Knowledge is a Dangerous Thing: Collective fluctuations in pilot intent

Philip Goetz

Intelligent Automation, Inc.
7519 Standish Place
Rockville, MD 20855
phil@i-a-i.com

Abstract

It has been shown that, in an abstract model of 2 limited shared resources, when abstract "agents" with similar utility functions for these resources attempt to maximize their private utilities based on obsolete information about what other agents are doing, oscillations or chaotic behavior can arise spontaneously, resulting in lower utility for everyone (Kephart, Hogg, & Huberman 1989). We have seen the same oscillations appear in more realistic simulations of a future air-traffic-management scenario.

Motivation

NASA and the FAA are considering ways to redesign the US national airspace (NAS) to operate in a more distributed fashion known as Free Flight. More responsibility for route determination and possibly separation would be given to flight crews (RTCA 1997, DAG-TM Team 1999), in order to reduce fuel consumption and delays, and to cope with air traffic densities above the capacity of the present ground-control-based air traffic control (ATC) system. Since the new paradigm is to have collaborative decision-making rather than centralized control, it is called air traffic management (ATM).

In today's ATC system, controllers develop a plan for their airspace and position aircraft accordingly. An ATM system, however, is more like a market. The state of the airspace is not dictated by some master plan, but is a compromise between the desires of the agents involved. Like a market, the system should generally move towards an efficient equilibrium, but it may also oscillate about that equilibrium, in cyclic or even chaotic fluctuations of over- and under-valuations of particular resources.

Analytical work with multi-agent systems, including highly abstract simulations, has studied the conditions under which such fluctuations can arise (Huberman & Hogg 1988, Kephart, Hogg, & Huberman 1989). The model is one in which N agents each must use 1 of 2 resources. The utility U_r that each agent derives from resource *r* is the same for each agent, and is a function of the total number of agents using that resource. For analytical purposes, this is expressed as a preference probability function $\rho(f)$, where *f* is the fraction of agents using resource 1, and $\rho(f)$ is the probability that an agent choosing which resource to use will choose resource 1. $\rho(f)$ is a function of $U_1(f) - U_2(f)$.

Fluctuations arise when each agent re-evaluates which resource to use at a similar rate, and each agent has information that is delayed by a similar amount. We let α represent the average evaluation rate for each agent, and τ be how old their information is when they make an evaluation. In addition, the analysis adds normally-distributed noise with zero mean and standard deviation σ to each use of U_r by an agent to estimate a payoff. The system can be characterized by the parameters σ and $\beta = \alpha\tau$. (Two systems with identical values of β behave identically modulo a time scaling factor.)

The mechanism causing strong fluctuations to arise spontaneously is that when resource *r* is under-utilized to an extent that $U_r(f+n) = U_{(1-r)}(f-n)$ (*n* more agents could switch from resource (1-*r*) to resource *r* before the two resources had equal payoffs), more than *n* agents may switch to resource *r* in time τ. Similarly, if resource *r* is over-utilized, more agents may switch to resource (1-*r*) than would if a controlling agent were optimizing the system.

Chaotic fluctuations require that $\rho(f)$ be non-monotonic. That means that for some *f*, a resource will begin to appear more attractive as more other agents decide to use it. This seldom if ever happens with air traffic resources; hence we do not expect to see chaotic fluctuations.

Regular oscillations around an equilibrium point f_{eq}, or oscillations damping to f_{eq}, occur when $\rho'(f_{eq})$, the slope of the preference function at f_{eq}, lies between two negative values, the larger of which approaches zero from below as β increases (Kephart, Hogg, & Huberman 1989).

The agents we are considering are airplane flight crew (whom we will refer to as pilots). The resources, in this study, are the regions of airspace (called *sectors*) that they wish to fly through. A sector is, approximately, a slice of airspace covering a few miles in altitude and ten to a hundred miles in horizontal extent, with a polygonal appearance from above, all of which is controlled by one to three en-route controllers. In today's ATC system, a sector may handle in the region of 10-15 flights at a time.

At present, pilots are unaware when filing a flight plan of what sectors it will take them through, or how much traffic may be in those sectors. The supposition is that, in the future, pilots will have and make use of that information. If we then suppose that pilots attempt to go around sectors that are overbooked to avoid delays, then we may wonder if oscillations of the sort seen in the simple 2-resource analysis and simulations will occur in the skies with airplanes carrying several hundred passengers each.[1]

Simulation

Intelligent Automation, Inc. (IAI) has developed an agent-based NAS simulator to study the possible behavior of pilots and airlines when they are given more autonomy and more information. We are using it to study the question outlined above.

The supposition that different agents evaluate their resource usage at similar rates is reasonable, especially if (as seems most likely) this re-evaluation is done automatically by software. The supposition that the age τ of the information available is the same for all agents is also reasonable if we suppose that this information is reported through the ATM hierarchy and then made available to all flight crew. Thus we did not model the creation of individual values for α and τ, but coded them directly into our pilot and controller models. (In fact, we restricted α more strongly than in the analytical work, making evaluations regular and periodic rather than probabilistic.) We were more interested in the fact that there are many resources, and each pilot has a different utility function for each resource.

In our simulation, s_{Xt} is the number of spaces that will be available in sector X that a plane wants to enter when the plane arrives there at time t. The plane will be admitted iff $s_{Xt} \geq 1$. The pilot has available $s_{Xt}^{p-\tau}$, the estimate of s_{Xt} made by the en-route controller at τ time units before the present time p. Every $1/\alpha$ time units, each pilot computes, for each remaining sector on his flight plan, the probability $P(s_{Xt} \geq 1)$, as the area to the left of $s_{Xt}^{p-\tau}$ in the normal distribution with mean .5 and variance σ^2. If $P(s_{Xt} \geq 1)$ drops below some threshold value for any number of sectors on his flight plan, the pilot computes an alternate flight plan going around those sectors, and then compares the utilities of the old flight plan and the new flight plan. Utility is defined as the probability of success divided by the length. The probability of success of a flight plan is defined as the product of $P(s_{Xt} \geq 1)$ over all sectors in the flight plan. If the new flight plan has higher utility, the pilot switches to it, and declares his intent to all of the en-route controllers on both the old and new flight plans.

We used a simplified airspace with square sectors. Flights are generated randomly, with no bias towards particular regions or routes in the airspace, so that all sectors are, on average, comparable.

Predictions

Regions of stability

Suppose there are two sectors side-by-side, A and B, A being west of B. Consider all south-to-north and north-to-south flights in an area centered on A+B (the area of A joined with B). Suppose every sector is allowed to contain up to m flights at any one time. To make our equations use the f notation above, we will write them in terms of $f_{Xt} = (m - s_{Xt}) / m$. A and B have different utilities to each individual flight, but suppose they have an equal average utility U to all these flights. Expected utility $E(X) = U \times P(X)$, where $P(X) = P(s_{Xt}^{t} \geq 1)$. Let $\rho_A(f_{At}^{p-\tau}) = P(E(A) > E(B) \mid f_{At}^{p-\tau})$. Then $\rho_A(f_{At}^{p-\tau}) = P(P(A) > P(B) \mid f_{At}^{p-\tau}) = P(P(A) - P(B) > 0)$. $P(A) - P(B)$ has a normal distribution with mean $s_{At}^{p-\tau} - s_{Bt}^{p-\tau} = m f_{Bt}^{p-\tau} - m f_{At}^{p-\tau}$ and variance $\sigma_{A-B}^2 = 2(\sigma^2)$.

We are looking for flights that oscillate back and forth between intending to go through sector A and sector B. The rate of change of the expected proportion of flights in sector A at time s is the rate at which planes from B decide to switch to A, minus the rate at which planes from A decide to switch to B. This equation, then, is the equivalent, for our application, of equation 7 in (KHH 1989). Note that now s is the time we are making a prediction for; t is the time at which the prediction was made, to be consistent with the notation in (KHH 1989). We write $\rho_A(f_{As}^{t})$ instead of simply ρ_A to remind us that it is a function of f_{As}^{t}.

$$df_{As}^{t}/dt = \alpha f_{Bs}^{t} \rho_A(f_{As}^{t-\tau}) - \alpha(f_{As}[1-\rho_A(f_{As}^{t-\tau})] = \alpha[\rho_A(f_{As}^{t-\tau}) - f_{As}^{t}] \quad (1)$$

Following (KHH 1989 equation 8), we find a stability equation for the values of β which lead to non-oscillatory convergence to the equilibrium f_{As}^{∞}, in terms of $\gamma = \rho_A'(f_{As}^{\infty})$. This involves the following assumptions and approximations:

- Both sectors have, on average, equal utilities.
- Switches of planes between A, B, and other sectors will cancel out, so $f_{At}^{t-d} + f_{Bt}^{t-d} = 1$.
- $f_{As}^{\infty} = f_{Bs}^{\infty} = ½$
- $\rho_A(f_{As}^{t-\tau})$ is approximated linearly around the fixed point f_{As}^{∞}, using $\delta(t) = f_{As}^{t} - f_{As}^{\infty}$, $\gamma = \rho_A'(f_{As}^{\infty})$:

$$d\delta(t)/dt = \alpha[\gamma\delta(t-\tau) - \delta(t)] \quad (2)$$

[1] Delays are more commonly caused by limited acceptance rates at airports than by sector congestion. However, airplanes do not commonly route to alternate airports when their destination airports are busy, so we did not study this case.

- Following (KHH 1989 equation 9), we make the dubious[1] approximation $\delta(t) = e^{\alpha\xi t}$, as suggested in (Bellman & Cooke 1963). Recall that $\beta = \alpha\tau$:
$$\gamma e^{-\beta\xi} - \xi - 1 = 0 \quad (3)$$
Equation 3 has infinitely many complex solutions $\xi = \xi_r + i\xi_i$. Where it has no real solution, the behavior cannot converge to a stable equilibrium.
- We approximate $e^{-\beta\xi}$ with a Taylor series expansion up to the quadratic term:
$$\xi = [(1+\gamma\beta) \pm \mathrm{sqrt}((1+\gamma\beta)^2 - 4\gamma\beta^2(\gamma-1)/2)] / \gamma\beta^2 \quad (4)$$
- This has real solutions iff
$$(1+\gamma\beta)^2 - 4\gamma\beta^2(\gamma-1)/2 \geq 0$$
$$(\gamma-\mathrm{sqrt}(2\gamma^2-2\gamma))/(\gamma^2-2\gamma) \leq \beta \leq (\gamma+\mathrm{sqrt}(2\gamma^2-2\gamma))/(\gamma^2-2\gamma) \quad (5)$$

To have stability, we require that Equation 5 holds. For air traffic, $\gamma < 0$, because it is the slope of a decreasing function. The following graph shows the region in γ and β for which it does, between the top curve and the line $\beta = 0$. The lower boundary for β is graphed to convince the reader that it does not rise above $\beta = 0$.

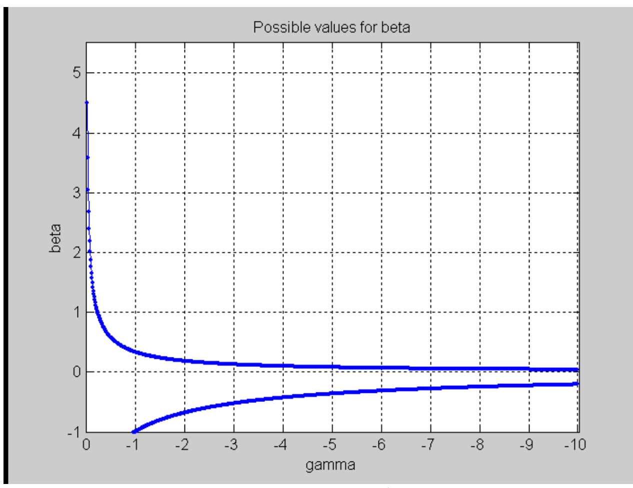

Figure 1: Possible stable values for β as a function of γ

When two sectors have equal value, γ is the negative of the slope of the integral of a normal probability distribution at its mean, so its actual value is $-1/A$, $A = \sigma\mathrm{sqrt}(2\pi)$. This gives us the stability equation

$$-\sigma^2 2\pi[\mathrm{sqrt}(2+2A) + 1] / (A + 2) \leq \beta,$$
$$\beta \leq \sigma^2 2\pi[\mathrm{sqrt}(2+2A) - 1] / (A + 2) \quad (6)$$

This allows us to give a maximum "information age" β for stability given σ. (Note that β is dimensionless.) Note that stability is impossible when there is no uncertainty. The area below the curve in the graph below shows the ranges for σ and β for which equation 6 holds, subject to $\beta \geq 0$.

[1] This approximation is dubious because it makes $\delta(t) > 0$, whereas we wish to investigate the behavior of $\delta(t)$ around zero.

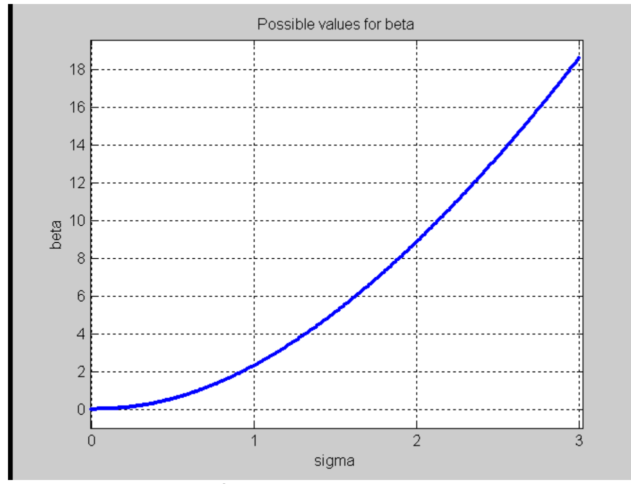

Figure 2: Graph of $\beta_{max}(\sigma)$. Values below the blue line are possible stable values for β, subject to $\beta > 0$

Oscillation frequency

To determine whether observed oscillations are the result of this mechanism, we wish to know the frequency to look for. Recall equation 1: $df(t)/dt = \alpha[\rho(f(t-\tau)) - f(t)]$, where $f(t)$ is the fraction of agents choosing resource 1 (the sector under consideration) over all other resources at time t, and $\rho(f(t-\tau))$ is the probability that an agent making a choice will choose resource 1 given that agent believes that $f(t-\tau)$ agents are currently using it.

The tops and bottoms of the oscillations are where $df/dt = 0$, so we try to find the zeros of this equation. Denote any arbitrary value of t for which $df/dt=0$ as t_0. Let f_{eq} be the equilibrium value of f in the case where $\tau = 0$, at which $\rho(f_{eq}) = f_{eq}$. $df/dt = 0$ when the (old) information available, $f(t-\tau)$, indicates that the system is at an equilibrium; thus $df/dt = 0$ when $\rho(f(t-\tau)) = \rho(f_{eq}) = f_{eq}$. Thus $f(t_0-\tau) = f_{eq}$. The time between crossing f_{eq} at time $t_0-\tau$, and reaching the next maximum or minimum at t_0, is thus τ. If you consider one full cycle of the oscillation, it crosses f_{eq} two times. Thus f rising from f_{eq} to a maximum, and falling from f_{eq} to a minimum, adds time 2τ to the wavelength.

We need to add in a time for the falling from a maximum to f_{eq}, and for rising from a minimum to f_{eq}. We can assume, not too unreasonably, that our utility function $U(f)$ is symmetric around f_{eq}, and hence these times are the same due to the symmetry of the situations. Denote this time as u. We will say that the slope at time $t_0 + u/2$, halfway between a maximum and f_{eq}, is approximately the average slope over that interval, $(f_{eq} - f_{max}) / u$. Thus $(f_{eq} - f_{max}) / u \approx \alpha[\rho(f(t_0+u/2-\tau)) - f(t_0+u/2)]$. We will further suppose, not entirely unreasonably, that $f(t_0+u/2) \approx (f_{eq} + f_{max})/2$, and $\rho(f(t_0+u/2-\tau)) \approx [\rho(f(t_0-\tau)) + \rho(f(t_0))]/2 = f_{eq}/2$. Then

$$(f_{eq} - f_{max}) / u \approx \alpha[f_{eq}/2 - (f_{eq} + f_{max})/2] = \alpha[-f_{max}/2]$$
$$2(f_{eq} - f_{max}) \approx -\alpha u f_{max}$$
$$u \approx (2/\alpha)(1 - f_{eq}/f_{max})$$

The time to add to the wavelength from two intervals of length u is $(4/\alpha)(1 - f_{eq}/f_{max})$. We expect that $1 > f_{eq}/f_{max} > ½$, since $f_{max} - f_{eq}$ is probably similar to $f_{eq} - f_{min}$, and $f_{min} > 0$. This leads us to expect our wavelength to be between 2τ and $2\tau + 2/\alpha$. A more intuitive way of explaining this is that agents choose resource 1 at time t if most agents were using resource 2 at time t-τ, and it takes time roughly $1/\alpha$ before every agent has had an opportunity to switch.

Results

To test whether these equations tell us something useful about airspace stability, we ran simulations with different values for σ and β. For each value of σ, we want to test one value of β below the threshold β_σ, and one above it. We expect to see no oscillations for $\beta < \beta_\sigma$, and oscillations of period $2\tau + 1/\alpha$ for $\beta > \beta_\sigma$. To choose our test values, we chose a few values of σ, calculated β_σ, chose values of β on both sides of β_σ, and for each β we found values of α and τ to solve the simultaneous equations $2\tau + 1/\alpha = 5$ minutes, $\alpha\tau = \beta$.

Plotting the resulting sector schedules clearly showed oscillations of the expected frequency in some cases, as shown in Figure 3. There were no significant oscillations at significantly different frequencies. The schedule for a sector is divided into timeslots, with 4 timeslots per hour. On the x-axis of Figure 3 is the time of prediction t_p. The graph shows a set of functions $flights_{slot}(t_p)$ = <# of flights predicted, at time t_p, to be in the sector during timeslot *slot*>. Each differently-shaded line represents a different timeslot *slot*, graphed for $0 < t_p < endtime(slot)$. Predictions are made based on pilot intent. (Pilots inform sector controllers immediately when their intent changes.)

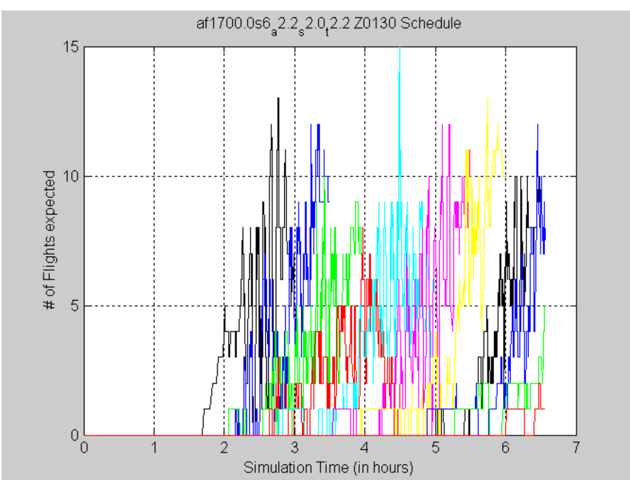

Figure 3: Graphs of expected number of flights by timeslot.

In order to test our prediction, however, it would not do to have humans examine thousands of such graphs. We designed a machine-computable score to tell whether a sector had oscillations of the expected frequency.

We began by throwing away all timeslots in the first 3 hours, because our simulations begin with zero flights in the air and so the initial timeslots are not valid. For each remaining timeslot, we computed the start time of the interesting part of the signal as the later of {3 hours after start of simulation, first time that the schedule reaches the sector maximum minus σ}. We looked at the predicted number of planes in that sector during that timeslot from that signal start time up to that timeslot's starting time. We added the power spectra of that signal over all timeslots, and computed an "oscillation score" from it. The oscillation score for each run was computed as the average of all the oscillation scores of all the timeslots of all its 16 central sectors.

The difficulty is in computing an oscillation score. At first, we took the power spectrum's area between $\tau + 1/\alpha$ and $3\tau + 1/\alpha$, and gave that as the "oscillation score". However, the oscillation frequency could not be predicted accurately enough for this to be a good measure. The actual frequency of oscillations was often slightly lower than the predicted frequency. Figure 4 gives an example of a sector schedule graph with many oscillations at approximately the right frequency, and the average power spectrum of each of the lines graphed, with vertical lines at $\tau + 1/\alpha$ and $3\tau + 1/\alpha$. Note that the peak in the power spectrum is at a frequency slightly lower than the expected frequency. This results in a low score. Also, results at different frequencies were not comparable, because the underlying noise does not have a flat spectrum.

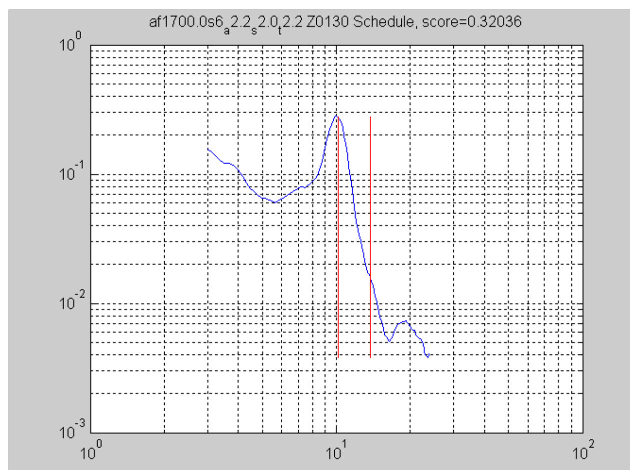

Figure 4: Lower-than-predicted frequency peak leads to low score.

We then concentrated on finding a way to measure oscillation strength without knowing ahead of time the frequency of oscillation. We wanted to do this by searching the power spectrum for maxima, but first had to do some processing in order to make values at different frequencies comparable.

The idea is to divide out from the power spectrum the underlying shape that it would take without oscillations. ("Dividing out" *pow* by *baseline* means that pow_{new}(freq) =

pow(freq) / *baseline*(freq) for every datapoint index *freq*.) In all cases we normalized both spectra so that the area under the spectrum within ($\omega_{exp}/4$, $4\omega_{exp}$) = 1, because of noise at high and low frequencies due to our sampling process.

Our first attempt was to model the "null-hypothesis" power spectrum as the power spectrum that would be expected if there were no collective effects. This model generated a random schedule record according to a model, then found its average power spectrum.

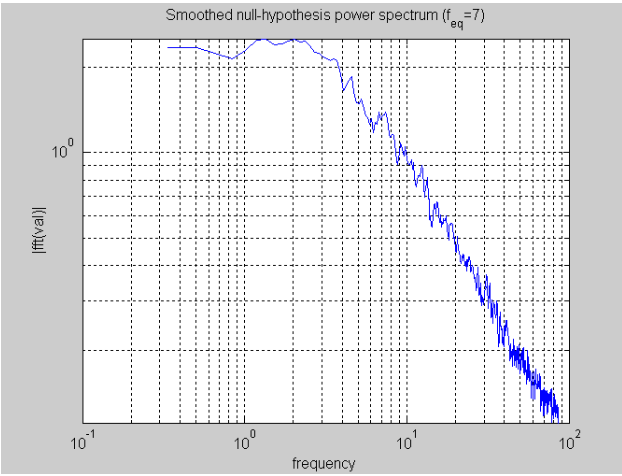

Figure 5: Null-hypothesis power spectrum.

This model produced null-hypothesis power spectra such as that shown in Figure 5. However, dividing out the experimental power spectra by the null-hypothesis power spectrum produced flatter, but not completely flat, spectra. Rather than refine our model, we obtained better results by fitting an exponential function to the experimental power spectrum, and dividing out by that exponential function. The fit was made by finding λ and c to minimize the sum of the values

$$[(\text{signal(freq)} - e^{\lambda \log(\text{freq}) + c}) / e^{\lambda \log(\text{freq}) + c}]^2$$

with *freq* varying over all datapoint frequencies.

This technique produced modified power spectra so flat that we could reliably say that the maximal value in such a spectrum indicated the frequency that was the most over-represented in the signal. The oscillation frequency found matched the predicted frequency over all parameter values. However, there were a few cases in which the harmonics of the expected signal had peaks comparable to the signal itself, and averaging these frequency values in as we were doing produced meaningless "average frequencies". Thus, we restricted our search for the real maximum ω_{max} to within ($\omega_{expected}/1.7$, $1.7\omega_{expected}$). The strength of the oscillation was taken as the height of the maximum in this adjusted power spectrum. This approach worked well except in degenerate cases where τ approached zero and very-high-frequency oscillations appeared.

The results generally support the prediction that stronger oscillations occur at higher values of β. However, they do not appear to support the prediction that oscillations disappear for $\beta < \beta_\sigma$, nor that increasing σ reduces oscillations, as seen in Table 1.

		σ	.399	1	2	3
β	α	τ				
.2	.28/min.	.714 min.	2.74	3.76		
1	.6/min.	1.67 min.	3.07	3.45	3.27	
5	2.2/min.	2.27 min.	3.55	4.88*	3.70	
12	5/min.	2.40 min.		4.21	4.72	5.87
25	10.2/min.	2.45 min.				6.27

Table 1: Oscillation scores for β below and above $\beta_{max}(\sigma)$. (* indicates an outlier more than 5 standard deviations out was discarded.)

A peculiarity of the way in which the simulator was programmed is that pilots filed their initial flight plans before checking on the density of traffic in the sectors it passed through. If they received notification back that sectors were expected to be overcrowded, they might then revise their flight plans. This caused situations where new flights could heavily overload a sector, leading to a crash as they and other flights then tried to switch their flight paths around that sector, setting up an oscillation. We revised the code so that pilots checked on the traffic density expected along their flight plan before filing their initial flight plans. The results were remarkable. The regular oscillations, for the most part, disappeared.

Interpretation

There are many possible reasons why oscillations do not disappear below a critical value of β. The equation for β_{max} involved several approximations and idealized assumptions. The simulation is very "messy", incorporating real-life complications such as collision avoidance. The measurement of oscillation scores is problematic. However, the lack of even a sharp rise in oscillation scores near β_{max} makes it difficult to make any specific recommendations that could be used to avoid oscillations in Free Flight.

The disappearance of oscillations when pilots check before filing does not mean that the oscillations in previous runs consisted entirely of new flights filing for and then retracting filings from full sectors. Inspection of the data on previous runs clearly shows that most oscillations dip well below what the pilots involved would regard as a problematic density. Also, oscillations began in sectors long before they were overcrowded, indicating that they had spilled over from neighboring sectors. Thus the oscillations involve more than just the new flights. Apparently the oscillations are triggered by overloading. When all pilots check before ever filing a flight plan for a sector, oscillations occur much less frequently. When just a few overload a sector, however, this can trigger oscillations involving other pilots. The potential to oscillate at a certain frequency exists within the system in both cases, but it requires a small triggering event to set it off.

Future Work

The primary objective of this project is to determine whether selfish optimizing behavior might produce sub-optimal collective behavior in a predictable way.

Simulations will be run with the actual U.S. sector map, and actual ETMS flight data, rather than the square sectors and random flights, to see if oscillations can still be triggered, and to determine the size of a triggering stimulus relative to the size of the resulting oscillation. Simulations will be run over a wider range of values for β and σ to see if the transition to oscillations lies elsewhere.

The major question regarding the issue of avoiding triggering events is whether it is advantageous to individual pilots to do so. This will be tested by running simulations in which some pilots are extremely optimistic, assuming that they will be able to get into sectors even if there are negative spaces predicted to be available. What might happen is that such pilots will trigger oscillation events, chasing other pilots out of their desired sectors, and get to enter all their desired sectors simply by anticipating the oscillations and not re-routing when it appears their sectors will be overbooked. In that case, the system needs some safeguards against selfish behavior. Additional work could be done studying different mixes (percentages) of "selfish" (or optimistic) agents, to see which mix is most efficient. It may be that, although the system is vulnerable to parasitism by some selfish agents, its behavior will alter so that such selfish behavior is no longer self-advantageous when more than a certain number of selfish agents exist, protecting the system from chaos. This is similar to work in (Kephart, Hogg, & Huberman 1989) on "smart" agents that attempt to anticipate oscillations, in which it was found that a mix of 10% smart and 90% dumb agents was more efficient than 100% dumb agents, but higher percentages of smart agents led to chaotic behavior.

Behaviors that could lead to instability in the manner studied here are generally caused when resources are available on a first-come, first-serve basis with no cost to the user. Various solutions could be imagined that would require reservations or payment for services, so that the fluctuations and settling to equilibrium would occur with prices rather than with airplanes. However, they are difficult to implement in the NAS. Requiring reservations without cancellation costs encourages airlines to overbook resources intentionally. Requiring payment for services would be politically difficult to implement. Thus we cannot rely on market mechanisms to resolve these issues once Free Flight is implemented. We must understand them while the procedures for Free Flight air traffic management are being drawn up.

Acknowledgements

This work is funded by the NASA Ames Research Center. Thanks go to Renato Levy for preliminary work, to the NASA FACET team for access to their data and code, and to Dr. George Meyer of NASA Ames for his support.

References

Bellman, R., & Cooke, K.L. 1963. *Differential-Difference Equations.* NY: Academic Press.

DAG-TM Team 1999. Research Plan for Distributed Air/Ground Traffic Management (DAG-TM) version 1.01. Available from http://www.asc.nasa.gov/aatt/dag.html

Huberman, B. & Hogg, T. 1988. The Behavior of Computational Ecologies. In B. Huberman (editor), *The Ecology of Computation*: 177-205. Amsterdam: North Holland Press.

Kephart, J.O., Hogg, T., & Huberman, B.A. (1989). Dynamics of Computational Ecosystems. *Physical Review A* 40(1): 404-421.

RTCA Select Committee on Free Flight Implementation (Dec. 1, 1997), Government/Industry Operational Concept for the Evolution of Free Flight. Washington DC: RTCA (www.rtca.org).