Soft Walls - Frequently Asked Questions

Draft 1

July 21, 2003

Technical Memorandum UCB/ERL M03/31,
University of California , Berkeley, CA 94720
http://ptolemy.eecs.berkeley.edu/papers/03/softwalls/

Edward A. Lee, UC Berkeley, eal@eecs.berkeley.edu

Soft Walls

Background

Since its introduction shortly after September 11, 2001, the Soft Walls concept has generated considerable controversy and discussion. This paper collects frequently raised objections to the concept and presents a discussion of the objections. I cannot claim that this discussion is unbiased, but I have made every attempt to be fair and objective. If you are not familiar with the Soft Walls concept, please see http://ptolemy.eecs.berkeley.edu/projects/softwalls or [6] . Here, I simply state the questions, followed by a discussion. A glossary is provided at the end.

In brief, modern aircraft all have electronics on board that is involved with the control and navigation of the aircraft. Many of the newer planes have computers on board that mediate the commands issued by the pilot and translate those commands into action, for example to bank and turn to the right. It is possible to modify the software in the computers in such a way that an airplane will refuse to enter pre-specified regions. We call these regions "no-fly zones" and we call the boundaries of these regions "Soft Walls." If an aircraft is equipped with the Soft Walls system, then if the pilot attempts to enter a no-fly zone, the airplane will be diverted. This happens gently at first, but if the pilot does not cooperate, then the system becomes more assertive. The key principle is to give the pilot as much control over the aircraft as is consistent with the constraint that the airplane does not enter the no-fly zone.

1. Hasn't this terrorism problem solved itself?

Prior to Sept. 11, airline personnel were trained to deal with hijackers by cooperating, and the prevailing wisdom was that passengers should also cooperate. It was believed that this would maximize the likelihood of getting the aircraft safely on the ground. However, the suicidal intent of the September 11 hijackers changed everything, and it is very unlikely that passengers will ever again passively accept a hijacking. If there is a recurrence, then passengers are likely to use what is now called the "let's roll" defense, where they will fight the hijackers. This appears to argue that a system like Soft Walls is not needed.

Apparently, the Pentagon does not agree, since critical sites in Washington DC are now protected by antiaircraft batteries. The mere presence of this protection scheme poses significant risk to pilots, crew, and passengers, possibly more than the risk of another September 11-style hijacking.

Moreover, the "let's roll" defense does not apply to, for example, cargo aircraft, which can be just as lethal as passenger aircraft. In the October 2001 issue of *Forbes*, Peter Huber says [5]:

> "A fully fueled jumbo jet is about as lethal an instrument as ever gets entrusted to civilian hands. Nuclear power plants and big hydroelectric dams are far safer from the get-go--they don't have to fly, so they can be encased in vast excesses of concrete, and indeed they are. Assaults with nuclear or biological weapons can't begin with cardboard cutters. They require substantial factories somewhere in the background, which can, we must hope, be identified and knocked out well before they get up and running. There are many other potential instruments of terror, but none quite so essential, ubiquitous and--now--terrifying, as civilian jets."

2. Can Soft Walls enhance air safety, absent terrorism?

Aviation experts use the euphemistic term "controlled flight into terrain" for accidents like that of the American Airlines Boeing 757 in Colombia in 1995 that crashed into a mountain and the Korean Air 747 that crashed short of a runway in Guam in 1997. In the Soft Walls system, the database of no-fly zones can include not just critical infrastructure and urban areas, but also terrain obstacles such as mountains, and, of course, the ground. Thus, the same system that keeps the airplane from flying into buildings can keep it from flying into mountains.

3. Isn't reducing pilot authority dangerous?

Pilots need the authority to respond to emergencies, including unexpected weather conditions, possible collisions with other aircraft or other obstacles, turbulence, on-board equipment failures, fires, or other problems. A pilot's responsibility, however, extends beyond the craft, crew and passengers to the people on the ground. No on-board emergency is severe enough to justify endangering large numbers of people on ground. The principle of Soft Walls is that pilots should be given the maximum maneuvering room subject to the constraint that the aircraft not enter the no-fly zones. If the no-fly zones are defined with some restraint, then they represent exactly those regions where no on-board emergency can be severe enough to justify entering them.

A New York Times article in April of 2002 examined this issue [9]:

> "A Boeing 737 pilot for a major airline recalled approaching Reagan National Airport from the south a few years ago and facing a microburst, a rainstorm that includes sudden changes in wind direction. Such a condition can lead to a crash if a plane is at low altitude and low air speed, as it is on approach.
>
> He broke off the approach and turned east. "It was the only way to go," he said. However, if he had been a little deeper into the approach, he said, "I'd be flying right toward the protected area," the forbidden zone that includes the White House. A system that prevented him from turning that way would be unsafe, said the pilot, whose airline, like most, has been reluctant to discuss security changes."

Today, that plane would be shot down. So this pilot was wrong. The absence of the system is far more unsafe. No microburst is as dangerous as a modern surface-to-air missile. With Soft Walls, this pilot would have maximum maneuverability, and there would be no need to shoot down the plane (assuming that the military has confidence in the system).

Again, Peter Huber, in *Forbes* advocates approaches like Soft Walls [5]:

> "Giving a computer certain powers over the controls curtails pilot autonomy, of course--even during emergencies. But no emergency, however grave, justifies a trajectory into the heart of Manhattan, and F-16s will end up enforcing no-fly zones with air-to-air missiles if cockpit computers don't. Link the computers with on-board collision avoidance systems, and eventually with those that superintend flight paths from the ground, and air travel will end up safer still--even safer than it already is, at its best, today."

4. Isn't a more complex flight control system more likely to fail?

Anytime the complexity of engineering system increases, so does the difficulty of maintaining reliability. To many people, it seems that more complex systems are more likely to fail. However, there are many counterexamples. Automobiles, for example, are far more complex

and far more reliable today than they were 30 years ago. Much of that reliability stems from the use of computer control. The safety of aircraft has also continually improved, despite (or, in fact, partly because of) increased complexity.

Soft Walls is a computer controlled system. Regrettably, the direct interaction most people have with computers is through desktop machines, which hardly meet even the most lax reliability requirements. Failures are extremely common and are widely tolerated. However, the methods used to design embedded software in safety critical systems are (and have to be) significantly different from the methods used to design application software for desktop computers. The software systems in your car only very rarely fail, much more rarely than the purely mechanical systems they replace used to fail.

Again, from New York Times article in April of 2002 [9]:

> "Another pilot, Stephen A. Luckey, the chairman of the Air Line Pilots Association's security committee, said computers were still too prone to failure to allow them to override human pilots without recourse. Captain Luckey, who flew a Boeing 747-400 until he retired recently, said that in the last three years, he had seen three instances of computer failure on planes he was flying. In two of those, he said, as he looked at a navigation screen, the ground appeared to rotate 120 degrees."

Yes, computers and software can fail. However, the benefits of computer control overwhelm the negatives, and "fly-by-wire" aircraft will eventually overwhelmingly dominate the fleet. In these aircraft, there are no mechanical or hydraulic couplings between the cockpit controls and the control surfaces of the aircraft. All pilot commands are mediated by computers. A significant portion of the commercial fleet is already fly-by-wire, and very likely, nearly all passenger aircraft that are designed in the future will be computer controlled.

Of course, if the computers or software on such an aircraft fail, then the craft becomes completely uncontrollable. The premise in fly-by-wire design is that the computers and software can be made more reliable and more robust than the mechanical systems they replace. This requires careful engineering, of course, and, appropriately, these systems are subject to far more stringent validation than desktop software. But such validation is done, and experience so far indicates that fly-by-wire airplanes perform extremely well. They are the wave of the future.

5. How does Soft Walls relate to fly-by-wire?

Soft Walls is easiest to deploy in fly-by-wire aircraft because it is "just" a software change. Of course, software changes in such aircraft are far from easy because of the extreme safety concerns and the resultant validation requirements, but nonetheless, it makes the change technically simpler.

In fly-by-wire aircraft, pilots control the craft only through computers. There are no direct mechanical or hydraulic linkages between the cockpit controls and the control surfaces of the aircraft. Airbus describes it well on their web page, http://www.airbus.com/media/fly_by.asp, in an article entitled, tellingly, "Our Advantages – Fly-by-wire:"

> "Fly-by-wire is an electronically managed flight control system, which uses computers to make aircraft easier to handle while further enhancing safety. First introduced on a commercial jetliner on the Airbus A320 in 1988, it has become an industry standard.
>
> Pilots manoeuvre their aircraft by controlling the moveable parts, known as flight control surfaces, on the aircraft's wings and tail plane. Fly-by-wire replaces the mechanical

linkage between the pilot's cockpit controls and the moving surfaces by lighter electrical wires – hence its name.

At the heart of the system are computers that convert the pilot's commands into electrical impulses delivered to the control surfaces.

When this technology, already used extensively on fighter aircraft, was first used on the A320, it was a major achievement for several reasons. Firstly it reduced the weight of aircraft – and therefore the amount of fuel consumed. This in turn lowered operating costs for airlines and benefited the environment by reducing exhaust gas.

Fly-by-wire technology also provided a considerable safety enhancement with the introduction of hard protection. Indeed, the pilot's commands to the control surfaces are monitored to ensure the aircraft is kept within a safety margin, called the 'flight protection envelope'. Thus, the pilot can always get the maximum out of the aircraft in an emergency without running the risk of exceeding the flight envelope or over-stressing the aircraft.

Finally, fly-by-wire technology has also made it possible for Airbus to develop a true family of aircraft, from the 107-seat A318 to the 555-seat A380, with near identical cockpit designs and handling characteristics. This makes crew training and conversion shorter, simpler and highly cost-effective for airlines and allows pilots to remain current on more than one type simultaneously."

The Boeing 777 is also a fly-by-wire aircraft. Many military aircraft are fly-by-wire. The Concorde was a fly-by-wire aircraft. Most future aircraft are likely to be fly-by-wire.

6. How does Soft Walls relate to flight envelope protection?

As explained above, fly-by-wire aircraft have efficiency advantages over more conventional mechanical and hydraulic control systems. But because control is mediated by computer, such systems can also be made more intelligent. Airbus systems impose flight envelope protection schemes, where the computers ensure that the pilot does not force the aircraft beyond its safe performance parameters. For example, the computers can prevent the pilot from stalling the aircraft.

Flight envelope protection works very synergistically with Soft Walls. In particular, Soft Walls works by introducing a bias into the commands issued by the pilot when the aircraft approaches too close to a no-fly zone. To ensure that the aircraft does not enter the no-fly zone, the bias needs to increase as the craft gets closer until the bias overwhelms the commands that the pilot can issue. For instance, when the aircraft has penetrated the boundary sufficiently to be very close to the no-fly zone, the pilot may be commanding a hard turn to the right, but the bias will nonetheless force the aircraft to turn to the left, away from the no-fly zone.

In aircraft with flight envelope protection, as for example most Airbus planes, the limits on pilot induced maneuvers are known (because they are imposed by the on-board computers). Thus, the extent of the bias that must be applied is known.

Not all fly-by-wire aircraft have flight envelope protection. The Boeing 777, in particular, does not. The computers will permit the pilot to make maneuvers that exceed the safety specifications of the aircraft. Boeing argues that this is safer than flight envelope protection because these safety specifications conservative anyway, so allowing the pilot to exceed them gives the pilot the authority to consider and compare the risks in responding to an emergency.

Both approaches have their merits, but Boeing's approach requires that a Soft Walls system be more aggressive. In particular, for example, since there is no fixed limit on bank angle, there is no single amount of bias on bank angle that is guaranteed to exceed the pilot command. This complicates the design of the Soft Walls system, which must ensure that the bias it introduces does not take the aircraft outside the safety specifications.

To some degree, a Soft Walls system must realize some flight envelope protection. For example, if an aircraft is flying above a no-fly zone, then the Soft Walls system must prevent the pilot from stalling the aircraft. If it does not, then it cannot ensure that the aircraft will not enter the no-fly zone (because the stall could lead to loss of control).

7. Can Soft Walls be deployed on non-fly-by-wire aircraft?

In fly-by-wire aircraft, Soft Walls is "just" a software change. However, only a fraction of the fleet today is fly-by-wire. From the New York Times, April 2002 [9]:

> "In November, the F.A.A. counted about 2,300 fly-by-wire planes among Boeing and Airbus models, the two most popular among big jets; another 8,700 planes in those fleets had conventional mechanical systems.
>
> Herman A. Rediess, director of the Office of Aviation Research at the F.A.A., said in a paper representing his own views: "For the near future, no airline will have the financial resources to even modify the F.B.W. aircraft. It's not clear that they would even have sufficient funds to retrofit the non-F.B.W. aircraft."
>
> Adding fly-by-wire ability to older planes would be wildly expensive. George K. Muellner, an Air Force veteran and president of Boeing's research and development arm, called the Phantom Works, recalled that the Air Force had taken some of its oldest F-4's and converted them into pilotless drones, for use as target practice. The conversion, he said, cost more than the plane did new."

Converting older aircraft to fly-by-wire is clearly out of the question. However, there is an alternative, which is to modify the autopilot systems in older aircraft to implement fly-by-wire. The effectiveness of this strategy is still an open question (see the next question).

8. Can Soft Walls be realized as part of the autopilot system?

One option for deploying Soft Walls on aircraft that are not fly-by-wire is to modify the autopilot system. However, this implies some significant changes. An auto-pilot system is typically on or off, and if it is on, the pilot does not attempt to directly control the aircraft. To get the effect of the Soft Walls bias, the auto-pilot system would have to be modified to blend a bias with pilot commands.

In the Soft Walls proposal, the no-fly zones are enforced by adding a bias into the pilot control commands. A similar system could be constructed that enforces the no-fly zones in a slightly different way, by using motors to drive the pilot controls (such as the stick and pedals), applying a force that moves them in the desired direction. A cooperative pilot would simply yield to the force. An uncooperative pilot would attempt to counteract the force, but as the aircraft penetrate further into the buffer zone around the no-fly zone, the force increases until the pilot cannot overwhelm it.

This approach is likely to work well with a cooperative pilot. However, an uncooperative pilot (or one who is either confused or attempting to evade an obstacle or weather problem) would end up arm wrestling the controls. This may result in highly erratic aircraft behavior as the

pilot applies greater force, and may by itself cause an accident. In particular, when the force feedback being applied is substantial, it may become difficult (or impossible) for the pilot to exercise sufficiently precise commands for example to evade another aircraft.

There have been a number of experiments with force feedback in military aircraft, and some observers have attributed at least one crash to the pilot fighting with the force feedback system. It may be possible to avert this eventuality by giving the pilot better information about the situation. This approach is being explored by a research team in The Netherlands at the Technical University in Delft (see [2][4]). This team is using pilots in a simulator to evaluate the approach.

9. Doesn't the crew need an override?

The surest way to make the Soft Walls system effective is to prohibit override in any form. Manual override on the aircraft is certainly out of the question. Override from the ground is perhaps doable, but the security of the communications becomes a problem, and the human authorization of the override creates a vulnerability.

Some pilots argue that every cockpit system needs an override. What if, for example, there is a fire, and power needs to be cut to some subsystem? In fact, current aircraft design permits pilots to turn off most aircraft electronics. However, this can go too far.

On September 11, 2001, the hijackers turned off the transponders in their hijacked airplanes. This delayed detection of their intent, possibly preventing a successful intercept. The transponder is a device carried by essentially all modern aircraft that identifies the aircraft to the air traffic control system and specifies its location and velocity. In retrospect, it is clear that the risk posed by allowing pilots to turn off the transponder greatly exceeds the risk posed by the transponder itself. With 20-20 hindsight, it is unconscionable to allow pilots to turn off the transponder.

More fundamentally, with the trend towards fly-by-wire aircraft, turning off cockpit electronics becomes impossible without causing a crash. Fly-by-wire aircraft cannot fly without electronics. It is incumbent on aircraft engineers to make these electronics systems sufficiently robust that the risk they pose is smaller than the risks posed by the systems they replace. Indeed, this has been their charter, and all evidence points to success, so far. Fly-by-wire aircraft appear to be very safe indeed.

10. Does Soft Walls increase the risk of collisions between aircraft?

Anytime the pilot's control of an aircraft is impaired, collision with other aircraft in the vicinity becomes more likely. However, this risk occurs only if two or more aircraft are simultaneously approaching a no-fly zone. Under normal circumstances, it would ideally be years between events where an aircraft approaches a no-fly zone. This makes the increased risk of collision very much smaller.

Moreover, the Soft Walls principle maximizes pilot authority, subject to the constraint that the aircraft not enter the no-fly zone. This allows the pilot to still retain fine-grain maneuverability. This maneuverability may be sufficient to prevent collisions in the unlikely event that two aircraft are simultaneously approaching a no-fly zone.

11. How is Soft Walls related to collision avoidance systems?

Most passenger aircraft today carry ACAS or TCAS systems (airborne/traffic collision avoidance system). These are exemplary of a family of advisory systems that improve aircraft safety by monitoring the situation around the aircraft and recommending to the pilot evasive maneuvers when a threat is detected. ACAS and TCAS rely on properly equipped transponders in other aircraft, using the information provided by those transponders to identify situations in which mid-air collisions are imminent.

Unlike Soft Walls, ACAS and TCAS are advisory. They simply recommend action to the pilot. They do not interfere with the control of the aircraft. In fact, this advisory nature has been blamed for at least one crash. On July 1, 2002, a Bashkirian Airlines Tupolev Tu-154 collided with a DHL Boeing 757 over southern Germany. According to Aviation International News [7]:

> "According to initial reports on the 757/Tu-154 collision, the DHL 757 pilots followed TCAS resolution advisories to descend, but the Russian pilots ignored the commands of their ACAS to climb and instead obeyed the Swiss controller's instruction to descend– with tragic results."

One result has been a re-examination of the policy of making such systems advisory. Future systems may impose evasive actions on pilots rather than advising them. If this comes about, then it may be practical to permit such collision avoidance systems to override the Soft Walls systems for brief periods of time, long enough to avoid a collision. By nature, collision avoidance systems deal with smaller targets (other aircraft) than Soft Walls (no-fly zones), and hence operate at finer granularity.

12. How is Soft Walls related to ground proximity warning systems?

In order to address what aviation experts call "controlled flight into terrain," thousands of planes have been equipped with a system called an enhanced ground-proximity warning system ("groundprox"), which includes much of the Soft Walls idea. Groundprox systems rely on GPS (the global positioning system) or other localization information available to the system to compare the airplane's location to a database identifying where the ground is. If the system determines that the craft is headed for a mountain or towards the ground, it creates a display for the pilot showing the terrain and recommends to the pilot evasive maneuvers.

However, like TCAS and ACAS, groundprox systems are advisory. They recommend actions but do not enforce them. Nonetheless, it seems obvious that starting with groundprox systems and adapting them to realize Soft Walls is a reasonable strategy.

13. Wouldn't control from the ground be preferable?

It is technically possible to control aircraft from the ground. Northrop Grumann's Global Hawk aircraft is an unoccupied air vehicle (UAV) that is controlled from the ground. It flies without a pilot, and played a significant role in the recent Afghan and Iraq wars. Northrop Grumann has argued that the control system of Global Hawk could be adapted to permit controllers on the ground to take over an airplane and fly it safely to landing.

While technically feasible, this approach is probably more complex than Soft Walls, and it opens new vulnerabilities. For one, it creates the possibility of a hijacking from the ground, which suggests that sites equipped to take over aircraft would require serious protection, and personnel with access would be have to be severely vetted. Moreover, it creates a truly scary prospect of a wholesale hijacking of an entire fleet.

A second problem is that communication delays and lack of visibility into conditions on the aircraft make fine-grain control much more difficult. For example, collision avoidance maneuvers would be hard to execute. The system would therefore probably require non-advisory collision avoidance systems to be installed.

Moreover, authorization for takeover of an aircraft would have to be very carefully granted. This is likely to create human-in-the-loop delays that may make it impossible to prevent, for example, an aircraft approach Reagan National Airport from diverting to hit the Pentagon.

14. Wouldn't fully automatic control be preferable?

It is technically possible for an airplane to fly without any human intervention at all. It could be programmed with a sequence of waypoints to follow. Fully automatic landing systems are already available in many aircraft, although they are rarely used.

An extreme proposal is to dispense with the pilot altogether and have all passenger aircraft completely controlled by computer. However, the technology is not sufficiently advanced for such systems to be adequately adaptable, for example to changing weather conditions. This proposal is not a near term solution.

However, it is not as far fetched to switch to fully automatic control only in extreme circumstances. For example, the switch could be triggered from the ground or by a panic button in the cockpit. Note, however, that once an aircraft has switched to fully automatic control, it is in a critical state. It may not be able to evade other aircraft or bad weather, so the air traffic control system will have to clear away traffic in its path. To prevent this happening by accident, if there is a panic button in the cockpit, then it has to be hard enough to push to make it extremely unlikely that it gets done by accident. It can't be a big read button in the middle of the console with a label "Panic Button." It would need some sort of elaborate interlock or some authentication of the pilot. This is at odds with enabling the pilot to throw the switch as his throat is being cut.

Moreover, the scheme is considerably more complex than Soft Walls; it requires trajectory planning and automatic landing, neither of which Soft Walls requires. It is doable, of course (automatic landing technology is already deployed on many aircraft), but extra complexity will, at a minimum, imply longer development times.

15. Can pilots tolerate a reduction of navigable airspace?

Among the more extreme ideas circulating include restricting aircraft to narrowly defined air lanes, making, in effect, tunnels in the sky. This greatly reduces flexibility in the system, making it much more difficult to adapt to unusual weather or traffic conditions, for example. If Soft Walls is deployed, the regulatory bodies that define the no-fly zones will have to exercise restraint to not unnecessarily reduce the navigable airspace. Ideally, Soft Walls does not reduce legally navigable airspace at all, since regulatory bodies already restrict the airspace around inhabited areas. As such, Soft Walls only reduces navigable airspace by removing the space where flying is unacceptable anyway.

But there is a significant difference between regulatory no-fly zones (what we have now) and regions into which an aircraft will not fly (what Soft Walls will impose). Some pilots argue that there are emergencies on an aircraft that would justify flying through regions of airspace where flight is forbidden. However, the pilot who does this is choosing to override the regulatory bodies, putting people on the ground at risk in an effort to protect the people in the craft. Should the pilot have a right to make that decision? Soft Walls means that the decision is made by the regulatory bodies. There is no aircraft emergency grave enough to justify an

attempt to land on Fifth Avenue, and no pilot should have the right to choose to take that risk. Soft Walls can enforce that policy.

Of course, it is not new that there are regions into which aircraft will not fly. No aircraft, for example, can fly through a mountain, no matter how grave the on-board emergency that makes the pilot want to be on the other side of the mountain. Soft Walls creates no-fly zones where enforcement is gentler than that defined by mountains, but the constraint is equally strong. The aircraft simply cannot fly there.

16. Would Soft Walls prohibit engine cutoff in an emergency?

An objection frequently cited by pilots is that a Soft Walls system would have to regulate engine throttle along with other controls on the aircraft. Otherwise, a malicious pilot could fly over a no-fly zone and cut the engines. Engine throttle is particularly problematic because pumping fuel into a malfunctioning engine could prevent the pilot from recovering from, for example, an engine fire. This is a valid objection, and it creates an engineering challenge.

Aircraft engines are already equipped with sensors that detect a wide range of malfunctions. This sensor data should be provided to the Soft Walls system to help it choose the recovery strategy. Of course, there may be circumstances in which there is no workable recovery strategy. In this case, the Soft Walls system will choose the strategy that is most likely to protect the no-fly zone, even if it puts the airplane and its passengers at risk. This course of action may be much more difficult for the pilot to choose, but may well be the right course of action.

17. Isn't GPS vulnerable to attacks?

The Soft Walls system relies on localization information. The aircraft computers have to reliably know where the aircraft is. Avionics systems today already include localization systems, which are required for navigation (and for more advanced safety systems, like ground proximity warning systems).

The principle source of localization information today is the global positioning system (GPS), which uses signals emitted by a suite of 24 satellites. A GPS receiver performs a simple triangulation calculation to determine the location of the receiver. However, most aircraft have at least two backup systems. First, an inertial navigation system (INS) measures acceleration to determine when the aircraft is turning, ascending, or descending, and continually calculates the new location based on its knowledge of the previous location. Second, a variety of radio beacons are also used to triangulate the aircraft location. Radio beacons are particularly common around airports, and automatic landing systems rely on them.

Most radio signals can be jammed. This means that a malicious party transmits a radio signal that swamps the one of interest, making it impossible to receive reliably. GPS signals are vulnerable to jamming. During the second Iraq war, Russian-made GPS jamming devices were sold to the Iraquis to use against smart munitions, many of which rely on GPS.

Some radio signals can also be spoofed. This means that a malicious party transmits a radio signal that masquerades as the radio signal of interest, hoping that it will be picked up instead of the legitimate signal. Spoofing can be prevented by encryption techniques if the encryption key can be kept private. That is, it can be made extremely difficult (in today's technology, essentially impossible) to construct a legitimate signal without having knowledge of a key that can be very closely guarded.

GPS signals currently contain encrypted channels that make spoofing by synthesizing a signal extremely difficult. Radio beacons can be both spoofed and jammed, and hence probably cannot be relied upon in a hostile environment. INS systems cannot be either spoofed or jammed, since they do not use communications of any kind.

If a radio signal cannot be spoofed, then jamming can be reliably detected. Hence, if the GPS system is being jammed, then the Soft Walls system will know that it is being jammed, and instead of begin confused by random data, would switch to backup systems, primarily INS.

Without knowledge of the encryption key, GPS cannot be spoofed by constructing an artificial GPS signal. However, it may be technically feasible to pick up a GPS signal at one location and rebroadcast it to another location in such a fashion as to confuse a GPS receiver at the second location into thinking it is actually at the first. However, this technique would be difficult to use in a hijacking scenario. To go undetected, it would require that a second aircraft start at the same place and at the same time as the aircraft to be hijacked, and then slowly diverge so that over time it is at a different location. That second aircraft would have to rebroadcast what it receives from the GPS satellites at high enough power that the first aircraft picks up its signals rather than the ones coming directly from the satellites. Even if this highly unlikely scenario could be pulled off, the transponders of the two aircraft would report the same locations to air traffic control, which will certainly raise suspicion. Air traffic control would determine that the aircraft had collided, but were still flying.

A real vulnerability lies in the protection of the encryption key used to construct the GPS satellite signals. If this key leaks to a malicious party, then we are unlikely to hear about it until the key is misused in some disastrous way. This vulnerability is not limited to commercial aviation, but much of military operations would also be compromised. For this reason, it would be wise to adapt GPS so that the encryption key can be periodically changed.

18. Don't inertial navigation systems drift?

Inertial navigation systems (INS) serve as the backup for GPS for localization. However, inertial navigation systems drift over time. A localization system based on INS will have an accumulating error, where the size of the error depends on the amount of time that has elapsed since the system had a known good fix. That is, the error will depend on the amount of time that has elapsed since either the GPS system failed or the aircraft took off from a known airport.

Fortunately, INS has formed the cornerstone of aviation navigation for decades, long predating GPS, and the technology has gotten very good. Drift rates are small, but cannot by ignored in the design of the Soft Walls system. Fortunately, different airports have different requirements for precision. An airport in the middle of a city, such as San Diego or Reagan National, requires that approaching aircraft have precise localization information. An airport in the open, such as Washington Dulles or Reno, does not require as much precision. After GPS fails, precision will degrade over time. Thus, for a given airport, an aircraft would have only a certain amount of time (which depends on the airport) to land there after GPS fails. Back of the envelope calculations indicate that the time would be about half an hour for Reagan National airport, but several hours for Dulles.

Of course, air traffic control would have to be alerted when the GPS in an aircraft failed, and the pilot would have to be issued a revised flight plan. That flight plan would have to be such that if the pilot refused to follow it, there remained enough time to intercept the aircraft with a more forceful response.

19. To be effective, don't all aircraft have to be equipped with Soft Walls?

To be practical, it has to be possible to phase in the Soft Walls system. We cannot require overnight that all aircraft be equipped with it. One strategy would be to prioritize by geographic region. In the U.S., for example, Washington DC would likely be a high priority area.

Would it be practical for the U.S. to require even foreign carriers to be equipped with Soft Walls? In fact, it is more practical to do this than to require some sort of vetting of the pilots. In the October 2001 issue of *Forbes*, Peter Huber says [5]:

> "In late 1999 a demented Egyptian pilot deliberately flew his plane and its 217 passengers into the ocean soon after its takeoff from Kennedy Airport in New York. A Silkair jet that crashed into an Indonesian jungle in 1997, killing 104, was most probably a suicide, too. We can't put our well-vetted citizens in every cockpit, and no amount of advance screening of the pilots can or should reassure us about many of the foreign air carriers that fly planes to and from our shores. But between them, the U.S. and western Europe do control the software and hardware in every last one of those cockpits."

In fact, it can be made extremely difficult to tamper with the software and hardware in those cockpits. Consider the infinitely simpler case of cars with computerized controllers. How many hobbyists no longer work on their cars because it is no longer practical? Fly-by-wire aircraft surely raise the level of sophistication required to tamper with the machine.

20. What if aircraft without Soft Walls get in?

Soft Walls does not eliminate the need for other defenses like military aircraft and anti-aircraft batteries. It does, however, make it far clearer when those should be used. If regulation requires, say, that all aircraft approaching within 200 miles of Washington DC be equipped with Soft Walls, then an aircraft that is not so equipped must be met forcefully. But there is time to do so in a reasoned and careful way.

Correspondingly, an aircraft that is so equipped but deviates from its approach path while approaching Reagan National Airport need not be shot down unless it somehow manages to penetrate the no-fly zones. Since this should only occur if the unlikely event of a failure of the Soft Walls system, the anti-aircraft batteries can safely remain silent.

What if the Soft Walls system itself fails, so an aircraft that would normally be suitably equipped is not? Failure of the Soft Walls system will have to be detected, and this aircraft will have to be diverted (by air traffic control) to an airport that does not require Soft Walls. If the pilot does not cooperate with the diversion, then once again, force will be necessary.

21. How can air traffic control determine whether an aircraft has Soft Walls?

A key question is how can the air traffic control system reliably detect whether an aircraft is equipped with Soft Walls? Could an aircraft that is, say, maintained by a malicious organization, masquerade as one equipped with Soft Walls when in fact it is not?

In practice, it can be made very difficult to tamper with the software in a fly-by-wire aircraft. One simple strategy is build into the hardware a validation circuit that refuses to run the software if it is not appropriately digitally signed. Then we can be sure that a fly-by-wire aircraft that can fly will be flying using software that has been certified. But there is still the problem of how the air traffic control system will know that a particular aircraft is a properly

functioning fly-by-wire aircraft. Couldn't some other aircraft spoof the system by transmitting false transponder data?

Today, this would be easy (although it would require far more technical sophistication than the September 11 perpetrators had). However, it would not take much to make it much more difficult to falsify the identity of an aircraft. At a minimum, it could be made prohibitively expensive, requiring for example cannibalizing the entire electronics system of one aircraft to carry it on a second aircraft. Even this could be defeated by, for example, running self tests that verify that certain control actions alter localization information in predicted ways. But it seems doubtful that such extreme paranoia is justified. (Of course, before September 11, 2001, this whole discussion would have seemed extremely paranoid).

In any circumstance, the Soft Walls system will have to perform self-tests to determine that it is correctly operational, that its software and its database of no-fly zones have not been tampered with, and that it is able to bias the control of the aircraft. And it would need a mechanism for certifying to authorities that such tests have been passed.

22. Can the database be hacked?

Soft Walls assumes that the aircraft carries on board a database that defines the no-fly zones. The idea is that this database has fairly coarse grain information, not including individual buildings, for example, but including urban areas and nuclear power plants. Consequently, updates need not be frequent, and could coincide with periodic recertification of the aircraft. But even infrequent updates can create a vulnerability. Would it be possible for a malicious party to tamper with the data in the database?

First, it is important to realize that there is no reason to keep the data in the database private. Pilots can (and should) know where all the no-fly zones are. But the database must be tamper-proof. The standard method for accomplishing this is a digital signature, which functions like encryption, but backwards. Instead of hiding the data, it certifies the data. This technique requires that the authority that creates the database sign it using a private encryption key. A hacker would have to know the encryption key to create a variant of the database. The decryption key would be public, and would be used on board the aircraft to extract the database information.

23. Could maintenance crews install systems with Soft Walls disabled?

In fly-by-wire aircraft, the Soft Walls system would be integrated with the basic flight control software, and hence would be extremely difficult to disable or replace. In older aircraft, where for example Soft Walls may be part of the autopilot system, it may much more difficult to ensure that maintenance crews do not replace it with another system. In this case, we would have to rely on air traffic control being able to detect an aircraft without Soft Walls (see above).

24. How should the no-fly zones be defined?

Aviation authorities already restrict navigable airspace. A pilot that violates these regulations can lose his or her license. These regulations seem like a reasonable starting point for defining no-fly zones, but very likely are more conservative than what is really necessary. In defining no-fly zones, regulatory agencies will have to exercise restraint so as to not unduly restrict the options that a pilot has to respond to emergencies. The surest way to ensure that this happens is to involve the pilots in the regulatory process.

25. How long would it take to deploy Soft Walls?

In fly-by-wire aircraft, Soft Walls is "just" a software change. However, software changes in such aircraft are extremely expensive and time consuming because they (appropriately) force revalidation. Consequently, it could be expensive and take some time to deploy.

26. What about general aviation and cargo aircraft?

In the very long term, even the smallest aircraft are likely to be fly-by-wire and can implement Soft Walls at minimal additional cost. But this is very long term. In the near term, it seems unlikely to be cost effective to deploy this technique on all general aviation aircraft. For this reason, restrictions on the navigable airspace of such aircraft that were instituted after September 11 (some airports were even closed) are likely to be permanent.

27. Has the concept been applied to other problems?

The Soft Walls concept is related to the idea of *virtual fixtures*, where soft or hard constraints are applied in surgical assistance systems (see for example [1] and [8]). The objective of virtual fixtures is to place no-cut zones around delicate structures to prevent the surgeon from accidentally contacting them during robot-assisted minimally invasive surgery and microsurgery.

28. Why is Soft Walls the best option available for pilots?

Clearly, restricted control is better than being shot down. If Soft Walls does no more than reduce the likelihood of an accidental shooting, then we have accomplished a lot. But there are several other competing approaches that have gotten a lot of traction and are far worse from the pilot's perspective. Forced automatic landing systems, control from the ground, and fully automated flight clearly require the pilot to cede more authority than Soft Walls does.

The principle in Soft Walls is maximally generous to the concept of pilot authority. The pilot has as much control over the aircraft as is possible, subject to the constraint that the aircraft does not enter the no-fly zones.

29. What if pilots refuse to accept the Soft Walls system?

A pilot comes from a 2000-year-old tradition of the ship's captain, where even the authority to marry the passengers is granted. The captain is responsible for the ship, its crew, and its passengers, and tradition dictates absolute control over all elements of the craft. But since September 11, the safety of the people on the ground trumps pilot authority.

Despite all efforts to maximize pilot authority, it is still likely that some pilots will refuse to accept this system without an override in the cockpit. An implementation or phase-in plan will have to take into account the burden on the airlines to replace the pilots that refuse to accept the change.

30. Why is the system called Soft Walls?

There are two reasons. First, the no-fly zones are surrounded by a soft boundary, in the sense that an aircraft can penetrate the boundary, partly. As the penetration gets deeper, the bias of

the Soft Walls system increases until no further penetration is possible and the aircraft is diverted. Second, the Soft Walls system would be implemented in software.

Glossary

In the context of this article, a number of technical and non-technical terms are used with the specific meanings given below.

Actuator: A device that accepts commands from an electronic device (such as a control computer) and translates those commands into physical action (such as raising the temperature of something by turning on a heater, or the moving the flaps on a wing, to bank an aircraft).

Advisory systems: Avionics systems, such as ground proximity warning systems and collision avoidance systems that advise the pilot of dangerous conditions and recommend evasive action.

Authority: The term used in the aviation community for the ability that a pilot has to control all aspects of his or her craft.

Autopilot: A system available on most aircraft that will, at a minimum, keep an aircraft on a specified heading at a specified altitude. More sophisticated versions can steer the aircraft through a series of waypoints.

Avionics: Aviation electronics.

Beacons (see localization): A set of radio transmitters that emit signals that can be used on board an aircraft to triangulate to determine the precise location of the aircraft relative to the beacons.

Bias: An offset in a value. For example, if the pilot executes a command to turn at 5 degrees per second but the aircraft turns at 3 degrees per second, then there is a bias of -2 degrees per second.

Collision avoidance system: An avionics system designed to prevent collisions between aircraft. Prime examples today are TCAS and ACAS.

Computer control: A control system where actuation of physical systems is mediated by a computer.

Control surfaces: The parts of an airplane that affect the pitch, roll, and yaw of an aircraft. These include the rudder, the ailerons, and the flaps on the wings.

Digital signature: A number that can be used to verify that a particular body of data has not been tampered with since it was produced by an authenticated source.

Embedded software: Software that engages the physical world through sensors and actuators.

Flight control systems: The avionics systems that control what the aircraft does in response to pilot commands.

Flight envelope protection: The principle of interpreting commands from a pilot in such a way that the aircraft always remains within pre-specified performance parameters.

Fly-by-wire: The use of electronics (and, in particular, computers) to convert pilot commands into control of the aircraft flight surfaces, rather than mechanical or hydraulic linkages from the cockpit to the flight surfaces.

Force feedback: Force applied by a control system to the physical controls used by the pilot.

Global positioning system (GPS) (see localization): A satellite based system for localization.

Ground proximity warning system (groundprox): A system that advises the pilot when the aircraft is dangerously approaching terrain.

Hacking: The art of gaining unauthorized access and modifying commands and/or software.

Inertial navigation system (INS) (see localization): A localization system based on measuring acceleration.

Jamming: Creating interference for a radio signal by emitting a similar radio signal. For example, GPS signals can be jammed by emitting a radio signal in their frequency spectrum.

Localization: Technology for determining where an aircraft is.

Navigation: The act of planning where the aircraft will be.

Public key encryption: An encryption technique where the key used to encrypt the data is public, but the key used to decrypt the data is private. To get you to send me encrypted data, I provide you with a public key. There is no harm in someone else seeing that key, as they are free to also send me encrypted data if they wish. However, only I can decrypt that data because I hold the private key.

Sensor: A device the measures a physical phenomenon (such as acceleration, temperature, or humidity) and provides that information to an electronic device (such as control computer).

Spoofing: Fooling a system by pretending to be a legitimate source of information, but supplying invalid information. For instance, radio beacons can be spoofed by emitting an identical radio signal from a different location.

Transponder: A device carried by aircraft that alerts the air traffic control system (and other aircraft) of their identity, location, and velocity.

Unoccupied air vehicle (UAV): An aircraft that flies without a pilot on board.

Waypoints: Positions in space that form the destination of flight segments. For instance, to get from one city to another, an airplane will proceed through a series of waypoints according to a flight plan that is filed with the air traffic control authorities.

References

[1] A. Bettini, S. Lang, A. Okamura and G. Hager, "Vision Assisted Control for Manipulation Using Virtual Fixtures," IEEE/RSJ International Conference on Intelligent Robots and Systems, 2001, pp. 1171-1176 (http://www.me.jhu.edu/~allisono/publications.html).

[2] D.I.K. Brouwer, F.H. Grootendorst, M. Mulder and R. van Paassen, “Evaluating the Safety Augmentation System,” ,” *to appear* in AIAA Conference on Guidance, Navigation and Control, Providence (RI), 2004.

[3] J. Adam Cataldo, Edward A. Lee, and Xiaojun Liu, "Preliminary Version of a Two-Dimensional Technical Specification for Soft Walls," Technical Memorandum UCB/ERL M02/9, University of California, Berkeley, CA 94720, April 17, 2002.

[4] F.H. Grootendorst, D.I.K. Brouwer, M. Mulder and R. van Paassen, "Design of the Safety Augmentation System," *to appear* in AIAA Conference on Guidance, Navigation and Control, Providence (RI), 2004.

[5] Peter Huber, "Disable the Pilots," *Forbes Magazine*, October 15, 2001.

[6] Edward A. Lee, "Soft Walls - Modifying Flight Control Systems to Limit the Flight Space of Commercial Aircraft," Revised from Technical Memorandum UCB /ERL M001/31, University of California, Berkeley, CA 94720, October 3, 2001.

[7] Nigel Moll and Charles Gilson, "Midair collision reveals gap in international ACAS procedures," *Aviation International News*, Farnborough, July 22-28, 2002.

[8] M. Li and A. M. Okamura, "Recognition of Operator Motions for Real-Time Assistance Using Virtual Fixtures," in *11th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE Virtual Reality, 2003, pp. 125-131 (http://www.me.jhu.edu/~allisono/publications.html).

[9] Matthew L. Wald, "Can Computers Foil Air Pirates?" *New York Times*, Print Media Edition: Late Edition (East Coast), New York, N.Y. (ISSN: 03624331) April 11, 2002.