



## **A NEW ARCHITECTURE FOR FLIGHTGEAR FLIGHT SIMULATOR**

AJ MacLeod, Ampere K. Hardraade, Michael Koehne, Steve Knoblock

MVC architecture,  
FDM Server,  
FDM Instance,  
Client

To continue improving existing features and add new ones, FlightGear must make better use of computing power. Preparing for the widespread adoption of multi-core CPU architectures is an important step in FlightGear's development. Today, CPU clock rate has reached a plateau. The old idea, that features can be added without regard to their effect on performance because computers will become ever faster, have ceased to hold true. In addition, as more features are added, developers are increasingly bumping up against existing limitations in the current FlightGear architecture. Now would be a good time to begin the process of restructuring FlightGear to address the above issues. This proposal describes a new architecture for FlightGear, one which would greatly improve FlightGear's efficiency and flexibility by making extensive use of parallel processing. It is also hoped that this new architecture will improve the quality of multiuser sessions, as well as providing truer support for the simulation of time-critical systems.

## ABSTRACT

To continue improving existing features and add new ones, FlightGear must make better use of computing power. Preparing for the widespread adoption of multi-core CPU architectures is an important step in FlightGear's development. Today, CPU clock rate has reached a plateau. The old idea, that features can be added without regard to their effect on performance because computers will become ever faster, have ceased to hold true. In addition, as more features are added, developers are increasingly bumping up against existing limitations in the current FlightGear architecture. Now would be a good time to begin the process of restructuring FlightGear to address the above issues. This proposal describes a new architecture for FlightGear, one which would greatly improve FlightGear's efficiency and flexibility by making extensive use of parallel processing. It is also hoped that this new architecture will improve the quality of multiuser sessions, as well as providing truer support for the simulation of time-critical systems.

## INTRODUCTION

Since its inception, FlightGear has proved to be nothing short of a revolution in flight simulation. For perhaps the first time, a serious flight simulator has been available which allows a vast (and theoretically unlimited) degree of flexibility and extensibility. For instance, several FDMs are supported and the output data, both aeronautical and graphical, can be transmitted, utilized and displayed in a remarkable number of ways. Such flexibility and extensibility have allowed FlightGear to have an impressive following in the academic and aero-engineering world.

FlightGear has also found popularity among hobbyists. As such, the multiuser implementation and the AI traffic system have begun to progress into a very usable and important

enhancement to the simulator. In fact, the multiuser feature has proven to be so popular that it has caused an unprecedented boost in FlightGear's user base.

These developments and the simultaneous developments in the abilities of commonly available computing hardware have, however, begun to highlight some fundamental limitations in the current structure of the FlightGear simulator. As computing hardware places more and more emphasis on multiple processor cores and the threading necessary to take advantage of these, the current "main loop" of FlightGear begins to pose severe limitations on the growth of FlightGear. This can only be addressed by a restructuring of FlightGear, which, although certainly far from trivial, will offer large benefits in many different areas.

This report discusses both the current problem areas and a proposed solution, and highlights some of the benefits which will follow on from such a move.

## BACKGROUND

FlightGear is an open source flight simulator originally developed as a single user application capable of running on multiple platforms. The architecture of FlightGear is based on an infinite loop called the "main loop". Unlike the main loop of typical event-driven programs, the main loop in FlightGear is not a loop which runs in idle, polling for the user's input. In FlightGear, many components work in the background regardless of inputs. Thus, the main loop in FlightGear is responsible for periodically telling these components to update themselves. The main loop was established at the beginning of the FlightGear project; however, recent work such as the expansion of multiuser functionality and the sophisticated Nasal scripting have raised some questions regarding the current architecture of FlightGear.

Currently, the following tasks are done

sequentially in the main loop (see Figure 1):

- ATC simulation
- control of AI objects
- update of other aircrafts in a multiuser environment
- flight dynamics calculations
- scenery update
- audio scheduling
- rendering

As the loop is run, the control of the program is passed to different sections of the simulator for handling different tasks. This control is not returned to the main loop until the task is finished. To elaborate, the FDM cannot be run until the renderer is finished updating the scene; Nasal scripts cannot be run before other components finish running and the entire sim is frozen while a script is running. Consequently, the performance of each component greatly affects all others.

The most resource intensive process in FlightGear is rendering. A rendering call is made for every iteration of the main loop. Since control is passed to the renderer and is not returned until the renderer finishes with the scene update, the frequency at which the main loop can run is slaved to the framerate. It would therefore be more correct to say that the loop is run once for every rendering update. This, in turn, affects the frequency at which other components can run.

Since framerate is a function of scene complexity and not a constant value, all other components in the simulator may be updated at irregular intervals. As a pseudo closed loop control, the autopilot is one of the systems that is most sensitive to the fluctuations caused by the framerate (see Figure 2). When the framerate is low, the autopilot will overcompensate for an attitude change. This overcompensation leads to a rapid attitude change, which is in turn overcompensated by the autopilot. The result is a positive feedback ending with the autopilot plunging the aircraft into the ground. Should the autopilot be replaced with code that drives a full flight simulator, low framerate situations could result in such unpredictable behavior that not only

could cause injury to the occupants, but might also shake the simulator apart. As a result, FlightGear is undesirable for use as a serious flight simulator, and fail one of the primary goals of the project.

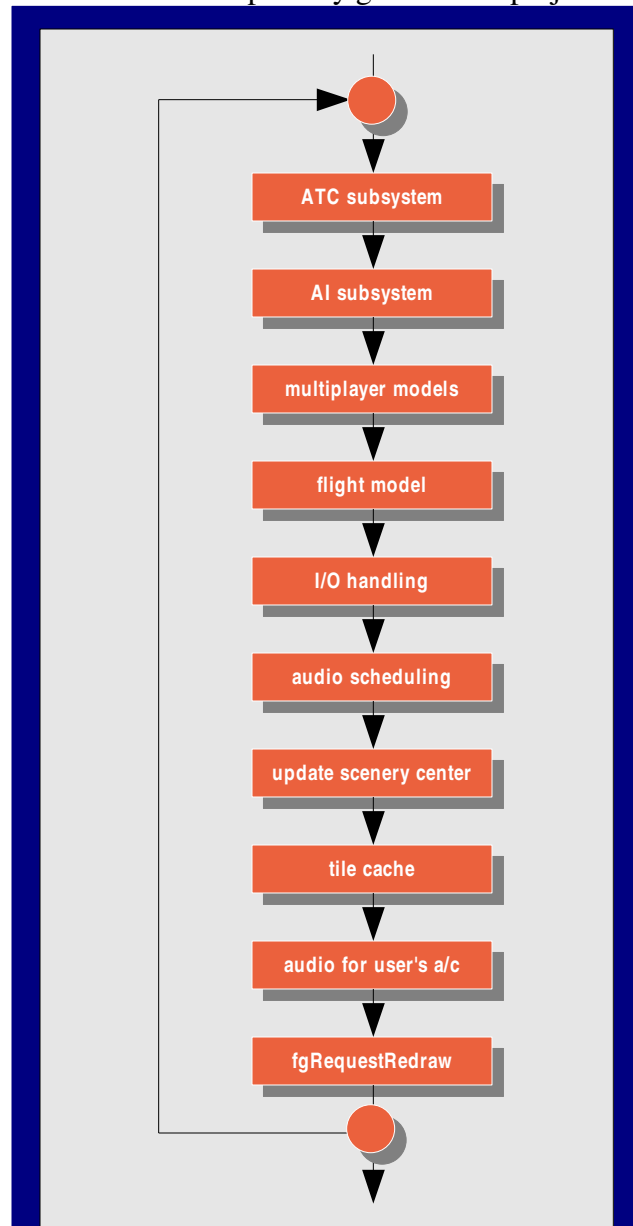


Figure 1

This condensed flowchart illustrates some of the tasks handled by the mainloop. Notice the amount of tasks that one thread has to handle. Also notice the rendering request made at the end, essentially tying the frequency of the loop to the framerate.

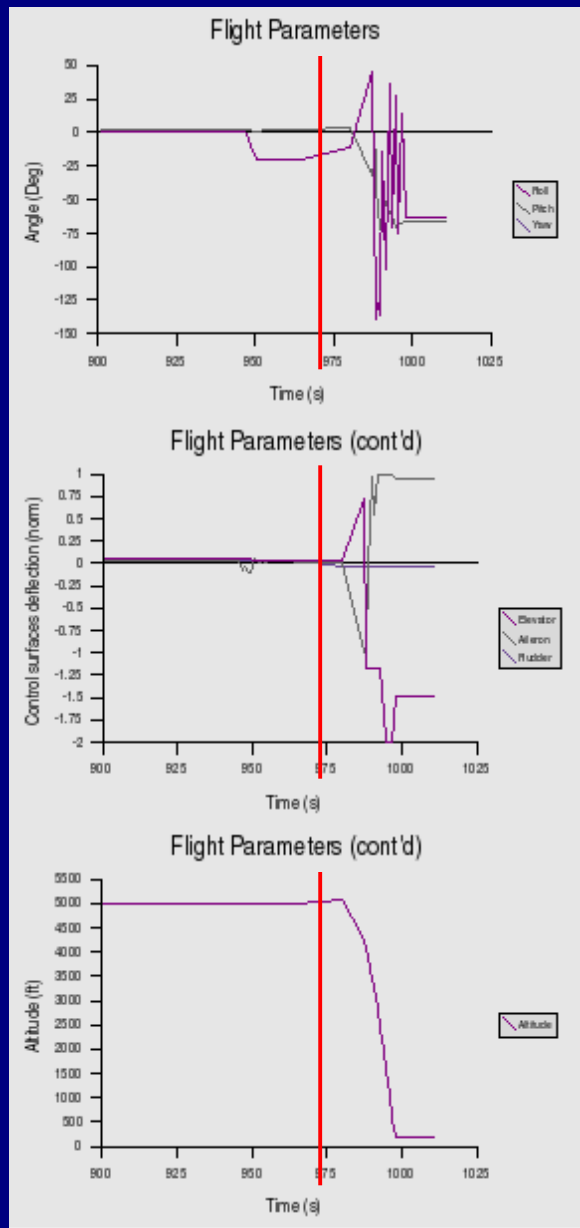
Scripts within FlightGear are not immune to the irregular update interval either. Theoretically, a Nasal script should be able to run

at a frequency of 100Hz easily. However, since every component in the simulator is tied to the framerate, the practical frequency at which a script can run is that of the scene updates.

framerate cannot be throttled, because doing so would affect the performance of the FDM engine as well. Moreover, in order to make everything run smoothly, the loop must run as quickly as possible, leading to high loading on a single processor.

Another consequence of being effectively single threaded is that all but one CPU sits idle on multi-processor platforms, forsaking other resources that could be used to improve FlightGear's performance. This problem would only become more prominent in the future, as CPU manufacturers have shifted their priority from increasing CPU clockrates to developing multi-core processors<sup>1</sup>. As more features are added and the main loop becomes ever larger, the current architecture of FlightGear would only cause the performance of future versions to decrease rather than increase.

It should also be noted that FlightGear is not currently oriented toward multiuser capabilities. This can be seen from the way movable objects in the simulator are handled. In FlightGear, the user's aircraft, AI models, and aircrafts of other users are all considered as separate classes. Feature implementation would require a great deal of redundant coding, as each feature would need to be duplicated and implemented differently for different object types. The current multiuser system also exhibits many limitations. For instance, the inner workings of AI models do not allow AI models to be updated at more than 3Hz<sup>2</sup>, thus aircraft in the multiuser environment suffer noticeable jitter. Limitations also come in the form of constraints on what can be implemented in a multiuser environment. Sophisticated features such as multi-pilot capability are extremely difficult, if not impractical to add because of the jitter and latency involved with the current system.



**Figure 2**  
 Three graphs showing the effects of framerate hit on the autopilot. The red line indicates the time of framerate drop.

With all components controlled from within one loop, FlightGear is essentially a single thread application. The present configuration makes it impossible to control the performance of individual components. For example, the

## METHODOLOGY

In designing this solution, the foremost concern is shielding the simulator from the effects of fluctuating frame rates. For FlightGear to be a

serious flight simulator in the truest sense, the core components must be run in a time-critical manner. Therefore, one of the objectives of this project is to remove the rendering process from the main loop and place it in a separate thread. Focus was also placed on restructuring FlightGear from a single-user application to a fully network-aware program, so as to benefit the development of multiuser and other complex simulator arrangements.

Emphasis was also placed on ensuring that packets be transferred and received at well regulated intervals and to unifying how objects are handled within FlightGear. The design should incorporate support for desirable features (some reserved for a future time), such as the ability for users to switch aircraft while the simulator is running, showing the same instance of an AI object across the entire network or the implementation of multi-pilot capabilities.

## PROPOSED SOLUTION

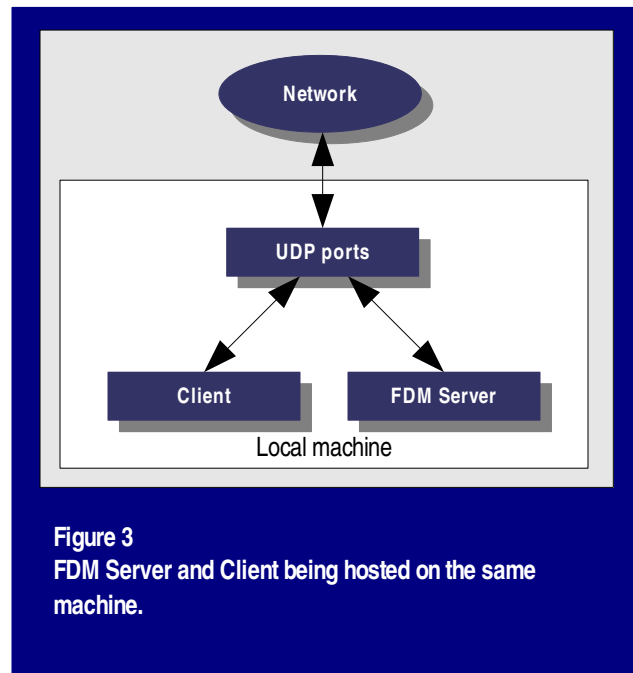
### Overview

This document proposes solutions to the various obstacles standing in the way of further FlightGear development. Specifically, the existing code base of FlightGear should be refactored using MVC (Model-View-Controller) architecture<sup>3</sup>, a software architecture in which the data model, user interface, and control logic are separated into distinct components. Under this scheme, FlightGear would be refactored into two primary independent components, one being an FDM server hosting the data model, the other a client handling the viewing and controlling components. After the refactoring process, FlightGear's architecture would bear similarities to that of the X-window system<sup>4</sup>.

While not going into too much detail, the following is a brief description introducing the model of an FDM server and client. The purpose of the FDM server is to provide services to simulate the flight dynamics and systems of multiple aircraft; in essence, a stripped down

version of FlightGear, without the graphics-related components. The client is a controller/viewer without any simulation components, which would have two responsibilities – to process inputs from the user and forward them to the server, and to listen to the servers and display objects in the scene accordingly.

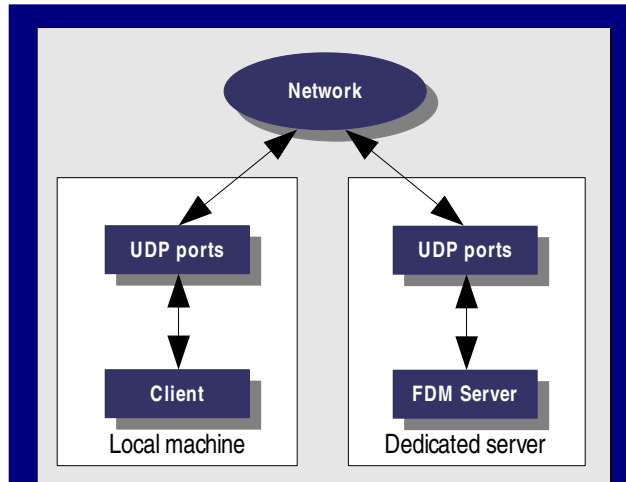
The FDM server and client would communicate property changes through UDP ports. Typically, the server and client would be hosted on the same machine (see Figure 3), although they could also be hosted on different machines (see Figure 4). Inputs from the user are first processed by the clients, and then sent to the server. The server would then make adjustments to the FDM based on the information it has received from the client.



**Figure 3**  
FDM Server and Client being hosted on the same machine.

To reduce the chance of pilot-induced-oscillation, the frequency at which the client uploads the user's inputs must be higher than the typical human reaction time of 0.25s. A suggested uplink frequency would be above 20Hz for the setup where the server and the client are hosted on different computers, but this frequency could be higher when the server and the clients are being hosted on the same computer. The frequency of downlink – the frequency at which the server broadcasts updates - would be dependent on the aircraft being flown, although for most aircraft,

this value would be 10Hz.



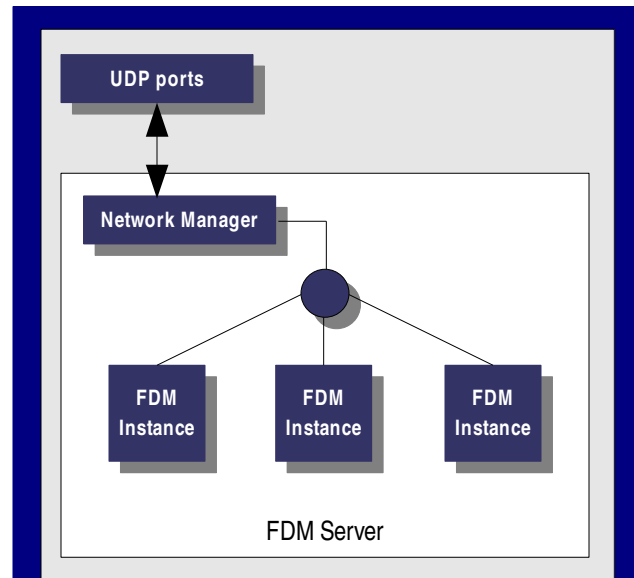
**Figure 4**  
FDM Server and Client being hosted on different machine.

### FDM Server

The FDM server is a component dedicated mainly to providing flight dynamics simulation services to the clients. Its responsibility would include handling communication between the clients and the FDM, as well as communicating the position of various aircraft to a multi-player server. As a server, it would be capable of running multiple FDMs simultaneously, and if necessary, merge or average the outputs of these FDMs.

The actual FDM would be encapsulated in a component called the FDM instance within the server (see Figure 5). This FDM instance (see Figure 6) is the “model” portion in the MVC architecture, and each would have an independent property tree hosting inputs, outputs, environmental variables, and other information about the aircraft being flown, as well as information about the clients that are related to the FDM being run. Multiple instances could be run on the server so that services could be provided to multiple clients. Since the instances are completely independent from one another, partitioning could also be achieved. For example, a corrupted property tree from one instance would not pose problems to other instances.

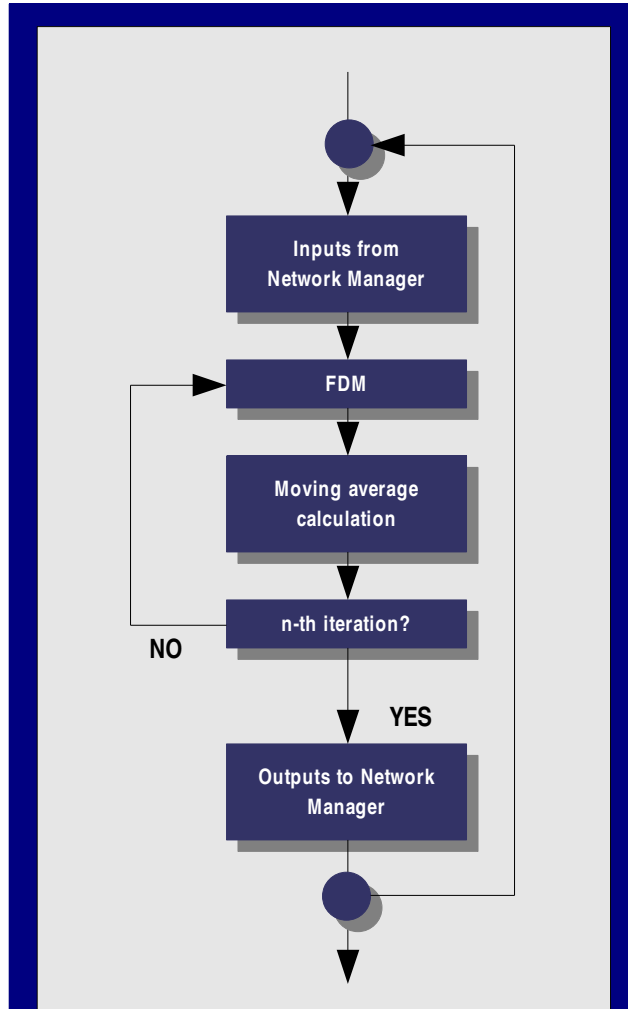
Each FDM instance would carry a unique ID, which would serve two purposes: to associate the correct AI model in the viewer with the FDM, and to identify the aircraft being simulated in a multiuser environment. The ID of two instances could be swapped so as to allow the flight dynamics of an aircraft flown by the client to change. This feature could allow a user to change aircraft while the simulator is running, or to drastically change the handling characteristics of an aircraft in order to simulate the effects of structural failures.



**Figure 5**  
The FDM Server would be capable of running multiple FDM Instances. The circle and lines represent communication channels between various threads.

An FDM server could be dedicated to hosting AI objects such as the existing aircraft carrier, which would allow all users to see the same instance of the AI objects in a multiuser environment. In addition, the FDM server could provide a very unique feature: planes that are not being used could be represented by null FDMs, with an external “AI server” updating their positions. An airliner flown by a user could be turned around and send back to its home airport automatically. Finally, hosting FDMs on a dedicated server could allow damage to aircraft to persist across multiple sessions. As an example, if

a pilot landed a plane too hard, the landing gear of this plane could collapse on its next landing in a different session.



**Figure 6**  
A conceptual work flow of the FDM Instance. The purpose of the “moving average calculation” is to filter out fluctuations from the FDM, so as to reduce the amount of jitter observed in multiuser mode.

### Robustness of FDM Server

Since the server would be expected to provide service to many users simultaneously, the server must be very robust and capable of surviving as well as recovering from multiple failures.

Due to the amount of uncertainty involved,

one cannot trust the robustness of the code encapsulated by an FDM instance. In other words, an assumption has to be made that the encapsulated code could eventually pose problems. One method of mitigating such problems is to avoid them.

Let's assume that a memory leak occurs in the code being run by one instance. Repeatedly creating and destroying this instance would result in pointers being created and not subsequently freed. Eventually memory would run out, and the FDM server would likely be killed by the operating system. This would not only affect the client hosted by the particular instance, it would also affect all clients whose FDMs are being hosted on the server. Therefore, to contain the effects of memory leaks, only a fixed number of instances would be initialized for any session of the server. In other words, no FDM instance would be created or destroyed while the server is running. In such a scheme, the mistake of creating a pointer and not freeing it would only be made once instead of multiple times. The unused FDM instances could lie dormant and be skipped by the server.

Of course, errors are impossible to avoid completely. When an unhandled exception occurs, the FDM server must be able to recover from it: specifically, methods would be needed to contain the errors and prevent them from migrating outward from the encapsulated instance. One example of a mitigation technique would be to restart an erroneous component from its last error-free state. When multiple errors have occurred, the entire FDM instance could be restarted. Likewise, if an FDM instance has an error, the FDM server could restart the instance and let it continue from the last error-free state. After multiple failures, the problematic FDM instance could be shutdown permanently.

The FDM server would also need the capability to transfer the FDMs it is running to another trusted server, for many reasons. A client might want to switch to another server to avoid network latency, or the server may have encountered too many errors and is no longer

suitable for providing service. Whatever the reason, the changeover should not be noticeable at the client end.

## Client

The client (also known as viewer/controller in MVC terminology) is the portion with which the user would interact. The client would provide visual and audio cues to the user and manage communications between the user and the FDM server.

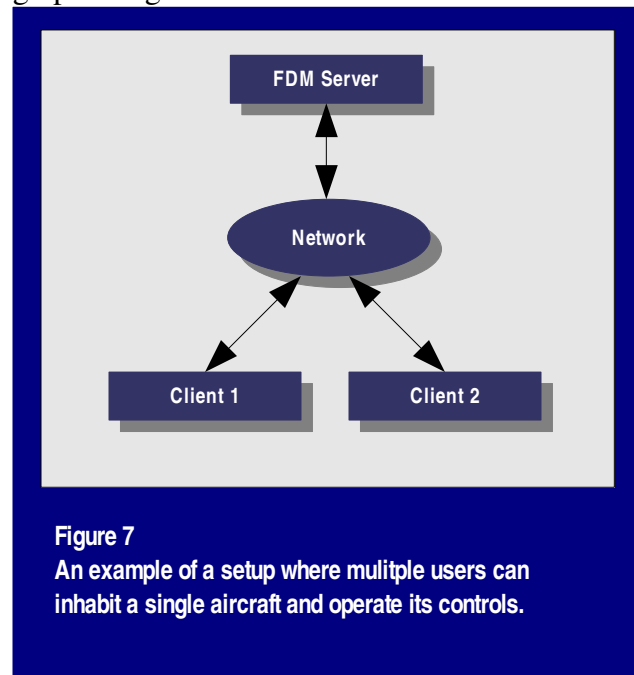
The client would require two threaded loops, one handling I/O and another rendering. Leaving aside the I/O handler for the moment, one can simply view the client as an “observer” that displays aircraft and objects currently existing in the multiuser “world”. The point of this design is to unify into a single object framework, with the user's aircraft, the AI aircraft and other users' aircraft in the multiuser environment represented as a single class of objects. This framework would create the potential for the sub-classing of any specialized movable object types, which will inherit shared characteristics from the parent class.

The I/O loop handles the communication between client and server. In addition, the I/O loop would also manage property changes, as well as performing audio scheduling. The rendering loop would be dedicated solely to scene updates. Having two loops would avoid interference to data transfer posed by fluctuating framerates. With the renderer within its own thread, framerate throttling could also be implemented. A hard limit could be set (using the sleep() function), which would maintain the framerate below a certain peak value. As an alternative, a “soft limit” could be provided, which would vary the framerate by adjusting the complexity of the rendered scene.

Separating and encapsulating all visual rendering into a “viewer” allows for multiple cameras. Cameras could be assigned different views into the FlightGear environment or assigned to follow a number of particular aircraft, thus opening up the possibility of allowing multiple users to inhabit a single aircraft and operate its

controls (see Figure 7). This latter feature is critical to realistic simulation of certain aircraft operations, such as the B-29 or Clipper. Moreover, the capability to have multiple viewers interacting in a multiuser environment would allow FlightGear to be used as an online teaching tool. An instructor may fly along with the student pilot, by flying within the same aircraft (using the copilot's seat view.)

Although it may be something best left to the future, the new architecture would make it possible for each client to utilize a different graphic engine.



**Figure 7**  
An example of a setup where multiple users can inhabit a single aircraft and operate its controls.



## SUMMARY

FlightGear is experiencing a rapid proliferation of sophisticated features. However, the current architecture of FlightGear prevents the simulator from utilizing available resources on multi-processors platforms, as well as on future multi-core CPUs. In addition, the framerate has influence on every aspect of FlightGear, causing the outputs of the simulator to be unpredictable and making the assumption of correct simulation of time-critical systems an impossibility.

The proposed new architecture would address these problems through extensive use of parallel processing. By moving to a Model-View-Controller architecture, code division within FlightGear would become more clear and better organized. Flexibility would also increase, since through networking, many new features could be “added” without having to tamper directly in FlightGear's source code. Fewer changes to the FlightGear's source code would also mean that the simulator would be less error prone.

The multiuser experience would also be enhanced, as very little effort would be needed to utilize the new architecture to allow multiple users to control a single aircraft. AI simulation could be hosted from a central location on a network, ensuring the same AI objects are displayed in all clients that are connected. In the future, ATC and weather systems could be hosted in the same manner, ensuring that all users fly in a consistent weather and air traffic environment.

FlightGear has clearly demonstrated how well an Open Source Flight Simulator can work. The past decade of development has allowed us to gain much knowledge and experience in the area of flight simulation. Post-1.0 release would be an ideal time to make use of the knowledge gained to provide a step change in the flexibility and opportunities offered by FlightGear.

## REFERENCES

1. Herb Sutter (2005). The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software. Retrieved March 19, 2006.  
Website:  
<http://www.gotw.ca/publications/concurrency-ddj.htm>
2. pigeond (2006). A list of videos captured from FlightGear's multiuser sessions. Retrieved March 19, 2006.  
Website:  
<http://www.pigeond.net/photos/FlightGear/videos/>
3. Wikipedia (2006). Model View Controller. Retrieved March 19, 2006.  
Website:  
[http://en.wikipedia.org/wiki/Model\\_view\\_controller](http://en.wikipedia.org/wiki/Model_view_controller)
4. Wikipedia (2006). X Window System protocols and architecture. Retrieved March 19, 2006.  
Website:  
[http://en.wikipedia.org/wiki/X\\_Window\\_System\\_protocols\\_and\\_architecture](http://en.wikipedia.org/wiki/X_Window_System_protocols_and_architecture)