

A Brute-Force Approach to Fiber-Optic Sensors: Achieving High-Precision Results with Low-Resolution SLED-Based Spectrometers

LAKSH BHASIN¹ AND PATRICK RALL¹ - SUMMER 2013

¹*Institut für Robotik und Mechatronik, DLR Oberpfaffenhofen, Münchner Straße 20, 82234 Weßling, Germany*

Abstract

With the increasing importance of robot-assisted, minimally-invasive medical surgery, the ability to precisely measure physical parameters with fiber-optic sensors has become a significant objective in medical technology. In particular, the currently achievable pixel resolution of diffraction-grating spectrometers is only in the range of several tens (if not hundreds) of pm, while many fiber Bragg grating (FBG) sensors have linear temperature coefficients of $\sim 5 \frac{\text{pm}}{\text{K}}$ and strain coefficients of $\sim 1 \frac{\text{pm}}{\mu\text{m}\cdot\text{m}^{-1}}$ [1]. Thus, it is often necessary to use subpixel detection techniques (e.g. Gaussian fits, centroid detection) on FBG spectra to determine their Bragg wavelength, and thereby ascertain information about physical parameters. Unfortunately, as we discuss, these techniques do not produce linear regression models with sufficiently low residuals of parameters like temperature or strain. In a novel approach, we employ a brute-force algorithm that seeks to replicate a spectrometer’s measurement response to FBG spectra. Using previously-collected high-resolution optical spectrum analyzer (OSA) data for a given Bragg wavelength, as well as spectrometer-specific measurements, we implement a process of “profile” generation to produce a bin-by-bin estimate of the spectrometer’s expected output. By fitting these pre-generated profiles to low-resolution real spectrometer measurements, we can achieve higher levels of precision in Bragg wavelength measurement, better linear regression models, and significantly lower residuals compared to existing methods such as the LPO4 (linear phase operator) technique described in [1]. Moreover, while we primarily investigate one specific type of spectrometer (manufactured with optics from Ibsen Photonics and custom-made electronics), the technique can be adapted to other spectrometers by taking into account their light source and their bin-by-bin response to a tunable laser.

I. INTRODUCTION

As fiber-optic sensors become more valuable in medical surgery and other fields due to their small size, manipulability, versatility, and immunity to electromagnetic interference, it is increasingly important to develop systems with high levels of measurement precision. A particularly important type of sensor in the field of fiber optics is the Fiber Bragg Grating (FBG), which can respond to applied strains (ϵ) as well as changes in temperature (ΔT). These gratings have, over the past few decades, emerged as one of the world’s most widely used fiber-optic devices due to their linear response to various stimuli.

FBGs consist of periodic perturbations of a fiber core’s refractive index along its length (see Figure 1). These variations are formed by exposing the core to intense ultraviolet light through a periodically corrugated phase mask (approximating a square wave in profile), which is usually constructed from silica glass. When this ultraviolet light comes in contact with the optical fiber, the fiber’s refractive index is changed almost permanently at wavelengths much higher than that of the irradiating light; this phenomenon is known as “photosensitivity” [2]. While the mechanisms behind photosensitivity are not completely understood, the refractive index change ($\Delta n \equiv n_2 - n_1$) has been shown to be correlated with the structural rearrangement of the glass matrix [3] and the photoinduced paramagnetism of donor-vacancy pairs (particularly with germanium-doped cores) [4].

The index modulation of FBGs is interesting due to the reflective properties that result from it. At a certain FBG-dependent wavelength λ_B - known as the Bragg wavelength - each reflection from a crest in the index variation is in phase with the next crest coming in at λ_B . As a result, a very narrow band of the incident optical field is reflected




Figure 1: Variation of refractive index between n_1 and n_2 as a function of z (the fiber axis) for a non-apodized, uniform FBG without chirping (i.e. where the period of index modulation is constant).

by coherent scattering, with the strongest mode-coupling interaction occurring at the wavelength λ_B . This reflectivity response is shown in Figure 2, and features a distinct peak that occurs at the Bragg wavelength. For the specific case of a non-apodized, non-chirped grating (with the assumption of two-mode coupling), it is a well-known result in the literature that the power reflection coefficient at the incident end of an FBG is given by [1]:

$$R = \frac{\kappa^2 \sinh^2 (L\sqrt{\kappa^2 - \hat{\sigma}^2})}{\kappa^2 \cosh^2 (L\sqrt{\kappa^2 - \hat{\sigma}^2}) - \hat{\sigma}^2} \quad (1)$$

Where L is the FBG’s length and κ and $\hat{\sigma}$ are inversely λ -dependent coupling coefficients for the AC component and common modes respectively. In practice, however, we do not generally observe as symmetrical or trigonometrically well-behaved a response as the theory suggests; thus, an FBG’s reflectivity spectrum cannot easily be described analytically.

The Bragg wavelength is particularly useful since it responds linearly to various external stimuli such as temperature and strain. This response manifests itself as a shift in the reflectivity spectrum, and is entirely independent of optical intensity (when measuring reflected power as a percentage




Figure 2: An example of various reflectivity responses for the same FBG, collected at different temperatures through a narrow-bandwidth (5-pm resolution) optical spectrum analyzer. In this case, we see a Kaiser-like peak with asymmetrical sidelobes. The shift in Bragg wavelength over this 40.0°C range is ~ 1.2 nm.

of incident power). In Figure 2, we see, however, that this response is relatively minuscule; with a 40.0°C increase in temperature, we only observe a λ_B shift of ~ 1.2 nm.

As many spectrometers are incapable of producing a wavelength resolution of < 10 pm while still delivering data at a sufficiently frequent rate, it is still very challenging to extract high-precision information from FBGs. Given the difficulties in describing FBGs analytically, several alternative subpixel detection algorithms (SDAs) have been developed to determine Bragg wavelengths from the bins in spectrometer data. These include:

1. A Gaussian fit as an approximation for the FBG’s overall shape (including the sidelobes). Earlier, the Gaussian was assumed to be the actual shape of FBG spectra due to apodizations in the data caused by lower laser coherence lengths (and hence decreased fringe contrast).
2. A Kaiser window function fit, defined in Section VIII, taking into account the fact that actual spectra exhibit a shorter decay factor compared to the Gaussian shape [1].
3. A simple, weighted-average centroid detection algorithm (CDA).
4. The linear phase operator (LPO) technique described in [1]. This algorithm applies a finite impulse response (FIR) filter to points around the peak bin to produce an output sequence, which serves as an approximation to the differentiator of the FBG spectrum. By applying linear interpolation on the intensities around the zero-crossing of this differentiator, we achieve a subpixel estimate of the spectrum’s peak.

However, as we see in Section VIII, these three techniques fail to produce a linear regression model with sufficiently small

standard deviations, and often indicate systematic patterns in their residual plots. Moreover, since spectrometer bins are not always evenly spaced in the wavelength domain, the subpixel approach taken by these methods is not always appropriate. While interpolation between bins can serve as a possible fix to this issue, this can - for instance - significantly change the transfer function used in the LPO technique for approximating the spectrum’s differentiator.

In this article, we instead apply a new approach that effectively emulates how the spectrometer actually collects data for a given FBG. We begin by taking a high-resolution (1-pm) optical spectrum analyzer (OSA) spectrum of our FBG at a given reference temperature. By taking into account the SLED light source of our spectrometer and each bin’s response to a narrow-bandwidth, constant-intensity laser, we can estimate how each of the 80 channels of our spectrometer would respond to this FBG spectrum. This results in a spectrometer “profile” for a given Bragg wavelength (we do not determine the actual λ_B however). We then take our OSA data and shift it (by some multiple of 1 pm) to simulate a change in the Bragg wavelength, and repeat the process of profile generation. By repeating this process over a shift range of ± 1.1 nm, we are usually able to generate enough profiles for a given FBG. We can then take these profiles and compare them against incoming spectrometer data; this in turn allows us to ascertain how much our FBG’s Bragg wavelength has shifted relative to our reference temperature. If we use spectrometer data measured at a known temperature, then we can additionally produce a $\Delta\lambda$ vs. ΔT plot through this method, and can subsequently apply regression to produce a linear model for our FBG. Since we are primarily interested in the shift of the Bragg wavelength, this approach successfully allows us to determine temperature changes and strain without actually determining the Bragg wavelength.

This process succeeds in producing a model with lower standard deviations in temperature prediction than the LPO4 technique described in [1] (previously assumed to provide the best overall resolution gain factor of all SDAs). For instance, with one specific FBG, we were able to surpass LPO4 and demonstrate standard deviations of ~ 0.1 K in our temperature estimation over a range of $T \in [10^\circ\text{C}, 50^\circ\text{C}]$, which is sufficiently low for many uses in medical technology.

The contents of this article are organized as follows. In Section II, we briefly examine the theory of Bragg wavelengths and step through some important assumptions that have been made to deduce the FBG’s linear response. Next, in Section III, we describe the process we used to generate our spectrometer profiles from OSA data. We then elaborate further on our equipment and data collection methods in Section IV, before describing our profile fitting process and results in Sections V and VI respectively.

In order to optimize our fitting algorithm, we also take steps to minimize comparisons and arithmetic operations. Specifically, we use layered searches in place of a single sequential search, and try to make our profiles as small as possible (in terms of number of bins) while still retaining accuracy. Some

noticeable results related to this optimization are discussed in Section VII. Lastly, in Section VIII, we compare our algorithm with existing SDAs and examine the relative benefits and drawbacks of the results we achieve.

II. BRAGG WAVELENGTH THEORY OVERVIEW

In this section we briefly overview the theory behind the FBG's linear response to temperature and strain. Essentially, the periodic perturbations of the refractive index of the fiber core create a band-stop filter for a narrow wavelength band of incident light; this results from coherent scattering due to the index variation. The strongest such interaction occurs at a specific wavelength λ_B , known as the Bragg wavelength. We can describe this effect more precisely with coupled mode theory [1]. The propagation of light in fibers can be described by fiber modes resulting from Maxwell's field equations in a dielectric medium with perturbed permittivities. Let the z axis define the propagation axis along the length of the fiber. Applying coupled mode theory, we assume that the transverse component of the electric field \vec{E}_t can be expressed as a superposition of i modes with amplitude A_i propagating in the positive z direction and amplitude B_i propagating in the negative z direction. If we additionally assume that these oppositely propagating modes have the same real propagation constant β_i , then we have for the transverse electric field along the fiber [1]:

$$\vec{E}_t(x, y, z, t) = \sum_i [A_i(z) \cdot e^{j\beta_i z} \cdot \vec{e}_{it}(x, y) \cdot e^{-j\omega t} + B_i(z) \cdot e^{-j\beta_i z} \cdot \vec{e}_{it}(x, y) \cdot e^{-j\omega t}] \quad (2)$$

Where $\vec{e}_{it}(x, y)$ describes the transverse mode field in the plane normal to the mode's propagation. The actual values for A_i and B_i of a specific mode are linked to the values A_k and B_k of all of the other modes through a transverse coupling coefficient $C_{ki}^t(z)$ (the longitudinal component is assumed to be negligible). To simplify the coupled differential equations that arise in determining A_i and B_i , we assume a uniform grating with constant coupling coefficient over its entire length, and additionally reduce the number of modes to a single dominant transverse mode propagating in the positive and negative directions along the length of the fiber (with propagation constants of β_{01} and $-\beta_{01}$). This mode is designated in literature as the LP₀₁ mode, as it is the dominant mode solution to the three-dimensional scalar wave equation for which the electric field is linearly polarized. Note that we use the scalar wave equation to arrive at an approximate modal field description instead of solving the full set of Maxwell's equations; this is a valid approximation under the weakly guiding condition, where we assume that the difference between the refractive indexes of the core and cladding is small (usually less than 1%).

With the above simplifications, the strongest mode-coupling then occurs at a specific wavelength λ_B where the incident mode is optimally coupled with the counter-propagating mode

and thus reflected; these two modes must meet the phase-matching condition specified by the core's grating period Λ in order to interact coherently, and so we arrive at the simple condition that:

$$\frac{2\pi}{\Lambda} = \beta_{01} - (-\beta_{01}) = 2\beta_{01} \quad (3)$$

In order to satisfy the conditions for a mode in a vacuum, we must have $\beta_{01} = \frac{2\pi}{\lambda_B}$. For an FBG, we instead introduce a term n_{eff} as our "effective" refractive modal index (defined as a frequency-dependent [1] combination of the indices n_1 and n_2 from Figure 1), in order to state $\beta_{01} = \frac{2\pi n_{\text{eff}}}{\lambda_B}$ and thus:

$$\lambda_B = 2n_{\text{eff}}\Lambda \quad (4)$$

Equation (4) signifies the dependence of the Bragg wavelength on changes in fiber properties, since any variation in strain, temperature, or polarization will vary the modal index n_{eff} and grating period Λ , and will thus change the Bragg wavelength. Moreover, we see that an FBG sensor's response to external stimuli is entirely contained in its Bragg wavelength; shifts in the narrow-band reflection spectrum are independent of optical intensity and depend solely on the manufacturing process for the gratings. We can use a linear approximation with (4) to find this change in the Bragg wavelength with respect to strain and temperature.

$$\frac{\Delta\lambda_B}{\lambda_{B0}} = \frac{\Delta\Lambda_B(\epsilon_i; \Delta T)}{\Lambda_{B0}} + \frac{\Delta n_{\text{eff}}(\epsilon_i; \Delta T)}{n_{\text{eff},0}}, \quad i = 1 \dots 6 \quad (5)$$

Where the index 0 signifies the FBG's properties at some arbitrary reference temperature T_0 and strain 0. Since the strain tensor is symmetric for our fiber (oriented along the z -axis as before), the index i is used in accordance with Voigt notation such that $\epsilon_1 = \epsilon_{xx}$, $\epsilon_2 = \epsilon_{yy}$, $\epsilon_3 = \epsilon_{zz}$, $\epsilon_4 = 2 \cdot \epsilon_{yz}$, $\epsilon_5 = 2 \cdot \epsilon_{xz}$, and $\epsilon_6 = 2 \cdot \epsilon_{xy}$.

Since the FBG is oriented along the z axis, a change of the spacing Λ must be affected by the strain ϵ_3 , giving $\epsilon_3 = \frac{\Delta\Lambda_B}{\Lambda_{B0}}$. This strain can be caused by thermal expansion (in which case we have $\epsilon_3 = \alpha \cdot \Delta T$, where α is the linear temperature elongation constant), or by mechanical strain ϵ^m (in which case we have $\epsilon^m \cdot E = \frac{F}{A}$, where A is cross-sectional area, F is applied force, and E is the tensile modulus of the fiber).

The response of the effective modal index to changes in the strain is determined by the Pockel's coefficients p_{11} and p_{12} of the fiber. In the case of pure mechanical strain (without temperature shifts), we find that $\frac{\Delta n_{\text{eff}}}{n_{\text{eff},0}} = \frac{-n_{\text{eff},0}^2}{2} [p_{11}\epsilon_t + p_{12}(\epsilon_3 + \epsilon_t)]$ [2], where ϵ_t represents strains transverse to the z axis. Thus, when purely mechanical strain is applied, we have:

$$\frac{\Delta\lambda_B}{\lambda_{B0}} = \epsilon_3 - \frac{n_{\text{eff},0}^2}{2} [p_{11}\epsilon_t + p_{12}(\epsilon_3 + \epsilon_t)] \quad (6)$$

We now assume that the strain is axial and isotropic. In order to account for lateral contraction due to the axial strain ϵ_3 , we introduce the Poisson ratio ν such that $\epsilon_t = -\nu\epsilon_3$, giv-




Figure 3: The response of λ_B to various pressures for an FBG coated with a thick Hysol jacket, adapted from [2].

ing $\frac{\Delta\lambda_B}{\lambda_{B0}} = \epsilon_3 \left(1 - \frac{n_{\text{eff},0}^2}{2} [p_{12} - \nu \cdot (p_{11} + p_{12})] \right)$. Introducing the effective photoelastic coefficient p_{eff} [1], we arrive at the desired linear result:

$$\frac{\Delta\lambda_B}{\lambda_{B0}} = (1 - p_{\text{eff}}) \epsilon_3 \quad (7)$$

In practice, this strain response is indeed linear for temperatures as high as 370°C [5]; this is demonstrated in Figure 3.

Lastly, the response of the effective modal index to changes in temperature (in the absence of mechanical strain) depends on the thermo-optic coefficient $\alpha_n = \frac{1}{n_{\text{eff},0}} \cdot \frac{dn_{\text{eff}}}{dT}$. Substituting in $\frac{\Delta n_{\text{eff}}}{n_{\text{eff},0}} = \alpha_n \Delta T$ and using the fact that $\frac{\Delta\lambda_B}{\lambda_{B0}} = \alpha \Delta T$ (where α is the thermo-elastic coefficient as found above), we arrive at a linear result for pure temperature shifts:

$$\frac{\Delta\lambda_B}{\lambda_{B0}} = (\alpha + \alpha_n) \Delta T \quad (8)$$

However, unlike the case with mechanical strain, this sensitivity response does become slightly nonlinear at higher temperatures [1]. In these cases, a quadratic polynomial approximation can work well instead.

Overall, the theory of Bragg wavelengths, assuming our abovementioned simplifications, indicates that we expect a linear response in measuring λ_B as a function of strain or temperature. Thus, our algorithms should also reflect this with no systematic pattern in our residuals. Nonetheless, since we are exclusively looking at temperature responses in this paper, we do expect to see a slight U-shaped pattern in our residuals, which would indicate that a quadratic fit might be slightly more suitable.

III. PROFILE GENERATION PROCESS

To generate spectrometer profiles for a given FBG with respect to a given reference temperature, we must account for all of the factors that go into the spectrometer’s measurement process. In order to do this, we use the following steps, which are described in further detail in subsections I through III:

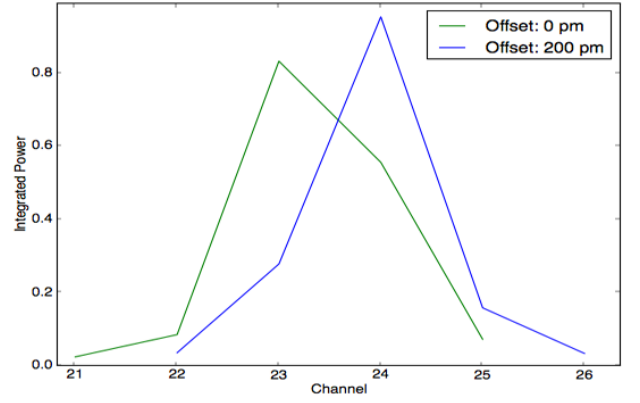


Figure 4: Emulated spectrometer profiles generated for FBG 8 (see Section IV) with a reference temperature of 25°C corresponding to offsets of 0 pm and 200 pm.

1. Collect high-resolution (picometer-precision) OSA data for a given FBG and reference temperature over a sufficiently large wavelength domain of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$.
2. To investigate how each bin of the spectrometer responds to a constant-intensity beam of light at a given wavelength, collect the “sweep factor” data described in subsection I.
 - This must span over a wavelength domain of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ and should have the same resolution as the OSA data (interpolation may be necessary).
 - We refer to the i th bin’s sweep factor as $P_{\text{sweep},i}(\lambda)$.
3. Furthermore, to account for the uneven optical density produced by our superluminescent light-emitting diode (SLED) source (see subsection II), we measure its broadband optical spectrum.
 - Once again, this must span over a wavelength domain of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ and should have the same resolution as the OSA data (interpolation may be necessary).
 - We refer to this SLED spectrum as $P_{\text{SLED}}(\lambda)$.
4. Apply a Gaussian fit on the ten highest points of the OSA data. This will allow us to determine a ballpark estimate of the Bragg wavelength λ_B (in nm) at this temperature.
 - Note that we do not attach too much importance to the actual value of this Bragg wavelength; instead, we simply use it to determine relative shifts ($\Delta\lambda$) and give each spectrometer profile a somewhat accurate absolute label for the Bragg wavelength.
 - We also make use of this Bragg wavelength to determine which bins the profile will be centered around (see subsection III).
 - Overall, the choice of the Gaussian fit was arbitrary.
5. Shift the OSA data by x pixels (in any direction) to simulate a shift in the Bragg wavelength. Since the data

were collected with picometer precision, this gives us a new Bragg wavelength of $\lambda_{B,new} = \lambda_B + \frac{x}{1000}$ nm.

- We refer to this shifted OSA spectrum as $P_{\text{shifted},x}(\lambda)$.
- Using the bin boundaries of the spectrometer, determine the bin number that will contain $\lambda_{B,new}$. The process of determining the middle 78 bin boundaries of our spectrometer is described in subsection III.
 - To generate a profile of width w (where w is an odd integer), consider the $\frac{w-1}{2}$ bins on either side of the abovementioned central channel. Let N be the set of all of these bin numbers.
 - The “integrated profile value” corresponding to the n th bin for a shift of x pm is then given by the following integral:

$$\text{Profile}(n, x) = \int_{\lambda} P_{\text{shifted},x}(\lambda) \cdot P_{\text{sweep},n}(\lambda) \cdot P_{\text{SLED}}(\lambda) d\lambda \quad (9)$$

- Where λ represents the largest possible wavelength domain that is covered by all three of the above power functions. This is no longer simply $[\lambda_{\text{min}}, \lambda_{\text{max}}]$ because of data shifting.
 - Calculate (9) for all $n \in N$ to generate this profile.
- Lastly, repeat steps 5 through 8 for $x \in [-1100, 1100]$ pm with a 1-pm step size. For most cases, this range will cover all profiles of interest.

Two example profiles generated through this process are shown in Figure 4. Qualitatively, the shift in λ_B between these two profiles is apparent.

For the remainder of this section, we will discuss further details about the sweep factors, SLED light source, and bin boundary data used in the above procedure. The procedure used to fit profiles with actual (low-resolution) spectrometer data is presented in Section V.

I. “SWEEP FACTOR” GENERATION

The spectrometer we use in this article was manufactured with 80 bins of uneven width in the wavelength domain. In order to emulate how the spectrometer would react to a given FBG spectrum, we first observed how each channel responded as we swept over a range of $\lambda \in [1528, 1562]$ nm with a constant-intensity 1-pm bandwidth tunable laser. This resulted in the generation of several “sweep factors” for each channel, as demonstrated in Figure 5 for the particular cases of channels 0 and 1.

While we expected these sweep factors to be smooth, we observed a constant jitter in our results, as demonstrated by the jagged edges of the curves in Figure 5. Since this jitter was present at regular intervals of ~ 10 pm in our data, we hypothesized that our tunable laser had systematic flaws that




Figure 5: Picometer-resolution sweep factors for channels 0 (black) and 1 (blue) collected with a tunable laser. Note that these data are unfiltered.

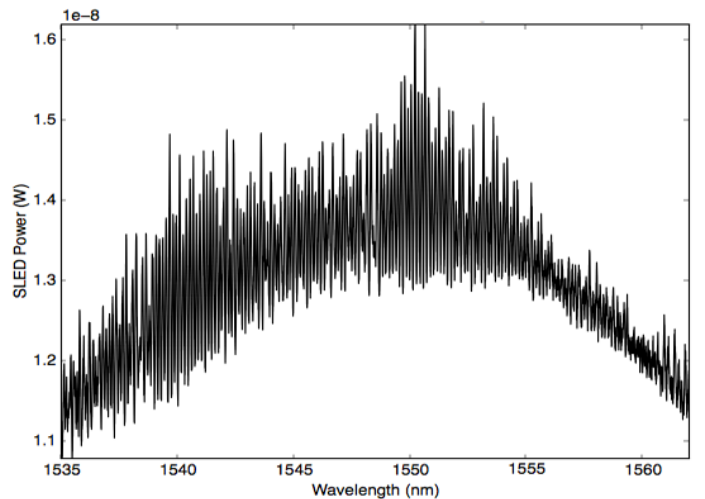


Figure 6: The optical spectrum of the SLED source used in this paper, showing the expected inverted U-shaped pattern.

prevented it from reaching its theoretical 1-pm resolution. In fact, further analysis with external optical reference systems (i.e. an interferometer and power meter) demonstrated that our tunable laser is unstable with an uncertainty of ~ 1 pm when tuned to certain wavelengths.

In order to remove this high-frequency noise from our data, we applied a tenth-order low-pass Butterworth filter with a cutoff frequency of 30.0 nm^{-1} . This allowed us to generate 80 filtered sweep factors that accurately and precisely captured each bin’s response to constant-intensity light.

II. ACCOUNTING FOR THE SLED SOURCE

Our spectrometer was used in conjunction with a superluminescent light-emitting diode (SLED) optical source. SLEDs are particularly advantageous for FBGs since their low temporal coherence prevents unwanted interference effects and allows for a broadband optical power spectrum; this in turn




Figure 7: *In red, we see the normalized and filtered sum of all 80 channels’ sweep factors as a function of wavelength. Relative minima are marked in blue and are used to determine bin “boundaries” for the middle 78 channels.*

allows for multiple FBGs to be addressed simultaneously. In addition, their high spatial coherence allows for low divergence angles, and therefore permits higher coupling efficiency into a single-mode fiber.

However, since the optical power from an SLED is emitted over a wide spectral range, we end up observing an inverted U-shaped pattern for optical power density as a function of wavelength; this pattern is to be expected for any SLED source. For our SLED, we have a manufacturer-specified central wavelength of ~ 1550 nm and a full width at half-maximum (FWHM) of 50 nm; Figure 6 is in agreement with these specifications.

In our setup, both the high-resolution OSA and the spectrometer use the same SLED source. However, the OSA can measure the light source’s spectrum, take it as a reference, and measure the FBG spectra against it. On the other hand, this is not possible for the spectrometer as there are inherent differences between how these two instruments collect data. Thus, any attempt at getting rid of the light source spectrum in the spectrometer data would end up adulterating the FBG spectrum. Instead, we take the high-resolution data from the OSA and re-apply the SLED optical density spectrum in order to emulate the spectrometer’s behavior and generate our profiles.

III. BIN BOUNDARY DETERMINATION

In order to connect generated profiles to actual channels in the spectrometer (as described in step 6 of the above procedure), it was necessary to determine the spectrometer’s bin boundaries. These bin boundaries would also be required to compare our algorithm with other SDAs (e.g. LPO heavily relies upon the idea of evenly spaced “pixels”). However, the notion of a bin boundary is not very well-defined; as we can see in Figure 5, channels 0 and 1 overlap significantly in wavelength space. One possible, unambiguous method of determining the boundary between these two channels is to sum their sweep factors for each wavelength and determine the relative minimum that occurs in the resulting graph. For the remainder of this article, it will be assumed that all bin boundaries were determined through this process.

Unfortunately, given the jitter in each individual sweep factor, the sum of all 80 channels’ sweep factors ended up being especially noisy. In order to actually extract meaningful information from these data, we once again applied a tenth-order low-pass Butterworth filter with a cutoff frequency of 30.0 nm^{-1} . This allowed us to more easily determine the boundaries between all 80 bins, as shown in Figure 7. Of course, this method cannot be used to assign lower and upper bin boundaries for the first and last channels (respectively) of the spectrometer, so our procedure does not work on the extreme ends of the spectrometer. Nonetheless, given the particularly broadband spectrum of the SLED and the sub-optimal positions of channels 0 and 79, we will assume that these extreme channels of the spectrometer will not be used in any case. Thus, the boundaries of the middle 78 bins ought to be sufficient.

As mentioned in subsection I, the bins of our spectrometer were manufactured with uneven “widths” in the wavelength domain. This was additionally verified through our boundary determination procedure, where we found an average bin width of 399.2 pm with a standard deviation of 25.5 pm and an overall range of $\Delta\lambda \in [349, 472]$ pm. Clearly, this degree of variation will not bode well for LPO, as the FIR technique used to determine the differentiator sequence relies on the even spacing of “pixels” [1]. Unfortunately, this discrepancy cannot easily (and consistently) be fixed across all existing spectrometers.

IV. EQUIPMENT AND DATA COLLECTION

In this section, we examine the equipment used to produce our results, in addition to some specifics about our data collection methods.

I. FIBER BRAGG GRATINGS

We use four different FBG sensors in our setup; for the remainder of this paper, these will be referred to as FBGs 1, 3, 5, and 8. All four gratings are 2.5 mm in length and are written into a Corning[®] SMF-28[™] single-mode optical fiber to a reflectivity of $\sim 50\%$. This fiber has a refractive index difference of 0.36% between its core and cladding, and so the weakly guiding approximation for single-mode fibers (see Section II) is valid. These gratings are then glued into the middle of metallic 20-mm-long tubes with an epoxy adhesive. The outer diameter of these tubes is around 0.7 mm, while the inner diameter is only slightly larger than the cladding’s diameter ($125 \mu\text{m}$).

Our four samples are distinguished not only by their Bragg wavelengths, but also by the alloy of the metallic tubes into which they are glued. One of the most significant differences between these alloys (for the purposes of our analysis) is their temperature response. In particular, these alloys not only have different coefficients of thermal expansion, but also demonstrate qualitatively different reactions (i.e. somewhat linear vs. very non-linear) to changes in temperature. For

example, the alloy for FBG 1 is known to have very non-linear temperature characteristics, as we will see in Section VI.

Any further details regarding these metallic alloys are not disclosed in this paper since they involve patent-pending technologies.

II. MEASUREMENT EQUIPMENT




Figure 8: Schematic of how the SLED and spectrometer (or OSA) interface with the various FBGs in our setup.

A general schematic of our setup can be found in Figure 8; we expand on the specific equipment and terminology of this diagram in the following paragraphs.

The temperatures of our FBGs’ metallic alloys are adjusted via a custom-made assembly controlled by a Newport[®] 8350 40W Thermoelectric Cooler (TEC) module with a $100\mu\text{A}$ thermistor. This module accommodates a temperature range of $[-100.0^\circ\text{C}, 240.00^\circ\text{C}]$ with a display resolution of 0.01°C and an accuracy of 0.1°C . Thus, any temperature residuals of $< 0.1^\circ\text{C}$ in our regression models can be partly attributed to measurement error.

In order to interrogate our FBGs, we use a near-infrared Thorlabs[®] SLD1005S SLED as our light source. This SLED has a central wavelength of ~ 1550 nm and a FWHM of 50 nm; since our spectrometer is concerned with a wavelength range of $[1528$ nm, 1562 nm], this is a sufficiently large range for our purposes.

The light from our SLED is channeled into our four different FBGs through the use of an Agilent[®] 86062C Lightwave Switch. The resulting FBG response is measured through either an OSA or a spectrometer.

In order to produce our high-resolution OSA data, we made use of an Agilent[®] Advantest Q8384 spectrum analyzer with built-in light source calibration. Our data was collected by setting the OSA’s sweep mode to “normal.” While this spectrum analyzer has a documented optimal resolution setting of 10 pm, we can improve this to a 5 pm resolution by using 1000 samples (as opposed to the default 500) for each 5 nm range. In order to achieve the picometer-precision shifts specified in Section III, we then apply linear interpolation as appropriate.

Our spectrometer was designed with optics from Ibsen Photonics and custom-made electronics; a raw data flow schematic is presented in Figure 9. After a steady temperature reading was reached, we collected 24-bit data from all 80 bins at a sampling rate of 3 kHz. The dimensions of these data

are specified as “arbitrary units” throughout this paper, yet they are related to the photocurrent (in pA) measured by the spectrometer via the following transfer function:




Figure 9: Raw data flow for our spectrometer (boxed on the left), adapted and modified from Figure 5.16 of [1].

$$\text{Spectrometer Reading} = \frac{13}{1.59} \cdot (\text{Photocurrent (pA)}) + 16384 \quad (10)$$

In order to remove noise, the spectrometer data used in this report were averaged over 250 samples. This resulted in a data production rate of 12 spectra per second.

One important distinction to note about our spectrometer is that we use a photodiode array instead of a CCD detector to collect our FBG’s reflectivity response. Photodiode arrays are advantageous since they allow for simple parallel readout of bins. On the other hand, CCDs often demonstrate charge bleeding across pixels when one attempts to conduct dynamic measurements.

Finally, our bin-wise sweep factors were collected by using an Agilent[®] 8190A compact tunable laser source in conjunction with our spectrometer. Through the use of external optical reference systems, the sweep mode of this picometer-precision laser was found to be unstable when sweeping at a rate of less than $5 \frac{\text{pm}}{\text{s}}$. This was visible in the interferometer output: almost every 10 pm, there seemed to be an in-built correction mechanism in the 8190A that caused the laser wavelength to suddenly jump up or down. Thus, a discrete one-picometer stepping mode was used instead. While we still noticed an uncertainty of ~ 1 pm with this mode when tuned to certain wavelengths, we were able to produce more reliable and accurate results relative to the sweep mode.

V. PROFILE FIT PROCESS

The procedure of finding a given Bragg wavelength will be run in real-time on a dedicated device connected directly to the spectrometer. The implemented algorithm must hence be fast enough to compute the λ_B in real time without sacrificing precision. Since we expect wavelength changes on the order of $5 \frac{\text{pm}}{\text{K}}$ for temperature and $1 \frac{\text{pm}}{\mu\text{m}\cdot\text{m}^{-1}}$ [1] for strain, we aim for picometer-level precision in our algorithm.

I. LAYERED SEARCHES TO REDUCE COMPARISONS

Over a variation range of ~ 1.2 nm (which is expected for a 40.0°C temperature range), this implies ~ 1200 profiles. One of these profiles must be selected at high precision with each iteration. With a fitting function, we assign each profile a “quality factor”, and search over this range of λ_B shifts to find

the profile with the best quality factor. To avoid fitting 1000s of profiles with every iteration – each fit requiring several possibly elaborate operations – we perform the search in “layers.” This involves splitting the wavelength range of ~ 1.2 nm into smaller “sections” and computing the quality factor for each section’s central profile. We then compare the quality factors of these profiles and use the best profile’s wavelength as an approximate indicator for the wavelength of the overall best profile. For our next iteration, we can thus limit our search range to wavelengths within one “section” of the most recent best profile. We can execute several search layers like this, before finally performing a picometer-resolution search over a comparatively small range. This simplification works provided the spectrum of qualities is single-peaked, which is usually the case.

For initial testing in a simulation, we employ our search algorithm on 400 measurements from the spectrometer, which were taken over a temperature range of 10.0°C to 50.0°C in intervals of 0.1°C . These were collected over time intervals on the order of minutes, which is longer than what is expected in application, meaning that the spectra in this article will contain less noise than in practice. A lower-noise sample is appropriate for this simulation as we wish to test the accuracy of the algorithm rather than the sensor as a whole. Additional noise would make correctly identified Bragg wavelengths difficult to distinguish from incorrect identifications.

We generate profiles using high-resolution OSA spectra for our four FBGs, each collected at reference temperatures of 15.0°C , 25.0°C , and 45.0°C . In this simulation, we search over a conservatively large profile offset interval of ± 1.1 nm, i.e. a total width of 2.2 nm, which gives us 2200 profiles to select from. A layered search is implemented, first performing a rough search with 40 evenly-spaced profiles. Then, a range of $\pm \frac{2.2\text{nm}}{40} = \pm 0.055$ nm is searched at pm precision on both sides of the best profile of the rough search. This reduces the 2200 fits to $40 + \frac{2 \cdot 0.055\text{nm}}{1\text{pm}} = 150$ fits, which is a much more manageable amount.

Our layered search procedure assumes that the 40 roughly selected profiles are representative of their respective ranges. However, if even more layers are performed before the picometer-resolution layer, or if a layer is too rough, this search algorithm risks completely missing the globally best profile. The full extent to which the total number of fits can be reliably reduced without incurring such “misses” is further discussed in Section VII, where various search layer arrangements are examined. For now, we note that a 40-profile rough layer followed by a picometer-resolution layer is found to reduce computational time to a reasonable level (for available simulation hardware) and never incur such misses.

II. OTHER SIMPLIFICATIONS AND ASSUMPTIONS

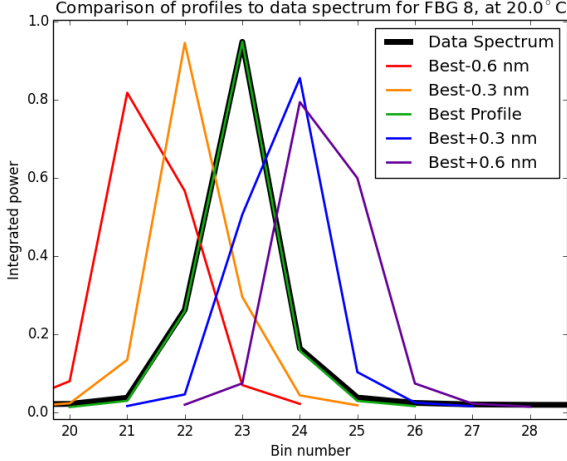
Profiles can contain intensity information on all 80 bins of the spectrometer. However, comparing all of these bins to a given data spectrum would not be practical for several reasons. An individual peak in the spectrum, as seen in

Figure 2, occupies a width on the order of 1.5 nm; with bin widths of approximately 0.5 nm, this translates to 3-4 bins depending on alignment. Comparisons between bins far away from the peak of a distribution would only add noise to the fit. We hence limit the comparison of profiles against data spectra to very few bins centered around the peak bin of the profile. In this simulation, we compare seven bins, aligning the fourth to the profile maximum. This should safely cover the peak in our spectrometer data. We do not worry about noise from bins far from the center as our spectrometer data in this article is expected to be less noisy than in practice. The effects of changing profile width are further discussed in Section VII. Omission of most of the 80 bins per profile also eliminates the need to store this information on a dedicated computing device, which may have limits in (for example) flash storage space. Note that, although bins are being omitted, the identities of the bins are still stored since the widths of the bins vary significantly; thus, a comparison should avoid misaligning bins (e.g. fitting profile bin 54 to data bin 56).

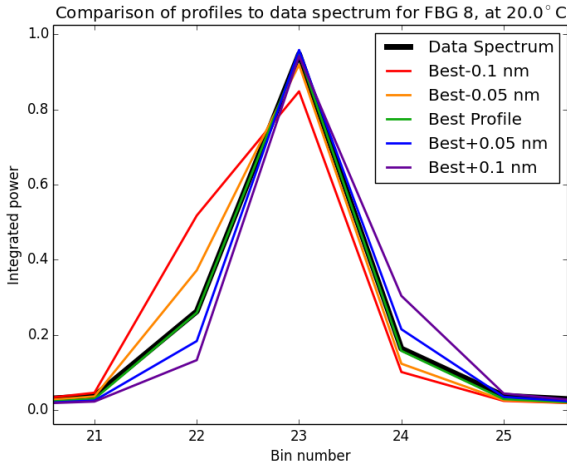
To further reduce the number of necessary computations, it is noted that a given profile’s precomputed Bragg wavelength rarely falls outside the profile’s highest intensity bin; in fact, the case where λ_B is outside this bin is never observed in any of the profiles used in this simulation. We can hence make the assumption that if a spectrum’s peak bin does not match a profile’s peak bin, then the profile cannot match the spectrum’s Bragg wavelength within one bin width (which is ~ 0.4 nm, much larger than the precision we desire). Accordingly, we add an additional condition to our search procedure to check if the peak bins of a data spectrum and a given profile match; if this is not the case, we ignore the profile completely. This comparison is easy to perform as the peak bin of the profile is always the center bin, and the peak pin of the data spectrum only needs to be computed once (namely, before the search begins). In simulation, this simplification is found to significantly increase performance, as occasionally up to a third of profiles can be skipped. However, in practice (or with other spectrometers or data), this does incur the risk that all profiles in a given rough search fail the peak-matching test and are skipped; in such a case, no information about the position of λ_B can be gathered. This is not observed in the simulation performed here, but in the investigation in Section VII, it is shown that sufficiently rough searches encounter this problem.

III. FITTING VIA INNER PRODUCTS

The number of profiles that need to be fitted is reduced as much as possible via layered searching and peak matching. After this initial culling procedure, we use a fitting algorithm that accounts for the fact that collected data from the spectrometer (which is in arbitrary units) may not be of the same magnitude as profile data (which is in integrated power units). A χ^2 algorithm is not suitable for this purpose, as the difference in shape of the distributions is the point of interest, rather than deviations solely dependent on the magnitude of



(a) *Clear progression of profile shape with wavelength offset. Note that all profiles shown in this figure except for the best profile would be ignored due to the peak-matching constraint.*



(b) *Small changes in profile wavelength offsets produce fluctuations in bin values with a less clear progression. However, inner product comparison is still effective.*

Figure 10: *Data spectrum, best profile, and surrounding profiles for FBG 8. Profiles were generated with a reference temperature of 25.0°C and are used to fit a data spectrum measured at 20°C.*

the bins. In initial testing, it was found that a χ^2 algorithm failed to even match simulated spectra (i.e. spectra generated using the same technique as the profiles), yielding results that were a few picometers off.

Instead, we use an inner product fitting algorithm. To begin, we normalize each of the profiles such that its inner product with itself, over the bins that it contains, is one. We do the same for the data spectrum before running our fitting procedure. This normalization allows us to use the inner product of the data spectrum with a profile as our “quality factor.” Denoting the data spectrum bin vector with \mathbf{D} and the profile bin vector with $\mathbf{P}(t_{\text{ref}}, \lambda)$, we let:

$$\mathbf{D}' = \frac{\mathbf{D}}{\sqrt{\mathbf{D} \cdot \mathbf{D}}}; \quad \mathbf{P}' = \frac{\mathbf{P}}{\sqrt{\mathbf{P} \cdot \mathbf{P}}}$$

to normalize. This normalization can be performed before the search procedure. Then we maximize the quality factor:

$$\lambda_{\text{best}} = \arg \max_{\lambda} (\mathbf{D}' \cdot \mathbf{P}'(\lambda))$$

so that λ_{best} is the best wavelength offset. If the profile were exactly the same as the data spectrum, this inner product would yield a maximum value of one. Any deviation in shape would yield a lower inner product.

This quality factor prioritizes shape in the fit and effectively deals with magnitude differences. Moreover, it is not very computationally intensive, as it requires a single inner product normalization on the data (which can be performed before the search procedure) and as many multiplications as there are bins to compare per profile. Finally, it is flexible to any number of bins per profile, provided that they all have the same number.

This fitting procedure was implemented in a series of Python scripts and executed. In Figure 10, we see an example of the described procedure, showing several profiles generated at different offsets as well as the best profile as found by the inner product. This figure was constructed in reference to spectrometer data collected at 20°C. The performance of this procedure with regards to extracting useful data from an FBG is discussed in the following section.

VI. PROFILE FIT RESULTS

Our search algorithm takes a data spectrum with unknown temperature or strain and returns the spectrum’s λ_B shift relative to a reference λ_B (technically, an actual wavelength is returned, but it is measured with respect to a Gaussian fit and hence is not as physically useful as the offset). From this offset parameter, we must still extract information on temperature or strain in real time in order to ascertain useful data.

In our analysis, focusing solely on temperature, this requires the preliminary construction of a linear fit of $\Delta\lambda$ vs ΔT using spectrometer data collected at known temperatures. Thus, we compute offsets for data spectra with temperatures in the range of 10.0°C to 50.0°C in intervals of 0.1°C.

I. CORRELATING BRAGG WAVELENGTH AND TEMPERATURE

Since we are only interested in Bragg wavelength offsets relative to a reference λ_B , we display it as a delta to the λ_B of the OSA spectrum used to generate the profiles. We refer to the difference in λ_B of a given data spectrum to that of the OSA spectrum as $\Delta\lambda$. Similarly, we subtract the the OSA spectrum's reference temperature from the measured temperature and refer to this variable as ΔT .

After obtaining data on ΔT vs $\Delta\lambda$ for the entire temperature range, we perform polynomial fits of linear and quadratic order on the results, which are later used to translate $\Delta\lambda$ to temperature units. The strength of our regression model reflects the performance of our overall procedure: the effectiveness of the search algorithm, the correct choice of degree of the polynomial fit, the procedure of profile generation as a representation of the spectrometer's distortions, and the actual correlation of $\Delta\lambda$ with the temperature variable.

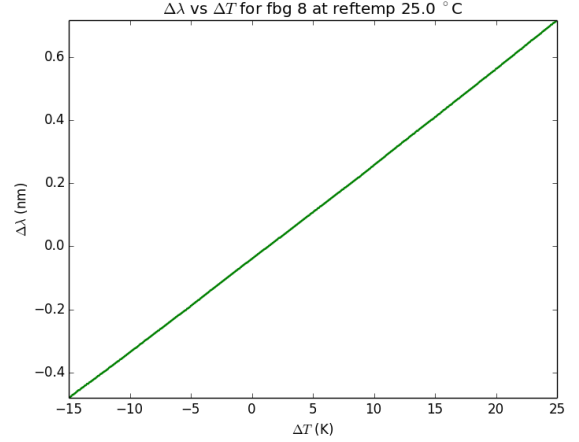
A result with seven bins per profile, working with FBG 8 and profiles generated with an OSA spectrum recorded at 25°C, is shown in Figure 11. This figure displays a $\Delta\lambda$ vs ΔT plot as well as difference plots for linear and quadratic fits.

The following conclusions are drawn from the results of the search algorithm and fits, referring to Figure 11. Figure 11a shows that, qualitatively, our procedure yields a consistent, linear correlation with no immediately visible jumps, dents, or other deviations; this suggests that the algorithm we apply is indeed a feasible technique.

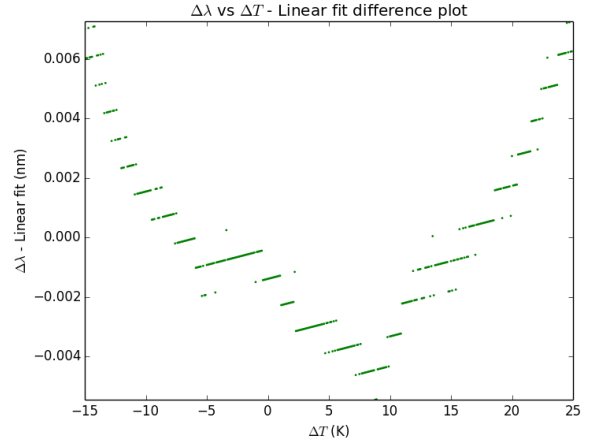
The linear fit result shown in Figure 11b has two properties of interest. First, the points are not distributed continuously along the shape of the curve but, instead, are grouped into lines. This can be explained by the fact that, in our algorithm, we chose a minimum search interval of one picometer, which groups the $\Delta\lambda$ component of the results into bins that are one picometer wide. As for the sloped lines in the figure, these result from computing a difference plot by subtracting a linear prediction from our points in Figure 11a. As a second point of interest, a quadratic systematic pattern is clearly visible with residual magnitudes on the order of 5-6 pm.

Figure 11c shows that a quadratic fit describes the correlation between $\Delta\lambda$ and ΔT to a high degree of precision, and additionally demonstrates that no higher-order polynomial components are evident in the correlation at this level of noise. Again, we see that the points are distributed in groups, this time in parabolic shapes that result from subtracting a second-degree polynomial.

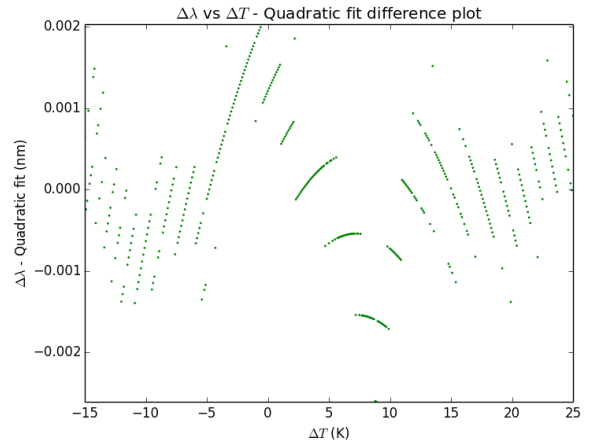
The fact that we don't have completely linear behavior in an FBG is unsurprising, as the theory predicted slightly nonlinear behavior for $\Delta\lambda$ vs ΔT at high temperatures. However, we do not make use of the more accurate quadratic fit because constraints on available hardware make a linear fit easier to implement. Since the quadratic component is minimal compared to the linear and constant components of the fit, the amount of accuracy lost is worth the gain in computational



(a) Qualitatively very linear correlation of $\Delta\lambda$ with ΔT .



(b) Linear fit residuals with a slightly quadratic pattern.



(c) Small, non-systematic residuals for a quadratic fit.

Figure 11: Results and fits of $\Delta\lambda$ vs ΔT from applying the search and fitting procedure on FBG 8 at reference temperature 25°C.

σ_{res} (K)	15°C	25°C	45°C	σ_{σ} (K)
FBG 1	1.3375	1.3556	1.3613	0.0101
FBG 3	0.4359	0.4215	0.4150	0.0087
FBG 5	0.2493	0.2471	0.2338	0.0069
FBG 8	0.1016	0.1001	0.1017	0.0007

Table 1: Table of standard deviations of residuals for all FBGs at all reference temperatures, as well as the standard deviations of σ_{res} over reference temperatures. Observe the very small variation in quality for different reference temperatures.

speed and simplicity. Furthermore, as we will see in analyses of other FBGs, even a quadratic fit sometimes fails to account for an FBG's nonlinearity (which can sometimes be very severe).

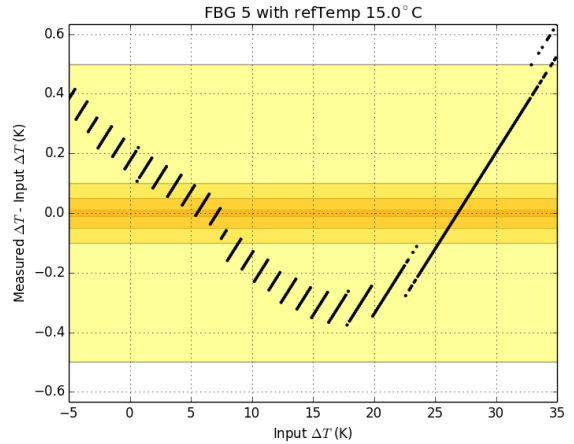
II. ALGORITHM PERFORMANCE AT DIFFERENT REFERENCE TEMPERATURES

We use the linear fit parameters to transform the $\Delta\lambda$ axis to a ΔT_{meas} axis (where $\Delta T_{\text{meas}} \equiv \text{Computed } \Delta T - \text{Actual } \Delta T$). Linear fit parameters generated in a similar fashion (i.e. by fitting the $\Delta\lambda$ vs ΔT results of a large number of data spectra with known temperatures) can be regarded as a calibration procedure for this algorithm; this procedure will also be useful in a final implementation.

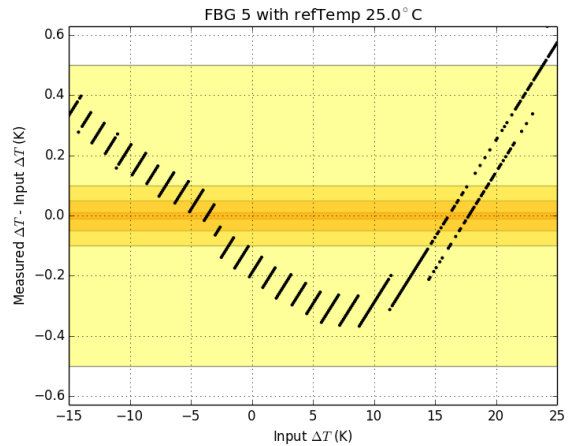
Spectrometer data was collected for our four FBGs for the temperatures 10.0°C to 50.0°C in intervals of 0.1°C as before. For each FBG, high-resolution OSA spectra were generated at “reference” temperatures of 15°C, 25°C, and 45°C, and these were used to generate profiles. Linear fits were applied to convert $\Delta\lambda$ to ΔT_{meas} , and difference plots were computed to show the quality of the linear fit and the behavior of each FBG.

First, we compare changes in algorithm performance for different reference temperatures. Table 1 shows the changes in standard deviations of residuals for different FBGs and reference temperatures. Figure 12 shows the difference plots of FBG 5 for different reference temperatures.


As we can clearly see, the reference temperature makes very little difference in algorithm performance according to standard deviation. Furthermore, we observe no correlation between the size of residuals and their closeness to $\Delta T = 0$ (on the horizontal axis). This means that the shapes of the high-resolution OSA spectra, exemplified in Figure 2, do not change significantly in shape with temperature; instead, these spectra mostly just shift in the wavelength domain. This is convenient as it means that we can reliably perform this entire procedure at only a single reference temperature. Otherwise, it would have been necessary to coalesce profiles generated at different reference temperatures for different intervals of temperature in order to maximize procedure performance. For the rest of the analysis, we hence only generate profiles with the OSA spectrum at reference temperature 25.0°C, chosen



(a) Residuals have standard deviation $\sigma_{\text{res}} = 0.2493$ K.



(b) Residuals have standard deviation $\sigma_{\text{res}} = 0.2471$ K.



(c) Residuals have standard deviation $\sigma_{\text{res}} = 0.2338$ K.

Figure 12: ΔT_{meas} vs ΔT , computed for FBG 5 at all three reference temperatures. Observe the minimal changes in shape. Color markers added for residuals of magnitude <0.5 K, <0.1 K, <0.05 K and <0.01 K. These markings are present in later figures.

arbitrarily because it is most central in our temperature range of interest.

III. ALGORITHM PERFORMANCE FOR DIFFERENT FBGs

Figure 13 shows the performance of our algorithm over the four different FBGs in our test procedure. We display standard deviations of the residuals again in Table 2, along with the slope of the linear fit with a 95% confidence interval.

We observe that FBG 1 behaves very nonlinearly, while other FBGs also show behaviors that deviate from a linear behavior. However, high-quality fitting with residuals within $\sim 0.1\text{K}$ is observed for FBG 8. At this level of precision, the error can largely be attributed to our TEC’s temperature control accuracy of $0.1\text{ }^\circ\text{C}$ as mentioned in Section IV, Subsection II.

The failure of some FBGs to correlate linearly with temperature may not necessarily result from a flawed search algorithm. These difficulties are most likely consequences of the FBGs’ manufacturing process, and could be eliminated with (for instance) a higher-order polynomial fit. We do not believe that the strange shape of the ΔT_{meas} vs ΔT graphs of some FBGs is a consequence of the search algorithm; to the contrary, it is possible that the search algorithm gives an accurate measurement of the FBGs’ nonlinearity. This postulate is somewhat supported by the fact that other well-performing algorithms, such as the LPO technique, yield similar shapes in ΔT_{meas} vs ΔT for the FBGs examined, as further explored in Section VIII.

VII. VARYING PROFILE WIDTHS AND LAYERED SEARCHES

The parameters for the brute-force algorithm discussed in this article are:

For generating profiles:

- High-resolution FBG-specific measurements
- Measurements on the spectrometer

For fitting profiles:

- Number of profile bins to compare (profile width)
- Arrangement of search layers

For calibrating linear fits:

- Measured spectra at known temperatures over the range of interest

The majority of these parameters are equipment-specific measurements, and the performance of our procedure will be dependent on their accuracy. This is an important difference between our algorithm and those discussed in Section VIII.

However, the two parameters related to profile fitting are not measurements, even though the best choice for each variable depends on our FBGs and spectrometer. This section explores the consequences of varying these parameters on overall procedure performance.

I. VARYING PROFILE WIDTH

As discussed earlier in Section V, Subsection II, we expect peaks in FBG spectra to occupy $\sim 1.5\text{ nm}$ in wavelength space, which – with a bin width of $\sim 0.5\text{ nm}$ – yields 3-4 bins per peak. However, to ensure minimal data loss, we chose to fit seven bins centered about the maximum of the spectrometer data. This extension to seven bins was acceptable in our simulation since our data was averaged over several seconds and hence contained little noise. In practice, bins outside the peaks are expected to be more noisy, and thus their inclusion in fits should be avoided. On the other hand, if we choose too few bins, they may not contain enough information to accurately determine the correct profile.

We perform our search algorithm on all FBGs with varying profile widths, with results shown in Table 3. We choose odd numbers of bins and generate profiles such that the bin containing the spectrometer maximum is at the center. As we can see, there is minimal performance variation over different profile widths.

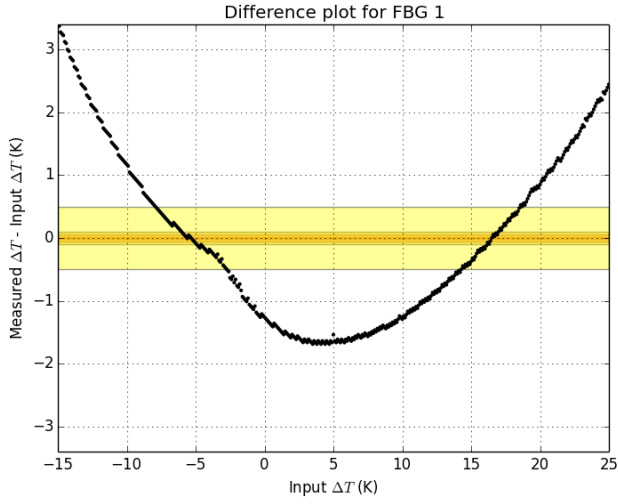
σ_{res} (K)	7 bins	5 bins	3 bins	σ_{σ}
FBG 1	1.3556	1.3551	1.3538	0.0008
FBG 3	0.4215	0.4200	0.4186	0.0012
FBG 5	0.2471	0.2546	0.2457	0.0039
FBG 8	0.1001	0.1005	0.1009	0.0003

95% ($\pm \frac{\text{K}}{\text{nm}}$)	7 bins	5 bins	3 bins	$\sigma_{95\%}$
FBG 1	0.8686	0.8683	0.8668	0.0008
FBG 3	0.7336	0.7314	0.7338	0.0011
FBG 5	0.2234	0.2303	0.2236	0.0032
FBG 8	0.0284	0.0286	0.0286	0.0001

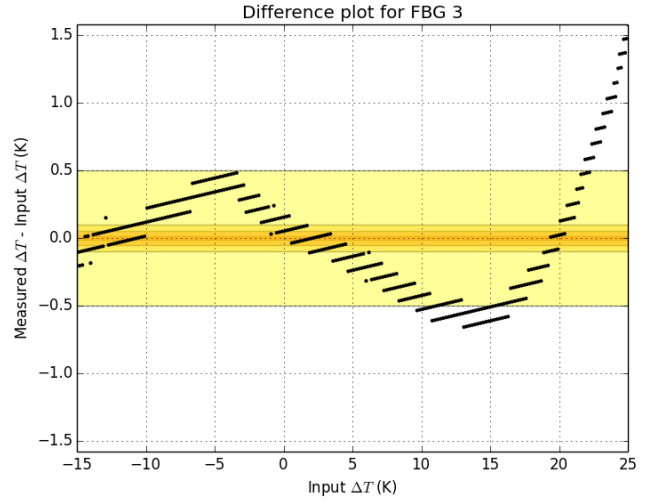
Table 3: Table of standard deviations of residuals as well as a table of the half-width of the 95% confidence interval for the linear fit’s slope parameter. These tables cover all four FBGs and three different profile widths.

Two conclusions can be drawn from this consistency in performance. First, the central bins contain enough information to reliably fit profiles most of the time, meaning that we can confidently reduce the profile width to three bins. Second, since large profile widths perform as well as small profile widths, we can confirm our hypothesis that our simulation data is relatively low-noise. This is true since comparisons with more bins will fit bins outside of the peak region, which – due to their low magnitude in spectrometer output – will be more influenced by noise. To be safe, we recommend using as few bins as possible, because actual data (collected in real time) will be more noisy than the data used in this simulation. Thus, for the remainder of this analysis, we use three bins for profile comparisons. As mentioned earlier, this also decreases profile storage space and reduces computation time, since fewer bins need to be stored and compared.

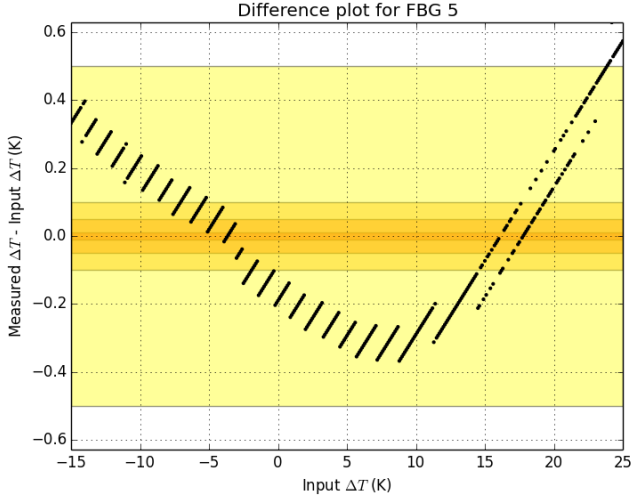
Despite minimal variation in the quantitative performance variables of σ_{res} and the slope parameter’s confidence interval, reducing the number of compared bins does produce misses in



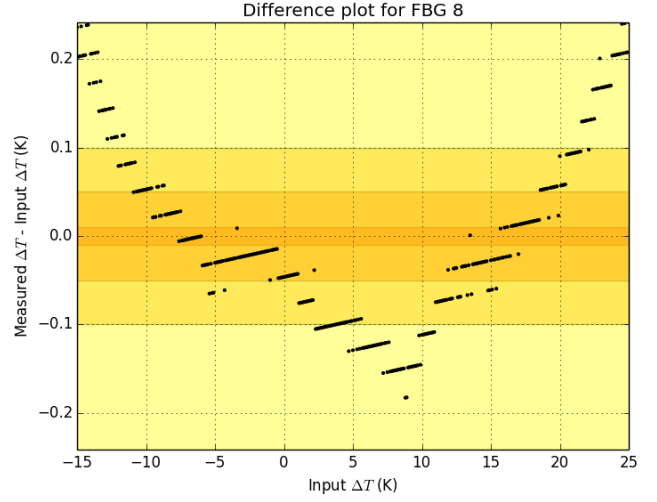
(a) Residuals have standard deviation $\sigma_{res} = 1.3556$ K.



(b) Residuals have standard deviation $\sigma_{res} = 0.4215$ K.



(c) Residuals have standard deviation $\sigma_{res} = 0.2471$ K.



(d) Residuals have standard deviation $\sigma_{res} = 0.1001$ K.

Figure 13: ΔT_{meas} vs ΔT , computed for all FBGs at reference temperature 25°C .

	σ_{res} (K)	slope (K/nm)	95% conf.
FBG 1	1.3556	74.95	± 0.87
FBG 3	0.4215	204.89	± 0.73
FBG 5	0.2471	106.46	± 0.22
FBG 8	0.1001	33.45	± 0.03

Table 2: Table of standard deviations of residuals, as well as slope of linear fit with 95% confidence interval, for all FBGs.

our search algorithm. Figure 14 gives examples of FBGs with misses in their ΔT_{meas} vs ΔT difference plots. However, these are regarded as insignificant since they are very few in number, especially in comparison to the misses that might occur due to extra noise when comparing more bins.

II. PERFORMING LAYERED SEARCHES

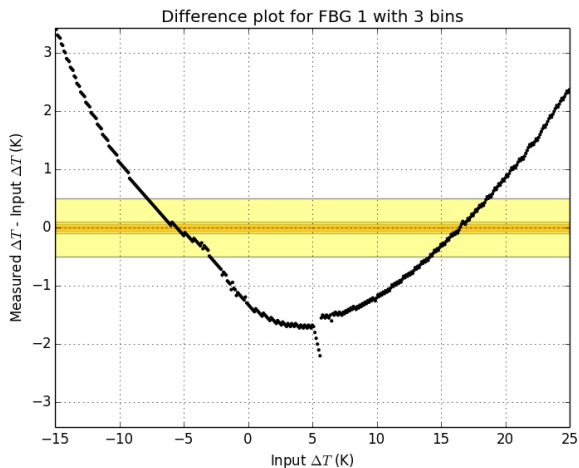
Until now, a layered search over the entire profile range was performed with an initial rough search of 40 profiles, followed by a picometer-resolution layer centered on the best profile of the rough search. Using this technique, the upper bound of 2200 profile comparisons was reduced to only 150 comparisons.

Performing multiple layers of rough searches before a high-resolution search, or varying the size of the rough searches, could decrease the number of necessary comparisons even further. However, performing too rough of a search may cause the search to miss the overall best profile. It is therefore necessary to find a balance between performance and reliability. Note that, as the effectiveness of a layer arrangement depends on the FBG, spectrometer, and temperature range in question, it is likely that the numbers determined in this section may not be applicable to an application of this algorithm to different devices. The technique used to determine the optimal arrangement, however, can be useful for other setups.

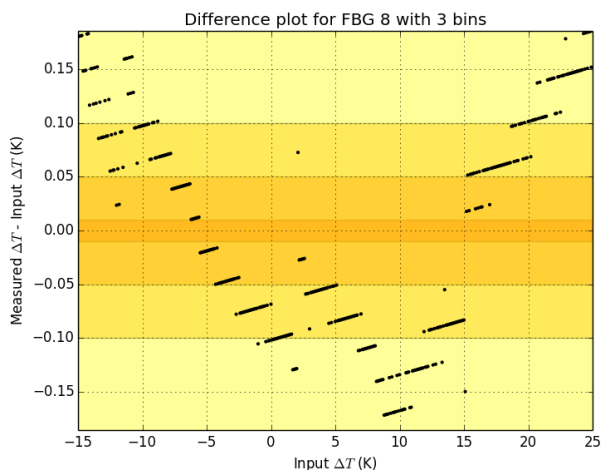
In order to establish a notation for the order of search layers, we will denote a search via a list of layer sizes surrounded by parentheses. Additionally, a layer size of “p” implies picometer-level precision. Thus, the search layers used so far would be denoted by (40, p). Furthermore, we apply a shorthand for successive layers of the same size using multiplicative notation. For example, an algorithm with a layer of 30 comparisons, followed by three layers each with five comparisons, followed by a picometer-level layer would be denoted by (30, 3×5, p). Occasionally, the expanded (30, 5, 5, 5, p) representation is chosen for readability.

The total number of comparisons is the sum of the number of comparisons over all layers; we intend to minimize this quantity. Moreover, we require search precision on the picometer level, i.e. over 2200 profiles, so the product of the number of comparisons per layer must be approximately equal to 2200.

The search algorithm that minimizes the sum of comparisons is a binary search, which in our case would be (11×2, p). However, this cannot be applied for two reasons: First, a binary search works best if the qualities of the profiles were mostly sorted, which is very much not the case. In reality, we observe patterns like the one shown in Figure 15. The fact that some profiles are skipped due to peak-matching, combined with the fact that the entire algorithm fails if all profiles in a certain layer are skipped, will cause such a binary search to not only produce erroneous results but also (possibly) produce no result at all. Secondly, and more importantly, we noted in Section V, Subsection I, that we choose a range of $\pm w/n$ around the best profile of a layer for the next search, where w is the previous search width and n is the size of



(a) FBG 1 shows some misses around the center that do not occur for comparisons with fewer bins (see Figure 13a).



(b) FBG 8 shows some misses and a small change in shape (see Figure 13d). This new shape may not be incorrect since noise may affect comparisons with more bins.

Figure 14: ΔT_{meas} vs ΔT for FBGs 1 and 8 with a profile width of three bins.

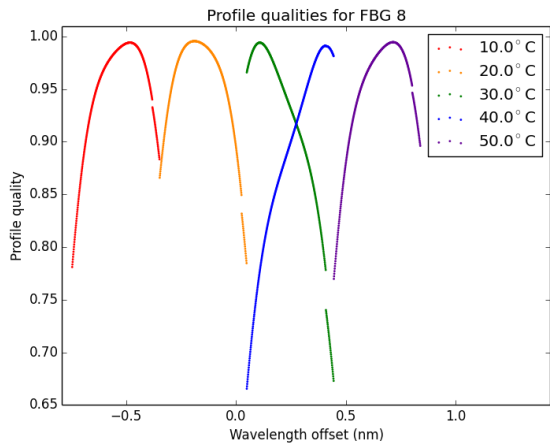


Figure 15: Profile qualities as a function of λ_B offset for (p) search on FBG 8 with data spectra at specified temperatures.

the next layer. This implies that the search width becomes $w_2 = w/(n/2)$; with a binary search (where $n = 2$), the search width would not decrease and the overall algorithm would not converge on a range at all.

For this reason, the product of the layer sizes does not represent the precision of the search algorithm. Instead, it is the product of each of the layer sizes divided by two. As an example, for (4, 6, 4), the search precision is $(4/2) \cdot (6/2) \cdot (4/2) = 12$.

To assess the performance of a given search layer arrangement, we count the total number of fitted profiles f and skipped profiles s . We compute the average of these values over all temperatures and all FBGs. To create an overall performance factor, we note that skipping a profile requires just a single comparison of the profile center to the data spectrum center, whereas the fit requires this same comparison as well as one multiplication and one sum per bin (to compute the inner product) and one comparison to check for improvement in the quality factor. Thus, for a profile width of three bins, the ratio of computation cost for skips to fits is 1:7, where the seven comes from the fact that a fit requires two comparisons, three multiplications, and two sums. Hence, the overall “score” r for a search is given by:

$$r = s + 7f$$

In order to find the algorithm with as many small layers as possible without missing the best profile, we compute the overall best profile beforehand using a (p) search. We can then use the following algorithm to find the best layer arrangement:

1. Set a lower bound on layer size, which we will call n_l .
2. Attempt to perform a layered search of (n_l, p) , and cross-check the results of the search with the previously computed offsets (from the (p) search).
3. If the search fails to identify the best profile, or stops because all profiles in a layer are skipped, increment the

first search to $(n_l + 1, p)$ and try again. Continue to increment until the search succeeds at $(n_l + k, p)$.

4. Append another layer just before the picometer-precision layer of size n_l , i.e. search with $(n_l + k, n_l, p)$.
5. Increment the penultimate search layer until the search succeeds.
6. Repeat steps four and five until the search precision before the picometer-resolution layer exceeds 2200.
7. Of all layer arrangements iterated, choose the one with the lowest score r .

We perform this check for several values of n_l , since – although the smallest possible value $n_l = 3$ might appear to be optimal – the actual performance of the search is also dependent on the skip-to-fit ratio and the increments made to the layers by the algorithm above to ensure accuracy. We test values of $n_l \in [3, 12]$ and show our results in Table 4.

n_l	Resulting search	f	s	r
3	(6, 9×3, 5, p)	37.40	8.06	269.88
4	(6, 4×4, 12, p)	33.54	7.92	242.69
5	(6, 4×5, 6, 5, 5, p)	31.65	7.81	229.63
6	(6, 6, 6, 6, 9, p)	32.07	8.31	232.93
7	(7, 7, 7, 8, 7, p)	30.53	8.93	222.68
8	(8, 8, 8, 9, p)	31.17	9.30	226.77
9	(9, 9, 9, 9, p)	31.25	10.22	228.98
10	(10, 10, 14, p)	36.29	11.16	265.21
11	(11, 11, 12, p)	36.63	10.83	267.22
12	(12, 12, 14, p)	35.09	12.36	258.03
-	(40, p)	113.61	36.47	831.74
-	(p)	400.27	1800.73	4602.61

Table 4: Performance of layered searches with various values of n_l , averaged over all temperatures and all FBGs. All of these searches reach the same accuracy as (p). The searches (p) and (40, p) are shown for reference.

We see that the best search arrangement is (7, 7, 7, 8, 7, p), which has an r value that is 26.8% of the (40, p) search arrangement’s r value – a 273.3% improvement.

There are a few points of interest to note about this procedure. First, note that reducing the search layers to the bare minimum risks incorrectly identifying the best profile or causing other overall failures in our procedure. Furthermore, it is recommended to run the analysis performed in this section with data that represents the final measurement setup. In our case, the result we achieve might not be useful for our final setup since this algorithm will be applied to data more noisy than the data in this analysis. Lastly, since implementation of these searches on an FPGA chip is highly parallelizable, it is feasible to adjust the layered search to contain fewer layers, with each layer containing at most the number of fits a chip can perform simultaneously. Note that the computational cost score r in this section is an assessment

of *sequential* computation time and hence does not correlate with parallelizable computation time. In such a scenario, the optimal layer arrangement can be found by performing the algorithm used above with a value of n_l close to the maximum parallelizable number of fits, while adjusting the score r for parallel execution.

VIII. COMPARISONS AGAINST OTHER SDAs

In order to assess our algorithm’s precision and speed, we must compare it against various existing subpixel detection algorithms (SDAs). As before, we apply these algorithms on previously collected spectrometer data (measured at a known temperature) and determine a Bragg wavelength offset. This allows us to find a linear model for ΔT as a function of $\Delta\lambda$ for a given FBG, from which we can later determine temperatures with real-time spectrometer data. We then compare the difference plots, residual standard deviations (in K), and (in some cases) the fit parameter 95% confidence intervals for each algorithm’s regression model against the results obtained in Section VI.

Furthermore, to get a relative idea of our algorithm’s speed, we compare the number and kind of operations required by each algorithm for a single Bragg wavelength (or wavelength offset) determination. By comparing computation costs, we get a sense of how fast our algorithm will run when implemented on an FPGA.

The contents of subsections I through V are organized as follows. We begin by devoting a subsection to each algorithm (i.e. the Gauss fit, Kaiser fit, CDA, and LPO4 FIR technique). In these subsections, we describe the theory behind each SDA and produce difference plots in a similar fashion to Figure 12; these can be used to compare accuracy and residual patterns against those of our algorithm. A summary of performances is given at the end in Table 6. Finally, in subsection V, we produce a summary table that compares the number (and kind) of arithmetic operations required for a Bragg wavelength (or offset) determination for each algorithm explored in this paper.

For convenience, we only look at data produced by FBGs 5 and 8; as described in Section VI, these FBGs produce strongly linear plots for $\Delta\lambda$ vs ΔT with low temperature residuals when our algorithm is applied, while FBGs 1 and 3 are known to be highly non-linear. Moreover, whenever our profile-fitting algorithm is mentioned for speed and accuracy comparisons in the following subsections, it is to be assumed that we use a reference temperature of 25°C and a profile width of three spectrometer bins. As described in Sections VI and VII, the choice of reference temperature is arbitrary, and three spectrometer bins ought to be sufficiently accurate. In addition, we use a layered search of (7, 7, 7, 8, 7, p), which was found to require the fewest operations in Section VII, subsection II.

Finally, it is important to note that, since all four SDAs require some notion of wavelengths that are tied to the ac-

tual spectrometer bin data, we use the centers of the bin boundaries (described in Section III) for this purpose.

I. GAUSS FIT

Initially, the reflectivity spectra of FBGs were assumed to be Gaussian in shape due to apodizations in the data caused by lower laser coherence lengths; these apodizations resulted in decreased fringe contrast, making it difficult to identify sidelobes. In cases where the characteristic Kaiser peak of the FBG spectrum is significantly more noticeable than its sidelobes, it is still somewhat true that the observed peak approximately fits a Gaussian profile.

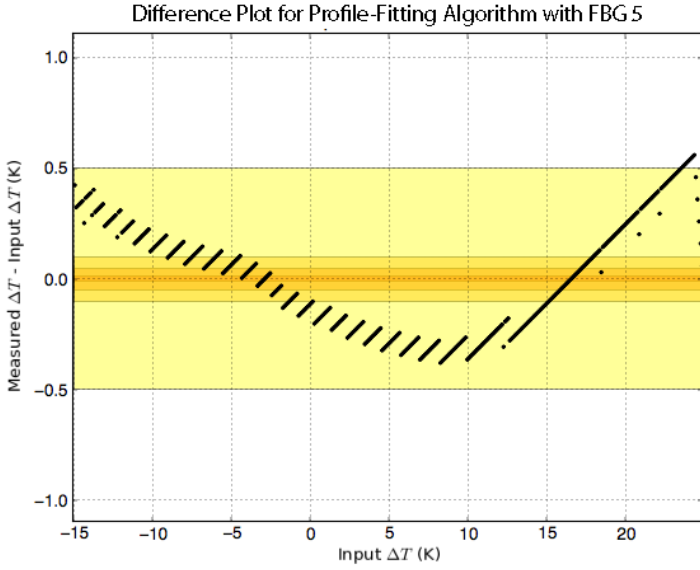
In our case, we are concerned with a three-bin Gaussian fit, which we will center around λ_p – the central wavelength of the bin B_p containing the peak power value in our spectrometer data. As mentioned in [1], if we let $R(B)$ represent the reflectivity of the spectrometer data as a function of bin number, we have the following sub-channel estimate B_s for the Bragg wavelength with a Gaussian fit:

$$B_s = B_p + \frac{1}{2} \left(\frac{\ln(R(B_p - 1)) - \ln(R(B_p + 1))}{\ln(R(B_p - 1)) - 2\ln(R(B_p)) + \ln(R(B_p + 1))} \right) \quad (11)$$

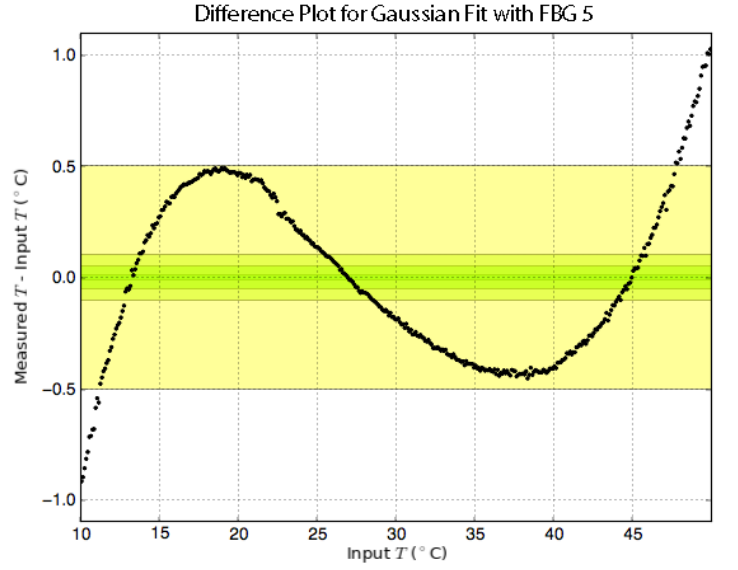
In order to actually convert this sub-channel estimate into a wavelength, we make use of the bin boundary data presented in Section III and interpolate as necessary. We are then able to compare the results of this Gaussian fitting procedure with our own algorithm on FBGs 5 and 8 in Figure 16, using three bins of spectrometer data as inputs in both cases.

For both FBGs, we see that our profile generation algorithm outperforms the Gaussian fit in terms of both its residuals and standard deviations. In Figures 16b and 16d, we clearly see oscillatory patterns in the Gaussian fit’s residuals; these patterns do not match the theoretical prediction of a systematic quadratic trend in the difference plot. On the other hand, our profile generation procedure does produce the expected U-shaped plot.

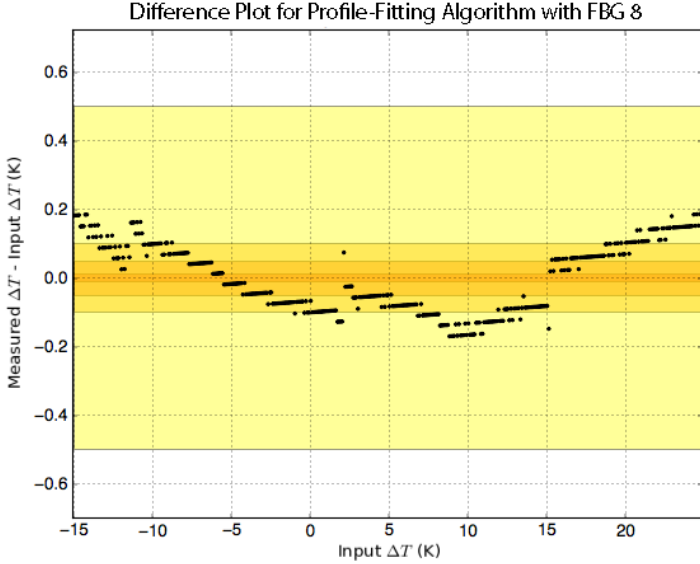
The reasons behind these oscillatory patterns can be easily explained by understanding how the spectrometer collects data. If we had achieved a perfectly linear fit with the Gaussian approximation, this would mean that the spectrometer data is best fit by a “traveling” Gaussian that moves a fixed distance in the wavelength domain for a given change in temperature. However, even though this might be approximately the case for the high-resolution OSA spectrum of an FBG, it is important to note that the spectrometer does not work in the same way as an OSA. After all, both the sweep factors and SLED optical density variations described in Section III end up biasing the spectrometer data in a way that cannot easily be predicted by a simple, evenly-shifting Gaussian fit. The end result is that the Gaussian function sometimes jumps too far ahead in λ_B for a given ΔT , and – to “correct” for changes in the bias – later jumps less far ahead in λ_B for the same ΔT . This would explain the volatilities in $\frac{d\lambda}{dT}$ that are suggested by the linear regression’s residual plots in Figures 16b and 16d.



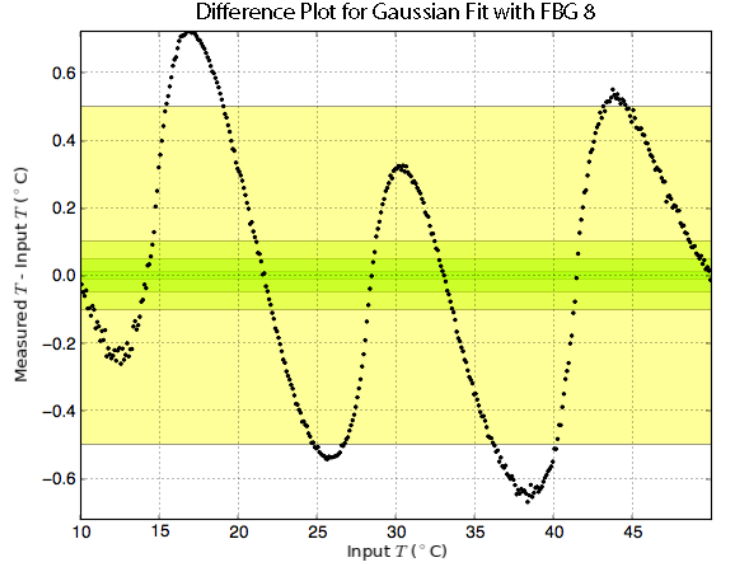
(a) Profile-fitting method with FBG 5 ($\sigma = 0.2457$ K).



(b) Gauss-fitting method with FBG 5 ($\sigma = 0.3835$ K).



(c) Profile-fitting method with FBG 8 ($\sigma = 0.1009$ K).



(d) Gauss-fitting method with FBG 8 ($\sigma = 0.3942$ K).

Figure 16: Comparisons between our profile-generation fitting method (orange) and Gaussian fitting method (green) for FBGs 5 and 8. The shades of green are used to demarcate temperature residuals in a similar manner to the color markers of Figure 12.

In addition, it is worth noting that we see more oscillations in the residual plot for FBG 8 than we do in the plot for FBG 5. This can also be explained by considering the slopes for $\Delta\lambda$ vs. ΔT for these FBGs, which are the reciprocals for the slopes we found in Table 2 with our profile-fitting algorithm. Since FBG 5 has a best-fit slope of $9.39 \frac{\text{pm}}{\text{K}}$ while FBG 8 has a best-fit slope of $29.90 \frac{\text{pm}}{\text{K}}$, we see that the spectrum for FBG 8 will shift more rapidly in λ_B over the same temperature range. This means that the spectrometer data for FBG 8 will cover more bins than will the data for FBG 5. Since transitions between bins result in significant changes in the sweep factor (see Figure 5) and some changes in the SLED spectrum (see Figure 6), we would expect to

see more instances of spectrometer bias affect the underlying Gaussian-like reflectivity curve for FBG 8. This is in exact agreement with what we observe.

Beyond the residual plot, our algorithm's standard deviations in Figure 16 consistently surpass those of the Gaussian fit, especially in the case of FBG 8 (which, based on all previous investigations, seems to most clearly follow a linear response to temperature). In summary, then, we see a major improvement when we use our profile-generating algorithm over the Gaussian fit.

II. KAISER FIT

The Kaiser window is a discrete tapering sequence that is used often in digital signal processing; in appearance, it can often look similar to a Gaussian curve, albeit being very differently defined. To further elaborate on this sequence, we first suppose that M points are desired for the output window, and let n be an integer index for these points such that $-\frac{M-1}{2} \leq n \leq \frac{M-1}{2}$. In this case, the Kaiser window over the M points is defined by:

$$K(n) = \frac{I_0\left(\beta\sqrt{1 - \frac{4n^2}{(M-1)^2}}\right)}{I_0(\beta)} \quad (12)$$

Where I_0 is the zeroth order modified Bessel function of the first kind and β is a parameter that affects the shape of the window (a higher β corresponding to a narrower shape) and the sidelobe attenuation of the Fourier transform of the window.

Investigations with FBGs have shown that a Kaiser peak more precisely models the real spectrum of FBGs that do not possess a sidelobe characteristic; this was observed to be the case since FBG spectra exhibit a shorter decay factor compared to the Gaussian shape [1]. While Kaiser windows are far too computationally expensive for FPGAs due to the presence of the modified Bessel function of the first kind (which, in its analytical expression, involves an infinite sum and the Γ function), we analyze them in this section to see how effective they are compared to the Gaussian fit.

Given the discrete nature of the Kaiser window, however, it is not as straightforward to “fit” a Kaiser peak to three points as it is to fit a continuous Gaussian curve. In order to circumvent this restriction, we use the following procedure:

1. As inputs from the spectrometer, we take a set of wavelengths λ corresponding to the centers of the bins of interest and a set P containing the power measurements of these bins.
2. To simulate continuity, we define a function that takes a parameter β and returns a 0.1-pm-resolution Kaiser window with M points.
 - This parameter M is a constant and is sufficiently large to cover our domain λ (and an extra padding region of 2.5 nm on either side for offsets) with a 0.1-pm resolution.
3. We then parametrize each Kaiser window with three inputs: a multiplicative amplitude A , the abovementioned parameter β , and an offset parameter Δ . We can generate these windows by taking the window from step 2, applying the offset Δ , and multiplying all points by A .
 - When $\Delta = 0$, the Kaiser peak is assumed to be centered in the middle of our wavelength domain λ . A positive Δ (by arbitrary convention) represents a leftwards shift of the Kaiser window.

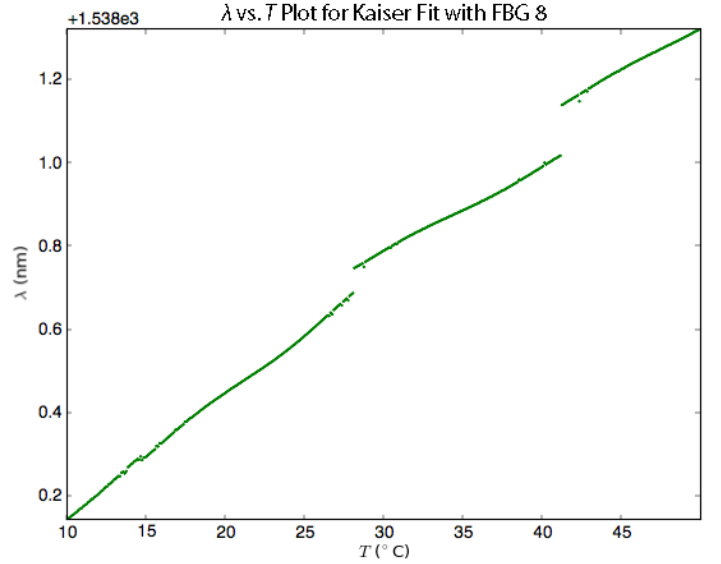


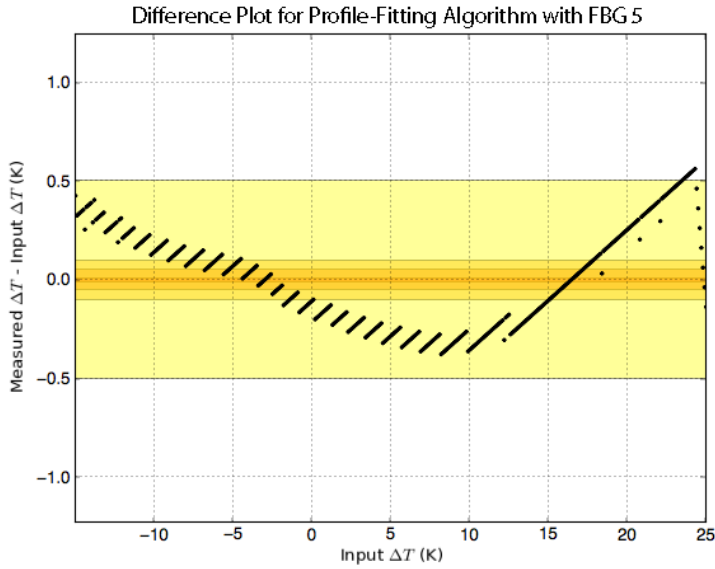
Figure 18: λ vs T plot for Kaiser-fitting method on FBG 8 over a temperature range of $[10.0^\circ\text{C}, 50.0^\circ\text{C}]$. We can clearly see jump discontinuities around $T \approx 28^\circ\text{C}$ and $T \approx 41^\circ\text{C}$.

- The offset (measured in nm) can thus be used to determine the Bragg wavelength with a 0.1-pm resolution.

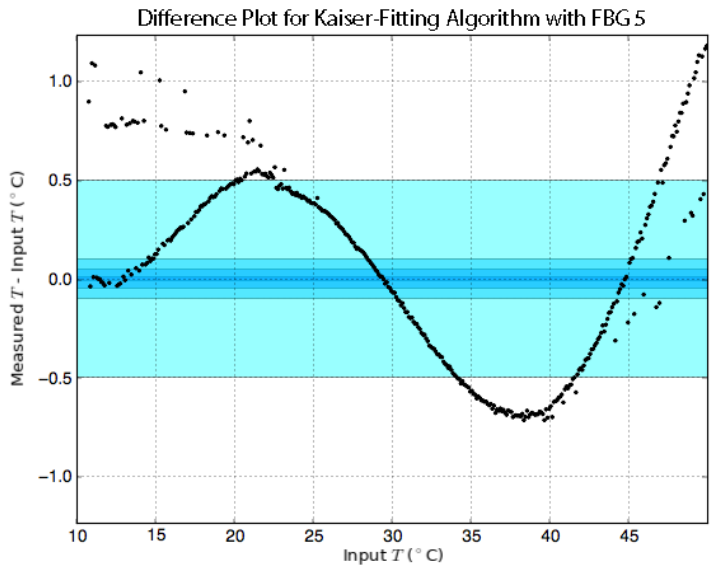
4. Lastly, we make use of SciPy’s “curve_fit” function (a least-squares procedure that works with multiple variables) to fit a window from step 3 to our spectrometer data from step 1. The parameter Δ of the best-fit window can be used to determine the Bragg wavelength.

We compare the results of this Kaiser fitting procedure with our own algorithm on FBGs 5 and 8 in Figure 17, using three bins of spectrometer data as inputs in both cases. For both FBGs, we see that our profile generation algorithm performs significantly better in terms of both its residuals and its standard deviations. In the case of FBG 5, we see an oscillatory pattern in the Kaiser fit’s residuals, which most likely results from spectrometer bias as explained above; this does not match the theoretical prediction of a systematic quadratic pattern. On the other hand, our profile-matching procedure does produce the expected U-shaped plot. Furthermore, our algorithm’s standard deviation of $\sigma = 0.2457$ K significantly outperforms the Kaiser fit’s standard deviation of $\sigma = 0.7290$ K.

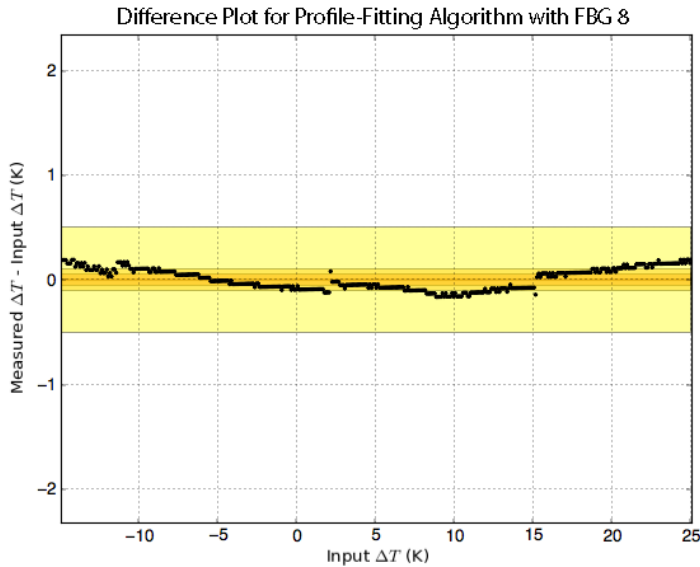
As for the FBG 8 data, our profile generation algorithm clearly demonstrates a more well-behaved residual plot and a lower standard deviation. However, we do observe some interesting patterns in Figure 17d that are worth noting. We can see from this figure that the Kaiser window fit seems to do a relatively good job for temperatures lower than 28°C ; yet over the remainder of the temperature range, our fit yields very inaccurate results. These inaccuracies cannot be explained by the number of bins that are used for the Kaiser



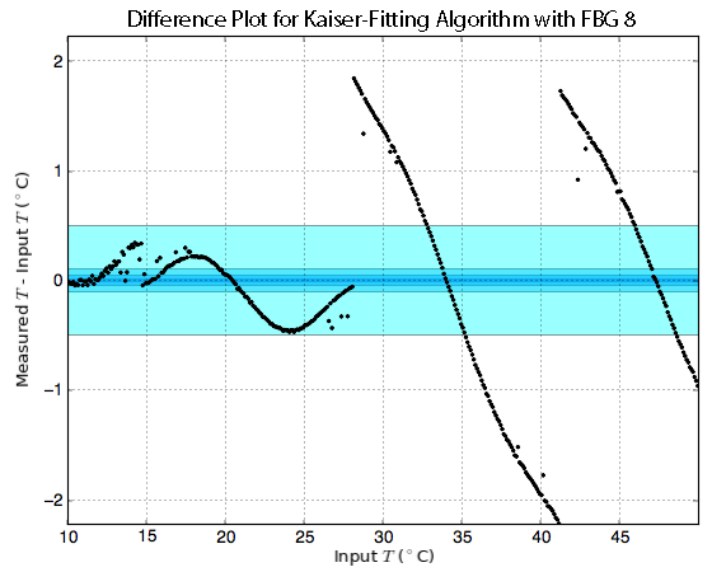
(a) Profile-fitting method with FBG 5 ($\sigma = 0.2457$ K).



(b) Kaiser-fitting method with FBG 5 ($\sigma = 0.7290$ K).



(c) Profile-fitting method with FBG 8 ($\sigma = 0.1009$ K).



(d) Kaiser-fitting method with FBG 8 ($\sigma = 0.8842$ K).


Figure 17: Comparisons between our profile-generation fitting method (orange) and the Kaiser fitting method (blue) for FBGs 5 and 8.

fit; an increase to five bins for the Kaiser fit yielded an even more oscillatory residual plot with a higher standard deviation of $\sigma = 1.0745$ K.


In order to explain these jumps in Figure 17d, we closely inspected the λ vs. T plot for the FBG 8 Kaiser fit, shown in Figure 18. We found that the plot was divided into three somewhat linear sections with discontinuities at $T \approx 28^\circ\text{C}$ and $T \approx 41^\circ\text{C}$; these correspond exactly to the jumps that we see in the residuals. By looking at the actual Kaiser windows generated around these temperatures, we found that these discontinuities occur at points in the spectrometer data where two neighboring bins have relatively high and similar values. At these “double-peaked” points in the data, the least-squares

fitting algorithm understandably makes a jump from a Kaiser window centered around the left peak to one centered around the right peak. This is a particularly noticeable issue with the Kaiser fit since the parameter β allows for a wide variety of different-looking Kaiser windows to be generated; for instance, $\beta = 0$ gives a completely rectangular Kaiser window, while large values of β give a nearly Gaussian curve.

With the clear systematic patterns that we observe in the Kaiser fit, as well as the high standard deviations and jump discontinuities, one might wonder why the Kaiser window is believed to be superior to the Gaussian fit. However, there are certain cases where the Kaiser fit can do a better job; if an FBG has a noticeable Kaiser shape and relatively dimin-



(a) Gaussian-fitting method with FBG 8 ($\sigma = 0.3893$ K).



(b) Kaiser-fitting method with FBG 8 ($\sigma = 0.1800$ K).

Figure 19: Comparisons between the Gaussian fitting method (green) and Kaiser fitting method (blue) for FBG 8 over a temperature range of $[10.0^\circ\text{C}, 28.0^\circ\text{C}]$.

ished sidelobes in its OSA spectrum, it is possible that the Gaussian fit might actually do worse over certain temperature ranges. For example, in Figure 19, we compare the Kaiser and Gauss difference plots for FBG 8 generated over the temperature range $[10.0^\circ\text{C}, 28.0^\circ\text{C}]$ (i.e. avoiding points of major jump discontinuity). From the resulting graphs, we can see that the Kaiser residuals do follow a somewhat regular oscillatory pattern, similar to the pattern in the Gaussian plot. Nonetheless, we see an improved standard deviation of $\sigma = 0.1800$ K in the Kaiser fit, as opposed to the deviation of $\sigma = 0.3893$ K achieved by the Gaussian algorithm over the same temperature range.

III. CDA5 FIT

A centroid detection algorithm (CDA) is a common method of extracting subpixel information from low-resolution data. The method for this algorithm is a very simple weighted average over a range of $(2a + 1)$ bins, where a is a positive integer. If we define B_p as the bin containing the spectrometer's peak data and $R(B)$ as the bin-wise spectrometer response, we have the following formula for the subchannel estimate B_s :

$$B_s = \frac{\sum_{i=B_p-a}^{i=B_p+a} i \cdot R(i)}{\sum_{i=B_p-a}^{i=B_p+a} R(i)} \quad (13)$$

It is common to use a fitting range of five bins (i.e. the CDA5 algorithm) to cover all bins of interest, achieve extra accuracy, and ensure relatively low residuals. Thus, we will follow this convention in our analysis.

In Figure 21, we compare the results of the CDA5 technique with our own algorithm on FBGs 5 and 8, using three bins of spectrometer data for our profile generation and five bins for



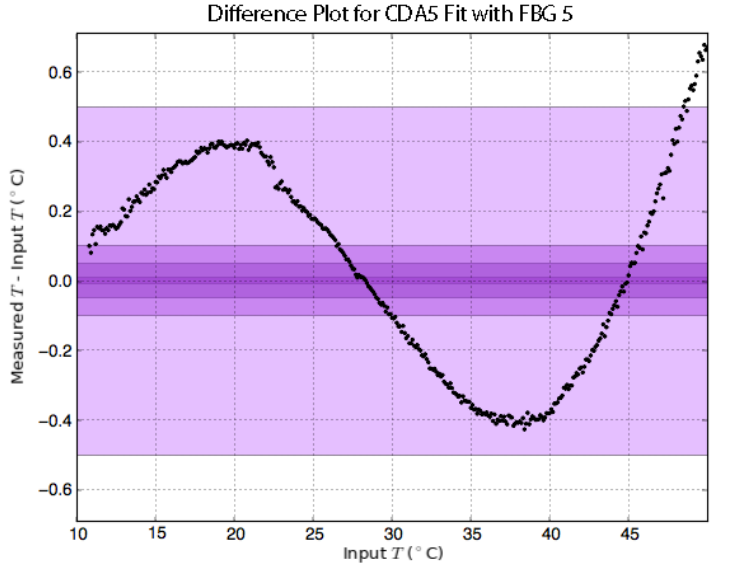



Figure 20: λ vs T plot for the CDA5 procedure on FBG 8 over a temperature range of $[10.0^\circ\text{C}, 50.0^\circ\text{C}]$. We can clearly see jump discontinuities around $T \approx 15^\circ\text{C}$, $T \approx 28^\circ\text{C}$, and $T \approx 41^\circ\text{C}$.




(a) Profile-fitting method with FBG 5 ($\sigma = 0.2457$ K).



(b) CDA5 procedure with FBG 5 ($\sigma = 0.4640$ K).



(c) Profile-fitting method with FBG 8 ($\sigma = 0.1009$ K).



(d) CDA5 procedure with FBG 8 ($\sigma = 0.4722$ K).

Figure 21: Comparisons between our profile-generation fitting method (orange) and the CDA5 procedure (purple) for FBGs 5 and 8. Note that CDA5 makes use of five bins centered around the peak bin B_p .

the CDA5 moving average. With both FBGs, we find that our profile matching method is more successful in terms of both its residuals and its standard deviations. In the case of FBG 5, we see an oscillatory pattern in the CDA’s residuals that is very similar to the pattern in Figure 16b. As one would expect from the simplistic “moving average” approach of centroid detection, this suggests that CDA5 is heavily influenced by spectrometer bias.

In addition, the FBG 8 data in Figure 21d clearly demonstrates jump discontinuities in CDA5’s linear regression model. On closer inspection (see Figure 20), we found that CDA5’s λ vs. T plot featured four somewhat linear sections with discontinuities at $T \approx 15^\circ\text{C}$, $T \approx 28^\circ\text{C}$, and $T \approx 41^\circ\text{C}$. This

pattern is very similar to what we observed in the Kaiser data (Figure 18), and occurs due to difficulties with applying centroid detection on double-peaked data. Nonetheless, one important distinction from the Kaiser plot is that the sections in Figure 20 are qualitatively much more linear than those in Figure 18.

Overall, in terms of both standard deviations and systematic residual patterns, CDA5 is surpassed by the three-point Gaussian procedure in subsection I; this is in agreement with the results concerning resolution gain factors in Figure 5.12 of [1]. Nevertheless, centroid algorithms have the advantage of being far more computationally efficient than Gauss fits.

IV. LPO4 TECHNIQUE

A novel, popular approach for Bragg wavelength determination is the linear phase operator (LPO) technique described in [1]. This method combines a finite impulse response (FIR) filter with linear interpolation to arrive at a subpixel estimate of the reflectivity spectrum's peak wavelength. In this paper, we specifically look at the LPO4 technique, which uses five points to produce its FIR filter (the order of an LPO technique is generally given as $N - 1$, where N is the number of points considered). We choose this particular technique as it uses a low number of bins and was shown to perform sufficiently well compared to higher-order operators [1].

To begin, we cover the theory behind the LPO technique. As is common practice in discrete FIR theory, we first describe an output sequence $y(n)$ as a weighted sum of the current and (a finite number of) previous values of the reflectivity. If we let $R(B)$ represent the bin-wise reflectivity data from the spectrometer, $h(i)$ the impulse response coefficients, and N the number of points considered in our FIR filter, we can define $y(n)$ generally as:

$$y(n) = \sum_{i=0}^{N-1} h(i) \cdot R(n-i) \quad (14)$$

It is a generally known result in FIR theory that a negative-symmetry FIR filter defined over an odd number of points N can be appropriately chosen to design a differentiator of an input sequence ($R(B)$ in this case). By the term “negative-symmetry filter,” we refer to a set of impulse coefficients $h(i)$ such that $h(n) = -h(N-1-n)$ (i.e. the coefficients are opposite-signed symmetric about the center coefficient). For the specific case of the LPO4 technique for FBGs, these coefficients were found in [1] to be $[1, 1, 0, -1, -1]$.

We also specifically choose an odd length to allow for straightforward accounting of the group delay of our FIR (i.e. the average delay of our filter as a function of frequency). When we use a negatively-symmetric linear phase FIR, we observe a constant group delay of $t_{gr} = \frac{N-1}{2}$ samples; this delay was found by observing the negative first derivative of the filter's phase response [1]. The result for t_{gr} determined in [1] is in accordance with linear phase response theory, and is demonstrated in Figure 22 for LPO8. Thus, an odd length N was required to produce an approximation $y(n)$ to the true derivative of any FBG spectrum with a delay of $\frac{N-1}{2}$ samples.

By determining the zero-crossing of $y(n)$ and taking into account the group delay of the LPO filter, we can arrive at a subpixel estimate for the Bragg wavelength. Based on the above reasoning, this zero-crossing has a delay of $\frac{N-1}{2}$ pixels with respect to the peak wavelength $\lambda_{B, \text{pix}}$ from the spectrometer data (this can be seen as a delay of four pixels in Figure 22 for LPO8). Suppose, for LPO($N-1$) (where $N-1$ represents the order), that the intensities around the zero-crossing of $y(n)$ are defined as $\xi_{\text{LPO}(N-1)}$ and $g_{\text{LPO}(N-1)}$. We can then approximate the slope of $y(n)$ as:




Figure 22: LPO8 output sequence $y(n)$ (solid line) for an FBG spectrum (dashed line), adapted from Figure 5.7 of [1]. We can clearly see the group delay (t_{gr}) of $\frac{N-1}{2} = 4$ samples in this image.

$$m = \frac{g_{\text{LPO}(N-1)} - \xi_{\text{LPO}(N-1)}}{\text{pixel}}$$

This allows us to use the secant method and find the subpixel zero-crossing of $y(n)$.

For LPO($N-1$) with a peak bin of B_p and a bin-wise reflectivity spectrum $R(B)$, we clearly have the following for $g_{\text{LPO}(N-1)}$ and $\xi_{\text{LPO}(N-1)}$:

$$\begin{aligned} g_{\text{LPO}(N-1)}(R, B_p) &= y\left(B_p + \frac{N-1}{2}\right) \\ &= -R\left(B_p - \frac{N-1}{2}\right) \cdots -R(B_p - 1) + \\ &\quad R(B_p + 1) + \cdots + R\left(B_p + \frac{N-1}{2}\right) \end{aligned} \quad (15)$$

$$\begin{aligned} \xi_{\text{LPO}(N-1)}(R, B_p + 1) &= y\left(B_p + \frac{N-1}{2} + 1\right) \\ &= -R\left(B_p - \frac{N-1}{2} + 1\right) - \cdots - R(B_p) + \\ &\quad R(B_p + 2) \cdots + R\left(B_p + \frac{N-1}{2} + 1\right) \end{aligned} \quad (16)$$

By applying the secant method, we then see that a sub-channel estimate B_s for the Bragg wavelength can be formulated by:

$$B_s = B_p + \frac{g_{\text{LPO}(N-1)}(R, B_p)}{g_{\text{LPO}(N-1)}(R, B_p) - \xi_{\text{LPO}(N-1)}(R, B_p + 1)} \quad (17)$$

Note that Equation (17) only applies for the condition $R(B_p - 1) < R(B_p + 1)$, in which case we have a group delay of $\frac{N-1}{2}$ samples. On the other hand, for the condition $R(B_p - 1) > R(B_p + 1)$, we would logically have a shorter group delay since the peak is closer to bin $(B_p - 1)$ than it is to bin $(B_p + 1)$. In this case, it was found in [1] that we simply have to decrease our group delay to $\frac{N-1}{2} - 1$ samples in

order to accurately reflect the data; this ends up transforming Equation (17) to:

$$B_s = B_p + \frac{g_{\text{LPO}(N-1)}(R, B_p - 1)}{g_{\text{LPO}(N-1)}(R, B_p - 1) - \xi_{\text{LPO}(N-1)}(R, B_p)} \quad (18)$$

Thus, unlike the previous SDAs that have been considered, a comparison (between $R(B_p - 1)$ and $R(B_p + 1)$) is required in order to implement LPO. To actually convert this sub-channel estimate into a wavelength, we once again make use of the bin boundary data presented in Section III and interpolate. This then allows us to compare the results of this LPO4 technique with our own algorithm on FBGs 5 and 8 in Figure 23. For our own profile generation process, we make use of three bins of spectrometer data. However, LPO4 requires five bins in order to work at all and, due to the caveat in Equation (18), requires at least two additional bins in order to offset appropriately. Thus, we use seven bins of spectrometer data to create a linear regression model for LPO4.

In spite of the fact that LPO4 uses more bins than our profile-fitting method, we can see from Figure 23 that our method is actually superior to LPO4 for both FBGs in terms of both its residuals and standard deviations. On first glance of the FBG 5 standard deviations, this is not immediately obvious as our fitting method only slightly outperforms LPO4 in this criterion. However, the residual plots of FBG 5 make this distinction much more clear. For example, in Figure 23a, we see a U-shaped pattern; this indicates that a quadratic fit might be more appropriate. While a very similar pattern exists in Figure 23b, we can clearly see that spectrometer bias creates additional oscillations in this U shape. These oscillations play a large role in making our standard deviation for FBG 5 better than that of the LPO4 technique.

As for FBG 8, we see from Figures 23c and 23d that our profile-fitting algorithm’s standard deviation surpasses LPO4 by a factor of ~ 1.9 . This large improvement in accuracy results from the fact that FBG 8 covers a large λ range for the same temperature range of [10.0°C, 50.0°C]. Since this translates to more bin transitions, we end up observing more instances of spectrometer bias due to sweep factors; this bias manifests itself in the relatively large number of oscillations in Figure 23d. Our algorithm, on the other hand, displays smaller oscillations since it accounts for much of the spectrometer bias. Nonetheless, we do see some small oscillations that are somewhat similar to Figure 23d. This undesired systematic pattern can in fact be removed if we use more than three bins, as we can see by comparing Figure 13d (seven bins) with Figure 14b (three bins).

In addition, LPO4 might produce higher magnitude residuals due to more inherent theoretical flaws in its approach. The approach outlined by [1] assumed that all spectrometer bins were evenly spaced in the wavelength domain; in such a case, the LPO4 FIR coefficients of [1, 1, 0, -1, -1] determined by Zeh might work more ideally. However, in reality, many modern spectrometers are manufactured with uneven “bin widths.” Combined with the fact that the notion of a bin boundary is not particularly well-defined, this could mean

	Slope (K/nm)	95% conf.
FBG 5 LPO4	107.48	± 0.24
FBG 5 Profiles	107.16	± 0.22
FBG 8 LPO4	33.56	± 0.05
FBG 8 Profiles	33.45	± 0.03

Table 5: Slopes of linear fits ($\Delta\lambda$ vs. ΔT) for FBGs 5 and 8 with 95% confidence intervals, using both profile fitting and LPO4 to produce regression models. Note that LPO4 computes across five channels of input data, while profile matching uses three channels.

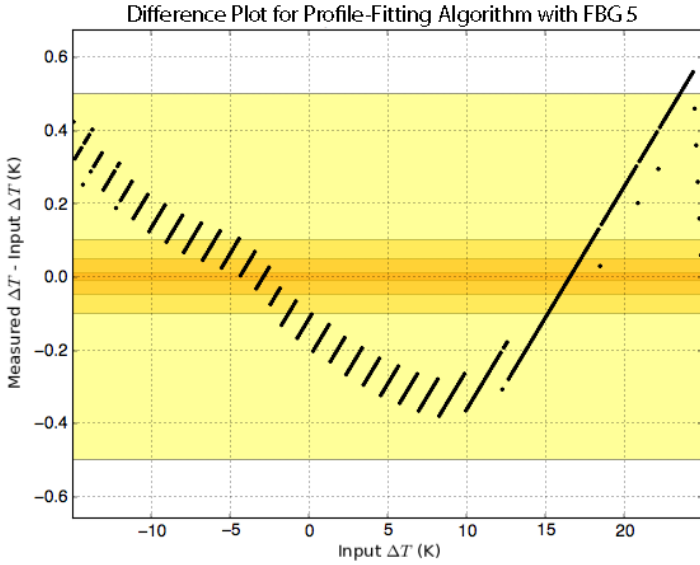
Algorithm by quality	Linear fit residuals (K)	
	FBG 5	FBG 8
Profiles	0.2457	0.1009
LPO4	0.2675	0.1887
Gauss	0.3835	0.3942
CDA5	0.4640	0.4722
Kaiser	0.7290	0.8842

Table 6: Summary of performances of profile-fitting, Gauss and Kaiser fitting, centroid detection, and linear phase operator algorithms.


that the FIR coefficients stated by [1] are inaccurate. Indeed, it might be necessary to apply a spectrometer-specific process that determines the best-fit coefficients for a first-order FIR differentiator; such an approach would also have to impose the constraint of negative symmetry on the resulting FIR filter. However, even though it may be possible to arrive at more suitable FIR coefficients, we still might see systematic oscillations in LPO4’s residual plot as spectrometer bias is difficult to remove.

As a further point of comparison, we examine the best-fit slope parameter for the ΔT vs. $\Delta\lambda$ plots in both LPO4 and our algorithm; 95% percent confidence intervals for these parameters are presented in Table 5. We see that both LPO4 and our profile-fitting algorithm generate similar confidence intervals for each FBG’s Bragg wavelength response; in fact, in the case of FBG 5, these intervals even overlap. This helps attest to the accuracy of our profile generation procedure. Due to the large, noticeably systematic oscillations in LPO4’s residuals, however, we have reason to believe that our best-fit parameters are in fact more accurate. Finally, it is worth noting that we arrive at a more precise estimate of the slope when we use our profile-matching method.


In spite of its oscillatory residuals and theoretical flaws, LPO still performs much more accurately compared to all of the other SDAs described in this section. As we can see from Figure 23, the LPO4 technique succeeds in producing the



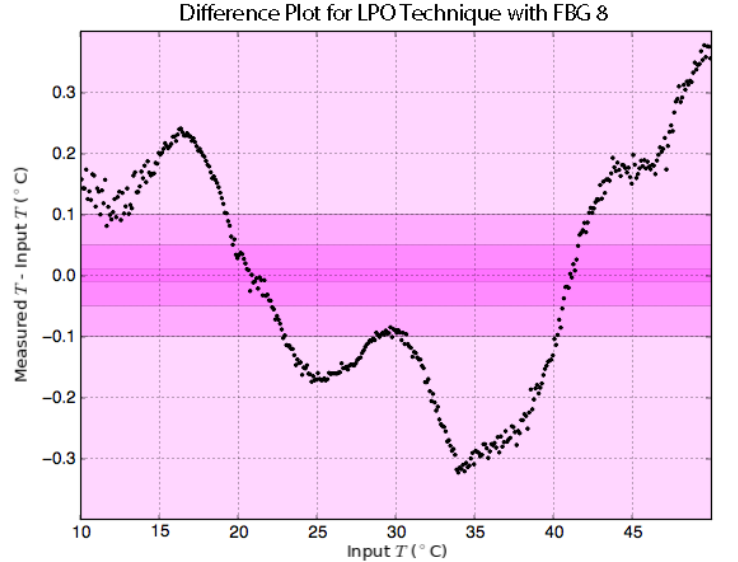
(a) Profile-fitting method with FBG 5 ($\sigma = 0.2457$ K).



(b) LPO₄ technique with FBG 5 ($\sigma = 0.2675$ K).



(c) Profile-fitting method with FBG 8 ($\sigma = 0.1009$ K).



(d) LPO₄ technique with FBG 8 ($\sigma = 0.1887$ K).

Figure 23: Comparisons between our profile-generation fitting method (orange) and the LPO₄ technique (magenta) for FBGs 5 and 8. Note that LPO₄ computes across five channels.

expected systematic U shape in its residuals, and additionally demonstrates relatively low standard deviations. This suggests that the approach of finding the zero-crossing of the differentiator $y(n)$ results in a more robust prediction for the peak of the FBG spectrum, less hindered by spectrometer bias due to our sweep factors and SLED.

A summary of the fit performances of all algorithms discussed in this section is given in Table 6.

V. COMPUTATION COST SUMMARY

To conclude our analysis of other SDAs, we compare the computation costs of determining a Bragg wavelength for

each algorithm presented in this paper; of course, we omit the Kaiser fit in this comparison due to the modified Bessel functions that it requires. For the Gauss and CDA5 algorithms, the number of operations required to calculate B_s can be easily determined by Equations (11) and (13) respectively. For LPO₄, on the other hand, Equation (17) can simplify significantly due to cancellation in the denominator. In the end, LPO₄ requires one comparison, two additions and two subtractions to determine g_{LPO4} , an additional two additions and one subtraction to determine the denominator $g_{LPO4} - \xi_{LPO4}$, and lastly one division and one addition to arrive at a subchannel estimate.

Profiles	Gauss	CDA5	LPO4
	5 × LN		6 × ADD
61 × ADD	3 × ADD	9 × ADD	5 × SUB
92 × MUL	4 × SUB	2 × SUB	1 × DIV
70 × COMP	1 × DIV	1 × DIV	1 × MUL
	3 × MUL	6 × MUL	1 × COMP

Table 7: Comparison of the average computation costs (rounded to the nearest integer) per Bragg wavelength (or offset) determination for subpixel algorithms and our profile-matching procedure. LN = natural logarithm, ADD = addition, SUB = subtraction, MUL = multiplication, COMP = comparison. Note that linear interpolation (in order to determine an actual wavelength) was included for the Gaussian, CDA5, and LPO procedures.

However, once we determine B_s , we must also apply linear interpolation to determine the actual Bragg wavelength due to the uneven bin boundaries. Suppose B_l is the bin immediately to the left of the subchannel estimate B_s , and $\lambda(B)$ is a mapping that takes a bin number and returns the wavelength at the center of that bin. Then the interpolated wavelength is simply given by:

$$\lambda_B = \lambda(B_l) + (B_s - B_l) \cdot (\lambda(B_l + 1) - \lambda(B_l)) \quad (19)$$

This adds an extra addition, two subtractions, and a multiplication operation to our Gaussian, CDA5, and LPO algorithms.

As for our own algorithm, we choose to use the layered search arrangement that produces the lowest r performance parameter (i.e. average number of operations) as defined in Section VII, subsection II; following a brute-force determination process across all FBGs at a reference temperature of 25°C, this was found to be a layered search of (7, 7, 7, 8, 7, p). This search procedure achieved an optimum average combination of 8.93 profile skips and 30.53 profile fits per Bragg wavelength offset determination. Since each profile skip requires one comparison (for peak matching), while each profile fit requires two comparisons (peak matching and maximum quality determination), two additions, and three multiplications (inner product), we were able to determine an average number of operations for each λ_B search.

The results of our computation cost analysis are shown in Table 7. While the values computed for our profile-fitting algorithm are only meant to serve as order-of-magnitude estimates (since the best layered search depends on the spectrometer, FBGs, OSA data, and countless other factors), we clearly see that our algorithm is much more computationally intensive than any of the SDAs we have examined. Moreover, Table 7 does not take into account the fact that we must read profile data from memory. Since these read operations are implementation-dependent and can have variable speeds depending on data storage choices (e.g. USB, RAM, or flash

memory), we have not included them in our table; nonetheless, the costs of these operations cannot be neglected.

In spite of the relatively high number of comparisons our algorithm has to perform to determine a Bragg wavelength shift, it is important to remember that our profile-fitting procedure can be parallelized in FPGA hardware. For example, the floating-point inner product calculations for a given layer can be distributed across DSP (digital signal processor) slices in our actual implementation. In the specific case of a Xilinx® Virtex-5 SX95T FPGA, we have 640 DSP48E slices available for computations. If we make use of the LogiCORE IP Floating-Point Operator (v5.0) to perform single-precision computations (making full usage of the DSP48E slices), we require only two DSP48E slices per addition or multiplication operation. Since our largest layer in (7, 7, 7, 8, 7, p) requires at most eight simultaneous inner product calculations (and hence 24 multiplications followed by 16 additions), we have more than enough DSP48E slices in the SX95T FPGA to parallelize our algorithm for an entire layer in an appropriate manner.

In addition, while the number of operations stated in Table 7 for our algorithm might seem very high, we are only performing a relatively small number of trivial additions, multiplications, and comparisons. With sufficiently fast DSP clocks and the addition of parallelization, it is conceivable that these processes will be performed in sub-millisecond time, which is more than sufficient for most applications. Most importantly, a sub-millisecond execution time will be more than fast enough to keep up with our spectrometer’s data output rate of 12 sample-averaged spectra per second. However, in order to arrive at a more precise estimate for computation time, further research in the field of hardware implementation will be necessary.

Lastly, the improvement in accuracy that results from our produced regression model is a very important factor; in the case of FBG 8, for example, our model has a standard deviation of ~ 0.1 K over a temperature range of [10.0°C, 50.0°C]. This will be very useful for high-precision temperature or strain applications over a large range of input stimulus values.

IX. SUMMARY AND FUTURE WORK

Fiber Bragg gratings are a modern technology used to measure temperature and strain on a fiber-optic cable based on changes in a Bragg wavelength parameter. These gratings have applications in several fields, including non-invasive surgery. Unfortunately, many techniques of extracting the Bragg wavelength parameter – and hence measuring the temperature or strain variables – are limited in accuracy due to low-resolution data with distortions and irregular binning (“spectrometer bias”).

Instead of fitting the data with no reference to spectrometer properties, we apply a “brute-force” approach where we emulate all possible spectrometer responses and store them in precomputed “profiles” that are later fitted to spectrometer output. We measure “sweep factors” representing the

contribution of a wavelength of light to each bin over the whole spectrum, a SLED spectrum that accounts for the light source of our experimental setup, and a high-resolution OSA spectrum of the FBG in question at a specific reference temperature within the temperature range of interest. Profiles are generated by offsetting the high-resolution OSA spectrum by a given wavelength and integrating over each bin's sweep factor, along with the SLED spectrum, to yield a bin-by-bin prediction for the spectrometer output for this wavelength offset. These profiles are generated for the expected range of wavelength variation and are then normalized and stored. Test measurements with the spectrometer were conducted over a temperature range of 10.0°C to 50.0°C and averaged over 250 samples, resulting in lower noise than in practice.

In order to obtain wavelength offsets in real time, incoming data spectra are normalized and compared to profiles, using the inner product of the spectrum and profile as a “quality factor.” This inner product is computed over three bins, centered on the bin containing the most photocurrent, to reduce the influence of noise in low-power bins further away from the 2-3 bins-wide peak. The range of 2200 profiles is searched via “search layers,” which, when optimally arranged, can narrow down on the best profile with few computations. Furthermore, a peak-matching constraint is applied, which requires a given profile's maximum-power bin to match that of the data spectrum. Using a brute-force approach, a layered search was determined that identified the best profile for any given spectrum with ~ 30 inner product comparisons and ~ 9 profiles skipped due to peak-matching.

The wavelength offset variable is a representation of the Bragg wavelength, and should hence be correlated with the temperature and strain variables. Temperature correlations were analyzed to be highly linear – in accordance with the approximate predictions of coupled-mode theory – with linear fit slope parameters having 95% confidence intervals ranging in half-width from ± 0.03 to ± 0.87 K/nm depending on the FBG in question. When using the linear regression parameters to convert the wavelength offset variable to temperature, linear fit residuals ranged from 0.1001 K to 1.3556 K, again depending on the FBG. An accuracy of 0.1 K is within the precision of the temperature controller used.

An analysis was conducted to compare our profile-fitting method against the Gaussian fit, Kaiser window fit, centroid detection algorithm (CDA), and the novel linear phase operator (LPO) technique. None of these algorithms are capable of taking into account distortions due to the spectrometer. Furthermore, they assume evenly-spaced binning and – in the cases of CDA5 and LPO4 – make use of five bins rather than just three; this implies greater susceptibility to noise in low-power bins. A summary of performances is given in Table 6. The profile-fitting algorithm outperforms all other SDAs, decisively beating the Gauss, Kaiser, and CDA5 fitting procedures. Moreover, it improves on LPO4 by up to a factor of ~ 2 depending on the FBG.

However, this improvement in accuracy comes with the drawback of comparatively large computation costs as shown

in Table 7. At the same time, profile fitting is the only algorithm to reach temperature measurement residuals on the order of 0.1 K. In practical applications like medical technology, a real-time implementation of this algorithm is still plausible and reliably gives high-precision measurements over a large range of temperatures.

In order to more thoroughly assess the performance of this algorithm in practice, a real-time implementation should be developed on a parallelizably-programmable device. This will allow for further performance comparisons (e.g. of computational costs and temperature/strain accuracies) against the SDAs discussed above.

X. ACKNOWLEDGMENTS

We would like to thank:

Patrick Leyendecker for mentoring us throughout our analysis and providing us with data and information on this subject. We wish him and his colleges the best of luck with their implementation of remotely-operating robots for non-invasive surgery.

Robert Haslinger for being the driving force behind the experimental setup of this analysis and providing us with measurement data, without which this analysis would have been impossible.

The German Aerospace Center (DLR) for giving us the exciting opportunity to contribute to modern developments in science and technology.

REFERENCES

- [1] Zeh, Thomas (2004). “Optical Fiber Bragg Sensors: Measurement Systems and Signal Processing.” *Dissertation for Technischen Universität München*.
- [2] K. O. Hill and G. Meltz (1997). “Fiber Bragg Grating Technology Fundamentals and Overview.” *Journal of Lightwave Technology*, IEEE: 10.1109/50.618320.
- [3] B. Poulmelic, P. Guenot, I. Riant, P. Sansonetti, and P. Niay (1995). “UV Induced Densification During Bragg Grating Inscription in Ge:SiO₂ Preforms.” *Opt. Mat.*, vol. 4, pp. 441-449.
- [4] T. E. Tasi, G. M. Williams, and E. J. Friebele (1997). “Index Structure of Fiber Bragg Gratings in Ge-SiO₂ Fibers.” *Opt. Lett.*, vol. 22, pp. 224-226.
- [5] W. W. Morey, G. Meltz, and J. M. Weiss (1995). “Recent Advances in Fiber Grating Sensors for Utility Industry Applications.” *Proc. SPIE, Self Calibrated Intelligent Optical Sensors and Systems*, vol. SPIE-2594, pp. 90-98.