

A closer Look at Vision

Tino de Rijk, The Netherlands, January 2007

Introduction

In November 2006 an experienced Inspiration diver reported a strange incident: during a dive in the Red Sea, his CO2 sensor TempStik readings suddenly flat lined, and his PO2 readings were frozen (as in: not changing anymore, stuck to the same value) on the wrist display.

It triggered a lot of E-mails with people wondering if the Inspiration would still be life-supporting in this stage. It occurred to me that a lot of reasoning about this was by Vision and non-Vision divers that still reason with the "old" Inspiration Classic (the one with the dual handsets) in mind. In a Classic, a frozen display typically indicates that that controller isn't properly functioning anymore. However, the total hardware setup of the Vision is so radically different from the Classic that these comparisons have little value anymore.

Having translated the Vision manuals into Dutch, I know that not all text about the inner workings of the Vision, and especially of the workings in case of component failure(s), is nicely put together in one place. So I compiled a set of questions, fired them off to the as usual friendly & helpful people at Ambient Pressure Diving, and as usual I got a very extensive reply back pronto.

The rest of this memo is split in two parts:

- 1) A closer look at Vision: which component of Vision controls which function?
- 2) A "what if component x fails?" section, digging into what you might see during a failure, and what it could mean to you in terms of appropriate reactions.

Here we go!

A closer look at Vision: which component controls which function?

Vision component sum-up

The vision Electronics consist of three controllers, three displays, two batteries, an intelligent bus, a buzzer and an optional TempStik scrubber monitor:

- 1 computer in the wrist display unit, further referred to by me as the WDC (Wrist Display Controller);
- 2 redundant, but independent controllers in the lid of the scrubber with as main function oxygen level (PO2) control, referred to by me further on as SPC (Setpoint Controllers);
- A serial databus that connects the WDC with the SPC's and also houses a real-time clock (RTC) to supply the logbook functions of the Vision with the correct date & time. It also contains a "hub" to provide electrical and functional isolation in the event of a failure of any one component hooked up to the databus;
- Physically, the databus bus is represented by the electric cables running through the rubber medium-pressure hose between scrubberlid and WDC;
- 1 LCD wrist display, built into the WDC that shows the detailed status of the system;
- 2 redundant, but independent HUD displays, comprising a green and red light each, placed on top of the mouthpiece. It is important to note that the actual LED's themselves are located in the scrubber lid, near the SPC's, and the light is transported to the end of the HUD through a fiber-optic cable. This was done to avoid having to use a thin, vulnerable electric cable all the way to the mouthpiece and having to use relatively big LED's there, as in most other HUD designs.
- 2 batteries (non-rechargeable CRP2 type Lithium), placed next to the SPC's in the scrubber lid, in a water-resistant housing.
- A "buzzer on a stalk", located to the left ear, and driven by the SPC's.

- An optional TempStik, that measure the progressing warmth front that travels through the scrubber bed as the Sofnolime progressively gets spent, causing an exothermic reaction in the process. The measurement is done by a number of serially placed temperature sensors in the central pole of the scrubber container.

Functions per component

So, after this sum-up, let's look at which component does what. This is VERY important in order to determine later on which functionality you may lose if a given component fails on you!

- Contrary to popular believe, and probably also driven by wrong comparisons between the Classic and the Vision, the WDC is in my view NOT the most important component of Vision. Yes, it gives you a detailed view of it's status, and it lets you control all settings of the system, but it also does NOT do some important things:
 - o it does NOT control the PO2 (setpoint maintenance),
 - o it does NOT control the solenoid,
 - o it does NOT control the two HUD's (that is done by the SPC's: each SPC runs its own HUD),
 - o it does NOT control the buzzer, and
 - o it does NOT control power management (switching between batteries, and combining them, in case of an empty battery)
- The WDC does some other important things though:
 - o It contains the pressure transducer (digital pressure gauge) for depth measurement. It is located in the bottom of the WDC, shielded by the bottom plate that holds the straps;
 - o It executes all CCR and OC decompression calculation functions;
 - o It does control the value of the two setpoints, and (auto-)switching between them;
 - o It contains the round-robin (FIFO) log memory for alter readout of dive profile and warnings.
- So, losing the WDC does NOT mean losing a reliable setpoint maintenance, or even the ability to verify that, as the HUD's are driven directly by the SPC's (each SPC has its own HUD). However, you would lose your decompression information, but more on that later. Realize also that for proper CCR deco calculations, you only need three external inputs:
 - o depth (coming from the built-in pressure transducer in the WDC),
 - o time (coming from its internal oscillator in the WDC; see below) and
 - o PO2 (supplied to the WDC by the SPC's, its datastream transported to the WDC over the databus).

For OC calculations you even only need depth and time and the gas mix, not more; the rest is algorithm, executed by the program code in the WDC.
- The RTC is located near the SPC's on the databus; it is only used for time-stamping logbook entries (dive profile, warnings and errors). However, each controller (WDC and the 2 SPC's) maintains its own internal tick count, based on its own built-in oscillators.
- The solenoid is controlled by both SPC's, also different from the Classic. The Master SPC controls the setpoint and thus the solenoid if everything is normal and o.k., but the solenoid is linked through a so-called "wired-or" configuration to both SPC's. The secondary SPC (slave is actually not a good word anymore in Vision terms) also independently monitors the PO2 (through its own cell readings), and starts acting intervening when the PO2 falls below 80% of the current target setpoint. This happens without the need to switch from master to slave, as on the Classic, and acts as a "catch all" safety-net function.
- A similar thing happens for power management. Normally the primary ("master") SPC will perform battery management (switching to the secondary battery in case of battery-low, and later combining the two parallel when both are low). But also here, the secondary SPC can intervene if the active "master" SPC is not able to perform this function.

- The SPC's also contain the calibration correction factors, established during calibration of the unit. As in the Classic, each SPC contains and retains its own set of correction factors: 3, one for each cell.

Now, let's look closer at the WDC:

- As mentioned before, it contains the log memory, on non-volatile, round-robin, FIFO memory. This means a "dead" battery will not cause a loss of logbook info, as unfortunately many diving computers do.
- The WDC receives the following data from the SPC's over the databus:
 - o Real-time PO2 info, supplied by the SPC's;
 - o Battery status info, also supplied by the SPC's;
 - o TempStik data, received directly from the TempStik (no SPC "interference" here);
 - o PO2- and cell-related warnings and errors, as generated by the SPC's.
- The WDC on its turn can send commands to the SPC's:
 - o Power up/down;
 - o Do calibration;
 - o Change setpoint. The SPC however does a validation of the new setpoint request. It will e.g. refuse a new setpoint of 2.0 if for some error reason the WDC would send that.
 - o Switch setpoint. This is effectively a setpoint change, initiated by the WDC, either on instruction of the user (long press of center button), or by the auto-switch function, driven by passing a certain depth. Remember that only the WDC 'knows' the current depth!
 - o Warnings and errors that are not PO2-related, e.g. decompression errors and warnings. They are sent from the WDC to the SPC with the sole purpose of showing them on the HUD's and the buzzer, as both of these are controlled by the SPC's.
- The WDC cannot trigger the buzzer by itself: it needs a SPC to do that. The same applies to displaying errors and warnings on the HUD's.

"What if component x fails?"

In this paragraph I will not go into what bail-out or other actions are appropriate if a component fails. That is up to the instruction agencies or your smart self to decide. E.g.: do you stay on the loop or not in case of a WDC or databus failure? In my PERSONAL view it is totally justified to do that, flying it on the HUD's, but at the same time for sake of "peace of mind" I could very well defend going OC, as you have less info to draw on.

So by now it should be clear that e.g. auto-switching, or in fact any change in setpoint, will not occur anymore if the WDC malfunctions, or the databus link is lost. The latter was the most likely reason of the failure reported by the diver in the introduction, by the way.

If the databus link is lost (e.g. broken cable in the wiring loom between WDC and SPC's, the following will occur:

- The WDC will not show the proper PO2 numbers anymore. In the current version of the Vision firmware (V02.01.02) this shows up as a seemingly frozen PO2 display. However, APD promised me that in a near future version of the software (likely already in the next V02.01.03 version), the WDC will show asterisks (*.**) instead of the last received value when it does not receive proper PO2 data anymore. This will make it much easier to spot, even though static numbers should be a clear tell-tale of trouble for a CCR diver.
 - o I stress the wording "frozen PO2 display" instead of "frozen computer", as the other display functions (e.g. decompression, time, depth) continue showing valid data if only the databus fails;
- The TempStik data will not show anymore, its black and white temperature progression bar replaced by dashes. This was also reported by the same diver;

- Depth and time continues to run, and is logged correctly in memory, as can be clearly seen in the logbook download file supplied by the affected diver on Internet. This includes things like ascent alarms till the very end of the dive, but at the same time will not contain proper PO2 data and/or PO2 related warnings and errors anymore.
- Switching off is not possible anymore: the WDC cannot instruct the SPC's to do this anymore.
- Setpoint changing (either auto or manual) cannot be performed anymore. The most recent setpoint is maintained;
- Setpoint maintenance however can be verified: a quick squirt on the manual O2 button, followed by a diluent flush, will run the HUD's and buzzer through all main functions: rapid blinking and red due to setpoint higher than 1.6, followed by green, followed by slow blinking and red due to setpoint low, followed by green again when the SPC instructs the solenoid to open till the setpoint is o.k. again. This is a quick set of actions that can easily be performed.
- Your CCR deco-calculations are however not linked to the actual PO2 value in the loop anymore: they will feed on the last received frozen PO2 values or, in the next version of the firmware, the current static setpoint value.
You could elect to switch the computer to OC-mode, even if you decide to stay on the loop. As this is in most cases safe, because OC decostop-times are typically longer, this would be a justified action. However, in the before mentioned upcoming firmware update (V02.01.03), the unit will default to the current setpoint value and so will act as if it is an unlinked deco computer and as such can safely be used. Very much like a Buddy Nexus or a VR3 without 4th cell!
- Similarly, in the current version of the firmware (V02.01.02), if the frozen PO2 values are close to the setpoint then the value used is again sufficiently close to act as an unlinked deco computer, and is again safe to use. Proper dive planning should however ensure you always have tables as backup.

Let's run through some other, more simple error scenario's, to be sure we're all on the same page by now:

- if a SPC fails, its HUD will go off (its lights will go off), and its values will not show anymore on the WDC. (currently shown as all zeros, newer version will show as asterisks)
- If you suspect a SPC of being frozen, do the "squirt with O2 + dil flush" test as described above. In that case the WDC can instruct the secondary SPC to take over. This SPC already has the above described "catch all" function that kicks in if the PO2 drops to 80% of the current setpoint.
- If the WDC or databus fails, deco- and ascent-related errors will not pop up on the HUD's, or trigger the buzzer.
- If the WDC or databus fails, and you do not feel comfortable flying it solely on the HUD's (which give less info than the WDC screen), you can still elect to do the ascent on OC, taking extra care when using hypoxic mixes, but go back to CC at 6 meters, effectively running it as an oxygen rebreather. You can then use the HUD as indication you should inject O2 again, as you likely strive to maintain a PO2 of 1.5 or 1.6, which will cause the HUD to blink fast. If the PO2 drops to 1.3 (assuming the last selected setpoint was 1.3), the HUD's will go steady green, indication you to inject O2 again. Nice feature!
- If only the databus link fails, that will show up as frozen or asterisk PO2 data and dashes in the TempStik data (if fitted), but depth and time will still run. Check it, to isolate the error. It also means the WDC is still usable as an unlinked multi-gas constant PO2 and OC computer, as described above. Or at the very least as depthgauge + timer.
- If the last selected setpoint was 1.3, and the WDC or databus fails, the unit will NOT switch back to low setpoint at 3 meters of depth. So: stay deeper than 3 meters to avoid wasting O2. Realize also that, contrary to the Classic, the Vision will keep the solenoid continuously open if the setpoint is too low; the classic does this in 17-second intervals, followed by 6 seconds "rest".
- Once you get to the surface, do NOT close the O2 valves: in the stress of the situation (after all, you experienced an annoying failure), you might keep the CC mouthpiece in and get hypoxic. Better to lose all of your O2 slowly than go hypoxic! Better still: go to OC once on the surface. Once on the boat, unit taken off, close the O2 supply. You have to, as you will also not be able to switch the system off, as this is done by the WDC instructing the SPC's.... you'll have to remove the batteries.

A closer look at power management

Now that we've looked at the "roles" of the WDC and SPC's, let's look a little bit closer to the power management functionality of vision and the way it deals with low and flat battery situations. This is an important area, as the whole "power grid" of Vision is majorly different from that in the Classic Inspiration. In the Classic, B1 powers handset 1, and B2 powers handset 2, and that's it. In the Classic the handsets communicate with each other over a bus, to see if they are still alive, and the slave handset can take the master-function over from the current master if that one runs out of power, but that's it: there is no capability to use each other's battery power.

In the Vision however, each SPC is technically capable of accessing both batteries, and even use them simultaneously (in parallel) if needed. I'll describe this functionality below in three separate scenarios.

Scenario 1: B1 running low and flat.

Let's first explore the situation where B1 is running out and starts generating lower than wanted voltages:

- In a normal situation, where B1 and B2 are full batteries, B1 powers SPC1 (and its HUD), and B2 powers SPC2 (and its HUD). SPC1 is the primary controller (i.e. "Master").
- B1, being the master battery, also additionally powers "the big consumers": the solenoid, the WDC (think about backlighting...) and the TempStik.
- The WDC can see if any SPC is still present and alive by regularly sending "pings" over the databus to both SPC's. The SPC's on their turn respond to these "pings", indicating "I'm alive & kicking!".
- If B1 now gradually loses power and drops below 4.7 volt, SPC1, being the current primary/master, will start to draw also from B2 for driving the solenoid and the WDC.
- So now it effectively draws from both batteries at the same time, but each battery is powering different functions:
 - o B1 now only powers the "core functions" of SPC1,
 - o B1 has stopped powering the solenoid and WDC,
 - o B2 now does the "heavy lifting", powering solenoid and the WDC,
 - o B2 of course also still powers the "core functions" of SPC2,
 - o and SPC1 is still the master controller.
- If B1 drops further (to around 3.2 volt), SPC1 cannot function properly anymore, and will stop answering the "pings" from the WDC, effectively shutting itself logically (but not yet physically) down:
 - o Even below 3.2V, SPC1 is still alive, but it deliberately stops responding to the databus "pings" it gets sent from the WDC to see if it is still up & running. To the WDC, SPC1 now appears as having shut down.
 - o When the B1 voltage falls below 3.2V, SPC1 doesn't want to be considered "alive" anymore by the WDC because at such a low voltage it's analogue measurements (specifically PO2 and battery voltages) can no longer be relied upon.
 - o When B1 eventually reaches 2.7V, SPC1 will really shut down. This is called the "brown out" voltage.
- Once B1 is below 3.2V, and subsequently SPC1 stops answering pings from the WDC, the WDC will instruct SPC2 to take over the primary controller function. SPC2 was and is still driven by its own B2 battery. So, now B2 powers all functions: controller, solenoid and WDC.
- B1 and SPC1 are now effectively taken out of operation in the whole Vision setup.
- In this case of B1 not properly powering SPC1 anymore, SPC2 can of course still answer the "pings" from the WDC, as it is independently powered by B2.
 - o If you think about it, this is logical; it is a typical "chicken & egg"-problem: if B1 would also have been the sole power-supplier of SPC2, it would not be able to get promoted to primary controller, as it would be running on a flat battery....

It is easy to spot that SPC1 powered off, because its two HUD lights will go out.

So, in the above scenario, with a degrading and subsequently flat B1, SPC2 eventually ends up running the show, controlling solenoid, WDC and TempStik, and using B2 for its power.

Scenario 2: B1 and B2 both running low, but not flat.

In another scenario, with B1 gradually running low (but not flat), and next B2 also gradually running low (but not flat), power management eventually ends up putting both B1 and B2 in parallel mode, and draws power from them in parallel. So now:

- B1 still powers SPC1 (and its HUD),
- B2 still powers SPC2 (and its HUD), and
- both batteries power the WDC, solenoid and TempStik.

Let's explore this scenario a bit further. Let's assume that earlier on B1 ran low (below 4.7V), so B2 became the primary battery (powering WDC and solenoid), but SPC1 is still the primary (master) controller.

This could then be the sequence of events in practice:

1. As described above, we assume that B1 ran low (drops below 4.7V), issuing a battery warning;
2. On request of SPC1, B2 then starts helping out, for doing the "heavy lifting", while at the same time it also continues serving SPC2's "core" functions;
3. As B2 now starts to degrade faster than before (because of it having to do the heavy lifting), B1 in contrast has gotten an extension of life, because it is now relieved from having to do the heavy lifting work;
4. If B1 and B2 now both end up running low before B1 runs flat, the "parallel" scenario as described above kicks in: both batteries will now drive the WDC, solenoid and TempStik.

This situation with both batteries being put in parallel is *quite likely* to happen, because the "relieved of its duty" B1 battery now only has to supply around 3 mA for powering the SPC1 core functions and HUD, while "heavy lifting" B2 now has to deliver up to 450 mA for solenoid, backlighting etc. as well as the power for SPC2 core functions and HUD.

Freak scenario 3: B1 running low while the WDC encounters an unrecoverable error.

Let's assume the quite unlikely "freak" scenario of B1 running low, while simultaneously for whatever reason (e.g. a broken cable in the databus) the WDC can not monitor and instruct the SPC's anymore. This means it can subsequently not switch the primary controller function over from SPC1 to SPC2 anymore.

This is what will happen:

1. B1 runs low (drops below 4.7V);
2. B1 and SPC1 will sooner or later not be able to operate the solenoid anymore, drop below 3.2V, switch off and its HUD lights will go out;
3. SPC2 however still has the earlier described autonomous "80% safety net" function!
4. This means it will start actively driving the solenoid when the PO2 drops below 80% of the expected setpoint. So, the setpoint will typically be maintained at a PO2 of $0,8 * 1,3 = 1,01$ bar. Good enough to survive!

Realize though that in this scenario you should re-calculate your deco from this point on, based on a setpoint of 1,0 instead of 1,3.

To me however the loss of a WDC, an SPC (as seen by its HUD switching off) and not being able to monitor the actual PO2 anymore would typically mean "bail-out" time....

So as you can see, it is not just under manual control or in case of malfunction that there is a switch from SPC1 to SPC2 as master.

Final remarks

I don't consider a low battery a malfunction. I consider it a lazy diver or a "penny wise, pound foolish" diver, or bad luck with a dodgy battery, but that's an entirely other discussion.....

Given my experiences with the two most unreliable parts in ANY rebreather (oxygen cells & batteries), the way power management in the Vision works is pretty sophisticated.

What remains however is that by far the best advice is to avoid power management having to kick in anyway in the first place! This is easily done by:

- Renewing batteries in time. I swap them after some 15 hours, assuming their normal lifespan should be around a bit more hours.
- Using good brands of batteries. Never go "cheap" in a life-support system. If I assume a new good battery costs around 6 euro (4 pound) or less, and assuming you make 15 dives of an hour with it, it will set you back a "whopping" 40 eurocents or 30 pennies for a dive. Need I say more....?

Also remember, with 2 fresh batteries you will have about 30 hours of dive time: 15 hours from B1, followed by another 15 hours from B2. The scrubber is rated at 3 hours (although generally speaking you can push out some more, using the TempStik as gauge) and the O2 on board will "only" last 8-10 hours on a full 200 bar fill.

As such you will burn the scrubber, or run out of O2 in one dive way before just even one battery gives out. So on a task-loaded long and deep dive, starting out with at least a fresh battery in B1 is a really good idea.

The power management system is like an airbag in a car: it is good to know it is there, because it may save your life in an unplanned and/or erroneous situation, but do you REALLY want to test if it works....? Me, I hope to never having to find out if it really works....

Having said that though, it is good practice to once in a while (e.g. after a software upgrade) test if it still works o.k., by allowing B1 to run low. However, in my opinion you should choose a quiet and unchallenging dive to do that.

That's all. I hope this contributes to a better understanding of the Vision electronics, and effects of possible failure modes.

Ciao,

Tino.