

The Most Misunderstood Fact behind Dissolved Oxygen Sensors (It's the Pressure, not the Concentration)

Mark Spencer, President of Water Analytics

A dissolved oxygen sensor ought to be simple to understand. Whether it is membrane or optically based, it gives a signal that is proportional to the concentration of oxygen in water. Like any other sensor, one calibrates it by immersing it in a standard for which the dissolved oxygen concentration is known. Are we in luck or what? We don't have to buy expensive and degradable standards. Ordinary, pure water works just fine. Just stick an air stone in or a really good mixer to make sure the water is saturated with air. We know that pure water at 25 °C and at sea level will hold 8.3 mg/l (or ppm) of D.O.

But wait... it gets even better. You don't even need water with an air stone. You can hold the sensor in air and get the same result. How can that be? Air is comprised of 210,000 ppm of oxygen. How can 210,000 ppm of oxygen give us the same reading as a sample of water containing 8.3 ppm of dissolved oxygen?

If you know the answer then go do something fun. You definitely earned it. But, if you're like me a few years ago, then you're scratching your head. It's one of those maddening paradoxes, for which I've never read an adequate explanation. So I wrote this little white paper. I hope it helps.



Most analytical sensors measure concentrations of chemical species. Whether they're complex organic compounds measured by a gas chromatograph or tiny protons measured by a pH probe, the signal is proportional to the concentration.

However there exist a class of chemicals that are measured by sensors for which this is simply not true. This class of species is dissolved gases, i.e. gases that persist in the atmosphere and dissolve in water. They behave different from dissolved solids, such as a salt, in one very important way: A dissolved gas is in equilibrium with the atmosphere above it.

The most important example is, of course, oxygen. Any sample of water, from a glass of water to the Pacific Ocean, will have a dissolved oxygen concentration in the sample that is in equilibrium with the oxygen in the air. (We can make similar arguments to free chlorine except that it differs in two important ways: Cl_2 reacts in water to form HOCl and, because it is added by humans, it's concentration in water is always out of equilibrium with any Cl_2 in the atmosphere—hopefully negligible).

A stable concentration of oxygen, nitrogen and other gaseous compounds persists in our atmosphere because the earth's gravity keeps it there. All of these gases exert pressure on any surface they border, such as the earth's surface or the walls of a flask.

Eighteenth century British scientist John Dalton first elaborated in the law that bears his name a fact that seems obvious: The fractional pressure of oxygen alone is equal to its

fractional composition. In other words, the fractional composition of oxygen in air is 21% and the total pressure of air at sea level is 1 atmosphere (atm) so the pressure from oxygen alone is 0.21 atm. Daltons' Law reflects the mostly-true-truism that gases behave independently of each other. We call the fractional pressure of a gas like oxygen its **partial pressure**.

The partial pressure of oxygen forces the gas into solution where it also exerts a pressure—albeit in a watery medium that is much denser and more viscous than air. The pressure of the oxygen dissolved in water has to be equal to the pressure of oxygen in air. So the partial pressure of oxygen in water must also be 0.21 atm. This is the mental hurdle we must clear. To acknowledge that oxygen exerts as much pressure dissolved in water as it does in the air above it just sounds counter-intuitive. But consider this: If the partial pressure of oxygen were greater in air than in water then the imbalance would force more oxygen into the water—and vice versa. Equilibrium demands that they be equal.

To understand why sensors respond to oxygen's partial pressure instead of concentration think small—really small. So small, in fact, that you can magically observe oxygen molecules zipping around water molecules and striking the surface of membrane of an electrochemical (Clark or Galvanic) or the luminescent cap of an optical sensor. The signal of the sensor is proportional to the number of oxygen molecules striking the surface and that number is proportional to pressure—not concentration.

So a D.O. sensor measures partial pressure. It knows what the pressure should be in an O₂-saturated sample since that is equal to the ambient pressure, i.e. 0.21 atm. It divides the *measured* partial pressure by the expected 0.21 atm to give the percent saturation. We users like units of %-saturation because they make more sense to us than units such as atm, Pascal, torr, bar or inches of mercury.

So how does the nebulous parameter of partial pressure translate to dissolved oxygen concentration, stated in units of mg/l or ppm? Another British scientist, William Henry answered this question in 1803. The two are simply proportional and the proportionality constant is appropriately called Henry's constant. So the DO analyzer basically does the following:

1. The sensor measures the partial pressure of dissolved oxygen.
2. The analyzer divides the measured partial pressure by the partial pressure of oxygen at 100% saturation (0.21 mg/l at sea level) to get the %-saturation.
3. The analyzer multiplies the partial pressure by Henry's constant to get the absolute concentration in mg/l or ppm.

So if both mg/l and % saturation are derived from the measured partial pressure why do some say that %-saturation is the natural output variable of the D.O. sensor and the absolute concentration is a derived variable? The answer is I lied. Henry's constant isn't constant at all. Some call it Henry's coefficient. It depends on both the temperature and salinity of the water. That dependence is complicated. The calculation of the %-saturation value depends on the atmospheric pressure, which is a very easy number to derive. Unless you are working on Pike's Peak, it's very close to 1 atm. Most analyzers do allow the user to input the ambient pressure.

Texts that state that D.O. sensors measure %-saturation and convert that value to a concentration value are not telling the whole truth. Both % saturation and concentration

are derived units but the computational pathway from partial pressure to the former is a lot simpler than that to the latter. The figure below illustrates this point.

Higher temperatures and dissolved solids result in lower D.O. concentrations but they don't change partial pressures. Think about it: Two samples of water that are fully saturated are at two different temperatures. We know the colder sample holds more oxygen (more mg/l). Yet the partial pressure and the %-saturation of the D.O. in the hot sample is exactly the same as those values in the cold one. Or two samples of water are fully saturated with oxygen but one is pure water and the other is brine—same partial pressure and same %-saturation but, again, the salt water sample has much less D.O.

I have neglected one minor variable that influences the DO reading and that is the relative humidity in the air. At 100% relative humidity at 25 °C the air is about 2% water vapor. That means less partial pressure of oxygen (and other gases) and both lower %-saturation and mg/l values. This is the reason why you will see D.O. saturation levels usually stated as 8.3 mg/l but sometimes as 8.5. The prevailing convention is that saturation levels are stated for air that is 100% saturated. For this reason, calibration of a sensor in air ought to be done in either a closed container with liquid water or with a wet cloth wrapped around the sensor.

The most prevalent use of D.O. measurement is in setting its level in the aerobic and anaerobic/anoxic environments of a wastewater treatment plant. Wastewater operators don't give a you-know-what about partial pressures or %-saturation. They care about mg/l. They know, for instance, that their aerobic tank, where nitrification occurs, should have a D.O. level of at least 1 mg/l. So the conversion from pressure units to concentration units via Henry's Law is of paramount importance. Knowing how that mg/l value came about is vital for troubleshooting a sensor that gives readings that arouse suspicion.

Perhaps the greatest benefit of understanding why a D.O. sensor responds to partial pressure and not concentration is to be able to answer the burning question I stated in the introduction—why one calibrates a sensor simply by holding it in air. The difference in concentration between D.O. in water and oxygen in air is 8 ppm vs. 210,000 ppm, or 26,000 to 1! Clearly holding the sensor in air would give a hugely distorted calibration measurement if it responded to concentration instead of partial pressure. Now you know the answer. It's the pressure.

Mark Spencer is the President of Water Analytics. He holds a PhD in physical chemistry from M.I.T. and purchased AquaMetrix in 2009 to have a lot of fun. His book explaining how water quality instrumentation works will be published by AWWA sometime in 2016—earlier if his editor cracks down on him.

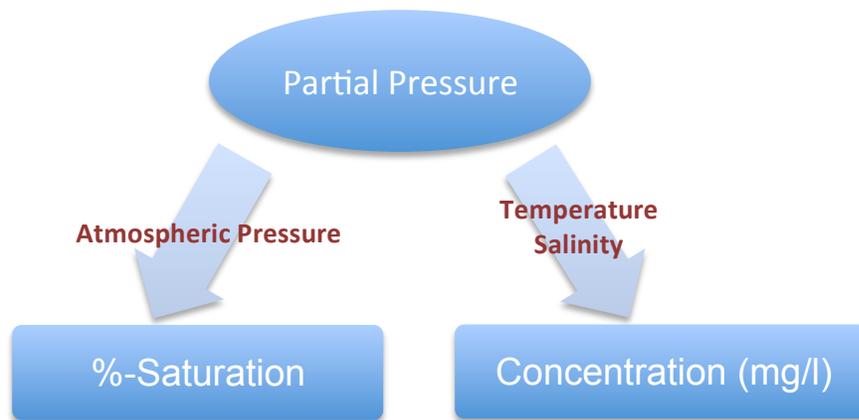


Figure 1 –The D.O. Sensor measures partial pressure. The analyzer converts that measurement to %-saturation simply by dividing the measurement by the atmospheric pressure. The analyzer converts partial pressure by a much more complicated relationship involving the temperature and salinity.