Enabling quantum optomechanics for phononic crystals and superfluid helium films

Wei, X.

Part I

Membrane cavity optomechanics

Chapter 2

Superconducting Nanowire Single Photon Detector

One of the essential components of quantum optomechanical experiments is the ability to detect extremely weak optical signals with high efficiency and negligible background noise. In particular, the observation of optomechanically generated photons requires detectors capable of resolving single-photon events with near-unity efficiency while maintaining ultra-low dark count rates. Superconducting nanowire single-photon detectors (SNSPDs) are among the most sensitive photon-counting technologies currently available [43, 44], offering high detection efficiency [45], excellent timing resolution, and exceptionally low background count rates.

The SNSPDs is optimized for operation at 1064 nm were obtained from *Photon Spot*. While these detectors provide a specified system detection efficiency of approximately 94%, achieving such performance in practice requires considerably more than the detector device itself. The overall detection efficiency and background count rate are strongly influenced by the complete optical and cryogenic environment, including fiber coupling, optical filtering, thermal-photon suppression, stray-light shielding, and detector bias conditions. Consequently, the successful implementation of the SNSPD system required extensive installation, integration, calibration, and optimization efforts to realize the detector performance necessary for quantum optomechanical measurements. Beyond optomechanics, SNSPDs have also attracted interest for dark-photon and sub-GeV dark matter searches due to their high efficiency and ultra-low dark count rates [46, 47].

This chapter presents the installation, characterization, and optimization of the SNSPD system. The detector performance is investigated as a function of temperature, bias current, and environmental shielding, with particular emphasis on identifying and suppressing background photon contributions. Master's student Rémi Claessen contributed to parts of the detector characterization presented in this chapter.

2.1 Detection mechanism of SNSPD

Superconducting nanowire single-photon detectors are among the most advanced technologies for single-photon detection, offering exceptional performance in terms of detection efficiency, time resolution, and low noise. This exceptional performance requires operation in the 1 Kelvin range, which is compatible with our experimental infrastructure. The fundamental operating principle of SNSPDs is based on a thin superconducting nanowire, typically made of materials such as niobium nitride (NbN) or tungsten silicide (WSi), patterned into a meander structure that covers the active detection area. These nanowires are on the order of 5 nm thick and approximately 100 nm wide [43, 45, 48]. In addition, optical anti-reflection coatings and a reflecting back mirror have to be applied for optimization for a particular frequency.

A bias current is run through the wire, tuned to be just below its critical current. When a photon is absorbed by the nanowire, it locally disrupts the superconducting state by breaking the Cooper pairs, creating a region of increased resistance which is commonly referred as a "hotspot". The super conducting current starts bypassing this region, which increases the local current density around the hotspot. If the local current exceeds the critical current density, additional Cooper pairs are broken, and a resistive barrier spanning the width of the nanowire rapidly forms. This transition results in a measurable voltage pulse, which signals the detection of a single photon. The detector then returns to the superconducting state by dissipating the excess energy into the thermal bath, typically at a temperature well below the critical temperature of the superconducting material. The recovery time, also known as the dead time, is primarily determined by the thermal relaxation dynamics of the nanowire and is typically on the order of 50 ns.

In our work, we utilize SNSPDs based on amorphous tungsten silicide (WSi), which is known for their high detection efficiency, particularly at infrared wavelengths [45]. Each SNSPD consists of a 20 μm $\times$ 20 μm active area composed of meandering WSi nanowires. These detectors are optimized for operation at 1064 nm and demonstrate excellent performance, including low timing jitter and high timing resolution.

As shown in Figure 2.1, a voltage pulse is measured when the superconducting nanowire absorbs a single photon. τ_r and τ_f are the rise and fall time of a photon detection, which are ~ 8 ns and ~ 50 ns, respectively. The fall time indicates the time needed to restore the superconducting state on the nanowire.

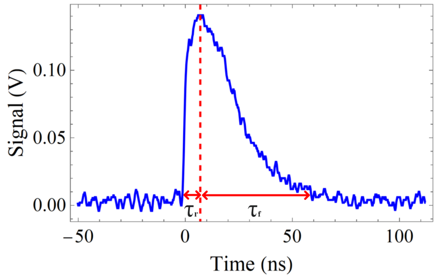

Figure 2.1: Single photon detection signal from the SNSPD as function of time. The typical voltage pulse is measured on an oscilloscope with a low-noise pre-amplifier after the SNSPD. The sharp rise time provides a ~ 1 ns timing resolution of a single photon detection event.

Not all the recorded voltage pulses are the result of actual photon detections. Even in a dark environment without sending in any photons, there can be pulses originating from thermal fluctuations, cosmic rays, or electronic noise. These false positives are collectively referred to as the device dark count rate (DDCR), which is influenced by parameters such as bias current, temperature, and external magnetic fields [48]. Unlike photon-induced events, dark counts cannot be optimized through optical solutions. Although the exact physical model of the origin of the dark counts remains incompletely understood, there is some evidence that vortices in the superconductor are essential to break the superconducting state, leading to a registration of a detection.

Cosmic rays, e.g. high energetic muons, are certainly also contributing to the dark counts. Due to their extreme sensitivity, it has even been suggested that the dark counts on well optically isolated superconducting nanowire detectors could be triggered by dark matter or dark energy. Several studies have proposed using optically shielded SNSPDs as probes of such phenomena [47, 49, 50].

Another important consideration in SNSPD operation is the phenomenon of latching, where the detector fails to reset to the superconducting state after the detection events. In this case, the superconducting nanowire does not recover by sufficiently dissipating the heat deposited during photon absorption. Instead, it relaxes into a stable resistive state and becomes insensitive to incoming photons. Latching is typically observed at high bias currents and at high count rates [43, 51]. Latching can be observed when the series resistor is too large, which acts as a limitation to the SNSPD response bandwidth [52]. Apart from latching, after-pulsing may also occur when the electrical-relaxation and thermal-relaxation time scales are comparable [53].

In our setup, latching becomes observable at bias currents of approximately 16 μA. We obtain a system detection efficiency of 94% at $I_\mathrm{b} \approx 14\,\mu$A, providing a reliable operating regime safe from latching (Figure 2.11).

2.2 Experimental singe photon detection setup

The four SNSPDs investigated in this chapter (WSi nanowire technology from the company *Photon Spot*) have a critical temperature T_c of approximately 3.5 K. To enhance the optical absorption efficiency, a dielectric Bragg mirror is integrated beneath the superconducting nanowires. This reflective layer provides a second chance for the unabsorbed photons to interact with the nanowire, effectively increasing the overall absorption probability [45]. The active area is approximately 20 μm $\times$ 20 μm, optimized for the operation near 1064 nm. The more detailed dimension of the nanowire and the precise structure is not provided from the manufacturer.

To ensure optimal detection performance and minimal dark count rates, all four SNSPDs were mounted onto the 1 K stage of a dilution refrigerator which is well below the superconducting transition temperature. The detectors are thermally mounted to a gold-coated copper base-plate, which serves two purposes: it provides robust thermal contact to the 1 K stage and prevents oxidation of the copper surface, thereby maintaining high thermal conductivity over extended operation. Figure 2.2(a) shows the mounted SNSPD array.

To further suppress background counts from stray light and to avoid mechanical damage, two custom-fabricated copper shields were installed around the detectors, as shown in Figure 2.2(b). These shields also ensure proper optical isolation, particularly around the fiber feedthroughs, which are common sources of scattered photons.

Each detector is coupled to the 1060XP optical fiber to optimize the detection for the 1064 nm laser-frequency used in the experiment. These fibers are spliced to single-mode fibers that are 40 m long, which are coiled onto mandrels. The fiber coils are thermalized to the 3 K stage of the dilution refrigerator. Thermalizing the long fiber coil is crucial for suppressing blackbody photons propagating through the fiber, thereby reducing background count rates. The mandrel diameter is carefully optimized to increase loss for higher-order spatial modes while maintaining single-mode transmission at 1064 nm. This filtering reduces thermal photon background

(a) SNSPD mount.

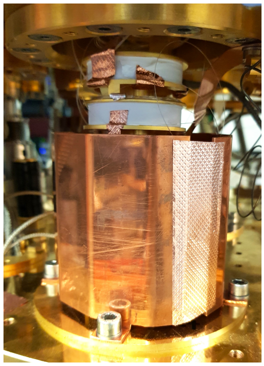

(b) Shield of the four SNSPDs

Figure 2.2: SNSPD inside the dilution refrigerator. (a) The SNSPDs are mounted to a copper plate and then fixed to the 1 K stage. (b) The copper shield is to protect fibers be accidentally touched and to reduce scattered photons from other experimental setups to couple into the detectors.

from upstream optical components, see Figure 2.3(c).

To minimize optical loss, bare fiber splicing is used instead of standard fiber connectors. The typical loss from a splice is $\approx 1\,\%$, whereas connector losses can range from 3–10 % and vary with each reconnection. All optical fibers are routed through rigid metal tubes filled with foam to minimize mechanical vibrations and maintain a low thermal load. To further reduce background light and vibration, a moderate vacuum level can be maintained inside the metal tubes.

The cryostat feedthrough system is shown in Figure 2.3(a–b). Bare fibers pass through a custom Swagelok feedthrough at the room temperature flange. A Teflon insert with four $400\,\mu$m diameter holes is compressed by the Swagelok connector, creating a vacuum-tight seal.

Following the fiber mandrels, each of the four fibers is individually enclosed in a thin Teflon tube for mechanical protection. These shielded fibers are routed through the cryostat feedthrough up to the room temperature plate, as shown in Figure 2.3(a). The same Teflon insert used for vacuum sealing also serves as the exit point for the fibers from the cryostat, as shown in Figure 2.3(b). After exiting the cryostat, the fibers pass through a flexible bellow and enter a metal tube. To minimize background noise and mechanical disturbances, the entire fibers is carefully thermalized and isolated from the external environment.

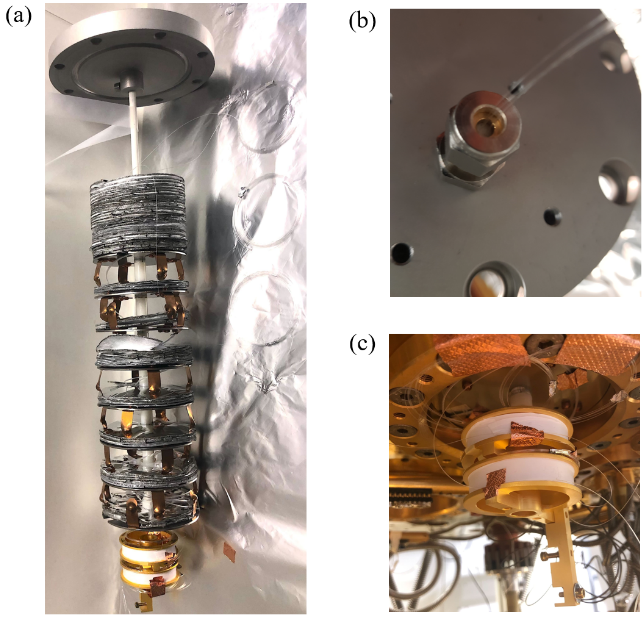

Figure 2.3: Fiber feedthrough. (a) The bare fiber is installed in this cryostat component to get through from room temperature plate to the 1 K plate. (b) A Swagelok feedthrough is installed at the room-temperature plate. A Teflon cylinder with four small holes is placed inside the connector. Tightening the connector compresses the Teflon, forming a vacuum-tight seal around the optical fibers. (c) Fiber mandrel thermalized to 3 K stage, used for spatial mode filtering and background suppression.

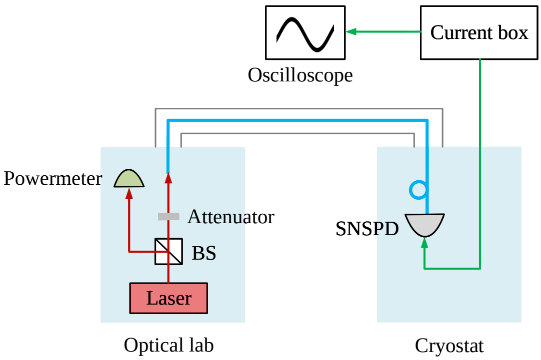

Figure 2.4: Simplified experimental setup diagram for characterizing and optimizing the SNSPDs.

The detectors used in this work were manufactured by the company *Photon Spot*, which labeled the devices as 923–927. To adapt them to our system we relabeled them as DP1, DP3, DP5, and DP7. This labeling convention is used throughout the manuscript. Of the four detectors, three (DP1, DP5, and DP7) are connected via optical fibers. The fiber connection to DP3 was accidentally broken off at the detector interface. Instead of having this one repaired, we decided to completely optically isolate it and use it as an important reference for the lowest dark count rate that can be possibly achieved in our system.

The electrical measurement setup is illustrated in Figure 2.4. In the optical laboratory, the laser beam is split by a beamsplitter (BS). The reflected portion is monitored by a power meter to track power fluctuations and to enable precise calibration of the photon number arriving at the SNSPD. An attenuator is inserted in the transmitted path before the laser is coupled into an optical fiber which is routed through a metal tube into the cryostat.

The bias current for the detector is supplied by a low noise current source. The electrical signal from the detector is extracted using a bias-tee and recorded on an oscilloscope with a $50\,\Omega$ input impedance. To suppress high frequency electrical noise, a ferrite ring is placed immediately next to the oscilloscope input.

2.3 Device dark count rate

The intrinsic noise floor of an SNSPD is commonly characterized by its device dark count rate, defined as the count rate measured in the absence of incident photons. This measurement is performed under conditions of complete optical isolation, ensuring that the observed counts originate from the detector itself rather than from external photon sources. This parameter is essential for evaluating the detector's ultimate noise floor and is primarily determined by the material properties, geometry, and fabrication quality of the SNSPD [43].

The performance of SNSPDs is determined by several critical parameters: the critical current I_c, the latching current I_l and the switching current I_s are often mentioned.

The critical current I_c is the maximum bias current that a superconducting nanowire can carry while remaining in the superconducting state; above this value the device switches to a resistive (normal) state [54]. In operation, the device is typically biased slightly below I_c to maximize detection efficiency while avoiding spontaneous resistive transitions.

The switching current I_s is the experimentally observed bias current at which the nanowire switches from the superconducting to the resistive state during a current ramp (often reported as the mean or most probable value of the switching distribution). In practice, I_s is typically smaller than the critical current I_c, since non-uniformities and external thermal or electromagnetic fluctuations can trigger switching before the intrinsic critical limit is reached.

The latching current I_l is the bias-current threshold for electrothermal latching. Above I_l, positive electrothermal feedback stabilizes a resistive state and prevents the nanowire from self-resetting after an event. Avoiding latching requires an appropriate balance between the nanowire thermal design and the electrical environment (e.g., load impedance), as discussed in Ref. [51]. In a well-designed detector and readout environment, I_l is engineered to lie above the operating bias range, so that electrothermal latching does not limit performance.

By varying the bias current applied to the SNSPDs, the practical switching current I_s is measured. All four detectors are rigidly mounted to the 1 K plate in the dilution refrigerator, ensuring sufficient thermalization and electrical isolation. For the initial DDCR measurements, only the manufacturer-provided 1 m fiber coupled to each SNSPD is used, and no additional fiber sections are spliced onto it. This avoids adding extra fiber length that can increase the optical background (e.g., thermal photons from warmer stages), which would otherwise increase the DDCR.

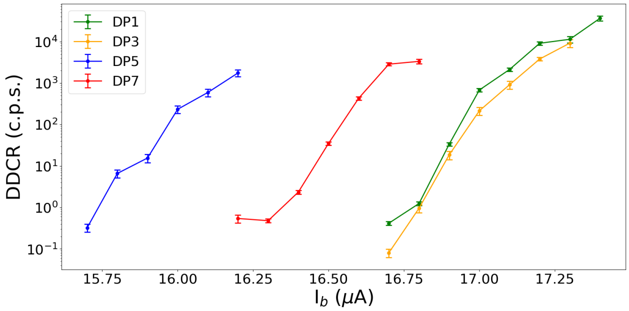

Figure 2.5: Device dark count rate (DDCR) of four SNSPDs measured at 600 mK as a function of bias current. All detectors are optically well isolated, and no additional fibers are connected to the detector.

The DDCR is measured by a triggered oscilloscope which is connected to the detector as shown in Figure 2.4. The bias current is applied to the detector via a low-noise current source. To reduce clicks from environmental electrical noise, ferrite rings are applied on the measuring cable of each channel.

As shown in Figure 2.5, the four detectors exhibit similar DDCR response curves but differ in their switching currents. Detector DP1 and DP3 show very close switching current which are also the highest. Detector DP5 consistently exhibits the lowest switching current and tends to latch more frequently. These variations can be attributed to minor difference in the nanowire fabrication process or material inhomogeneity. Despite these differences among all four detectors, the setup achieves a baseline DDCR on the order of $\sim 10^{-1}$ counts per second (c.p.s.), which represents a lower limit in our measurement configuration.

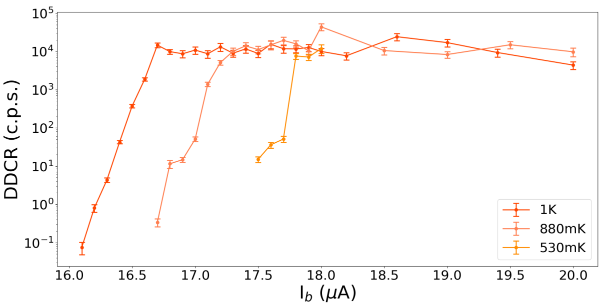

Figure 2.6: Device dark count rate (DDCR) of the optically isolated SNSPD DP3 measured as a function of bias current at 530 mK, 880 mK, and 1 K.

To further investigate the thermal dependence of DDCR, detector DP3 was selected due to its optical isolation—no fiber is attached—and mounted on the 50 mK stage near the mixing chamber of the dilution refrigerator, see Figure 2.7. The detector was additionally enclosed with copper shielding tape to block external photons and to ensure thermal equilibrium with the base temperature.

The reference dark count rate measured from DP3 is particularly valuable, as optical fibers can introduce additional counts through various mechanisms, including thermally generated photons and those induced by acoustic vibrations. Since DP3 is operated without any attached optical fiber, it provides a baseline measurement of the lowest achievable device dark count rate, serving as a reference for evaluating the performance of the three fiber-coupled SNSPDs.

As shown in Figure 2.6, DP3 shows different switching currents when temperature varies. The suggested operation temperature by the manufacturer is 1 K. With the measured data we can conclude that the lower the temperature, the larger working current range can be applied. It is helpful to lower the environmental temperature to below the switching temperature. The DDCR saturates at the order of 10^4 c.p.s. This limit comes from the design and the intrinsic properties of the SNSPDs.

To explore the DDCR temperature dependency, the DDCR is measured at different temperatures and at a higher bias current close to its expected critical current, see Figure 2.6. The DDCR increases rapidly when the bias current is getting closer to the critical current, because the probability of creating the resistive hotspot increases the closer the system operates to the critical current. The critical current also increases when the temperature drops, which agrees with the three curves displayed in the Figure 2.6.

At 1 K, the DDCR reaches the maximum of $\sim 10^4$ counts per second at 16.6 μA and the DDCR plateaued above this bias current. The plateau is not due to the oscilloscope/counter resolution or settings. One possible reason can be that the cooling power of the refrigerator is too small to cool the heated up non-superconducting nanowire after an event is triggered. However, this reason is not valid because the measured photon count rate in our system is up to $\sim 1.6 \times 10^7$ c.p.s. without significant temperature increase as shown in Figure 2.1. There is not much discussion so far about what causes the plateau of DDCR as the function of bias current [55].

Figure 2.7 shows the mounting configuration of the fiber-free reference detector DP3. The device is installed on the 50 mK plate, allowing operation at the base temperature of the refrigerator. With no optical fiber attached, DP3 provides a baseline DDCR without any contribution from fiber coupled thermal photons. To further sup-

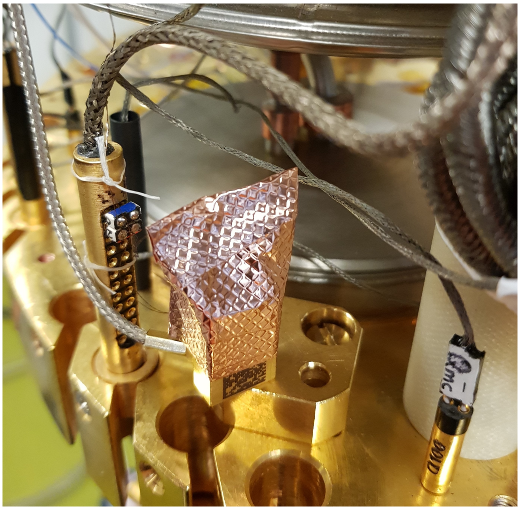

Figure 2.7: Photograph of the fiber-free SNSPD DP3 as mounted in the dilution refrigerator. DP3 is installed on the 50 mK plate and shielded with copper tape in the dilution refrigerator. This detector provides a reference for the lowest achievable count rate when compared with the three fiber-coupled SNSPDs.

press background light, the detector is additionally shielded with copper tape as a final light-tight enclosure.

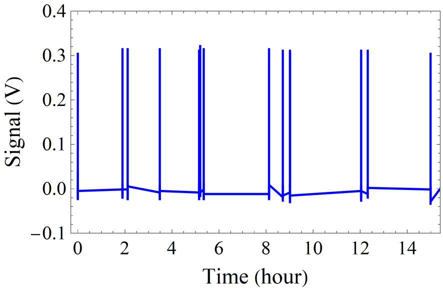

Figure 2.8: Device dark count measurement of a well-isolated SNSPD (DP3) operated in a dark configuration (optical input blocked, no fiber attached). Over 15 h at $I_\mathrm{b} = 12\,\mu\mathrm{A}$, 14 counts were recorded. During this cool-down the 50 mK stage was operated at 633 mK.

All recorded click segments are saved and shown in Figure 2.8. During a continuous 15 h overnight run, 14 clicks were observed at a bias current of $12\,\mu\mathrm{A}$ and an operating temperature of 633 mK, corresponding to an average count rate of $0.93\,\mathrm{h}^{-1}$. Under these conditions the recorded events provide a baseline for the detector-intrinsic background (including the readout chain) in the absence of any fiber-coupled optical input. Note that for this particular cool-down the 50 mK stage was operated only at 633 mK.

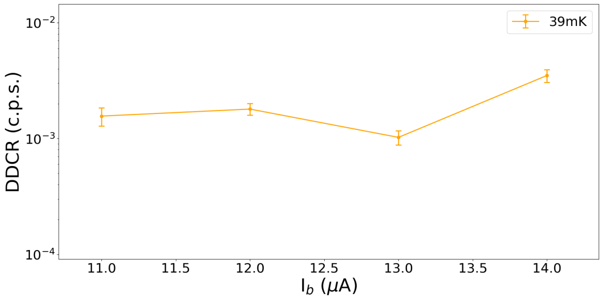

Figure 2.9: Device dark count rate as function of bias current on a well isolated SNSPD (DP3) operating at a base temperature of 39 mK.

Figure 2.9 shows the measured device dark count rate of DP3 as a function of bias current under otherwise identical operating conditions. All measurements were performed at a base temperature of 39 mK. To probe the ultra-low count-rate regime, we carried out multiple overnight measurements at reduced bias currents. Even at the lowest bias current investigated, the DDCR does not vanish but remains finite, corresponding to approximately one detection event every few thousand seconds.

DP3 is operated in a fully shielded configuration, and the observed counts therefore characterize the intrinsic lower bound of the DDCR at this temperature. Within the temperature range explored in this work, the DDCR exhibits a clear dependence on the applied bias current. However, the observed increase with bias current does not follow a simple exponential behavior. This deviation suggests that the physical mechanism responsible for triggering detection events in DP3 may differ from the dominant sources contributing to the general DDCR observed in other devices. This distinction is further illustrated by the bias-current dependence of the DDCR shown in Figure 2.6.

Notably, in the bias current range between 11 μA and 13 μA, the DDCR remains approximately constant within the experimental resolution. At the lowest measured level around 10^{-3} cps, it cannot be completely excluded that a fraction of the observed events is triggered by residual background photons, even though the detector is effectively isolated from ambient optical and thermal radiation. In general, background-photon-induced counts are expected to decrease with decreasing bias current due to the concomitant reduction in detection efficiency.

Consistent with this expectation, an additional overnight measurement performed at a bias current of $10\,\mu$A yielded no detection events over the entire acquisition time. A more definitive assessment of the ultimate DDCR limit would require operation of the detector in an environment with even stronger shielding against residual electromagnetic radiation, which could further suppress spurious photon-induced events.

2.4 Detection efficiency

The system dark count rate (SDCR) is defined as the total dark count rate of an SNSPD after it is integrated into the full experimental setup, including all fiber links and optical components. In practice, the SDCR is often limited by background photons coupled from the environment into the measurement fibers, as well as thermal photons generated within the fibers and all elements along the optical path. Since these counts are indistinguishable from photon-triggered detection events, a high SDCR directly degrades the signal-to-noise ratio and therefore impacts the practically achievable system detection efficiency (SDE), particularly at low detected count rates.

In the SDE measurements, the detectors are operated under optical input, and the measured click rate includes contributions from both photon-induced events and system backgrounds. The detector performance is characterized by the system detection efficiency, defined as the ratio between the measured click rate and the calibrated incident photon rate. The SDCR is primarily determined by the quality of optical shielding and fiber isolation in the system, and we quantify this effect for different shielding conditions in Figure 2.10.

Figure 2.10 shows the count rates of four detectors measured at a bias current of $14\,\mu$A under different shielding conditions. The highest count rates are observed when the optical fibers are left unshielded from the dilution refrigerator to the laboratory environment. In this configuration, background light coupled into the fibers constitutes the dominant source of photon-triggered detection events, highlighting the importance of maintaining the entire optical path in a dark and well-shielded environment.

Among the four devices, DP1 and DP7 are connected to fiber mandrels, while DP3 is not. Under the shielding conditions explored here, their measured count rates are comparable, and no clear reduction associated with the presence of a mandrel is observed. Fiber mandrels are generally introduced to suppress thermal photons propagating along the fibers. However, as shown in Figure 2.5 and Figure 2.11, the overall performance and detection efficiency of DP5 are lower than those of the other three SNSPDs. This reduced efficiency may limit the sensitivity of this measurement to changes in the background photon clicks, and could explain the absence of a pro-

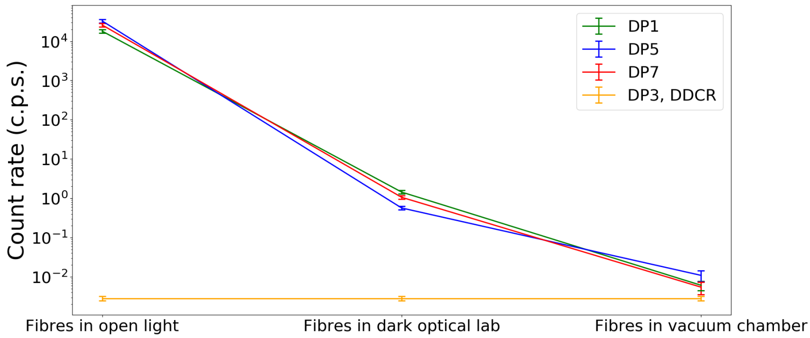

Figure 2.10: Count rates of four SNSPDs measured at a bias current of 14 μA under different optical shielding conditions of the fibers. DP3 was operated at a base temperature of 39 mK, while the other three SNSPDs were installed on the 1 K plate and operated at 500 mK.

nounced mandrel-related effect in this dataset. A more controlled assessment of the mandrel performance could be achieved by directly comparing otherwise identical detectors with and without mandrels, for example by using DP1 and DP7 as a compare pair.

Figure 2.11 shows the system detection efficiency (SDE) as a function of bias current for detectors DP1, DP5, and DP7. DP5 and DP7 were measured at two incident optical powers, 50 nW and 210 nW, while DP1 was measured at 210 nW, with the incident powers measured at the optical powermeter port shown in Figure 2.4.

All three SNSPDs exhibit a similar dependence of the SDE on the bias current. In the range from 8 to 12 μA, the SDE increases with increasing bias current and reaches its maximum at around 14 μA. In the bias current interval between 14 and 15 μA, the SDE remains close to its maximum value, while the corresponding device dark count rate reaches its minimum for this three detectors (see Figure 2.5). Taking all measurements into account, we identify 14-15 μA as the optimal operating bias current range for these detectors at temperatures around 500 mK.

For DP5 and DP7, the SDE curves measured at 50 nW and 210 nW overlap (see legend), indicating that the extracted SDE is consistent between these two optical input settings. In contrast, DP5 and DP7 exhibit different DDCR levels (see Figure 2.5), reflecting different baseline device characteristics. The agreement of the SDE curves at the two power settings suggests that these optical inputs lie within a regime where the extracted SDE is not affected by the applied optical power. Furthermore, the

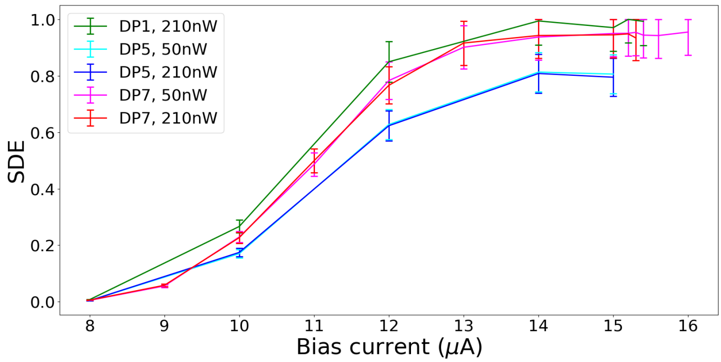

Figure 2.11: System detection efficiency as a function of bias current for detectors DP1, DP5, and DP7. DP5 and DP7 were measured at incident optical powers of 50 nW and 210 nW, while DP1 was measured at 210 nW, as measured at the optical power monitor port shown in Figure 2.4.

bias-current dependence of the SDE is similar for DP1, DP5, and DP7, even though their DDCR levels differ, indicating that DDCR variations mainly shift the absolute SDE level rather than changing its overall bias-current dependence.

Figure 2.12 shows the SDE measured while increasing the incident optical power at a fixed bias current of 14 μA. For convenience, the incident power is represented on the horizontal axis as the corresponding expected number of clicks. The SDE increases as the expected clicks rise from roughly 10^3 to 10^4. It then reaches its maximum level and stays near this value up to about 10^6 expected clicks. For expected clicks above this level, the SDE decreases. These observations identify a broad operating window in which the detector maintains near-maximal efficiency before performance degrades at the highest count rates.

The decrease in SDE at the highest optical powers can be explained by practical limitations of the experimental setup. At high photon arrival rates, the detector recovery time is a limiting factor. Photons that arrive during the recovery time after a detection event cannot be registered, which reduces the SDE. In addition, higher optical input power adds heat to the detector, and the finite cooling power of the dilution refrigerator stage can then lead to a higher local temperature. This shifts the detector operating point and lowers the system detection efficiency. Taken together, these considerations define an optimal range of photon count rates for each detector in the present setup.

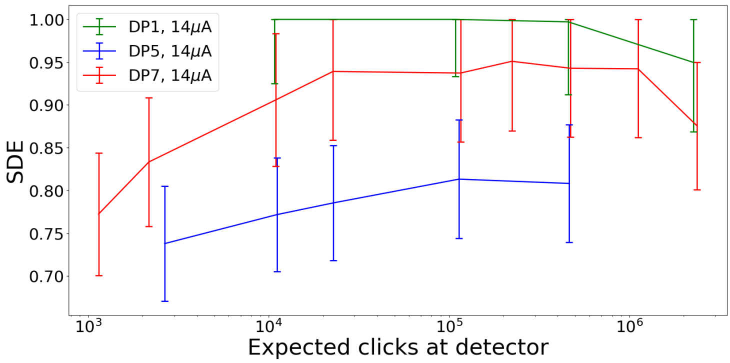

Figure 2.12: System detection efficiency as a function of incident optical power for detectors DP1, DP5, and DP7, measured at a fixed bias current of 14 μA. The incident optical power is measured at the optical power monitor port.

2.5 Conclusions

In this chapter, the SNSPD system was installed, characterized, and optimized for operation at 1064 nm. The detector performance was characterized through measurements of the DDCR and the SDE under controlled operating conditions. Particular attention was paid to suppressing background photon counts through optimization of the optical coupling, detector environment, and readout electronics. As a result, a system detection efficiency of 94% was achieved together with ultra-low dark count rates (about 1 count per 200 seconds) at a bias current of 14 μA. For a completely optically isolated detector, the dark count rate was as low as 1 count per hour.

The observed dependence of the SDE on optical power highlights practical limitations associated with detector recovery time and thermal loading at high photon flux. Importantly, these effects allow a well-defined and stable operating window to be identified, within which the detector performance is both efficient and reliable.

The planned experiments on investigating an optomechanical system in the quantum regime will require first single photon detection to project the mechanical state in a particular phonon state and subsequently will require single photon detection to probe the time-evolved state (on the time scale of the mechanical decoherence time). Therefore, at least the square of the overall detection efficiency is essential. Furthermore, each experimental run will have to start with optical cooling to near the quantum ground state, and the locking of 5 optical cavities is required. Together, this implies that a single experimental run will require on the order of 1-3 second (on the same order of the ring down time of our mechanical membrane structures). In order to collect enough data to make statistically meaningful statements for the quantum system, we need on the order of 1000 successful experimental runs. With an overall detection efficiency (after passing all the 5 cavity systems and corresponding optics for locking and mode matching, and taking the detector efficiencies into account) of say 10% (which is quite typical for such a complex optical system), we would need on the order of 100 hours to perform a single experiment. With 94% efficient single photon detectors and with special optics we aim for an overall single photon detection efficiency of around 60% to 90%, reducing the time for a single experiment to a couple of hours, which is much more manageable than week long experiments.