
Electronics/software functionality requirements

	Name
Owner	O. M. Røhne
Reviewer	Erik Adli and Håvard Gjersdal
Approver	

1 Scope

The overall functionality requirements of the image processing system are to provide a synchronous next-pulse veto signal and to supply operational- and diagnostic beam imaging (System specification, ESS-0066913). The required functionality and processing steps are summarized in Table 1, adapted from the working document at [Software and FPGA functionality summary \(ESS Confluence\)](#).

Table 1: Image processing steps, adapted from ESS Confluence.

Functionality	SW	FW
Display the image - 14 Hz - Also raw data	x	o
Processing/correction - Distortion - Background subtraction - Dead pixel mask - Vignetting correction - Median noise filter - Motion deconvolution	x x x x ? ?	? ? ?
Parameter extraction - Centroid position - Peak intensity - Percentage outside footprint - Normalization to beam current - Position of fiducials - Correction for temp./yield - Histogram and intensity contours - Beam diagnostic pulse	x x x ? ?	x x x o ? ?

2 Input assumptions

2.1 Photon source

From SNS experience (T. Shea, private communication) the photon yield per thickness of flame sprayed chromium doped alumina is similar to that of the sintered ceramic version of the material, e.g. standard AF995 "Chromox". The photon attenuation length has been seen to be well in excess of a typical coating thickness of 100 μm , so the full coating is assumed to contribute to the external luminous yield.

Studies done for the CERN LHC Dump Line quotes a specific yield of 10^4 ph/MeV for $\text{Al}_2\text{O}_3 : \text{Cr}^{3+}$ (Lefèvre et al 2007, cds.cern.ch/record/1045239).

The stopping power of alumina for 2 GeV protons is 1.66 MeV cm^2/g (NIST PSTAR, [PhysRefData/Star/Text/PSTAR.html](https://physrefdata.star.text.pstar.html)) and is fully dominated by electronic ionization. Assuming a standard alumina density of 3.95 g/cm^3 the specific photon yield per proton is found to be

$$y_{\text{ph}} = 6.56 \text{ ph/proton}/\mu\text{m}.$$

Inserting the number of protons per pulse as $1.12 \cdot 10^{15}$ (System overview and schedule, ESS-0066912) and assuming the active coating is 100 μm thick, the expected number of luminescent photons per pulse is around $7 \cdot 10^{17}$.

Before onset of radiation damage the spectral composition is dominated by the narrow ruby lines around $\lambda = 695$ nm. The comparatively long fluorescent lifetime (3 ms), combined with the rotational speed of the target wheel will cause a smearing of the image on the Beam Entrance Window (BEW), akin to photographic motion blur. The image on the Proton Beam Window (PBW) is not affected.

The luminescent yield is known to be affected by elevated temperatures, at 200 C the narrow spectral lines have almost disappeared but much of the light output is retained in a broadband spectrum. This effect is reversible, and normal fluorescence resumes when the coating is returned to standard temperature.

Due to radiation heating the coating temperature will have significant excursions during operation. A system of online temperature measurements and yield-versus-temperature calibration curves is required in order to provide charge-normalized beam measurements.

In a somewhat less benign scenario the luminescent coating develops a temperature gradient across its surface. If the effect of the gradient turns out to be significant, a 2D-map of the expected temperature response is required to predict the local temperature from the global temperature measurement. The predicted local temperature is then used with the yield calibration curve, for a pixel-by-pixel charge normalization.

The effect of radiation damage on the other hand is permanent. From SNS experience the eventual light yield is expected to level off at a about 5 % of the initial yield (Blokland 2013, [IBIC2014/papers/tupd13.pdf](#)). A charge-normalization procedure obviously requires a method for in-situ light yield re-calibration, as radiation damage gradually sets on.

2.2 Optical system

Reference: Target optical system, design document. Chess ESS-0149764

The over-all resolution of the imaging system is assumed to be 1 mm in the object plane, but with sizeable geometric distortions. The numeric aperture is $NA = 0.001$, which is to say only one in four million emitted photon will reach the imaging sensor.

2.3 Final focus and camera

The task of the final focus optics is to project the virtual image created by the mirror system onto the camera sensing chip. Given the over-all geometrical constraints and typical sensor sizes an objective of focal length $f \geq 600$ mm and an aperture of $f/6$ is assumed. The operational focal distance is fixed by the location of the virtual image at about $l \simeq 15$ m. In order to cross-check mirror alignment it might be requiquired to focus also at closer distances, e.g. to see fiducial marks at intermediate planes.

The fundamental requirements on the camera are that the pixel resolution does not compromise the 1 mm resolution of the imaging system, and that the triggering- and frame rate capabilities are compatible with two times beam synchronous operation at 14 Hz; one background pulse with fiducials are to be taken per beam pulse.

Facing a diminishing light yield due to radiation damage, a state-of-the-art high sensitivity camera is required. A candidate model, the [ORCA-Flash4.0 V3](#) from Hamamatsu, combines a 4 MPix resolution with a readout noise of $1.4 e^-$. Key specifications and performance parameters are summarized in Table 2. Considering a field of view of about 250×250 mm² a 4×4 pixel binning would be sufficient to match the 1 mm resolution requirement.

Table 2: ORCA-Flash4.0 V3 Digital CMOS camera specifications. Excerpt from Hamamatsu website

Sensor type	Hamamatsu sCMOS
Quantum efficiency	82 % @ 560 nm
Pixel size	$6.5 \times 6.5 \mu^2$
Pixel count	2048×2048
Full well capacity	30 ke^- (typ)
Readout noise	1.4 e^- (rms)
Interface	Camera Link, USB 3.0
Frame rate	30 – 100 fps

3 Algorithmic pipeline

The processing steps leading up to the synchronous next-pulse veto must be completed within a hard deadline of order 10 ms. The baseline strategy is to implement the time-critical steps in hardware (FPGA), and also to use the hardware output video stream as input to further software processing and preparation for presentation. Additionally, a full archive of raw video stream is required.

Since the key decision variables e.g. centroid and beam-outside-box are defined on the undistorted beam intensity distribution, all the processing steps leading up to the distortion correction and the subsequent parameter extraction are considered time-critical and will be implemented in hardware. These steps include noise filtering and intensity normalization, the geometric remapping itself and the integration of the beam-outside-box estimates.

3.1 Preprocessing

The purpose of the preprocessing step is to suppress noise, normalize the input data, and provide rapid feedback for e.g. exposure adjustments or fiducial illumination.

Sources of image artifacts

The contributions to *undesirable* pixel-level image artefacts are broadly described as either *noise* or *defects*. Pixel defects are typically *dead*, *hot*, or *noisy* pixels, and their number is expected to increase during the camera lifetime. Image noise is characterised as either stochastic or patterned, which will have different effects in downstream image processing.

Uniform or *uncorrelated* noise originate from pixel leakage currents, from electronics thermal noise sources, and from the fundamental Poissonian distribution of photons per pixel. While uncorrelated noise will degrade image quality and parameter measurements, the degradation is mostly gradual and predictable.

Patterned or *correlated* noise, e.g. from digital cross-talk, typically occurs in blocks or with distinct spatial frequencies. Unchecked correlated noise above a certain level can have a sharply increasing unpredictable adverse effect on measurements, in particular in conjunction with advanced/non-linear downstream image processing.

Background subtraction and pixel mapping and

Background subtraction employs regular non-beam frames, subtracting the pixel values frame-by-frame. A simple ADC pedestal subtraction maps both dead- and hot pixels defects to black, and it might be beneficial to treat these uniformly. Excessive noise defects will remain noisy, while regular non-defect pixels will have their noise increased by a factor $\sqrt{2}$.

An alternative approach to managing pixel defects is by means of maintaining an explicit list of bad pixels. For the corresponding position the pixel values are interpolated from neighboring pixels. Such a *pixel map* can be derived in-situ using scheduled captures of white frames (dead pixels) dark frames (hot pixels) or multiple grey frames (noise pixels).

Median filter

The median filter employs a sliding window, replacing each pixel value with the median of the surrounding window values. The obvious benefit of the median filter is to suppress outlier values, which could come from isolated pixel defects as well as uncorrelated noise. The downside is a degradation of the ultimate spatial resolution, some reduction in contrast, and possibly even an amplification of multi-pixel defects and patterned noise.

Vignette correction and overall scaling

The purpose of vignette correction is to compensate for the variation of the luminous yield and the effective numerical aperture across the field of view. Key questions are how strong a correction is needed and if the correction factor can be accounted for by a simple parameterization or if a full field map is necessary.

A possible concern with the non-uniform scaling is that it will give rise to a non-uniform noise, which might be a problem for subsequent/downstream image processing.

3.2 Distortion correction

The combination of the doubly-curved target wheel surface and an uncorrected imaging system means that the raw image will be severely distorted with respect to the beam-transverse reference system. Since the compound distortion is the result of two geometrically unrelated off-axis contributions the net correction cannot be expected to be expressed in terms of the familiar barrel/pincushion parameterizations. Instead, parameterizations of each contribution are constructed, and the combination of the two parameterizations is used to extract an explicit per-pixel distortion map:

$$(x_{\text{beam}}, y_{\text{beam}}) \mapsto (x_{\text{proj}}, y_{\text{proj}}) \mapsto (x_{\text{imag}}, y_{\text{imag}})$$

The different contributions to the beam image distortion affect the intensity normalization in subtly different ways. First, the angle between the incident beam and the doubly-curved surface modulates the effective thickness of the luminous coating. Second, the intermediate camera projection preserves the *apparent* surface luminance. Finally, the uncorrected distortion of the curved mirror preserves the *integral* intensity. A naive correction remap, using e.g. a bilinear interpolation in the input pixels, will by nature preserve the *apparent* luminance and implicitly rescale the *integral* intensity. A more careful approach might require a companion intensity rescaling map to be determined by in-situ calibration.

Target wheel surface and pinhole camera projection

For the purpose of modeling the geometric distortion an isometric beam-transverse $200 \times 100 \text{ mm}^2$ grid is first projected onto the rim of the target wheel. This doubly curved surface is approximated using a torus of minor radius $r_1 = 58 \text{ mm}$ and major radius $r_0 = 1250 \text{ mm}$,

$$\left(r_1 - \sqrt{z^2 + x^2}\right)^2 + y^2 - r_0^2 = 0.$$

Solving the torus-equation yields the intersection z along the beam axis in terms of the beam-transverse (x, y) coordinates,

$$z = \sqrt{2r_1 \sqrt{r_0^2 - y^2 - y^2 - x^2 + r_1^2 + r_0^2}}.$$

The resulting curved grid is projected by means of an ideal pinhole camera placed at the nominal position of the first mirror. The resulting input image distortion is shown in Figure 1.

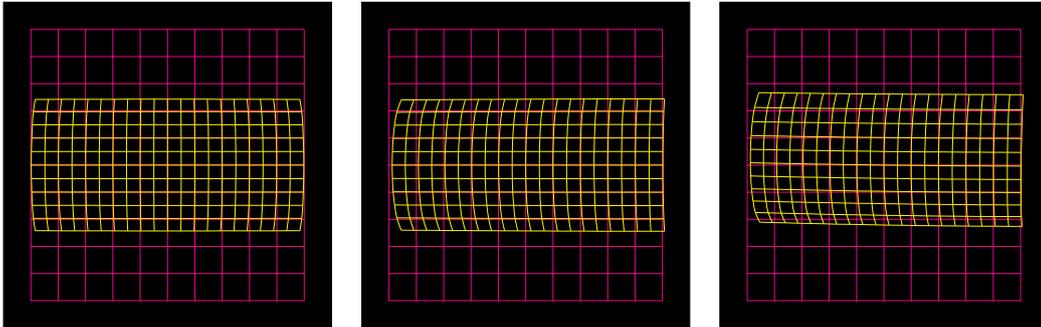


Figure 1: Beam-transverse system (yellow) projected onto the target wheel and seen from various vantage points: on the beam axis (left), nominal first mirror (center), and 150 mm above the nominal (right).

Imaging system intrinsic distortion

The distortion intrinsic to mirror system itself is readily evaluated by passing a rectilinear test image through Zemax simulations or by directly using the functional prototype. The required parameterization shown in Figure 2 is found by fitting the pixel position of the indicated with a second degree multinomial in the reference coordinates.

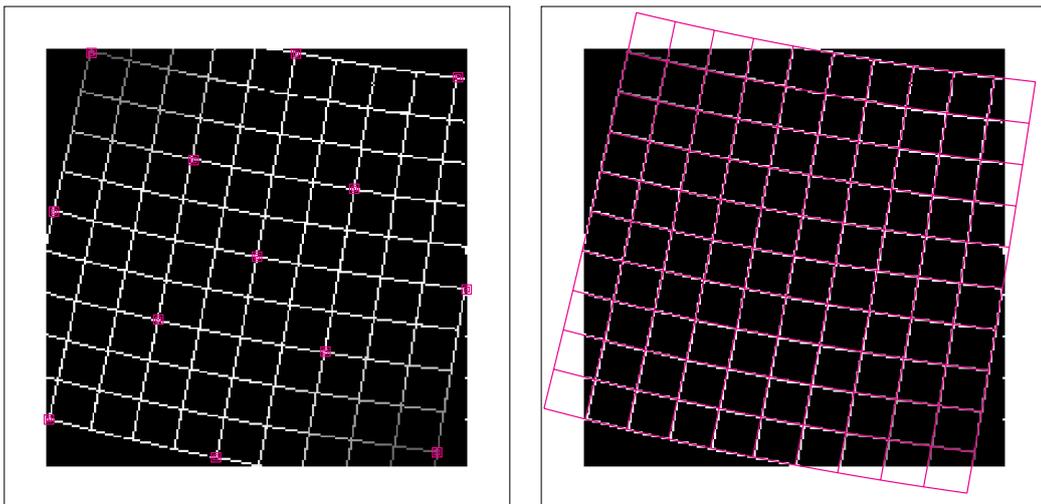


Figure 2: Image system distortion parameterization (purple) superimposed on a simulated image. Reference points (left) are fitted with a second-degree multinomial (right).

Extracted pixel map correction

The procedure to construct the compound correction map proceeds in two stages. First, the sequential composition of the two parameterized maps are fitted with fourth degree multinomials, which if needed can be used directly by a parameterized correction scheme. Second, an explicit pixel map is generated from the fitted multinomial. The sequentially composed map is shown in Figure 3.

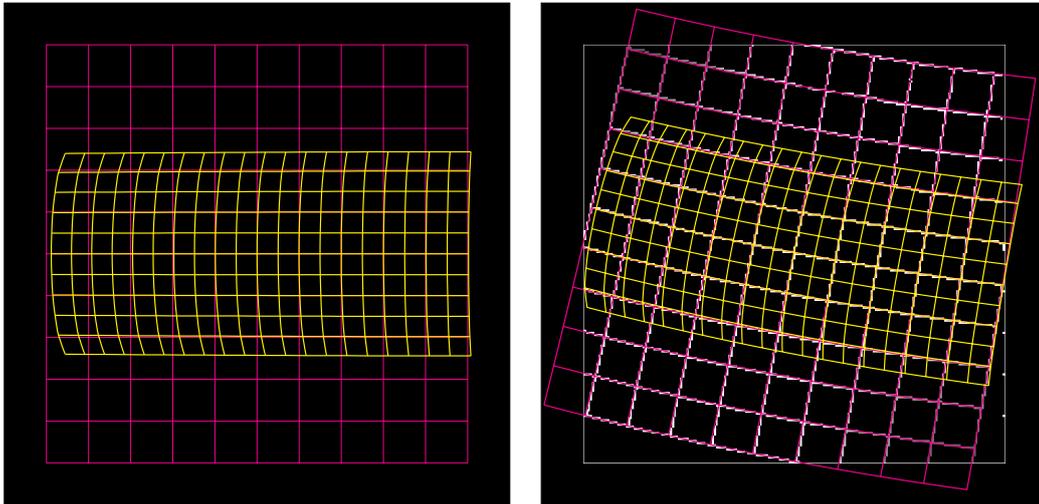


Figure 3: Image distortion contributions from curved-surface projection (left), intrinsic to the imaging system, and the resulting distortion (right). The yellow grid represents the nominal beam-transverse system, the purple grid is the pinhole-camera reference.

An example application of the partial- and full distortion correction is shown in Figure 4. Using as starting point a synthetic image exposing each contribution to the distortion the partial (imaging-only) correction as well as the full correction is demonstrated by means of the remap-function from the [OpenCV](#) library.

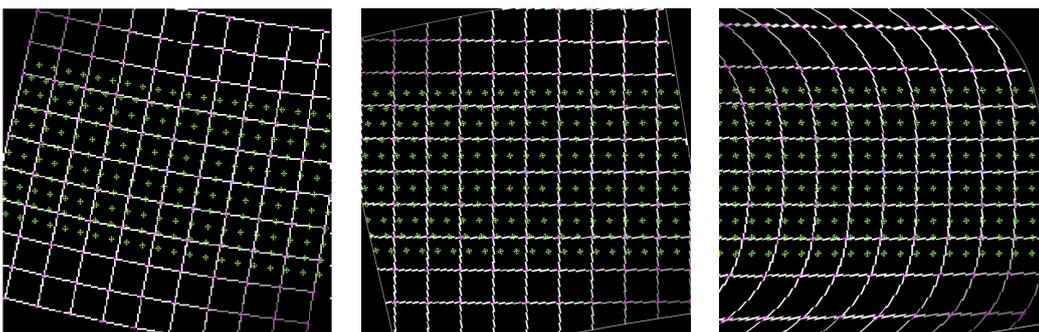


Figure 4: Synthetic image (left), partially corrected (center) and fully corrected (right) using an explicit pixel map. The green/yellow markers represent the nominal beam-transverse system, the purple/magenta markers are the pinhole-camera reference.

3.3 Motion blur deconvolution

A possible need to unfold motion blur arises because of the finite lifetime of the luminous source, coupled with the constant angular velocity of the target wheel.

Phosphor lifetime, rotational speed

For sake of a proof-of-principle implementation a luminous lifetime of 1 ms is assumed. The 2.5 m diameter wheel rotates at a speed of 14/36 revolutions per second. The resulting point-spread-function (PSF) will be an exponential with decay length of about 3 mm.

Richardson-Lucy software implementation

The Richardson-Lucy deconvolution procedure (Richardson 1972, [doi:10.1364/JOSA.62.000055](https://doi.org/10.1364/JOSA.62.000055) and Lucy 1974, [doi:10.1086/111605](https://doi.org/10.1086/111605)) is a method for reconstructing the latent image u from the observation d using a known PSF p . The method is iterative, with each step yielding an improved estimate $u^{(t+1)}$ from the previous estimate,

$$u^{(t+1)} = u^{(t)} \cdot \left(\frac{d}{u^{(t)} \otimes p} \otimes \hat{p} \right).$$

The iterative procedure is seeded using the raw observation as the first estimate, $u^{(0)} = d$.

In order to ensure a reasonably uniform PSF it is assumed that the deconvolution is applied after any imaging distortion has been corrected, and in order to minimize the computational bandwidth it is assumed that the image scan lines are aligned with apparent direction of motion.

An example reconstruction is shown in Figure 5. The starting point is a 6 mm diameter luminous disk, smeared with a Gaussian blur ($\sigma = 0.5$ mm), and rendered with a resolution of 2 pix/mm. A horizontal motion blur with characteristic length of 3 mm is then applied. Finally, a Poisson noise contribution corresponding to 100 photons/pixel is added. The reconstructed image up to six iterations is shown.

4 Parameter extraction and image presentation

Careful consideration should be given to whether key parameters should be extracted from raw or processed images. Spatial parameters extracted from raw images would themselves need correction. However, given the inherently non-linear nature of the distortions, such post-extraction corrections might not be generally possible. Either way, it is assumed that noise removal and response equalization is always applied first.

4.1 Beam centroid and RMS

A general image map does not commute with the n th-order moment operation. In other words, there is no general method to obtain the centroid of a *corrected* image by post-correcting the centroid of the *uncorrected* image. Consequently it is assumed that correcting the images prior to calculating the centroid and RMS will yield more reliable results.

4.2 Beam peak intensity

The beam peak intensity is calculated as the maximum of a sliding-window pixel intensity average. Unlike the centroid and RMS, the spatial position of the peak is locally defined, and can be transformed to the beam-transverse system by means of the inverse distortion map.

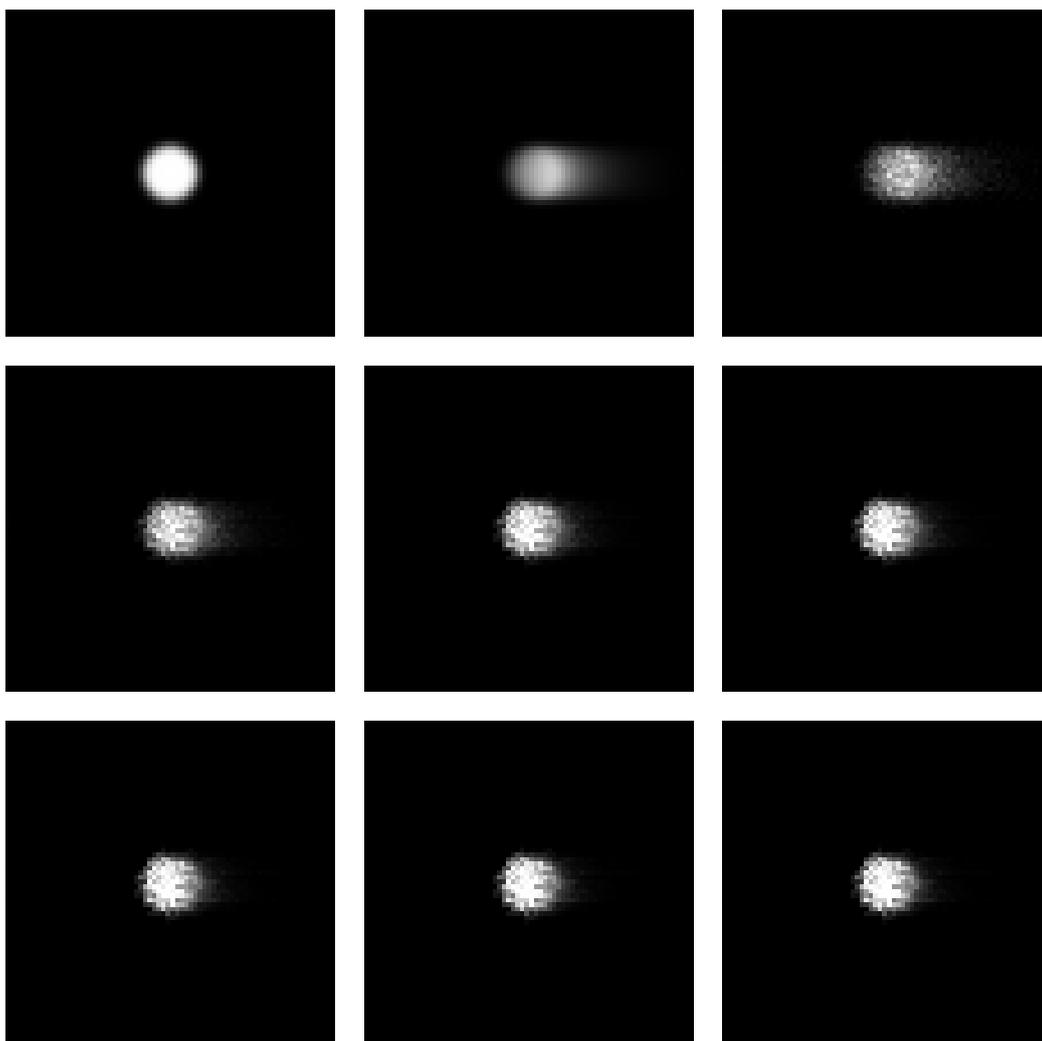


Figure 5: Synthetic test image and Richardson-Lucy deconvolution reconstruction. First row: smeared disk (left) added motion blur (center) added poisson noise (right). Subsequent rows left to right: successively improved reconstructed images.

4.3 Beam outside box

Based on simulated beam profiles similar to the one shown in Figure 6 an envelope corresponding to e.g. 99.9% of the nominal beam intensity can be defined as a rectangular box.

In practical terms, setting the threshold on the ratio of the inside- and outside ratios will cancel any global intensity normalization factor.

The statistical error on the measurement of the integrated beam intensity *outside* the envelope will have two contributions; the summed uncorrelated read-out noise of the relevant pixels, and the intrinsic fluctuations of the photon yield and the sensor quantum efficiency.

For a practical example consider a state-of-the-art low-noise camera like [Hamamatsu ORCA-Flash4.0 V3](#) as referenced above.

Combining the previous numbers for the expected photon yield with the numerical aperture of the imaging system, and assuming a conservative sensor quantum efficiency of 20 %, the image of the full pulse is expected to contain about $3.7 \cdot 10^{10}$ photo-electrons total, or about 18 ke^- per pixel if the image covers 50 % of the 4 MPix sensor.

Assuming somewhat arbitrarily that about 10 % of the pixels contribute to the beam-outside-box calculation, the uncorrelated read-out noise of those pixels will be $\sqrt{0.4 \text{ M}} \cdot 1.4 \text{ e}^-$ or about 900 e^- . Considering 5σ an acceptable confidence level, the lowest achievable threshold would be about 4.5 ke^- , or less than one part per million of the full pulse.

4.4 Image presentation

Following hardware- and software processing the resulting image is provided as an EPICS Area Detector stream. Standard formatting and presentation post-processing are applied downstream, e.g. logarithmic pseudo-color maps and overlay printing of extracted features and parameters.

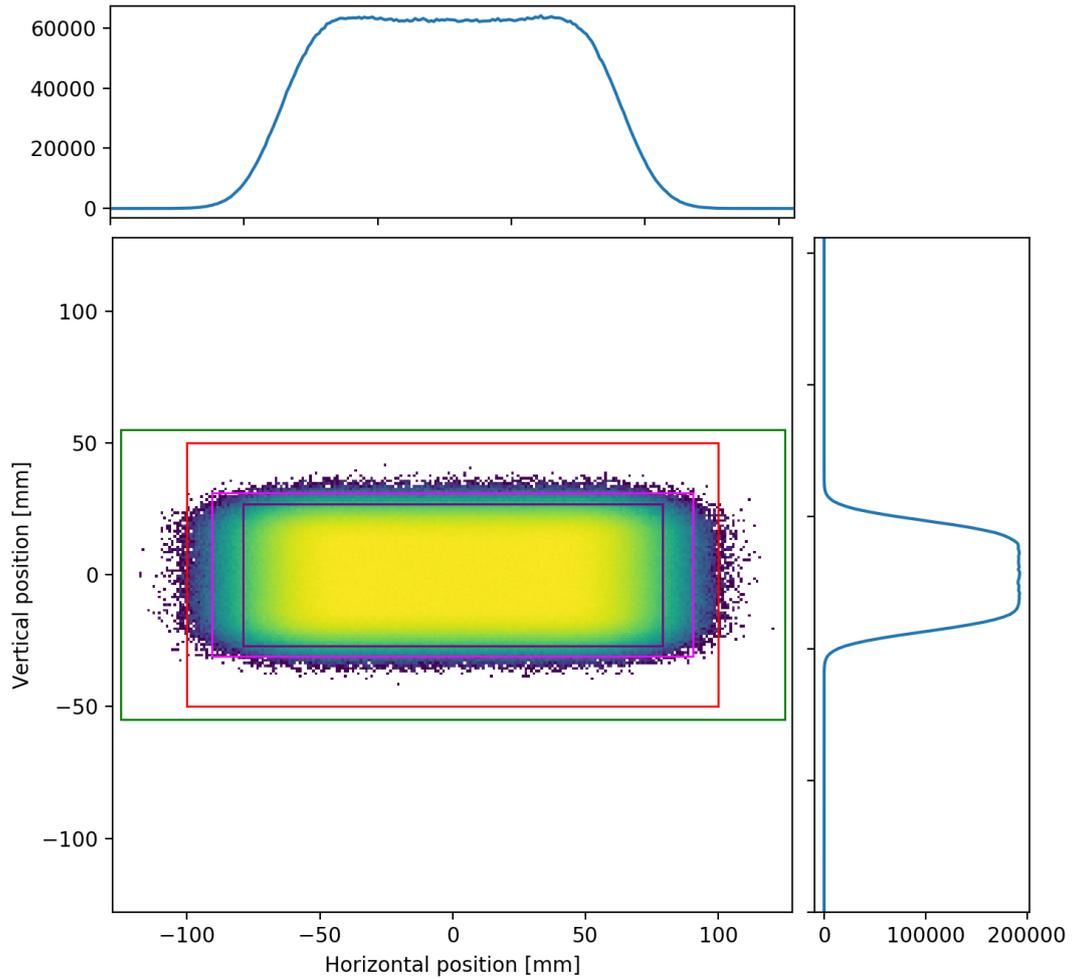


Figure 6: Rastered beam profile, 8 million simulated protons. The 2-dimensional profile uses a logarithmic intensity map. Reference rectangles shown, green: nominal field-of-view, red: beam entrance window, pink: 99.9% nominal beam, purple 99.0% nominal beam. Proton simulation courtesy of Heine Dølrath Thomsen, Aarhus University.