AGENT-BASED MODEL OF AERIAL AD-HOC NETWORK MARKET POTENTIAL

Christopher B Watkins, Cihan H Dagli, Missouri University of Science and Technology
Systems Engineering Graduate Program, Rolla, MO

Abstract

Mobile Ad hoc NETworks (MANETs) have seen significant growth as a research topic in recent years. As opposed to a traditional network with a fixed node infrastructure, the nodes in a MANET are free to move and arbitrarily organize themselves while maintaining wireless communication with each other. The major benefit of MANETs is the lack of reliance on a fixed network infrastructure. The challenge is to dynamically route traffic through a set of mobile network nodes.

MANETs represent a novel solution for providing internet access in aircraft. Currently, each aircraft maintains its own gateway to the internet using expensive satellite data links. A MANET solution would allow multiple aircraft to communicate with each other in-flight and share an internet gateway, thus reducing the user cost for internet service. This is what we call an "Aerial Ad-Hoc Network."

Significant research has been focused on protocols that can efficiently route traffic in this dynamic environment of mobile network nodes. This paper focuses on the evaluation of the market potential for an aerial ad-hoc network. The benefits of low cost internet access could grow the in-flight internet user base. However the benefits are balanced with greater network failures due to bandwidth overload conditions. Each aircraft carries hundreds of people, and therefore the number of internet users across a group of aircraft sharing an internet gateway could easily overload the network.

This paper presents a Monte-Carlo simulation that characterizes the network failure rates that could be expected in an aerial ad-hoc network. In addition, we construct an agent-based model of the aerial ad-hoc network market. This simulation is used to determine the steady state user levels that can be expected for the system given individual user satisfaction with the internet service.

Introduction

The internet has experienced tremendous growth in its user base since the mid-1990s. Figure 1 shows that the internet user base has grown over 123 times within 15 years (December 1995 - September 2010) and consists of almost 30% of the world's population [1].

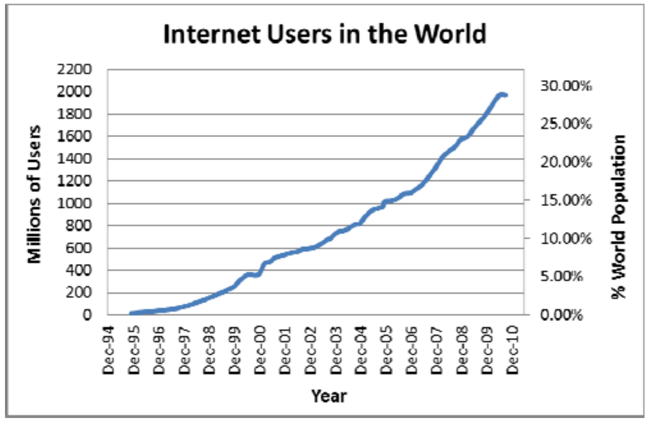

Figure 1 – Internet Users in the World [1]

However internet access in commercial aircraft has been unavailable until 2008. At the end of 2010, only about 1000 commercial jets were equipped with Wi-Fi internet access in the US, which is less than a third of the fleet [2]. Therefore the user base is relatively small. Analysts estimate that less than 10% of the passengers equipped with a Wi-Fi device actually use the internet while in-flight [2]. Virgin America has claimed a higher-than-average internet user base: 12-15% of passengers across their fleet. This has been attributed to the fact that Virgin America maintains a high proportion of cross-country flights between technology-industry locations (i.e. Boston and San Francisco) [3].

The current aerial internet access architecture requires each aircraft to maintain a high-cost satellite link to provide an in-flight internet gateway. Under this model, it is estimated that 8-10% of the travelers need to pay for Internet access in order for the service to be profitable within five years [3].

Aerial Ad-Hoc Networks

Aerial ad-hoc networks represent a new, low-cost architecture for aerial internet access. The utilization of MANET technology would allow multiple aircraft to share a single internet gateway. The gateway could be a satellite or ground link in the aerial ad-hoc network as shown in Figure 2.

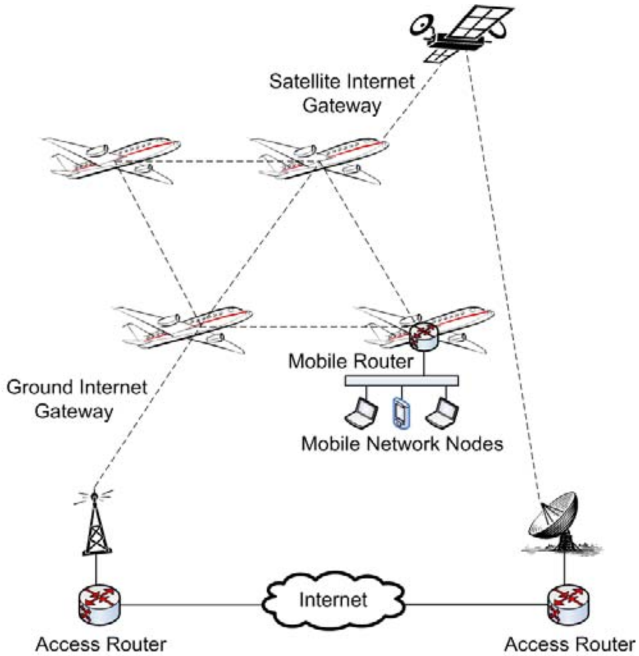

Image adapted from Hoffmann, Medina, Wolisz [4]

Figure 2 - Aerial Ad-Hoc Network Concept

Aerial ad-hoc networks rely upon the assumption that they will be within communication range of a sufficient number of other aircraft to maintain a single or multi-hop route to an internet gateway. The density of air traffic over the US was analyzed in previous work [5] in order to determine the probability of being within range of other aircraft in a low density region during non-peak travel times (worst case scenario). It was shown that there is nearly a 100% probability that there will be at least 12 aircraft within communications range in the US, making an aerial ad-hoc network a viable solution. It was also proposed in previous work [5] that Doppler shift of packets can be used to identify relative stability of mobile routers when a routing protocol dynamically selects a route. There is a significant amount of previous work that evaluates and proposes routing protocols for MANETs and aerial ad-hoc networks [4-10].

Potential to Grow Aerial Internet Market?

Instead of focusing on the technical feasibility of the aerial ad-hoc network technology, this study evaluates whether an aerial ad-hoc network would grow the in-flight internet user base.

It was shown that the internet industry is booming, but the market for in-flight users has been slow. This begs the question: How many users can be expected to adopt a lower cost internet service based on an aerial ad-hoc network? Can we expect this technology insertion to grow the user base beyond current levels?

It is expected that a lower cost internet access will remove financial barriers and increase the user base. However the drawback of an aerial ad-hoc network is that multiple aircraft share a single internet gateway, and therefore are more likely to experience failures due to an overloaded network. These failures can be expected to create dissatisfaction and result in lost users. It is assumed that the combined addition and loss of users will result in a steady state user base, however it is not known if this critical user level is better or worse than current levels for traditional in-flight internet. In order to determine the market potential, we must answer two important questions:

- What failure rates can be expected for aerial ad-hoc networks?
- What steady state user level can be expected for the system given the failure rates and resultant individual user satisfaction?

Two models were created to address these questions. The first model simulated the random variables which affect aerial ad-hoc network performance (and resultant failures). The second model simulated human behavior in response to both successful and failed user experiences. These models were coded in PHP and hosted on the web for others to access and use:

www.watkinsplace.com/simulations/aerialnetwork/

A description of the model construction and analysis results are summarized in this paper.

Flight Behavior Modeled after the North Atlantic Corridor

The PHP-based models allow the investigator to select input conditions to simulate any flight behavior. For the purposes of this study, the simulation input conditions characterized the flights between the US and Europe in the North Atlantic corridor. This was selected because it represents a significant area of opportunity for an aerial ad-hoc network.

The flights are lengthy, typically 7-11 hours depending upon departure and arrival locations. Internet access is likely to be more attractive to a passenger on a long flight as compared to a short 1-2 hour domestic flight.

Additionally, the US-European flights in the North Atlantic Corridor generally fly together in the same direction (Figure 3). In the North Atlantic evening, the flights fly westward from the US to Europe. Likewise, the flights fly eastward from Europe to the US in the North Atlantic day. Since the aircraft generally fly in the same direction, the network packet routing would remain more stable. This places less demand on the dynamic routing protocol and allows for earlier, less complex routing algorithms to be implemented and tested.

Specifically, the following simulation inputs were used to characterize the flights in the North Atlantic Corridor [12]:

- Number of concurrent flights = 300-500
- Number of passengers per flight = 100-300

Expected Failure Rates

There are two main types of failure rates for an aerial ad-hoc network:

- Loss of link between source aircraft and aircraft providing the internet gateway
- Network overload when the internet gateway bandwidth exceeds practical limits for throughput performance

The first failure type, loss of link, is not addressed by this study. It is assumed that the routing protocol used in an aerial network can sufficiently minimize link loss by proactively selecting new network routes before the old route is lost, and by evaluating the goodness of a potential route's stability before selecting it [5].

Figure 3 – 2008 Snapshot of Air Traffic (courtesy of Zurich University of Applied Sciences) [11]

It is also assumed that the aircraft will maintain a traditional dedicated internet gateway link that could be used as a backup if an acceptable ad-hoc network route cannot be located. Given these assumptions, a loss of link failure can be considered a statistically insignificant event.

The second failure type, network overload, is addressed by this study. A Monte-Carlo simulation was created to determine the expected network failure rate as a function of the number of internet users on the network. There were four random variables included in the simulation:

- Number of concurrent flights
- Number of passengers per flight
- Internet bandwidth utilized per user
- Number of flights providing internet gateways in the ad-hoc network

These random variables represent a large number of coupled degrees of freedom in the network loading equation. The Monte-Carlo simulation is an appropriate method to address this problem space since it manages the uncertainty by performing repeated experiments based on random behavior. Based on these experiments, it calculates the average resultant network failure rate.

Monte-Carlo Simulation Assumptions

The PHP model allows the investigator to define their own assumptions for the simulation. The analysis results presented here are based on the following assumptions:

- Number concurrent flights (*f*) = 30-50
- Number passengers per flight (*p*) = 100-300
- Internet bandwidth utilized per user (*b*) = 10-90 Kbps
- Number of flights providing internet gateways in the ad-hoc network (*g*) = 2-10
- Internet gateway bandwidth (gb) = 30 Mbps
- Max network loading before failure (ml) = 40%

The first two assumptions are based on traffic profiles in the North Atlantic corridor as described earlier. Due to memory and processing time constraints on the webserver running the PHP script, the analysis was scaled to address 10% of the full set of air traffic (30-50 flights instead of 300-500). Since the number of internet gateways was also scaled by 10%, the analysis results for the subset of flights is identical to the analysis of the full set of flights.

The internet bandwidth utilized per user is based on actual experience for users of satellite broadband internet service. These users averaged 50-60kbps at peak hours [13]. However it is worth noting that this assumption does not allow for extended high-resolution video streaming, or frequent downloading of large files.

The number of internet gateways was selected to allow anywhere from 3 to 25 aircraft to share a single internet gateway.

The internet gateway bandwidth (30Mbps) is based on bandwidth currently available to aircraft with satellite links [14].

The max network loading allowed before declaring a failure was set to 40% based on a CISCO LAN Switching reference stating that failures can be declared when network load is between 30-50% [15]. Beyond this network load, there will be significant packet collisions, resulting in unacceptably slow performance for the internet user.

Monte-Carlo Simulation Method

The Monte-Carlo simulation repeats the experiment multiple times as defined by the investigator. In this study, 2000 experiments were completed. For each iteration, an integer value for each of the four variables (*f*, *p*, *b*, *g*) is randomly selected between the ranges defined by the simulation assumptions. The distribution of random parameters conforms to a normal distribution according to the Box-Muller Transformation method [16]:

```
$parameter = round($avg_parameter_value
+(sqrt(2*log(UniformRandomNumber()))*si
n(2*pi()*UniformRandomNumber())*(($max_
parameter_value - $avg_parameter_value)
/4.25)),0);
```

Equation 1 – PHP Representation of the Box-Muller Normal Random Distribution Function

The average parameter value is assumed to be half-way between its assumed minimum and maximum values.

Next, the number of flights per internet gateway (*fpg*) is calculated based on the assumption that the aircraft are evenly distributed among available internet gateways:

fpg = floor (f / g)

Finally, the bandwidth used by all internet users sharing a single internet gateway is calculated and compared to the maximum allowable bandwidth.

*If [(fpg*p*user_percent* b) > gb*ml] then failure=true*

This failure check is repeated 100 times, once for each internet user percentage (*user_percent*), 1 to 100%. Note that each aircraft can contain a different amount of users, but the percentage of all passengers in the group of aircraft sharing an internet connection is equal to *user_percent*. This experiment is completed 2000 times. Then the total number of failures as a function of user percentage level is dived by 2000 in order to calculate the failure rate for each user level.

Monte-Carlo Simulation Results

The Monte-Carlo simulation calculates a failure rate based on the percentage of internet users across the set of aircraft sharing a single gateway (Figure 4).

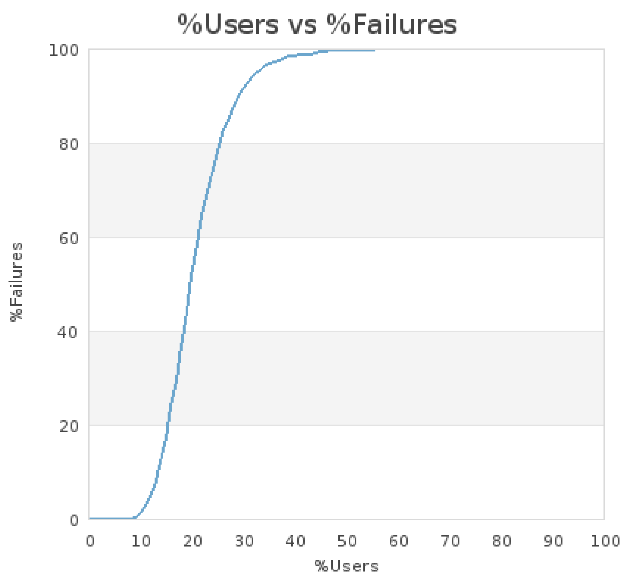

Figure 4 – % Users vs. % Network Failures

Given the stated simulation assumptions, there are no network failures when the user base is 10% or less. There is a sharp rise in failures when the user base grows from 10-40%, and failures are nearly certain when the user base grows above 40%.

Expected Steady State User Levels

After the failure rate table has been calculated using Monte-Carlo analysis, an agent-based simulation was constructed to model human response to both a successful and failed internet user experience. The agent-based simulation relies on simple rules of human behavior that guide whether each individual user (an autonomous agent) is a repeat internet user, lost user, refers new friends to use the service, or refers friends to drop the service. These are known as “happiness rules,” since they characterize the satisfaction (or happiness) of each individual internet user. The simulation addresses the simultaneous interaction of all user agents to model the emergent system level behavior. It was shown that simple rules that guide individual human behavior can lead to significantly different and complex outcomes for the aerial internet service system as a whole.

Agent-Based Simulation Assumptions

The PHP model allows the investigator to define their own assumptions for the simulation. The analysis results presented here are based on the following assumptions:

- Total number of potential passengers: 80,000
- Initial percentage of internet users (early adopters): 1%

The total number of potential passengers represents the pool of people from which a random subset of passengers is selected to fly each day. Not all potential passengers fly each day since this would not be a realistic behavior model. The group of potential passengers was set to 10x the number of actual passengers each day (one day = one iteration of the simulation). There are approximately 80,000 daily passengers that fly the Northern Atlantic Corridor each way (i.e. US-to-Europe or Europe-to-US). This is based on about 400 flights * 200 average passengers per flight [12]. Due to memory and processing time constraints on the webserver running the PHP script, the analysis was scaled to address 10% of the full set of air traffic, or 8,000 passengers (average of 40 flights instead of 400). Therefore the total number of potential passengers is set to 10 * 8,000 passengers = 80,000 people. Since the number of internet gateways was also scaled by 10%, the analysis results for the subset of flights are identical to the analysis of the full set of flights.

The other assumption, initial percentage of internet users, defines the initial condition for the simulation. This parameter does not significantly affect the steady-state system behavior for most cases of the happiness rules. As long as the happiness rules include a provision to add new internet users and loose existing users, then the initial condition of internet users only serves to define how quickly the critical system state is achieved.

Happiness Rules

The happiness rules are defined by the investigator to characterize the human response to a successful or failed user experience. The model was constructed to define a happiness rule for a combination of multiple successes and failures (Figure 5). Unique rules can be applied when an individual's count of successful or failed flights range from 1-4. The human response is assumed to be the same after 5 successes or failures.

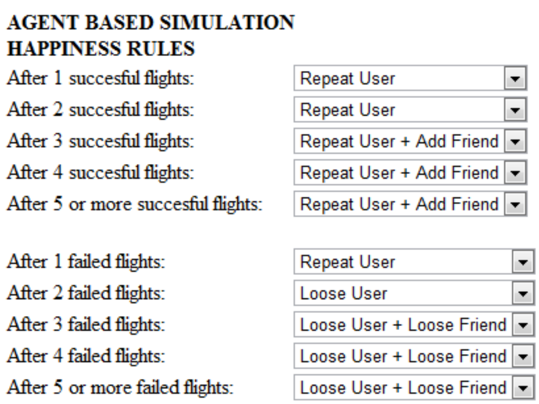

Figure 5 – Web Interface to Happiness Rules

There are two choices for the rules governing a response to successful flight experiences:

- "Repeat User" – the passenger is satisfied enough that they will choose to use the internet the next time they fly
- "Repeat User + Add Friend" – the passenger will use the internet the next time they fly, AND they are so satisfied that they convince a non-internet-user to use the service the next time they fly

There are three choices for the rules governing a response to failed flight experiences:

- "Repeat User" - the passenger remains satisfied enough that they will choose to use the internet the next time they fly
- "Loose User" – the passenger is dissatisfied and decides to not use the internet service the next time they fly
- "Loose User + Loose Friend" – the passenger decides to not use the internet service the next time they fly, AND they are so dissatisfied that they convince another internet user to also discontinue using the internet service

It should be noted that the first success for an individual in the simulation may actually be due to a referral based on their friend's success.

Agent-Based Simulation Method

Before running the main loop of the simulation, the array of potential passengers is created. An array of 80,000 people is initialized with two keys called 'success' and 'failure.' These keys are used to count the person's history of successful and failed in-flight internet experiences. The array is initially empty (all values set to zero). The initial condition (initial percentage of internet users) is created by randomly selecting 1% of the people in the array and assigning their 'success' key a value of 1.

Next, the main loop of the agent-based simulation is started. Based on the model assumptions used for the analysis, each iteration of the loop represents one day of flying one-way across the North Atlantic Corridor. The first step is to select a random group of passengers from the total pool of potential passengers. 8,000 passengers are selected from the pool of 80,000 people. Next, these passengers are reviewed to index the people who are internet users ('success' >= 1). The number of internet users is counted and cross-checked in the Monte-Carlo failure rate table to determine the day's failure rate. In addition, the total number of internet users associated with a single internet gateway failure is calculated.

Num_Users_Per_Failure = fpg * p * Internet_User_Percentage;

The failure rate is applied to the internet gateways, and the total number of failed internet users is calculated:

Total_User_Failures = round(g*Failure_Rate) * Num_Users_Per_Failure

Now we are ready to apply the "happiness rules" as defined by the study investigator. The failures are applied to the internet-using passengers according to the behavior defined by the happiness rules. A random set of the day's internet-using passengers equal to Total_User_Failures is selected and their 'failure' key is incremented by one. The happiness rules defined by the investigator are based on the number of user failures experienced by each passenger (1 to 5+ failures):

- If the happiness rule specifies "repeat user" then the user's 'success' key is reset back to "1".
- If the happiness rule specifies "loose user" then the user's 'success' key is reset back to "0".
- If the happiness rule specifies "loose user + loose friend" then the user's 'success' key is reset back to "0" AND another user who didn't experience a failure has their 'success' key reset back to "0".

Next the successes are applied to the remaining passengers that used the internet service successfully and were unaffected by the failure rules. The happiness rules defined by the investigator are based on the number of consecutive user successes experienced by each passenger (1 to 5+ successes):

- If the happiness rule specifies "repeat user" then the user's 'success' key is incremented by "1".
- If the happiness rule specifies "repeat user + add friend" then the user's 'success' key is incremented by "1" AND another random passenger that was not an internet user ('success' key = 0) has their 'success' key set to "1".

The happiness rules are applied for the 8,000 passengers of the current day. Next the total number of internet users for the entire set of 80,000 potential passengers is summed up and saved as a historical data point for the simulation iteration. The total number of daily user successes, failures, added users and lost users is also recorded for each iteration. These statistics describe the historical state of the entire system of users. After completing all iterations of the simulation, the statistics are plotted to show trends and steady-state system performance.

Agent-Based Simulation Results

The first simulation case represents the baseline for this study and is based on the following model inputs:

Number of Monte Carlo experiments:	2000	
Number of simulation iterations:	500	
ASSUMPTIONS		
Total number of people:	80000	
% of initial internet users (early adopters):	1	%
Max Internet Gateway Bandwidth:	30	Mbps
Max Network Loading before failure:	40	%
Min Flights:	30	
Max Flights:	50	
Min Passengers per Flight:	100	
Max Passengers per Flights:	300	
Min Bandwidth Used per User:	10	Kbps
Max Bandwidth Used per User:	90	Kbps
Min Internet Gateways:	2	
Max Internet Gateways:	10	

AGENT BASED SIMULATION HAPPINESS RULES

After 1 succesful flights:	Repeat User
After 2 succesful flights:	Repeat User
After 3 succesful flights:	Repeat User + Add Friend
After 4 succesful flights:	Repeat User + Add Friend
After 5 or more succesful flights:	Repeat User + Add Friend
After 1 failed flights:	Loose User + Loose Friend
After 2 failed flights:	Loose User + Loose Friend
After 3 failed flights:	Loose User + Loose Friend
After 4 failed flights:	Loose User + Loose Friend
After 5 or more failed flights:	Loose User + Loose Friend

Figure 6 – Baseline Simulation Inputs

These baseline happiness rules define that a user is not likely to refer a new friend to use the service until they have experienced three successful in-flight internet experiences. A failed experience is defined to more quickly affect the user. A single failure is assumed to cause the user to discontinue service and complain to others such that they also discontinue service even though they didn't experience a direct failure themselves.

Figure 7 – Baseline Simulation Results

The Monte-Carlo calculated failure rates are shown in the top-left corner of Figure 7. This result was discussed earlier in the paper. The main agent-based simulation result is shown in the bottom-left corner of Figure 7. This plot represents the running total of internet users from the full set of 80,000 potential passengers. As shown, there is an exponential growth in internet users until the system reaches steady-state. In this case, steady-state performance is achieved at 12,000 users, or about 15% of the total passenger population.

The other four smaller plots provide supporting views into the activity of user successes, failures, added users, and lost users across the run of simulation iterations.

So now that a baseline result has been defined, let's change the happiness rules to adjust the model of human behavior.

Free Voucher

What would happen if the airlines handed out a free internet voucher to passengers that experienced a failure? In this case we change the happiness rule such that one failed experience results in a repeat user due to the free voucher. However after two failed experiences users discontinue service even if it was free.

Only one happiness rule was adjusted from the baseline, as shown in Figure 8. This resulted in a steady-state internet user base of 17.5% of the total passenger population (14,000 users) as shown in Figure 9. This is a 17% increase in users as compared to the baseline. Therefore it is probably a good idea for the airlines to give out vouchers to increase the in-flight internet business. This is an example of how the agent-based simulation can be used to predict system level affects driven by individual incentives.

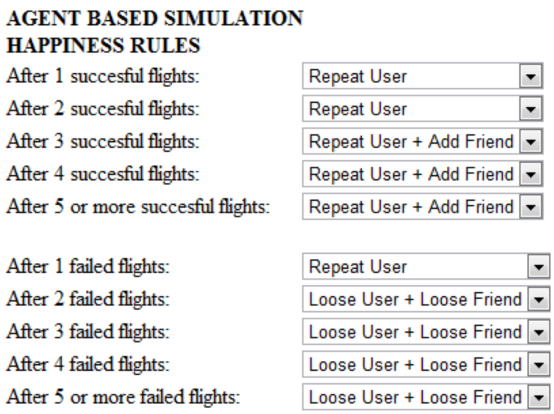

AGENT BASED SIMULATION
HAPPINESS RULES

After 1 succesful flights:	Repeat User
After 2 succesful flights:	Repeat User
After 3 succesful flights:	Repeat User + Add Friend
After 4 succesful flights:	Repeat User + Add Friend
After 5 or more succesful flights:	Repeat User + Add Friend
After 1 failed flights:	Repeat User
After 2 failed flights:	Loose User + Loose Friend
After 3 failed flights:	Loose User + Loose Friend
After 4 failed flights:	Loose User + Loose Friend
After 5 or more failed flights:	Loose User + Loose Friend

Figure 8 – Happiness Rules for Free Voucher

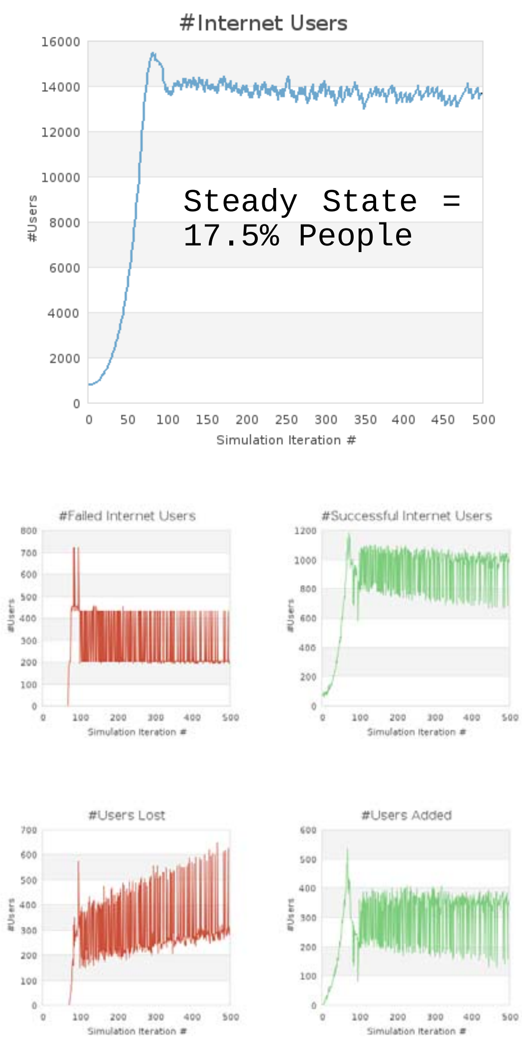

Figure 9 – Simulation Results for Free Voucher

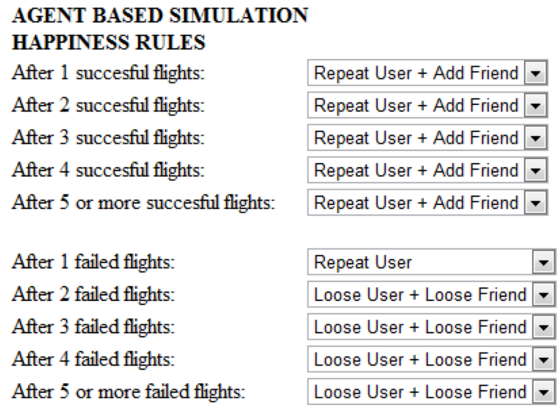

AGENT BASED SIMULATION
HAPPINESS RULES

After 1 succesful flights:	Repeat User + Add Friend
After 2 succesful flights:	Repeat User + Add Friend
After 3 succesful flights:	Repeat User + Add Friend
After 4 succesful flights:	Repeat User + Add Friend
After 5 or more succesful flights:	Repeat User + Add Friend
After 1 failed flights:	Repeat User
After 2 failed flights:	Loose User + Loose Friend
After 3 failed flights:	Loose User + Loose Friend
After 4 failed flights:	Loose User + Loose Friend
After 5 or more failed flights:	Loose User + Loose Friend

Figure 10 – Happiness Rules for Free Voucher + Free Friend Voucher

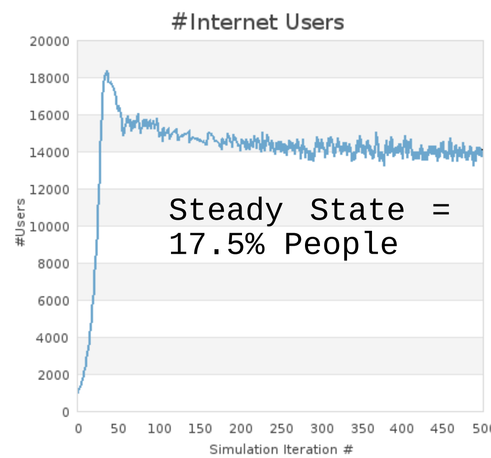

Figure 11 – Simulation Results for Free Voucher + Free Friend Voucher

Free Voucher + Free Friend Voucher

The free voucher appeared to be a good idea when failures occurred. What if the airline extended this idea by handing out free friend vouchers to their existing internet users? This would influence the users to refer a friend immediately instead of waiting for three successful experiences. This affects two more happiness rules as shown in Figure 10. While this resulted in an initial spike in internet users, the steady state performance remains at 17.5% of total passenger population (14,000 users) as shown in Figure 11. Even though friend vouchers seemed to be an intuitive incentive, the simulation shows that this would not impact the steady state system performance. It can be concluded that a free friend voucher should be avoided since it would cost the airlines money without real benefit.

Pessimistic Referral Behavior

For the next simulation run, let's go back and re-evaluate the human behavior model in the baseline. Perhaps it is too optimistic to assume that internet users will recruit other new users. Instead, let's assume that the airline will hand out a free friend voucher to first-time users, but that is the only time a user will refer a friend. In addition, let's assume the airline provides a free voucher to users that experience failures. The corresponding happiness rules for this case are shown in Figure 13. This simulation produced an interesting result with multiple critical states as shown in Figure 14. The first critical state occurs at a user level of 17.5% of total passenger population (14,000 users). Then there is a user fall-off that occurs around day 350 (iteration 350) and the second critical state follows at a level of 15% of total passenger population (12,000 users). This represents a 14% loss of users. Using these pessimistic assumptions, the simulation indicates that it would probably be good for the airlines to issue a free friend voucher on an annual basis, or whenever that user drop-off occurred in reality.

Optimistic Referral Behavior

For the next simulation, let's evaluate the system behavior based on an optimistic assumption for friend referrals. The baseline simulation was re-run, but the happiness rules were changed to optimistically assume a user would refer a friend after only two successful experiences, instead of three as assumed in the baseline (Figure 15). This simulation produced a result that exhibited an interesting emergent behavior. The internet user population was much less sporadic than in other simulations, and exhibited a saw-tooth profile as shown in Figure 16. The steady-state user level is shown at about 16.25% of total passenger population (13,000 users), but also shows a slight upward trend. It is not likely that a saw-tooth profile would be achieved in real life since the exact conditions that produced this result are not likely to be realized in precisely the same way. Nonetheless, the theoretical result is an example of unexpected emergent behavior, and practically shows that the optimistic behavioral assumptions might result in an 8% increase in users as compared to the baseline simulation.

Sensitivity of Assumptions

The minimum and maximum "bandwidth used per user" is probably one of the most important simulation assumptions to validate. The assumptions used in these simulations assume 10-90Kbps usage, with an average of 50Kbps. This is based on the referenced study of internet user behavior with satellite internet service [13]. However this does not allow for high resolution video streaming or frequent downloading of large files. Such activity is much more demanding and can easily average 100-200Kbps. Changes in this assumption have a great impact on the magnitude of the steady internet user base. To demonstrate this, the baseline simulation was repeated using the same happiness rules, but the "bandwidth used per user" was doubled: 10-210Kbps (average of 100Kbps). As shown in Figure 12, the failure rate curve shifted to the left, such that 100% failures occur at 20% user base instead of 40% (compare to Figure 4). This resulted in the internet user base dropping from 15% of total user population (12,000 users) down to about 6.9% (5,500 users). Over half the users were lost based on this single change to the assumptions! Therefore it is critical to validate the simulation and its assumptions before using the results to make significant business decisions.

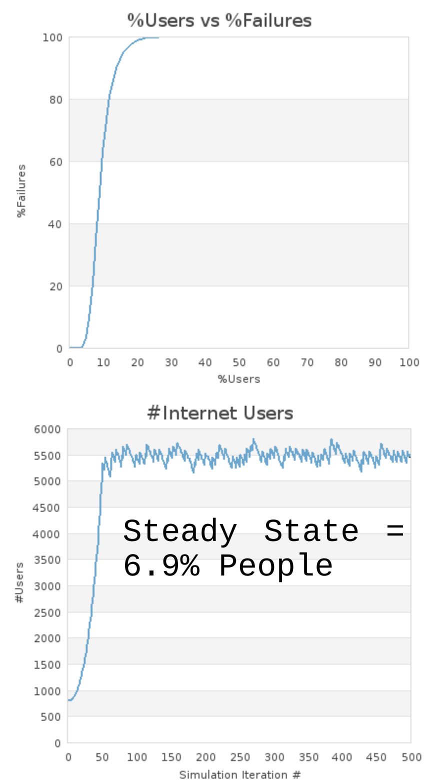

Figure 12 – Simulation Results for 10-210Kbps User Bandwidth

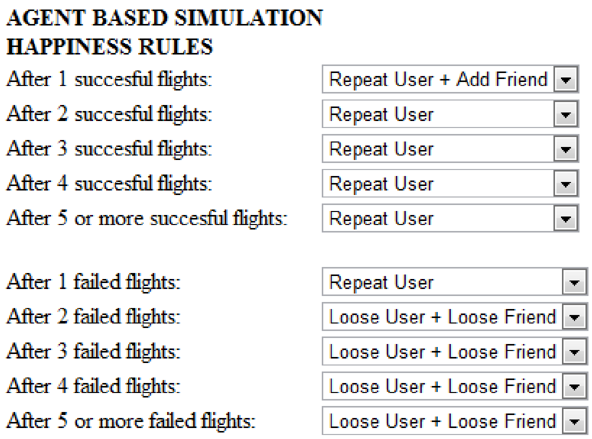

Figure 13 – Happiness Rules for Pessimistic Referral Behavior

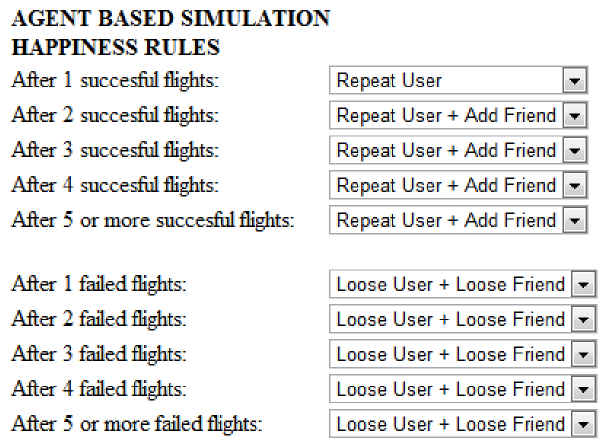

Figure 15 – Happiness Rules for Optimistic Referral Behavior

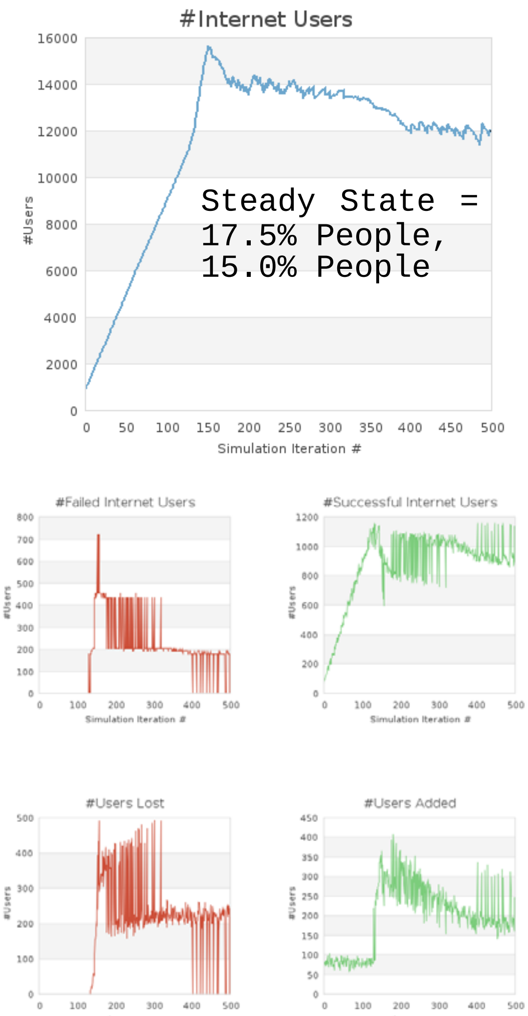

Figure 14 – Simulation Results for Pessimistic Referral Behavior

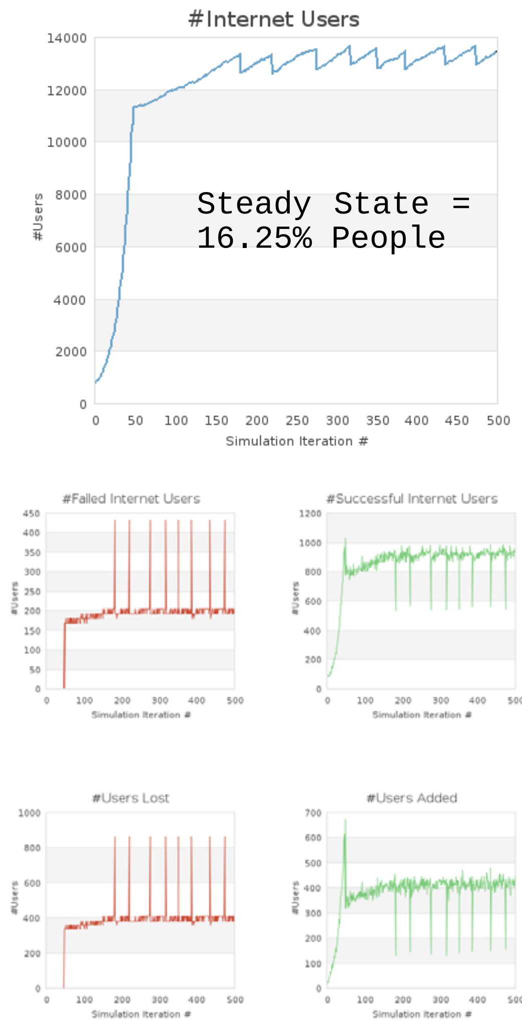

Figure 16 – Simulation Results for Optimistic Referral Behavior

Conclusions

The simulation successfully answered the main questions of this investigation. First, it calculated the failure rates that can be expected based on the size of the user base. Second, it calculated the steady state user level that can be expected for the system given these failure rates and resultant individual user satisfaction.

If the simulation assumptions are valid, then this study demonstrates that a low-cost aerial ad-hoc network is <u>not</u> likely to significantly increase the internet user base. Traditional in-flight internet access has realized a market user base of 10-15% [2-3]. The simulation in this study forecasts that an aerial ad-hoc network would realize a market user base of about 7-17%. This is an interesting conclusion. Aerial ad-hoc networks would likely drive the user price for the internet access down where financial barriers are removed. However the increase in users will tax the performance of the internet service leading to new technical barriers for the service. These financial and technical barriers appear to equally affect the system since the aerial ad-hoc network is forecast to bear about the same amount of users as traditional in-flight internet service.

Based on this conclusion, the only benefit that might be achieved is airline cost savings. An aerial ad-hoc network is expected to provide a lower operational cost to the airlines since each aircraft isn't required to maintain an expensive satellite link for the internet gateway. This savings can be shared between the airlines and the internet users. The airlines can reduce the user price down to a level where economic barriers are removed. Any additional cost savings are realized as additional airline profits. If these profits outweigh the non-recurring engineering and installation costs for an aerial ad-hoc network then the approach is still viable. Traditional in-flight internet requires a user base of 8-10% in order to be profitable within 5 years [3]. An aerial ad-hoc network may be able to provide profitability in less than 5 years due to lower operational costs. The airlines would need to complete this financial trade study, as it is outside the scope of this investigation. However this study supports the financial trade by demonstrating that the internet user base should be assumed to be about equal to what is experienced today.

Topics of Further Study

The simulation could be extended to model further complexities that impact the market potential for aerial ad-hoc networks. These represent new variables not considered by the existing model, but could be added to the base PHP model script.

Effects of Economy

The current model assumes that aerial ad-hoc networks would remove financial barriers for internet users. In actuality, the financial barriers are dependent upon the economy, and therefore the internet user base would be expected to vary in proportion to the overall health of the economy. This factor could be added to the model in order to align the simulation results with future economic forecasts. This is important since the implementation costs for an aerial ad-hoc network would require a business case that likely spanned 5-10 years.

Quality of Service (QoS)

The current model assumes that all users are provided with equal access to internet bandwidth. In the domain of computer networking, it is a common practice to implement Quality of Service (QoS) controls that constrain bandwidth on a per-user basis. This allows certain users to be guaranteed higher-than-average bandwidth at the cost of other users be constrained to lower-than-average bandwidth. In a paid internet service model, QoS can be used to create tiered pricing that grants more bandwidth to higher paying users. More specifically, QoS could be used to alleviate the concern that some users will want to consume 100-200Kbps bandwidth for activities such as video streaming. Average gateway bandwidth could be maintained without failure condition while allowing for these power users.

Model Loss of Link

As mentioned earlier, a loss of link failure is not considered in this simulation, but it could occur in the dynamic routing environment for aerial ad-hoc networks. It is assumed that the routing protocol selected would minimize this failure. If loss of link becomes a significant issue in reality, then this effect should be added to the simulation since it would lead to a higher failure rate and lower internet user base.

Uneven Aircraft Distribution among Gateways

The current model assumes that the aircraft are evenly distributed between the internet gateways. The validity of this assumption depends upon whether the routing protocol attempts to load balance the internet gateways. In reality the protocol may attempt to minimize network hops (aircraft hops) and resultant network latency at the cost of effective load balancing. In this case, a random variable could be added to the model to simulate an uneven distribution of aircraft.

Simulation Available to Public Investigators

This is a web-accessible simulation that has been made available to the public. This allows other investigators to set model assumptions and happiness rules which they believe are most valid or most interesting to study. The simulation is hosted on the principal author's website:

www.watkinsplace.com/simulations/aerialnetwork/

References

[1] www.internetworldstats.com, Nov. 2010, Internet Growth Statistics, Miniwatts Marketing Group, viewed on 24 November 2010, http://www.internetworldstats.com/emarketing.htm

[2] Reed, Dan, 6 July 2010, Less Than 10% of Fliers are Logging onto Wi-Fi in the Air, USA Today, http://www.usatoday.com/tech/wireless/2010-07-06-airlinewifi06_ST_N.htm?

[3] McCartney, Scott, 27 Aug. 2009, In the Air, Wi-Fi Gets a Ho-Hum Reception, Wall Street Journal, http://online.wsj.com/article/SB10001424052970203706604574374571364228440.html

[4] Hoffmann, Felix, Daniel Medina, Adam Wolisz, Oct. 2010, Protocol Architecture Analysis for Internet Connectivity in Aeronautical Ad Hoc Networks, IEEE/AIAA 29th Digital Avionics Systems Conference (DASC), Salt Lake City, Utah.

[5] Sakhaee, Ehssan, Abbas Jamalipour, Nei Kato, 2006, Aeronautical Ad Hoc Networks, IEEE Wireless Communications and Networking Conference, Las Vegas, Nevada

[6] Walter, Michael, Nico Franzen, Michael Schnell, October 2010, New Concepts for a Decentralized Self-Organizing Air-to-Air Radio Link, IEEE/AIAA 29th Digital Avionics Systems Conference (DASC), Salt Lake City, Utah

[7] C. E. Perkins, P. Bhagwat, Sept. 1994, Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers, in Proc. ACM SIGCOMM, pp. 234-244, London, UK.

[8] D. Johnson, D. Maltz, 1996, Dynamic Source Routing in Ad Hoc Wireless Networks, in T. Imielinski and H Korth, editors, Mobile Computing, Chapter 5. Kluwer Academic.

[9] V. D. Park, M. S. Corson, Apr. 1997, A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks, in Proc. IEEE INFOCOM, pp. 1405-1413, Kobe, Japan.

[10] Y. Ko, N. H. Vaidya, Oct. 1998, Location-Aided Routing (LAR) in Mobile Ad Hoc Networks," in Proc. IEEE/ACM MobiCom, pp. 66-75,. Dallas, Texas,USA.

[11] Zurich School of Applied Sciences, 2008, 24-hour Simulation of Actual Air Traffic, viewed on 12 October 2010, http://radar.zhaw.ch/

[12] Buyck, Cathy, 1 Mar. 2008, Open Skies at Last!, ATW Online, viewed on 5 November 2010, http://atwonline.com/print/128098

[13] Nagai, Yutaka, 11 Oct. 2002, Broadband Service via Satellite: Challenges at JSAT in Japanese Market, AIAA Japan Forum – Satellite Communications Colloquium, Houston, Texas, http://www.aiaa.org/pdf/industry/presentations/japan02nagai.pdf

[14] Row 44, 2010, Row 44 Giving Broadband Wings: Our Products, viewed on 2 November 2010, http://www.row44.com/products.php

[15] Clark, Kennedy, Kevin Hamilton, Aug. 1999, CISCO LAN Switching, CCIE Professional Development Series, Cisco Press

[16] Weisstein, Eric W., 2010, Box-Muller Transformation, From MathWorld--A Wolfram Web Resource, http://mathworld.wolfram.com/Box-MullerTransformation.html

Email Addresses

Christopher B Watkins: chris@watkinsplace.com

Cihan H Dagli: dagli@mst.edu

30th Digital Avionics Systems Conference
October 16-20, 2011