WPS1-PR(18) 21228

U Hoefel et al.

Bayesian Modelling of Microwave Radiometer Calibration on the example of the Wendelstein 7-X Electron Cyclotron Emission diagnostic

Preprint of Paper to be submitted for publication in
Review of Scientific Instruments

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

This document is intended for publication in the open literature. It is made available on the clear understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the Publications Officer, EUROfusion Programme Management Unit, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK or e-mail Publications.Officer@euro-fusion.org

Enquiries about Copyright and reproduction should be addressed to the Publications Officer, EUROfusion Programme Management Unit, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK or e-mail Publications.Officer@euro-fusion.org

The contents of this preprint and all other EUROfusion Preprints, Reports and Conference Papers are available to view online free at http://www.euro-fusionscipub.org. This site has full search facilities and e-mail alert options. In the JET specific papers the diagrams contained within the PDFs on this site are hyperlinked

Bayesian Modelling of Microwave Radiometer Calibration on the example of the Wendelstein 7-X Electron Cyclotron Emission diagnostic

Udo Hoefel,[1, a)] Matthias Hirsch,[1, b)] Sehyun Kwak,[1, 2] Andrea Pavone,[1] Jakob Svensson,[1] Torsten Stange,[1] Hans-Jürgen Hartfuß,[1] Jonathan Schilling,[1] Gavin Weir,[1] Johan Willem Oosterbeek,[1] Sergey Bozhenkov,[1] Harald Braune,[1] Kai-Jakob Brunner,[1] Neha Chaudhary,[1] Hannes Damm,[1] Golo Fuchert,[1] Jens Knauer,[1] Heinrich Laqua,[1] Stefan Marsen,[1] Dmitry Moseev,[1] Ekkehard Pasch,[1] Evan R. Scott,[1] Fabian Wilde,[1] Robert Wolf,[1, 3] and the W7-X team

[1)] *Max Planck Institute for Plasma Physics, Greifswald, Germany*
[2)] *Department of Nuclear and Quantum Engineering, KAIST, Daejeon, Korea*
[3)] *Department of Astronomy and Astrophysics, TUB, Berlin, Germany*

(Dated: 5 November 2018)

This paper reports about a novel approach to the absolute intensity calibration of the electron cyclotron emission (ECE) spectroscopy system at Wendelstein 7-X. The ECE radiometer consists of a 32 channel heterodyne radiometer covering the frequency range from 126 GHz to 162 GHz. An absolute calibration of the overall diagnostic including near plasma optics and transmission line is achieved with blackbody sources at LN_2 temperature and room temperature via a hot/cold calibration mirror unit. As the thermal emission of the calibration source is typically a few thousand times lower than the receiver noise temperature, coherent averaging over several hours is required to get a sufficient signal to noise ratio. A forward model suitable for any radiometer calibration using the hot/cold method and a periodic switch between them has been developed and used to extract the voltage difference between the hot and cold temperature source via Bayesian analysis. By using a variance scaling factor a model sensitive adaptation of the absolute uncertainties can be implemented, which will be used for the combined diagnostic Bayesian modelling analysis.

PACS numbers: 07.57.Kp

I. INTRODUCTION

Microwave radiometers are applied in a wide range of research areas, ranging from atmospheric physics[1], radio astronomy[2] to nuclear fusion research[3,4]. A common way to absolutely calibrate the radiometers is the hot/cold calibration method, which uses two reference temperatures to determine the calibration factors[4], although often only a relative calibration is done[5]. Due to the amount of components used in a radiometer it is challenging to get appropriate uncertainties for each component up to the data acquisition system (DAQ). Realistic uncertainties are especially important for Bayesian modelling of multiple diagnostics, as too small or too big uncertainties will artificially shift the result.

In magnetically confined plasmas, electron cyclotron emission (ECE) is widely used to measure the electron temperature with high spatial and temporal resolution. The significance of the measurement is limited by the finite optical thickness of the plasma and black body radiation statistics.

At the optimized stellarator Wendelstein 7-X[6] (W7-X), the Bayesian combined diagnostic modelling is done within the Minerva framework[7]. For the ECE system at W7-X, the hot/cold method mentioned above is used, with the whole calibration unit being located in the W7-X torus hall. A good estimate for the uncertainties of the ECE radiometer calibration is obtained by inverting the forward model of the calibration process described in this paper. This forward model incorporates a variance scaling factor used on the variance of the binned and coherently averaged data, thus finding the uncertainties matching the predictive capabilities of the model. The implementation allows easy automation of the whole calibration procedure.

This paper is structured as follows: Section II introduces the ECE hardware setup at W7-X and the calibration procedure, section III explains the Minerva concepts and classes used to evaluate the obtained calibration raw data, section V shows results from a hot/cold calibration, section IV generalises the model such that it is in principle applicable to any radiometer calibrated with a hot/cold source and a rotating mirror and section VI finally summarizes the results and discusses possible further enhancements.

II. HARDWARE SYSTEM

The ECE system at W7-X measures at the outboard, low field side of the torus at a toroidal angle of $\phi = 7.2°$ near the bean shaped plane. A schematic respresentation is shown in figure 1, wherein the vacuum magnetic flux surfaces are shown for the standard magnetic configuration. The line of sight is chosen perpendicular to the magnetic flux surfaces, therefore maximizing the radial resolution and minimizing Doppler shift from the motion of the electrons parallel to the fieldlines[8]. A wideband op-

[a)] Electronic mail: udo.hoefel@ipp.mpg.de
[b)] Electronic mail: matthias.hirsch@ipp.mpg.de

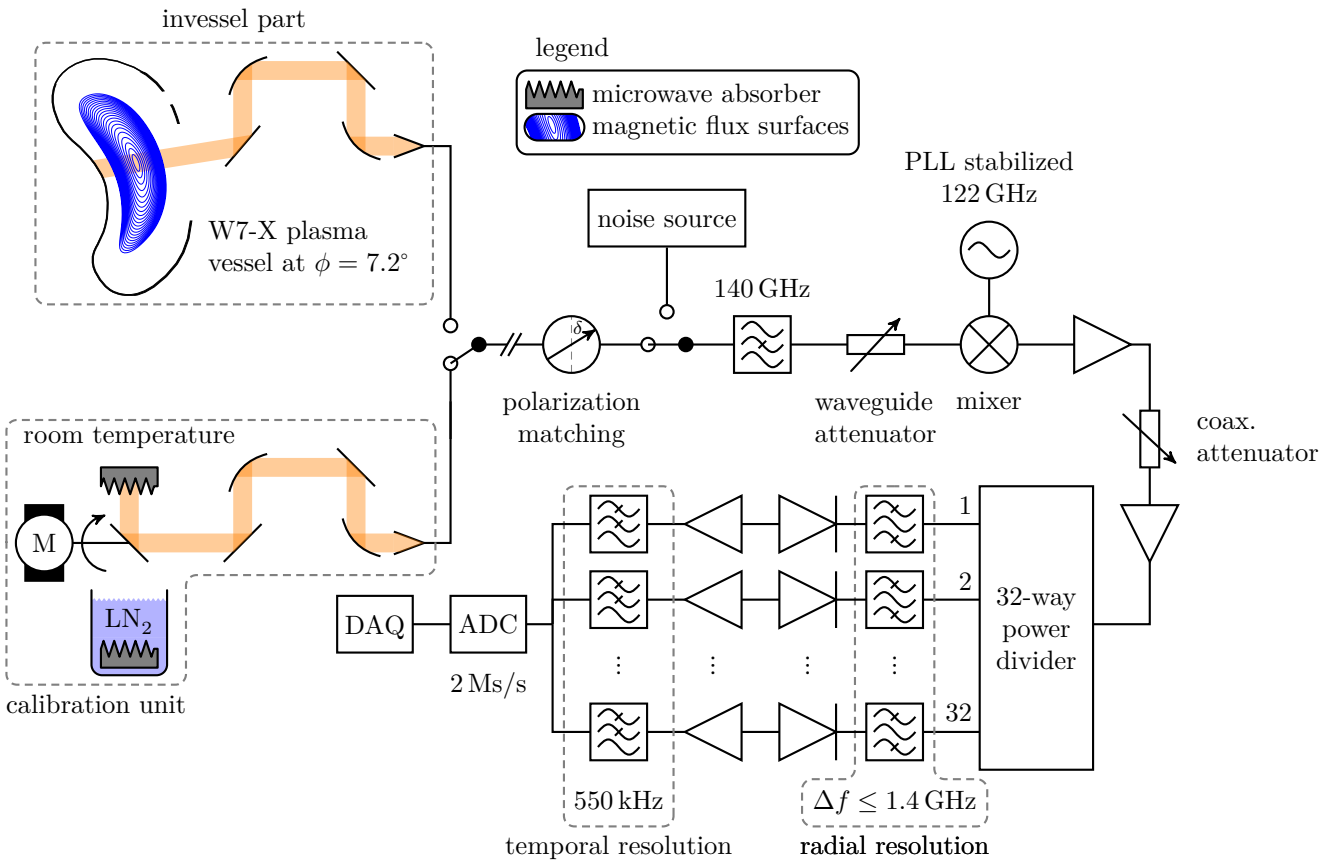

FIG. 1: Schematic drawing of the ECE diagnostic system at W7-X. The calibration unit is below the torus, with identical optical systems, except for the slight tilt of the last plane invessel mirror before the plasma to have the beam be perpendicular to the magnetic flux surfaces. Instead of the inner plasma vessel, the calibration unit has its sightline leading to a steadily rotating, gold-coated mirror, which is surrounded by a highly microwave absorbent foam (ECCOSORB®) at room and liquid nitrogen temperature, see also figure 4. An oversized waveguide transmits the radiation to the detection system, while a polarization matcher is used to get the desired mode. A noise source allows for a fast calibration due to the higher signal to noise ratio. After a notch filter, used to get rid of the 140 GHz ECRH radiation, the signal is mixed down via a phase locked loop (PLL) stabilized local oscillator (LO) to 4 GHz to 40 GHz. Thereafter the power is divided into 32 channels, with a channel specific band-pass filter with center frequencies between 4.4 GHz and 39.6 GHz and a bandwidth of 0.25 GHz $\leq \Delta f \leq$ 1.4 GHz. A 550 kHz low-pass filter ensures that no aliasing occurs, effectively determining the temporal resolution. Finally, the signal passes through an ADC and is recorded by the DAQ.

tical Gauss telescope system with a $1/e^2$ beam intensity width of approximately (20 ± 2) mm within the plasma volume is used. The variation of the width over the sightline in the plasma can be considered small. Exactly the same optical system including the vacuum windows is used below the W7-X experiment in the calibration unit, which is explained in more detail in section II A. The only difference between the two optical systems is the small tilt of the last plane in-vessel mirror to achieve a sightline perpendicular to the magnetic flux surfaces. As stated beforehand, this is necessary to suppress the Doppler shift. A detailed comparison of the in-vessel and calibration unit optics can be seen in figure 2 measured with the Gaussian beam horn with a rectangular monomode waveguide exit used in the first experimental campaign (internally called OP1.1). An oversized waveguide of approximately 23 m length, including two tapers, 11 mitre-bends and one polarization tuner transmits the radiation from either the calibration unit or the plasma to the detection system outside the torus hall, allowing easy access during operation. The overall loss of this transmission line is 13.3 dB. The mode (extraordinary, ordinary or a mixture thereof) can be selected via the polarization tuner. A calibrated noise source can be selected by a waveguide switch instead of the transmission line, which allows to calibrate with a higher signal to noise ratio at the expense of not taking the influence of components in front of the noise source switch into account. A Bragg reflection notch filter with at least 55 dB insertion loss within (140±0.5) GHz and approximately 5.3 dB

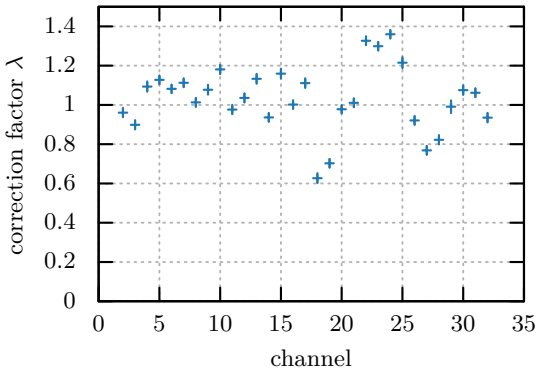

FIG. 2: Factors and estimated uncertainties required to match the rough estimated sensitivities from a hot/cold source calibration. This measurement has been obtained using the in-vessel optic in a lab setup, measuring each channel separately. The origin of the deviations is currently not understood. Tests do not indicate an origin caused by standing waves.

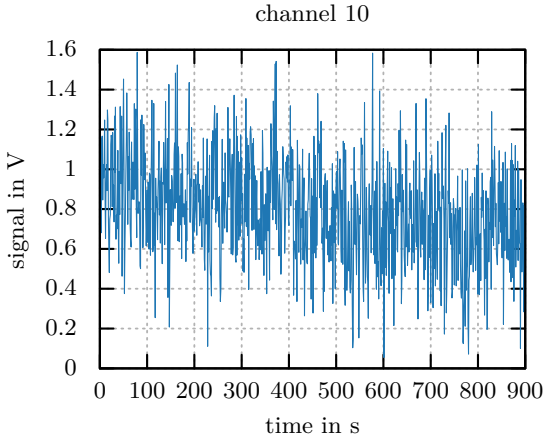

FIG. 3: Example of the drifts of an ECE channel during a few minutes. One can see a notable drift of the average signal, thus measuring hot and cold temperatures separately for several minutes would yield considerably larger uncertainties if the drifts are not strictly linear. Channel 10 has a signal difference between the two reference temperatures of around 4.7 mV. Only every 1000^{th} point is shown.

insertion loss outside was used to block non-absorbed 140 GHz electron cyclotron resonance heating (ECRH) radiation[9]. Afterwards, the signal is down-converted to 4 GHz to 40 GHz via a phase locked loop (PLL) stabilized local oscillator (LO) at 122.06 GHz[10]. Via power dividers the power is split into 32 channels, which subsequently are band-pass filtered with a bandwidth between 0.25 GHz and 1.4 GHz, chosen to adapt the radial resolution to the expected typical optical thickness of the plasma resulting in a resolution of about 1 cm[8]. After that, the signal passes through predetection amplifiers and the detection diodes. Highly linear postdetection amplifiers with a variable gain and an adjustable DC offset allow to choose a reasonable signal amplitude for each plasma discharge, therefore making maximum use of the range of the analog-to-digital converter (ADC). Subsequently, a low-pass filter with a 3 dB point at 550 kHz determines the maximum temporal resolution and ensures that no aliasing occurs. Finally, a 16 bit ADC with a sampling rate of up to 2 Ms/s is used, before the data is stored in the central, immutable W7-X database by the data acquisition system (DAQ).

A. Calibration Unit

Many ECE radiometers that are absolutely calibrated use a rotating blade to switch between two reference temperatures[5,11]. Another method to switch between the reference temperatures is given by a rotating mirror, as described by Hartfuß et al.[4], and is sketched in figure 4. The advantage of these two methods in contrast to just recording data at one reference temperature for several minutes and then at another reference temperature for several minutes (both without using a chopper) is the decreased sensitivity to drifts of the measurement signal, which, if the drifts are not strictly linear, would change the ratio of the bit signal corresponding to the reference temperatures over time. The rotating mirror and the subsequently applied conditional averaging act as a bandpass which suppresses drifts on timescales larger than a rotation period. As drifts on timescales > 10 s are not negligible at W7-X (see figure 3) despite having the electronics in a temperature controlled rack, the rotating mirror method has been chosen for W7-X as the temperature control allows only for temperature stability on the order of 1 K to 2 K. The advantage of the rotating mirror over a rotating chopper lies in the better symmetry of the intermediate temperatures that are measured when radiation from multiple radiation sources at different temperatures is collected. It should be noted that drifts on the magnitude observed here pose a considerable problem for long term plasma operation – either further measures to suppress drifts have to be taken, or a regularly repeated offset determination within a discharge has to be performed. For W7-X, it is planned to repeatedly close the shutter in the planned 30 min plasma discharges to correct the offset.

The calibration unit of W7-X contains a gold-coated brass mirror rotating with approximately 3.6 Hz, see figure 4. Arranged cylindrically around the mirror, a microwave absorber guarantees a black body emitter at room temperature T_{RT}, which is kept at (294.45 ± 3.5) K in the torus hall. However, a small part at the lower side of this cylinder is cut to allow the observation of a stainless steel container thermally insulated by styrofoam. The inner wall of the stainless steel container is

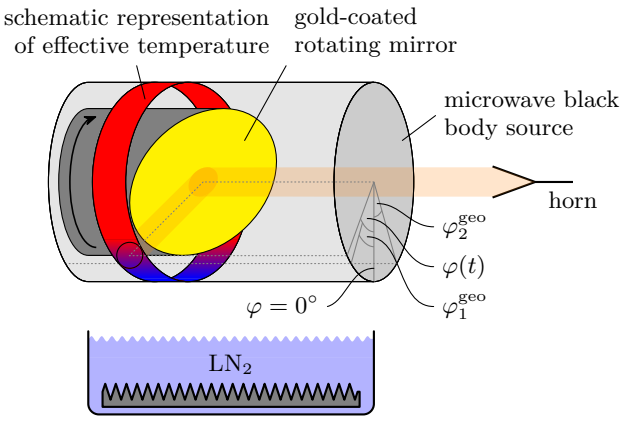

FIG. 4: Schematic drawing of the ECE calibration unit. Not shown is the optical system between the rotating mirror part of the calibration unit and the microwave antenna, which is identical to the invessel optical system. The path that is "seen" by the microwave antenna is shown in light orange. The microwave antenna is characterized by a Gaussian beam. The beam is reflected at the gold-coated mirror and finally "sees" either room or liquid nitrogen temperature, the latter being produced by a liquid nitrogen tank underneath the rotating mirror. Due to the finite size of the beam, the effectively measured temperatures are smeared out at the hot/cold edges, as different parts of the beam "see" different temperatures.

lined with a microwave absorber. The cold reference temperature is not directly given by the temperature of the liquid nitrogen, $T_{\mathrm{LN_2}} = (77.2 \pm 0.5)\,\mathrm{K}$, as water vapour (assumed to be at $T_{\mathrm{H_2O}} = (280 \pm 10)\,\mathrm{K}$ with an uniform emissivity of $0.01 < \varepsilon_{\mathrm{H_2O}} < 0.03$) accumulates above the liquid nitrogen reservoir. Moreover, the temperature of the mirror needs to be taken into account (corresponding to room temperature described above, with an emissivity $0.01 < \varepsilon_{\mathrm{mirror}} < 0.03$). This leads to an effective temperature difference of about 205 K that is used in further calculations. Details in the effective temperature estimation are given in section IV A.

III. MINERVA IMPLEMENTATION

To obtain the physical quantities that are of interest for the evaluation of an ECE calibration, namely the calibration factor of each individual channel (that is the inverse sensitivity) and the effective beam width, the calibration procedure itself is modeled. To evaluate multiple ECE channels in a consistent way, a forward model predicting the different channel sensitivities has to be used. Using Bayes' formula[12] the sensitivities and the effective beam widths can be inferred by

$$\underbrace{P(F|D)}_{\text{posterior}} = \frac{\overbrace{P(D|F)}^{\text{likelihood}}\overbrace{P(F)}^{\text{prior}}}{\underbrace{P(D)}_{\text{evidence}}}, \tag{1}$$

wherein $P(\cdot)$ represents a probability, F the free parameters to be inferred and D the data. A priori knowledge about the free parameters is encoded in the *prior* probability distributions. The probability of the measured data, given the free parameters, is called the *likelihood*. The normalization factor in the denominator, often called *evidence*, is important for model comparison. The *posterior* is the term we want to calculate, as it yields the probability distribution of the sought-after free parameters given the data. As the evidence does not change the shape of the posterior, one can neglect that term if one is not interested in comparing models explicitly, for which the different models would have to be normalized with a penalty on complexity (that is, by applying Occams razor).

Minerva is a Java based general Bayesian modelling framework used among several other large scale experiments around the world, for example at the Joint European Torus[13] (JET) and the Mega-Ampere Spherical Tokamak[14] (MAST) and, as the main inference framework, at W7-X.

Graphical models are a powerful tool to describe the conditional dependency structure of a probabilistic model[15]. A Minerva graphical model is a Bayesian network, more precisely, a directed acyclic graph. It consists of nodes, which are connected via arrows, that are used to evaluate, for example, the calibration processes in dependence of their parent nodes. Do note that these models can be used to generate sample predictions that may be used for example for neural net training sets[16] as well as inversions based on different techniques, such as maximum a posteriori (MAP) and Markov Chain Monte Carlo (MCMC) methods.

The following paragraphs shed some light at some of the basic concepts used for the implementation of the data analysis at W7-X.

A. Datasource

A Minerva datasource is a node in a graphical model. It should be the only experiment specific node in the whole graphical model to allow an easy adaptation of the model to another experiment by exchanging the datasource. More specifically a datasource loads the data from an immutable database and does some preprocessing, if necessary. Furthermore, it provides metadata like lines of sight or frequencies.

In the case of the datasource used for the evaluations presented here, the preprocessing consists of the following steps: In the beginning, the chopper channel data is

used to automatically identify the periods, subtract the mean of each period to avoid an overestimation of the uncertainties and bin each period to avoid influences by variations of the motor speed. Finally a coherent averaging is done to get a signal with a better signal to noise ratio.

Due to the small signal to noise ratio, many rotations of the rotating mirror need to be acquired to extract enough information for meaningful results. The preprocessing is necessary, as working with the full data is unpractical due to the sheer volume of raw data. An example of such a binned signal for a channel with good sensitivity can be seen in figure 8.

B. Graphical Models

For comparison of different modelling approaches, different descriptions of the calibration process were implemented. The following properties of the model have been combined, creating a total of eight structurally different models: i) Evaluation of individual channels respectively multiple channels simultaneously (indicated from here on with the keywords 'single' respectively 'multi'), ii) with and without the use of a channel specific scaling factor α_i to scale the variance of the observed conditionally averaged signal and iii) by having a channel specific beam width or a beam width fixed for all channels by a scaling following

$$w(f) = w(140\,\mathrm{GHz})\sqrt{140\,\mathrm{GHz}/f}, \tag{2}$$

as expected from broadband Gauss telescope optics, indicated from here on with the keywords 'shared' respectively 'individual'.

It should be noted that technically the models with a variance scaling factor α_i describe a continuum of assumed uncertainty models. The variance scaling factor α_i is indicative of the explanatory power of the assumed model. A value close to 1 means that the uncertainty model explains on average most of the observed data spread, larger values are indicative of underestimated uncertainties and vice versa. Do note that if the uncertainty model is assumed to be identical to the one that produced the data and there are still deviations from the observed values, they stem from structural inaccuracies of the model predicting the data. While this method does not allow arbitrarily shaped uncertainty distributions to be inferred, it does allow for a scaling of the uncertainty distribution, thus providing *uncertainties that match the predictive power of the whole model*. Too large uncertainties are penalized by Occams razor as the model can explain more data and the probability density is thus lower. The deviation of the predicted data from the observed data puts a penalty on the smaller uncertainty side. In principle, this scaling factor could be applied to each datum individually, which would allow, for the models presented here, for an angularly resolved information on how well the model predicts the observed data. This could potentially be used to automatically detect the structural parts of the model that have the largest impact on badly predicted data, and adjust them.

The advantage of evaluating all channels simultaneously is the consistency gained for channel independent parameters, namely the beginning and ending of the hot temperature source, and, depending on the model, the beam width. The disadvantage that comes with evaluating multiple channels simultaneously lies in the curse of dimensionality – the evaluation time increases notably. Exemplarily, the graphical model for the multiple channel case with scaling factors α_i and channel specific beam widths is shown in a simplified way in figure 5. The single channel evaluation will be compared with the multi channel evaluation in section V. The evaluation strategy in general is the following: first get reasonable estimates for the starting parameters via the datasource, then do a maximum a posteriori (MAP) inversion (for example using the Hooke and Jeeves pattern search algorithm[17]), of which the result in turn is used as a starting point for the MCMC inversion with a sufficient burn in and an adaptive Metropolis adapter[18] that is deactivated once the MCMC chain is stable. The burn in is a period at the beginning of an MCMC in which all samples are discarded from further analysis. It is not strictly necessary, but it helps in the interpretation of the result, as otherwise, if one starts in an unlikely position of the posterior distribution and does not run the chain long enough, the result will be biased towards these unlikely values. Do note that the MCMC in use is essentially a black box MCMC, thus the only way to ensure convergence instead of just pseudo-convergence is by using sufficiently long runs[19]. Therefore, all MCMCs have been run with a burn in of at least 1 million iterations, until the MCMC traces of the free parameters did not show notable drifts of the running average of the logarithm of the probability density function. This allows for a realistic estimation of the uncertainties of the drop of the measured signal associated with the decrease in radiation temperature as produced by the switch from room temperature to the cold temperature. It has to be noted that the measured signal stored on the database is the raw bit signal, which is also what the model predicts. The logarithm of the probability density function, which is a measure for how probable a specific realisation is, of an MCMC chain of the full model, and example plots of each the calibration and scaling factor traces are shown in figure 6. This evaluation is fully automatized and writes the results to the central W7-X database, making automatic calibration overnight feasible.

C. Limitations

Practical limitations are given for example by the number of bins as well as the number of channels. If the sampling rate notably exceeds 1 kHz the evaluation slows down considerably as well, as the data that has to be

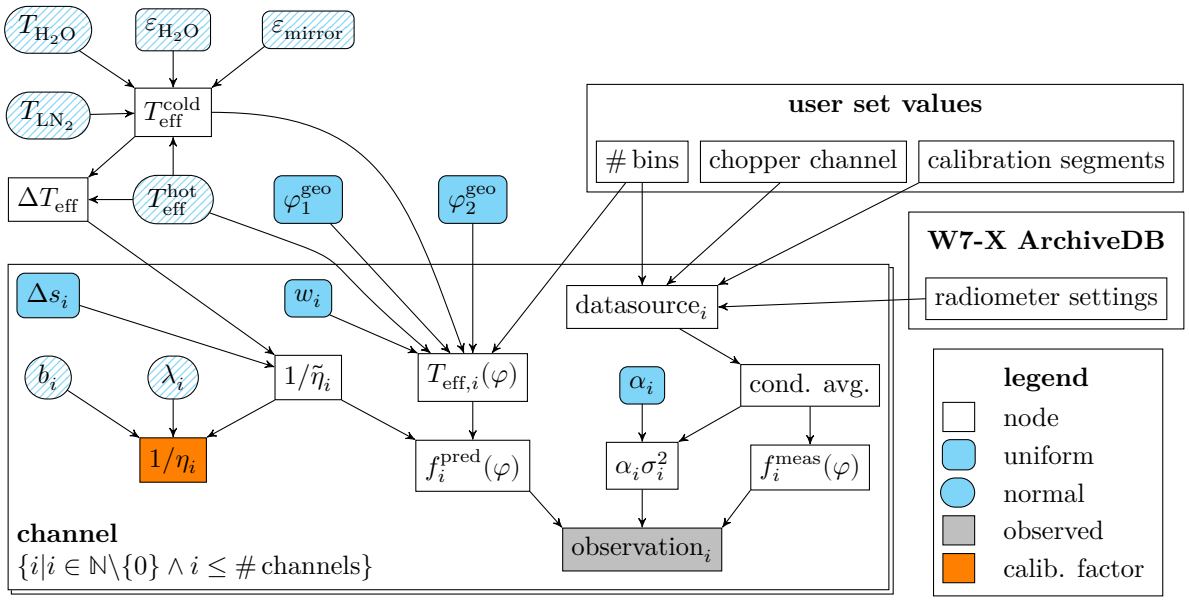

FIG. 5: The simplified Minerva graph showing the dependencies used in the model. Free parameters have a cyan background, while the shape specified in the legend corresponds to the used prior distribution. Do note that the nodes with a hatched background do sample completely randomly from their distribution, thus propagating the corresponding uncertainties without increasing dimensionality by parameters which do not influence the prediction. This approach was used for the temperatures and emissivities (see appendix). T_{LN_2} represents the liquid nitrogen temperature, T_{H_2O} the temperature of the water vapour above the liquid nitrogen, ε_{H_2O} the emissivity of the water vapour, T_{RT} room temperature and ε_{Mirror} the emissivity of the mirror. ΔT is the effective temperature difference between the hot and the cold source. The start of the hot source phase of the signal is described by φ_1^{geo}, the end by φ_2^{geo} (see figure 4. Together with the beam width of the microwave antenna characteristic w_i, the number of bins and the effective temperatures the predicted effective temperatures $T_{eff,i}(\varphi)$ can be calculated. The calibration sensitivity $1/\tilde{\eta}_i$ as calculated from the bit change Δs_i and the temperature difference ΔT_{eff} can be multiplied with $T_{eff,i}(\varphi)$ to obtain the measurement signal prediction $f_i^{pred}(\varphi)$ for channel i. The observation consists of the conditionally averaged and binned signal $f_i^{meas}(\varphi)$ in bit and the variance scaled by a factor α_i. In combination with the bit to Volt conversion factor, b_i, and the factor taking the measured differences between the invessel and calibration optics into account, λ_i, the calibration factor $1/\eta_i$ can be calculated.

loaded for the preprocessing increases accordingly. A practical number of bins can be determined automatically, which gives a value close to the average number of data points per mirror rotation. As there are computationally expensive steps involved for each bin, increasing the number of bins also leads to an increase of required computation time. The full 98D model includes all 32 channels and the scaling of the variance, and takes roughly 230 hours with a Intel® Xeon® CPU E5-2660 v4 @ 2 GHz on a virtualized linux server. By evaluating each channel independently the required CPU time is reduced to about 21 hours, such that with full parallelization the evaluation time can go down to around 40 minutes. This simplified model comes at the price of a generally smaller consistency and larger uncertainties for the weaker channels. However, the differences are negligible for reasonably strong channels (differences in the calibration factor are typically below 1 %). For weaker channels the difference can reach about 10 %. This is due to the stronger channels keeping the geometrical factors more or less fixed, such that the impact on channels where the geometrical information is more concealed in noise profit the most. The single channel evaluation routine provides a pragmatic approach to obtain calibration factors if time requirements prohibit the full model use.

IV. MODEL FORMALISM

This section introduces the required formulæ, first for W7-X and then for a generalized model with arbitrary geometries and effective temperatures.

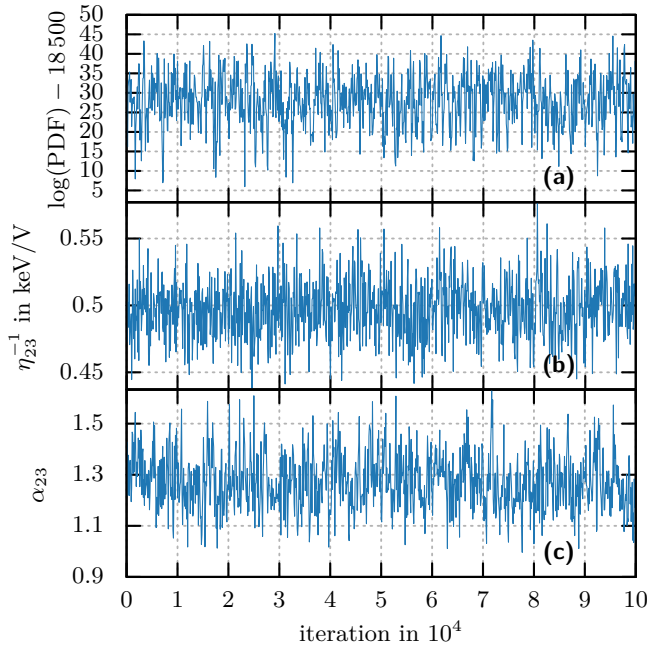

FIG. 6: **(a)** shows the logarithm of the probability density function of an MCMC chain of the full model with variance scaling factors. **(b)** and **(c)** show the MCMC chain of the calibration factor η_{23}^{-1} and the variance scaling factor α_{23}. One can see that the values do not show jumps to a notably different phase space part, indicating that the chain converged. Only every 100^{th} point is shown to keep the plot data volumes at a reasonable size.

A. Model for W7-X

To calculate the temperature difference between the room temperature and the effective cold temperature, one needs to take the influence of the water vapour emissivity and the mirror emissivity into account. The effective temperature after the radiation passes through the water vapour forming above the liquid nitrogen is calculated via

$$T_{\text{eff}}^{\text{vapour}} = T_{\text{LN}_2} + \varepsilon_{\text{H}_2\text{O}}(T_{\text{H}_2\text{O}} - T_{\text{LN}_2}), \tag{3}$$

where $\varepsilon_{\text{H}_2\text{O}}$ is the emissivity of the water vapour. The effective cold temperature after the mirror is given by

$$T_{\text{eff}}^{\text{cold}} = T_{\text{eff}}^{\text{vapour}} + \varepsilon_{\text{mirror}}(T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{vapour}}). \tag{4}$$

The emissivity of the mirror only adds to the effective cold temperature and not the effective hot temperature, as the mirror is already at the same temperature as the effective hot temperature. In case the hot reference temperature would be at a temperature different from the mirror temperature it would have to be taken into account there as well. Note that this works because the signals are recorded in the frequency range where the Jeans law approximation holds. This allows the calculation of the channel specific calibration sensitivity (respectively, the sensitivity η_i scaled with the attenuation and post-detection amplification as chosen for this channel during the calibration)

$$\tilde{\eta}_i = \frac{\Delta s_i}{\Delta T_{\text{eff}}} = \frac{\Delta s_i}{T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{cold}}}, \tag{5}$$

wherein Δs_i represents the signal change in bits caused by the temperature difference ΔT_{eff}. The index i denotes the channel specific parameters. Moreover, the channel specific expected effective temperature, depending on the mirror position, is given by

$$T_{\text{eff},i}(\varphi) = \frac{\int \hat{T}_{\text{eff}}(\hat{\varphi}) g_\varphi(w_i, \hat{\varphi}, \varphi)\, \mathrm{d}\hat{\varphi}}{\int g_\varphi(w_i, \hat{\varphi}, \varphi)\, \mathrm{d}\hat{\varphi}}, \tag{6}$$

where $\hat{T}_{\text{eff}}$ is the effective temperature at a given mirror position without taking the finite width of the Gaussian beam, defined by the horn characteristic, into account,

$$\hat{T}_{\text{eff}}(\hat{\varphi}) = \begin{cases} T_{\text{eff}}^{\text{cold}} & \text{for } \hat{\varphi} < \varphi_1^{\text{geo}} \\ T_{\text{eff}}^{\text{hot}} & \text{for } \varphi_1^{\text{geo}} \leq \hat{\varphi} \leq \varphi_2^{\text{geo}} \\ T_{\text{eff}}^{\text{cold}} & \text{for } \hat{\varphi} > \varphi_2^{\text{geo}} \end{cases}, \tag{7}$$

and g_φ the weight of each $\hat{T}_{\text{eff}}$ assuming a perfect Gaussian beam horn characteristic.

This in turn allows to predict the measured bit signal in dependence of the mirror angle,

$$f_i^{\text{pred}}(\varphi) = \frac{T_{\text{eff},i}(\varphi)}{\tilde{\eta}_i}. \tag{8}$$

In practice, it can be useful to subtract the mean of the predicted signal to avoid dependencies on the effective temperatures where not necessary, see appendix.

Scaling the variance σ_i^2 with a channel specific parameter α_i on the observation node allows for a realistic estimation of the uncertainties by taking the predictive capability of the model into account. The reason for not blowing up the variance as much as possible can be understood in the following way: Practically, changing the variance changes the assumed uncertainty model. If the uncertainties get bigger, the model could explain more data, thus decreasing the probability density of the whole model. Therefore, this should also work with a uniform prior, given that one does an MCMC. As the variance appears in the likelihood, this also changes the maximum posterior. Further information to the variance scaling can be found in section III B.

The calibration factor is given by

$$\frac{1}{\eta_i} = \frac{\lambda_i b_i G_i}{\tilde{\eta}_i} 10^{-(R_{\text{RF}} + R_{\text{IF}})/10}. \tag{9}$$

This takes the following quantities into account: The measured differences between the invessel and calibration optics, λ_i, the measured bit to volt conversion factor, b_i, the post detection amplification chosen during calibration, G_i, the setting of the waveguide attenuator right in

front of the radiometer in decibel, R_{RF}, and the setting of the attenuator at the intermediate frequency device in decibel, R_{IF}. For calibrating plasma measurements one needs to rescale this factor with the appropriate gains and attenuator settings used during the measurement. The reason for including this branch in the model is twofold: i) it allows direct extraction of the sought after quantity, without having to implement separate uncertainty propagation for λ_i and b_i and ii) preparing for future evaluations of multi diagnostic calibration factors, in which case one can simply extend this branch by supplying ECE raw plasma data to get an electron temperature profile that can also be supplied for example via Thomson scattering[20]. This would not be a simple cross calibration, but would rather combine the diagnostic specific calibration models and their corresponding plasma forward models, thus taking all information optimally into account. The calibration factors obtained that way are inherently consistent within the frame of the model.

In the Bayesian formalism this leads to (for the model shown in figure 5)

$$P(\boldsymbol{\varphi}^{\mathrm{geo}}, \boldsymbol{\Delta s}, \mathbf{w}, \boldsymbol{\alpha}|\tilde{\mathbf{D}}) = \frac{P(\tilde{\mathbf{D}}|\boldsymbol{\varphi}^{\mathrm{geo}}, \boldsymbol{\Delta s}, \mathbf{w}, \boldsymbol{\alpha}) P(\boldsymbol{\varphi}^{\mathrm{geo}}, \boldsymbol{\Delta s}, \mathbf{w}, \boldsymbol{\alpha})}{P(\tilde{\mathbf{D}})}, \quad (10)$$

with

$$\boldsymbol{\varphi}^{\mathrm{geo}} = (\varphi_1^{\mathrm{geo}}, \varphi_2^{\mathrm{geo}}), \quad (11)$$
$$\boldsymbol{\Delta s} = (\Delta \mathrm{s}_1, \ldots, \Delta \mathrm{s}_n), \quad (12)$$
$$\mathbf{w} = (w_1, \ldots, w_n), \quad (13)$$
$$\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n), \quad (14)$$
$$\tilde{\mathbf{D}} = (\mathbf{D}_{\mathrm{ECE},1}, \ldots, \mathbf{D}_{\mathrm{ECE},n}), \quad (15)$$

where n corresponds to the number of ECE channels, $\boldsymbol{\varphi}^{\mathrm{geo}}$ to the angles at which the central line of sight switches from the hot source to the cold source and vice versa, $\boldsymbol{\Delta s}$ to the channel specific change in the bit signal observed when switching the temperature sources, $\mathbf{w}$ to the channel specific Gaussian beam width, $\boldsymbol{\alpha}$ to the channel specific variance scaling factors and $\tilde{\mathbf{D}}$ to the channel specific conditionally averaged measured data. Note that the other models differ, for example by not using the variance scaling factors.

B. Generalised model

The model described previously uses some simplifications that can easily be dropped to generalise the model. For instance, one can drop the assumption that the problem is one dimensional and that there are only two reference temperatures. This allows easy extension to three (or more) reference temperatures for example by adding a hot ceramics hot source. Switching to cylindrical coordinates is a sensible approach for a geometry similar to the one presented here, thus introducing z along the horizontal axis of the cylinder shown in figure 4, and r as the radius. Assuming $i \times j$ reference temperatures leads to a definition of the effective temperature $\hat{T}_{\mathrm{eff}}(\hat{\varphi}, \hat{z}) = T_{\mathrm{eff},ij}$, where $T_{\mathrm{eff},ij}$ is the effective temperature valid for $\varphi_i^{\mathrm{geo}} \leq \hat{\varphi} \leq \varphi_{i+1}^{\mathrm{geo}}$ and $z_j^{\mathrm{geo}} \leq \hat{z} \leq z_{j+1}^{\mathrm{geo}}$. The second dimension is represented by $\hat{z}$. As in the model described above, $T_{\mathrm{eff},ij}$ can be the result of multiple layers contributing to the effective temperature at the selected coordinates, such that a dependency on r might occur as well. The calibration sensitivity $\tilde{\eta}(T)$ does not necessarily have to be linear, however, one will have to use free parameters for the temperatures and emissivities in this case, as the prediction will no longer be independent of these parameters. In general, any instrument function g can be used to calculate the appropriate weighted effective temperature that the radiometer would see by looking at (ϕ, z), therefore

$$T_{\mathrm{eff}}(\varphi, z) = \frac{\iint \hat{T}_{\mathrm{eff}}(\hat{\varphi}, \hat{z}) g(\hat{\varphi}, \varphi, \hat{z}, z)\, \mathrm{d}\hat{\varphi}\, \mathrm{d}\hat{z}}{\iint g(\hat{\varphi}, \varphi, \hat{z}, z)\, \mathrm{d}\hat{\varphi}\, \mathrm{d}\hat{z}}, \quad (16)$$

which in combination with the calibration sensitivity allows the calculation of the prediction.

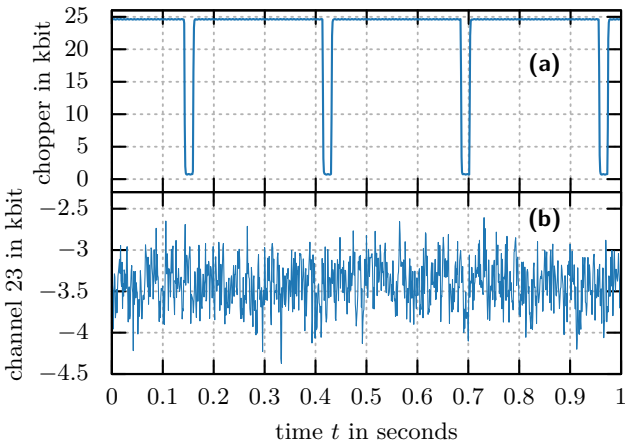

FIG. 7: **(a)** shows the chopper signal measured by a photo diode, **(b)** the raw signal of channel 23, which is a sensitive channel. The signal-to-noise ratio is typically of the order of 1/50. The background signal results from the broadband noise of the IF amplifiers right after the mixer, subsequently measured for each frequency bin by the detector diode.

V. RESULTS

A typical excerpt of the calibration timetrace for a sensitive channel can be seen in figure 7b, while the corresponding chopper signal is shown in figure 7a. There are three points to consider: i) the chopper signal does not correspond to the full width of the cold phase, the real hot/cold duty cycle is approximately 0.2, given by the calibration unit geometry, increasing that value further

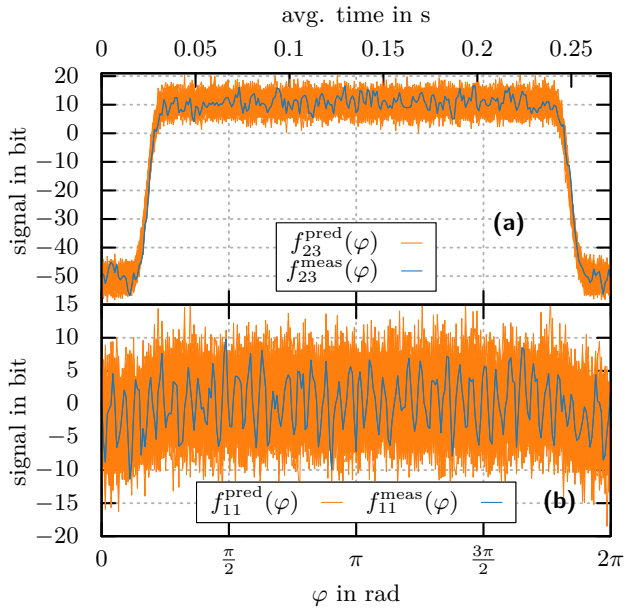

FIG. 8: **(a)** shows the sensitive channel 23, while **(b)** depicts the weakly sensitive channel 11. $\varphi = 0$ corresponds to a mirror position "looking" at the center of the liquid nitrogen cold source. One can see the measured (blue) conditionally averaged and binned signal for both channels. The predictions (orange) are 100 MCMC samples each, taken after the free parameters of the graph have been set to their mean values.

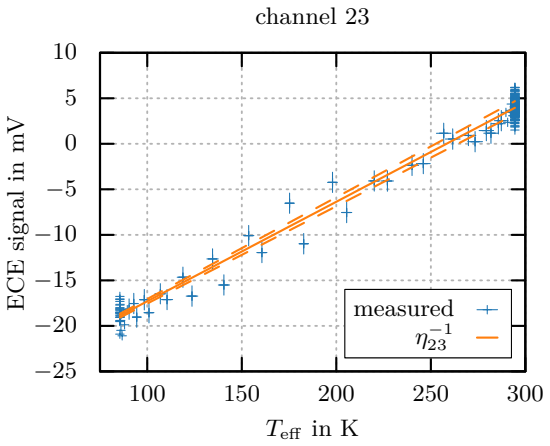

FIG. 9: The voltage values in dependence of the effective temperature. Do note that the plotted voltage uncertainties correspond to $b_{23}\sqrt{\alpha_{23}\sigma_{23}^2}$ with b_{23} the bit to Volt conversion factor. The temperature uncertainties correspond to the uncertainties given by the priors.

would require a significant modification of the calibration unit, ii) even for the most sensitive channel the signal difference associated with the chopper channel cannot be seen directly, confirming that more elaborated analysis techniques are necessary and iii) no relevant drift within one rotation period can be observed. As said before, notable drifts were seen in some cases on timescales on the order of 10 s (see figure 3). The conditional averaged signal in figure 8 supports the conclusion that there is no relevant drift within one period. The blue curve in figure 8a) corresponds to the measured and subsequently conditional averaged signal of the strong channel 23, $f_{23}^{\mathrm{meas}}(\varphi)$, while the orange curves are samples from the graph that has been set to the mean values obtained from the previously run MCMC inversion, so these are the predictions, $f_{23}^{\mathrm{pred}}(\varphi)$. Correspondingly, figure 8b) shows the weakly sensitive channel 11. It is important to note that these samples are calculated from the model that allowed the scaling of the prediction variance, so that the uncertainties match the predictive capability of the model. Each orange point in figure 8 corresponds to a predicted effective temperature scaled by the calibration sensitivity, with the offset of a whole period being removed.

The measured bit values at each φ value are illustrated in figure 9, where the bit values have been scaled to represent voltages. It should be noted that the uncertainties on the voltage axis are scaled with the channel specific variance scaling factor α. For the sake of clarity error bars are only shown in the intermediate temperature range. The orange curve shows the sensitivity with its uncertainties as calculated from the graphical model. Remaining deviations might be caused by 50 Hz noise (or its higher harmonics) that are not completely notched out by the bandpass filter properties of the conditional average. The plot highlights the advantage of this analysis method: While no other radiometer calibration approach known to the authors uses the data that is taken when the horn pattern collects radiation from more than one reference temperature, this method allows to estimate the effective temperature (and corresponding uncertainties) reducing the overall uncertainty and predicting the frequency dependent beam width. However, one has to keep in mind that this is valid only as long as the Jeans law approximation is valid. For the radiometer and the reference temperatures used here this is a very good approximation.

The kernel density estimation of the calibration factors given by the MCMC is shown in figure 10. One of the strengths of an MCMC based evaluation is that one can get posterior distributions that are non-Gaussian as well. However, in the case shown here a Gaussian fit is a reasonable approximation to the posterior for the strong channel 23, while for channel 11 the posterior deviates notably from a Gaussian distribution. For the sake of simplicity the results of a Gaussian fit are used for all higher level analysis, although it would mark a gain in consistency, especially for the weak channels, to use the asymmetric uncertainties originating from the MCMC. As the model predictive capability is taken into account,

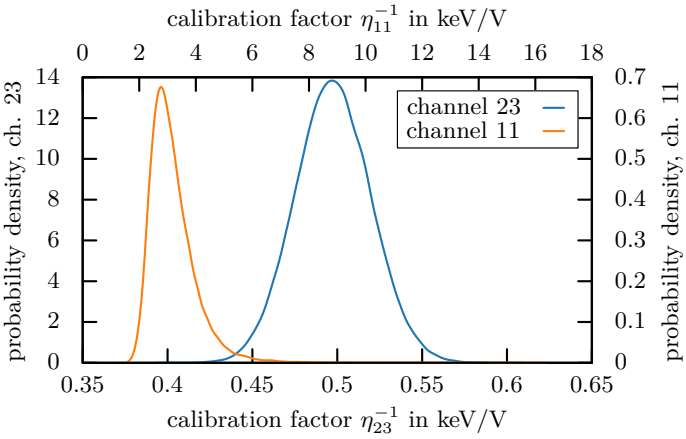

FIG. 10: The kernel density estimates of the MCMC samples of the calibration factors of the strong channel 23, η_{23}^{-1}, and the weak channel 11, η_{11}^{-1}. It is clearly visible that the distribution of the weak channel deviates from a normal distribution by having a pronounced tail towards larger calibration factors. In contrast, the stronger channel closely follows a normal distribution.

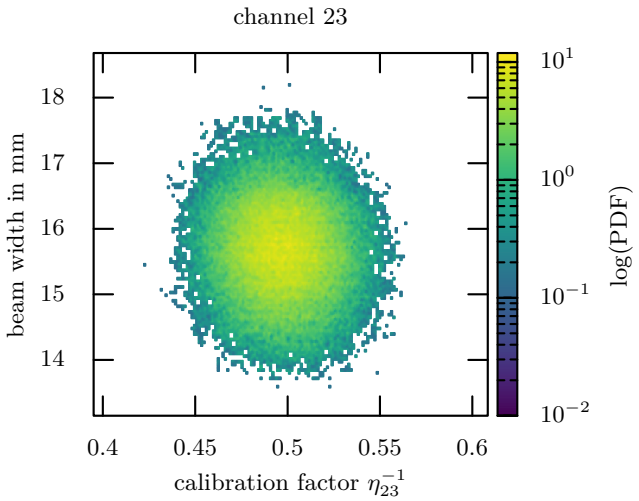

FIG. 11: The posterior distribution for the sensitive channel 23 with the single channel evaluation model. One can see that the distribution is reasonably close to a Gaussian distribution.

an appropriate estimation of the calibration factor uncertainties is achieved.

Figure 11 shows the posterior distribution for the sensitive channel 23 for the two physically interesting quantities, the beam width and the calibration factor. The distribution is reasonably close to a Gaussian distribution.

An example of a posterior revealing more complex relationships between two parameters can be seen in figure 12. Therein, the PDF given the angle that determines the beginning of the hot reference temperature measurement phase and the calibration factor is plotted. One can clearly see the deviation from a Gaussian distribution. It should be noted that the channel chosen here shows a relatively low signal to noise ratio. For channels with a better signal to noise ratio, the distribution resembles again a Gaussian distribution.

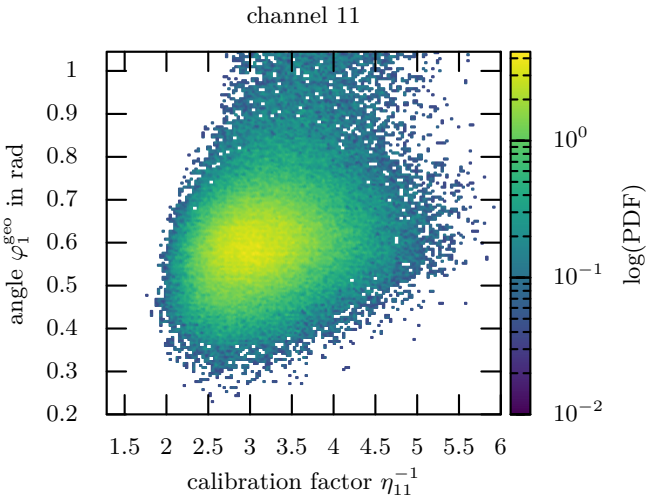

FIG. 12: Example for a posterior revealing the complex interaction between the beginning of the hot temperature measurement phase φ_1^{geo} and the calibration factor η_{11}^{-1} using the single channel evaluation.

A. Comparison of single and multi channel evaluation

Figure 13a shows the calibration factor, which is the inverse sensitivity, for each channel. One can see that the calibration factors vary over more than two orders of magnitude. A single mixer is used for the whole spectrum to allow for a better correlation analysis[10]. This is unusual as many ECE systems use multiple mixers to avoid frequencies above 18 GHz after mixing[21]. The single mixer approach leads to intermediate frequencies up to 40 GHz which need to be detected. The low sensitivities for higher frequencies might at least partially originate from different cables used for frequencies below 18 GHz (corresponding to frequencies below 140 GHz in the shown spectrum) and above 18 GHz. The cable frequency response damps higher frequencies more. The conversion efficiency of the extreme broadband mixer also drops for frequencies above 140 GHz. Individual diode sensitivities are expected to play an important role as well.

Figure 13b shows the uncertainties of the different

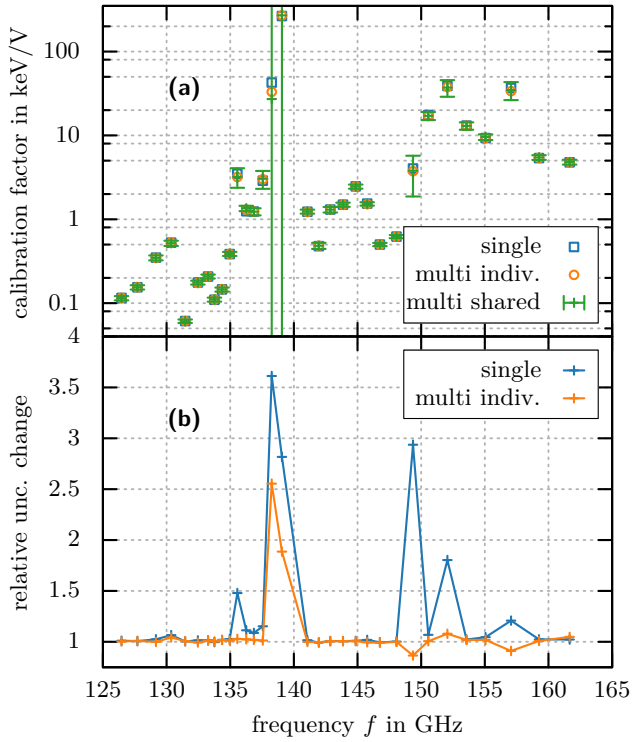

FIG. 13: **(a)** shows the frequency dependent calibration factors as obtained for single channel evaluations and the multi-channel analysis with and without a shared free parameter for the beam width. The uncertainties are only shown for the multi-channel analysis with a shared (and appropriately scaled) beam width, as this is the reference value for the relative uncertainty changes shown in **(b)**. There, the standard deviation of the specified models is normalised by the standard deviation of the multi shared model. The given uncertainties correspond to one standard deviation as calculated from the MCMC samples. One can see that the single channel evaluation has the largest uncertainties, as expected. For channels with low sensitivities the reductions in the uncertainties that occur when switching from single channel evaluations to multi channel evaluations can be significant, as the more sensitive channels provide information about φ_1^{geo} and φ_2^{geo}. Using a common w reduces the uncertainties further, but less drastically.

models, normalised to the multi shared model. Going from the single channel evaluation to a combined model yields substantial decreases in the calibration factor uncertainties of insensitive channels, although the uncertainties for these channels remain very large. This phenomenon is most likely caused by the additional information about the geometrical properties φ^{geo} that is mainly provided by stronger channels, helping the less sensitive channels to determine the begin and end of the hot/cold phases. Using a single beam parameter leads for a few channels to a small shift of the calibration factor, also

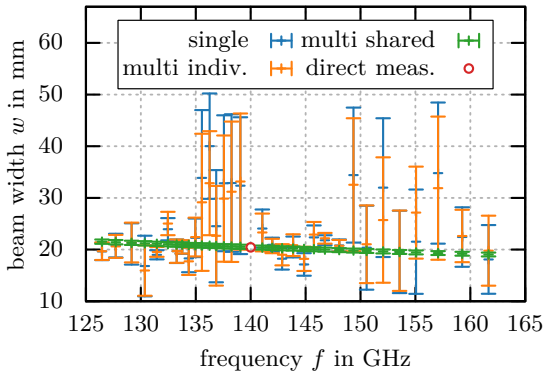

FIG. 14: The inferred Gaussian beam width of the microwave antenna characteristic for each channel. The direct measurement at 140 GHz was done with a 440 point 2D scan of the beam intensities roughly 37.5 cm away from the minimum beam waist.

reducing the uncertainties slightly, but less drastic than the switch from the single channel evaluation to a multi channel evaluation model.

Figure 14 shows the inferred intensity Gaussian beam width for each ECE channel. The beam width has been measured in the lab with a 140 GHz source attached at the receiver end of the antenna to be roughly (20.4 ± 0.2) mm at a distance of 37.5 cm away from the minimum beam waist, inferred with 10 million iterations of an MCMC of a simple Gaussian squared forward model that adapts the prediction uncertainties as well. This does not exactly match the distance at which the microwave foam is located relative to the minimum beam waist (which would be roughly 26 cm to 32 cm), but due to the small divergence of the beam width the introduced error is small. One can see that switching from the single channel analysis to a combined channel analysis slightly decreases the uncertainties for some less sensitive channels. In most cases the beam width shifts slightly towards values closer to the directly measured width. If only a single beam width, scaled according to equation 2, is used, the uncertainties get drastically reduced. The measured width is roughly 40 µm away from the predicted value by the model with a single beam parameter, with prediction uncertainties on the order of 0.4 mm.

The variance scaling factors for the different channels are shown in figure 15. One can see that the values are not too far away from 1, which indicates that the most relevant physic effects are considered. These values were reduced from values typically around 2.6 at the begin of the first W7-X experimental campaign by two changes of the setup: i) the horn was changed to include all polarizations and ii) the container for the liquid nitrogen had a round aperture that was changed to a rectangular aperture. The round aperture system was more sensitive to misalignments of the mirror-antenna system. Indeed an offset of around 3.5 cm of the beam on the mirror

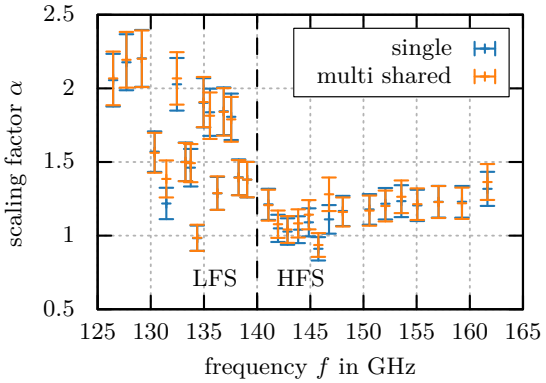

FIG. 15: The inferred variance scaling factors α for the covariance is a measure of the model uncertainty for each channel. For the 'multi individual' case there is no visible difference to the 'multi shared' case, thus it has been omitted for the sake of clarity.

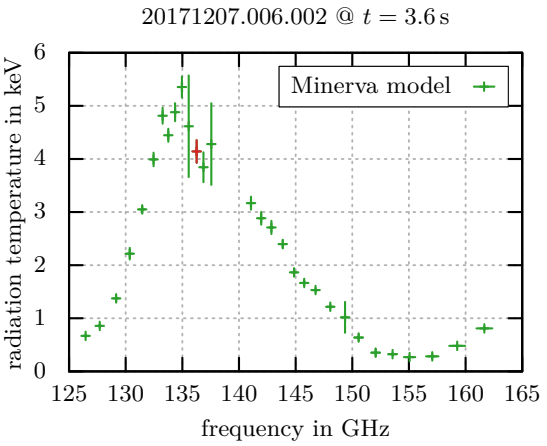

FIG. 16: Exemplary radiation temperature spectrum from an on axis ECRH W7-X plasma discharge, calculated with the calibration factors as obtained from the Minerva model that incorporates variance scaling. The channel marked with red corresponds to the timetrace shown in figure 17.

was measured for the calibration used for the first experimental campaign, but the asymmetry introduced in the signal was not visible to the naked eye. As the physics for each channel should be similar, α should have similar values for the different channels. A notable difference between the α values is an indication that different physics effects play a role, or at least that these different effects are of different importance. One could expect that for channels with a small sensitivity electronics effects are of larger importance, therefore changing the variance scaling. Notable differences between the models that couple the beam width of different channels directly and those that do not, and between the individual and multi-channel evaluations are not observed. The values on the low field side (below 140 GHz) scatter more and tend to be larger. This indicates that the uncertainties are underestimated for these channels, respectively, that the explanatory power of the model is *smaller* than for channels on the high field side. A potential source for this behaviour can be found in hardware issues, respectively, implicit assumptions that are violated more strongly for low field side channels, although currently no such problem is known to the authors.

B. ECE Spectra

From the calibration procedure radiation temperature spectra can be derived from the measurements done during a plasma discharge. An example is shown in figure 16. The data originates from a 3.8 s long plasma discharge that was heated on axis with electron cyclotron resonance heating[23] (ECRH). The ECRH power was 2.5 MW in the first phase and was increased to roughly 5 MW shortly after pellet fuelling started. The line averaged electron density as measured by a single channel dispersion interferometer[24] rose during pellet fuelling up to about $7 \times 10^{19}\,\mathrm{m}^{-2}$.

Channel 15 (138.26 GHz) and 16 (139.06 GHz) show a very low sensitivity, leading to radiation temperatures above 20 keV and uncertainties of several hundred percent. Consequently, they were omitted in this plot. For channel 16 this is expected, as the channels frequency band locates it in the slope of the notch filter. Above 155 GHz contributions from the third harmonic X mode emission start to play a role. The reconstruction of the electron temperature profile from the spectrum will be discussed in a separate publication. Figure 17 shows a comparison of the ECE timetrace of a channel close to the core to a Thomson scattering channel close to the core. The deviations in the first second are probably caused by the filters of the Thomson scattering system that lead to big uncertainties for electron temperatures above around 7 keV. The remaining seconds of the discharge show a good agreement for a wide range of electron densities and multiple power levels.

VI. CONCLUSIONS

The Bayesian Minerva model developed for the multi channel consistent calibration of a generic microwave radiometer provides the sensitivities in an automated fashion, insensitive to signal drifts on timescales > 1 s. By using an explicit model for the calibration, the analysis can be done in a more formalized way. This also allows to obtain feedback on how well the modelled physical system is understood.

Moreover, it allows to obtain non-Gaussian posterior distributions for the calibration factors, although for the sensitive channels in the radiometer studied a Gaussian

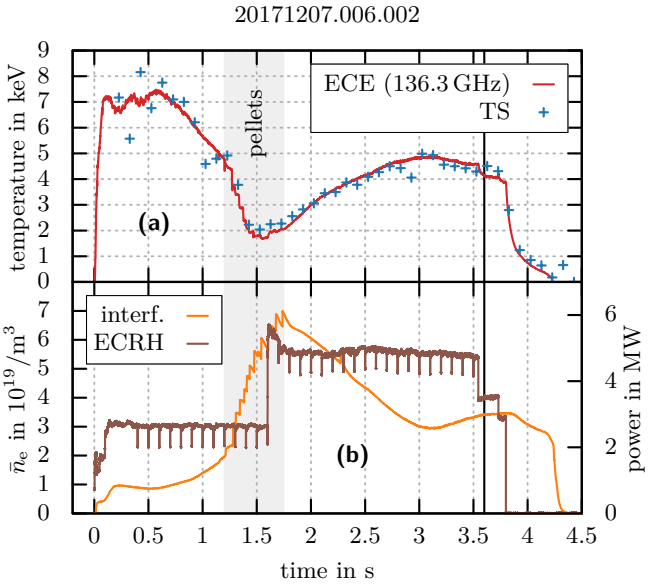

FIG. 17: **(a)** shows the calibrated signal of an ECE channel close to the core, compared with a central channel from the Thomson scattering system[20] for the plasma discharge also shown in figure 16. The line at 3.6 s indicates the spectrum shown in figure 16. **(b)** depicts the ECRH and line averaged density from the single channel dispersion interferometry. The ECRH blips are necessary for the collective Thomson scattering diagnostic[22].

distribution provides an excellent fit to the posterior distribution.

The use of a variance scaling factor allows to obtain uncertainties matching the predictive capability of the model, which, to the authors knowledge, has not been used before at all.

The ECE spectra obtained from typical W7-X plasmas do not show unexpected or unphysical features. The uncertainties are for reasonably sensitive channels typically on the order of 6 %. While the high dimensional (98D) model provides in principle the highest consistency, a comparison to the much quicker and parallelized single channel evaluation (4D for each channel) yields only little differences, justifying the use of the simpler approach.

Further improvements could be achieved by applying neural networks, as described in Pavone et. al.[16]. It also would be interesting to use the calibration model in combination with a plasma model containing predictions for Thomson scattering and ECE, guaranteeing consistent calibration factors across different diagnostic systems. Forward models of plasma ECE that use radiation temperatures provided by a hot/cold load calibration method can be found for example in Rathgeber et al[25].

ACKNOWLEDGMENTS

The authors wish to thank Karsten Ewert for his technical work support. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

Appendix: Temperature dependence of the prediction

For a specific channel i the predicted signal in bit is described via equation (8). Given a linear sensitivity, this equation can be rewritten to

$$f_i^{\mathrm{pred}}(\varphi) = T_{\mathrm{eff},i}(\varphi)\tilde{\eta}_i \tag{A.1}$$

$$= \frac{\Delta s_i T_{\mathrm{eff},i}(\varphi)}{T_{\mathrm{eff}}^{\mathrm{hot}} - T_{\mathrm{eff}}^{\mathrm{cold}}}. \tag{A.2}$$

As we do subtract the offset of the measured signal, we end up with

$$f_i^{\mathrm{pred}}(\varphi) - \bar{f}_i^{\mathrm{pred}} = \Delta s_i \left(\frac{T_{\mathrm{eff},i}(\varphi)}{T_{\mathrm{eff}}^{\mathrm{hot}} - T_{\mathrm{eff}}^{\mathrm{cold}}} - \underbrace{\left(\beta - \frac{\gamma}{n}\right)}_{\text{offset}} \right), \tag{A.3}$$

where

$$\beta = \frac{T_{\mathrm{eff}}^{\mathrm{hot}}}{T_{\mathrm{eff}}^{\mathrm{hot}} - T_{\mathrm{eff}}^{\mathrm{cold}}},$$

γ : number of entries for the cold vector,

n : number of entries for the cold and the hot vector.

As all effective temperatures are between $T_{\mathrm{eff}}^{\mathrm{cold}}$ and $T_{\mathrm{eff}}^{\mathrm{hot}}$, they can be expressed by $\beta - \delta$ with $\delta \in [0, 1]$. Thus, the predicted signal can be written as

$$f_i^{\mathrm{pred}}(\varphi) - \bar{f}_i^{\mathrm{pred}} = \Delta s_i \left(\beta - \alpha - \left(\beta - \frac{\gamma}{n}\right) \right) \tag{A.4}$$

$$= \Delta s_i \left(\frac{\gamma}{n} - \alpha \right). \tag{A.5}$$

$$\tag{A.6}$$

Thus, the predicted signal does not depend on the absolute values of the effective temperature for a linear sensitivity.

[1] R. H. Dicke, R. Beringer, R. L. Kyhl, and A. B. Vane, Physical Review **70**, 340 (1946).
[2] W. A. Imbriale, S. Weinreb, and H. Mani, in *2007 IEEE Aerospace Conference* (2007).
[3] E. Mazzucato, *Electromagnetic Waves for Thermonuclear Fusion Research* (World Scientific Publishing Company, 2014) Chap. 9.
[4] H. J. Hartfuß and T. Geist, *Fusion Plasma Diagnostics with mm-Waves: An Introduction* (Wiley-VCH, Weinheim, 2013).
[5] X. Liu, H. L. Zhao, Y. Liu, E. Z. Li, X. Han, C. W. Domier, N. C. L. Jr., A. Ti, L. Q. Hu, and X. D. Zhang, Review of Scientific Instruments **85**, 093508 (2014).

[6]C. Beidler, G. Grieger, F. Herrnegger, E. Harmeyer, J. Kisslinger, W. Lotz, H. Maassberg, P. Merkel, J. Nührenberg, F. Rau, J. Sapper, F. Sardei, R. Scardovelli, A. Schlüter, and H. Wobig, Fusion Technology **17**, 148 (1990), https://doi.org/10.13182/FST90-A29178.

[7]J. Svensson, A. Werner, and JET EFDA contributors, in *2007 IEEE International Symposium on Intelligent Signal Processing* (2007).

[8]S. Schmuck, H.-J. Hartfuss, M. Hirsch, and T. Stange, Fusion Engineering and Design **84**, 1739 (2009), proceeding of the 25th Symposium on Fusion Technology (SOFT-25).

[9]D. Wagner, W. Kasparek, F. Leuterer, A. Marcor, F. Monaco, M. Munich, E. D. Rijk, H. Schutz, J. Stober, M. Thumm, and A. Zeitler, in *2015 40th International Conference on Infrared, Millimeter, and Terahertz waves (IRMMW-THz)* (IEEE, 2015).

[10]C. Fuchs and H. J. Hartfuss, Review of Scientific Instruments **72**, 383 (2001), https://doi.org/10.1063/1.1309005.

[11]J. L. Ségui, D. Molina, G. Giruzzi, M. Goniche, G. Huysmans, P. Maget, and M. Ottaviani, Review of Scientific Instruments **76**, 123501 (2005).

[12]D. Sivia and J. Skilling, *Data Analysis*, 2nd ed. (Oxford University Press, 2006).

[13]J. Svensson, O. Ford, D. C. McDonald, A. Meakins, A. Werner, M. Brix, A. Boboc, M. Beurskens, and JET EFDA Contributors, Contributions to Plasma Physics **51**, 152 (2011).

[14]M. J. Hole, G. von Nessi, D. Pretty, J. Howard, B. Blackwell, J. Svensson, and L. C. Appel, Review of Scientific Instruments **81**, 10E127 (2010).

[15]C. M. Bishop, "Graphical models," in *Pattern Recognition and Machine Learning*, Information Science and Statistics (Springer-Verlag New York, 2006) Chap. 8, pp. 359–362, 1st ed.

[16]A. Pavone, J. Svensson, A. Langenberg, N. Pablant, U. Hoefel, S. Kwak, R. C. Wolf, and the W7-X team, Review of Scientific Instruments **89**, 10K102 (2018).

[17]R. Hooke and T. A. Jeeves, Journal of the ACM **8**, 212 (1961).

[18]H. Haario, E. Saksman, and J. Tamminen, Bernoulli **7**, 223 (2001).

[19]C. J. Geyer, in *Handbook of Markov Chain Monte Carlo*, Handbooks of Modern Statistical Methods, edited by S. Brooks, A. Gelman, G. Jones, and X.-L. Meng (Chapman and Hall/CRC, 2011) Chap. 1.

[20]S. A. Bozhenkov, M. Beurskens, A. Dal Molin, G. Fuchert, E. Pasch, M. R. Stoneking, M. Hirsch, U. Höfel, J. Knauer, J. Svensson, H. Trimino Mora, and R. C. Wolf, Journal of Instrumentation **12**, P10004 (2017).

[21]X. Han, X. Liu, Y. Liu, C. W. Domier, N. C. Luhmann, E. Z. Li, L. Q. Hu, and X. Gao, Review of Scientific Instruments **85**, 073506 (2014).

[22]J. van den Berg, I. Abramovic, N. J. Lopes Cardozo, and D. Moseev, Review of Scientific Instruments **89**, 083507 (2018), https://doi.org/10.1063/1.5035416.

[23]Stange, Torsten, Laqua, Heinrich Peter, Beurskens, Marc, Bosch, Hans-Stephan, Bozhenkov, Sergey, Brakel, Rudolf, Braune, Harald, Brunner, Kai Jakob, Cappa, Alvaro, Dinklage, Andreas, Erckmann, Volker, Fuchert, Golo, Gantenbein, Gerd, Gellert, Florian, Grulke, Olaf, Hartmann, Dirk, Hirsch, Matthias, Höfel, Udo, Kasparek, Walter, Knauer, Jens, Langenberg, Andreas, Marsen, Stefan, Marushchenko, Nikolai, Moseev, Dmitry, Pablant, Novomir, Pasch, Ekkehard, Rahbarnia, Kian, Mora, Humberto Trimino, Tsujimura, Toru, Turkin, Yuriy, Wauters, Tom, and Wolf, Robert, EPJ Web Conf. **157**, 02008 (2017).

[24]K. Brunner, T. Akiyama, M. Hirsch, J. Knauer, P. Kornejew, B. Kursinski, H. Laqua, J. Meineke, H. T. Mora, and R. C. Wolf, Journal of Instrumentation **13**, P09002 (2018).

[25]S. K. Rathgeber, L. Barrera, T. Eich, R. Fischer, B. Nold, W. Suttrop, M. Willensdorfer, E. Wolfrum, and the ASDEX Upgrade Team, Plasma Physics and Controlled Fusion **55**, 025004 (2013).