Chapter 4

The Crash of Korean Air Lines Flight 007

For want of the nail the horse-shoe was lost;
For want of the shoe the horse was lost;
For want of the horse the rider was lost;
For want of the rider the battle was lost;
For want of the battle the kingdom was lost;
And all for the want of a horse-shoe nail.

— George Herbert (1593-1633), "Horse Shoe Nail"

We are now ready to make a leap. Quite a big one, from watches to high-flying aircraft, showing that the same underlying structures that render the watch confusing also exist in the world of complex safety-critical systems. In fact, this is one of the central messages of this book. But please don't be intimidated by the perceived complexity of the aircraft systems, autopilots, and navigation systems—you will see that, basically, they are not different from fans, climate control systems, and digital watches. We will be using the same descriptive language with which we are already familiar to look at and understand the inherent problems that cause confusion and error.

Flight 007

One of the most tragic and perplexing civil aviation disasters in the twentieth century was the loss of a Korean jetliner in 1983. Korean Air Lines Flight 007 (KAL 007), a Boeing 747 jumbo jet flying from Anchorage, Alaska, to Seoul, South Korea, deviated more than 200 miles into Soviet territory and was subsequently shot down. There were no survivors; all 240 passengers and 29 crewmembers perished in the crash.

When we hear about such a horrific disaster, all of us ask, Why? Why would a well-equipped aircraft piloted by an experienced crew veer more than 200 miles off course? And why was a harmless passenger aircraft shot from the sky? This set of questions was, for many years, a puzzling mystery that tormented the victims' kin, baffled the aviation industry, haunted many, and gave rise to more than seven books and hundreds of articles, ranging from technical hypotheses and covert spy missions, to a handful of conspiracy theory offerings. But as you will come to see, the truth, as always, is cold, merciless, and needs no embellishment.

Alaska

The date: August 31, 1983. Time: 4:00 A.M. Location: Anchorage Airport, Alaska. Flight 007, which originated in New York City, was now ready and fueled up for the long transpacific flight to Seoul. After a long takeoff roll, the heavy aircraft, full of cargo, fuel, and passengers, pitched up and climbed slowly into the gloomy morning sky. After reaching 1,000 feet, the white aircraft rolled gently to the left and began its westbound flight. Leaving Anchorage behind, Korean Air Lines Flight 007 was given the following air-traffic-control instruction: fly directly to the *Bethel* navigational waypoint and then follow transoceanic track R-20 all the way to Seoul (see figure 4.1).

The aircraft followed the instructions and changed its heading accordingly. But as minutes passed, the aircraft slowly deviated to the right (north) of its intended route, flying somewhat parallel to it, but not actually on it. It first passed five miles north of *Carin Mountain*, a navigation point on the way to *Bethel*. Then, instead of flying straight over *Bethel*, a small fishing hamlet in Western Alaska, it passed 12 miles north of it. With every consecutive mile, the rift between the actual location of the aircraft and the intended route increased. Two hours later, when the aircraft reached *Nabie*—an oceanic waypoint about 200 miles west of the Alaska coast—it was already about 100 miles off track (see figure 4.1). In the early morning darkness, the weary passengers on this westbound flight were fast asleep. In the cockpit, the flight crew reported to air traffic control that they were on track, flying toward *Nukks, Neeva, Ninno, Nippi, Nytim, Nokka, Noho*—a sequence of navigation waypoints on route to Seoul.

Everything looked normal.

But it wasn't. As the flight progressed and the divergence between the actual aircraft path and the intended flight route increased, the lone jumbo jet was no longer flying to Seoul. Instead, it was heading toward Siberia. An hour later, Flight 007, still over international waters, entered into an airspace that was closely monitored by the Soviets. In the very same area, a United States Air

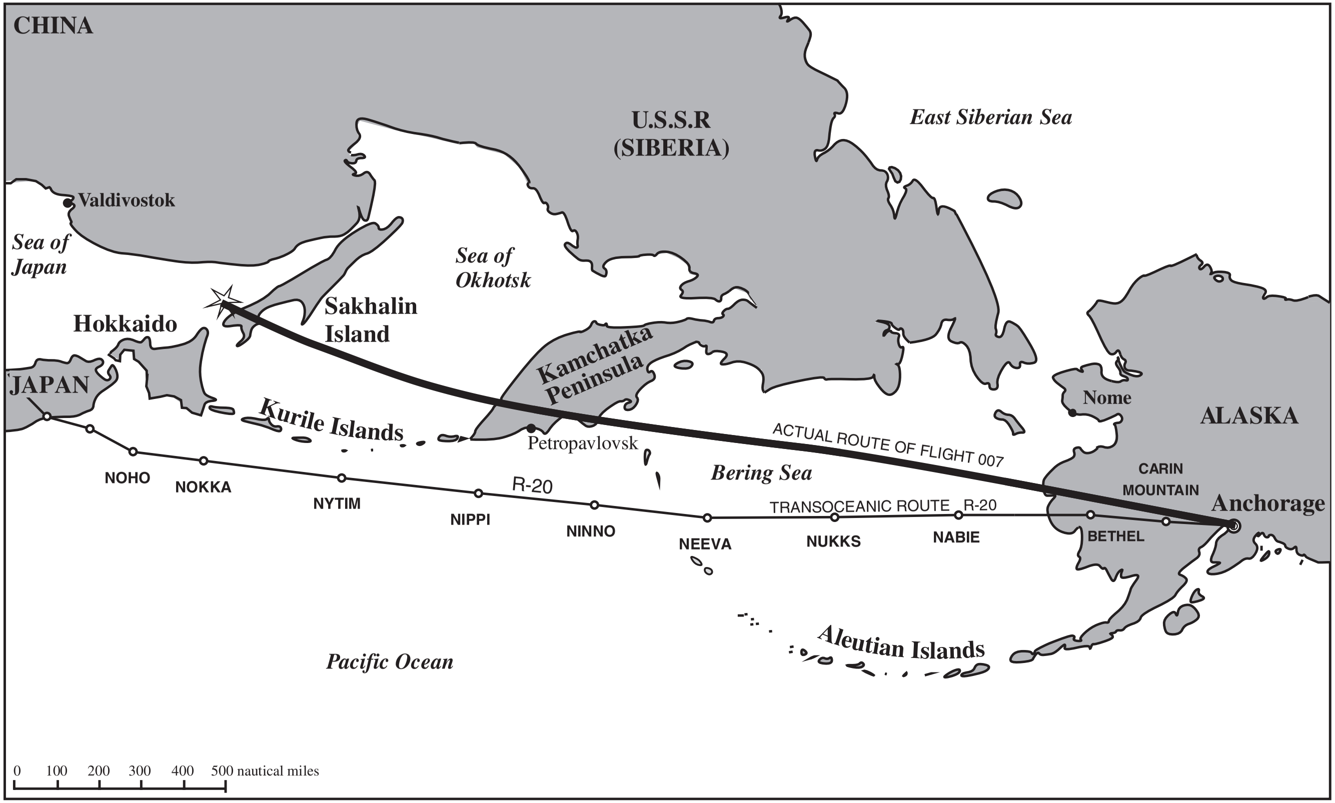

Figure 4.1. The flight path of Korean Air Lines Flight 007.

Force Boeing RC-135, the military version of the commercial Boeing 707 aircraft, was flying a military reconnaissance mission, code named "Cobra Dane." Crammed with sophisticated listening devices, its mission was to tickle and probe the Soviet air defense system, monitoring their responses and communications. The U.S. Air Force aircraft was flying in wide circles, 350 miles east of the Kamchatka Peninsula.

Kamchatka

Kamchatka is a narrow and mountainous peninsula that extends from Siberia to the Bering Sea. It has many active volcanoes and pristine blue, yet acid-filled, lakes. In those mountains, the Soviets installed several military radars and command centers. Their purpose was to track U.S. and international flight activities over the Bering Sea. At the height of the Cold War, under the leadership of Yuri Andropov, the General Secretary of the Communist Party, the Soviets were obsessed with keeping their borders secured and tightly sealed. As the U.S. Air Force Boeing 707 aircraft was circling over the frigid water in the dark of the night, purposely coming in and out of radar range, Soviet radar operators were monitoring and marking its moves. And then, during one of the temporary disappearances of the reconnaissance aircraft from the radar screen, the Korean airliner came in. The geographical proximity between the two large aircraft led the Soviet air-defense personnel sitting in front of their radar screens to assume that the target that reappeared was the military reconnaissance aircraft. They designated it as an *unidentified* target.

The Korean airliner continued its steady flight toward Kamchatka Peninsula. In the port town of Petropavlovsk, on the southern edge of the peninsula, the Soviets had a large naval base with nuclear submarines. KAL 007 was heading straight toward it. But the pilots could not see Kamchatka, because although the night sky above them was clear, everything below them was pitch dark. When the jetliner was about 80 miles from the Kamchatka coast, four MiG-23 fighters were scrambled to intercept it. The fighter formation first flew east for the intercept, then turned west and started a dog chase to reach the fast and high-flying Boeing 747. Shortly after, low on fuel, the fighters were instructed to return to base. The Korean jetliner, now 185 miles off track, crossed over the Kamchatka Peninsula and continued into the Sea of Okhotsk. Over international waters, safe for the moment, the large aircraft was heading, unfortunately, toward another Soviet territory—Sakhalin Island—a narrow, 500-mile-long island off the Siberian coast, just north of Japan.

Sakhalin

On the radar screen inside a military control center on Sakhalin Island, the approaching blip was now designated as a *military* target, most likely an

American RC-135 on an intrusion mission. Because of this military designation, the rules for identification and engagement were those reserved for military action against Soviet territory (and not the international rules for civil aircraft straying into sovereign airspace). Apparently, there had been more than a few such airborne intelligence-gathering intrusions by U.S. aircraft into Soviet airspace in the preceding months, not to the liking of Soviet air-defense commanders.

As the target approached Sakhalin Island from the northeast, two Soviet Su-15 fighters, on night alert, were scrambled from a local military airbase toward the aircraft. The Soviet air defense system had fumbled in its first encounter with the intruding aircraft, but was unlikely to miss the second chance. A direct order was given to local air-defense commanders that the aircraft was a *combat* target. It must be destroyed if it violated state borders.

About 20 minutes later, Flight 007 crossed into Sakhalin Island. The flight crew—sitting in their womb-like cockpit, warm and well fed—had no idea they were flying into the hornet's nest. At 33,000 feet, while the pilots were engaged in a casual conversation in the lit cockpit, monitoring the health of the four large engines, and exchanging greetings and casual chat with another Korean Air Lines aircraft also bound for Seoul—a gray fighter was screaming behind them to a dark intercept. The fighter pilot made visual contact, throttled back, and was now trailing about four miles behind the large passenger aircraft. The fighter pilot saw the aircraft's three white navigation lights and one flickering strobe, but because of the darkness was unable to identify what kind of an aircraft it was.

The identification of the target was a source of confusing messages between the fighter pilot, his ground controller, and the entire chain of air-defense command. When asked by his ground controller, the pilot responded that there were four (engine) contrails. This information matched the air-defense commanders' assumption that this was an American RC-135, also with four engines, on a deliberate intrusion mission into Soviet territory. The Soviets tried to hail the coming aircraft on a radio frequency that is reserved only for distress and emergency calls. But nobody was listening to that frequency in the Boeing 747 cockpit. Several air-defense commanders had concerns about the identification of the target, but the time pressure gave little room to think.

Completely unaware of their actual geographical location, the crew of KAL 007 were performing their regular duties and establishing routine radio contact with air traffic controllers in Japan. Since leaving Anchorage they were out of any civilian radar coverage. After making radio contact with Tokyo Control, they made a request to climb from 33,000 feet to a new altitude of 35,000 feet. Now they were waiting for Tokyo's reply.

Meanwhile, a Soviet air-defense commander ordered the fighter pilot to flash his lights and fire a burst of 200 bullets to the side of the aircraft. This was intended as a warning sign, with the goal of forcing the aircraft to land at Sakhalin. The round of bullets did not include tracer bullets; and in the vast and empty darkness the bullets were not seen or heard by the crew of KAL 007. The four-engine aircraft continued straight ahead. Flying over the southern tip of Sakhalin Island, Soviet air defense controllers were engaged in stressful communications with their supervisors about what to do. The aircraft was about to coast out of Soviet territory back into the safety of international waters; the target was about to escape clean for the second time. The Sea of Japan lay ahead—and 300 miles beyond it, mainland Russia and the naval base of Vladivostok, the home of the Soviet Pacific fleet.

Sea of Japan

The air-defense commander asked the fighter pilot if the enemy target was descending in response to the burst of bullets; the pilot responded that the target was still flying level. By a rare and fateful coincidence, just as the aircraft was about to cross into the sea, KAL 007 received instructions from Tokyo air traffic control to "Climb and maintain 35,000 feet." As the airliner began to climb, its airspeed dropped somewhat, and this caused the pursuing fighter to slightly overpass. Shortly afterward, the large aircraft was climbing on its way to the newly assigned altitude. The fighter pilot reported that he was falling behind the ascending target and losing his attack position. This otherwise routine maneuver sealed the fate of KAL 007. It convinced the Soviets that the intruding aircraft was engaging in evasive maneuvers.

"Engage afterburners and destroy the target!"

The fighter pilot moved in to attack. He climbed back to behind the target and confirmed that his missiles were locked on and ready to fire. He could see the target on his green radar screen and also visually; the steady white light from the tail of the aircraft was shining through and the rotating strobe flickered against the cruel arctic darkness. With the morning's first light behind him and the lone aircraft ahead of him, he was in a perfect attack position against his foe. The only problem, of course, was that the high-flying aircraft, with 269 innocent souls onboard, was neither a foe nor a military target. But the fighter pilot, as well as the entire chain of air-defense command, did not know that, and if any officer, including the fighter pilot, had any doubts about the target's designation, they certainly did not voice them.

Seconds later, the fighter aircraft launched two air-to-air missiles toward the target. One missile exploded near the vulnerable jet. The aircraft first pitched up. The blast burst a hole in the aircraft skin and caused a loss of pressure inside the large cabin. The public address speakers gave repeated orders to don the yellow oxygen masks, as the aircraft started to fall while rolling to the left. The captain and his copilot tried helplessly to control the aircraft and arrest its downward spiral. Two minutes later, the aircraft stalled out of control and then plummeted down into the sea. It impacted the water about 30 miles off the Sakhalin coast.

The Many Whys of This Accident

Following the downing of the aircraft, Russian military vessels searched for the aircraft wreckage. Two month later, deep-sea divers operating out of a commercial oil-drilling ship brought up the aircraft's "black boxes"—the cockpit's voice recorder and the digital flight data recorder. The Soviets analyzed the recorders to find clues about the accident, but kept their findings secret. Ten years later, after the Iron Curtain fell, the tapes were handed over to the International Civil Aviation Organization (the civil aviation arm of the United Nations Organization). The accident investigation was completed in 1993.

With that, let's look into the first "why" of this disaster. Why did the aircraft deviate from its intended flight route? To answer this question, we need to trace the crew's interaction with the aircraft's autopilot and navigation systems from the very beginning of the flight. Two minutes and ten seconds after takeoff, according to the flight data recorder, the pilots engaged the autopilot. When the pilot engaged the autopilot, the initial mode was HEADING. In this mode, the autopilot simply maintains the heading that is dialed in by the crew. The initial heading was 220 degrees (southwesterly). After receiving a directive to fly to *Bethel* waypoint, the pilot rotated the "heading select knob" and dialed in a new heading of 245 degrees (see figure 4.2). The autopilot maintained this heading until the aircraft was shot down hours later.

Using HEADING mode is not how one is supposed to fly a large and modern aircraft from one continent to another. In HEADING mode, the autopilot maintains a heading according to the *magnetic* compass, which, at these near arctic latitudes, can vary as much as 15 degrees from the actual (*true*) heading (not to mention that this variation changes and fluctuates as an aircraft proceeds west). Furthermore, simply maintaining a constant heading does not take into account the effect of the strong high-altitude winds so prevalent in these regions, which can veer the aircraft away from its intended route. The aircraft was also equipped with an Inertial Navigation System (INS). This highly accurate system encompasses three separate computers linked to a sophisticated gyro-stabilized platform. Every movement of the aircraft is sensed by the gyros and then each of the three computers separately calculates the corresponding change in latitude and longitude. The three computers

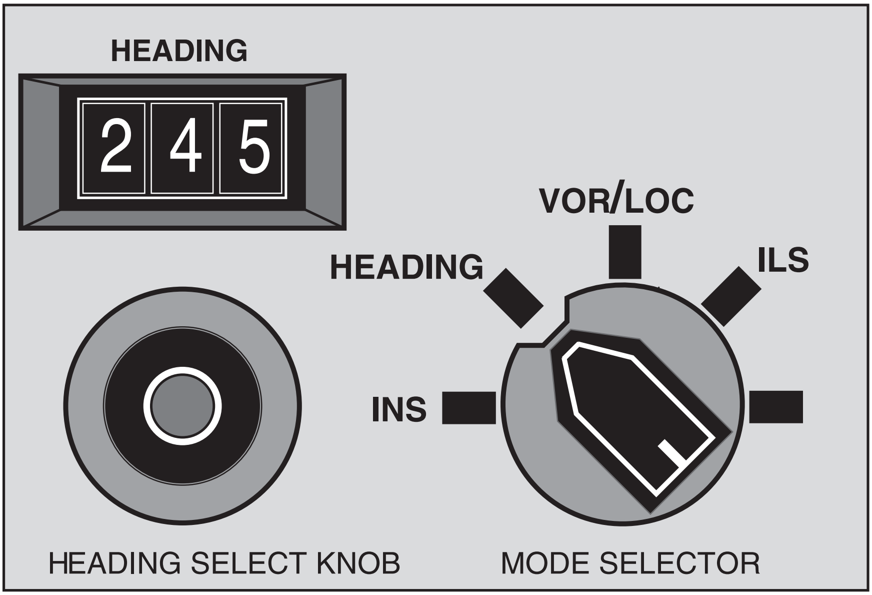

Figure 4.2 Heading indication and autopilot mode selector switch. The autopilot is in HEADING mode.

constantly crosscheck one another, making sure that their calculations agree. The system thus provides triple redundancy and is usually accurate within a mile of actual location. Throughout the flight, the pilots can see their latitude/longitude location on the digital readout of each computer, and know their exact location. In addition, the pilots can enter their entire flight route into the inertial navigation computers, and let them give steering commands to the autopilot, thereby creating a fully automatic navigation system. The pilots select inertial navigation (INS) mode by rotating the autopilot mode selector to the left (see figure 4.2).

On this flight, the route of flight was transoceanic track R-20 with some 10 waypoints (*Bethel, Carin Mountain, Nabie, Nukks, Neeva, Ninno, Nippi, Nytim, Nokka, Noho*), which passes very close to, yet shy of, Soviet-monitored airspace. To comfortably fly the aircraft along R-20 to Seoul, the pilots should have engaged the autopilot in INERTIAL NAVIGATION mode and let the system do the rest. That's what they were trained to do and that's what they have done in every transoceanic crossing with this aircraft. But they did not do it on this flight.

Why?

To answer this question it is first necessary to understand how the autopilot works and how pilots interact with it. The initial autopilot mode was HEADING. Then the pilot rotates the switch and selects INERTIAL NAVIGATION mode, and the autopilot should guide the aircraft along the route of flight (e.g., R-20 to Seoul).

Sounds simple enough. But actually, it's a bit more complicated, and this was part of the problem. It turns out that several things have to take place before the autopilot *really* goes to INERTIAL NAVIGATION mode: one is that the aircraft must be close to the flight route. Specifically, the distance between the aircraft's position and the flight route must be within 7.5 miles for the INERTIAL NAVIGATION mode to become active. That's one condition; necessary, but not sufficient. The other condition is that the aircraft must be flying *toward* the direction of the flight route (and not, for example, 180 degrees away from it). Only when these two conditions are met, or become True, will the autopilot engage the INERTIAL NAVIGATION mode.

Guards

There is a software routine that constantly checks for these two conditions. It is called a guard. A guard is a logical statement that must be evaluated as True before a transition takes place. Sound complicated? Not so. You and your friend go into a downtown bar, and there is a beefy guy there, leaning to the side and blocking the entrance with his wide body, checking young barflies' ages—no ID, no entry. Automatic teller machines, or ATMs, same story. You can't make an ATM transaction unless you punch in your personal identification number (PIN). All said and done, the PIN, just like the beefy guy, is a guard. Unless you supply an age-valid ID or the correct PIN, you are locked out.

Back to the autopilot: the first condition says that the lateral distance between the aircraft and the nearest point on the route must be within 7.5 miles. What is meant by this is that the distance should be "less than or equal to" 7.5 miles. Right? As long as the aircraft is 1, 2, 3, 4, 5, 6, 7, or 7.5 miles from the route, then we are allowed in. We write this as,

The distance between the aircraft and the route is "less than or equal to" 7.5 *miles.*

This logical statement, once evaluated, will produce either a True or False result. The second condition says

aircraft is flying toward the direction of the route.

You can see these two conditions written on the transition in figure 4.3.

The guard—the software routine that evaluates whether this transition will fire or not—cares not only about the first and second condition, but also about the relationship between them. In our case, the relationship is the logical

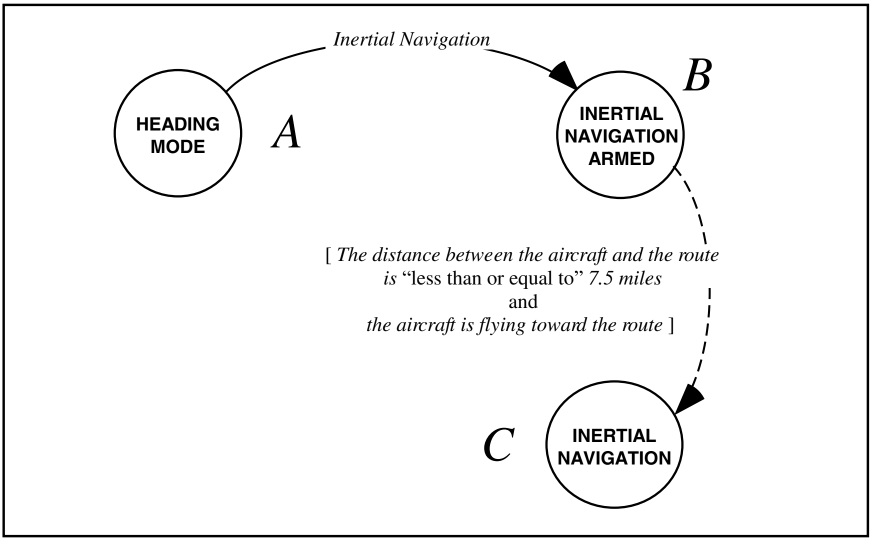
Inertial Navigation
HEADING MODE
A
INERTIAL NAVIGATION ARMED
B
[*The distance between the aircraft and the route is* "less than or equal to" *7.5 miles* and *the aircraft is flying toward the route*]
C
INERTIAL NAVIGATION

Figure 4.3. Machine model (of the autopilot's logic).

"and," which means that both conditions must be evaluated True for the transition to fire. In figure 4.3, the guard on the transition to the INERTIAL NAVIGATION constantly checks that *the distance between the aircraft and the route is* "less than or equal to" 7.5 *miles* and that the "*aircraft is flying toward the direction of the route*." Unless the two conditions are True at the same time, the guard will hold us back and no transition will take place.

So let's consider this scenario for engaging the INERTIAL NAVIGATION mode. Look again at the machine model in figure 4.3. Initially we are in HEADING mode. This is state *A*. We then decide to engage INERTIAL NAVIGATION. Very well—we reach forward and rotate the mode selector, and initiate event *Inertial Navigation*. This fires the transition to state *B*, ARMED, which is like a waiting room. We wait there until the two conditions are True before we can move on. If they are, we automatically transition to state *C*, which is what we wanted in the first place. Kind of awkward, isn't it?

This business of a waiting room, where we await patiently while the guard checks our entry credentials, is indeed awkward because if any one of the conditions is False, we have to wait until it turns True. But this begs an interesting question: What will the system do while we wait? According to this system's internal logic, the autopilot will stay in an awaiting state called INERTIAL NAVIGATION ARMED. What will happen from now on is that, each second, the guards evaluate the conditions. True, you're in, False, stay out. This will happen over and over (and stop only when the transition to INERTIAL NAVIGATION finally takes place).

All right, we are stuck in this ARMED state, but which autopilot mode is now active and flying the aircraft? This problem of what to do when conditions have not been met or when conditions are no longer valid is a common concern in the design of automated systems (and is discussed in detail in chapter 9). In our case, the designers of this autopilot chose to stay in a mode in which the autopilot keeps on maintaining the current heading of the aircraft. When in INERTIAL NAVIGATION ARMED, the autopilot is actually in HEADING mode. Note what is happening here: the mode is INERTIAL NAVIGATION ARMED, the display says "INS" (with an amber light), but the autopilot is in HEADING mode. Was all this a factor in the KAL accident? We'll soon see.

Automatic Transitions

By now we have seen two kinds of transitions that exist in automated systems. The majority were *user-initiated*. The user made the system move from one state to another in a very direct and deliberate way; the user pushed a button and the machine changed its mode. We have also seen *timed* transitions, such as the transition from LIGHT-ON to LIGHT-OFF in the digital watch. After three seconds in LIGHT-ON, the system automatically switched to LIGHT-OFF. The transition from ARMED to INERTIAL NAVIGATION mode is also *automatic*. But this one is not triggered by some predetermined (e.g., 3 second) timed interval; rather, it takes place when the guard becomes True. That is, it may stay in its current mode or switch depending on the two conditions (the distance between the aircraft's position and the route of flight). In a way, an automatic transition has a life of its own, switching and changing modes by itself. A guarded transition is the basic architecture of any automated system, and we shall therefore consider this structure very carefully in this book.

So who is involved in making an automatic transition take place? It's not only the pilot, is it? No, the transition happens without any direct user involvement; no buttons are pushed and no levers are moved. So who is it? For one, it is whoever wrote the software routine for the autopilot—this is where it all begins. And then there are the events themselves, some of which are only indirectly controlled by the pilot (such as flying toward the route) and some that are triggered by events that may be beyond the pilot's immediate control. More complex automated systems, with ever more sophisticated routines and capabilities, have authority to override the pilot's commands, and in some extreme cases the automated system will *not* give the pilot what he or she asked for.

And who bears the responsibility when something goes wrong? The legal responsibility of operating a system has resided traditionally with the user—the aircraft's captain in this case. But is that fair? The answer is not so clear-cut as it used to be in the many centuries of maritime tradition and one century of powered, human-controlled flight. And this indeed is a serious legal matter that is still pending in several courts of law, all involving cases of faulty user interaction with complex aircraft automation.

User Model

For now, let's leave this problem of responsibility and legality aside. What we need to remember is that there are two things that govern automatic transitions: the *guards*, which are conditions that are written by the software designer, and the *events* themselves. Therefore, to track and monitor what the autopilot is doing, the pilot must have a model in his or her head of specific modes, transitions, conditions, and events. And we recognize that this model may at times be incomplete and perhaps also inaccurate because of poor training and lousy manuals, or simply because, over time, the pilot might forget the minute details, especially the numbers. So let us assume, for the sake of this discussion, that the pilots of KAL 007 had an inaccurate user model about the behavior of this autopilot. Say they knew about going toward the flight route, and that the aircraft needed to be west of Anchorage, but they only knew that the aircraft must be *near* the flight route for the autopilot to transition to INERTIAL NAVIGATION. And let us quantify this by saying that "near" is a distance of up to 20 miles.

Our user model of figure 4.4 is quite similar to the machine model. There are three modes: HEADING, ARMED, INERTIAL-NAVIGATION. Transition from HEADING

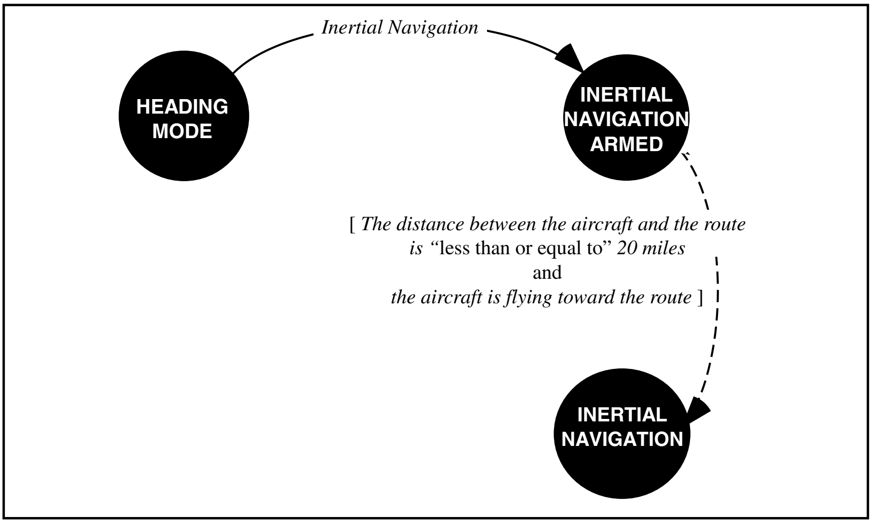
Inertial Navigation
HEADING MODE
INERTIAL NAVIGATION ARMED
[The distance between the aircraft and the route is "less than or equal to" 20 miles and the aircraft is flying toward the route]
INERTIAL NAVIGATION

Figure 4.4. User model.

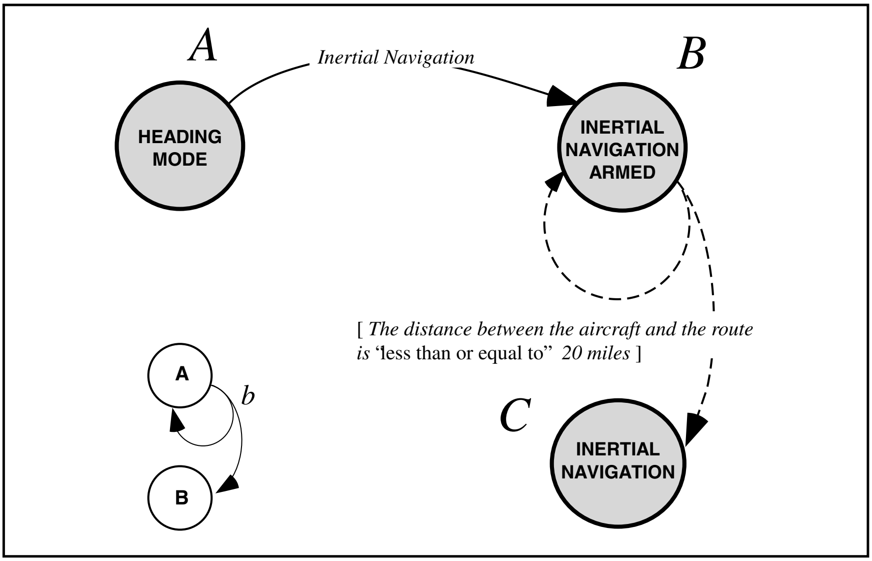
A
Inertial Navigation
B
HEADING MODE
INERTIAL NAVIGATION ARMED
[The distance between the aircraft and the route is 'less than or equal to" 20 miles]
A
b
B
C
INERTIAL NAVIGATION

Figure 4.5. The composite model (and a non-deterministic structure in the insert).

to INERTIAL NAVIGATION ARMED is accomplished by rotating the mode selector. But according to this model, the transition from ARMED to INERTIAL NAVIGATION occurs when the aircraft's distance is equal to or less than 20 miles away from the route and the aircraft is flying toward it.

Now we have a machine model (figure 4.3) and a user model (figure 4.4). Sounds familiar? Just as we did with the remote control of chapter 2 and the watch in chapter 3, we combine the two models to create a composite model (figure 4.5). Quick observation: this composite model looks different from either the machine model or the user model of the autopilot. Let's look at the differences one at a time. First, note that the *aircraft is flying toward the route* condition is omitted. I did this on purpose, because we want to focus our attention on the distance condition only (which is where the problem resides). With that out of the way, you can see in figure 4.5 that there are two transitions out of INERTIAL NAVIGATION ARMED (state *B*): one is going to INERTIAL NAVIGATION (state *C*) and the other one loops back to state *B*.

Why?

We know for a fact that if the aircraft's distance from the route is up to 7.5 miles, the transition to INERTIAL NAVIGATION will take place. The pilot, however, thinks that if the aircraft is up to 20 miles from the route, the transition to INERTIAL NAVIGATION will take place. As long as the aircraft is up to 7.5 miles from the route, the pilot model will not cause any problem. However, if the aircraft is more than 7.5 miles from the route, the pilot thinks that the autopilot will transition to INERTIAL NAVIGATION—but actually the autopilot will stay in ARMED.

From the pilot's point of view, the autopilot is very confusing; sometimes the autopilot will transition to INERTIAL NAVIGATION and sometimes not. Think about it this way: say that we do repeated flights and try to engage the inertial navigation system. On the first flight, the distance between the aircraft and the route is 5 miles. On the second flight it's 10 miles, and on the third flight it's 15 miles. Our pilot, who has the "up to 20 miles" model, will get completely confused. The pilot expects that the transition to INERTIAL NAVIGATION will take place in all three cases, when in fact it will only engage on the first flight (5 miles).

The source of this problem, as you already know, is non-determinism. In the composite model (figure 4.5), the same event generates two different outcomes. The structure of the composite model is the same as the non-deterministic structure of the vending machine. You can see the generic structure at the lower left corner of the figure. It's the same old story.

Back to the Air

Now that we better understand the essence of automatic transitions, we can return to the tragic saga of KAL 007 and finish the story, although from here on we are navigating in uncharted territory, because what really happened in that cockpit will never be known completely. From the data recovered from the black box, we know that the crew engaged the autopilot two minutes and nine seconds after takeoff. It was initially in HEADING mode. Afterward, two possible sequences of events could have occurred: one, that the crew completely forgot to select the inertial navigation system; two, that the crew did select inertial navigation, but it never activated.

We'll follow the more probable scenario, the second one. The crew of Korean Air Lines Flight 007 most probably selected INERTIAL NAVIGATION on the mode selector panel once the aircraft was *near* the route. This, after all, was the standard operating procedure on every transoceanic flight. They anticipated that the transition to INERTIAL NAVIGATION mode would take place and the autopilot would fly route R-20 all the way to Seoul.

Based on radar plots recorded while the aircraft was in the vicinity of Anchorage, we now know that the actual distance between the aircraft and the flight route, at the time the pilots could have switched to INERTIAL NAVIGATION mode, was already greater than 7.5 miles (see figure 4.6). If that was indeed the case, the first condition (*the aircraft's distance from the route is* "equal or less than" *7.5 miles*), was evaluated False. As the aircraft started to deviate from the intended flight route, the distance between the aircraft and the route only grew and grew. Every second the guards evaluated the condition, and every time it

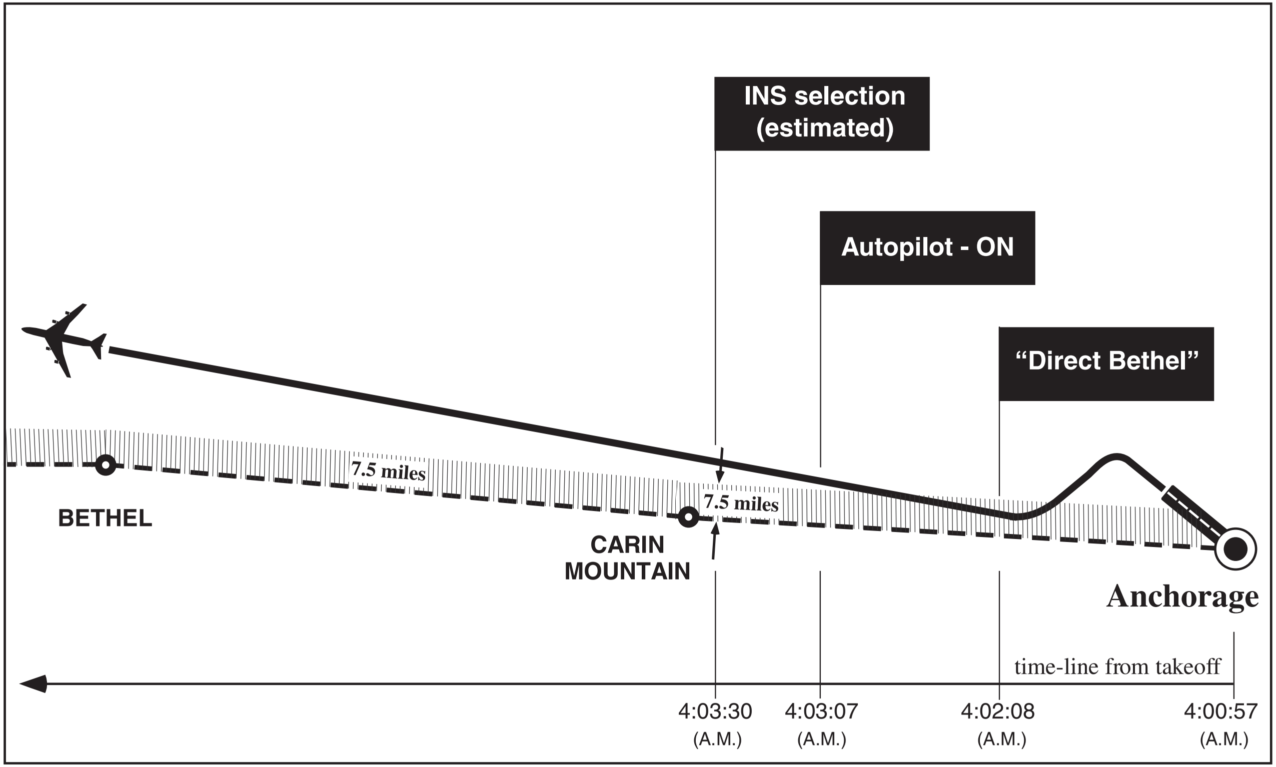

Figure 4-6. The flight path of KAL 007 shortly after takeoff.

came back False. The system was perpetually stuck in ARMED and the autopilot was on a straight, 245 magnetic heading, all the way to Siberia. The pilots probably had a very different mind-set. According to their model, the autopilot had indeed transitioned to INERTIAL NAVIGATION and was following transoceanic route R-20 to Korea.

Autopilot Displays

But why did the crew of Korean Air Lines Flight 007, sitting in the cockpit for five long hours, fail to detect that the autopilot was not following the intended flight route and that it wasn't in INERTIAL NAVIGATION mode? Fatigue, boredom, complacency? Many factors may have come into play, but it is a fact that on their way home, the crew was weary and tired after a long five-day trip. For centuries, captains and navigators have always feared homeward-bound carelessness on the final leg of a long and protracted journey. But lest we get lost in speculations, let's stay on our primary topic, pilot interaction with the autopilot. You see, the only hope of recovering from such a divergence between what the mind imagines and what the autopilot is actually doing is by getting some kind of feedback from the machine about its active mode—something that will alert the pilot that there is a problem. Were there indications about the active mode in the cockpit? Yes and No.

The inertial navigation system had a dedicated display to show the pilot the active mode. There is an indicator labeled INS. If the inertial navigation system is ARMED, the indicator lights up with an amber color; if the inertial navigation is engaged, the indicator is green. This is the "Yes" part of our answer. However, there was no indication anywhere in the cockpit that the autopilot was in HEADING mode. Why? It's hard to tell, but we are talking here about the early generation of automated flight guidance systems, when the importance of human factors was only remotely considered, if at all. Fact: when the autopilot was in INERTIAL NAVIGATION ARMED, there was no indication that the autopilot was actually maintaining HEADING. This is the "No" part of our answer.

Did this lack of mode indication play a role here? Most probably, yes. Did the pilots mistake INERTIAL NAVIGATION ARMED for INERTIAL NAVIGATION engaged? Perhaps; after all, it's the same indicator, and who knows what were the exact lighting conditions in the cockpit, and we all know how the mind sometimes sees what it wants to see. Did the pilots understand the subtle difference between these two modes? Did they know about guards and triggering events and how automatic transitions take place? Maybe, but we'll never know.

What we do know is that the lack of indication about the autopilot's active mode deprived the crew of an important cue. Such a cue might have drawn their attention to the fact that the INERTIAL NAVIGATION was *not* engaged and that the aircraft was actually flying on HEADING mode. Following the accident, all autopilots were modified to include this information. In today's autopilot systems, the display also shows the currently engaged mode as well as the armed mode; this is now a certification requirement. But this, as is always the case in accidents, is the blessing of hindsight. This design change came too late to help the crew and passengers of Flight 007.

Conclusion

The story behind Korean Air Lines Flight 007 was one of the greatest mysteries of the Cold War era (see endnotes). It is unnerving to consider that the sequence of actions and coincidences that finally led to the chilling shoot-down of a commercial aircraft in which 269 innocent people died, began with something that *didn't* happen—a tiny mode transition from HEADING to INERTIAL NAVIGATION. The autopilot was probably stuck in INERTIAL NAVIGATION ARMED mode, and that prompted the deviation over the initial waypoint (*Bethel*) that, for whatever reason, was not queried by Anchorage air traffic controllers. Then we have the similarity, at least initially, between the aircraft heading and the intended flight route. That heading, 245 degrees, took the aircraft toward the orbiting military RC-135 plane, which, coincidentally, was leaving the Soviet radar screen just as the 747 was coming in. Flight 007 was first labeled as *unidentified* and then, when it crossed into Soviet territory, it was designated as a *military* and then *combat* target because of the stringent rules of engagement that the Soviets employed at the height of the Cold War. The early morning darkness, which may have prevented any real visual identification of the white aircraft with the red Korean Air Lines logo on its large tail as an innocent civilian airliner, was another factor in this complex tragedy. Finally, what sealed the fate of KAL 007 was the climb from 33,000 feet to 35,000 feet, just after the fighter aircraft fired a warning burst in order to force the airplane to descend and land on Sakhalin. The climb was mistaken for an evasive maneuver and the aircraft was shot out of the sky and crashed into the sea. And all for the want of a little automatic transition.

interface for the digital watch was okay. But in a practical sense it was not—the traveler acted quite rationally but still was surprised by the clock. When it comes to user–machine interaction, the usability test is not just an issue of strict formalities but of practice. The proof of the pudding is in the eating.

I wore the digital watch described in this chapter for several years before I was finally able to figure out why it would *sometimes* fail me on wakeups. A somewhat similar problem in another digital watch is told by Donald Norman in *The Design of Everyday Things* (Basic Books, 2002):

> I had just completed a long run from my university to my home in what I was convinced was a record time. It was dark when I got home, so I could not read the time on my stopwatch. As I walked up and down the street in front of my home, cooling off, I got more and more anxious to see how fast I had run. I then remembered that my watch had a built-in light, operated by the upper right-hand button. Elated, I depressed the button to illuminate the reading, only to read a time of zero seconds. I had forgotten that in stopwatch mode, the same button (that in normal, time-reading mode would have turned on a light) cleared the time and reset the stopwatch. (page 110)

The aircraft accident mentioned in the beginning of the chapter took place on April 26, 1994, in Nagoya, Japan. After the copilot inadvertently engaged the GO-AROUND mode, which automatically increases engine thrust to maximum power and initiates an aggressive climb, the copilot pushed the control wheel down to force the aircraft to descend. The captain also instructed the copilot to push the control wheel down and make the landing.

As the copilot was pushing the wheel to force the aircraft to land, the autopilot—engaged in GO-AROUND mode—was countering by adjusting the control surfaces to make the aircraft climb. The crew continued with the landing, unaware of this abnormal situation (there was no warning to alert the crew directly to the onset of this abnormal condition). By the time the captain recognized the abnormal situation, and took control of the aircraft, he was unable to recover in time. The aircraft pitched high, stalled, and crashed tail-first on the runway. Of the 271 passengers and crew, only 7 passengers survived.

The captain and copilot assumed that if they would exert heavy nose-down forces to the control wheel, they would be able to override the autopilot (in GO-AROUND mode) and continue the descent to the runway. This was true for most autopilots, but on this particular model, it was different. The feature that allows the pilot to override the autopilot was disabled when the aircraft altitude was less than 1,500 feet. (The intent was to prevent the pilots from inadvertently pushing or pulling on the wheel and thereby overriding the autopilot). It also appears that the aircraft manual had unclear descriptions of this autopilot feature, which contributed to the pilots' assumption that the manual override always works. Subsequently, all autopilot models were modified such that this manual override feature is never disabled.

Chapter 4

The discussion in this chapter is primarily based on the report of the International Civil Aviation Organization (ICAO) investigative team and transcripts of recorded conversations within the Soviet Air Defense System. Additional information was obtained from the article "Closing the File on Flight 007" by Murray Sayle (*The New Yorker*, December 13, 1993, pages

90–101), a chapter by David Beaty titled "Boredom and Absence of Mind" in his book *The Naked Pilot: The Human Factors In Aircraft Accidents* (1995 Airlife Publishing: Shrewsbury, England), as well as videotapes of a press conference in Moscow several days after the disaster featuring the Red Army's chief of staff and a large and detailed diagram of the flight route, the first interception over Kamchatka, and the final attack on Flight 007.

Although the inertial navigation system with its triple redundant computers was hailed as the technological solution that would eliminate navigation errors, similar navigation snafus were no news to the commercial aviation community. There were more than a dozen similar incidents that were investigated (and Lord only knows how many never got reported, to save face), in which flight crews selected INERTIAL NAVIGATION mode but did not detect that the navigation system was not steering the autopilot. Other incidents occurred in situations in which flight crews failed to re-engage INERTIAL NAVIGATION after temporarily using HEADING mode to fly around thunderstorms. Most of the incidents involved track deviations of less than 60 miles. One incident, however, involved a 250-mile off-track deviation, and another (110 miles) almost resulted in a mid-air collision between an off-course Israeli El-Al Boeing 747 and a British Airways 747 over the Atlantic Ocean.

The presence of military reconnaissance aircraft, flying close to commercial air routes and collecting electronic and communications information, is not a new practice; the first intelligence-gathering reconnaissance missions of the Cold War era were flown during the airlift to Berlin in 1946. The practice is still going on today. On April 1, 2001, a U.S. Navy aircraft was flying a communication-gathering mission over the South China Sea. A Chinese fighter aircraft was sent to intercept and track the Navy EP-3 reconnaissance aircraft. The fighter aircraft came too close to the reconnaissance aircraft and collided with it. The Chinese fighter crashed in the sea and the pilot died; the U.S. Navy EP-3 aircraft made an emergency landing on the Island of Hainan. The crew was released after long and tense diplomatic negotiations between China and the United States. The aircraft was dismantled and then brought back to the United States.

Commercial aircraft straying into military zones, causing confusion, and at times strife, has occurred in the past. In 1978, MiG fighters strafed a Korean Air Lines plane that mistakenly wondered into Soviet airspace. The aircraft made an emergency landing on an icy lake in the Kola Peninsula (in northeast Russia, close to Finland), and fortunately only a few passengers were injured. In 1985, two years after the shooting down of Korean Air Lines 007, a Japan Airlines Boeing 747 also had its autopilot in HEADING mode and strayed 60 miles into the same (Siberian) airspace. By this time, the Soviets had made changes in their rules of engagement and were more cautious. The aircraft was escorted out.

The shooting down of Korean Air Lines Flight 007 was used by both the Soviets and the United States for political propaganda and internal politics. President Ronald Reagan called it a "heinous act" and used it to press Congress to continue with the arms race and the mystifying "Star Wars" weapons system. Yuri Andropov, in turn, denounced the event as "a sophisticated provocation masterminded by the US special services with the use of a South Korean plane" and drove the USSR into a counter arms race that eventually broke the Soviet Union's economic back, giving rise to the reforms of his successor, Mikhail Gorbachev. Cold War politics aside, the disaster shook every nation on both sides of the Iron Curtain and became an international tragedy. It eventually led Reagan to offer the world's civil aviation operators free use of the then military-only, Global Positioning System (GPS), a constellation of more than three dozen orbiting satellites that allows precision navigation for ships and aircraft. Many have argued that if KAL 007 had the highly accurate global positioning system, such a blunder could not have happened. But as we will see in chapter 8, which details the story behind the grounding of the cruise ship *Royal Majesty*, which was equipped with a global positioning unit, accuracy is not the only factor that must be considered in the design of highly automated navigation systems.

Chapter 5

The *Statecharts* description language, which embodies these three characteristics of automated systems (concurrency, hierarchy, and synchronization), was conceived in the early 80s to deal with the growing complexity of automated systems. It has since been used to describe a variety of systems such as avionics, automotive, medical technology, as well as for specifying and understanding computer security systems. *Statecharts* is similar to the finite state machine description that we have been using in the previous chapters, but extends the description to include concurrency, hierarchy, and synchronization. In this book we will stick with the original *Statecharts* language as described by David Harel and Michal Politi in their book *Modeling Reactive Systems with Statecharts* (McGraw Hill, 1998). *Statecharts* (and its many derivatives) are commonly used in the process of designing and specifying computer code, and is one of the components of the Unified Modeling Language (UML), a structured approach for software development (see *The Unified Modeling Language Reference Manual* by James Rumbaugh, Ivar Jacobson, and Grady Booch published by Addison Wesley, 1999).

The leap from *Statecharts* as a specification language for code generation into using it for user interfaces was taken by Ian Horrocks (*Constructing the User Interface with Statecharts*, Addison-Wesley, 1999). Horrocks, a professional designer working for British Telecom, shows a systematic methodology for specifying interfaces using the *Statecharts* language. His is a book about implementation—how to specify the behavior of the user interfaces and how to generate from this specification a well-structured code.

The notion of hierarchy casts some light on the topic of abstraction, which, as I have mentioned earlier, is the essence of any interface design. By going up in a hierarchy of states and super-states one can abstract information that resides within a super-state, and by going down one can refine. For a down-to-earth treatise on these issues, see a small book, titled *Science of the Artificial* (MIT Press, 1996), by the late Nobel prize winner, Herbert Simon.

Chapter 6

The technical term for the dual switch we discussed here is a *three-way* switch. The reason it is called a three-way, even though there are only two switches, is that for this arrangement to work, each light switch must have three terminals (see the figure on page 291); one "traveler" wire goes from switch A to switch B, another from switch B to A, and then there is a "common" wire that runs from the electrical source to the switches and from there to the light fixture. The light fixture is ON when there is a continuous pathway from the electrical source (via the common wire) to the light. When either switch creates a gap in that pathway, the light fixture is OFF.

On a more general level, the dual, triple, or multiple switch poses an interesting quandary: how do we know what should be the state of the light fixture, for any number of switches? What if we have four switches—how can we tell?

Let's work it out together; it is rather illuminating. We already know how a system with two light switches works, right? Here is a table for a system with two switches: