

VPM - THE INNER WORKINGS

Marcin Kaluza (www.hlplanner.com)

November 2005

INTRODUCTION

The primary purpose of this text is to explain in layman's terms the inner workings of the Varying Permeability Model that played an important role in our understanding of decompression and laid the foundations for other decompression bubble models that followed. Texts available publicly, either on the internet or in print, focus primarily on the mathematical side of the algorithm ignoring in most cases the principles that govern its foundations.

This can leave the average diver with a black box, the inner workings of which are a total mystery. This also means that the vast majority of divers are susceptible to misinformation and have not got the knowledge required to make informed choices and decisions regarding their own safety. The following text tries to explain and illustrate the basics, without going too deeply into the mathematical side of things as most of the texts quoted in the reference section do just that, and sadly not much more. This paper contains an absolute minimum of formulas and they have been included only to allow readers interested in the underlying maths to be able to reproduce graphs presented in further sections. Those not interested in crunching numbers may safely ignore them.

Physicists may still argue which of the models available today is more accurate or "better", but from a diver's perspective the principles, not the intricacies, are what matters the most. From this point of view, the importance of works of Prof. David E. Yount cannot be underestimated.

It would be very difficult to compress the entire decompression theory into a single paper, and therefore the reader is expected to understand at least the principles of the neo-Haldanean models. Reference material can be found in the reference section of the HLPlanner [website](http://www.hlplanner.com).

M FOR M-VALUE

All decompression models used currently in technical diving use the same type of model to calculate tissue on and off-gassing: the body is divided into set of hypothetical "compartments" with varying on/off gassing rates known as halftimes. The term "halftime" indicates how long it would take for the tissue to halve or double its current contents (gas loading). If the tissue halftime is 4 minutes, and it is completely empty, to make it half full we'll need 4 minutes. We'll need another 4 minutes to fill half of the remaining half (i.e. a quarter) and so on. In practical terms a tissue is considered fully loaded within 5-6 halftimes which in the case of our 4 minute tissue is equal to 20-24 minutes and 97-98% of the possible load. Most of the models assume a varying number of compartments, with the most popular, Buhlmann's ZHL16, assuming 16 of them with halftimes of between 4 and 640 minutes.

Compartment based gas loading is common for all the models being used today, or at least the author does not know of any model being used in recreational and technical diving that would be different in this particular aspect. The number of "tissues" and their halftimes may vary; the concept however remains the same.

The purpose of the decompression model is to control the gas loading within each compartment in such a way that would prevent DCS in diver. The pressure of the gas dissolved within a compartment is often referred to as **tension**, and that is exactly what decompression models have tried to control since the beginning of the 20th century. In techno speak the upper limit of gas tension (for any one compartment) is known as its **M-Value**, with the “M” standing for maximum.

All neo-Haldanean models use simple set of linear equations to describe the M-Value which generally takes the form of:

$$M = M_0 + \Delta M \times d$$

where M_0 and ΔM are both tissue specific parameters and d denotes depth.

The above equation uses the so-called Workman form, as opposed to a slightly different approach adopted by Buhlmann. It is important to note that given the M-Value it is then easy to calculate allowable tissue supersaturation gradient P_{ss} , i.e. the difference between tissue tension and ambient pressure:

$$P_{ss} = M - d = M_0 + d \times (\Delta M - 1)$$

Fig. 1 illustrates allowed supersaturation gradients for the slowest and the fastest compartment in Buhlmann’s ZHL16C algorithm.




Fig. 1 ZHL16C Maximum supersaturation gradients vs. depth.

There are several important observations that can be deduced by looking at this graph that have very important implications when it comes to the shape of the decompression curve generated by a neo-Haldanean algorithm:

- The faster tissues (the ones that controls initial deep deco stops) are allowed to be supersaturated to a higher degree than their slower counterparts;
- Greater depths allow for higher supersaturation; and
- Allowed supersaturation is dependent only on depth.

These three characteristics of the neo-Haldanean models have serious implications when it comes to deco profiles generated: because fast compartments are allowed to be saturated to a higher degree and even higher at depth, the first deco stops controlled by these

compartments are relatively close to the surface as there is no need to stop at depth. Because slow compartments are allowed to be supersaturated to a lesser degree and they off-gas slowly, shallow stops are bound to be long.

Such an approach to decompression, as we know it today is suboptimal (to say the least) as over a period of years it has been observed that “deep stops” not strictly required by the model make the diver “feel better” after the dive.

DCS, BUBBLES & THEIR MECHANICS

As we already said, all current deco models use similar concept of “compartments” to calculate the gas loading during the dive. Where they differ is the way the maximum allowed tissue tension is calculated. In the case of neo-Haldanean models, a set of linear equations is used, whereas for VPM, or bubble models in general, the situation is far more complex, but the principles are relatively simple and we’ll tackle them one by one. Let’s start with the most important one.

DCS is caused by free gas released in our bodies ([4]). As long as the gas remains dissolved, it is not a problem, as soon as it forms bubbles, it can cause all kinds of trouble: obstruction to the blood flow, clots, mechanical damage to surrounding tissue, etc. which, depending on the extent of the damage, may result in fatigue, pain, paralysis or even death. One undisputable fact remains certain however: when the volume of the free gas in our bodies exceeds certain limits, we’re likely to suffer from DCS. The question is then what is the limit and how do we calculate it? Before we get there however we have to discuss the mechanics of the bubbles or at least the bit concerning their growth.

HOW BUBBLES GROW?

VPM assumes that our bodies are full of tiny micro bubbles called gas nuclei that subjected to supersaturation act as seeds for bigger, “proper” bubbles that cause all sorts of problems commonly known as DCS. The interesting thing is that upon decompression, only some of them will grow. To fully understand this paradox imagine the following: you take a set of differently sized balloons underwater, as soon as you start ascending to the surface, some of them grow due to the change in ambient pressure and gas expansion, but some of them... shrink. Yup, that’s the way bubbles behave. If we were to repeat this experiment on a micro scale using jelly samples as they were originally used by the inventors of VPM, we could observe that growth of the gas nuclei is dependent on two factors:

- nucleus size – only bubbles bigger than a certain size grow; and
- level of supersaturation – following the same dive, the number of growing bubbles will depend on the maximum supersaturation achieved during the decompression.

These two facts have massive implications for all bubble models: **by varying the allowed degree of tissue supersaturation, we can control the number of growing bubbles.** In principle if we wanted to keep the number of bubbles small, we have to keep the supersaturation levels low and vice versa. It is also very important to note the relationship between supersaturation gradient and bubble sizes and the so-called “Laplace’s formula” does just that. Given the supersaturation gradient P_{ss} we can calculate which bubbles will grow, or in other words it is possible to calculate **critical radius** (r_c) that separates growing nuclei from those that will collapse:

$$r_c = \frac{2 \times \gamma}{P_{ss}}$$

where P_{ss} denotes supersaturation gradient and γ is the surface tension of the surrounding fluid. Even more importantly we can reverse the question and ask: given the critical radius

what is the allowed gradient that will prevent smaller bubbles from growth? Being able to answer this question is fundamental as we'll see clearly in a moment.

IN SEARCH OF THE M-VALUE

In the good old days M-Values and deco models in general, were based on experiments. Both divers and researchers knew that our bodies can easily deal with certain amount of inert gas in the tissue. J.S. Haldane almost a century ago observed that you can keep a diver at 10m "forever" and bring him to the surface immediately without any ill effects. The same applied to a dive to 6 ATA and ascent to 3 ATA. From there Haldane derived an initial 2:1 ratio of tissue gas tension to ambient pressure that was later refined by Workman to be 1.58:1 due to only inert gases being responsible for the development of DCS. The model was refined later by the introduction of depth dependent gradients and different M-Values for different tissues. All in all, the limits were verified empirically and there was no real "science" behind them. And it all changed with the arrival of VPM.

THE VARYING PERMEABILITY MODEL

VPM assumes that there is indeed a certain volume of free gas that our bodies can deal with, and that this volume is related to the number of bubbles excited into growth during decompression, however instead of saying "the number of bubbles allowed to grow is 1,057" it uses a concept called distribution. Distribution describes the relationship between bubble size and number of bubbles present in the population:

$$N(r) = N_0 \times e^{-k \times r}$$

Where N_0 is the total number of bubbles, r is the radius and k is some arbitrary parameter. The general shape of the distribution curve is illustrated in Fig. 2.

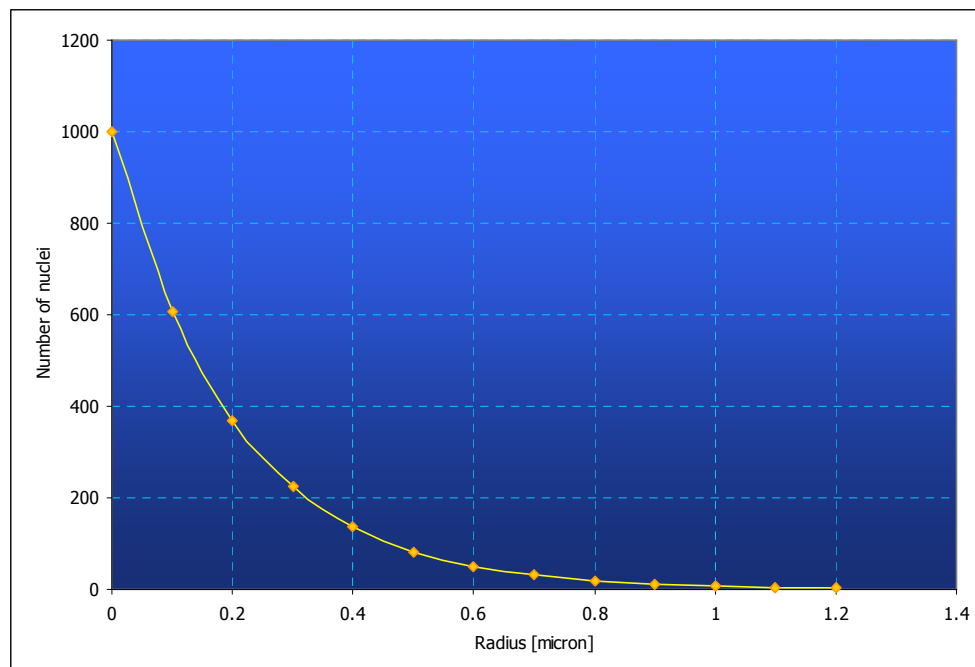


Fig. 2 Sample distribution of gas nuclei vs. size.

Given this hypothetical distribution we can say that there are roughly 7 bubbles bigger than $1\mu\text{m}$, 600 bubbles bigger than $0.1\mu\text{m}$ and 1,000 bubbles in total (all the bubbles bigger than 0). Therefore instead of saying "600 bubbles" we can say "all the bubbles bigger than $0.1\mu\text{m}$ ".

This way of expressing numbers of bubbles may seem convoluted at first but that's what physicists like to do and there is a very good reason for it.

Let's recap what we know by now:

- 1) DCS is caused by free gas released in the form of bubbles.
- 2) There is a certain volume of free gas that our bodies can deal with, and that volume is related to the number of bubbles allowed to grow during decompression.
- 3) We can express the number of bubbles by using distribution and bubble size.
- 4) From bubble mechanics (Laplace's formula) we know that we can calculate a supersaturation gradient (and hence a deco schedule) such that only bubbles bigger than a certain size will grow.
- 5) Combining 3 and 4 gives us a recipe for safe decompression, and that's the foundation of VPM (or any other bubble model for that matter).

To rephrase this lengthy explanation we can say simply that the primary job of VPM is to prevent excessive numbers of bubbles from growing and "excessive" translates to "no bubble **smaller** than X". In order to achieve this goal, the program has to calculate for each compartment a supersaturation gradient such that only bubbles **bigger** than X will grow. The magical number X is derived experimentally and is one of the fundamental parameters of VPM – it is called **initial critical radius** and two of these (one for Helium and one for Nitrogen) are maintained by any VPM based program.

If you reached this point and understood what was said so far you'll have no problems in understanding the rest of it. There are some important details still to be explained but from now on, everything becomes relatively simple.

WHY "VARYING PERMEABILITY"?

So far we only dealt with the decompression phase of the dive, and we already know how the supersaturation gradients are calculated. What we have to deal with now is the "compression" or descent phase. Before we get there however, there is one more of VPM's tricks that has to be mentioned: the **ordering hypothesis**. The ordering hypothesis simply states that if bubble A is smaller than bubble B, this relationship will be true even if we subject the bubbles to any compression-decompression cycle. This way if we're given the decompression requirement that states "prevent bubbles smaller than X from growth", although all bubbles including X change their size during the dive, the number of bubbles smaller than X remains the same, even though the value of X changes. And that is another foundation of VPM.

Considering a simple dive scenario, any VPM based program has to track what happens with gas nuclei in each compartment. To simplify the job it is enough to track bubbles of a size equal to the initial critical radius, as we know that anything smaller than the critical radius will collapse and anything bigger will grow. Bubble growth or shrinkage is driven by the difference between the tissue tension and the ambient pressure - the pressure gradient. Because tissue tensions depend among other things on their halftimes, critical radii are tracked individually for each compartment. On ascent, given the critical radii, supersaturation gradients are calculated for each compartment and that is what drives the deco phase. However the compression of bubbles during the descent phase has serious implications for the model and that's why it's worthwhile to look at it.

As we said earlier, compression of bubbles will determine their size at the end of the dive. VPM defines two different gradient ranges in which bubbles behave differently: the permeable and impermeable range. According to VPM, bubbles are surrounded by surfactant molecules forming a "skin". In normal conditions this skin is permeable, i.e. gas can freely diffuse between the bubble and the surrounding tissue. When compressed to a certain degree however, the skin becomes impermeable and bubbles behave like air inflated balloons, i.e. their compression is driven by Boyle's law and ambient pressure. This difference in

behaviour (permeable/impermeable) gave VPM its name. As you can imagine, depending on the way one descends to a given target depth, bubbles will be compressed to a greater or lesser degree, depending on the descent rate. This is one of the reasons why VPM does not support “instantaneous” descents. The descent rate is important because bubbles are crushed not by the ambient pressure but by the difference between ambient pressure and tissue tension (inward gradient). If the descent is slow, the tissue tension at the end of it will be higher than in the case of a fast descent when tissues will have no time to take up gas and the inward gradient, or crushing pressure, will be higher. Fig. 3 illustrates this phenomenon.




Fig. 3 Crushed nuclei radius vs. descent rate for a 5 min compartment (40msw air dive, 0.55 micron initial radius)

Since smaller nuclei require higher supersaturation gradients to excite them into growth, from Fig. 3 we can deduce that initial supersaturation gradients will be higher for higher descent rates.




Fig. 4 Supersaturation gradient vs. descent rate for a 5 min compartment (40msw air dive)

In other words faster descents mean (potentially) shorter deco as illustrated in Fig. 4, and conversely, slower descent rates may cause longer decompression times. This may be a major inconvenience in cases where real descent rate differs from the plan. A slower descent however usually means less time on the bottom (e.g. given a fixed dive runtime before ascending) and therefore lower tissue saturation. The net effect is that although the allowed supersaturation gradient is lower, the tissue contains less gas and differences in total deco times are usually minor. Apart from this, slower compartments that control shallow stops are slow enough not to be dependent on variations of descent rate.

So far we have discussed the behaviour of crushed nuclei in the permeable range. For the sake of completeness it is worthwhile to look at the impermeable range as well. In this range as said before nuclei are subjected to Boyle-like compression and therefore reduction of their size is slower resulting in bigger gas nuclei at the end of the descent phase. Fig. 5 illustrates this phenomenon: first series (Always permeable) has been calculated assuming that gas nuclei are always permeable, while the second one was calculated assuming that once the inward (crushing) gradient reaches 37msw gas nuclei become impermeable.




Fig. 5 Gas nuclei compression during descent (5min compartment, 20msw/min air dive) for both “always permeable” and “permeable/impermeable” skins.

From Fig. 5 above we can clearly see the difference in radii. As we discussed earlier, supersaturation gradient during decompression is based on “bubble” size and Fig. 6 illustrates initial supersaturation gradients calculated for the same example.




Fig. 6 Initial supersaturation gradient for “always permeable” and “permeable/impermeable” gas nuclei.

As we can see from Fig. 6, the introduction of impermeability reduces the allowed supersaturation gradient.

We may ask now what are the practical implications of all this? What is important is that bubbles get compressed during descent. The faster the compression the better as the gas nuclei get reduced in size and will require higher supersaturation to excite them into growth, which again in practical terms means that it will be more difficult to develop “bubble trouble”.

Finally it is worth noting that unlike neo-Haldanean models, VPM’s supersaturation gradients (and thus the M-Values) are dependent on initial bubble compression. This dependency may not be significant in the case of “normal” dives, but it is important to keep it in mind when planning multilevel dives. Executing “deep” portion of a multilevel dive first will cause greater initial compression of nuclei and therefore will either shorten the decompression or make it safer when its time remains constant.

In case of VPM (and derived models) the parameter that controls the onset of impermeability is called (what else) “gradient for onset of impermeability”. Reducing its value leads to more conservative deco schedules for deep dives, although the differences are not great as slow compartments controlling shallow (and lengthy) stops are pretty much invariant to changes of this parameter.

The relationship between crushing pressure applied to gas nuclei and allowable supersaturation gradients postulated by VPM has many practical implications. First of all it explains to some extent Haldane’s pressure ratio theory. Since nuclei are crushed to a higher extent during deep dives, initial allowable supersaturations are higher as well and therefore a diver can ascend to “relatively” shallow depth after a deep dive ([2]).

Another interesting observation is that in VPM, unlike in neo-Haldanean models, allowable supersaturation levels for “fast” tissues are lower than for the “slow” ones. This has massive implications when it comes to the decompression curve: in such a case, the first deco stops are bound to be deep and shallow stops will be relatively short. What had to be done using “fudge factors” in the case of the Buhlmann algorithm, in VPM comes as standard.

This also explains why deep trimix dives have shorter deco schedules when using VPM (or VPM-B for that matter) as compared with Buhlmann [modified] with gradient factors.

Compression of the gas nuclei during descent causes relatively high allowable supersaturation during decompression and thus the lengthy shallow stops of ZHL16 are shorter in VPM.

CRITICAL VOLUME HYPOTHESIS

Interestingly enough, deco schedules calculated using the method described so far tend to be overly conservative for short and shallow dives. To overcome this problem Prof Yount adopted Critical Volume Hypothesis developed by Hempleman and Hennessy. The hypothesis states that a primary decompression-limiting criterion is the product of the gradient and the decompression time. In mathematical terms it is expressed as an integral of supersaturation gradient over time which in practical terms means that “short” decompressions can be executed more aggressively.

To illustrate how the Critical Volume Algorithm (CVA) works let’s have a look at Fig. 7 which illustrates two different deco profiles for the same dive calculated with and without the CVA. Decompression programs using the CVA start with an initial conservative decompression schedule and using an iterative process calculate a series of more aggressive schedules until they do not get any shorter, or to be painfully precise, until the “phase volume” converges.




Fig. 7 Deco profiles for 27min at 30msw using EAN35

In most of the VPM-based programs you can observe the differences in deco schedules calculated with and without the CVA by simply turning it on or off. Whilst the Critical Volume Algorithm influences mainly schedules of (relatively) short and/or shallow dives, there are circumstances however when it is relevant in tech diving as well. When calculating deco schedules involving one or more deco mixes, loss of a mix will cause a longer schedule and this in turn may lead (due to CVA) to lower allowed supersaturation gradients. Lower supersaturation gradients in turn may cause the decompression to start deeper and the deep stops to be longer. This is often perceived as major inconvenience when calculating bailout tables: a diver that discovers at 6m/20ft that his O₂ cylinder is not available (for whatever reason) is likely to also discover that in this case he should have started his decompression 3m/10ft deeper and he has already “cut” some time from his schedule. Luckily in the majority

of cases these differences are only minor, and what to do in such circumstances can be left to the diver's discretion.

VPM & REPETITIVE DIVING

All neo-Haldanean decompression models trace residual nitrogen (or generally inert gases) between the dives and based on the surface interval calculate initial tissue tension for repetitive dives. Generally these tensions will be higher than tensions in non-diving individuals and this will cause deco profiles of the repetitive dive(s) to be more conservative compared with the first dive in the series.

In the case of VPM however we have to deal with one extra decompression parameter: the size of the gas nuclei. The whole concept of VPM is to control the growth of gas bubbles so theoretically only those bigger than critical radius should grow. Unfortunately the CVA, as described above, will often increase allowable gradients calculated using bubble growth criteria. This leads to more bubbles being excited into growth and therefore initial critical radius for repetitive dives has to be recalculated in cases when application of the CVA leads to higher supersaturation gradients. In practical terms it means that VPM based schedules for repetitive dives will be more conservative not only because of residual inert gases, but also because of a reduction in initial critical radii.

It is worth noting here that the physics of repetitive dives is one of the greatest mysteries of diving. As Wienke noted in [3] some level of adaptation has been observed in caisson workers over time: those who just started the job were more susceptible to DCS. This phenomenon may be related to gas nuclei being either crushed or eliminated as bubbles during decompression from the bodies of those working for a long time. Subsequently, following exposures produce fewer bubbles and therefore total volume of free gas related to onset of DCS symptoms is lower.

On the other hand in recreational diving it is commonly known that divers are more susceptible to DCS while participating in repetitive, multi-day diving. Obvious difference between the two is the length of the surface interval and number of hyperbaric exposures per day.

In the end it is currently impossible to reliably predict the behaviour of the population of gas nuclei after decompression and therefore calculation of safe deco schedules for repetitive dives remains a major challenge.

VPM WITH BOYLE'S LAW COMPENSATION (VPM-B)

VPM as described so far revolutionised diving: if not directly by providing a physical model for a variety of phenomena observed empirically, it was also the foundation for development of other models (RGBM).

The deco schedules produced by VPM were however perceived as aggressive and on some occasions led to DCS. For this reason Erik C. Baker in 2002 modified original VPM and introduced a modification to the algorithm known today as VPM-B. The general idea behind this modification was to reduce allowable deco gradients towards the surface. Erik could have used his gradient factor (GF) approach but instead he decided to follow a "scientific" route of a reduction of gradients by "expansion" of gas nuclei according to Boyle's law.




Fig. 8 Expansion of gas nuclei during decompression (5min compartment, 40msw air dive)

The general idea was that gas nuclei, similarly with air-filled balloons, will expand when the ambient pressure (depth) decreases. Since supersaturation gradients are related to the size of these nuclei, they will decrease as well and schedules will generally require more hang time in the shallows. The following graphs illustrate both reduction in the size of gas nuclei (Fig. 8), as well as corresponding reduction in supersaturation gradient (Fig. 9).




Fig. 9 Reduction of allowed supersaturation gradient during decompression (5min compartment, 40msw air dive)

SUMMARY

Currently VPM-derived VPM-B is one of the most popular decompression algorithms; partially due to the availability of low cost software and partially due to the fact that it produces reasonable schedules. But the real importance of VPM lies not in its popularity, but in the fact that it finally provided a model for many of the phenomena taking place during decompression. The model may be inaccurate or even inadequate as postulated by some but it is undoubtedly a milestone in the history of decompression modelling.

VPM-B PARAMETERS

Both VPM and VPM-B are driven by handful of parameters described briefly in the following table:

Table 1 VPM(B) parameters

Parameter name	Description	Increase of the value causes...
Initial critical radii	Initial critical radii control bubble size and thus allowable tissue supersaturation levels. Since there are two of them (one for helium and one for nitrogen) it is possible to allow higher supersaturations for helium rich mixes.	Longer deco
Critical volume parameter lambda	Controls "Critical Volume Algorithm". In essence it controls the "severity" of the decompression required to start lowering of the supersaturation gradients.	Longer deco
Pressure of other gases	VPM takes into account presence of metabolic gases (CO ₂ , water vapour, etc.) that contribute to bubble growth. This parameter (available in some programs) controls the content of metabolic gases in venous blood.	Longer deco
Regeneration constant	Controls how quickly crushed nuclei restore their original size. Pretty much irrelevant for anything other than several days long saturation dives. In case of repetitive dives controls how quickly nuclei restore their sizes.	Shorter deco
Surface tensions (gamma and gamma c)	Control behaviour of gas nuclei under pressure (contraction and growth).	Longer deco
Gradient for onset of impermeability	Controls "depth" where skins of gas nuclei become impermeable. Influences mainly deep dives.	Shorter deco

ACKNOWLEDGEMENTS

I would like to thank Joe Hesketh and Joao Delgado from the HLPlanner Beta Group for their invaluable help and comments that helped to shape this documents.

REFERENCES

- [1] E.C. Maiken (1995) [Bubble Decompression Strategies](#)
- [2] D.E. Yount (1988) Theoretical Considerations of Safe Decompression. In Y. Lin, A. Niu eds. *Hyperbaric Medicine and Physiology*. San Pedro, CA: Best Publishing Co. 1988
- [3] B. Wienke (2001) *Technical Diving in Depth*. Flagstaff, AZ: Best Publishing Co.
- [4] C. Edmonds, C. Lowry, J. Pennefather, R. Walker (2005) *Diving and Subaquatic Medicine* 4th edtn. London: Hodder Arnold.
- [5] E.B. Baker (2000) [Derivation with Explanation of the VPM Dynamic Critical Volume Algorithm of Yount & Hoffman \(1986\)](#).
- [6] D.E. Yount, E.C. Maiken, E.B. Baker (1999) [Implications of the Varying Permeability Model for Reverse Dive Profiles](#),