

Reduction from multi-aircraft to single-aircraft scenario

We will hereafter focus on the problem of estimating and/or predicting the position of a single aircraft that does not interact with any potential surrounding traffic. This approach is justified by the evolution of two GA aircraft being only weakly coupled. Indeed, they can in principle operate freely and are solely responsible for determining their own flight path.

Once that problem is solved for a single aircraft, the information about the current and/or future location of each aircraft can be broadcasted to each agent, enabling the electronic visibility of current and future traffic configurations, i.e. traffic-monitoring and traffic-prediction functionalities. However, the last step is not covered hereby since it involves considerations about the system architecture that are still under investigation within the ProGA project.

C. Typical strategies for tackling the single-aircraft prediction problem, state-based predictions

Stochastic Filters (SFs) are often employed in applications demanding estimates and/or predictions of unknown quantities that may evolve over time. In the case of an aircraft, the set of those quantities generally comprehend, but is not limited to, aircraft position and speed. In other words, SFs can be used to obtain estimates/predictions of the aircraft kinematic state.

SFs return estimates/predictions in the form of a Probability Density Function (PDF) on the feasible values of the aircraft kinematic state. A SF is an iterative update scheme: the PDF describing the aircraft kinematic state at time t is projected some steps ahead in time through a *dynamical model* (a stochastic equation for the aircraft motion) to obtain the so-called *prior*; next, an *observation model* (an equation for the noisy measurements of the aircraft kinematic state) is used along Bayes' Theorem [4]-[7] to update the prior and obtain the so-called *posterior*. Fig. 1 outlines the general design of a SF.

In a SF the dynamical model is appointed to keep the last available PDF and the next measurement of the aircraft kinematic state up-to-date; technically, the filter prediction is the prior obtained this way. However, the process of projecting ahead in time a PDF will generally spoil its qualities of smoothness or sharpness. Therefore, the prior must be further processed to try and make it as smooth and sharp as possible. That operation is performed via Bayes' rule and a measurement of the aircraft kinematic state. Since in principle that measurement may be affected by some error (noise) and/or not cover all the components of the kinematic state, the method is completed by an observation model to clear out the noise and/or compensate the incompleteness of the measurement.

As previously explained, SFs work according to the following paradigm: PREDICTION > PRIOR > UPDATE > POSTERIOR. The prediction phase computes the prior based on a dynamical model of the flight evolution and the last computed posterior. Since the current kinematic state of the aircraft is in general the only ingredient of the dynamical model, we will hereafter refer to those predictions as *state-based predictions*.

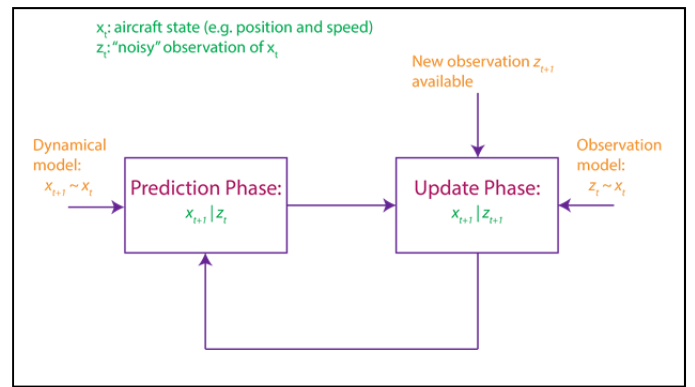


Figure 1. Iterative update scheme of a SF.

D. State-based estimates/prediction in air traffic applications

The large majority of research papers about SFs applied in air transportation contexts are dedicated to commercial aviation. On the one hand, this is symptomatic of the scarce attention the scientific community has paid so far to the GA domain; on the other, it suggests that the characteristics of GA flights make the prediction of their future evolution a very hard task to perform.

Sequential Monte Carlo techniques are SFs extensively used in air traffic applications to address the non-linearity of the aircraft dynamics, see [8] and [9]-[12]. A very interesting algorithm for multi-aircraft trajectory prediction is the so-called Sequential Conditional Particle Filter (SCPf), introduced in [13] and [14], which is able to provide very accurate estimates of both aircraft positions and global wind field.

None of the filters previously referenced can be used in light GA contexts, because they assume an aircraft dynamics specific to commercial flights – for example, SCPf includes the action of a high-precision flight management system. To the best of our knowledge, the only application of stochastic filtering to the trajectory prediction of a light aircraft is [15].

III. FROM STATE-BASED TO INTENT-BASED PREDICTIONS

A. Time limitations to state-based predictions

In Sec. II.B we have illustrated the benefits that large prediction timeframes can have for the GA community. However, it is quite challenging to infer the position of an aircraft over time intervals of such an extension.

As mentioned previously, the output of a state-based prediction will not be in general a smooth and sharp PDF over the aircraft kinematic states. As the prediction timeframe increases, the prior will rather spread over a larger and larger region of kinematic states. In other words, the uncertainty that affected the previous estimation of kinematic state will propagate over successive time steps and grow ever larger.

Figure 2. shows an aircraft executing a manoeuvre, the trajectory it is expected to follow in the near future, and the uncertainty that affects the predicted trajectory. More in detail, Figure 2. assumes as in [15] that the aircraft is flying at (nearly) constant altitude and that its position follows a

Gaussian law; the sequence of means and 3σ -regions at successive steps are represented by solid and dotted lines, respectively.

Let us now suppose an aircraft is flying at fixed altitude and with constant speed; suppose that both position and speed are known with infinite precision so that the only uncertainties about the future evolution of the flight are represented by the potential manoeuvres that the pilot may perform. In absence of further information, the best assumption is that the aircraft motion will not change, because for each manoeuvre there exists another one that produces the opposite effect, and there is no reason to infer that a manoeuvre is more probable than its opposite. However, the pilot is actually bound to perform at least one manoeuvre somewhere in the future, i.e. when he/she will eventually land.

Although for short time-periods it is perfectly legitimate to assume that the aircraft will keep on flying in the same way it has done in the recent past, for large time-horizons the very same assumption is likely to produce large errors. Figure 2. sketches the situation: the true future evolution of the aircraft is drawn in purple, whereas green and orange represent two trajectory predictions casted under the assumption that the aircraft kinematic-state will not change. As expected, the error, i.e. the distance between true and forecasted position, increases with the size of the prediction timeframe.

All in all, state-based predictions may prove themselves quite unsatisfactory if casted over a large timeframe. Focusing only on the expected future location of the aircraft, we may experience that true and expected position are quite distant from each other; conversely, if we consider the uncertainty that affects the forecasted position, we may obtain a volume occupied by the aircraft with high probability, e.g. 90%, that has the unfortunate property of being so large to be practically meaningless.

B. Extending the prediction timeframe, intent-based predictions

In the previous subsection we have explained the time limitations to state-based predictions. In particular, we have highlighted how casting a prediction over an increasing time horizon may quickly lead to large errors, or, if volumes are considered to compensate the increasing uncertainty, to expected occupied-volumes that are so large to be meaningless. Both issues do not arise from any inadequacy of the dynamical model used in state-based predictions, but rather from the high uncertainty that is introduced by the complete lack of information about the pilot intent.

The motion of a flying aircraft can be described with different levels of precision, e.g. by a simple coordinated-turn model or as the motion of a three-dimensional object in a fluid (Navier-Stokes equations). Complex dynamical models may lead to more precise estimates/short-term predictions of the flight because they can capture more characteristics of the aircraft physical motion. However, all those details are like noise added to a sort of flight plan when the whole duration of the flight is considered.




Figure 2. True future trajectory of the aircraft (top) and state-based prediction casted over short and long timeframes (resp. middle and bottom).

Shifting the focus from short-term to long-term predictions, the aircraft dynamics becomes less and less significant while the pilot's intent gains importance. If observed over a short time-scale, the evolution of a flight looks in fact like driven by its kinematic state; on the contrary, over a large time-scale the pilot's input appears as the dominant contribution to the evolution of the flight.

Earlier stages of ProGA project demonstrated how planning a flight in advance is a common habit among the GA community when flying VFR in uncontrolled airspace. In particular, what emerged from the discussion of distinctive features common to all GA flown paths is that, at least during the planning phase, a GA path is a "sequence of straight legs", see [16]. Due to external factors as wind, air traffic, or visibility conditions, the realisation of a flight can be in fact quite different from the initially planned route, which will be hereafter called *flight intent*.

The existence of flight intents is of crucial importance to the design of a system that could cast predictions about the flight execution over a large timeframe. Indeed, let us imagine implementing a SF where the dynamical model embeds a notion of flight intent. In this respect, one component of the aircraft kinematic state is the flight intent, and we no longer speak of state-based predictions but of *intent-based predictions*. We expect that intent-based predictions will suffer less from the issues detailed in Sec. III.A, allowing an extension of the prediction timeframe.

C. Flight intent, declarations vs. estimation

Casting intent-based predictions requires knowledge of the flight intent an aircraft is following. That information can be either known in advance (e.g. if the pilot compiles and shares it), but it may as well be not available (e.g. if the pilot compiles but does not share it). Thus, we need to find a surrogate to the intent whenever the pilot does not make it available. Yet the flight intent may be available but prove itself not trustworthy (e.g. the pilot partially or totally disregards it, either prior to take-off or during the flight execution). Therefore, we also need a way to check whether the declared intent, if any, is in accordance with the kinematic state of the aircraft.

In our view, the missing information of non-declared flight intent can be replaced by a statistical characterisation of the paths typically flown in a given area by GA aircraft, see Sec. IV. Checking whether an aircraft is following a given flight intent (or its statistically-derived counterpart) is performed within a Bayesian framework, see Sec. V.A. The performance of the new prediction concept will be studied through numerical experiments in Sec. V.B.

We would like to stress that estimating the flight intent is a very difficult problem, highly dynamic and complex at the same time. Moreover, the flight intent is planned by the pilot in such a way to connect special geographical points and/or elements (e.g. the location of VORs as well as the direction defined by rivers, motorways, etc.); therefore, the flight intent is an aspect of a GA flight that is strongly local, in the sense that it is highly dependent on the geography of the area overflown. Thus, the idea of attaining the flight intent through a statistical survey of the typical GA local paths may prove itself a valid trade-off between the complexity and the efficacy of the overall intent-based prediction-concept.

IV. FLIGHT INTENT AND GA RECURRING PATTERNS

A. Recurring patterns in GA

In Sec. III.C we have introduced the idea of replacing the flight intent with a statistically derived analogue whenever the former is not available or cannot be relied upon. Where does that idea come from? Ref. [16] shows an analysis of the way GA flights are typically conducted and, in particular, it suggests the existence of recurring GA patterns. Such patterns are the consequence of the common navigation methods that any pilot learns when training in preparation for getting his Private Pilot License [22]. These definite paths, typically flown by a local GA community, arise as a direct consequence of the airspace structure and the location of the reference elements (VORs, geographic landmarks, rivers, roads, etc.) commonly used by pilots to outline the flight intent and help the navigation.

B. Statistical database of recurring patterns

The relationship between flight intents and GA recurring patterns is twofold. On the one hand, any flight intent can be thought of as a recurring pattern because it is planned following the very same criteria that let recurring behaviour arise. On the other hand, recurring pattern can be regarded as flight intents because, if no information is available about the pilot intentions then, statistically speaking, the best assumption about the evolution of a flight is that it will follow one of the paths commonly flown by the local GA community.

The idea underlying ProGA is then the following: if we manage to acquire sufficiently many flight traces recorded by a Global Navigation Satellite System (GNSS) then we may build up a statistical description of GA recurring patterns in a given area. That statistical description is a set of couples (t, ϵ) , where t is a reference trace, i.e. a statistically inferred analogue of a flight intent, and ϵ is a measure of how much a flight trace typically deviates from t ; Sec. IV.C will show a potential

method to obtain such a statistical description of GA recurring patterns.

Once the statistical database of known GA patterns is available, the reference traces are a valid surrogate of the flight intent whenever the latter is not made available or cannot be relied upon. Moreover, should the deployment of ProGA introduce a flight-trace recording functionality, the capability of the statistical database to describe recurring GA patterns could be continually improved and increased. However, this last point may present some privacy and acceptance issues, so it is still under investigation.

C. Statistical analysis of recurring patterns

This section briefly introduces some statistical techniques that may lead to the creation of a statistical database out of a set of GNSS traces. The idea is based on a clustering of the observed trajectories, and [17] applies the technique to radar tracks in a terminal radar approach control. Unfortunately, the scarceness of available GA flown traces makes not possible any quantitative validation of the method illustrated hereafter.

Let us consider the case study of a GNSS trace flown in France between Saint-Cyr-l'École (LFPZ) and Bernay (LFPD). The cruise altitude is nearly constant throughout the flight, so the following analysis is carried out as if the aircraft motion occurred in a plane. Figure 3. presents a change-point analysis [18] of the aircraft track (measured in radians with respect to the East) over the whole flight duration; the red horizontal line has the meaning of average aircraft track, thus discontinuities in this line are associated to turning points. Superimposed to the initial path, the position of the turning points and the straight legs connecting them are depicted in Figure 4. . The last analysis can be generalised to the case of an aircraft not flying at constant altitude using multivariate¹ change-point techniques, e.g. [19]-[20].

Once the GNSS trace have been processed like in the example above, it is possible to classify a whole dataset as follows. After running a clustering analysis [21] of the turning points, each class is formed by the traces that can be represented as a sequence of turning points that fall within the same clusters. Next, the representative of each class, i.e. the reference trace, can be chosen, for example, as the polyline constructed by joining the centroids of consecutive clusters of turning points. Further, by measuring the distance between the elements in each class and the class representative, a probabilistic characterisation of the typical deviations from the reference trace can be derived.

The classification above is a way to produce a statistical description of GA recurring patterns, i.e. conduits of airspace that are likely to be flown in. The classification procedure may be refined to include as many factors as desired, e.g. period of the day, season, wind conditions, traffic density, etc.; it is more than reasonable to expect that such a refinement should improve the statistical description of GA recurring patterns.

¹ It is indeed well known that two angles are sufficient to characterise an arbitrary direction in the three-dimensional space.

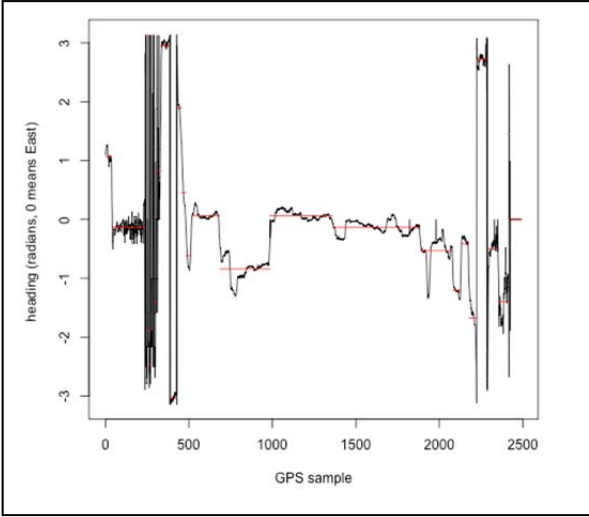


Figure 3. GNSS trace used as a case study.




Figure 4. Estimation of the flight intent resulting from the change-point analysis.

V. PROGA INTENT-BASED PREDICTIONS

A. Prediction concept

ProGA introduces a modular prediction concept, which combines all the elements discussed above, namely, a statistical database of GA recurring patterns, a unit for casting state-based predictions, and a Bayesian estimator of the flight intent. The latter deserves some more elaboration.

As mentioned in Sec. III.B, we consider a kinematic state that is composed by the aircraft position, speed, and flight intent. The flight intent is not bounded, in fact, to be constant over time, so it makes sense to include it in the aircraft state. As a consequence, the PDF describing the aircraft state is a joint distribution $P(x_k, t_k)$, where x_k and t_k are the proper kinematic state (position and speed) and the flight intent at time k , respectively. Please note that the variable t_k takes values on the union of the declared flight intent, if any, and the recurring GA patterns contained in the statistical database we have discussed in Sec. IV.

Let z_k be the observation of the aircraft kinematic state, x_k (e.g. a GNSS signal carrying the aircraft position and speed). The formula of conditional probability yields

$$(1) \quad P(x_k, t_k | z_k) = P(x_k | z_k) P(t_k | x_k, z_k).$$

Since the observation z_k is a function of the true state x_k ,

$$P(t_k | x_k, z_k) = P(t_k | z_k).$$

Therefore,

$$(2) \quad P(x_k, t_k | z_k) = P(x_k | z_k) P(t_k | z_k).$$

The importance of (2) is that $P(x_k | z_k)$ is the posterior at time k returned by a SF like the one presented in [15]. The factor $P(t_k | z_k)$ is computed using

$$(3) \quad P(t_k | z_k) = \frac{P(z_k | t_k) P(t_{k-1} | z_{k-1})}{P(z_k)}$$

and the model

$$(4) \quad P(z_k | t_k) = \frac{\exp[-\beta_1 \Delta_k - \beta_2 \Theta_k]}{Z},$$

where $\beta_{1,2}$ are suitable constants, Z is a normalisation factor, Δ_k is the Euclidean distance between the aircraft observed position and t_k , and Θ_k is the angle between the aircraft observed track and the leg of t_k that is closest to the aircraft position.

Model (4) penalise those flight intents that are *too distant* from the current kinematic state by assigning them an exponentially small weight. In other words, if the aircraft position is too far away from a potential flight intent t_k , or if the track of the aircraft and that of the intent are diverging, model (4) will consider x_k too large a deviation from the intended trajectory t_k and penalise the event that the aircraft is actually following t_k .

Figure 6. summarises the ProGA prediction concept. GNSS traces of flown paths are acquired and processed offline as discussed in IV.B-C; the result is a statistical characterisation of GA recurring patterns. During the in-flight phase, the unit deputed of casting state-based predictions takes care of continually estimating the aircraft position and speed by filtering a signal-in of GNSS type. Since $P(z_k)$ in (3) can be treated as a normalisation factor, it is possible to sample from the joint prior $P(x_k, t_k | z_k)$ using (2)-(4). Each sample can be propagated ahead for as many steps as desired by the dynamical model outlined in Figure 5. and detailed below. If sufficiently many samples are drawn and simulated, the expected future trajectory can be reconstructed via Monte Carlo techniques.

The dynamical model used to simulate the aircraft motion embeds the flight intent, which drives the evolution of the flight. At time k , let p_k , q_k , and t_k be respectively the aircraft position and speed (three dimensional vectors), and the flight intent (a three dimensional polyline); let the flight intent t_k be a sequence of adjacent segments, and let w_k be the segment composing t_k that minimise the distance from the aircraft position, p_k ; finally, let \hat{w}_k be the unit length vector aligned with the segment w_k .

The aircraft motion is described by the following discrete-time model:

$$\begin{aligned} p_{k+1} &= p_k + q_k + v_k \\ q_{k+1} &= |q_k|(1 + \omega_k) \cdot \hat{w}_k \end{aligned}$$

where v_k is a Gaussian perturbation acting on the plane orthogonal to \hat{w}_k , $|q_k|$ is the module of vector q_k , and ω_k is a zero-mean Gaussian noise.

The evolution of the aircraft position, p_k , and role played by the term v_k are illustrated in Figure 5. . The velocity q_k produces a reference displacement that is in accordance with the flight intent, in the sense that it is along the direction of the closest leg composing the intent. The length of the reference displacement is variable because of the presence of the Gaussian noise ω_k , which models small intrinsic fluctuations in the cruise speed. The reference displacement is summed to a Gaussian perturbation, acting orthogonally to the displacement itself, which models natural fluctuations from the originally intended trajectory.

B. Numerical Experiments

In this subsection we present a preliminary assessment of the ProGA prediction concept through some numerical evidences. The experiments presented in this section focus on

1. intent-based predictions when the flight intent is declared and actually flown;
2. the capability of equation (3) above to
 - a. discover unknown flight intents;
 - b. cope with changes in the intent being flown.

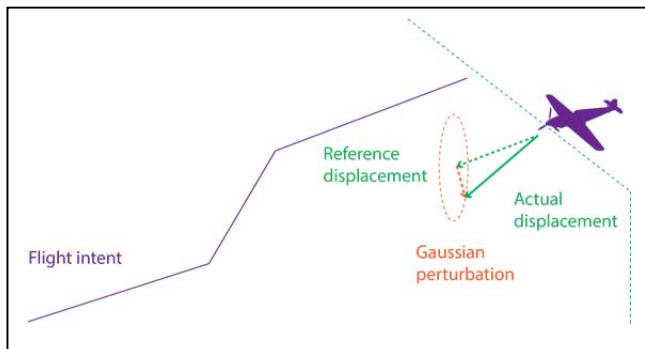


Figure 5. Intent-based prediction: dynamical model.

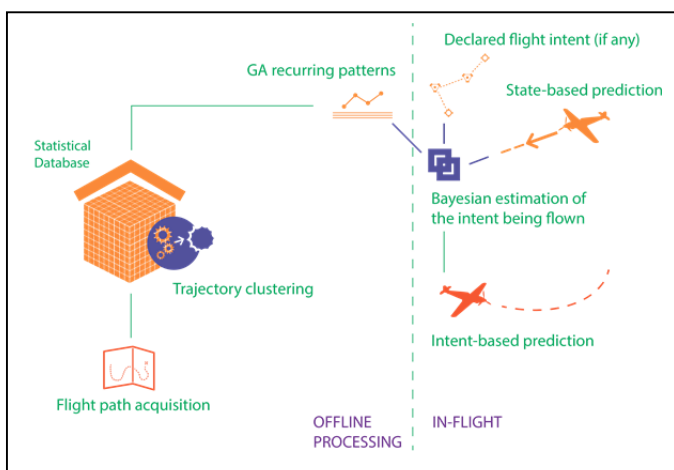


Figure 6. Outline of ProGA prediction concept.

Figure 7. shows a GNSS trace from a flight that took place in the South of Belgium between Courtrai and Liège. An expert GA pilot has estimated the flight intent used for the experiment through operative considerations. Figure 8. displays the outcome of the prediction algorithm when initialised with a state picked at random from the trace in Figure 7. . The path flown is drawn in black; the yellow point marks the starting position at time t_0 ; the particle clouds in green, red, cyan, and magenta shows the volume that is expected to be occupied by the aircraft after 5, 10, 15, and 20 minutes, respectively.

A thick magenta dot, partially hidden by the cloud of particles, marks the true position of the aircraft after 20 mins. The green, red, and cyan clouds hide the corresponding true position, which is well inside the forecasted volume. Repetitions of the experiment outline that the probability of the true position being far from the centroid of the forecasted volume less than 1 km is about 50%. Moreover, *the prediction algorithm works well in presence of a turn*, which is a critical issue according to the analysis put forward in Sec. III.A.

Experiments 2a and 2b consider a set of six GNSS traces recorded west of Paris between Saint-Cyr-l'École and Bernay (Figure 9.). Associated to the traces there are four different flight intents: "Ref1", "Ref2", "Ref3", and "Ref4", which refer to the flight intent of the dark red, dark blue, green, and orange traces, respectively. Those intents have been validated by means of the change-point analysis illustrated in Sec. IV.C and through the operative considerations detailed in [16].

The result of experiment 2a is presented in Figure 10. . Here the pilot flies according to intent "Ref3" without declaring it. The estimator is started with uniform knowledge (25% of probability assigned to each flight intent). "Ref2" is immediately penalised being "far" from the flown trace; accordingly, the probability of "Ref2" being the true flight intent quickly vanishes. The remaining flight intents are initially given roughly the same probabilistic weight. After a while, "Ref1" starts diverging from "Ref3" and "Ref4" and the probability of the red intent quickly vanishes as well. The subsequent yellow probability peak is due to a local, transient configuration of "Ref3" and "Ref4" with respect to the flown path. This situation can be partially avoided by a careful tuning of the estimator parameters β_1 and β_2 (see Sec. V.A). Finally, when "Ref4" diverges from "Ref3", the estimation converges to the true value. The experiment hence shows the capability of the estimator to discard with relative ease those flight intents that are obviously unlikely of being flown. At the same time, the behaviour of the estimator is sufficiently conservative not to exclude those flight intents that are similar to the flown path.

Experiment 2b (Figure 11.) shows what happens if an imaginary pilot declares intent "Ref3" but flies "Ref4" instead. Due to the initial knowledge of the flight intent, the prior is set to 97% of probability mass to "Ref3" and 1% of probability mass to the remaining intents. As at the beginning "Ref3" and "Ref4" are very close and run parallel to each other, the estimator gives high confidence to the hypothesis that the pilot

is really flying “Ref3”. However, as soon as the flight path gets away from “Ref3”, the algorithm quickly catches the change and adjusts the prediction giving around 100% of probability to “Ref4”, the intent actually flown.

All in all, experiments 2a and 2b demonstrate the capability of the algorithm to cope well with missing or wrong information about the flight intent. Moreover, the estimator is quick in the response and adjusts itself to sudden changes in the flight intent actually followed.

VI. HOW PROGA MAY SUPPORT PILOT’S OPERATIONS

The ProGA prediction concept has the great advantage of being highly modular. Each module displayed by Figure 6. contributes to the prediction of the future evolution of a single-aircraft, the core functionality of ProGA. Sec. II.B has illustrated why it is important for a GA pilot to have knowledge of the likely evolution of nearby flights. Some modules may have standalone uses, though, which we list in the following.

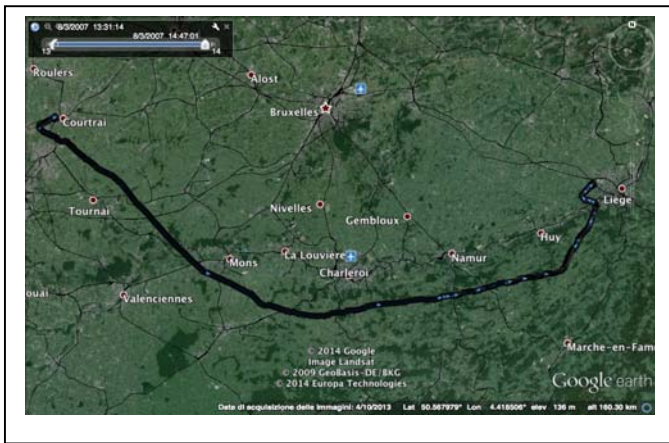


Figure 7. GNSS trace used for experiment 1.




Figure 8. Outcome of the particle filter for the long-term prediction. The yellow dot is the initial position of the aircraft. Four clouds of particles show the forecasted occupancy volume after 5, 10, 15, and 20 (resp. green, red, cyan, and magenta). The true position is always well inside the forecasted volume. Repeated trials show that the median distance between true and expected position (cloud centroid) is about 1 km.

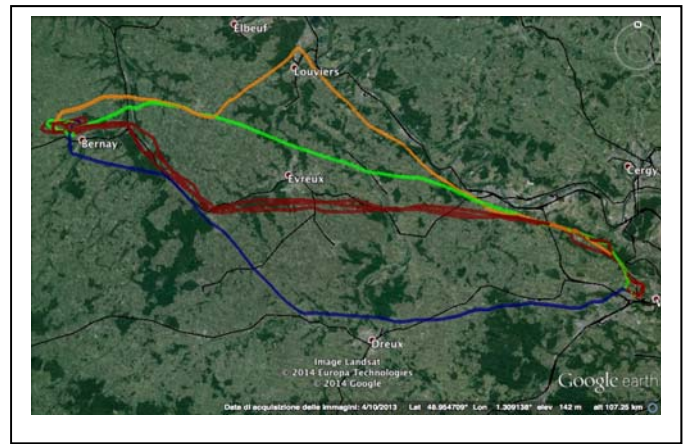


Figure 9. GNSS traces used for experiments 2a and 2b.

With reference to Figure 6. , the unit devoted to casting state-based predictions can be also used to monitor the traffic situations since it continually estimates the kinematic state of each aircraft. The position and track of each aircraft can be therefore employed to bring in the cockpit electronic visibility of surrounding traffic; further, this unit can detect airspace infringements, which may be communicated to the ATC via SWIM or recorded for offline safety investigations. The unit that estimates the flight intent being flown can be set to flag non-typical behaviours, which may be symptomatic of a dangerous situation. Finally, the statistical database and the GA recurring patterns can be queried during the planning phase of the flight to get an idea of which routes are typically flown by the local GA community, where the major traffic flows intersect, and what airspace volumes are typically more congested. In this way the pilot may decide to avoid those hotspots or knowingly fly them; in both cases, this is expected to result in an increased level of safety.

VII. CONCLUSIONS

The present work has detailed some aspects of ProGA, a SESAR WP-E project. At the core of ProGA there is the collection, the processing and the use of flight data to cast long-term predictions about the future evolution of one or more flights. The ProGA prediction concept is highly modular, and the units contributing to the computation may be exploited to support GA pilots both in the planning phase and during the execution of a flight. The paper has also presented a preliminary assessment of the prediction concept realised using real flight data. The experimental results reveal that the prediction concept proposed in this paper is really promising and worth of further investigation in later stages of the project.

ACKNOWLEDGMENT

This work is co-financed by EUROCONTROL acting on behalf of the SESAR Joint Undertaking (the SJU) and the EUROPEAN UNION as part of Work Package E in the SESAR Programme. Opinions expressed in this work reflect the authors' views only and EUROCONTROL and/or the SJU shall not be considered liable for them or for any use that may be made of the information contained herein.

C.L. gratefully thanks R. M. Castro (TU/e) and M. Roth (KNMI/Eurandom) for countless discussions and advice. All authors thank H. Brants (NLR) for sharing the trace in Figure 7. , and A. Ferreira (Deep Blue) for helpful feedbacks.




Figure 10. Experiment 2a: no flight intent declared, "Ref3" actually flown.




Figure 11. Experiment 2b: flight intent "Ref3" declared, "Ref4" flown instead.

REFERENCES

- [1] M. Robinson, *Whatever happened to the Single European Sky?* In Eurocontrol Publication 2013: *the aviation year in Europe*: 71-72, 2012
- [2] IAOPA, *The GA/R need for innovation for the Single European Sky*
- [3] Australian Transport Safety Bureau, *Limitations of the See-and-Avoid Principle*, 1991 (reprinted 2004)
- [4] S. Geisser. *Predictive inference: an introduction*. Vol. 55. CRC Press, 1993.
- [5] Y. Bar-Shalom, X. Rong Li, and T. Kirubarajan. Estimation with applications to tracking and navigation: theory algorithms and software. John Wiley & Sons, 2004.
- [6] M. S. Grewal and P. A. Angus. Kalman filtering: theory and practice using MATLAB. Wiley, 2011.
- [7] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. Signal Processing, IEEE Transactions on, 2002.
- [8] T. Schon, F. Gustafsson, and P-J. Nordlund. Marginalized particle filters for mixed linear/nonlinear state-space models. Signal Processing, IEEE Transactions on, 2005.
- [9] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan. Interacting multiple model methods in target tracking: a survey. IEEE Transactions on Aerospace and Electronic Systems, 1998.
- [10] H. A. P. Blom, E. A. Bloem . Joint IMM-PDA particle filter. *Sixth International Conference of Information Fusion*, Cairns, Australia, 2003.
- [11] H. A. P. Blom, E. A. Bloem. Tracking multiple maneuvering targets from possibly unresolved, missing or false measurements. *Eighth International Conference of Information Fusion*, Philadelphia, 2005.
- [12] X. Lin, T. Kirubarajan, Y. Bar-Shalom Y, S. Maskell. Comparison of EKF, Pseudomeasurement and particle filters for a bearing-only target tracking problem. *Proceedings of SPIE*, Orlando, FL, vol. 4728, 2-4 April 2002; 240-250.
- [13] I. Lymeropoulos and J. Lygeros. Sequential Monte Carlo methods for multi-aircraft trajectory prediction in air traffic management. *International Journal of Adaptive Control and Signal Processing*, 2010.
- [14] I. Lymeropoulos, J. Lygeros. Improved ground trajectory prediction by multi-aircraft track fusion for air traffic control. *AIAA Guidance, Navigation and Control Conference and Exhibit*, Chicago, 2009.
- [15] U. Maeder, M. Morari, T. I. Baumgartner. Trajectory prediction for light aircraft. *Journal of Guidance, Control, and Dynamics*, 2011.
- [16] C. Le Tallec. General aviation aircraft operation and related safety issues, D1.1 of ProGA. Edition 00.01.00, September 2013.
- [17] M. Gariel, A. N. Srivastava, E. Feron. Trajectory clustering and an application to airspace monitoring. *Intelligent Transportation Systems, IEEE Transactions on*, 2011.
- [18] R. Killick, I. A. Eckley. Changepoint: an R package for changepoint analysis. *R package version 0.6*, <http://CRAN.R-project.org/package=changepoint>(2011).
- [19] E. M. Maboudou-Tchao, D. M. Hawkins. Detection of multiple change-points in multivariate data. *Journal of Applied Statistics*, 2013.
- [20] M. Lavielle, G. Teyssiere. Detection of multiple change-points in multivariate time series. *Lithuanian Mathematical Journal*, 2006
- [21] A. K. Jain, M. N. Murty, P. J. Flynn. Data clustering: a review. *ACM computing surveys (CSUR)* 31.3 (1999): 264-323.
- [22] Federal Aviation Administration. *Navigation*, Chapter 15 in *Pilot's Handbook of Aeronautical Knowledge*. Skyhorse Publishing Inc., 2009.