

Horizon Sensing Attitude Stabilisation: A VMC Autopilot

Presented to the 18th International UAV Systems Conference, Bristol, UK, 2003

B. Taylor C. Bil S. Watkins
School of Aerospace, Mechanical and Manufacturing Engineering
RMIT University
Melbourne, Australia.
btaylor@labyrinth.net.au

G. Egan
School of Electrical and Computer Systems Engineering
Monash University
Clayton
Melbourne, Australia

1. Abstract.

This paper covers a very simple, low cost absolute attitude stabilisation system based on thermal horizon detection. This stabilisation system is very fast and is ready for operation from a cold start within 20 milliseconds (mS). It consumes very low power and occupies a total code space of less than 6000 words. The system works reliably, day and night, in Visual Meteorological Conditions (VMC) using readily available thermopile sensors in the thermal infrared band between 7.5 and 13.5 microns.

One application is a simple photo reconnaissance unmanned aircraft that can be operated by one person. These are flown almost exclusively in VMC as the cameras can rarely function in worse conditions. A second application is to stabilize air vehicles used for research applications by non-pilot students. The operator can concentrate on their own project in machine vision, optic flow, telemetry, video compression, wireless LANs, whatever, without needing to become a competent pilot.

One interesting feature of the HSAS is its ability to hold a sensor plate level in both pitch and roll during turns, without the need for any gyros or accelerometers. This greatly simplifies geo-referencing a down looking photograph as the aircraft GPS coordinates are automatically at the centre of the image, regardless of bank angle.

A stabilized platform allows development of a much simplified user interface where non-pilots may safely control a UAV with only a few minutes familiarisation. The command system controls the aircraft by one joystick and the ground operator can select proportional right or left turns and whether to climb, hold height or descend. During takeoffs and landings the ground controller merely steers right or left and the climb out or descent is automatic.

2. Biography

Brian Taylor is an electronics engineer currently studying for a Masters Degree by Research in Aeronautical Engineering at RMIT University. He is also an Honorary Research Associate at Monash University. He has worked for several years with Aerosonde and has extensive experience in the electronics and telecommunications industries. He has over 1000 hours of sailplane pilot experience and flies RC models for fun.

3. Introduction and aim

This work grew from an analysis of the cost structure of what were meant to be low cost UAVs that turned out to need a fleet of highly qualified operators, expensive hardware, complex software and where even simple photographic missions were costing over \$1,000 US dollars per flying hour.

Until operating costs can be brought well below that of small General Aviation (GA) aircraft there is little hope that the rich field of UAV applications so often quoted as being 'just around the corner' will materialise. Regulatory issues, often cited as a major problem to widespread UAV operation, will not be a real problem until UAV operating costs are low enough to stimulate a market for UAV services.

The overall aim of this work is to build a minimal complexity system capable of achieving autonomous flights in excess of 24 hours (or lifting modest payloads for less time) with a complete hardware cost of the airframe, control system, radio, propulsion and digital camera under \$US 5,000. It must be able to be operated by non-pilots with very basic training.

Low hardware costs alone are not enough to reduce overall operation costs. Attrition costs, launch and recovery costs, operating crew numbers and their qualifications must all be driven down. Activities must be automated as far as possible so that take-offs, landings, flight plan input and control during the mission are cut to the bare minimum.

This paper describes one element of this work, the flight stabilizer, a number of which have been built and successfully flown many times.

4. Human control strategies

Sailplane pilots, in most countries of the world, are initially taught to fly without instruments and are routinely checked for this skill. They are taught to fly by visual reference to the horizon alone. Speed in a turn is maintained by constant bank angle and ensuring the nose does not rise or fall relative to the horizon. Straight and level flight is maintained by keeping each wing at the same position relative to the horizon. Although humans are conscious of substantial detail in their visual field and this might imply detailed pixel by pixel analysis is needed, very simple, wide field of view, lens-less sensors are able to fly a UAV very accurately.

Our eyes have evolved to exploit the sun's dominant radiation spectrum, in the range from roughly 350 to 750 nanometres (nm) peaking at 550 nm. The brain removes the sun and cloud brightness anomalies, works out where the horizon is and the hand applies the corrective controls.

Stancliff and Nechyba [1] describe a strong body of work done in transferring human control strategies to road vehicles but lament that little similar work has been done with aircraft.

5. Absolute attitude vs rate detection

Rate based systems rely on detecting and integrating movement, typically rotation, acceleration, air speed and changing altitude. By referring to a complex model of the vehicle dynamics the stabiliser can then derive corrective control movements. Rate based systems do not work on stationary vehicles, require a significant time to initialize from a cold start or an in-air reset and need accurate stability derivatives. Gyros are significantly more expensive than thermopiles and typically rather fragile. Rate based systems need significant computing power as described by F. Valentinis, C. Bil and P. Riseborough in their paper "Development and trials of an autonomous Flight Control System for UAVs" where five Motorola 68000 class processor modules and two I/O boards were deemed to constitute a relatively simple fuzzy logic control system. [2]

By detecting a known vector such as the earth's magnetic field, GPS carrier phase, or the infrared horizon, absolute attitude detection can be achieved. With the exception of GPS carrier phase detection, which is slow to acquire lock, absolute attitude systems can be very fast and can determine corrective control surface deflections even while the vehicle is stationary. Resolving the 3D magnetic field requires considerably more computing resources than finding the infrared horizon which is the basis of this system.

Human pilots in VMC primarily use absolute attitude information. A quick glance out of the canopy instantly shows pitch and roll status. Humans are not very good at sensing small rates of change in rotation, acceleration, speed or altitude. Up until very recently, practical, affordable and small absolute attitude sensors were not practical for small UAVs but new Micro Electro Mechanical Systems (MEMS) thermopiles are now readily obtainable and affordable.

6. Horizon sensing.

The sky above is typically brighter than the earth below and a simple visual spectrum, photoelectric, optical brightness based system could conceivably resolve the horizon position in most daytime circumstances. Attempts to exploit the visible spectrum with simple sensors have had significant troubles with direct sunshine, bright cumulous clouds, patchy snow cover and cannot work at night. Early

stabilisers based on this approach, such as the Royal Aircraft Establishment albedo horizon sensor of 1971 [5] and the later Ripmax HAL-2100 and the Futaba Pilot Assist PA-1 and PA-2 modules, are very easily fooled by the sun or bright clouds nearby. [4]

In the period up to the mid 1960's there were numerous NASA research programs in the infrared sensing field. One, which is not referenced here, looked at ways of pinpointing the re-entry of a spacecraft into the earth's upper atmosphere. It turns out that 3-5 microns is the ideal spectrum for this. This is both a quiet part of the spectrum and a peak in the blackbody radiation from the glowing underside of the vehicle. A second study looked at stabilising spacecraft, telescopes and antennae by reference to planetary horizons. This technique has been in continuous use since the TIROS-2 weather satellite (launched November 23, 1960) and remains a common stabilisation technique today. It is equally applicable to aircraft only tens of metres above the earth and to spacecraft a million kilometres above the earth.

Theoretical and practical papers covered the conceptual design and construction of infrared horizon sensors. Planetary emissions were also studied by sub orbital X-15 aircraft. Early IR horizon sensors used wide spectrum sensors and were unduly influenced by water vapour in the upper atmosphere. Water vapour varies with latitude and season giving rise to small errors in the apparent position of the horizon. The most accurate earth alignment system used filtered infrared peaking at 15 microns to detect the emissions from carbon dioxide which is very constant in the earth's atmosphere regardless of latitude and season. [6, 7, 8, 9, 10]. A recent US patent has been issued (US6,181,989) covering aircraft infrared stabilizers. The patent does not appear to acknowledge the prior NASA work [11].

Much more complex pixel or vision oriented systems can analyse a TV image and find the horizon with very high accuracy. Ettinger, Nechyba, Ifju and Waszak [3] describe a horizon sensing system aimed at a micro air vehicle but this requires a TV camera, extensive computing power and cannot work at night. Their demonstration system requires a video downlink since the computing requirements far exceeds onboard capabilities. A small fully autonomous aircraft using this technique is currently not possible.

7. Atmospheric windows.

Figure 1 shows several parts of the electromagnetic spectrum where radiation to or from space passes freely through the atmosphere. The radio spectrum, with wavelengths roughly 1 cm (30 GHz) to 1000 metres (300 KHz) is transparent but there is no

significant source of such radiation from space, except from the GPS satellites, nor any uniform radiation source from the earth. There are only geographically sparse point source emitters from radio stations. A major window appears around 10 microns in the mid or thermal infrared and another window is in the visible spectrum around 550 nanometres. Figure 2 shows the 10 micron window in more detail.




Figure 1. Useful atmospheric windows





Figure 2. Thermal or mid infrared window with effects of precipitable water vapour

For stabilising a spacecraft or an aircraft, mid infrared radiation in the 7.5 to 13.5 micron wavelength is ideal. This is the wavelength emitted by black bodies with temperatures commonly found on earth. NASA has used infrared horizon sensing since the 1950's to stabilise satellites and for antenna pointing.


Modern high sensitivity MEMS thermopiles are able to detect 0.05 degree differences. On clear days with a blue sky, there is a significant temperature difference between zenith and nadir. Using a non-contact infrared thermometer, a blue sky shows a temperature of 255 Kelvin or less while the ground has a temperature of 295 Kelvin or more. This 40+ degree difference is very easily detected, even with low sensitivity thermopiles or thermistors. There is a very sharp temperature change at the horizon which is easily found, even by wide angle detectors. On cloudy days there is still a substantial temperature difference and on days of low cloud and drizzle or light fog, the temperature difference drops to around 1-3 degrees, still very detectable. This HSAS system has flown with temperature differences as low as 2 degrees with complete success.

8. Black body radiation.

All objects with temperatures above absolute zero emit radiation according to Plank's Law. Figure 3.



Expressed graphically this looks like:-



Planck's Law
Figure 3.

Wavelengths of 7.5 to 13.5 microns correspond roughly to the blackbody radiation from objects in the +75 to -75 degrees Centigrade range. A silicon filter prevents these thermopiles from responding to short wavelengths from the sun (6000 degrees Kelvin – around 550 nm peak).

9. Pitch and roll stabilisation.

A pair of matched thermopile sensors arranged 180 degrees apart and having their outputs feeding a differential amplifier makes a sensitive horizon detector. When both sensors receive equal energy, their axis must be parallel to the horizon and the amplifier output for that channel is zero (or the neutral value). Two sets of sensors, arranged laterally and longitudinally make a roll and pitch detector.

When banked, one sensor sees more cold sky, the other more warm ground and the amplifier output swings accordingly. This signal is used for feedback to the control servos. Essentially identical systems are used for pitch and roll control with different gains to match the different response rates for pitch and roll.

To establish a turn with the fixed sensors, an offset is applied to the neutral value. Figure 4 shows the general arrangement.




Figure 4

In practice, the output of the amplifier is read by an Analog to Digital Converter (ADC) and the offsets are applied digitally. Two variants of this concept have been flown, one with fixed wingtip sensors and the second with a movable sensor plate.




Figure 5. Fixed wingtip roll sensors



Figure 6. Fixed wingtip sensors

Determining the required amount of offset to apply can be handled several ways. In an earlier system, the calibration procedure was to point each wingtip up to vertical before each flight. This established the peak signal available between the zenith and the nadir and the maximum offset allowable was limited to a percentage of this peak signal. This may be acceptable for small hand launched aircraft but it adds an extra step before flight, is impractical for larger aircraft and limits the flight duration to periods of essentially similar thermal conditions. On bright days, a relatively large offset is needed to match the large zenith/nadir signal. That same offset if applied on dull days would be excessive.




Figure 7. Movable roll plate sensors



Figure 8. Movable Roll plate sensor.

Establishing a turn with the movable sensor plate (fig 8) is a matter of rotating the sensor plate by a small servo controlled from the ground, or from the onboard mission controller, in the opposite direction to the desired bank angle. The aircraft responds to the plate movement by adjusting its roll angle to keep amplifier output at its neutral value and so keeping the sensor plate horizontal.

10. Calculating offset for a given bank angle.

Assume we want a bank angle of 30 degrees, what offset should be applied to the sensors to achieve this?

Our sensors have a Field of View (FoV) of 100 degrees in a nominally square shape. In level flight each sensor sees 50 degrees of ground and 50 degrees of sky.

A typical zenith temperature is about 258 degrees Kelvin (−15 C) and that for the nadir is 288 degrees K (+15 C). This temperature differential causes an output to the Analog to Digital Converter (ADC) of X units.

From the geometry of the sensors and our desired 30 degree bank, the lower wing sensor will be filled with 80 degrees of ground and 20 degrees of sky. The temperature average = ((degrees seeing sky)*(sky temperature) + (degrees seeing ground)*(ground temperature))/ (total field of view)

The average temperature seen by the hotter sensor is $(80 \cdot 288 + 20 \cdot 258) / 100 = 282$ degrees K. The upper wing (cooler sensor) sees $(80 \cdot 258 + 20 \cdot 288) / 100 = 264$ degrees K. The average difference between the two tip sensors in this case is 18 degrees. The desired offset ADC count value that we need to inject is $18/30 \cdot X$ counts.

Regardless of what value X happens to be, and it does vary through any long flight as we fly over different ground temperatures or under changing skies, an offset of 60% of the current X will yield a bank angle of 30 degrees for a 100 degree square FoV sensor.

Bank angle degrees	% of Zenith - Nadir difference
5	10
10	20
15	30
20	40
25	50
30	60
35	70
40	80

When using only wingtip sensors for roll, a limit must be placed on the maximum offset applied. Commanding a bank angle where one sensor sees only sky and the other only ground leads to loss of stability. To maintain control in real world conditions, the maximum allowable offset is limited to 80% of the peak sky/ground reading for a maximum bank angle around 40 degrees.

If the offset is limited to 80% of the peak reading and turbulence pushes the aircraft bank to 90 degrees or more, recovery is automatic since the aileron deflection is always in the direction to lower a cold wing and raise a warm wing.

11. Adaptive self calibration & higher bank angles

By adding a third orthogonal channel, the system becomes adaptive and can self calibrate and command stable flight through a complete diurnal cycle or rapidly changing sky conditions (figure 10). By making the vertical and horizontal channels overlap, the system can command bank angles in excess of the 40 degree limit imposed with only wingtip sensors.




Figure 10.

As the bank angle approaches 45 degrees, the vertical and the horizontal channels both see ground and sky. In normal circumstances this reduces the peak vertical channel signal ("X" in the discussion above) and so limits the peak bank angle but by having the software interchange the vertical and horizontal sensors, bank angles from 45 to 90 degrees can be maintained.

12. Implementation

Several demonstration systems have been built around low cost Radio Controlled (RC) Almost Ready to Fly (ARF) trainers. The electronics are shown here in prototype form.

There are two basic variants, one with a movable sensor plate and the second with fixed, wingtip mounted sensors. The fixed wingtip sensor system is simpler but has roll limits while the movable sensor plate has advantages with cameras, compasses or payloads that need to be stabilized relative to the horizon. The control algorithms for both are largely similar.




Figure 11. Movable pitch sensor controlled by a servo under the black tape..



Figure 12. Controls ground test. Note the aileron & elevator deflection while the vehicle is stationary.



Figure 13. Infrared analog section and RC receiver



Figure 14. Digital section with ASI and Altitude sensors, servo, RS-232 interface and battery pack.

Electronic design strategy.

The design approach taken in this system was to eliminate complexity wherever possible. The typical design approach of [2] requires a complex multiprocessor real time operating system (RTOS) to harness the power of the five processors. Large single processor systems have functionally different code modules all sharing the same memory and software engineers can inadvertently compromise code written by other engineers. This makes code validation very difficult and execution times may be highly variable.

The programming approach taken with the HSAS system to guarantee execution time by using in-line code, no interrupts and dedicated micro-controllers for functionally separated tasks.

This leads to very fast execution times and minimal code interaction. Provided the i/o specification is obeyed (and this is readily validated), modifications to one functional block cannot affect the integrity of other controllers.

Software for each functional block is carefully debugged, burnt into a PIC micro-controller and thereafter treated as a piece of hardware with a known transfer function.

New functions can be handled by adding extra dedicated processors if there is any danger of code interaction or exceeding timing limits. Mid-range PIC micro-controllers from Microchip Incorporated are used. These are low cost, fast, low power, capable and reliable. They need few support components making for a compact and affordable design. No interrupts means timing can be tightly specified and guaranteed in all operating conditions. Three PICs (P1, P2 & P3) fly the aircraft, all running asynchronously with clearly defined i/o protocols between each chip. P1 looks after received data from the ground controller. P2 is the master controller, stabilising the aircraft in flight. P3 is the Mission Controller handling the flight plan, communications and some data logging. P2 is capable of flying the aircraft even if P1 and/or P3 stop functioning. Payloads are controlled by their own separate PICs so that rogue code cannot possibly compromise the flying function.

Code revision can be performed in the field if necessary via an RS-232 boot-loader while the processors remain installed inside the aircraft.

The next printed circuit board iteration will reduce the system down to a single unit approximately 50 x 75 mm and eliminating half the connectors.

Receiver.

The receiver is currently a 36 MHz hobby RC receiver but the architecture allows any receiver capable of sustaining around 9600 bps. The system currently uses Pulse Position Modulation (PPM) although encrypted channels, packetised or continuous data at any desired frequency and data rate can be handled by changing P1 with no changes needed for P2 or P3.

Analog card.

This amplifies and buffers the signal from the six thermopiles and the receiver. The thermopiles are arranged in three orthogonal axes for pitch, roll and sky/ground.

Receiver decoder.

The buffered receiver output is fed to a dedicated PIC 16F84A micro-controller (P1) which decodes the PPM information and assembles a packet of data where the eight bytes in the packet represent the pulse widths of the 8 RC channels. P1 also detects loss of signal or interference and generates failsafe packets in those cases. P1 presently uses about 320 words of code. This design approach allows radio substitution with only minor changes to P1 and no impact on other processors in the system.

Main processor.

The second PIC (P2) is a 16F877 with both analog (10 bit) and digital channels. This is the main flight controller. The software architecture is a single loop with no interrupts for a guaranteed cycle time less than 20 mSecs in all operating modes to ensure the servos are always refreshed at full rate and the ground pilot experiences negligible control delay. P2 starts the cycle requesting a packet of data from P1 describing the 8 RC channels. P2 reads the IR pitch, roll and sky sensors then the altitude and airspeed channels. P2 requests a data packet from P3 if used and then computes the servo values to drive the control surfaces.

Currently eight channels can be driven by the hardware and V tail, delta, flying wing and conventional servo arrangements are supportable.

Operating modes.

There are three operating modes. Local radio Control (LRC), Stability Augmented LRC (SALRC) and fully autonomous operation (AUTO).

In **LRC** mode, all ground pilot commands are passed straight to the servos within 40 mS of transmission. The aircraft behaves as a normal RC model with no discernable delays.

SALRC is selected by commanding the wheels up position. In this mode P2 flies the aircraft, keeping it within a safe operating envelope at all times. The ground operator can issue proportional left and right turn commands up to a 40 degree bank angle via the aileron stick. The operator can issue climb, hold height or descend commands via the throttle stick. Takeoffs are selected by commanding full throttle. The pilot can steer left or right during the ground run. The take off and climb out at a fixed pitch attitude is controlled automatically. Landings are selected by closing the throttle and the aircraft descends rapidly until an altitude of 20 metres is detected by the onboard altimeter where the descent rate is slowed. Again the ground pilot just steers left and right and the pitch and bank are handled automatically. Landings will be improved with the

addition of a ground sensor for final flare out. RF, Ultrasonic and IR ground detectors have been experimented with for this task. The smallest and simplest so far has been the IR system.

AUTO is selected by setting both the gear up and an auxiliary channel to maximum pulse width. A preloaded flight plan stored in P3 is then executed following GPS coordinates.

Mission controller.

A third PIC 16F877 (P3) serves as a mission controller. This reads current GPS position and executes the stored flight plan. It sends turn and height commands to P2 in essentially the same way as they are sent from the ground controller.

Range safety & flight termination.

In its current form, the system is always under supervisory control from the ground. At any time the ground controller can issue a command to revert to manual control. If P1 detects no signal (flown out of range perhaps) or significant interference then it issues a failsafe sequence to P2. This idles the engine, slows the aircraft for minimum kinetic energy then commands a gentle turn and descends to the surface. Should P1 or P3 stop functioning, P2 will see anomalous data from these controllers and P2 will execute its own failsafe strategy.

The system is capable of significant distance fully autonomous flight, with appropriate Civil Aviation Safety Authority permits.

13. Semi stabilised platforms.

Large UAVs carry sophisticated gyro stabilised Electro Optic sensor systems. These are usually well outside the payload, power supply and cost constraints of small UAVs. A significant problem with photos from small UAVs is knowing the GPS coordinates of the image. Photos taken during a turn, where the slant angle and heading at the time of the exposure are not known, requires considerable photo interpretation to geo-reference any object found.

A nose mounted movable sensor plate remains closely aligned with the horizon at all times and so a down looking camera will always have the centre of the photo at the current aircraft GPS coordinates making for simple interpretation. Such a platform, while not as accurately stabilised as a full E-O system, is a very useful location for a magnetic compass, small antenna or beacon or any payload needing horizon referencing at a tiny fraction of the cost.

14. Cloud and terrain avoidance.

The current system operates in VMC. Ground tests indicate it should work in fog down to perhaps 200 metres visibility where there is still a small vertical temperature gradient but this has not been flight tested yet. It will not function inside cloud where all the droplets surrounding the aircraft are of essentially the same temperature and the vertical temperature gradient has been disturbed by turbulence.

The target application of this very low cost system is mainly photo reconnaissance where the camera payload will also not work in cloud and so the flight is unlikely to take place. Typically the UAV will be despatched to a relatively nearby location, say within 100 kms, where the condensation level can be expected to be similar to local conditions and the flight programmed above or below cloud as appropriate.

By steering a narrow field of view sensor and taking a series of absolute temperature readings, a thermal map of the horizon can be built up. This requires lensing the sensor to reduce the angle of view and scanning the horizon during flight to build up, say, a 32 pixel low resolution map of the surrounding horizon. This can point to a cloud or terrain in the path of the aircraft. The system can be a useful complement to a gyro stabilised system, where a thermal image of surrounding terrain or cloud could reduce collision incidents.

15. Conclusion and future directions

The infrared Horizon Sensing Attitude Stabiliser (HSAS) is a highly effective and very low cost system for small to medium sized fixed wing UAVs operating in both day and night Visual Meteorological Conditions. It potentially has problems in canyons and so is not recommended for flying at low level in cities but for almost any flights more than a few tens of metres above terrain it has proven to be a satisfactory system.

Driving down UAV operating costs requires all components of the system to be low in cost, simple to operate and robust. Airframe & avionic build costs, crew size and skill requirements must be reduced if we are to compete against general aviation alternatives. The launch, in-flight management and landing workload must be eliminated where possible.

Automating the takeoff and landing, where a great many incidents and accidents normally occur, reduces attrition costs and protects valuable payloads.

Delivering a stable flying platform that researchers or students can use without extensive training improves experimental productivity.

The Horizon Sensing Attitude Sensor is a viable building block for a range of very low cost UAVs for civil, military and academic uses in VMC.

Future effort will focus on simplifying flight planning and improving autonomy, particularly in the landing phases. Takeoff and landing are the high cost phases for most UAVs as they need skilled operators.

Our goal is a system that can safely and gently carry third party payloads up to 10 kg through an entire flight sequence from takeoff to landing with a one man crew and no manual intervention during the flight.

16. References.

- [1] S. B. Stancliff, M. C. Nechyba "Learning to fly: Modeling Human Control Strategies in an Aerial Vehicle."
 - [2] F. Valentini, C. Bil, P. Riseborough "Development and trials of an autonomous Flight Control System for UAVs"
 - [3] S.M. Ettinger, M.C. Nechyba, P.G. Ifju, M. Waszak "Vision guided flight stability and control of micro air vehicles" <http://www.mil.ufl.edu/~nechyba/mav/iros2002.pdf>
 - [4] Futaba web site (<http://www.futabarc.com/radioaccys/futm0999.html>)
 - [5] A.K. Brookman. RAE TR 71153 July 1971 An albedo horizon sensor using hybrid circuitry.
- Apart from the mechanical arrangement and number of sensors used, this appears to be the forerun-

ner work of the Ripmax and Futaba optical stabilisation systems.

- [6] NASA TN D-1005 Langley January 1962 Study of a proposed IR horizon scanner for use in space orientation control systems. N.M.Hatcher, E.F.Germann
- [7] NASA TN D-2995 Langley September 1965 Development & Testing of a proposed IR horizon scanner for use in spacecraft attitude Determination. N.M.Hatcher, A.L.Newcomb, N.J.Groom
- [8] TN D-4045 Goddard Space Flight Centre August 1967 Theoretical computations of the outgoing IR radiance from a planetary atmosphere. V.G.Kunde
- [9] TN D-4654 Langley July 1968 Radiometric measurements of the earth's IR horizon from the X-15 in three spectral intervals. A. Jalink, R.E.Davis, D.E.Hinton.
- [10] TN D-6616 Langley Feb 1972 Conceptual Design and analysis of an infrared horizon sensor with compensation for atmospheric variability. A. Jalink, R.E.Davis, J.A.Dodgen
- [11] US Patent US6181989: "Aircraft attitude sensor and feedback control system" Inventor: Gwozdecki, Joseph Andrew; Houston, TX 77002-1135 Patent issued 30 January 2001

If we substitute 'aircraft' for 'spacecraft' in the introduction to NASA Tech Note TN D-6616, the NASA sentence becomes "A simple method to determine aircraft attitude is to sense the atmospheric radiation gradient of the earth's horizon". This describes the working principle and possibly the design origins of the Gwozdecki/FMA patent US6181989, probably invalidating that patent on the grounds of prior disclosure.