

Shining Light on Trace Impurities in Gaseous Xenon: Commissioning of a Spectroscopy-Based Water Monitoring System for Xenoscope



Author

Kevin Philip Nauer

supervised by

Prof. Dr. Laura Baudis

Dr. Florian Jörg

Department of Physics
University of Zurich

August, 2026

Abstract

Large-scale liquid-xenon (LXe) detectors have become highly sensitive instruments for rare-event searches addressing fundamental questions in modern physics, including the particle nature of dark matter and the possible existence of neutrinoless double-beta decay. A central part of their detection principle is the collection of xenon scintillation light in the vacuum-ultraviolet (VUV) regime. Water is an important impurity in LXe because it absorbs VUV light and can thereby reduce the light-collection efficiency of a detector.

This thesis describes the integration and commissioning of a cavity-ring-down-spectroscopy-based HALO moisture analyzer in the Xenoscope gas-handling system at the University of Zurich. Xenoscope is a liquid-xenon detector demonstrator developed to test technologies for future large-scale xenon time-projection chambers. The system was designed to measure trace concentrations of H_2O in gaseous xenon while allowing for the sampled xenon to be recovered.

The conditioning behaviour of the sampling system was first investigated using nitrogen and argon. During flushing, the measured water concentrations decreased from the parts-per-million (ppm) range to the parts-per-billion (ppb) range, reaching minimum values of 446 ppb in nitrogen and 259 ppb in argon. The time evolution of the concentration was described using an effective model accounting for rapid gas exchange and the slower release of water from the internal surfaces of the sampling lines and gas-handling system.

Four gaseous-xenon measurements were subsequently performed. Three measurements obtained under stable operating conditions yielded mutually consistent asymptotic water concentrations between approximately 21 and 23 ppb. The mean fitted asymptotic concentration of these three measurements was $\bar{c}_{\infty, GXe} = (22.2 \pm 1.0)$ ppb, where the overbar denotes the mean over the three stable gaseous-xenon measurements. A fourth measurement, affected by a malfunction of the main Xenoscope xenon circulation pump, yielded a significantly higher asymptotic concentration of (109 ± 2) ppb. The resulting unstable sampling-flow conditions prevented the concentration from reaching the low-ppb range observed in the first three measurements.

At a xenon temperature of 178.7 K and a pressure of 2.08 bar, the representative HALO result corresponds to a gas-phase water partial pressure of (4.62 ± 0.21) mPa. This agrees within approximately 10 % with the water-ice vapour pressure of 4.21 mPa calculated using the Murphy–Koop parametrization, which describes the equilibrium vapour pressure of water over ice as a function of temperature. At the stated temperature and pressure, this vapour pressure corresponds to a saturated gas-phase water concentration of 20.3 ppb. This agreement provides an independent consistency check of the measurement and demonstrates that the commissioned HALO system can determine physically meaningful water concentrations in gaseous xenon at the low-ppb level.

Acknowledgments

I would like to thank Prof. Dr. Laura Baudis for giving me the opportunity to carry out my Bachelor thesis in her group and for providing the possibility to work on the Xenoscope facility. I am especially grateful to Dr. Florian Jörg for his supervision, guidance, and continuous support throughout this project. His explanations, feedback, and help in the laboratory were essential for the successful integration and operation of the HALO moisture analyzer.

I would also like to thank the members of the Astroparticle Physics group at the University of Zurich for the welcoming working environment, as well as the physics workshop for welding our components and prioritizing some of our orders.

Finally, I would like to thank my family and friends for their encouragement and support throughout my studies and during the writing of this thesis.

AI Declaration

During the preparation of this thesis, AI-based tools were used for language improvement, LaTeX support, Python code assistance, and discussion of data analysis methods. All texts created with the help of AI were checked by me. I take full responsibility for the text and for all scientific content, interpretations, calculations, and final decisions.

Contents

1	Introduction to dark matter	1
1.1	Early evidence for dark matter: The Coma galaxy cluster	1
1.2	Present evidence for dark matter	2
1.3	Dark matter candidates	3
1.4	Hunt for dark matter using direct detection	4
2	Xenoscope and liquid-xenon detector principles	5
2.1	Purpose of Xenoscope	5
2.2	Operating principle of a dual-phase liquid-xenon TPC	5
2.3	Production of VUV scintillation light in liquid xenon	7
2.4	Influence of water impurities on VUV-light transport	8
2.5	Xenoscope detector and xenon-handling system	10
2.5.1	Xenon circulation and purification	11
3	HALO-Monitor	13
3.1	Cavity ring-down spectroscopy for trace impurity detection	14
3.2	Internals of HALO-Monitor	16
3.3	HALO Piping and Integration	17
3.3.1	Pipe Installation	17
3.3.2	Leak Checking	18
3.3.3	Panel integration	20
3.4	Measurement Procedure	21
4	Data Analysis and Discussion	23
4.1	Analysis of water content in N ₂ and argon	23
4.1.1	Qualitative behaviour of the flushing measurements	23
4.1.2	Double-exponential description of the decline	26
4.1.3	Implications for the xenon measurements	28
4.2	Analysis of water content in GXe	28
4.2.1	Overview of the xenon measurements	28
4.2.2	Double-exponential dry-down model	29
4.2.3	Fit and residual diagnostics for GXe runs 1, 2 and 3	30
4.2.4	Special behaviour of GXe run 4	33
4.2.5	Representative xenon concentration	34
4.3	Measurement of the water-ice vapour pressure at LXe temperature	34
4.4	Discussion and Outlook	37
5	Conclusion	39
6	Appendix	41

List of Figures

1	Working principle of a dual-phase TPC	6
2	Influence of water impurities on liquid-xenon scintillation light.	9
3	Xenoscope facility	10
4	Xenoscope purification loop	12
5	Used HALO-monitor	13
6	Working principle of the HALO moisture analyzer	14
7	Absorption feature used in CRDS	15
8	Flow diagram HALO-monitor	16
9	Flow diagram of HALO configuration	18
10	Initial HALO panel	19
11	Final HALO panel	20
12	Assembly of expansion vessel	22
13	Double-exponential fit and residual for nitrogen	24
14	Double-exponential fit and residuals for argon	25
15	Double-exponential fit and residuals for GXe run 1	30
16	Double-exponential fit and residuals for GXe run 2	31
17	Double-exponential fit and residuals for GXe run 3	32
18	Double-exponential fit and residuals for GXe run 4	33
19	Comparison of the measured and predicted water-ice vapour pressure	36

List of Tables

1	Summary of nitrogen and argon measurements	26
2	Double-exponential fits for nitrogen and argon	27
3	Summary of xenon measurements	29
4	Double-exponential fits of xenon measurements	29

1 Introduction to dark matter

The nature of dark matter remains one of the most fundamental open questions in modern physics. A wide range of astrophysical observations, including galaxy rotation curves, gravitational lensing, and the large-scale structure of the Universe, provide compelling evidence for the existence of an invisible form of matter that does not interact electromagnetically [1]. The dark matter hypothesis has become a cornerstone of modern cosmology, motivating extensive experimental efforts to detect it directly.

1.1 Early evidence for dark matter: The Coma galaxy cluster

The first indication of such missing mass dates back to 1933, when Fritz Zwicky studied the velocity dispersion of galaxies in the Coma cluster and concluded that the visible matter alone was insufficient to account for the observed dynamics [2]. In his work on the redshift of extragalactic nebulae, Zwicky compared the observed velocities of individual galaxies within the Coma cluster with the mass that could be estimated from the visible, luminous matter. Assuming that the cluster was in a stationary, gravitationally bound state, Zwicky applied the virial theorem. The virial theorem describes the relationship between the average kinetic energy $\langle T \rangle$ and the average potential energy $\langle U \rangle$ of a bound system at equilibrium:

$$\langle T \rangle = -\frac{1}{2}\langle U \rangle. \quad (1)$$

Since the gravitational potential energy is negative, the virial theorem relates the characteristic velocities of the galaxies to the gravitational binding energy of the cluster. For a cluster with total mass M , characteristic radius R , and velocity dispersion σ_v , this relation can be written approximately as

$$\sigma_v^2 \sim \frac{GM}{R}, \quad (2)$$

where the exact numerical factor depends on the mass distribution within the cluster. The velocity dispersion expected from the visible matter can therefore be estimated by inserting the luminous mass into (2), whereas the observed velocity dispersion is determined from the Doppler shifts of the galaxies' spectral lines.

This revealed that the velocities expected from the visible matter were significantly smaller than the actual measured velocity dispersions of the galaxies in the Coma Cluster [2]. Zwicky found that velocity differences of at least 1500 to 2000 km s^{-1} occur in the Coma Cluster. However, based on the luminous mass of the cluster assumed at the time, this would result in a characteristic velocity dispersion of only about 80 km s^{-1} . To account for the observed Doppler velocities of about 1000 km s^{-1} or more, the average density of the Coma Cluster would have to be, according to Zwicky's estimate, at least several hundred times greater than the density derived only from the luminous matter [2].

From today's perspective, it should be noted that a large fraction of the baryonic mass in galaxy clusters is contained in the intergalactic medium, which is mainly observable through

its X-ray emission and was therefore not accessible to Zwicky at the time. Nevertheless, even when this additional baryonic component is taken into account, the total visible and baryonic matter is still not sufficient to explain the observed dynamics of galaxy clusters. From this, he concluded that a large proportion of the matter in the cluster must be invisible or non-luminous. He referred to this non-luminous component as *dunkle Materie* [2], i.e. dark matter. Zwicky's analysis of the Coma galaxy cluster is therefore considered one of the first astrophysical pieces of evidence that visible matter alone is insufficient to explain the dynamics of large cosmic structures.

1.2 Present evidence for dark matter

Today, the evidence for dark matter is not based on a single observation, but on several independent measurements covering scales from individual galaxies to the entire observable Universe. At galactic scales, one of the most important observations is the approximately flat shape of galaxy rotation curves. For a test mass moving on a circular orbit at a distance r from the centre of a galaxy, the orbital velocity is given by

$$v(r) = \sqrt{\frac{GM(< r)}{r}}, \quad (3)$$

where $M(< r)$ is the mass enclosed within the radius r . If most of the mass were concentrated in the visible stellar disc, $M(< r)$ would approach a constant value outside the luminous region and the velocity would decrease approximately as $v(r) \propto r^{-1/2}$. Instead, observed rotation curves generally remain approximately flat at large radii. This indicates that the enclosed mass continues to increase beyond the visible disc and is commonly interpreted as evidence for extended dark-matter haloes. Measurements of gravitational lensing and the dynamics of galaxy clusters provide further evidence that the total gravitating mass is considerably larger than the mass contained in stars and gas alone [1].

Independent evidence is obtained from measurements of the cosmic microwave background (CMB). The positions and relative amplitudes of the acoustic peaks in the CMB temperature and polarization spectra depend differently on the densities of baryonic matter and cold dark matter. Within the standard cosmological model, the final Planck analysis obtained

$$\Omega_c h^2 = 0.120 \pm 0.001, \quad \Omega_b h^2 = 0.0224 \pm 0.0001, \quad (4)$$

where Ω_c and Ω_b denote the present cold-dark-matter and baryonic-matter density parameters, respectively, and h is the reduced Hubble constant [3]. These results imply that the cosmological dark-matter density is more than five times larger than the baryonic-matter density.

Big Bang nucleosynthesis (BBN) provides an independent constraint on the amount of baryonic matter. The predicted primordial abundances of light elements, particularly deuterium and helium, depend on the baryon-to-photon ratio and therefore on the cosmological baryon

density. Measurements of primordial deuterium give a baryon density consistent with that inferred from the CMB [4]. However, this baryonic density is far below the total matter density required by the CMB, galaxy dynamics, gravitational lensing, and the observed large-scale structure of the Universe. The missing matter therefore cannot consist solely of ordinary baryonic matter.

Dark matter also plays an important role in the formation of cosmic structure. Before recombination, baryonic matter was coupled to photons and could not efficiently collapse into gravitational potential wells. A non-baryonic matter component that did not participate in this electromagnetic interaction could begin forming structures earlier and provide the gravitational seeds from which galaxies and galaxy clusters subsequently developed [1, 3]. The agreement between these independent observations provides strong evidence for a dominant non-luminous and non-baryonic matter component.

1.3 Dark matter candidates

Although the gravitational evidence for dark matter is strong, its particle nature remains unknown. Since the standard model does not provide a reliable non-relativistic dark matter candidate¹, the dark matter problem is commonly interpreted as an indication of physics beyond the standard model [1]. A large number of non-baryonic candidates has therefore been proposed. These include axions, sterile neutrinos, supersymmetric particles such as the lightest neutralino, and Kaluza-Klein particles arising in models with universal extra dimensions [1, 5]. The proposed candidates span a very wide range of masses, from very light axions in the μeV – meV range to weak-scale particles with masses from several GeV up to the TeV scale [5].

A particularly important class of dark-matter candidates is weakly interacting massive particles, or WIMPs. WIMPs are hypothetical, electrically neutral particles that are stable on cosmological time scales and interact with ordinary matter through interactions of approximately weak strength. They are especially well motivated because particles with masses and interaction strengths near the electroweak scale can naturally produce a relic abundance of the observed order through thermal freeze-out in the early Universe [1, 5]. Prominent examples are the lightest neutralino in supersymmetric models and the lightest Kaluza–Klein particle in theories with universal extra dimensions [1]. This makes WIMPs a central target for direct-detection experiments, which search for rare scattering events between dark-matter particles from the Galactic halo and ordinary detector material. Other candidates are motivated by different theoretical considerations. Sterile neutrinos would not participate in Standard Model weak interactions, but could mix with active neutrinos and form a warm-dark-matter component [5]. Warm dark matter consists of particles that retained non-negligible velocities in the early Universe, causing free streaming that suppresses the formation of small-scale structures compared with cold dark matter. The diversity of possible candidates illustrates that the dark-matter problem connects cosmology, astrophysics, and particle physics, and that experimental searches must cover a broad range of masses and interaction strengths.

¹Although the standard model neutrino satisfies all requirements to be a dark matter candidate, its relativistic velocity in the early universe makes it unsuitable to constitute a major fraction of the dark matter in the Universe.

1.4 Hunt for dark matter using direct detection

Direct detection experiments are one of the main experimental approaches in the search for particle dark matter, complementary to indirect searches and collider experiments [1]. Their goal is to observe the scattering of dark matter particles from the Galactic halo with ordinary matter in terrestrial detectors. In the commonly assumed WIMP scenario, dark matter particles may occasionally scatter elastically off atomic nuclei, depositing small recoil energies typically in the keV range [6]. Since WIMPs interact only very weakly, the expected event rates are extremely low, requiring detectors with very low backgrounds and high sensitivity. The expected signal rate depends on astrophysical quantities, such as the local dark matter density and velocity distribution, as well as on particle and nuclear physics parameters including the WIMP mass, scattering cross section and nuclear form factor. For spin-independent interactions, the cross section approximately scales with the nuclear mass number squared (A^2), making heavy nuclei such as xenon attractive target materials [6]. A major challenge in direct detection is the suppression of backgrounds from radioactivity and cosmic rays. Experiments are therefore operated deep underground and use radiopure materials together with passive and active shielding systems [6]. Different detector technologies measure signals such as scintillation light, ionization charge or phonons produced by particle interactions.

2 Xenoscope and liquid-xenon detector principles

Xenoscope is a liquid-xenon research facility operated at the University of Zurich for the development and validation of technologies required by future large-scale xenon detectors. To understand the relevance of water monitoring in Xenoscope, this chapter first introduces the operating principle of a dual-phase liquid and gas-xenon time projection chamber and the microscopic processes responsible for the production of vacuum-ultraviolet scintillation light. The influence of water impurities on the propagation of this light is then discussed. Finally, the Xenoscope detector and its xenon-handling and purification systems are described, providing the experimental context for the integration of the HALO moisture analyzer.

2.1 Purpose of Xenoscope

Xenoscope was developed as a demonstrator for technologies required in future multi-tonne liquid-xenon experiments, particularly those used in rare-event searches such as the direct detection of dark matter. The facility consists of a cryogenic liquid-xenon detector together with dedicated systems for xenon circulation, purification, liquefaction, storage, and recovery [7].

One of the principal goals of Xenoscope is to demonstrate the transport of ionization electrons over a drift distance of 2.6 m in liquid xenon. Electron transport over such large distances requires a highly uniform electric field and a sufficiently long electron lifetime, since electronegative impurities can capture drifting electrons before they reach the liquid–gas interface. Xenoscope therefore provides a platform for studying detector operation, high-voltage stability, xenon purification, and long-distance electron transport under conditions relevant for future large detectors.

The technologies investigated with Xenoscope are relevant for the development of XLZD, a proposed next-generation liquid-xenon dark-matter experiment [8]. Whereas Xenoscope was designed to operate with approximately 360 kg of liquid xenon, the baseline design of XLZD foresees an active xenon mass of approximately 60 t, with the possibility of increasing the total mass to approximately 80 t [8]. The large target mass and low-background design of such an experiment would substantially improve the sensitivity to rare interactions over a wide range of dark-matter particle masses and interaction cross sections.

2.2 Operating principle of a dual-phase liquid-xenon TPC

Time projection chambers (TPCs) are widely used in noble-liquid dark-matter experiments because they allow both for the deposited energy and the position of a particle interaction to be reconstructed. In a dual-phase xenon TPC, an interaction in the liquid produces excitation and ionization of xenon atoms. These processes give rise to two measurable signals, commonly referred to as $S1$ and $S2$.

The $S1$ signal is the prompt scintillation light emitted in the liquid xenon shortly after the interaction. At the same time, electron–ion pairs are produced along the particle track. A fraction of the ionization electrons recombines with xenon ions and contributes to the scintillation signal, while the remaining electrons are transported upwards by an applied electric

drift field. At the liquid–gas interface, a stronger extraction field transfers the electrons into the gaseous xenon. There, they produce a delayed proportional scintillation signal through electroluminescence, known as $S2$ [9]. The operating principle is illustrated in Fig. 1.




Figure 1: Working principle of a dual-phase noble-liquid time projection chamber. A particle interaction in the liquid produces prompt scintillation light ($S1$) and ionization electrons. The drift field E_1 transports the electrons upward through the liquid, while the extraction field E_2 transfers them across the liquid–gas interface. In the gas phase, the electroluminescence field E_3 accelerates the electrons and produces proportional scintillation light ($S2$). The time delay between $S1$ and $S2$ provides the depth of the interaction. Reproduced from Ref. [9].

The combination of $S1$ and $S2$ allows for the interaction position to be reconstructed in three dimensions. The vertical position is determined from the time delay between the prompt $S1$ signal and the delayed $S2$ signal, since the electrons require a finite time to drift from the interaction position to the liquid surface. The horizontal position can be reconstructed from the distribution of the $S2$ light across the photosensors. The sizes of the two signals additionally provide information about the energy deposited in the detector.

The $S1$ and $S2$ signals are anti-correlated. A larger recombination fraction increases the scintillation signal but leaves fewer free electrons available for the ionization signal. Conversely, reduced recombination produces a larger $S2$ signal and a smaller recombination contribution to $S1$. Combining the two signals therefore improves the reconstruction of the deposited energy [10, 11].

The relative amounts of scintillation and ionization also depend on the type of interaction. The ratio $S2/S1$ differs for nuclear recoils, such as those expected from dark-matter scattering, and electronic recoils caused by many electromagnetic backgrounds. This provides an important method for distinguishing possible signal events from background interactions [9, 12].

For the present thesis, the important aspect of the TPC operating principle is that the $S1$ signal is transported through the liquid xenon in the form of vacuum-ultraviolet photons. The microscopic processes responsible for producing these photons are described in the following subsection.

2.3 Production of VUV scintillation light in liquid xenon

When ionizing radiation interacts with liquid xenon, the deposited energy is distributed mainly among excitation, ionization, and heat. A fraction of the deposited energy produces excited xenon atoms, or excitons, denoted by Xe^* , while another fraction creates electron-ion pairs, $Xe^+ + e^-$. Both channels can contribute to the production of scintillation light [11].

In the direct excitation channel, an excited xenon atom interacts with neighbouring xenon atoms and forms an excited xenon dimer, or excimer:



A second contribution arises from recombination. Ionization electrons may recombine with xenon ions, producing highly excited states that subsequently relax to excitons:



The excitons produced through recombination can then form excimers through the same process as directly excited xenon atoms (see Eq. (5)). The excited dimer Xe_2^* exists in a bound excited state, while its ground state is repulsive. It therefore decays into two ground-state xenon atoms while emitting a photon:



This process is the microscopic origin of the prompt $S1$ scintillation signal in liquid xenon.

The wavelength of the emitted photons lie in the vacuum-ultraviolet region. A high-accuracy measurement of the liquid-xenon emission spectrum by Fujii et al. found an approximately Gaussian spectrum centred at 174.8 nm, with a full width at half maximum of 10.2 nm [13]. Oliver-Mallory et al. measured the scintillation spectrum in a dual-phase xenon detector and obtained a peak wavelength of 177.1 nm, with a full width at half maximum of 11.3 nm [14]. The small difference between the measured peak positions may arise from differences in the experimental configurations and measurement conditions [14]. The exact wavelength distribution is relevant because optical properties such as absorption, Rayleigh scattering, reflection, refraction, and sensor response vary with wavelength in the VUV region [13, 14]. In this thesis, the liquid-xenon scintillation wavelength is therefore referred to approximately as 175–178 nm.

The xenon excimer can occupy singlet and triplet states, giving rise to fast and slow components of the scintillation pulse. Recombination additionally contributes to the time structure of the signal. The relative importance of these components depends on the electric field and on the ionization density of the interacting particle [10, 15]. Since the detailed scintillation pulse shape is not central to the present work, it is sufficient to note that these microscopic processes produce the prompt VUV photons that form the $S1$ signal.

Efficient detection of $S1$ requires these photons to propagate through the liquid xenon before reaching the photosensors. Their transmission can therefore be degraded by absorbing impurities, in particular water [10].

2.4 Influence of water impurities on VUV-light transport

In sufficiently pure xenon, intrinsic self-absorption of the scintillation light is expected to be small because the xenon emission spectrum is separated from the principal absorption spectrum of ground-state xenon. Molecular impurities can, however, introduce additional absorption bands in the VUV region [16]. If these bands overlap with the liquid-xenon scintillation spectrum, even low impurity concentrations can reduce the number of photons reaching the photosensors.

The attenuation of scintillation light over a propagation distance d can be described by

$$I(d) = I_0 \exp(-\alpha_{\text{att}} d) = I_0 \exp\left(-\frac{d}{\lambda_{\text{att}}}\right), \quad (8)$$

where I_0 is the initial intensity, α_{att} is the attenuation coefficient, and $\lambda_{\text{att}} = 1/\alpha_{\text{att}}$ is the attenuation length [17]. The attenuation coefficient contains contributions from absorption and Rayleigh scattering,

$$\alpha_{\text{att}} = \alpha_{\text{abs}} + \alpha_{\text{R}} = \frac{1}{\lambda_{\text{abs}}} + \frac{1}{\lambda_{\text{R}}}, \quad (9)$$

where λ_{abs} and λ_{R} are the absorption and Rayleigh-scattering lengths, respectively. Absorption removes photons from the scintillation signal, whereas Rayleigh scattering changes their propagation direction without destroying them.

For water, the absorption coefficient at the xenon scintillation wavelength depends linearly on the water mole fraction x_{H_2O} ,

$$\alpha_{H_2O} = n_{H_2O} \sigma_{H_2O}(178 \text{ nm}) = x_{H_2O} n_{Xe} \sigma_{H_2O}(178 \text{ nm}), \quad (10)$$

where n_{Xe} is the xenon number density and σ_{H_2O} is the absorption cross section of water. The corresponding water-induced absorption length is $\lambda_{\text{abs}, H_2O} = 1/\alpha_{H_2O}$ and therefore decreases inversely with increasing water concentration [17].

The effect of water impurities on the propagation of xenon scintillation light is illustrated in Fig. 2. The wavelength-dependent absorption of water overlaps with the liquid-xenon scintillation spectrum, as shown in Fig. 2(a). Consequently, water can absorb a significant fraction of the emitted VUV photons. The resulting concentration dependence is illustrated in Fig. 2(b), which shows the calculated remaining scintillation-light intensity as a function of propagation distance for different water concentrations. Increasing the water concentration leads to a progressively stronger reduction of the transmitted light and therefore to a shorter effective absorption length.

In a large detector, where photons may travel considerable distances before being detected, this absorption can cause a position-dependent reduction in the measured $S1$ signal and impair the reconstruction of the deposited energy.




Figure 2: Influence of water impurities on liquid-xenon scintillation light. (a) Absorption coefficient for a water concentration of 1 ppm, with the liquid-xenon scintillation spectrum superimposed. (b) Calculated remaining scintillation-light intensity I/I_0 as a function of propagation distance for different water concentrations in LXe. Figure adapted from Ref. [17].

Rayleigh scattering in liquid xenon is primarily caused by microscopic density fluctuations rather than by water impurities. A Rayleigh-scattering length of approximately 30 cm at 174 nm has been calculated for liquid xenon [16]. The concentration dependence introduced by water therefore enters mainly through the absorption coefficient in Eq. (10).

Early measurements in noble liquids reported attenuation lengths of only a few centimetres. Improved material selection and continuous purification later enabled attenuation lengths approaching the metre scale. Continuous circulation through a purifier was reported to reduce the water concentration to below approximately 100 ppb, accompanied by a substantial increase in the attenuation length [16]. This behaviour is consistent with the strong concentration dependence illustrated in Fig. 2.

Water can also be released continuously from the internal surfaces of tubing, vessels, seals, and detector components through desorption and material outgassing. The water concentration is therefore not determined solely by the initial purity of the xenon. It also depends on the cleanliness of the gas system, the exposed surface area, temperature, gas flow, purification rate, and the duration of detector operation.

Continuous purification and monitoring are consequently important for maintaining suitable detector conditions. The purification system removes reactive molecular impurities from the circulating xenon, while a moisture analyzer provides direct information about the water concentration in a sampled gaseous-xenon stream. The concentrations shown in Fig. 2(b) refer to water in liquid xenon, whereas the HALO analyzer used in this work measures the water concentration in gaseous xenon. The measured GXe concentrations therefore cannot be directly converted into an absorption length using Eq. (10) without knowledge of the corresponding liquid-phase water concentration.

2.5 Xenoscope detector and xenon-handling system

The Xenoscope facility is built around a cryogenic detector volume and an external xenon-handling system. The detector provides the low-temperature environment required to condense xenon and operate the liquid-xenon target. The external system controls xenon circulation, purification, liquefaction, storage, and recovery [7]. These functions are essential because xenon is both valuable and sensitive to contamination and must therefore remain contained and recoverable throughout detector operation. A photograph of the Xenoscope facility at the University of Zurich is shown in Fig. 3.



Figure 3: The Xenoscope liquid-xenon research facility at the University of Zurich. The central stainless-steel cryostat houses the liquid-xenon time projection chamber (TPC), while the surrounding aluminium support structure provides mechanical support and access for maintenance. The gas-handling and purification infrastructure is located around and behind the cryostat, and the detector is operated and monitored via the slow-control system displayed on the monitors. Photograph by author.

The detector consists of a cryogenic vessel containing the liquid-xenon target and the associated internal detector structures. The cryostat provides the thermal insulation and mechanical containment required to maintain xenon in the liquid phase. Components exposed to the xenon include the inner vessel, tubing, valves, seals, cables, feedthroughs, and detector-support structures. Water adsorbed on these surfaces or contained in materials can be released through desorption and outgassing, thereby introducing water into the xenon during operation.


The cryogenic detector and the gas system are closely connected. Xenon can be transferred between storage vessels and the detector, condensed into the liquid phase, circulated through the purification system, and recovered after detector operation. Continuous circulation is required to counteract the release of impurities from internal surfaces and system components.

2.5.1 Xenon circulation and purification

The purification loop used in Xenoscope is shown in Fig. [4](#). During circulation, xenon is extracted from the detector and transported through the external gas-handling system. In the configuration relevant to this work, gaseous xenon is withdrawn from the gas phase above the liquid xenon before entering the main purification line (see red line in Fig. [4](#)). A heat exchanger transfers thermal energy between the outgoing and returning xenon streams. This reduces the cooling power required to recondense the purified xenon before it is returned to the detector. After leaving the heat exchanger, the gaseous xenon passes through flow-control and pumping components that establish and regulate the circulation.

The xenon is then passed through a heated getter. The getter removes reactive molecular impurities, including water and oxygen, through chemical binding. Such purification is particularly important for liquid-xenon TPCs because different impurity species affect different detector signals. Electronegative impurities can capture drifting electrons and reduce the $S2$ signal, while optically absorbing impurities such as water can reduce the transmission of the VUV scintillation light and thereby affect the $S1$ light. After purification, the xenon is transported back towards the detector, cooled in the heat exchanger, and returned to the cryogenic volume. Continuous recirculation allows for impurities released from the detector and gas-system surfaces to be transported to the getter and removed progressively.

The observed purity of the circulating xenon depends not only on the getter performance but also on the circulation rate, system volume, surface outgassing, dead volumes, leaks, and the cleanliness of the piping. Water adsorbed on stainless-steel surfaces such as the piping or inner cryostat may be released gradually during operation, particularly after a system has been exposed to ambient air. As a result, the measured water concentration can evolve over time even when the xenon is circulated continuously through the getter.



3 HALO-Monitor

The measurement of trace water concentrations in gaseous xenon requires a highly sensitive and reliable detection system. For this purpose, a Tiger Optics HALO KA trace gas analyzer [18] was employed and can be seen in Fig. 5. The instrument is based on cavity ring-down spectroscopy (CRDS), an absorption-based technique that enables precise and calibration-free determination of molecular concentrations. According to the manufacturer, the measurement precision and accuracy of the HALO KA-M7500-S depend on the measured water concentration and on the carrier gas [18]. For each carrier gas, the manufacturer specifies a lower detection limit (LDL), defined at the 3σ level. For example, the LDL is 300 ppt for nitrogen and 130 ppt for argon [18]. The specified 1σ precision is the maximum of 0.75 % of the measured water concentration and one third of the corresponding LDL [18]. Thus,

$$\sigma_{\text{prec}} = \max \left(0.0075 c_{H_2O}, \frac{\text{LDL}}{3} \right), \quad (11)$$

where c_{H_2O} denotes the measured water concentration. The factor 0.0075 corresponds to the specified relative precision of 0.75 %, while the second term represents the concentration-independent precision floor at low water concentrations. The manufacturer lists the 1σ precision at zero concentration as 100 ppt for nitrogen and 45 ppt for argon [18]. Similarly, the specified accuracy is the maximum of 4 % of the measured water concentration and the corresponding LDL [18],

$$u_{\text{acc}} = \max (0.04 c_{H_2O}, \text{LDL}) . \quad (12)$$

The factor 0.04 therefore represents the specified relative accuracy of 4 %, whereas the LDL sets the minimum accuracy uncertainty at low concentrations.



Figure 5: HALO-monitor used in this experiment. Photograph by the author.

3.1 Cavity ring-down spectroscopy for trace impurity detection

The measurement of trace water concentrations in gaseous xenon is performed using CRDS, a highly sensitive optical absorption technique. The method is based on the absorption of near-infrared laser light by water molecules present in the gas sample. In CRDS, light from a continuous-wave diode laser is coupled into an optical cavity formed by two highly reflective mirrors. The cavity is filled with the sample gas containing trace amounts of water. The specific trace-gas moisture analyzer used in this experiment was the Tiger Optics HALO KA-M7500-S [18]. According to the manufacturer, the laser operates in the near-infrared wavelength range of 1200–1800 nm [18]. The exact operating wavelength used for the water measurement is not specified in the manual. However, the absorption spectrum presented by the manufacturer shows a water absorption feature in the vicinity of $\lambda \approx 1392$ nm, indicating that the laser is tuned to this spectral region during water measurements [18].

Due to the high reflectivity of the cavity mirrors, the light is reflected multiple times between them. This increases the effective optical path length through the sample gas and thereby enhances the sensitivity to small absorption effects. The separation of the mirrors is denoted by the cavity length d ; a numerical value for d is not specified in the manufacturer documentation.




Figure 6: Schematic illustration of the working principle of the HALO moisture analyzer. A continuous-wave (CW) near-infrared diode laser is coupled into an optical cavity containing the sample gas. The light undergoes multiple reflections between two highly reflective mirrors separated by the cavity length d , while a fraction of the light transmitted through the second mirror is detected by a photodiode. Adapted from Ref. [18].

A small fraction of the circulating light passes through the second mirror and is detected by a photodiode [18]. Once the detector reaches a predefined light level, the incident laser light is diverted from the cavity. The remaining light then decays due to losses at the cavity mirrors and absorption by molecules in the sample gas. This exponential decrease in light intensity

is referred to as the “ring-down” [18] signal and can be described by

$$I(t) = I_0 \exp\left(-\frac{t}{\tau}\right), \quad (13)$$

where I_0 is the light intensity at the beginning of the ring-down and τ is the characteristic ring-down time [18]. The value of τ depends on the total optical losses in the cavity and is therefore sensitive to additional absorption by trace water molecules.

In order to determine the concentration of impurities, the measured ring-down time $\tau(\nu)$, at laser frequency ν corresponding to an absorption line is compared to a reference measurement denoted by τ_0 at a wavelength without absorption.




Figure 7: Example of the wavelength dependence of the ring-down time around a water absorption feature. The on-peak and off-peak positions correspond to different laser wavelengths λ and are used to determine the contribution of molecular absorption to the cavity losses. Adapted from Ref. [18].

The impurity concentration is then given by

$$N = \frac{1}{c \sigma(\nu)} \left(\frac{1}{\tau(\nu)} - \frac{1}{\tau_0} \right), \quad (14)$$

where c denotes the speed of light, $\sigma(\nu)$ is the absorption cross section of the molecule at the laser frequency ν , and τ_0 is the off-peak reference ring-down time [18]. The high sensitivity of CRDS results from the effective optical path length produced by the repeated reflections inside the cavity together with the measurement of the ring-down time. Concentrations of water can be detected with a precision between ppm and ppb range. In practice, the laser wavelength is

tuned to a specific absorption line of water in the infrared region. By comparing measurements at absorption and non-absorption wavelengths, the influence of the cavity geometry and mirror reflectivity can be separated from the absorption due to impurities.

This technique is implemented in the HALO moisture analyzer, which enables continuous monitoring of water impurities in gaseous xenon systems. The method is particularly suited for applications requiring high sensitivity and non-invasive measurements, as no sample gas is consumed during the process.

3.2 Internals of HALO-Monitor

Internally, the analyzer consists of both an optical measurement system and a gas handling system. The gas handling system ensures controlled delivery of the sample gas into the measurement cavity. As shown in Fig. 8, the sample gas enters the analyzer through the sample inlet and first passes a particulate filter. This filter is designed to remove impurities introduced during installation or operation, which could otherwise contaminate the optical components [18]. After filtration, the gas flows through an internal pressure regulator, which maintains a stable cell pressure of approximately 0.4 barg, where barg denotes pressure in bar relative to atmospheric pressure. While the CRDS technique itself is largely insensitive to pressure variations, maintaining a stable pressure ensures consistent gas density and reproducible measurement conditions [18].




Figure 8: Schematic of the internal gas-flow path of the HALO moisture analyzer. The sample gas passes through a $2\ \mu\text{m}$ particulate filter and a pressure regulator before entering the CRDS sample cell where it interacts with the gas and the absorption due to water molecules is quantified via the ring-down time. It then leaves the analyzer through an orifice and the sample outlet and is subsequently vented or returned to the system. Adapted from Ref. [18].

Although the CRDS measurement does not require a precise flow rate, the transport of moisture within the system is strongly flow dependent. Higher flow rates lead to faster response times, while lower flow rates result in longer equilibration times, particularly during dry-down phases [18]. Additionally, pressure fluctuations can influence the propagation of moisture in the system and thus indirectly affect the measured concentration. It is important to note that the analyzer does not actively control the gas flow in most configurations. Instead, gas flow is driven by external pressure differences, and gas may continue to flow through the analyzer even when it is powered off. Therefore, careful control of the upstream and downstream conditions is essential for reliable measurements.

3.3 HALO Piping and Integration

The integration of the HALO moisture analyzer into the Xenoscope gas system required the installation of a dedicated sampling line and valve panel. The purpose of this integration was to allow gaseous xenon to be extracted close to the cryostat, passed through the HALO analyzer, and recovered after the measurement without disrupting the main xenon gas system. Since the expected water concentrations are in the ppb range, particular care had to be taken to minimize leaks and water outgassing from internal surfaces of its upstream piping.

The final configuration was designed to fulfil three main requirements. First, the HALO analyzer had to be connected in such a way that a stable gas flow through the measurement cell could be established. Second, the gas passing through the analyzer had to be recoverable, which is especially important for xenon operation, due to its significant price. Third, the sampling line had to be leak-tight and sufficiently clean to avoid artificially increased water readings caused by ambient air leaks or water desorption from contaminated surfaces.

3.3.1 Pipe Installation

The HALO sampling line was constructed using a 6m long stainless-steel 1/4-inch tubing divided into three 2m segments connected through VCR connections. These are suitable for high-purity gas systems because they provide metal-sealed connections with low leak rates and reduced outgassing compared to polymer-based fittings. Stainless steel was chosen in order to minimize water permeation and to ensure compatibility with xenon operation.

Before installation, the 2 m long 1/4-inch stainless-steel tube segments were cleaned after the VCR glands had been welded on. Each tube was flushed with ELMA60 cleaning solution, first at a concentration of approximately 10 % and subsequently with freshly prepared 5 % solution. Between and after the cleaning steps, the tubes and connectors were thoroughly rinsed with deionized water, followed by a final ethanol rinse and drying under a continuous N_2 5.0 purge for roughly 12 hours. This step is particularly important for moisture measurements, since water adsorbed on the inner surfaces of the piping can be released slowly during operation and lead to elevated background readings. The cleaned tubes were then installed between the heat exchanger outlet on the top floor and the HALO-monitor panel. The installed sampling line was then baked with a heating tape at 70 °C whilst being vacuumed in order to eliminate residual water. The sampling line was kept as short and simple as possible in order to reduce dead volume and limit the number of additional VCR connections.

The installed piping provides an inlet path from the Xenoscope gas system to the HALO analyzer and an outlet path from the analyzer to a recovery vessel with a volume of roughly 100 liters. The recovery vessel is initially evacuated in order to provide the positive pressure gradient needed to establish gas flow through the HALO. This configuration allows gas to flow through the HALO measurement cell while preventing direct venting of xenon to the laboratory. In addition, a bypass line was included so that the line could be flushed when not in use without going through the HALO-monitor. The detailed schematic of how the HALO-monitor was integrated in the system can be seen in Fig. 9:




Figure 9: Flow diagram of the HALO configuration within the Xenoscope gas-handling system. The blue line indicates the sampling path with the 0.05 standard liters per minute (SLPM) through the HALO moisture analyzer. MV20 provides a bypass around the analyzer, PT09 monitors the line pressure, and MV22 connects the sampling line to the vacuum vessel used for xenon recovery. The recovery vessel is connected to the xenon storage array. We can see the fully assembled panel in Fig. 11. Figure taken from [19].

3.3.2 Leak Checking

After installation, the HALO sampling line and the associated valve panel were leak checked. This step was essential because even small leaks to atmosphere would introduce water vapour and oxygen into the gas stream. Since the HALO analyzer is sensitive to water concentrations down to the ppb range, a leak that would be negligible for ordinary gas handling could still lead to a significant systematic error in the measured moisture concentration.

The leak checks were performed using a helium leak detector. Individual connections were tested by spraying helium locally onto the fittings while monitoring the leak detector signal. If an elevated leak rate was observed, the corresponding connection was re-tightened or replaced and tested again. This procedure was repeated until the leak rate of the sampling line was sufficiently low for high-purity operation.

The final measured leak rate was below approximately $4.4 \times 10^{-10} \text{ mbar L s}^{-1}$ which was considered acceptable for the subsequent xenon measurements. This result confirmed that the installed HALO branch did not introduce a significant leak path into the Xenoscope gas system.




Figure 10: Initial HALO panel. The pressure sensor PT09 is still missing. Photograph by the author.

3.3.3 Panel integration

For an optimal setup of the HALO gas system a valve-panel was installed in the lower floor of xenoscope. The HALO analyzer was integrated into the gas handling system in such a way that a quasi-continuous gas sample can be extracted and captured after measurement.




Figure 11: Final configuration of the HALO monitor panel corresponding to the schematic shown in Fig. [9]. The panel includes the PT09 pressure sensor, stainless-steel VCR tubing and fittings, the sampling and bypass valves, and the connections to the HALO moisture analyzer. Photograph by the author.

This setup was chosen as initially the pressure differential was not high enough for the HALO-monitor. Prior to this setup, when the panel sat between the cryostat and purification line, the inlet and outlet side had a difference of roughly 400 mbar, which was substantially lower than the minimum amount of 900 mbar for flow to happen through the internal pressure regulator and orifice of the HALO-monitor. Thus, an evacuated vessel was implemented such that the outlet pressure was at vacuum and could be tuned with MV22. The full vessel can be seen in Fig. [12]. Similarly, the panel was flushed several times over in order to minimize moisture buildup within the pipes. This setup allowed a quasi-continuous measurement of 1-2 hours per sitting.

3.4 Measurement Procedure

The HALO measurement procedure was designed to establish a controlled gas flow through the HALO analyzer while keeping the xenon recoverable. Since the flow through the analyzer is driven by pressure differences in the gas system, the valve sequence and pressure conditions had to be controlled carefully during each measurement. In addition, the gas leaving the HALO analyzer was collected in an evacuated vessel placed on a scale, which allowed the amount of extracted xenon to be monitored during the measurement. If necessary, the scale was tared before the start of the measurement. The HALO analyzer was checked to be in measure mode before the gas flow through the analyzer was established.

The measurement was started by slowly closing the outlet valve MV21 on the HALO panel. During this step, the inlet valve MV19 at the purification unit was left unchanged, since its opening position had previously been adjusted to provide the desired bypass flow of approximately 0.5 SLPM. A second person monitored the relevant system parameters during this operation, in particular the pump inlet pressure, the recirculation flow, and the cryostat pressure. After MV21 was closed, the HALO bypass valve MV20 was closed as well in order to constrain the gas flow through the HALO-monitor. The valve to the recovery vessel was then closed, and the pressure on the HALO panel was checked using PT09. At this stage, PT09 typically read approximately 2.1 bar. It was also verified that both MV20 and MV21 were fully closed.

The start time of the measurement was recorded before opening the extraction path. The valve MV22 at the bottom of the HALO panel was then opened slowly until a small and controlled pressure decrease was observed on PT09. Once a suitable pressure drop had been established, the valve to the recovery vessel was opened. During normal operation, the pressure measured by PT09 decreased to approximately 1200 mbar, providing the pressure difference required to drive gas through the HALO analyzer and into the recovery vessel.

During the measurement, both the pressure PT09 and the mass of the recovery vessel were monitored continuously. As the vessel filled, the pressure conditions changed, so MV22 continuously had to be opened further in order to maintain a sufficient pressure difference and therefore a stable flow through the analyzer. The water concentration reported by the HALO was checked regularly. An increase in the measured concentration could indicate that the gas flow had decreased or stopped, since stagnant gas in the sampling line can lead to an enhanced contribution from outgassing.

To stop a measurement, the stop time and the final mass increase of the recovery vessel were first recorded. The valve MV22 at the bottom of the HALO panel was then closed, followed by the valve at the recovery vessel. Afterwards, the HALO bypass valve MV20 was reopened. Finally, the HALO outlet valve MV21 towards the purification system was opened slowly. During this step, the system parameters were again monitored by a second person to ensure that the purification outlet flow did not increase by more than approximately 1 SLPM, with a preferred increase of about 0.5 SLPM.



Figure 12: Assembly of the approximately 100 L expansion vessel used to collect the xenon gas after it passed through the HALO moisture analyzer. Photograph by author.

4 Data Analysis and Discussion

The HALO moisture analyzer determines the water concentration from the measured cavity ring-down signal and internally converts the spectroscopic response into a concentration in ppb. The resulting datasets therefore consist of time series of the measured H_2O concentration, recorded at approximately one-minute intervals and given as a CSV datasheet. In the following, the temporal evolution of the water concentration is analysed for nitrogen, argon, and gaseous xenon.

4.1 Analysis of water content in N_2 and argon

Before performing the xenon measurements, the gas system and the HALO moisture analyzer were tested with nitrogen and argon in order to study the behaviour of the measured water concentration during extended flushing. The nitrogen and argon were supplied from 50 L high-pressure gas cylinders through a pressure reducer and were connected to the HALO-monitor panel by a 1.5 m-long, 1/2-inch VCR bellows. These measurements were not intended to provide the final xenon purity result, but rather to investigate how the measured water concentration evolves during prolonged gas flushing. The water concentration was recorded as a function of time with the HALO moisture analyzer.

4.1.1 Qualitative behaviour of the flushing measurements

The nitrogen and argon measurements both show the same qualitative behaviour. At the beginning of each measurement, the water concentration is comparatively high and lies in the ppm range. During flushing, the concentration decreases rapidly and reaches a minimum value. At late times, both measurements show a renewed increase in water concentration. The measurements therefore indicate that the observed concentration is not determined only by the removal of water from the gas phase, but is also influenced by the operating conditions of the gas system. Since the nitrogen and argon gas bottles were gradually depleted, the inlet pressure and flow conditions changed over time. This could have reduced the flushing efficiency, such that outgassing from internal surfaces and less efficiently flushed volumes became increasingly important.

For nitrogen, the concentration initially started at approximately 8.4 ppm and decreased rapidly shortly after the start of the measurement. It fell below 1000 ppb after approximately 1.8 h, below 700 ppb after 3.8 h, and below 500 ppb after 5.9 h. The minimum concentration reached during the nitrogen measurement was 446 ppb which occurred after approximately 6.1 h. Compared with the maximum concentration, this corresponds to a reduction of 94.7 %. After reaching this minimum, the concentration did not remain constant. It increased again towards the end of the measurement, where a final value of approximately 630 ppb was measured. This indicates that the long-term nitrogen measurement is not governed only by flushing, but also by continuous water release from internal surfaces and changing operating conditions.

The nitrogen fit and the corresponding residual are shown in Fig. [13](#). The first few measurements and the late-time part after the selected fit interval are shown in gray and were

excluded from the fit. The black curve shows the data used in the fit and the red curve shows the fitted double-exponential model which is further motivated in Section 4.1.2.




Figure 13: Water concentration measured during the nitrogen flushing measurement, together with the double-exponential fit. The concentration axis is logarithmic. Data excluded from the selected fit interval are shown in gray, fitted data are shown in black, and the fitted model is shown in red. The lower panel contains the residuals $c_i - c_{\text{fit}}(t_i)$.

The argon measurement showed a similar but more extended flushing behaviour as the measurement was run for longer duration. The maximum concentration at the beginning of the measurement was approximately 4 ppm. The concentration decreased below 1000 ppb after about 1.0 h, below 700 ppb after 1.8 h, below 500 ppb after 3.9 h, and below 300 ppb after 12.9 h. The lowest measured value was 259 ppb, which was reached after approximately 27.8 h. This corresponds to a reduction of 93.5 % relative to the maximum concentration. After the minimum was reached, the concentration increased again and reached a final value of approximately 690 ppb at the end of the measurement. In the final part of the argon measurement, after approximately 45 h, the increase in water concentration was approximately linear with a slope of roughly

$$\frac{dc_{H_2O}}{dt} \approx 21 \text{ ppb/h.} \quad (15)$$

This rise suggests that water was still being released into the gas stream, most likely from internal surfaces or volumes that were not efficiently flushed.

The argon fit and residual are shown in Fig. 14. The late-time increase is shown in gray and was not used to determine the dry-down parameters.




Figure 14: Water concentration measured during the argon flushing measurement, together with the double-exponential fit. The concentration axis is logarithmic. The two peaks in the beginning could be attributed to unstable flow conditions during which the concentration increased. Data excluded from the selected fit interval are shown in gray, fitted data are shown in black, and the fitted model is shown in red. The lower panel contains the residuals $c_i - c_{\text{fit}}(t_i)$.

The reduction was calculated from the maximum and minimum measured concentrations according to

$$R_{\text{red}} = \left(1 - \frac{c_{H_2O, \text{min}}}{c_{H_2O, \text{max}}} \right) \cdot 100 \%. \quad (16)$$

Both gases therefore reduce the measured water concentration by more than 90 % and bring the gas system from ppm-level water contamination down to the few-hundred-ppb range. However, the later increase in both measurements shows that water is continuously released from the system and that the complete time evolution cannot be described by a monotonic dry-down model.

Table 1: Summary of the nitrogen and argon flushing measurements. The time is given relative to the beginning of each measurement. The maximum expected concentration is taken from the vendor certificate for the specified gas purity and represents the upper limit of the total impurity concentration, not specifically the water concentration. The values are reported without uncertainty as they reflect a single data point taken by the HALO-monitor.

Gas	$c_{\max, \text{exp.}}$ [ppb]	Duration [h]	Gas flow [SLPM]	$c_{\max}$ [ppb]	$c_{\min}$ [ppb]	$t(c_{\min})$ [h]	Reduction [%]
N_2 (5.0)	$\leq 10\,000$	33.1	0.5	8372	446	6.1	94.7
Ar (6.0)	$\leq 1\,000$	62.2	0.9	3976	259	27.8	93.5

4.1.2 Double-exponential description of the decline

The removal of water from the gas system and the depletion of water released by outgassing can, to first approximation, be described by linear first-order differential equations. Their solutions are exponential functions. The presence of two characteristic removal time scales therefore motivates the use of a double-exponential model. Here, data taken at later times when the concentration increased was excluded from the fit.

The selected dry-down intervals were fitted with the double-exponential model

$$c_{H_2O}(t) = c_{H_2O, \infty} + A_{\text{fast}} \exp\left(-\frac{t}{\tau_{\text{fast}}}\right) + A_{\text{slow}} \exp\left(-\frac{t}{\tau_{\text{slow}}}\right). \quad (17)$$

The double-exponential model was chosen because a single-exponential model did not provide a sufficient description of the measured dry-down behaviour.

Here, A_{fast} and A_{slow} are the amplitudes of the fast and slow components, τ_{fast} and τ_{slow} are their characteristic time constants, and $c_{H_2O, \infty}$ is the fitted asymptotic water concentration. The corresponding component half-lives are

$$t_{1/2, i} = \tau_i \ln(2). \quad (18)$$

The relative residuals were calculated as

$$r_i = \frac{c_{\text{fit}, H_2O}(t_i) - c_{H_2O, i}}{c_{H_2O, i}} \cdot 100 \%. \quad (19)$$

The goodness of fit was quantified using the coefficient of determination R^2 , which measures how well the fitted model reproduces the variation in the measured water concentration. Values of R^2 close to one indicate that a large fraction of the observed variation is described by the model. It is calculated as

$$R^2 = 1 - \frac{\sum_i (c_{H_2O, i} - c_{\text{fit}}(t_i))^2}{\sum_i (c_{H_2O, i} - \bar{c}_{H_2O})^2}, \quad (20)$$

where $c_{H_2O,i}$ is the measured water concentration at time t_i , $c_{\text{fit}}(t_i)$ is the corresponding value predicted by the fit, and $\bar{c}_{H_2O}$ is the mean measured water concentration. The coefficient of determination is used here as a quantitative measure of how well the chosen fit function describes the observed concentration decrease.

The fit parameters are summarized in Table 2. The quoted uncertainties are the statistical parameter uncertainties obtained from the covariance matrix of the nonlinear least-squares fit.

Table 2: Double-exponential fit parameters for the selected nitrogen and argon dry-down intervals. Amplitudes and asymptotic concentrations are given in ppb, while the time constants are given in minutes.

Gas	A_{fast} [ppb]	τ_{fast} [min]	A_{slow} [ppb]	τ_{slow} [min]	c_{∞} [ppb]	R^2
N ₂	5210 ± 50	9.80 ± 0.12	1503 ± 13	103.5 ± 0.9	501.8 ± 0.8	0.9934
Ar	2589 ± 13	21.65 ± 0.18	622 ± 5	252 ± 3	271.3 ± 0.8	0.9947

For nitrogen, the fit yields a fast time constant of

$$\tau_{\text{fast},N_2} = (9.80 \pm 0.12) \text{ min} \quad (21)$$

and a slow time constant of

$$\tau_{\text{slow},N_2} = (103.5 \pm 0.9) \text{ min.} \quad (22)$$

These values correspond to approximately 0.163 h and 1.72 h, respectively. The associated half-lives are $(6.79 \pm 0.08) \text{ min}$ and $(71.7 \pm 0.6) \text{ min}$. For argon, the fitted time constants are

$$\tau_{\text{fast},Ar} = (21.65 \pm 0.18) \text{ min} \quad (23)$$

and

$$\tau_{\text{slow},Ar} = (252 \pm 3) \text{ min,} \quad (24)$$

corresponding to approximately 0.361 h and 4.21 h. The associated half-lives are $(15.01 \pm 0.12) \text{ min}$ and $(174.9 \pm 1.8) \text{ min}$. The two time constants in each measurement differ by approximately one order of magnitude. This supports the use of two characteristic time scales to describe the global dry-down behaviour. A possible interpretation is that the fast component is dominated by rapid replacement of high humidity gas in the directly flushed volume, whereas the slow component reflects delayed release of water from surfaces, tubing, or less efficiently flushed regions. These assignments are plausible but are not uniquely established by the fit itself.

The fitted asymptotic concentrations are

$$c_{\infty, N_2} = (501.8 \pm 0.8) \text{ ppb} \quad (25)$$

and

$$c_{\infty, Ar} = (271.3 \pm 0.8) \text{ ppb}. \quad (26)$$

These quantities should not be confused with the lowest individual measured values of 446 ppb and 259 ppb. The measured minima refer to single data points, whereas c_{∞} is the asymptotic level inferred from the fitted dry-down interval.

The values $R^2 = 0.9934$ for nitrogen and $R^2 = 0.9947$ for argon show that the double-exponential curves reproduce most of the overall concentration variation. Nevertheless, the residual panels in Fig. 13 and Fig. 14 still contain systematic structures, especially during the early transient and during local disturbances. Therefore, the high R^2 values do not imply that the model captures every short-term process or that the two fitted components can be assigned uniquely to specific physical mechanisms.

4.1.3 Implications for the xenon measurements

The nitrogen and argon measurements provide important information for the interpretation of the subsequent xenon measurements. They show that flushing can reduce the water concentration inside the pipes from the ppm range to the ppb range. They also show that the measured concentration is sensitive to outgassing, dead volumes, and changing flow conditions. Thus, a bypass flow should be present inbetween measurements to prevent the accumulation of outgassing components. For the xenon measurements, quantitative fits should therefore be applied only to well-defined measurement intervals. Measurements disturbed by technical issues, in particular the pump problem observed during one of the runs, are treated separately.

4.2 Analysis of water content in GXe

The xenon recirculation was started on April 8, 2026. The four HALO measurements were performed on May 11, May 18, May 20, and May 21, corresponding to 33, 40, 42, and 43 days after the start of recirculation.

4.2.1 Overview of the xenon measurements

After the nitrogen and argon flushing measurements, the water concentration was measured in gaseous xenon. Four xenon measurements were recorded with the HALO moisture analyzer. The first three measurements show a clear dry-down from the hundred-ppb range to approximately 20 ppb. The fourth measurement shows a different behaviour and was affected by a pump problem of the main xenon circulation pump during the measurement. It is therefore treated separately as a diagnostic measurement and is not used to determine the representative water concentration under stable conditions.

A summary of the raw concentration ranges is given in Table 3. The concentration reduction was calculated according to

$$R_{\text{red}} = \left(1 - \frac{c_{H_2O,\text{min}}}{c_{H_2O,\text{max}}}\right) \cdot 100 \text{ \%}. \quad (27)$$

Table 3: Summary of the four gaseous-xenon measurements. The specified upper limit for the water concentration in the xenon based on the vendor certificate was 2000 ppb. Continuous purification had already reduced the concentration well below this value before the measurements began.

Run	$c_{\text{max,exp.}}$ [ppb]	Duration [min]	c_{max} [ppb]	c_{min} [ppb]	Reduction [%]
GXe, run 1	< 2000	45.0	276.2	23.1	91.6
GXe, run 2	< 2000	89.1	239.9	18.6	92.2
GXe, run 3	< 2000	60.1	291.0	20.6	92.9
GXe, run 4	< 2000	76.0	240.7	118.0	51.0

The first three xenon measurements showed decreases of more than 90 % of the initial value and reach the low-ppb range. Run 4 reaches only approximately 118 ppb, even though its initial concentration is comparable to the other runs. This makes run 4 distinct from the stable measurements. The gaseous-xenon measurements were shorter than the nitrogen and argon runs because the evacuated xenon vessel limited the available gas volume and therefore the achievable measurement duration.

4.2.2 Double-exponential dry-down model

All four xenon measurements were fitted with the double-exponential model introduced in equation (17). The fitted asymptotic concentration c_{∞} was treated as a free parameter together with both amplitudes and both time constants. The analysis therefore provides a fast and a slow characteristic time scale for each measurement.

The fit parameters and R^2 values are summarized in Table 4. The quoted uncertainties are the statistical uncertainties obtained from the covariance matrix of the fit.

Table 4: Double-exponential fit parameters for the four gaseous-xenon measurements. Amplitudes and asymptotic concentrations are given in ppb, while the time constants are given in minutes. Run 4 is included as a diagnostic measurement but is not used to determine the representative stable xenon concentration.

Run	A_{fast} [ppb]	τ_{fast} [min]	A_{slow} [ppb]	τ_{slow} [min]	c_{∞} [ppb]	R^2
GXe, run 1	166 ± 5	0.96 ± 0.06	88 ± 4	9.8 ± 0.7	23.3 ± 1.0	0.9965
GXe, run 2	113 ± 3	0.65 ± 0.05	105 ± 3	8.6 ± 0.3	21.5 ± 0.3	0.9956
GXe, run 3	1120 ± 70	0.54 ± 0.02	105 ± 2	7.7 ± 0.2	21.8 ± 0.3	0.9989
GXe, run 4	29 ± 11	8 ± 2	107 ± 9	29 ± 3	109 ± 2	0.9984

4.2.3 Fit and residual diagnostics for GXe runs 1, 2 and 3

The fit for GXe run 1 is shown in Fig. 15. The double-exponential model follows the concentration decrease over the complete measurement and yields

$$c_{\infty,1} = (23.3 \pm 1.0) \text{ ppb.} \quad (28)$$

The fast and slow time constants are

$$\tau_{\text{fast},1} = (0.96 \pm 0.06) \text{ min} \quad (29)$$

and

$$\tau_{\text{slow},1} = (9.8 \pm 0.7) \text{ min}, \quad (30)$$

corresponding to component half-lives of $(0.67 \pm 0.04) \text{ min}$ and $(6.8 \pm 0.5) \text{ min}$, respectively. The fit has $R^2 = 0.9965$. The relative deviations are largest during the first few minutes and then remain close to zero, although a weak systematic structure is still visible.




Figure 15: Water concentration measured during GXe run 1, together with the double-exponential fit. The concentration axis is logarithmic. The lower panel shows the fit residuals.

The fit for GXe run 2 is shown in Fig. 16. It yields

$$c_{\infty,2} = (21.5 \pm 0.3) \text{ ppb}, \quad (31)$$

with

$$\tau_{\text{fast},2} = (0.65 \pm 0.05) \text{ min} \quad (32)$$

and

$$\tau_{\text{slow},2} = (8.6 \pm 0.3) \text{ min}. \quad (33)$$

The corresponding half-lives are $(0.45 \pm 0.03) \text{ min}$ and $(5.99 \pm 0.18) \text{ min}$, and the fit has $R^2 = 0.9956$. The residuals show a smooth time-dependent structure: they are negative after the initial transient, become positive over the middle of the run, and turn negative again towards the end. Therefore, the high R^2 value shows that the model reproduces the overall decrease very well, but the residuals indicate that the transient is not described perfectly at every time.




Figure 16: Water concentration measured during GXe run 2, together with the double-exponential fit. The concentration axis is logarithmic. The lower panel shows the fit residuals and reveals a remaining systematic time dependence despite the high R^2 value.

The fit for GXe run 3 is shown in Fig. 17. The short non-monotonic transient preceding the main decrease was excluded, and the double-exponential model follows the concentration decrease over the selected dry-down interval and yields

$$c_{\infty,3} = (21.8 \pm 0.3) \text{ ppb.} \quad (34)$$

The fast and slow time constants are

$$\tau_{\text{fast},3} = (0.54 \pm 0.02) \text{ min} \quad (35)$$

and

$$\tau_{\text{slow},3} = (7.74 \pm 0.19) \text{ min}, \quad (36)$$

corresponding to component half-lives of $(0.376 \pm 0.014) \text{ min}$ and $(5.36 \pm 0.13) \text{ min}$, respectively. The fit has $R^2 = 0.9989$. The residuals remain small over most of the selected interval.




Figure 17: Water concentration measured during GXe run 3, together with the double-exponential fit. The lower panel shows the fit residuals.

Runs 1–3 are the most mutually consistent xenon measurements. Their fitted asymptotic concentrations lie between approximately 21 and 23 ppb. This reproducibility supports the conclusion that the system reaches a stable low-ppb water concentration under normal operating conditions.

4.2.4 Special behaviour of GXe run 4

The fourth xenon measurement differs clearly from the first three runs. The water concentration decreases from 241 ppb to 118 ppb, but does not reach the low-ppb range. The double-exponential fit shown in Fig. 18 gives

$$\begin{aligned} c_{\infty,4} &= (109 \pm 2) \text{ ppb}, \\ \tau_{\text{fast},4} &= (8 \pm 2) \text{ min}, \\ \tau_{\text{slow},4} &= (29 \pm 3) \text{ min}. \end{aligned} \tag{37}$$

The corresponding half-lives are $(5.3 \pm 1.5) \text{ min}$ and $(20 \pm 2) \text{ min}$, and the fit has $R^2 = 0.9984$. Although the observed trend is reproduced closely, the fitted asymptotic concentration is approximately five times higher than the representative value of runs 1–3.

Run 4 was affected by a pump problem, and stable operating conditions could therefore not be guaranteed. It was consequently excluded from the determination of the representative GXe water concentration. The result nevertheless indicates that the measured concentration is sensitive to the gas recirculation flow rate and pumping configuration. Reduced or disturbed flow may increase the contribution from stagnant gas volumes and surface outgassing, preventing the system from reaching the lower concentration observed during stable operation.




Figure 18: Water concentration during GXe run 4 together with the double-exponential fit and residuals. Although the fit has a high R^2 , the fitted asymptotic concentration remains far above the stable low-ppb range reached in runs 1–3.

4.2.5 Representative xenon concentration

The first three xenon measurements give fitted asymptotic concentrations of

$$c_{\infty,1} = (23.3 \pm 1.0) \text{ ppb}, \quad (38)$$

$$c_{\infty,2} = (21.5 \pm 0.3) \text{ ppb}, \quad (39)$$

and

$$c_{\infty,3} = (21.8 \pm 0.3) \text{ ppb}. \quad (40)$$

The unweighted mean of the three fitted central values is

$$\bar{c}_{\infty, \text{GXe}} = (22.2 \pm 1.0) \text{ ppb}, \quad (41)$$

where the quoted uncertainty is the standard deviation between the three measurements. The run-to-run standard deviation is used because it reflects the reproducibility of the complete measurement procedure better than the formal covariance uncertainty of a single fit.

The fast time constants of runs 1–3 range from 0.543 min to 0.96 min, while the slow time constants range from 7.74 min to 9.8 min. The unweighted mean of the slow time constants is approximately

$$\bar{\tau}_{\text{slow, GXe}} = (8.7 \pm 1.0) \text{ min}. \quad (42)$$

All three stable xenon fits have $R^2 > 0.995$, showing that the double-exponential model reproduces nearly all of the overall concentration variation. Nevertheless, the residual plots must still be considered together with R^2 , since a high value does not guarantee that the residuals are fully random or that the fitted components correspond uniquely to individual physical processes.

4.3 Measurement of the water-ice vapour pressure at LXe temperature

The HALO moisture analyzer measures the molar water concentration in gaseous xenon. Although this value does not directly provide the water concentration dissolved in liquid xenon, it can be used to determine the partial pressure of water in the sampled gas and thereby to estimate the vapour pressure of water ice at the LXe operating temperature.

The representative asymptotic GXe water concentration determined from the stable xenon measurements in the previous subsection is

$$\bar{c}_{\infty, \text{GXe}} = (22.2 \pm 1.0) \text{ ppb}, \quad (43)$$

According to Dalton's law of partial pressures [20], the partial pressure p_i of a component (i) in a gas mixture is related to its molar fraction x_i and the total gas pressure (P) by

$$p_i = Px_i. \quad (44)$$

The xenon pressure during the measurement was approximately

$$P_{Xe} = 2.08 \text{ bar} = 2.08 \times 10^5 \text{ Pa}. \quad (45)$$

The experimentally determined partial pressure of water in the sampled xenon gas is therefore

$$\begin{aligned} p_{ice}^{\text{exp}} &= P_{Xe} x_{H_2O}^{\text{exp}} \\ &= (2.08 \times 10^5 \text{ Pa}) \left[(22.2 \pm 1.0) \times 10^{-9} \right] \\ &= (4.62 \pm 0.21) \text{ mPa}. \end{aligned} \quad (46)$$

Here, the stated uncertainty results from the uncertainty of the fitted water concentration. A possible uncertainty in the xenon pressure is not included since no quantified uncertainty for the pressure measurement was available. The vapour pressure of water ice has been investigated extensively in the literature. Murphy and Koop provide the following parametrization for its temperature dependence [21]:

$$p_{ice}^{\text{theory}}(T) = \exp \left(9.550426 - \frac{5723.265}{T} + 3.53068 \ln(T) - 0.00728332T \right), \quad (47)$$

where T is given in Kelvin and p_{ice}^{theory} is obtained in pascal. At the xenon operating temperature during the measurement

$$T_{Xe} = 178.7 \text{ K}, \quad (48)$$

the parametrization predicts

$$p_{ice}^{\text{theory}}(178.7 \text{ K}) = 4.21 \times 10^{-3} \text{ Pa} = 4.21 \text{ mPa}. \quad (49)$$

For comparison with the HALO concentration measurement, the theoretical vapour pressure can be converted into the corresponding saturated gas-phase molar fraction by rearranging Dalton's law:

$$x_{\text{sat}}^{\text{theory}} = \frac{p_{ice}^{\text{theory}}}{P_{Xe}}. \quad (50)$$

Using the measured xenon pressure gives

$$\begin{aligned}
 x_{\text{sat}}^{\text{theory}} &= \frac{4.21 \times 10^{-3} \text{ Pa}}{2.08 \times 10^5 \text{ Pa}} \\
 &= 2.03 \times 10^{-8} \\
 &= 20.3 \text{ ppb}.
 \end{aligned} \tag{51}$$

The experimentally determined water-ice vapour pressure of (4.62 ± 0.21) mPa is in remarkably close agreement with the theoretical value of 4.21 mPa predicted by the Murphy–Koop parametrization. The relative difference is only approximately 10 %. Equivalently, the measured gas-phase water concentration of (22.2 ± 1.0) ppb closely matches the predicted saturated concentration of 20.3 ppb. For comparison, the temperature dependence predicted by the Murphy–Koop parametrization is shown in Fig. 19. The figure also includes independent experimental measurements of the water-ice vapour pressure reported by Bielska et al. [22]. These measurements extend down to approximately 175 K and have reported relative standard uncertainties of 0.4–0.7 % [22]. As these uncertainties are small on the scale of the figure, they are not shown. For the Xenoscope point, the vertical error bar represents the uncertainty of the inferred water partial pressure.




Figure 19: Water-ice vapour pressure as a function of temperature. The solid line represents the Murphy–Koop parametrization [21]. The crosses show independent experimental measurements reported by Bielska et al. [22]. The point labelled “Xenoscope, this work” represents the water partial pressure inferred from the HALO measurement at 178.7 K.

This agreement provides a non-trivial consistency check of the measurement and strongly suggests that the extracted xenon gas was close to thermodynamic equilibrium with water ice at a temperature near the LXe operating temperature. Furthermore, the result demonstrates that the newly commissioned HALO system is capable of measuring physically meaningful water concentrations in gaseous xenon at a low-ppb level.

4.4 Discussion and Outlook

The measurements demonstrate that the HALO moisture analyzer can resolve the temporal evolution of water contamination in the Xenoscope gas system over a wide concentration range. The measured concentrations extend from several ppm during the initial nitrogen and argon flushing measurements down to the low-ppb range during stable xenon operation.

The nitrogen and argon measurements both showed a strong initial decrease in water concentration, followed by a slower evolution and a renewed increase at later times. The initial decline is consistent with the replacement of humid gas by dry flushing gas. The slower component can be attributed to the water desorption from surfaces and dead volumes. The late-time increase of the concentration is an effect of reduced gas-flow due to depletion of the gas bottles, which results in a higher concentration of desorped water in the gas. To better control for the latter point, a dedicated low-range flow controller was acquired and will be integrated in the near future. This will enable to set more reproducible flow conditions during the measurement.

The double-exponential model reproduces the overall dry-down behaviour well, with $R^2 > 0.99$, indicating that more than 99 % of the observed variation in the measured water concentration is described by the fitted model. The fast and slow components may be associated with rapid gas exchange and delayed water release from internal surfaces, respectively. However, the systematic structures visible in the residuals show that the model does not describe all features of the measurements. It should therefore be regarded as an effective parametrization of the observed decline rather than as a complete physical model of the gas system. In particular, the fitted time constants depend not only on water removal and desorption, but also on the gas flow, sampling volume, mixing, valve configuration, and changing boundary conditions during the measurements.

The first three xenon measurements reached reproducible fitted asymptotic water concentrations between approximately 21 and 23 ppb. Their unweighted mean is

$$\bar{c}_{H_2O,\infty,GXe} = (22.2 \pm 1.0) \text{ ppb}, \quad (52)$$

where the uncertainty represents the run-to-run standard deviation and therefore quantifies the reproducibility of the three measurements. A separate systematic uncertainty is not included, since the manufacturer does not specify an accuracy for water measurements with xenon as the carrier gas. The corresponding slow time constants lie between 7.7 min and 9.8 min, indicating a reproducible approach to the low-ppb water concentration under stable operating conditions.

The fourth xenon measurement behaved differently and yielded an asymptotic water concentration of (109 ± 2) ppb. Although the mathematical fit quality was high, the measurement was affected by a pump problem and was therefore excluded from the determination of the representative water concentration in xenon. This illustrates that a good fit alone does not guarantee that the underlying physical conditions were stable. The fitted value $c_{H_2O,\infty}$ must therefore always be interpreted together with the operating conditions of the gas system.

A particularly significant result is the close agreement between the measured water concentration and the expected saturated water concentration derived from the water-ice vapour pressure. At the operating temperature of 178.7 K, the Murphy–Koop parametrization predicts a water-ice vapour pressure of 4.21 mPa, corresponding to a saturated gas-phase concentration of 20.3 ppb at 2.08 bar. The water concentration measured by the HALO monitor, (22.2 ± 1.0) ppb, corresponds to a partial pressure of (4.62 ± 0.21) mPa. The relative difference between the measured and predicted values is approximately 10 %.

This agreement is noteworthy because the theoretical concentration is derived from the operating temperature and pressure, whereas the experimental value is obtained independently from the HALO measurement. It therefore provides a non-trivial consistency check of the measurement. The result suggests that the xenon gas extracted at the sampling location was close to thermodynamic equilibrium with water ice at a temperature near the LXe operating temperature.

The comparison further demonstrates that the newly commissioned HALO system is capable of measuring trace water concentrations in gaseous xenon at the tens-of-ppb level with physically plausible results. For the propagation of the $S1$ scintillation light, however, the relevant quantity is the concentration of H_2O dissolved in the liquid xenon. Water impurities can absorb VUV scintillation photons and thereby shorten the optical attenuation length, reducing the amount of $S1$ light reaching the photosensors. The measured value of approximately 22 ppb refers only to the sampled gas phase and therefore cannot be directly interpreted as the water concentration in the bulk liquid xenon. A quantitative estimate of the resulting reduction in $S1$ light would require knowledge of the partitioning of water between the liquid and gas phases as well as its absorption in liquid xenon.

Overall, the measurements show that the HALO integration provides valuable information on water contamination, gas-system conditioning, and changing operating conditions. Future measurements with simultaneous gas-flow monitoring would allow the dry-down behaviour to be interpreted more quantitatively. Estimating the H_2O concentration in liquid xenon from the measured GXe values would require a reliable phase-equilibrium model. A future upgrade therefore foresees a dedicated HALO sampling line probing the main circulation flow more directly.

A measurement immediately downstream of the purifier could probe both the getter performance and the lower measurement range of the HALO analyzer. Comparison with an upstream measurement would additionally allow the water-removal efficiency of the purification system to be quantified.

5 Conclusion

The goal of this thesis was to integrate a HALO moisture analyzer into the Xenoscope gas system and to use it to monitor trace water concentrations, with particular focus on gaseous xenon. Water is an important impurity in liquid-xenon detectors because it can absorb xenon scintillation light in the VUV region and thereby reduce the light-collection efficiency. Monitoring water at low concentrations is therefore relevant for the operation and development of large liquid-xenon detector systems.

A major part of this work concerned the integration of the HALO analyzer into the existing Xenoscope gas-handling system. For this purpose, a dedicated stainless-steel sampling line and valve panel were designed, constructed, cleaned, and installed. Particular attention was paid to the cleanliness of the internal surfaces and the leak tightness of the connections, since water desorption or small atmospheric leaks could significantly affect measurements at the ppb level. The completed sampling line reached a final leak rate below approximately 4.4×10^{-10} mbar L s⁻¹. An evacuated recovery vessel was additionally implemented to establish the pressure difference required for flow through the analyzer while allowing the sampled xenon to be recovered. The HALO system was successfully commissioned and subsequently operated during measurements with nitrogen, argon, and gaseous xenon.

The nitrogen (purity grade 5.0) and argon (purity grade 6.0) measurements were used to study the conditioning and dry-down behaviour of the sampling system. During flushing, the water concentration decreased from the ppm range to minimum values of 446 ppb for nitrogen and 259 ppb for argon, corresponding to reductions of 94.7 % and 93.5 %, respectively. These values are found to be consistent with the minimum purity specifications given in the gas data sheets of < 10 ppm for nitrogen and < 1 ppm for argon. The declining intervals were reproduced well by double-exponential functions, describing the rapid initial gas exchange followed by a slower contribution that is attributed to the water release from surfaces or less efficiently flushed volumes.

The first three xenon measurements reached reproducible fitted asymptotic concentrations between approximately 21 and 23 ppb. Their unweighted mean was

$$\bar{c}_{\infty, GXe} = (22.2 \pm 1.0) \text{ ppb}, \quad (53)$$

where the uncertainty represents the run-to-run standard deviation. The corresponding slow time constants ranged from 7.7 min to 9.8 min, demonstrating a reproducible approach to the low-ppb concentration under stable operating conditions.

The fourth xenon measurement yielded a considerably higher asymptotic concentration of (109 ± 2) ppb. Since this run was affected by the failure of the pump, with a subsequent interruption of the purification and a disruption of detector equilibrium, it was excluded from the determination of the representative xenon concentration. This demonstrates that stable and well-defined pump, flow, and valve conditions are essential for reliable moisture measurements, even when the mathematical fit quality is high.

A particularly important result is the close agreement between the measured water concentration and the expected saturated water concentration at the liquid-xenon operating temperature. At 178.7 K, the Murphy–Koop parametrization predicts a water-ice vapour pressure of 4.21 mPa. At the measured xenon pressure of 2.08 bar, this corresponds to a theoretical saturated gas-phase water concentration of

$$x_{\text{sat}}^{\text{theory}} = 20.3 \text{ ppb.} \quad (54)$$

This independently predicted value is remarkably close to the representative HALO result of (22.2 ± 1.0) ppb. The measured and predicted concentrations differ only by approximately 10 %.

This agreement suggests that the sampled xenon gas was close to thermodynamic equilibrium with water ice at a temperature near the liquid-xenon operating temperature. It also provides an independent consistency check of the newly commissioned HALO setup and demonstrates that the system is capable of measuring physically meaningful water concentrations in gaseous xenon at the tens-of-ppb level.

Overall, the integration developed in this work enables the water concentration in the Xenoscope gas system to be monitored over a range extending from several ppm during initial flushing to approximately 22 ppb during stable xenon operation. The HALO analyzer measures the water concentration only in the sampled gaseous-xenon phase. A direct liquid-phase measurement or a validated phase-equilibrium model would therefore still be required to determine the concentration of water dissolved in the bulk liquid xenon and to quantify its direct influence on VUV-light transport in the detector. Future work should include measurements upstream and downstream of the getter, systematic studies at different flow rates, pressures, and temperatures, and simultaneous monitoring of the relevant operating parameters. A measurement immediately downstream of the getter would provide information about its water-removal performance and would allow for the lower measurement range of the HALO analyzer to be investigated. Integrating the HALO readout into the Xenoscope slow-control system would additionally enable continuous long-term monitoring.

6 Appendix

The Python code and datasets used for the analysis presented in this thesis are publicly available in the following GitHub repository:

`https://github.com/Kevin-Nauer/xenoscope-halo-water-analysis`

The repository contains the Python analysis script, the measurement datasets used for the nitrogen, argon, and gaseous-xenon analyses, and the corresponding fit-diagnostic files. It also includes the code used to generate the comparison between the Murphy–Koop parametrization and the experimental water-ice vapour pressure data reported by Bielska et al. [22].

References

- [1] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: Evidence, candidates and constraints. *Phys. Rept.*, 405:279–390, 2005. doi: 10.1016/j.physrep.2004.08.031.
- [2] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helv. Phys. Acta*, 6: 110–127, 1933. doi: 10.5169/seals-110267.
- [3] N. Aghanim et al. (Planck Collaboration). Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910. Erratum: *Astron. Astrophys.* 652, C4 (2021).
- [4] R. J. Cooke, M. Pettini, and C. C. Steidel. One Percent Determination of the Primordial Deuterium Abundance. *Astrophys. J.*, 855(2):102, 2018. doi: 10.3847/1538-4357/aaab53.
- [5] L. Bergström. Dark Matter Candidates. *New J. Phys.*, 11:105006, 2009. doi: 10.1088/1367-2630/11/10/105006.
- [6] T. Saab. An Introduction to Dark Matter Direct Detection Searches & Techniques. In *Theoretical Advanced Study Institute in Elementary Particle Physics: The Dark Secrets of the Terascale*, pages 711–738, 2013. doi: 10.1142/9789814390163_0011.
- [7] L. Baudis, Y. Biondi, M. Galloway, F. Girard, A. Manfredini, N. McFadden, R. Peres, P. Sanchez-Lucas, and K. Thieme. Design and construction of Xenoscope—a full-scale vertical demonstrator for the DARWIN observatory. *JINST*, 16(08):P08052, 2021. doi: 10.1088/1748-0221/16/08/P08052.
- [8] J. Aalbers et al. (XLZD Collaboration). The XLZD Design Book: towards the next-generation liquid xenon observatory for dark matter and neutrino physics. *Eur. Phys. J. C*, 85(10):1192, 2025. doi: 10.1140/epjc/s10052-025-14810-w.
- [9] V. Chepel and H. Araujo. Liquid noble gas detectors for low energy particle physics. *JINST*, 8:R04001, 2013. doi: 10.1088/1748-0221/8/04/R04001.
- [10] E. Aprile and T. Doke. Liquid Xenon Detectors for Particle Physics and Astrophysics. *Rev. Mod. Phys.*, 82:2053–2097, 2010. doi: 10.1103/RevModPhys.82.2053.
- [11] L. Baudis. Dual-phase xenon time projection chambers for rare-event searches. *Phil. Trans. Roy. Soc. Lond. A*, 382(2266):20230083, 2023. doi: 10.1098/rsta.2023.0083.
- [12] D. S. Akerib et al. (LUX Collaboration). Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector. *Phys. Rev. D*, 97(11): 112002, 2018. doi: 10.1103/PhysRevD.97.112002.
- [13] K. Fujii, Y. Endo, Y. Torigoe, S. Nakamura, T. Haruyama, K. Kasami, S. Mihara, K. Saito, S. Sasaki, and H. Tawara. High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region. *Nucl. Instrum. Meth. A*, 795:293–297, 2015. doi: 10.1016/j.nima.2015.05.065.
- [14] K. C. Oliver-Mallory, A. M. Baker, E. Jacquet, T. J. Sumner, and H. M. Araujo. Spectroscopy of VUV luminescence in dual-phase xenon detectors, December 2025.

- [15] D. Cichon, G. Eurin, F. Jörg, T. Marrodán Undagoitia, and N. Rupp. Scintillation decay-time constants for alpha particles and electrons in liquid xenon. *Rev. Sci. Instrum.*, 93(11):113302, 2022. doi: 10.1063/5.0087216.
- [16] E. Aprile, A. E. Bolotnikov, A. L. Bolozdynya, and T. Doke. *Noble Gas Detectors*. Wiley, 2008. ISBN 978-3-527-40597-8, 978-3-527-61002-0. doi: 10.1002/9783527610020.
- [17] A. Baldini et al. Absorption of scintillation light in a 100 l liquid xenon gamma ray detector and expected detector performance. *Nucl. Instrum. Meth. A*, 545:753–764, 2005. doi: 10.1016/j.nima.2005.02.029.
- [18] *HALO KA Trace Gas Analyzer Instruction Manual, M7500-S Series*. Tiger Optics LLC, 2020.
- [19] F. Jörg. Internal communication. Xenoscope Wiki, University of Zurich, 2026. URL https://www.physik.uzh.ch/groups/baudis/xenon/atpwiki/doku.php?id=darwin:xenoscope:halo:ideas_for_flow_problem. Accessed 9 August 2026.
- [20] International Union of Pure and Applied Chemistry. Pressure. IUPAC Compendium of Chemical Terminology, the Gold Book, 2025. URL <https://goldbook.iupac.org/terms/view/P04819>. Version 5.0.0, accessed 14 July 2026.
- [21] D. M. Murphy and T. Koop. Review of the vapour pressures of ice and supercooled water for atmospheric applications. *Quarterly Journal of the Royal Meteorological Society*, 131: 1539–1565, 2005. doi: 10.1256/qj.04.94.
- [22] K. E. Bielska, D. K. Havey, G. E. Scace, D. Lisak, A. H. Harvey, and J. T. Hodges. High-Accuracy Measurements of the Vapor Pressure of Ice Referenced to the Triple Point. *Geophysical Research Letters*, 40(23):6303–6307, 2013. doi: 10.1002/2013GL058474.