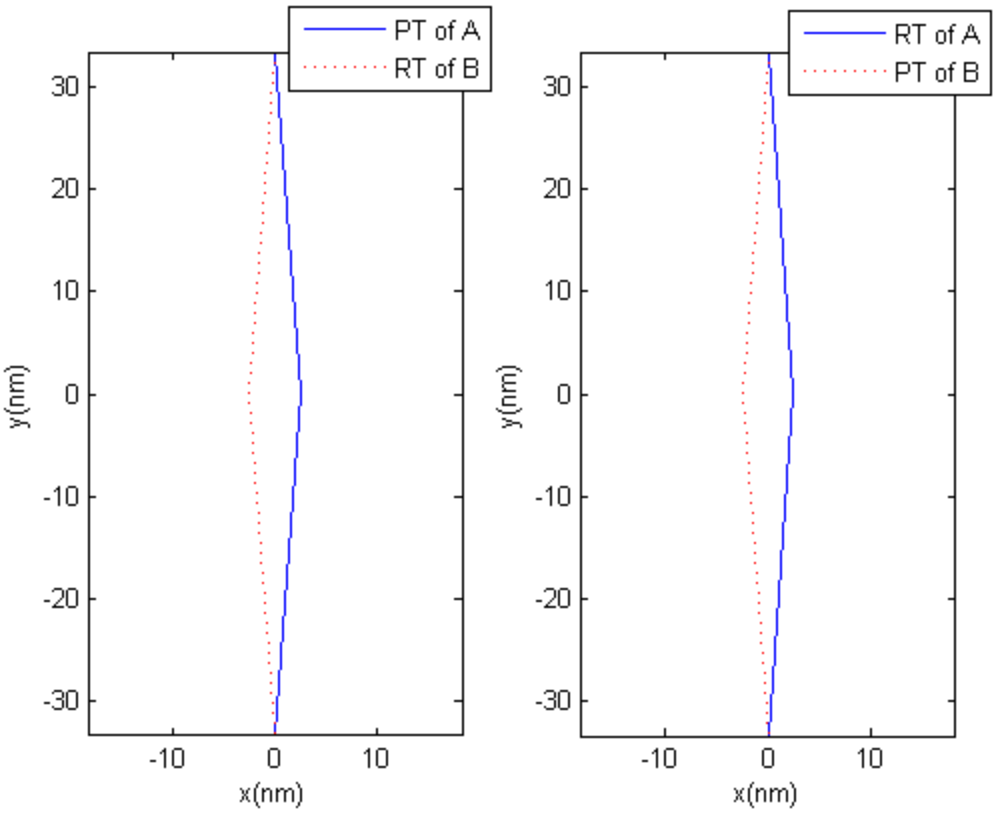

FIGURE 16: RESOLUTIONS EXPLORED AFTER STEP 3

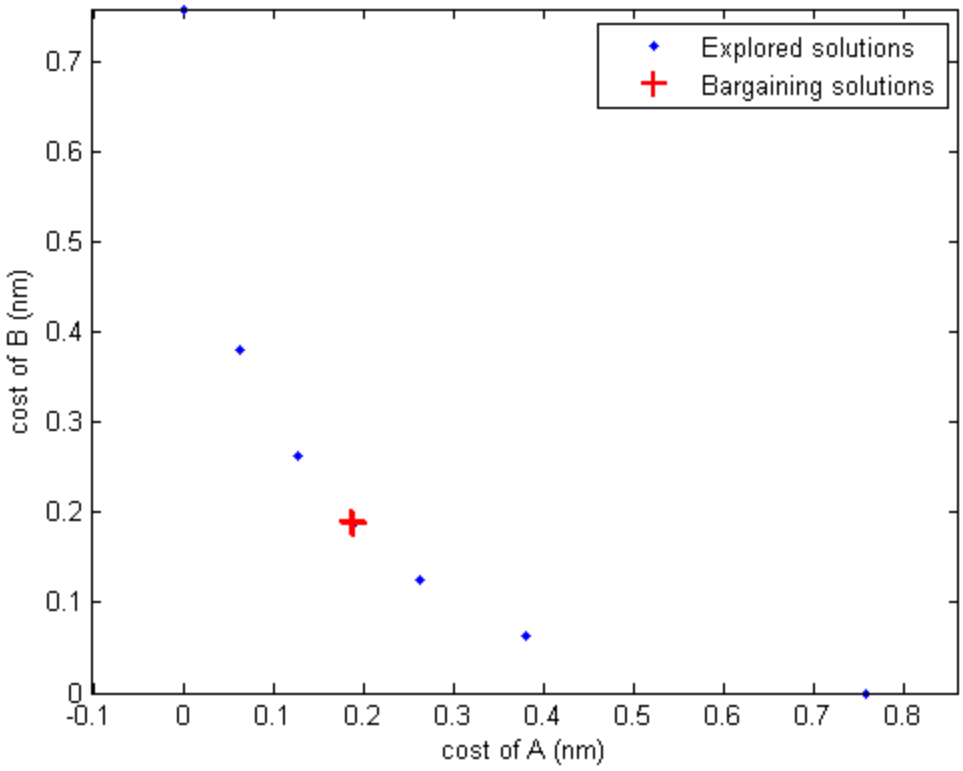

FIGURE 17: COST PLAN WITH N=12

In the end, all the points of the cost plan explored by the bargaining process in this simple situation are depicted in Figure 17. The curve created by these points corresponds to the Pareto frontier. The solution found by the bargaining process is fair in that it almost lies on the first bisector, i.e., both aircraft incur, in this symmetric case, the same cost.

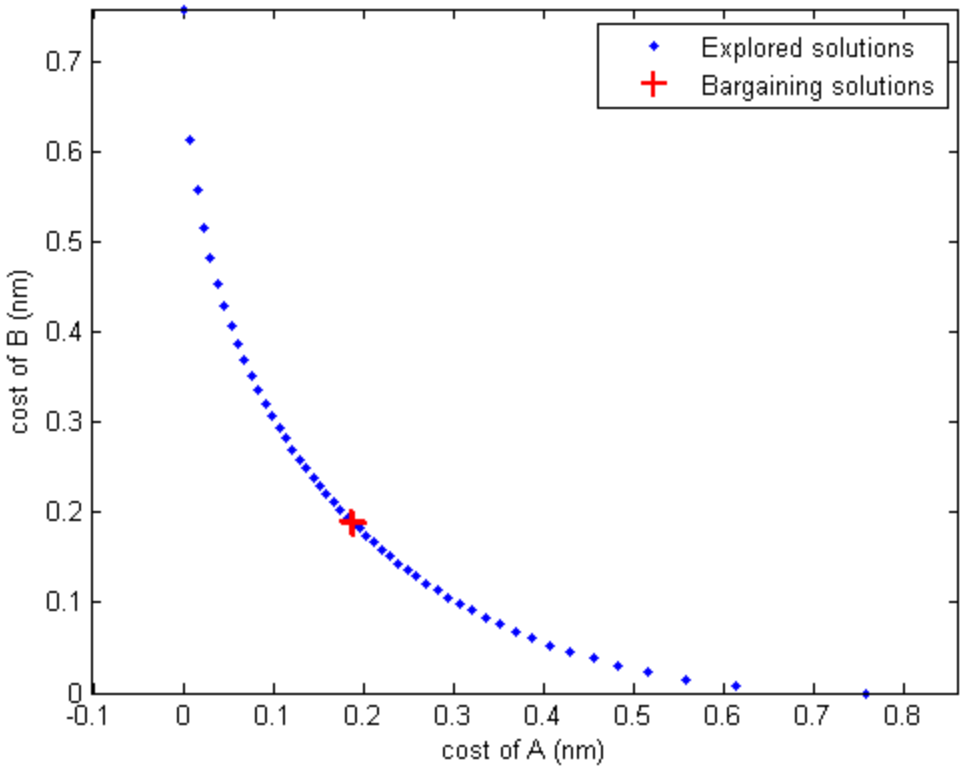

FIGURE 18: COST PLAN WITH N = 100

With $n = 100$ instead of $n = 12$, more points of the cost plan are explored, resulting in a higher precision, as depicted in Figure 18. This shows the impact of the choice of n. If it is too low, the final solution found by the bargaining process could cross further over the first bisector, impacting the fairness of the process. If it is too high, too many steps will be achieved before converging to the solution, which represents a higher computational effort.

3.6 Summary: Formal Definition

Using the constants λ_A and λ_B as described in 3.4, the i-th step has the following sub-steps:

1. A (resp. B) chooses personal trajectories $\{t_{A_i}^k\}_{k=1..6}$ (resp. $\{t_{B_i}^k\}_{k=1..6}$) in each of the six dimensions with their associated costs $\{f_A(t_{A_i}^k)\}_{k=1..6}$ (resp. $\{f_B(t_{B_i}^k)\}_{k=1..6}$) such that $f_A(t_{A_i}^k) \geq i\lambda_A$ (resp. $f_B(t_{B_i}^k) \geq i\lambda_B$).
2. A and B exchange their personal trajectories.
3. A (resp. B) responds to B's offered personal trajectories (resp. A's personal trajectories) by computing response trajectories $\{t_{Ar_i}^k\}$ (resp. $\{t_{Br_i}^k\}$) and their associated cost $f_A(t_{Ar_i}^k)$ (resp. $f_B(t_{Br_i}^k)$).

At the end of these three steps, we have the potential resolutions:

$$\begin{cases} (t_{A_i}^k, t_{Br_i}^k) & \text{Unconditionally accepted by A} \\ (t_{Ar_i}^k, t_{B_i}^k) & \text{Unconditionally accepted by B} \end{cases} \quad (7)$$

If $f_A(t_{Ar_i}^k) \leq f_A(t_{A_i}^k)$, A has no reason not to agree with B on selecting $(t_{Ar_i}^k, t_{B_i}^k)$. Similarly, if $f_B(t_{Br_i}^k) \leq f_B(t_{B_i}^k)$, B has no reason not to agree with A on selecting $(t_{A_i}^k, t_{Br_i}^k)$. When both inequalities $f_A(t_{Ar_i}^k) \leq f_A(t_{A_i}^k)$ and $f_B(t_{Br_i}^k) \leq f_B(t_{B_i}^k)$ are verified at the same step i, the aircraft can arbitrarily select one of the two potential resolutions as an agreement. Otherwise, the process moves to the $(i+1)$ step.

Unless prevented by performance constraints bounding the feasible set, $\{f_A(t_{A_i}^k)\}_i$ is an increasing sequence such that $f_A(t_{A_0}^k) = 0$ and $f_A(t_{A_n}^k) \geq n\,\lambda_A = c_{max}^A$, and

assuming $\{f_A(t_{Ar_i}^k)\}_i$ is a sequence such that $f_A(t_{Ar_0}^k) = c_{max}^A$ and $f_A(t_{Ar_n}^k) = 0$, there will automatically exist $i \in [\![0, n]\!]$ such that $f_A(t_{Ar_i}^k) \leq f_A\left(t_{A_i}^k\right)$. We can reason in the same way with B and replace A by B in the previous sentence. This proves the convergence of the defined bargaining process as long as the performance constraints do not excessively constrain the feasible set or make the cost functions too nonlinear.

CHAPTER 4: SIMULATION FRAMEWORK

WMC, "Work Model that Computes", is a simulation engine developed at the Cognitive Engineering Center at Georgia Tech. This engine permits the analysis of complex systems both at a local scale and at a larger system-wide scale.

The bargaining process described in the previous chapter was implemented by developing four core models within WMC. The first one is the outer loop aircraft dynamics given by *DCRAircraft*. The second one is the flight plan cost calculation, responsible for assessing the goodness of flight plan for a given aircraft. Then, the third core element is the Conflict Detector, essential to detect potential conflicts that need to be solved. Finally, the last core element developed is the Conflict Solver, which implements the bargaining process.

4.1 Outer Loop Aircraft Dynamics Model with Flight Plan Following

An outer loop aircraft model is a simple point-mass dynamics model that uses first-orders controllers to directly regulate eight states to follow a given trajectory: latitude, longitude, altitude, true airspeed, thrust, roll, heading, and flight path angle [15]. It was chosen for its relative dynamic simplicity to reduce simulation runtime. A fourth order Runge-Kutta adaptive step size integration algorithm (Cash-Karp method) is used to integrate the differential equations. Thus it reports its next update time to WMC as the current simulation time plus an additional adaptive time step, which is constructed using a 10^{-7} fractional error bound.

Aircraft performance limits are calculated using Eurocontrol's BADA (Base of Aircraft Data) performance models. BADA is a database of aircraft performance values

that can be used to model aircraft performance profiles throughout all flight regimes [16]. The performance values given for each available aircraft type include thrust, drag, and fuel coefficients as well as speeds and maximum altitudes. These values can be used with the Total Energy Model (TEM), a reduced point-mass model relating thrust, drag, acceleration, velocity, and vertical speed of an aircraft, to create performance profiles. This experiment uses the TEM to calculate all performance values from basic principles so as to parameterize the outer loop aircraft dynamics models, each according to its aircraft type. These values are also used in the flight plan cost calculation: they are used to compute the performance constraints of a given aircraft and the rate at which it burns fuel while cruising at a given altitude and speed.

The optimizer sets the initial trajectory of the aircraft to optimality by solving for the speed and altitude that minimize the cost of the trajectory. By noticing that optimal costs are always reached at maximum speeds, this problem was reduced to a one dimensional optimal problem that employs a golden ratio search.

4.2 Flight Plan Cost Calculation

The flight plan cost calculation examines the fuel burnt and time spent on each leg of the trajectory. These values are weighted according to the aircraft's cost index. The cost index is a parameter defined between 0 and 1 that is directly linked to the aircraft preference between minimizing fuel burn or delay.

To evaluate a given flight plan, the fuel burnt and time spent are computed between consecutive waypoints, and their weighted sum given by the cost index is stored. A cost

index close to 1 corresponds to weighting more delay, whereas a cost index close to 0 corresponds to weighting more fuel burn. If d corresponds to the delay between two consecutive waypoints, and f corresponds to the fuel burn between these two same consecutive waypoints, the flight plan cost c established between these two consecutive waypoints will be given by:

$$c = cost\_index * normalization\_factor * d + (1 - cost\_index) * f$$

The normalization factor is taken from [15] and is here as a bridge allowing the comparison between fuel burn and delay.

To implement the "finite" and "infinite" cost functions described in 3.2, linear and hyperbolic terms, i.e., penalty costs, are added when the aircraft is in a range of 30 knots of its maximum or minimum speed and in a range of 100 feet of its maximum or minimum altitude (arbitrarily set to 17000 feet). These ranges accounts for the situations in which the aircraft is flying close to the performance limits. When adding linear terms, flying at the performance limits of the aircraft yields an additional finite cost corresponding to the "finite" cost function; when adding hyperbolic terms, flying at the performance limits yields an additional infinite cost corresponding to the "infinite" cost function.

The cost index is set arbitrarily and not revealed to the other aircraft. Thus the cost functions defined for a given aircraft are personal, and allow only to assess the cost of a trajectory for this specific aircraft. Figure 19 shows the cost sensitivities to speed for an Airbus A320 flying at 25000 feet with a cost index of 0.5; the three different cost functions are shown.

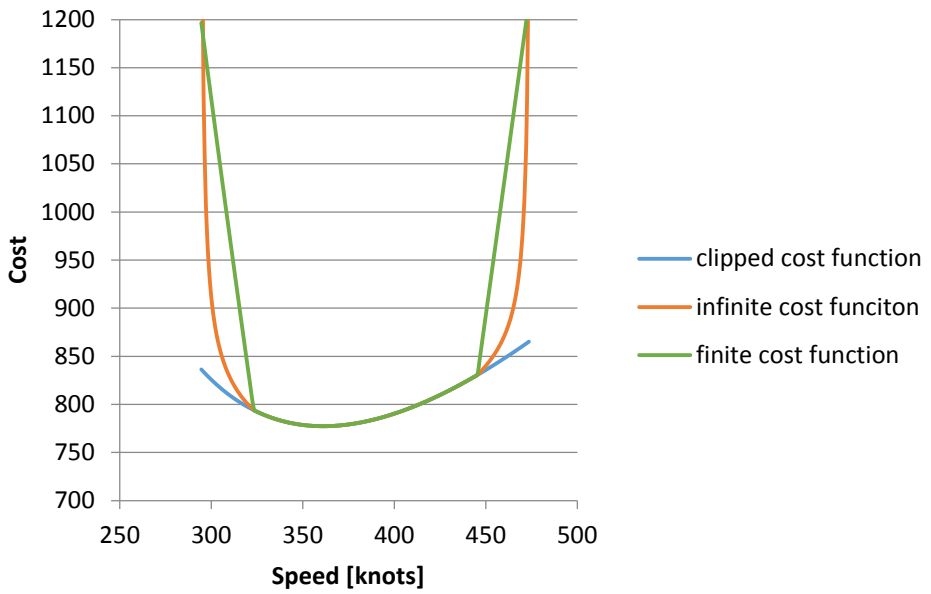

FIGURE 19: COST VERSUS SPEED AT 25000 FEET WITH A COST INDEX OF 50% FOR THE A320

4.3 Conflict Detector

The conflict detector first projects flight plans as lists of four dimensions waypoints into a Cartesian space. The distance between aircraft is evaluated every five seconds along their planned trajectories. If the horizontal or vertical separation is not respected, then the pair of aircraft is marked as conflicting and the conflict solver is scheduled three hundred seconds before the conflict is predicted to happen.

4.4 Conflict Solver

The conflict solver directly implements the bargaining via sequential notification described in CHAPTER 3. If no solution is found, the corresponding conflict is recorded as unresolved. When a solution has been found, conflict metrics are saved and the conflict

detector is run immediately to check for potential downstream conflicts. The time of closest distance between the aircraft is marked down and passed to the conflict solver to create additional waypoints for the two conflicting trajectories at that specific time. The resolution space is directly explored by shifting these new waypoints in the different dimensions. Solutions are then implemented as flight plans with added waypoints.

4.5 Summary

Four core models were developed in WMC to implement the bargaining process. The aircraft dynamics were simulated using an outer loop aircraft model given by *DCRAircraft*. Multiple instances of *DCRAircraft* can be run together, each of them able to follow their personal flight plan. Decentralized conflict detection is simulated by allowing aircraft to solve their conflicts with a local detection range corresponding to a look ahead time of three hundred seconds. The Conflict Solver directly implements the negotiating process detailed in the previous chapter.

CHAPTER 5: PAIRWISE CONFLICT RESOLUTION EXPERIMENTS

This experiment created a set of simulations of only two aircraft (i.e., pairwise simulations) that allowed for a full factorial design examining the three design variables in a range of conflicts. The following sections detail the experiment design, the metrics used to assess the bargaining process, and the results of the experiment.

5.1 Experiment Design

This experiment examines the overall performance of the bargaining process, and the impact of the three design variables described in 3.2, 3.3 and 3.4. The first design variable has three different methods for representing performance constraints in the feasible set and cost function. The second design variable corresponds to two different ways the multidimensional resolution space is searched, with a full or partial "within plane" exploration. The third design variable corresponds to the maximum number of iterations of the bargaining process: (n = 100, 200 or 500).

A full-factorial experiment design explored the eighteen combinations of these design variables in several different conflict conditions. Specifically, these conditions included the following.

Three different geometrical situations:

In [17], Andrews identified three main classes (or Rules) of horizontal conflict geometries that are each best resolved by different types of avoidance maneuvers. One conflict geometry per rule was selected to maximize the chances to see the aircraft using all the resolution space:

1. Rule A: 30 degrees with 0 miss distance. In this situation, aircraft are converging, almost following the same route. Andrews identified that the best horizontal resolution to this geometry eliminates closure rate.
2. Rule B: 150 degrees with 0 miss distance: aircraft are crossing. This corresponds to the rule B in [17]. Andrews identified here that the best horizontal maneuver to this geometry increases the existing miss distance [17].
3. 90 degrees with a miss distance of about 4 nautical miles. This corresponds to the rule C in [17]. Here, Andrew identified that the best horizontal maneuver to this geometry reinforces path crossing [17].

Three different pairs of cost indices:

1. 10% and 90%. In this situation, the first aircraft values more the fuel burn than the delay, while the second aircraft does the opposite.
2. 90% and 90%. Here, both aircraft use the same weight between fuel burn and delay, and give priority to minimizing the delay.
3. 10% and 10%. In this final situation, both aircraft use the same weight between fuel burn and delay again, but now give priority to minimizing the fuel burn.

Aircraft types:

1. Both aircraft are of the type Airbus A320 and have the same performance constraints.
2. Aircraft are of different types (A320 and A319) and therefore have different performance envelopes.

In the end, 18 different conflict conditions were making a total of 18x18 = 324 runs as shown in Table 1.

TABLE 1: PAIRWISE CONFLICT RESOLUTION EXPERIMENT DESIGN CONDITIONS

Independent Variable		**Number of Conditions**
1	Design variable for performance constraints	**3**
1.1	Clipped	1
1.2	Infinite	1
1.3	Finite	1
2	Design variable for multidimensions	**2**
2.1	In-plane	1
2.2	Full space	1
3	Design variable for convergence	**3**
3.1	100 iterations	1
3.2	200 iterations	1
3.3	300 iterations	1
4	Conflict geometry	**3**
4.1	Rule A	1
4.2	Rule B	1
4.3	Rule C	1
5	Cost indices	**3**
5.1	10% and 90%	1
5.2	90% and 90%	1
5.3	10% and 10%	1
6	Aircraft type	**2**
6.2	Same (A320)	1
6.3	Different (A320 & A319)	1
		Total: 324

5.2 Metrics

For each run, the following metrics were saved:

- Whether a resolution is found or not.
- The number of iterations.

- The dimensions of the negotiated conflict resolution trajectories.
- The costs of the negotiated conflict resolution trajectories, using an unmodified cost function and if used, any modified cost function accounting for performance constraints.
- The cost difference between the negotiated conflict resolution trajectories of the two aircraft.
- The margins of the final resolution trajectories from the maximum and minimum speeds and altitudes.
- The actual trajectories of the aircraft to examine the significant differences between the flight plans of the aircraft and the trajectories they actually flew, and to verify that the negotiated trajectories were effectively solving the conflict situation.
- The total cost of the final agreement point, summed across both aircraft computed using the clipped feasible set.

5.3 Results

5.3.1 Overall

5.3.1.1 Convergence

A first point to note is that the bargaining process was always successful in the sense that it converged in all 324 runs. This also means that the third design variable studied, corresponding to the maximum number of iterations, didn't impact the convergence of the bargaining process.

5.3.1.2 Impact of the Geometry of the Conflict

Figure 20 displays the different dimensions of the resolutions selected with the different conflict geometries (Rule A, B or C). It is interesting to note that the conflicts were usually resolved with maneuvers in the horizontal plane for the conflict geometries corresponding to rule B and C, whereas rule A was mostly solved with speed maneuvers, and some vertical maneuvers.

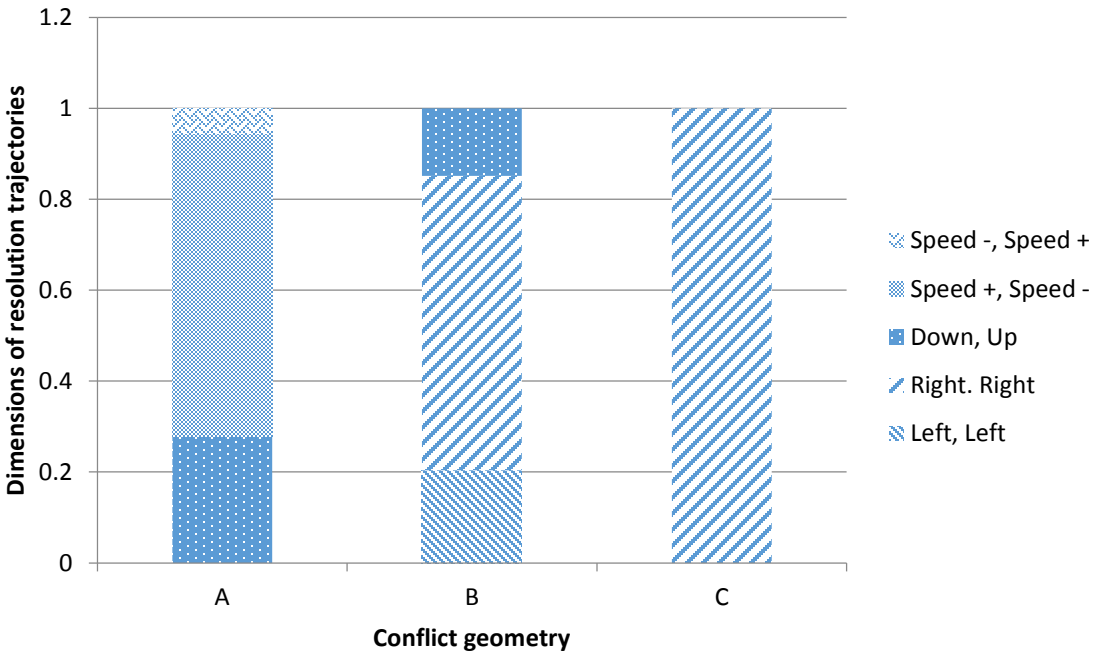

FIGURE 20: DIMENSIONS OF RESOLUTION USED FOR THE DIFFERENT CONFLICT GEOMETRIES

Rule A corresponds to a conflict angle of 30°, which means that aircraft are roughly going in the same direction and that the potential conflict period lasts longer because of a small relative velocity between aircraft. Thus, the choice to perform speed maneuvers in that situation can be interpreted qualitatively in the sense that this type of maneuvers will solve the conflict by progressively reducing the time span of the conflict to zero. The sort

of solutions used are different from the ones identified by Andrews for this Rule, as they belong to another dimension. In his study, Andrews only looked for horizontal maneuvers [17], which means that these maneuvers might be more efficient than the ones he identified.

Rule B corresponds to a conflict angle of 150°, which is the opposite situation of Rule A where aircraft are flying in opposite directions, have a high relative speed, and therefore a short period in which they are conflicting. In this situation, horizontal maneuvers seem to be costing less to the aircraft than speed maneuvers. The high relative speed between aircraft increase the efficiency of horizontal maneuvers as a change of direction will have a higher impact on the relative distance between aircraft across time. The horizontal maneuvers performed by the aircraft therefore increase the existing miss distance as identified by Andrews in [17].

Rule C corresponds to a conflict angle of 90° with a miss distance of 4 nautical miles. This means that Rule C is close to a no conflict situation because only one nautical mile of separation is missing. As for Rule B, horizontal maneuvers seem to be the cheapest choice for the aircraft. However, aircraft are not choosing to reinforce the path crossing predicted by Andrews [17]. It would have been the case if the aircraft had chosen opposite directions to maneuver, i.e., one aircraft going left, towards the other aircraft, while the other is going right, towards the first one. Instead, the aircraft systematically chose to maneuver in the same direction, both turning left, or both turning right.

5.3.1.3 Resolution of Symmetric Conflicts

In cases where aircraft had the same ability to maneuver (i.e., same aircraft type) and used the same cost index, and where performance constraints did not impact the resolution negotiated by the bargaining process, the global characteristics of the cost plan

seen earlier in the example in section 3.5 were observed again: the Pareto frontier defined by the feasible set explored during the bargaining appeared to be convex and symmetrical. Figure 21 and Figure 22 show a cost plan and a visualization steps graph corresponding to a situation where performance limits were not hit.

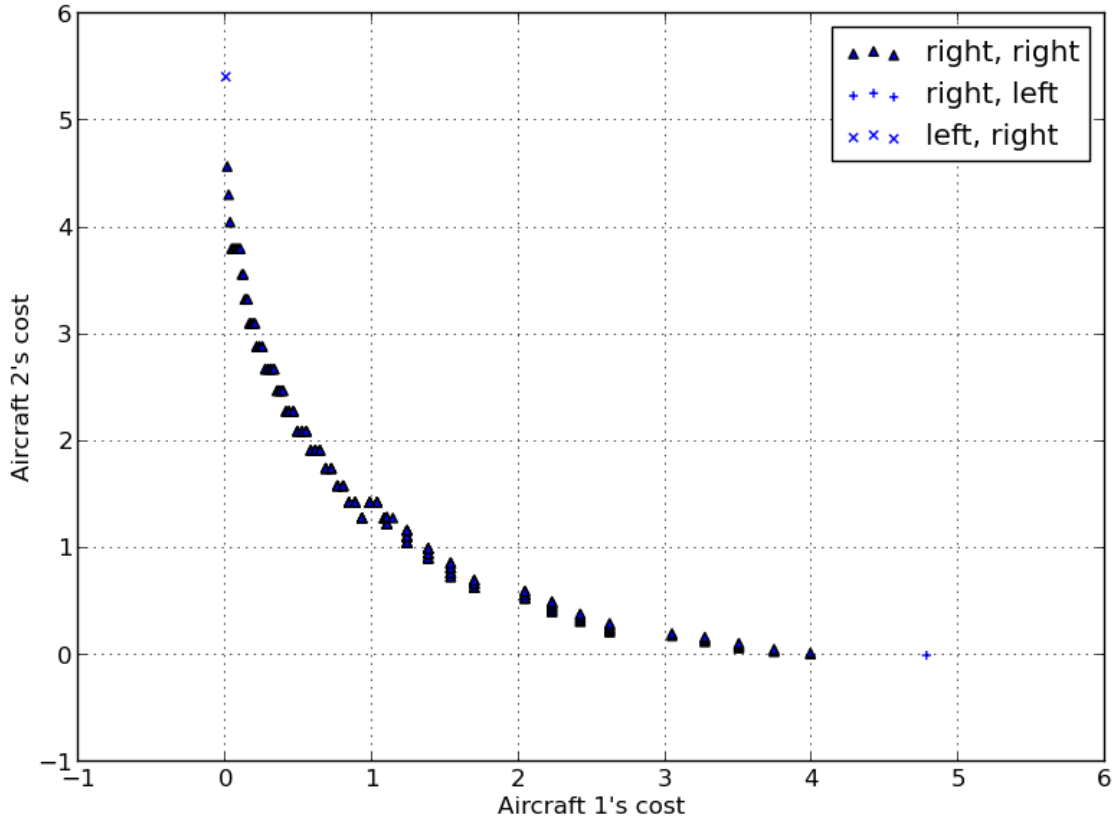

FIGURE 21: COST PLAN FOR A HORIZONTAL RESOLUTION NOT IMPACTED BY PERFORMANCE CONSTRAINTS: RULE C, COST INDICES 10% AND 10%, TYPES A320 & A320, CLIPPED FEASIBLE SET, FULL SPACE, N = 500

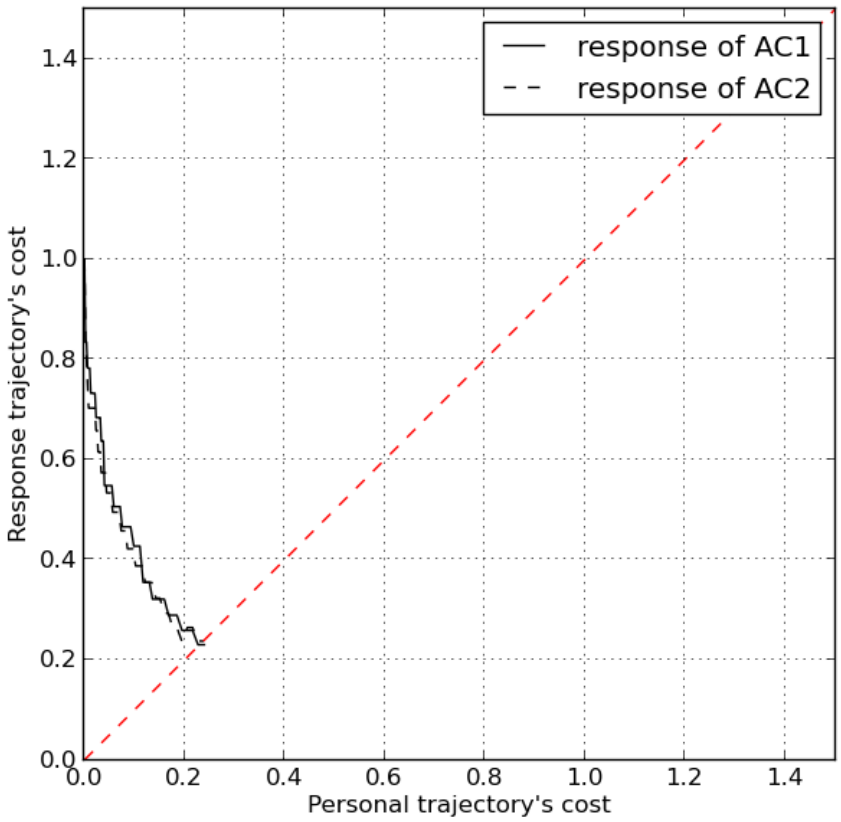

FIGURE 22: STEPS VISUALIZATION GRAPH FOR A HORIZONTAL RESOLUTION NOT IMPACTED BY PERFORMANCE CONSTRAINTS: RULE C, COST INDICES 10% AND 10%, TYPES A320 & A320, CLIPPED FEASIBLE SET, FULL SPACE, N = 500

Figure 22 shows the response trajectory costs as a function of the personal trajectory costs, where costs are normalized by the initial response costs. The evolution of the response trajectory costs can be seen by reading from left to right as personal trajectory costs increase at each step of the bargain. The curves of the response costs of the two aircraft are almost confounded, showing the symmetry of the resolution, and are constantly decreasing until the agreement is reached when the response trajectory costs get below the personal trajectory costs. Geometrically, it corresponds to the point where the curves cross

the first bisector of the plane. The assumption that the initial response trajectories costs are the maximum costs that could be incurred to the aircraft is verified here.

5.3.1.4 Impact of Asymmetry in the Cost Indices

Of the three different pairs of cost indices that were tested in the pairwise experiments, two pairs were the same (both at 0.1 or both at 0.9); the third created an asymmetry in which one aircraft prefers to reduce fuel burn (cost index of 0.1) while the other prefers to reduce the delay (cost index of 0.9). Figure 23 shows the cost difference between the negotiated resolutions as a function of the different pairs of cost indices. It can be observed that the asymmetric cost indices configuration has a significantly higher cost in the resolutions it negotiates.

Figure 24 compares the three costs plans resulting with the three different pairings of cost indices. It can be seen that, when identical cost indices are attributed to the aircraft, the solutions explored by the bargaining process define symmetrical Pareto frontiers. However, when different cost indices are given to the aircraft, the corresponding Pareto frontier is dissymmetrical and is located between the two symmetrical Pareto frontiers. Cutting the dissymmetrical Pareto frontier at the first bisector shows that it is formed by two legs that join the two symmetrical Pareto frontiers. The stars correspond to the costs of the negotiated trajectories finally obtained through the bargaining process. It can be seen that the resolution point associated with the asymmetrical pairing of cost indices is further away from the first bisector, which shows that the difference in cost between aircraft is higher.

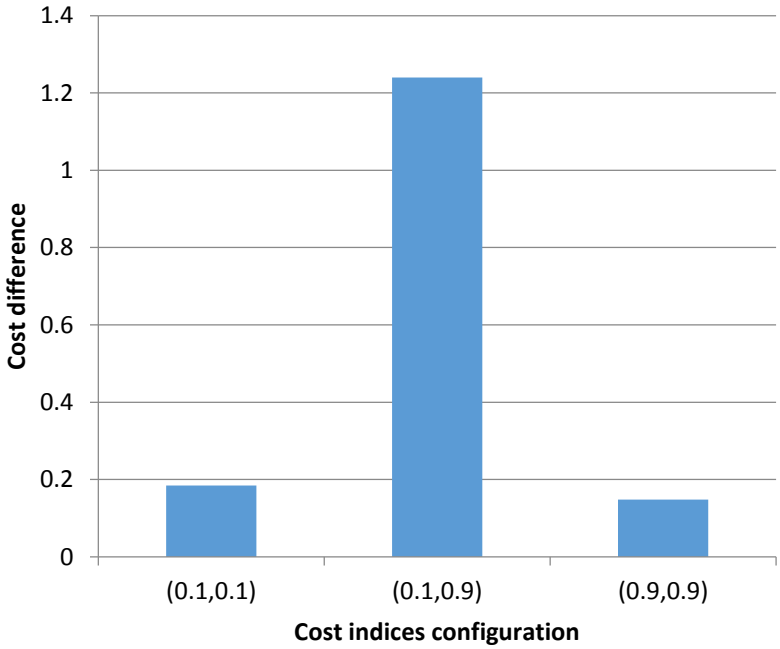

FIGURE 23: COST DIFFERENCES BETWEEN NEGOTIATED TRAJECTORIES FOR EACH PAIRING OF COST INDICES

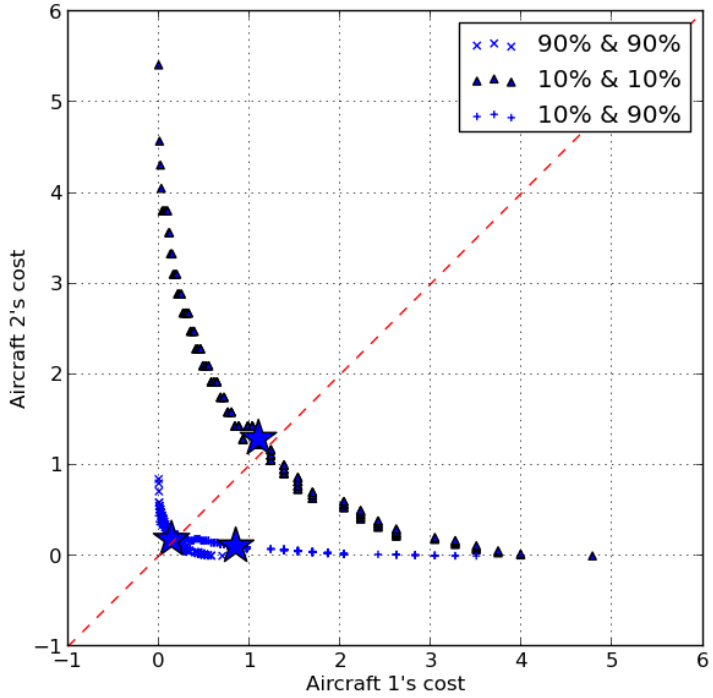

FIGURE 24: COST PLANS OVERLAID FOR EACH PAIRING OF COST INDICES. RULE C, TYPES A320 & A320, INFINITE COST FUNCTION, FULL SPACE, N = 500. THE STARS INDICATE THE AGREEMENT POINTS

5.3.2 Design Variable: Maximum Number of Iterations

The 'maximum number of iterations' design variable had no impact on the convergence of the bargaining process, since all of the 324 runs converged. However, as expected, Figure 25 shows that the average number of iterations required to solve the different pairwise conflicts increases linearly with the maximum number of iterations.

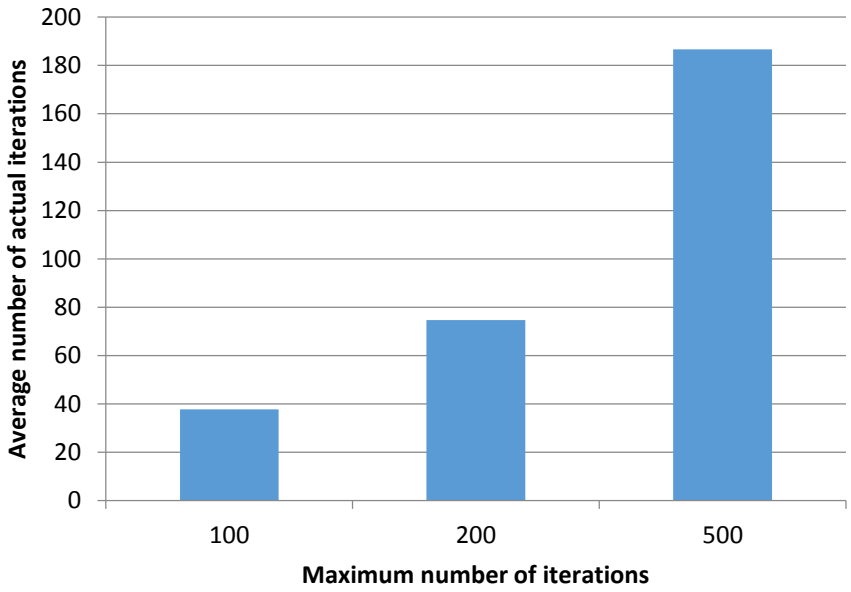

FIGURE 25: AVERAGE ACTUAL NUMBER OF ITERATIONS VERSUS MAXIMUM NUMBER OF ITERATIONS

The smaller the cost increment is, and the most precise the solution can be. However, no substantial gains were achieved by increasing the maximum number of iterations, and thus reducing the size of the cost increments of the personal trajectories. Figure 26 shows that the costs of the agreement trajectories were approximatively identical for the different maximum number of iterations tested. This shows that a value of 100 for

the maximum number of iterations yields small enough cost increments and therefore can be used for the large scale simulation presented in CHAPTER 6.

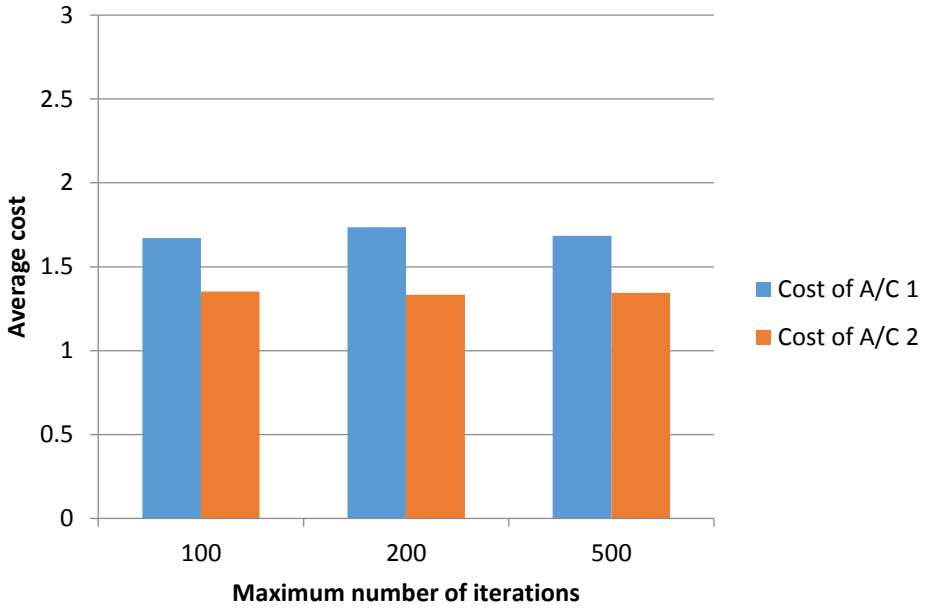

FIGURE 26: AVERAGE COSTS OF THE NEGOTIATED TRAJECTORIES FOR BOTH AIRCRAFT VERSUS THE MAXIMUM NUMBER OF ITERATIONS

5.3.3 Design Variable: Resolution Dimensionality

An important result is that the final trajectories given by the bargaining process were identical whether it was exploring solutions separately within the 3 dimensions (speed changes, horizontal maneuvers, vertical maneuvers), or across all directions, allowing pair of negotiated trajectories in different dimensions.

Thus, the full and partial 'within-plane' exploration of the feasible space gave the exact same outputs. An immediate consequence is that the computational efficiency of the bargaining process can be increased with no loss by only looking for pair of trajectories within the same dimension only.

5.3.4 Design Variable: Representing Performance Constraints

This section presents the overall impact of the performance constraints on the bargaining process. Performance constraints only impacted runs where the bargaining process explored solutions close enough to the performance limits that penalty costs were added when using the finite or infinite cost functions, or that the feasible set was clipped such that an aircraft couldn't offer higher cost personal trajectories. Two different situations are presented: where the performance constraints are hit but don't impact the final solution and where they are hit and actually impact the final solution.

5.3.4.1 Overall Impact

The performance constraints were handled by either clipping the feasible set, or by adding penalty costs with the 'finite' and 'infinite' cost functions. An intuitive result is that the negotiated trajectories calculated using the finite and infinite cost functions have higher margins with respect to the performance limits of the aircraft, as shown by Figure 27.

Another result is the fact that the total cost of the negotiated trajectories (assessed using the real, unmodified cost function), was on average lower for the resolutions negotiated when performance constraints were represented by clipping the feasible set than the ones representing performance constraints with the finite and infinite cost functions, as shown by Figure 28. Hence clipping the feasible set yields negotiated trajectories that are more efficient on a global perspective, but that have lower margins with the performance constraints.

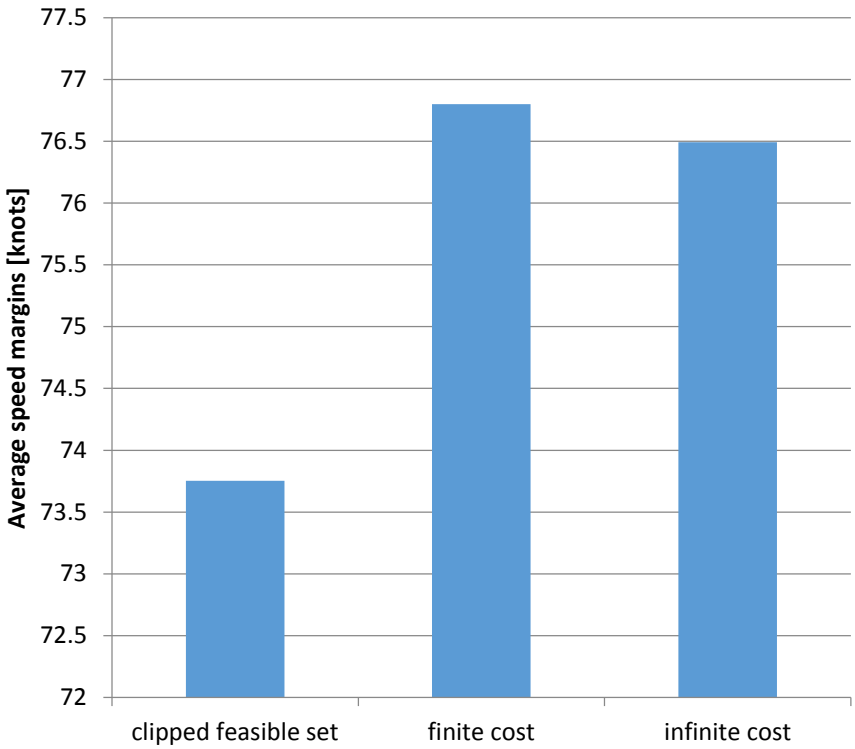

FIGURE 27: AVERAGE SPEED MARGIN WITH THE PERFORMANCE LIMITS VERSUS THE REPRESENTATION OF PERFORMANCE CONSTRAINTS

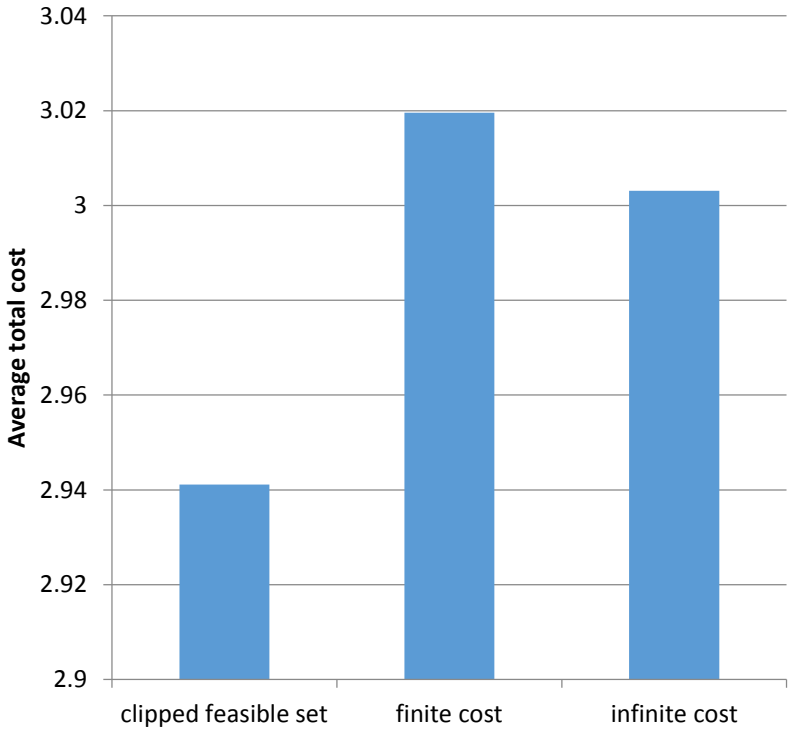

FIGURE 28: AVERAGE TOTAL COST VERSUS THE REPRESENTATION OF PERFORMANCE COSNTRAINTS

5.3.4.2 Performance Constraints Hit at the Beginning of the Bargaining

A first case is presented in which performance constraints are hit through the bargaining process but the final agreement point is far from the performance limits. Initially, the bargaining process started by selecting vertical maneuvers. However, these trajectories were close to the aircraft's maximum altitude and, after enough iterations in the bargaining process, the aircraft couldn't offer to go up anymore without hitting the constraint. At that point, they switched to feasible trajectories in the horizontal plane far from the constraints.

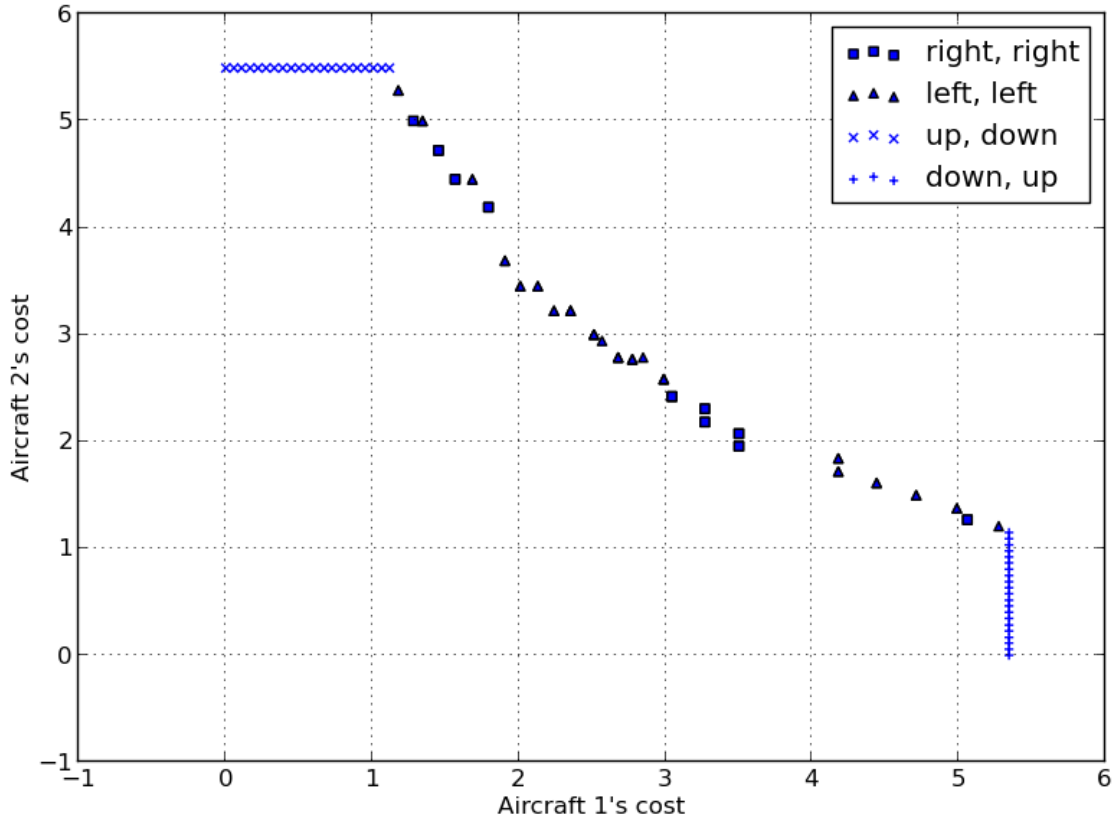

FIGURE 29: PERFORMANCE CONSTRAINTS HIT AT THE BEGINNING OF THE BARGAINING. RULE B, COST INDICES 10% AND 10%, TYPES A320 & A320, INFINITE COST FUNCTION. FULL SPACE, N = 100

This situation is shown on Figure 29, in a case using the 'infinite' cost functions to represent performance constraints. With the initial up/down resolutions, additional costs were introduced to penalize the proximity to altitude constraints, such that a small change in altitude incurred a high cost increment. Therefore, each aircraft could offer a trajectory with almost no significant changes that will meet the requirement for increased cost. This explains why the sequence of up/down proposed resolutions form nearly vertical and horizontal lines in the figure. Once the proposed vertical trajectories grew in cost to have the same costs as horizontal trajectories, the negotiation switched to the horizontal dimension and the subsequent bargaining follows.

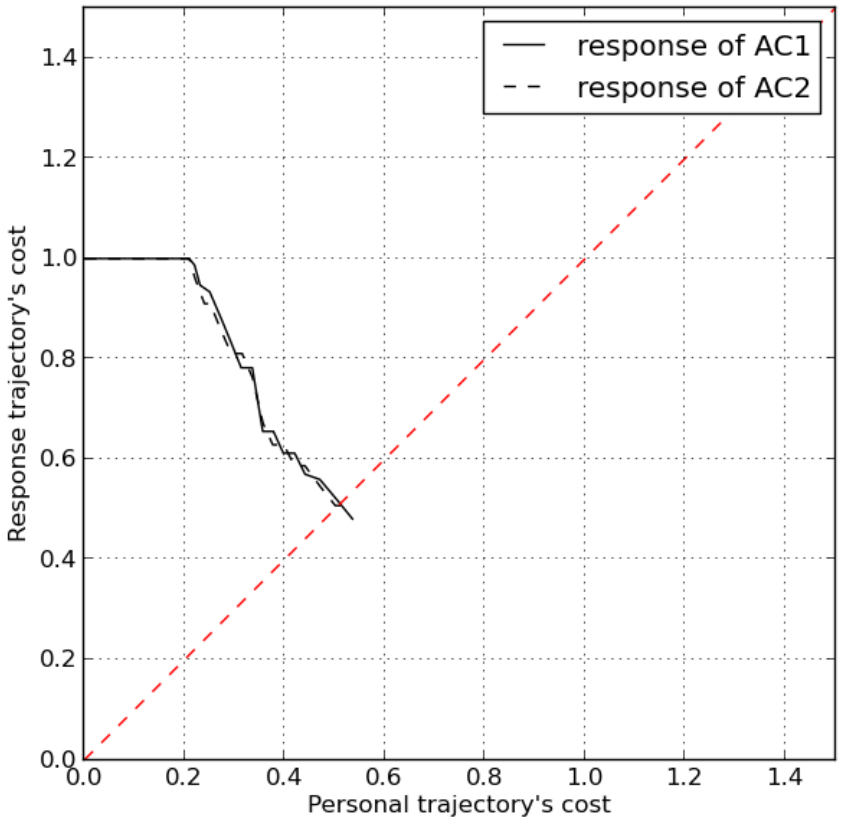

FIGURE 30: STEPS VISUALIZATION GRAPH. RULE B, COST INDICES 10% AND 10%, TYPES A320 & A320, INFINITE COST FUNCTION. FULL SPACE, N = 100

Figure 30 shows the response trajectories costs in function of the personal trajectories costs in that particular situation where constraints are hit at the beginning of the bargaining. Response trajectories costs remain high, at the same level as the initial response costs, until aircraft switch to horizontal maneuvers. The assumption that the initial response trajectories costs are the maximum costs that could be incurred to the aircraft is still verified here even though response costs are not anymore strictly decreasing as personal costs increase.

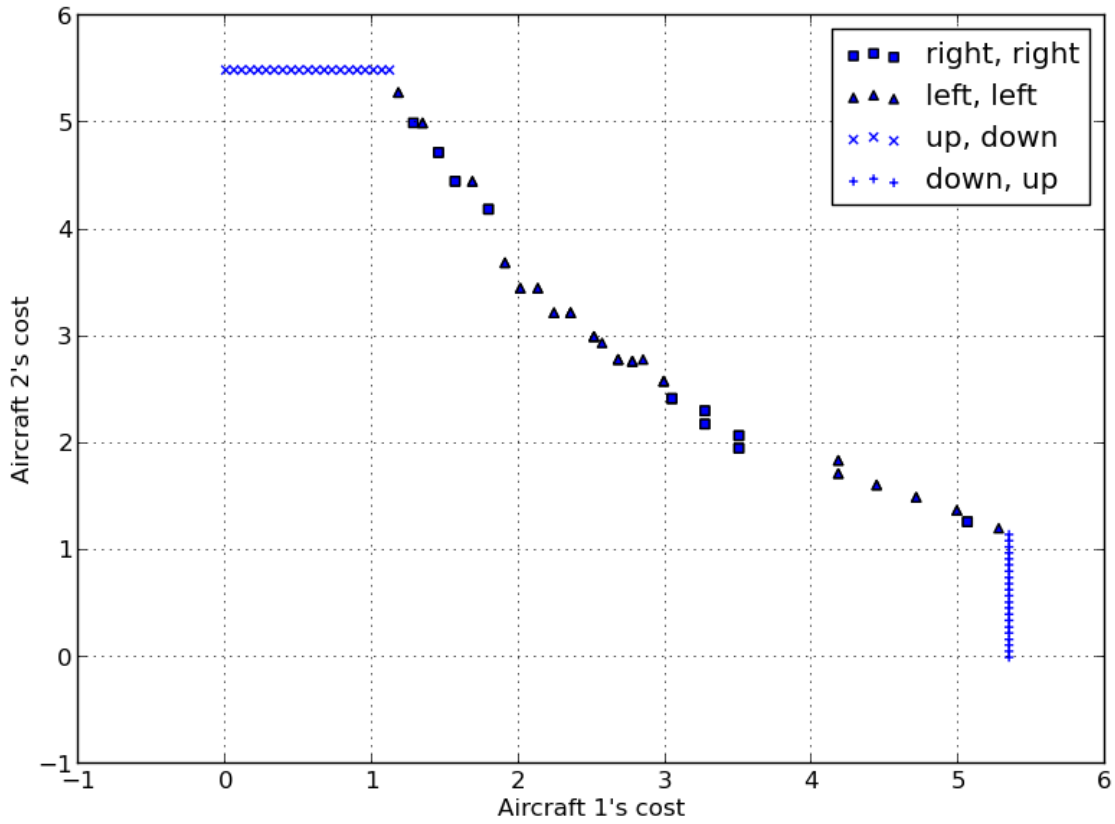

FIGURE 31: PERFORMANCE CONSTRAINTS HIT AT THE BEGINNING OF THE BARGAINING. RULE B, COST INDICES 10% AND 10%, TYPES A320 & A320, CLIPPED FEASIBLE SET. FULL SPACE, N = 100

Figure 31 depicts the same situation than before but this time using the 'clipped' feasible set. The only difference with Figure 29 corresponds to the points located close enough to the constraints that their costs differ when using the finite or infinite cost functions. The 'staircase pattern' that can be observed can be explained by the fact that response trajectories are computed by moving along the dimensions by slight geometrical increments of finite size. When successive offered trajectories don't provide large enough geometrical shifts, the other aircraft are forced to keep the same response trajectories, until there is enough space to gain one geometrical increment. Graphically, this corresponds to the gap between successive horizontal lines and vertical lines.

5.3.4.3 Agreement Point on the Boundaries of the Flight Envelope

A second case is presented here in which the final solution found by the bargaining process is located close to the performance limits of the aircraft. In this situation, the cost plans graphs are highly dependent on the representation of the performance constraints chosen.

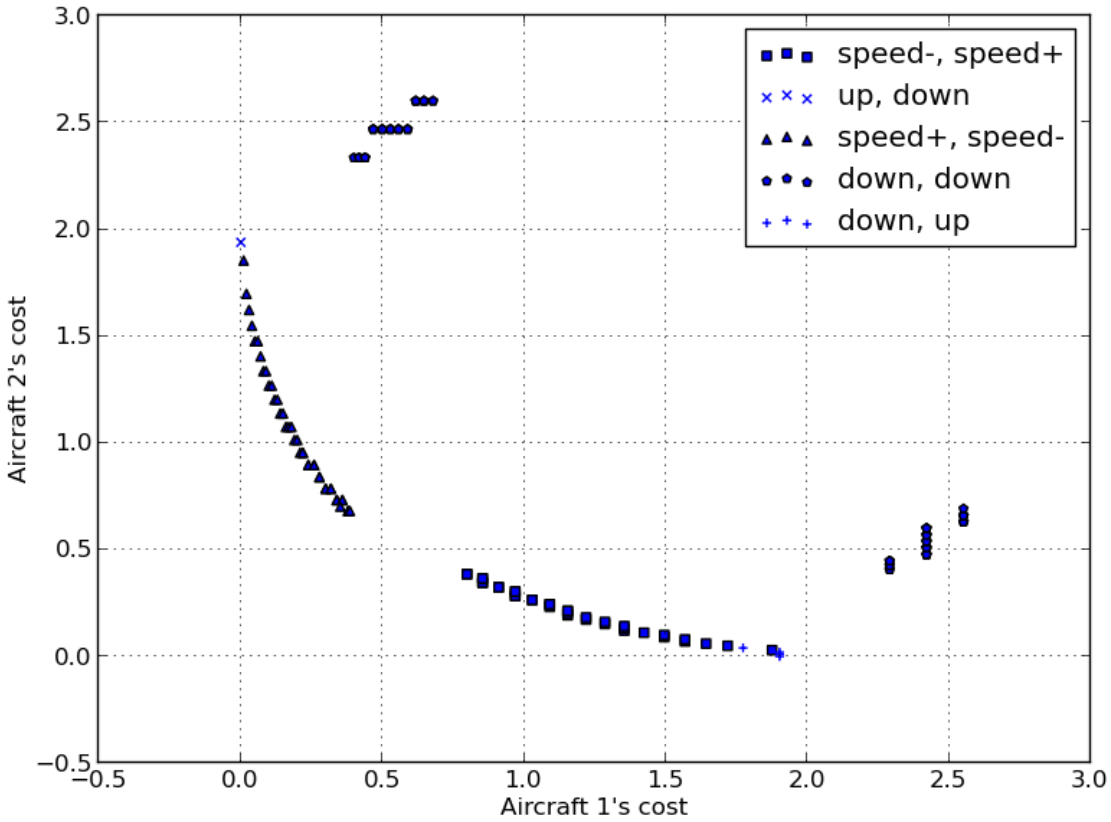

FIGURE 32: FINAL AGREEMENT POINT CLOSE TO THE PERFORMANCE CONSTRAINTS. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, CLIPPED FEASIBLE SET. FULL SPACE, N = 200

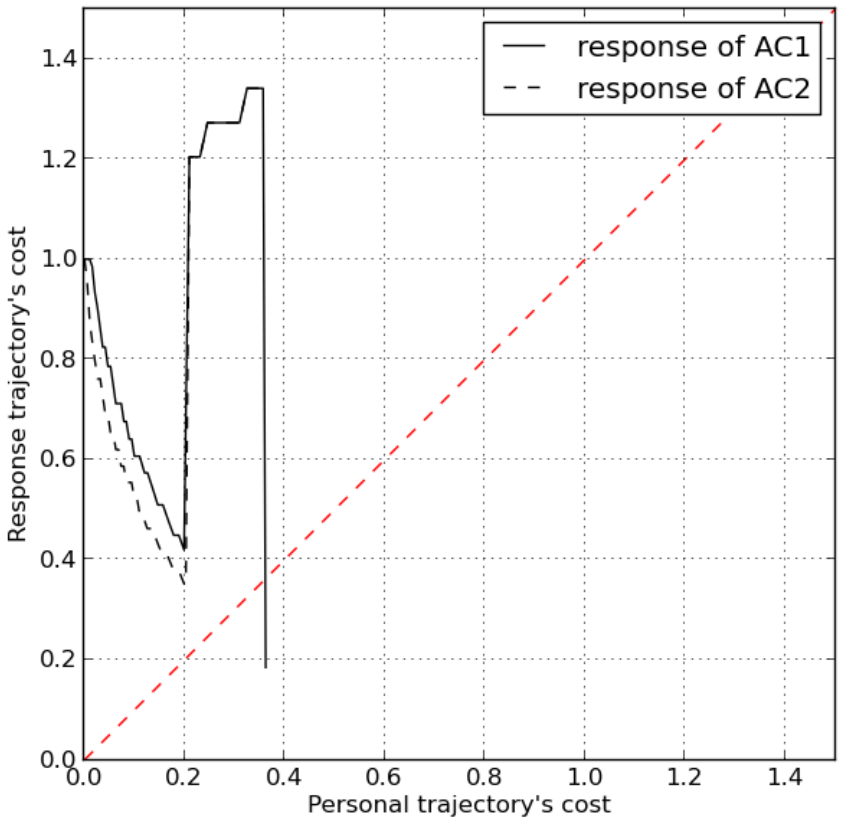

FIGURE 33: STEPS VISUALIZATION GRAPH. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, CLIPPED FEASIBLE SET. FULL SPACE, N = 200

In Figure 32, depicting the use of the clipped feasible set, aircraft first propose speed changes where one accelerates and the other decelerates. However before negotiating a solution, one or both aircraft reach their maximum speed and cannot propose to increase speed anymore. This means that aircraft can then only offer trajectories in five directions instead of six, forcing them to switch to other dimensions with higher costs. It is interesting to note that the assumption that the initial response costs are the highest ones is broken here: response costs keep increasing through the bargaining process goes, as shown on Figure 33. After each iteration, aircraft offer to descend more and more, and since they can't respond by climbing because they are already flying at maximum altitude, they are

forced to descend even more: they are ‘pushed down’. Even if the response costs incurred are increasing, they appear to be still lower than the response trajectories they could get by looking in other directions and that is why they are selected by the bargaining process. The cost requirements keep increasing until they finally offer personal trajectories where they slow down enough that they could respond by accelerating and exactly solve the conflict by flying at maximum speed: they switch back to the speed dimension, but this time, offered trajectories correspond to decelerations and response trajectories to accelerations. The agreement is reached at that point because at least one aircraft get a response trajectory cost lower than its offered trajectories costs but the symmetry of the solution is broken here. The cost of aircraft 2 is twice the one received by aircraft 1, even though the conflict situation considered here was symmetrical with identical aircraft and identical cost indices.

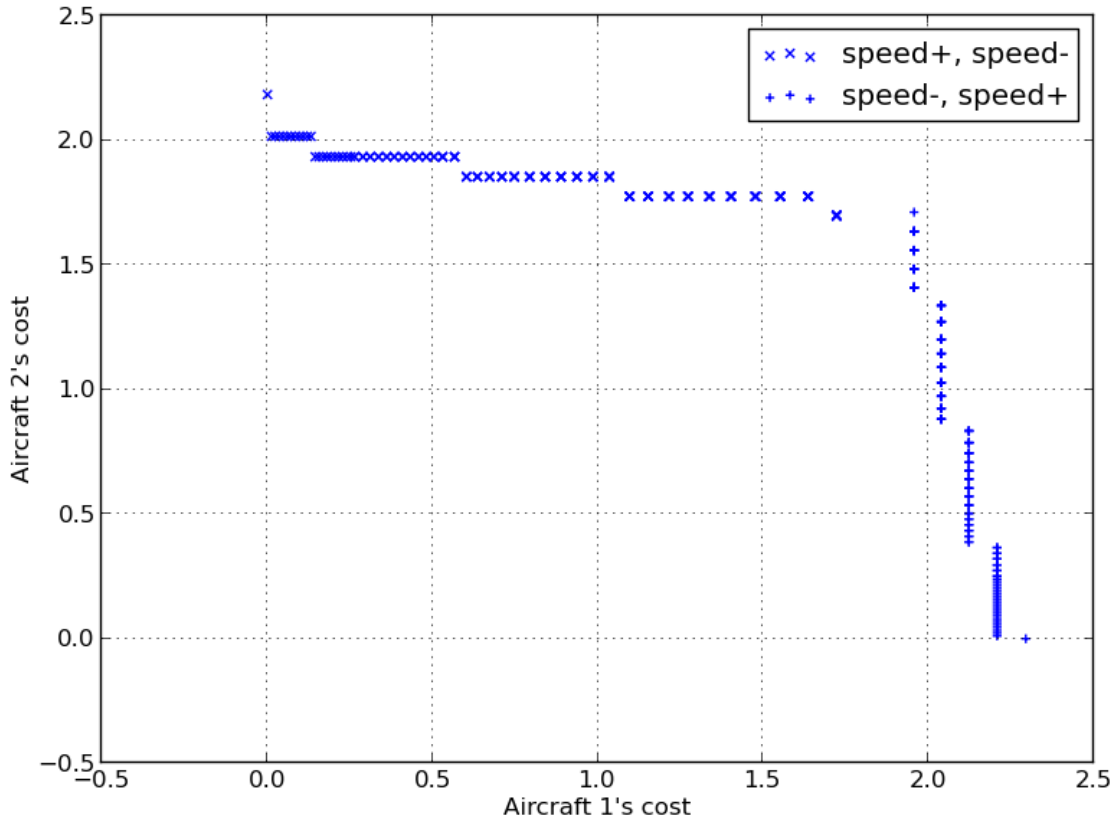

FIGURE 34: FINAL AGREEMENT POINT CLOSE TO THE PERFORMANCE CONSTRAINTS. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, INFINITE COST FUNCTION. FULL SPACE EXPLORATION, N = 200

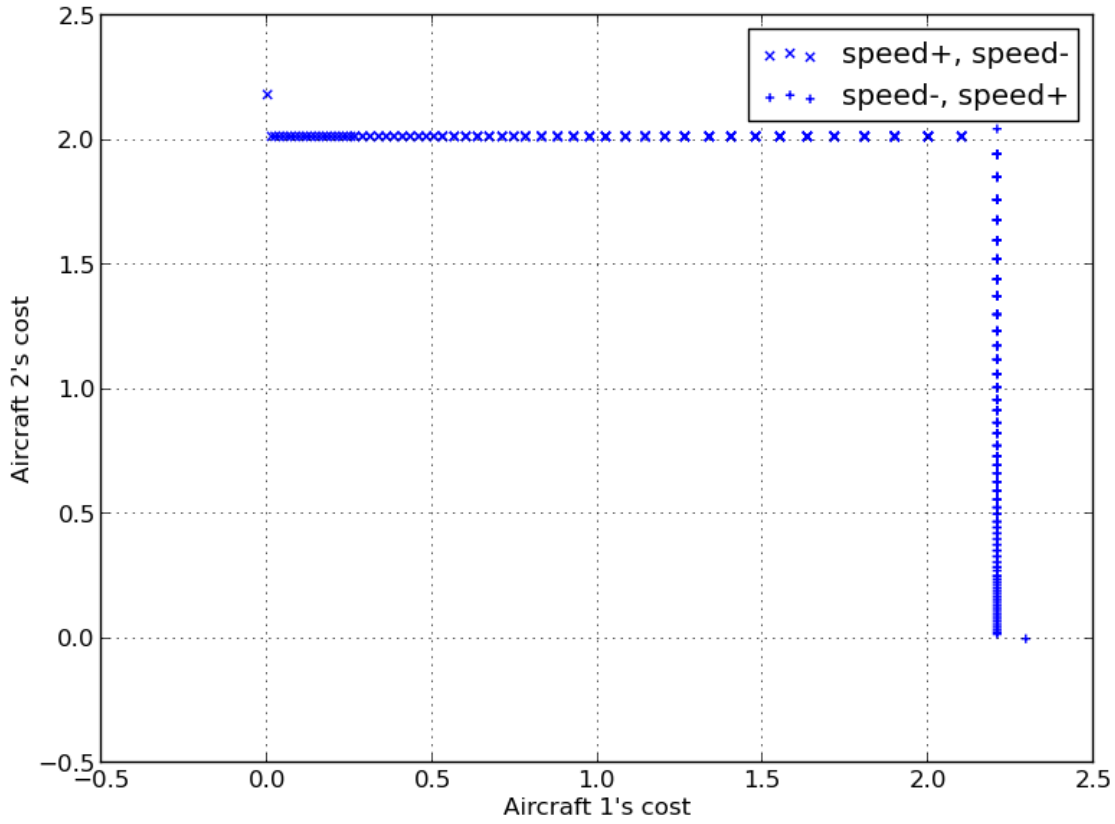

FIGURE 35: FINAL AGREEMENT POINT CLOSE TO THE PERFORMANCE CONSTRAINTS. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, FINITE COST FUNCTION. FULL SPACE, N = 200

On Figure 34 and Figure 35, the bargaining process uses respectively 'infinite' and 'finite' cost functions to solve the same conflict situation as before. Therefore, proximity to performance constraints is now penalized. The resulting cost plans of the bargaining are very similar. Aircraft only explore solutions in the speed dimension because they can keep offering trajectories by accelerating without reaching their maximum speed. Indeed, smaller and smaller velocities increases will keep satisfying the increasing cost requirement, and yield almost the same response trajectories as the geometry of the offered trajectories remain almost unchanged. An agreement is finally reached when the vertical and horizontal lines on the cost plan cross each other. The observed solution costs are close

for the two aircraft when using the modified cost function, but not at all when the negotiated trajectories are assessed using the real cost function. Indeed, one of the aircraft is almost not modifying its trajectory and the other aircraft has to do all the work by decelerating.

As shown in Figure 36, the speed margin of the solution obtained with the clipped feasible set is very low (2.52 knots). The margins obtained with the infinite and finite cost functions are much higher, meaning that aircraft stay further away from the performance limits, which was expected since constraints are penalized. In terms of performance, however, the solution obtained with the clipped feasible set is more efficient than the solutions obtained with the finite and infinite cost functions and has a lower total cost, as shown in Figure 37.

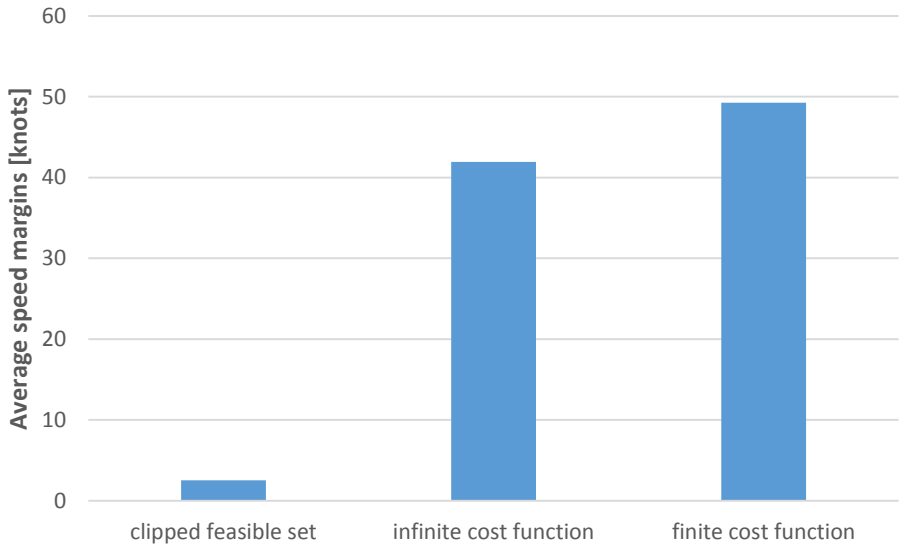

FIGURE 36: SPEED MARGIN WITH MAXIMUM SPEED VERSUS THE REPRESENTATION OF PERFORMANCE CONSTRAINTS. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, FULL SPACE EXPLORATION, N = 100

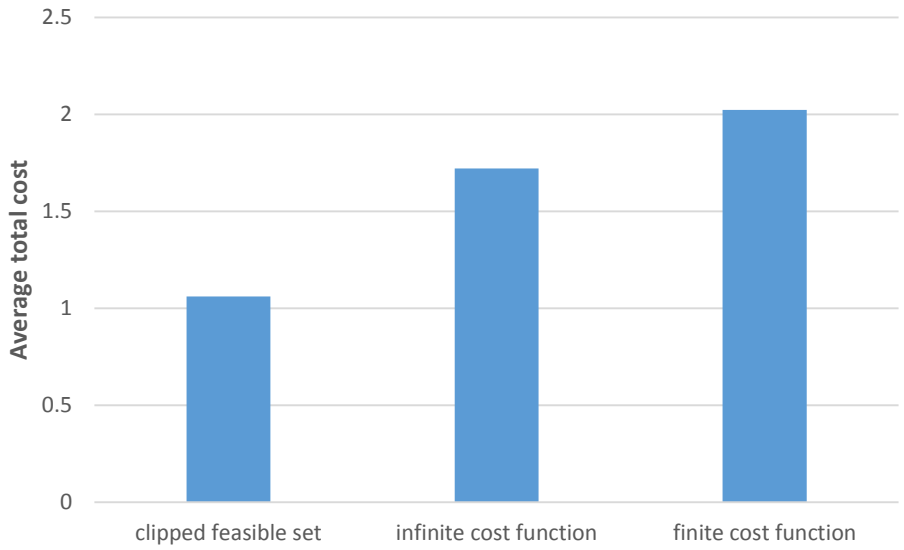

FIGURE 37: TOTAL COST VERSUS THE REPRESENTATION OF PERFORMANCE CONSTRAINTS. RULE A, COST INDICES 90% AND 90%, TYPES A320 & A320, FULL SPACE EXPLORATION, N = 100

5.3.5 Dependent Variable: Total Cost Analysis

The total cost of a feasible solution corresponds to the sum of the costs of the trajectories it requires of the two aircraft. To assess the efficiency of the bargaining process, it is interesting to compare the total cost of the final agreement point to the one that would be obtained if only one aircraft was maneuvering, which is common in current conflict resolutions selected by centralized air traffic controllers.

The expected total cost that would have been achieved if one aircraft had been chosen randomly to do all the maneuvers corresponds to the mean of the total cost in the two initial solutions explored by the bargaining process where aircraft simultaneously offer not to maneuver, leaving the other aircraft do all the trajectory modifications. Therefore, this expected total cost could be associated with a fictive solution point, corresponding to the mean of the initial feasible points explored by the bargaining process where only one

aircraft is maneuvering, as shown in Figure 38. Solving a conflict situation with the bargaining process is, for a global point of view, more efficient than a one aircraft maneuvering solution if the total cost of the negotiated trajectories is found to be lower.

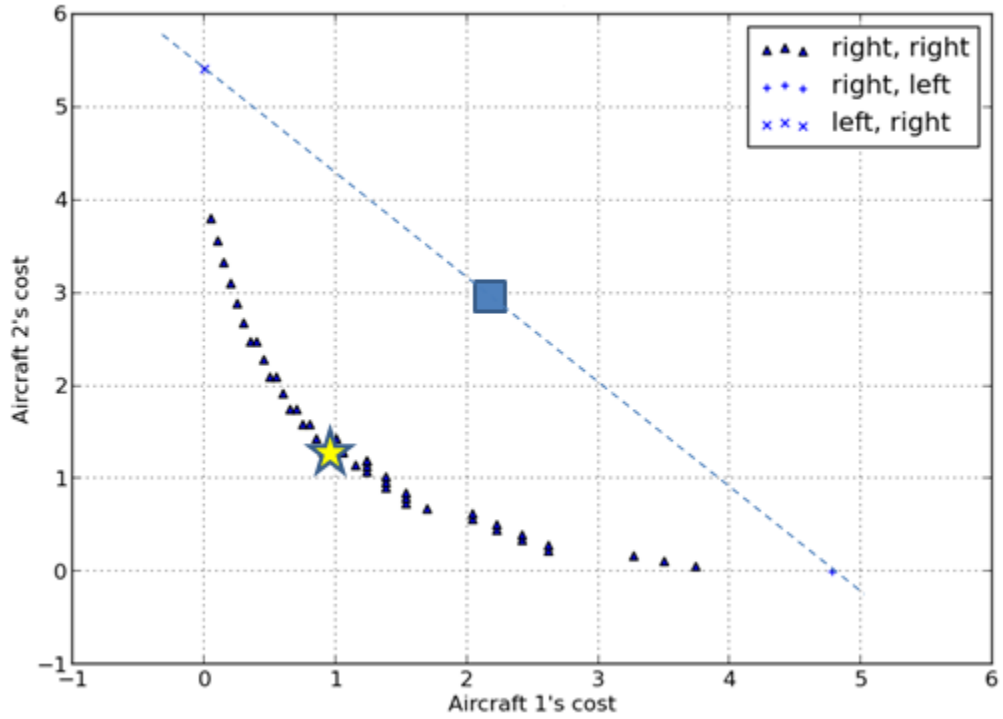

FIGURE 38: THE SQUARE INDICATES THE FICTIVE SOLUTION POINT ACHIEVED BY THE GROUND CONTROLLER, MEAN OF THE FEASIBLE POINTS EXPLORED AT STEP 0 DURING THE BARGAININNG PROCESS, WHILE THE STAR INDICATES THE SOLUTION ACHIEVED BY THE BARGAINING PROCESS

As shown in Figure 38, a geometrical sufficient condition for the solution provided by the bargaining process to be more efficient than the one aircraft maneuvering solution is to obtain a convex Pareto frontier on the cost plan. If the bargaining solution is 'above' the dashed line that links the initial explored feasible points, then the one aircraft maneuvering solution will be more efficient, limiting the interest of the bargaining process.

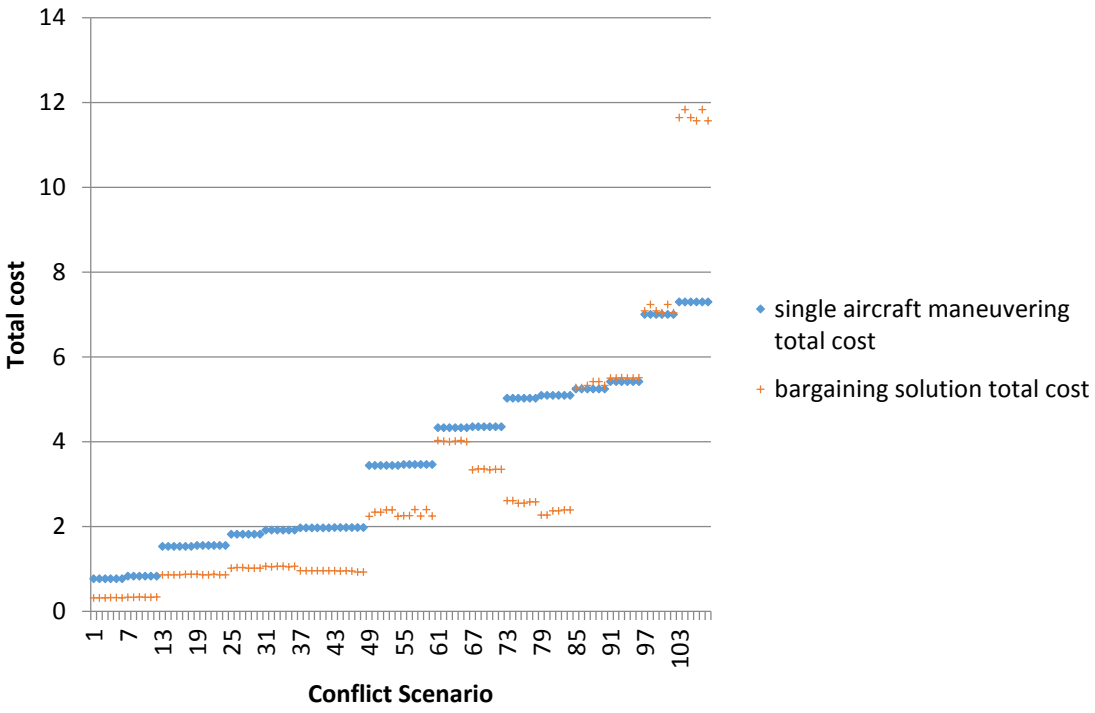

FIGURE 39: TOTAL COST COMPARISON ACROSS CONFLICT SCENARIOS USING THE CLIPPED FEASIBLE SET

Figure 39 shows that, for about 80% of the runs, the bargaining process was more efficient than the random choice of one aircraft doing all the maneuvers. Half of the conflict situations where the single aircraft maneuvering solution is more efficient than the bargaining solution don't display any significant difference in total cost. However the remaining cases show a strong difference where the one-aircraft-maneuvering solutions seem to be much more efficient in terms of total cost. These cases are those where performance constraints impact the bargaining process, such as the example case shown in Figure 40 in which the bargaining solution converged on a resolution requiring roughly twice the cost of either aircraft maneuvering alone. Across all the conflict scenarios, the bargaining solution is on average about 16% more efficient than the single aircraft maneuvering solution in terms of total cost.

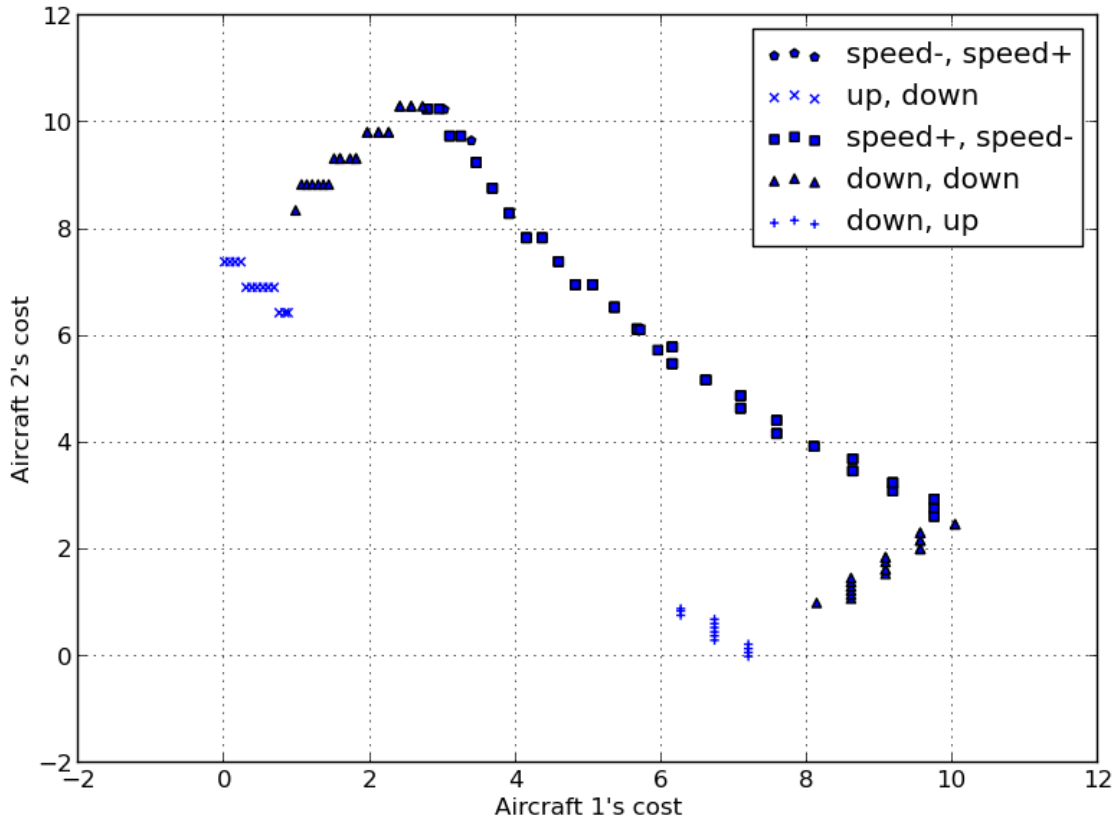

FIGURE 40: EXAMPLE OF CASE FOR WHICH THE SINGLE MANEUVERING AICRAFT SOLUTION HAS A LOWER TOTAL COST THAN THE BARGAINING SOLUTION. RULE A, COST INDICES 10% AND 10%, TYPES A320 & A320, CLIPPED FEASIBLE SET. FULL SPACE, N = 100

5.4 Summarized Results and Observations

No resolution was selected in multiple dimensions across all of the 324 runs. This is likely to be interpreted as the fact that, when a personal trajectory is offered in a given dimension, the space freed in the opposite direction in that same plane allows the other aircraft to find lower cost response trajectories in that particular direction. Therefore it is sufficient to only look for response trajectories in the same dimension as the personal trajectory considered. This means that a partial search of the resolution space leads to the same results as the full search where response trajectories are computed in all the dimensions. Instead of computing six response trajectories per offered trajectory, only two

are considered, increasing the computational efficiency by a factor of three for the exact same outcome.

The representation of performance constraints has a real impact on the bargaining solution only when a trajectory lies close enough to the performance limits, meaning that the different representations of performance constraints will yield different agreement points. In this case, penalizing the constraints by using the finite or infinite cost functions for the bargaining process will lead to solutions further away from the performance limits. This has one benefit and one drawback: it will yield increased safety margins but will provide negotiated trajectories that have, in general, higher total costs than the ones obtained using the clipped feasible set. Therefore it could be interesting for the operational community to assess the cost of flying too close to the performance limits to define a finely tuned penalty cost for the performance constraints. This could be done, for example, by trying to quantify the additional costs incurred by 'flying not safely'.

Looking at total cost, on average, the bargaining process led to more efficient trajectories than single aircraft maneuvering conflict resolutions. However, the conflict situations in which the single aircraft maneuvering solutions are more efficient than the bargaining solution could be more frequent in real operations. Thus it would be interesting to confront this bargaining process with a more realistic situation and make the same efficiency analysis. This is the purpose of CHAPTER 6 where a large scale simulation is performed. This also allows for the study of additional problems such as downstream conflicts that can only arise when more than two aircraft are considered.

CHAPTER 6: LARGE SCALE SIMULATION

This chapter describes the demonstration of the bargaining process in a large scale simulation that represents more 'realistic' conditions. Based on the conclusions drawn in 5.4, the resolution trajectories were only computed in the same dimension as the given personal trajectories, drastically improving the computational efficiency of the bargaining process. In addition, the performance constraints were represented by clipping the feasible set without modifying the cost function.

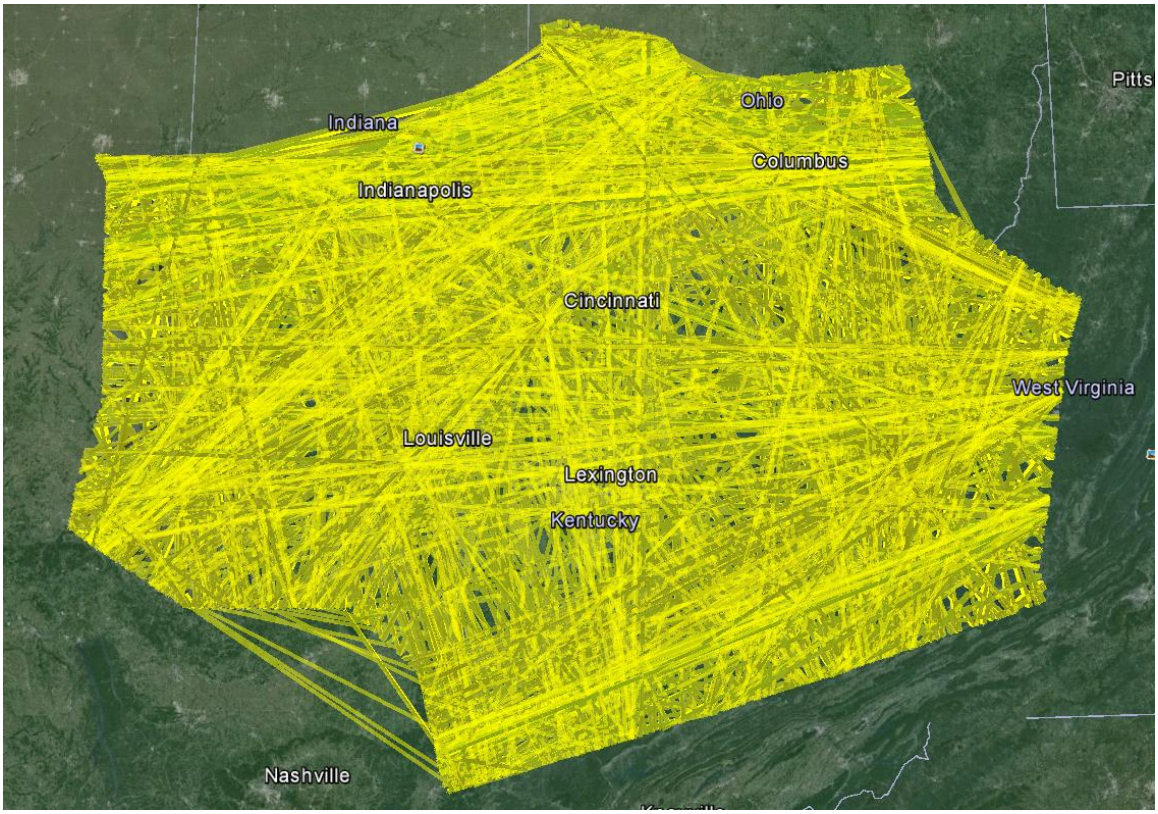

FIGURE 41: ALL THE TRAJECTORIES COMPUTED DURING THE SIMULATION VISUALIZED ON GOOGLE EARTH. THE BOUNDARIES OF THE INDIANAPOLIS CENTER WERE NOT DRAWN BUT ARE NATURALLY DELIMITED.

This large scale experiment simulated about 1200 aircraft as they fly through the Indianapolis Area Control Center. The Indianapolis Center, whose boundaries are depicted on Figure 41, is a facility responsible for controlling aircraft en route in the Indianapolis Flight Information Region (FIR) at high altitudes. The flight plans and aircraft types have been taken from real Enhanced Traffic Management System (ETMS) data over five hours. The initial traffic sets were filtered to only keep aircraft whose start and end waypoints were above 17000 feet, so that the aircraft could be considered as en route. In addition, initial trajectories were optimized to minimize their cost using of the golden ratio optimizer, so that any trajectory modification would increase cost. For some unusual types of aircraft, this process failed, leading to trajectories that the aircraft were not able to. About 300 aircraft were filtered out. Finally, 1184 aircraft and associated optimal flight plans were defined.

The conflict detector was used to maintain the list of all the future conflicts that would have occurred if nothing was done to solve them. Each conflict resolution was then triggered 300 seconds before the conflict and was guaranteeing conflict-free trajectories with a look ahead time of 600 seconds. By comparing the list of all the future conflicts before and after each resolution, the potential downstream conflicts that could be generated were identified and counted.

The goal of this large scale simulation was not to conduct once more the analysis of the different metrics presented in CHAPTER 5 but to gather a little information on concerns that arise when more than two aircraft are involved. Specifically, the risk of generating downstream conflicts is analyzed, as well as the possibility of getting unresolved conflicts. A total cost analysis is also once more conducted.

6.1 Overall Performance of the Bargaining Process

During the simulation, 538 conflicts were detected and a total of 532 were successfully solved using the bargaining process designed in this thesis. Examining the dimension of the resolution, 43% of the conflicts were solved in the horizontal plane, 47% were solved with vertical maneuvers, and about 10 % only were solved using accelerations and decelerations.

Compared to the results of the pairwise conflicts experiments, the fraction of conflicts solved with shifts of the cruising altitudes is substantially higher. The pairwise conflicts were designed so that the aircraft initial trajectories were always at the same altitude. In the large scale experiment however, each aircraft is flying at its optimal altitude according to its own performance model. This increases the likelihood that a conflict occurs where the two involved aircraft are not flying at the same altitude. Therefore, a smaller change of altitudes was often sufficient to resolve the conflict.

Another interesting fact is that most of the negotiated trajectories were found quite far from the performance constraints. About only 4% of the negotiated trajectories were found with less than a 50 knots margin from maximum and minimum speeds. On the other hand, about 2% only of the negotiated trajectories were found with less than a 100 feet margin from the minimum and maximum altitudes. This gives more credit to the representation of performance constraints by clipping the feasible set rather than consistently needing to apply the finite or infinite cost modifications to drive the trajectories away from the performance constraints.

6.2 Downstream Conflicts

The bargaining problem that was designed in this thesis only focuses on solving pairwise conflicts. An immediate issue with such an approach is that solving a conflict between two aircraft may create subsequent "downstream" conflicts with other aircraft. The worst case would be a chain reaction where for example each conflict resolution generates 2 or more additional conflicts that will themselves generate even more and more conflicts.

To detect downstream conflicts, the total number of conflicts detected after a given conflict resolution was compared to the total number of conflicts that occurred before the conflict resolution of the 538 conflicts during the simulation, 11.6% generated downstream conflicts. Then, 1.5% generated more than one downstream conflict.

6.3 Unresolved Conflicts

Examining the six unresolved conflicts, all of them were due to ill conditioning in how their trajectories were defined. In some cases, the aircraft trajectories started with the aircraft in conflict, and in other cases the aircraft trajectories were required to end as they exited the airspace at fixed locations which were so close to the conflict that they limited the resolution. General concerns with both of these issues relate to conflicts occurring near airspace boundaries or near fixed waypoints.

6.4 Total Cost Analysis

As previously done in 5.3.4, a total cost analysis compared the total cost of the bargaining solution to the total cost if one aircraft was maneuvering alone, as shown in Figure 42. The square marks, corresponding to the bargaining solutions are clearly lower than the circle marks, showing that bargaining process is more efficient in most cases. On average, using the bargaining process induces a 29% total cost reduction. The bargaining process appeared to be more efficient about 80% of the time, and only 3% of the conflicts were significantly better handled using single aircraft maneuvering, with more than a 10% total cost reduction.

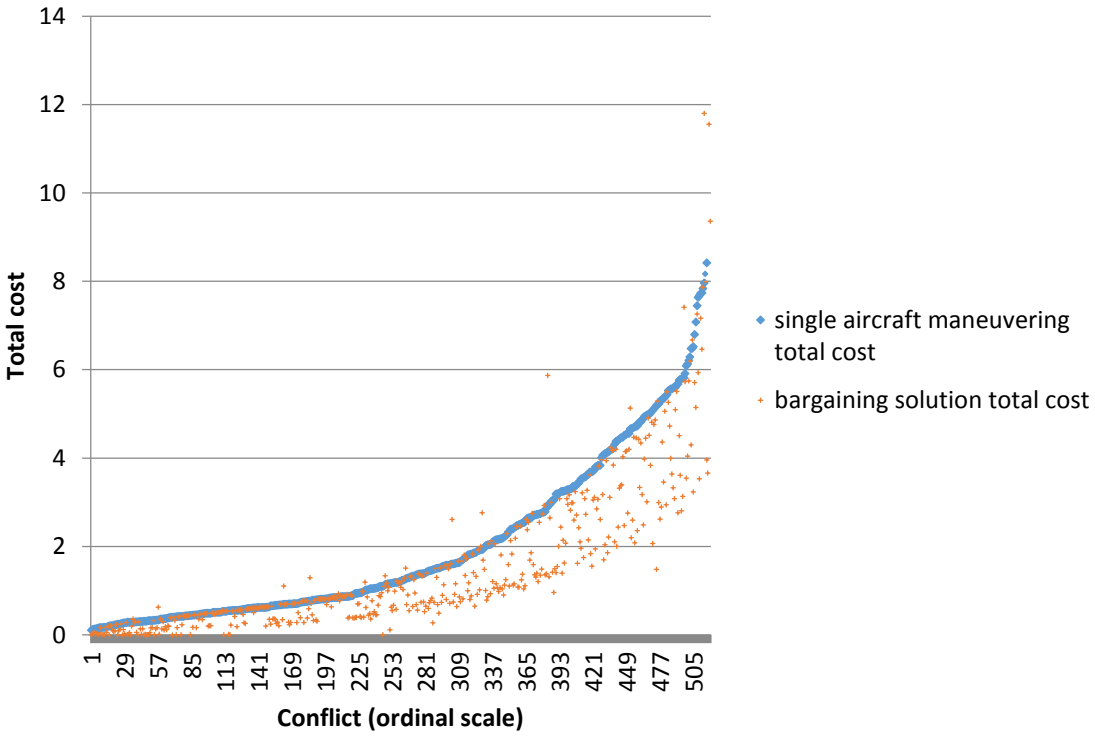

FIGURE 42: TOTAL COST COMPARISON ACROSS ALL OCCURING CONFLICTS IN THE LARGE SCALE EXPERIMENT

6.5 Summary and Observations

This large scale experiment assessed the performance of the bargaining process in a more realistic scenario. It was shown that the negotiating process was able to solve all the conflicts that occurred, apart from only six ill-conditioned conflicts, and the resolutions were significantly more efficient than single aircraft maneuvering conflict resolutions.

Since the resolutions only addressed an immediate pairwise conflict, the problem of downstream conflicts was analyzed. In this airspace, 10.1% of the resolutions generated a single downstream conflict, and a further 1.5% generated two or more downstream conflicts. Thus, this method of conflict resolution was – in this scenario – stable in terms of reducing the number of conflicts. Of course, other airspace configurations and different traffic densities may respond differently and would warrant examination.

CHAPTER 7: CONCLUSION

7.1 Summary

Starting from the observation that new technologies such as ADS-B and GPS can enable decentralized air traffic of operations designs, this thesis aimed at constructing a decentralized conflict resolution as a bargaining process building on game theory.

The first challenge addressed by this bargaining process is the fact that aircraft are not willing to share crucial information required to evaluate the costs of the trajectories. For example, aircraft don't communicate their cost index, which is however central in the cost calculation of a resolution trajectory. Thus, this pairwise sequential bargaining process was designed so that aircraft negotiate the final trajectories they will fly to avoid the conflict. Iteratively, they are required to compromise by offering trajectories of progressively greater cost to themselves until they find an agreement. The complexity of this design develops from the high dimensionality of the resolution space and the performance constraints of the aircraft limiting the feasible set of feasible resolutions.

The problems of performance constraints and multidimensionality of the resolution space were explored by two design variables in the bargaining process. The first design variable examined different methods of representing the performance constraints through the use of different cost functions or by clipping the feasible set of resolutions. The second design variable evaluated two different ways of exploring the dimensionality of the resolutions.

The pairwise conflicts experiment assessed the impacts of the two design variables on the outcome of the bargaining process. The main conclusions were:

- The negotiated trajectories were always in the same dimensions. Hence a partial search of the resolution space for only trajectories that are in the same dimensions is enough. This increases the computational efficiency of the bargaining process without losing performance.
- Representing performance constraints by modified cost functions that add costs with proximity to the flight envelopes of the aircraft created resolutions that are further from these constraints. However the total cost of these solutions increased.

This study was followed by a proof-of-concept experiment: a large scale experiment in which more than a thousand aircraft were simulated as they flew over the Indianapolis Center. The traffic sets were taken from real data and the performance of the aircraft was modelled as realistically as possible.

The overall performance of the bargaining process in this large scale experiment was given along with the analysis of more operational issues including downstream conflicts and unresolved conflicts. The main conclusions and observations drawn by this demonstration were the following:

- The bargaining process was quite effective in the sense that all the conflicts were resolved except for six ill-conditioned cases.
- Downstream conflicts were generated, but “chain reactions” were not generated.

- All the dimensions were used for the resolutions, but horizontal and vertical maneuvers were in general preferred to speed maneuvers.
- The comparison with single maneuvering aircraft solutions showed that the bargaining was significantly more efficient.

7.2 Contributions

This thesis laid down the foundations of a sequential bargaining process to solve pairwise conflicts in a decentralized air traffic concept of operations. Different specific issues related to conflict resolutions such as the representation of performance constraints of the aircraft or the high dimensionality of the resolution space were addressed and studied through a series of pairwise conflicts experiments.

This sequential bargaining process was demonstrated on a fairly realistic simulation given its large scale, where data came as much as possible from the real world. This study explored operational issues such as downstream conflicts. A comparison with single maneuvering aircraft resolutions was performed using total cost metrics, and showed that the bargaining process led to a substantial increase of efficiency.

7.3 Future Work

Several further aspects of the bargaining process warrant further research. First, the representation of performance constraints had significant impacts on the negotiation and resulting total cost of the resolution. It could be therefore interesting to refine as much as possible these representations. For example, cost barriers could be added only for specific performance limits considered to be critical, or tailored to reflect the severity of the constraint. Further, the constraints could be expanded to reflect not just the aircraft

performance, but also operational factors such as no-fly zones or limits imposed to mitigate downstream conflicts.

Another aspect worth investigating is the game theory optimization of private parameters such as cost indices. Let's consider, for example, two airlines with different private parameters settings. Airlines could potentially choose their private information to create asymmetries in the negotiations to their benefit. Airlines could therefore try to infuse the asymmetry of the conflict situations to get cheaper resolution trajectories.

The bargaining process was tested in this thesis in a large scale experiment that was meant to be closer to reality than simple pairwise conflicts. However, it should also be tested in other sorts of airspace and different traffic densities. Such a study could identify conditions that might lead to instability in terms of downstream conflicts.

The bargaining process could also be further optimized to be more computational efficient. A smaller number of cost increments, i.e., n, could for example be assumed at the first place, and augmented in case of an overshoot in the solution. A simple checker could examine for this overshoot.

A last interesting point would be to try to bridge the gap between the strategic approach that we used to build the bargaining process and the existing axiomatic bargaining solutions developed in the literature. Starting from the bargaining process designed in this thesis, this would require formal proof that the several axioms listed in 2.6 are verified, at least for an identified reduced set of conflicts, and then demonstration that the bargaining process defined does converge towards one of the defined axiomatic bargaining solutions.

REFERENCES

[1] Azarm, "Conflict-Free Motion of Multiple Mobile Robots Based on Decentralized Motion Planning and Negotiation," in *IEEE International Conference on Robotics and Automation*, Albuquerque, New Mexico, 1997.

[2] Krozel, "Free flight conflict detection and resolution analysis," in *AIAA, Guidance, Navigation and Control Conference*, San Diego, CA, 1996.

[3] Durand, "Optimal Resolution of En Route Conflicts," Laboratoire d'Optimisation Globale, Toulouse, France.

[4] Kuchar, "A Review of Conflict Detection and Resolution Modeling Methods," *IEEE Transactions on intelligent transportation systems, Vol. 1, No. 4*, pp. 179-189, 2000.

[5] J. John F. Nash, "The Bargaining Problem," *Econometrica, Vol. 18 No. 2*, pp. 155-162, 1950.

[6] H. Erzberger and W. Nedell, "Design of automated system for management of arrival traffic," NASA Ames Research Center, Moffett Field, CA, 1989.

[7] D. Sislak, P. Volf and M. Pechoucek, "Agent-Based Cooperative Decentralized Airplane-Collision Avoidance," *IEEE Transactions on intelligent transportation systems*, vol. 12, no. 1, pp. 36-46, 2011.

[8] S. Wollkind, J. Valasek and T. Ioerger, "Automated Conflict Resolution for Air Traffic Management Using Cooperative Multiagen Negotiation," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Providence, Rhode Island, 2004.

[9] G. Zlotkin and J. Rosenschein, "Negotiation and Task Sharing Among Autonomous Agents in Cooperative Domains," in *IJCAI*, Detroit, Michigan, 1989.

[10] Fatima, "Bargaining with incomplete information," *Annals of Mathematics and Artificial Intelligence 44*, pp. 207-232, 2005.

[11] Cramton, "Sequential bargaining mechanisms," in *Game Theoretic Models of Bargaining*, Cambridge University Press, 1985, pp. Chapter 8, 149-179.

[12] Kalai, "Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons," *Econometrica, Vol. 45, No. 7,* pp. 1623-1630, 1977.

[13] Kalai, "Other Solutions to Nash's Bargaining Problem," *Econometrica, Vol. 43, No. 3,* pp. 513-518, 1975.

[14] A. Rubinstein, "Perfect Equilibrium in a Bargaining Model," *Econometrica,* vol. 50, no. 1, pp. 97-109, 1982.

[15] M. Bigelow, "Examining the relative costs and benefits of shifting the locus of control in a novel air traffic management environment via multi-agent dynamic analysis and simulation.," Georgia Institute of Technology, Atlanta, GA, 2011.

[16] A. Nuic, "User manual for the base of aircraft data (bada) revision 3.9. Technical report," Eurocontrol, 2008.

[17] Andrews, "IPC Design Validation and Flight Testing Final Report," Lincoln Laboratory, Massuchusetts Institute of Technology, Lexington, Massachusetts, 1978.

[18] Ausubel, "Efficient Sequential Bargaining," *Review of Economic Studies 60,* pp. 435-461, 1993.