THE CONTINUAL IMPROVEMENT OF SAFETY NETS

by Rod Howell

Over the last twenty years or so, I have had the privilege of helping to improve the performance of numerous ground–based Safety Net systems, through design improvements and optimisations. This article draws on my experiences and explains some crucial analysis activities that should be undertaken to achieve continuous improvement. Following the advice offered here will save you time and money in the long run, and lead to better performing Safety Nets.

The four classic ground-based safety nets are:

Short Term Conflict Alert (STCA) - STCA assists the controller in preventing collision between aircraft by generating, in a timely manner, an alert of a potential or actual infringement of separation minima.

Minimum Safe Altitude Warning (MSAW) - MSAW warns the controller about increased risk of controlled flight into terrain accidents by generating, in a timely manner, an alert of aircraft proximity to terrain or obstacles.

Approach Path Monitor (APM) - APM warns the controller about increased risk of controlled flight into terrain accidents by generating, in a timely manner, an alert of unsafe aircraft altitude during final approach.

Area Proximity Warning (APW) - APW warns the controller about unauthorised penetration of an airspace volume by generating, in a timely manner, an alert of a potential or actual infringement of the required spacing to that airspace volume.

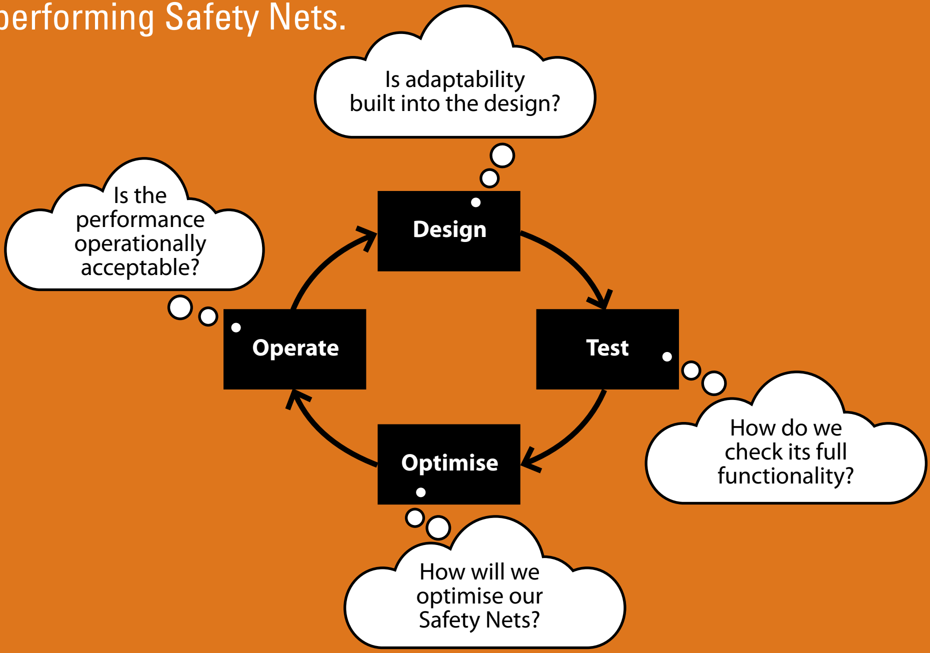

The "Continual Improvement" loop shows how the performance of Safety Nets should evolve – primarily as a result of testing, optimising and operational use of Safety Nets. For senior management, the procurement of a Safety Nets system should be seen not as a one-off event but as the beginning of an ongoing process of adaptation and improvement.

The 'Continual Improvement' loop shows that effective management of a Safety Net system requires regular monitoring of performance and the making improvements to design and parameter optimisations as necessary to maintain operational acceptability. There are numerous questions to ask at each stage of the process, some of which are shown in the diagram.

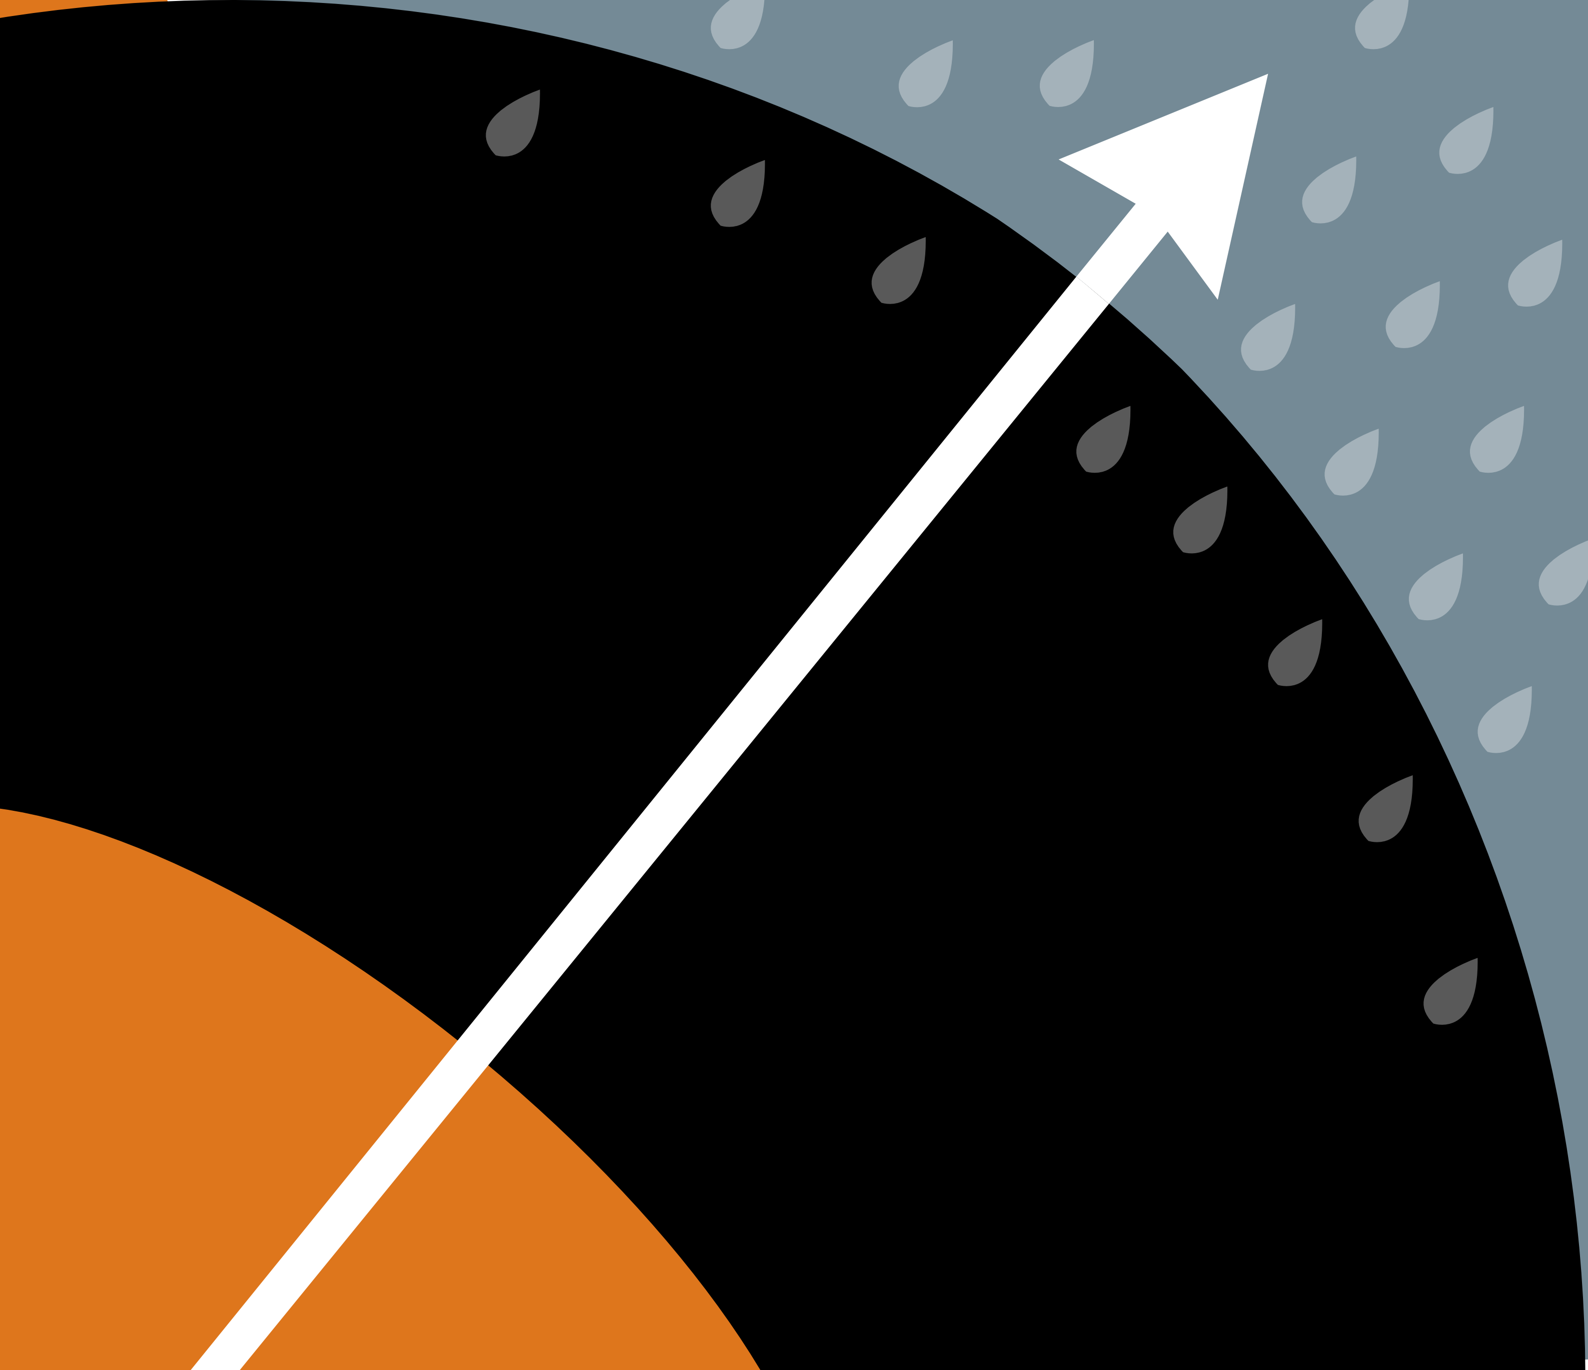

Is the Performance OK?

One thing that has become clear to me is that many ANSPs still do not have a clear understanding of how well a Safety Net system is performing or should perform. Performance figures are not generally widely published. Those that are can be very context specific, making them difficult to apply to somebody else's airspace.

Questions have been raised on numerous occasions over how many nuisance alerts are too many, and how much warning should a Safety Net provide. Precise numbers are impossible to provide because they depend so much on multiple local factors (including traffic levels and complexity, technical systems and HMIs), and the effect of these in concert with Safety Nets alerts on controller performance are not eadily quantifiable.

I therefore make a plea: If you have any concerns about the performance of a Safety Net, make sure that they are raised within your organisation. There are two very good reasons for this: firstly, these systems are there to add an additional layer of safety to Air Traffic Control. An underperforming Safety Net system (as well as being potentially annoying for controllers) is not providing the safety benefit that it should and secondly, virtually all of the operational issues that I have seen can be overcome either through algorithm improvements or through careful tuning of parameters.

Too many Tracks

The 'split tracks' issue can be a problem. The displayed system tracks are the fundamental data that the controller uses for Air Traffic Control, and that the Safety Nets use to determine if alerting conditions exist. A split track is essentially the occurrence of two (or sometimes more) tracks for only one actual aircraft. Surveillance errors are the main reason that split tracks occur, and whilst they can slightly clutter the controller's display, they can be much more distracting if they result in false STCA alerts.

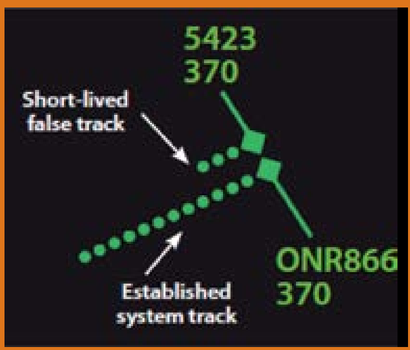

There are a number of reasons for split tracks. Whilst the Surveillance Data Processing System creates a split track, the root cause is usually due to erroneous radar data. These errors can include position errors, poorly extracted Mode A or Mode C (SSR data), or split plots (two radar plots where only one should exist).

As a case in point, an ANSP from an ECAC member state recently reported to EUROCONTROL's Safety Nets Performance Improvement Network (SPIN) Sub Group that it was experiencing a large number of nuisance STCA alerts in part of their airspace. Their own evaluations indicated that half of the alerts were from split tracks. Having verified their analysis, a novel and highly adaptable split track detection algorithm was designed, which the system supplier implemented in the STCA function. The result was a resounding success – halving the overall STCA alert rate overnight. Although removing split tracks themselves from the screen would inevitably take some considerable effort, modifying STCA to suppress the resultant false alerts was a quick and effective solution. In this particular airspace, some further suggested parameter tuning would then reduce the STCA alert rate to one third of its original value.

There have been a number of instances where ANSPs have reported issues which, on analysis, turn out to be due to inappropriate design of a Safety Net, either generally or which make it unsuitable for specific airspace. These cases are fortunately not common and, as Safety Nets gradually improve, the trend is for them to become increasingly so. In order to overcome such issues, it is of paramount importance that ANSP and ATC system suppliers work hand-in-hand to ensure that a Safety Net will be appropriate for the local operational environment by sharing information (including local procedures and system specifications), and by ensuring that full validation and verification are carried out before the Safety Net system goes into operational use. In my experience, the very best systems are produced when ANSP and manufacturer work collaboratively and are fully committed to achieving a common goal.

Good by Design

A good Safety Net design will typically incorporate a number of key features:

- Alerting rules (including prediction algorithms) appropriate for the operational and surveillance environment which should assess actual risk;
- Some protection against surveillance errors (especially split tracks);
- Protection against surveillance data items being missing (e.g. sudden drop of Mode A or Mode C);
- Flexibility for different parameters to be easily set in various types of airspace or for different types of flight;

SURVEILLANCE ERROR DETECTION	✔
ELIGIBILITY CRITERIA	✔
AIRSPACE DEFINITION	✔
ALERTING RULES	✔

There are a number of elements to consider in the design of a Safety Nets system. The design involves considerably more than just the rules for generating an alert. It includes testing whether input data is valid (including surveillance error detection), the determination of which tracks will be processed (eligibility criteria) and potentially the flexibility to allow different parameters to be applied in various types of airspace.

Much of the above will seem quite obvious. Nevertheless, I have found a number of problems with operational Safety Net systems which cannot be overcome by parameter tuning, and therefore require a modification to the Safety Net system itself. To mitigate this, the simplest and most cost-effective thing for an ANSP to do is to carefully examine all the available documentation including any system specifications and user manuals as early as possible in the procurement process. Some potential problems can be identified by having suitably-qualified staff check that the design of the Safety Net will be appropriate for the target operational environment. However, some issues may not be discovered until the system is trialled using real traffic data and an operationally realistic parameter adaptation.

It follows therefore that a completely off-the-shelf solution may not be appropriate in many cases, and it may be in an ANSP's interest to seek a flexible contracting mechanism which will allow some changes to be made to the supplier's standard product.

Testing Times

ANSPs will want to undertake some testing of any new Safety Net system to satisfy themselves firstly that it is functioning as specified and secondly that it's performance will be operationally acceptable.

In an ideal situation this testing and operational tuning will be undertaken in a similar time frame, so that a reasonable adaptation is already available on the day that the Safety Nets system is put into operational use, or perhaps earlier for pre-operational controller training.

For the purposes of system verification, the parameter adaptations (STCA volumes, MSAW alerting surfaces etc.) and the traffic scenarios do not have to be realistic; in fact they should be contrived in order to test as many aspects of the intended functionality as possible.

Separate parameter tuning will be required to assess the new system for operational acceptability. Airspace volumes and alerting thresholds must be set to operationally realistic values in order to make this assessment and parameter optimisation can only really be considered complete once the operational acceptability requirements have been met.

Optimisation Techniques

In the past, some ANSPs have activated Safety Net systems and expected them to be 'plug-and-play' by relying on the manufacturers default settings only to have to switch them off again for adjustment Nowadays, ANSPs and manufacturers alike understand that Safety Nets have to be configured for the local airspace and procedures before going into operation. Nevertheless, full optimisation can still take considerably more effort than many people realise.

Optimisation requires data and, ideally, plenty of it. The techniques used will vary depending on the particular Safety Net. Of significance will be whether or not the system relies on controller interaction to determine when an aircraft is under ATC and hence must participate in the system If controller interaction is necessary, then this can place a practical constraint on how much data can be realistically made available for alert analysis and tuning before the system goes operational. In this case, maximum benefit needs to be leveraged from whatever system track recordings and alert log files can be made available.

In my experience, the most powerful methods of optimisation involve the use of off-line models of Safety Nets – versions of the system which can be run repeatedly with different parameter sets. However, either way a full understanding of the algorithms and the role that each parameter plays in the alerting decision will make the optimisation process very much faster.

All optimisation starts with defining appropriate airspace volumes. These may be STCA volumes (where different conflict thresholds are used), MSAW volumes describing the alert surface, APW volumes describing danger areas and restricted airspace or APM approach funnels.

For some Safety Nets, such as MSAW and APM, the overall alerting performance is dictated by appropriate definition of these airspace volumes by the user. MSAW relies hugely on having a sufficiently fine resolution of the alerting surface combined with carefully-crafted inhibited areas which take account of the standard arrival routes. APM relies on having approach funnels defined to take account of all the various types of approach to and in the vicinity of each APM-protected runway. APM performance in particular benefits from some detailed technical and operational input and it is hard to imagine how one could easily optimise an APM system without recourse to an off-line model and analysis/visualisation tools.

One important thing to bear in mind is that it is very easy to tune Safety Net performance to match a particular set of traffic data. After a tuning exercise, it is important to compare the new tuning against the original parameters on a fresh traffic recording. This will provide confidence that the new parameters provide a benefit generally, rather than just for the traffic sample against which the Safety Net was tuned.

Closing the Loop

Once an optimisation is considered complete, an ANSP should be in 'monitoring' mode, making regular measurements to check that the performance of the new system has not degraded due to operational changes. They should also be seeking feedback from their controllers to help understand whether there are specific concerns or issues which might be grounds for restarting the 'Continual Improvement' loop.

In summary, the most effective Safety Net systems have been implemented when an ANSP and a supplier have worked collaboratively. This is not trivial and needs commitment from senior management on both sides, but it brings demonstrable safety benefits.

ROD **HOWELL**

is an expert in ground-based safety nets and surveillance tracking at QinetiQ (UK). His work has included R&D and design for the NATS' Enhanced STCA system, technical advice to the EUROCONTROL SPIN Sub Group, and technical support to help a number of ANSPs optimise their safety nets systems. He is the primary developer of an AV tool (STRACK), which is used for tracking analysis, and PolyGen which is used in the production of MSAW surfaces.