

Remove Before Flight

Defect-Free Software and Agile Development in SPARK 2014

Martin Becker

Chair of Real-Time Computer Systems (RCS)
Technical University of Munich

Presented at
Frama-C & SPARK Day 2017, Paris
May 30th, 2017



Lehrstuhl für
Realzeit-Computersysteme

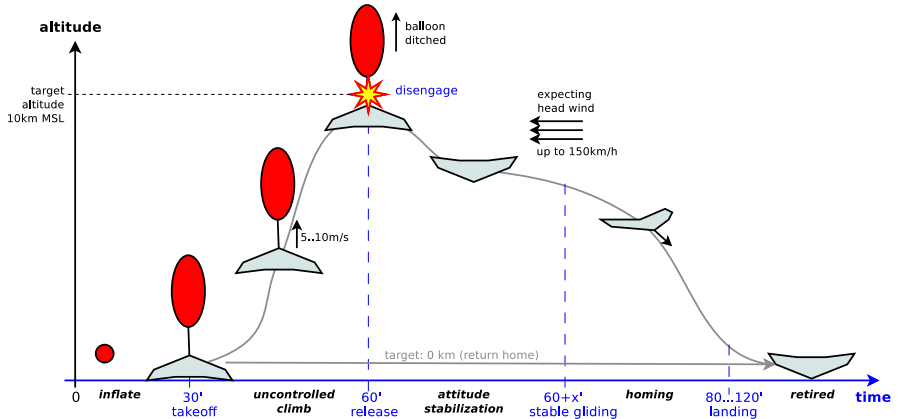


Does it work with “bugs”?
Which defects are removed, and when?

- 1 Briefing (Introduction)
- 2 Pre-Flight Checks (Development Process)
- 3 Request to Taxi (Verification Outcome)
- 4 Ready for Takeoff / En Route (Glider Launch)
- 5 Incident Report (SPARK Forensics)
- 6 Debriefing (Conclusion)

Briefing (Introduction)

Mission – A Novel Weather Balloon



Need to: monitor ascend, unhitch from balloon, return to take-off location (& record weather data)

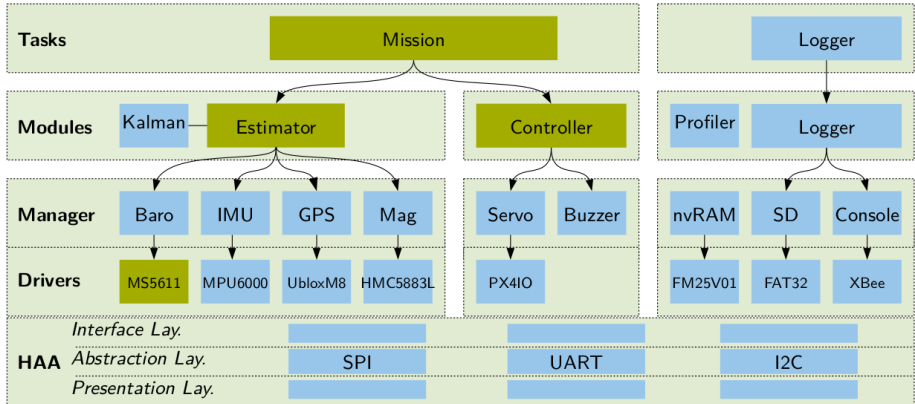


- high-altitude micro glider
 - > 10km altitude, 30-100km/h air speed
 - non-motorized vehicle, MTOW 400g
- need a flight stack
 - estimate attitude, stabilize flight, navigate w/ GPS, ...

- “Pixhawk” autopilot, 168MHz ARM Cortex-M4(F)
 - 2MB Flash, 256kB RAM, SD Card Interface, NVRAM
 - Co-Processor with “failsafe” functionality for servo output
 - (cross-verification)
- 1x GPS, 1x Barometer, 1xMEMS+Gyro, 1x Magnetometer, 2x Acuator, 1x Buzzer
- Ada/SPARK Runtime System: Ravenscar Small Footprint
 - determinsting & analyzable multi-threading with constraints
 - RTS must be ported from similar Cortex-M4 (F409)
- also: based on AdaCore’s Driver Library/Bareboard Code
- need: custom drivers for our hardware, attitude estimator, mission handling, homing functionality

3 months time, 2 developers

Anatomy of the Flight Stack



- Flight-critical (“mission”) and non-flight-critical (“logging”) tasks communicate via message queue (protected object)
- goal: full SPARK coverage everywhere except HAA + **isolation of tasks**

Pre-Flight Checks (Development Process)

- avoid loss of airframe (money and time)
- Germany: must not fly above 100m GND, final launch requires permission, insurance and manpower
- some low-altitude testing (risky and of limited use)

The Role of (Integration) Testing

- avoid loss of airframe (money and time)
- Germany: must not fly above 100m GND, final launch requires permission, insurance and manpower
- some low-altitude testing (risky and of limited use)



- avoid loss of airframe (money and time)
- Germany: must not fly above 100m GND, final launch requires permission, insurance and manpower
- some low-altitude testing (risky and of limited use)
- goal: “de-bug” before going for any flight
 - unit verification through static analysis (SPARK2014, GNATprove), separation of criticality,
 - at least: no exceptions during flight, keep flight-critical task alive

- avoid loss of airframe (money and time)
- Germany: must not fly above 100m GND, final launch requires permission, insurance and manpower
- some low-altitude testing (risky and of limited use)
- goal: “de-bug” before going for any flight
 - unit verification through static analysis (SPARK2014, GNATprove), separation of criticality,
 - at least: no exceptions during flight, keep flight-critical task alive

As little as possible, only when software is stable.

- language intended for formal verification, successor of SPARK
 - strongly typed, imperative, object-oriented
- adopted syntax of Ada 2012 \Leftrightarrow both languages can be mixed within an application
- SPARK 2014 = "constrained subset of Ada 2012":
 - no access (pointers), no aliasing, no exception *handling*, ...
 - need for proving the absence of exceptions
- user can specify functional contracts, data flow contracts and assertions
 - first-order logic, executable semantics, IEEE-754 Floats
- static analysis (i.e., deductive theorem proving)
 - proof for absence of errors (exceptions + failing contracts)
- Used Implementation: AdaCore's SPARK 2014, version GPL 2016 (partially Pro 17)

Ways of Finding Defects

- 1** most by static analysis (each developer & nightly runs)
 - replace unit testing (no harness + fixture, no env. sim.)
 - pinpoints defects and some amount of under-specification
- 2** few by integration testing
 - defects which were missed by static analysis
 - defects which require context beyond source code
 - logging of exceptions: pinpoint infections, no chasing through cause-effect chains
 - defects of three kinds:
 - masking defects during analysis / careless usage of verifier
 - ignoring failed proofs / wrong process
 - wrong “verification fixes” (saturation) / incomplete specification
- 3** none during operation
 - nevertheless: logging of exceptions & in-air reset

Nightly “deep” verification runs with long timeouts

- Jenkins server on Intel Xeon E5-2680 Octa-Core, 2.7GHz, 16GB RAM
- GNATprove works *all cores*, allowing for high steps

Output: Verification summary, sent to all developers

- complete verification log (`file.adb:13:12: overflow check might fail (e.g., when ...)`)
- `gnatprove.out` (type of property, % verified, which prover)
- **a custom verification summary in tabular form:**

Nightly Verification Table

compilation unit	props success	flows success	props proven	ents	cover	num props
px4io.protocol	100	100	54	1	100	54
bounded_image	100	100	35	11	100	35
magnetometer	100	100	18	7	100	18
servo	100	100	10	9	100	10
mpu6000.register	100	100	1	1	100	1
ublox8.driver	99.0	100	98	28	100	99
px4io.driver	98.4	100	124	19	100	126
main	97.4	100	38	4	100	39
mission	95.6	100	86	25	100	90
units.navigation	85.4	100	129	16	100	151
hil.devices.nvram	76.2	100	16	10	100	21
ms5611.driver	75.4	100	135	45	100	179
controller	67.3	86.9	111	46	100	165
kalman	65.0	100	91	18	100	140
imu	50.7	100	70	20	100	138
nvram	100	100	15	27	92.6	15
logger	96.9	100	62	21	90.5	64
ulog	98.5	100	133	14	85.7	135
mpu6000.driver	74.6	100	126	29	79.3	169
...

- different view than GNATprove report
 - success (% discharged)
 - coverage (%SPARK “on”) per unit
 - absolute counts
 - ...
- ordered by coverage, then success
- goal: promote flight-critical units to the top
 - developers work on critical parts first
- decision taken every morning
- code base grows quickly, but is kept in check

Latest version: separation by task (“flight-critical” table vs “non-critical” table)

Request to Taxi (Verification Outcome)

Catching Missing Requirements

An effective way to avoid overflows and range errors?

```
1 baro_alt : Altitude_Type;  
  if Float (alt_int) >= Altitude_Type'Last then  
    baro_alt := Altitude_Type'Last;  
  elsif Float (alt_int) <= Altitude_Type'First then  
    baro_alt := Altitude_Type'First;  
6 else  
    baro_alt := Float (alt_int16);  
  end if;
```

Even better (create a generic and use it all over the place):

```
function Sat_Cast_Alt is new Saturated_Cast (Altitude_Type);  
2 baro_alt := Sat_Cast_Alt (alt_int16);
```

Catching Missing Requirements

Consider this¹:

```
data_rx : Byte_Arr (0..91) := Read_From_Device (GPS_UART_0);  
-- ...  
3 subtype Lat_Type is Angle_Type range -90.0 .. 90.0;  
Lat : Lat_Type := Sat_Cast_Lat(Float(data_rx(28..31))*1E-7);
```

Static analysis: no range errors. However:

¹code simplified

Catching Missing Requirements

Consider this¹:

```
1 data_rx : Byte_Arr (0..91) := Read_From_Device (GPS_UART_0);  
  -- ...  
  subtype Lat_Type is Angle_Type range -90.0 .. 90.0;  
  Lat : Lat_Type := Sat_Cast_Lat(Float(data_rx(28..31))*1E-7);
```

Static analysis: no range errors. However:

- failure of GPS device \Rightarrow silently ignored
- missing error handling for component/subsystem failure
- majority of such cases was missing a requirement


Saturation as means to “silence the prover” should be avoided

¹code simplified

- using the GNAT dimensionality system:


```
1 angle : Angle_Type := 20.0 * Degree;  
   dt : Time_Type := 100.0 * Milli * Second;  
   rate : Angular_Velocity_Type := dt/angle; -- comp. err.
```

- majority of floats range-limited
- how hard can it be?




Attacking the Floats

Floats are hard to verify (but possible)



Attacking the Floats

Floats are hard to verify (but possible)



Ways out: SPARK lemma library, CodePeer

- CPU goes into reset loop during pre-flight checks
- log: exception in `units.vectors.rotate()`
- debugger (extremely hard to reproduce):

```
3  ▾ • package body Units.Vectors with SPARK_Mode is
4  •
5  ▾ •   procedure rotate
6  •       (vector : in out Cartesian_Vector_Type;
7  •         axis   :           Cartesian_Coordinates_Type;
8  •         angle  :           Angle_Type) is
9  •         result : Cartesian_Vector_Type := vector;
10 •
11 •   begin
12 ▾ •       case (axis) is
13 •         when X =>
14 → •         result (Y) := Cos (angle) * vector (Y) - Sin
```


`Units.Vectors.rotate`

Messages	Locations	Debugger Execution <6>	Debugger Console <6>	Debu
#1 0x08015a9c in units.vectors.rotate (vector=..., axis=x, angle=0.00429291604) at				
(gdb) info locals				
result = (0.00212701317, 0.0383495204, -2.02303554e-38)				
(gdb) print angle				
\$11 = 0.00429291604				


- CPU goes into reset loop during pre-flight checks
- log: exception in `units.vectors.rotate()`
- debugger (extremely hard to reproduce):
- static analysis had proven absence of runtime errors
- after some detour into assembly: run-time system configured incorrectly for target
 - config: FPU produces no denormals (“flush to zero”)
 - reality: FPU *does* produce denormals
 - internal `Float.Valid` check raises exception
- what happened: glider motionless long enough → gyro rates very small ($2E-38$) → rotating by small pitch/roll angle (glider level) → “numerical underflow”
- fix configuration of run-time system, mismatch gone

Reports for daily use

Success per Unit:




Unit SPARK coverage:




Reports for daily use



Why checks fail in unit:



Success vs. Time



Computational Effort




How we found defects

- 1 most by static analysis (each developer & nightly runs)
 - removed all “stupid bugs”
 - identified missing requirements (saturation)
- 2 few by integration testing
 - masking defects during analysis / careless usage of verifier
 - ignoring failed proofs / wrong process
 - incorrect config of run-time system / violating prerequisite for analysis
- 3 none during operation?

Ready for Takeoff / En Route (Glider Launch)

Final Flight (1)





Data Link


RSSI
Status

WAIT
Open

A/C Status

Lat	11.	0s
Lon	48.	0s
Alt	6650.1	0s
Hdg	0	0s
G'Vel	WAIT	2856
AltMSL	6650.1	0s
G'Speed	71.4131	0s

0 tiles remaining No GPS (No such file or directory) xbee



Data Link

RSSI	WAIT
Status	Open

A/C Status

Lat	11.	0s
Lon	48.	0s
Alt	5731.2	0s
Hdg	0	0s
G'Vel	WAIT	3041
AltMSL	5731.2	0s
G'Speed	83.7474	0s

0 tiles remaining No GPS (No such file or directory) xbee



- navigation failure, but stable flight
- glider did not return back home
- glider could not be recovered
 - ironically: stand-alone backup localization device had failed
- How to find out what went wrong?

Incident Report (SPARK Forensics)

- self-tests at startup had been passed
- launch and unhitch went according to plan
- Continuously flying **left circles** until landing:
 - stable flight is only possible with healthy airframe + controls
- Possible Explanations: another unexpected exception, or no GPS fix, or lost home position, or in-air reset with unsuccessful resume, or sensor error or numerical error, or, or, or ...

Luckily, high-level behavior of glider was encoded and verified in SPARK

```


procedure Update_Homing with
  Global => (Input => (G_Object_Pos, G_Target_Pos),
             In_Out => (G_state), Output => Ctrl_Mode),
  Depends => (G_state => (G_state, G_Object_Pos, G_Target_Pos)) is
5 begin
  G_state.once_had_my_pos := G_state.once_had_my_pos or have_my_pos;
  if have_my_pos and then have_home_pos then
    if G_state.distance_to_home < Config.TARGET_ZONE_RADIUS then
10      Ctrl_Mode := MODE_ARRIVED;
    else
      Ctrl_Mode := MODE_HOMING;
    end if;
  elsif have_home_pos and then (not have_my_pos and G_state.once_had_my_pos) then
    Ctrl_Mode := MODE_COURSEHOLD; — temporarily lost nav => hold course
15 else
    pragma Assert (not have_home_pos or not G_state.once_had_my_pos);
    Ctrl_Mode := MODE_POSHOLD; — don't know where to go => hold position
  end if;
end Update_Homing;

1 procedure Compute_Target_Attitude with
  Global => (Input => (G_state, Ctrl_Mode, G_Object_Att, Ada.Real_Time.Clock_Time),
             Output => (G_Target_Att)),
  Contract_Cases => ((Ctrl_Mode = MODE_COURSEHOLD) =>
6      G_Target_Att.Yaw = G_Target_Att_Prev.Yaw, — roll open
    (Ctrl_Mode = MODE_HOMING) =>
      G_Target_Att.Yaw = G_state.course_to_home, — roll open
    (Ctrl_Mode in MODE_POSHOLD | MODE_UNKNOWN) =>
      G_Target_Att.Roll = -Config.CIRCLE_TRAJECTORY_ROLL, — yaw open
11    (Ctrl_Mode = MODE_ARRIVED) =>
      G_Target_Att.Roll = Config.CIRCLE_TRAJECTORY_ROLL), — yaw open
  Post => G_Target_Att_Prev = G_Target_Att;

```

- enforcing exactly one out of four behaviors
`{MODE_HOMING, MODE_COURSEHOLD, MODE_POSHOLD, MODE_ARRIVED}`
- conditions for circling left:
 - `MODE_POSHOLD`: only if no home position or never had GPS fix
⇒ no, mission starts only with these (acoustic and optical pre-flight checks from glider)
 - `MODE_COURSEHOLD`: no, would fly straight ⇒ glider had GPS fix
 - `MODE_ARRIVED`: no, would circle right ⇒ no calculation error
 - `MODE_HOMING`: yes, could circle left
 - permanent in-air resets: extremely unlikely, as this causes twitches (trajectory smooth). Even if, last values would be restored from NVRAM.
- ⇒ Glider was in homing mode. Only two explanations left:
 - 1 mag/compass broken (would have been detected by BIST)
 - 2 mag/compass provided unexpected output (e.g., “cannot find north”)

Reviewing Previous Flight Logs...



Suggested explanation: Magnetic distortion in avionics bay

Debriefing (Conclusion)

Ways of Finding Defects

- 1 most by static analysis (each developer & nightly runs)
 - removed all “stupid bugs”
 - identified under-specification
- 2 few by integration testing
 - masking defects during analysis / careless usage of verifier
 - ignoring failed proofs / wrong process
 - wrong “verification fixes” (saturation) / incomplete specification
 - incorrect config of run-time system / violating prerequisite for analysis

Ways of Finding Defects

- 1 most by static analysis (each developer & nightly runs)
 - removed all “stupid bugs”
 - identified under-specification
- 2 few by integration testing
 - masking defects during analysis / careless usage of verifier
 - ignoring failed proofs / wrong process
 - wrong “verification fixes” (saturation) / incomplete specification
 - incorrect config of run-time system / violating prerequisite for analysis
- 3 *one during operation*
 - faulty but non-crashing behavior
 - missed during integration testing
 - unverified assumptions about sensor data
 - beyond context of source code

- writing code takes more time in SPARK 2014 (cmp.: C)
 - but some agility can be applied to speed up progress (write first → fix critical → fix uncritical)
- very little debugging work
 - practically no exceptions during system testing
 - no working through cause-effect chains
 - no minimization of problem cases
 - no isolation of failure causes
 - no reproduction issues
- concentrate on getting functionality right which is not modeled in SPARK
 - effect of Kalman Filter
 - evolution of sensor data
 - underspecification (what happens after GPS goes silent?)

- SPARK 2014 and tools work very well
 - priceless execution semantics of annotations
 - many defects are avoided with almost no additional effort
 - no excuses for uninitialized variables
 - floating-point numbers are usable, but need some work (instantiate lemmas, refactoring of code, longer time to verify)
 - testing can be significantly reduced (not system testing)
 - high-level behavior can and should be encoded
- Ravenscar+SPARK: easy and effective multi-threading, separation of criticality, must-have
- not evaluated in detail: OO (not for us) and flow dependencies (for being a cousin of “const hell”)
- code being released open source

Reached Parking Position – Questions?



To appear: *Development and Verification of a Flight Stack for a High-Altitude Glider in Ada/SPARK 2014*, M.Becker, E.Regnath, S.Chakraborty, Computer Safety, Reliability and Security (SAFECOMP), 36th International Conference, Trento, Italy, 2017.