Real-Time Detection of Overloads on the Plasma-Facing Components of Wendelstein 7-X

Aleix Puig Sitjes [1,*], Marcin Jakubowski [1], Dirk Naujoks [1], Yu Gao [1], Peter Drewelow [1], Holger Niemann [1], Joris Fellinger [1], Victor Moncada [2], Fabio Pisano [3], Chakib Belafdil [2], Raphael Mitteau [2], Marie-Hélène Aumeunier [2], Barbara Cannas [3], Josep Ramon Casas [4], Philippe Salembier [4], Rocco Clemente [4], Simon Fischer [1], Axel Winter [1], Heike Laqua [1], Torsten Bluhm [1], Karsten Brandt [1] and The W7-X Team [†]

1   Max-Planck-Institut für Plasmaphysik, Wendelsteinstr. 1, 17491 Greifswald, Germany;
     marcin.jakubowski@ipp.mpg.de (M.J.); dirk.naujoks@ipp.mpg.de (D.N.); yu.gao@ipp.mpg.de (Y.G.);
     peter.drewelow@ipp.mpg.de (P.D.); holger.niemann@ipp.mpg.de (H.N.); joris.felinger@ipp.mpg.de (J.F.);
     simon.fischer@ipp.mpg.de (S.F.); axel.winter@ipp.mpg.de (A.W.); heike.laqua@ipp.mpg.de (H.L.);
     torsten.bluhm@ipp.mpg.de (T.B.); karsten.brandt@ipp.mpg.de (K.B.)
2   Commissariat à l'Énergie Atomique et aux Énergies Alternatives (CEA), Institut de Recherche sur la Fusion
     par Confinement Magnétique (IRFM), F-13108 Saint Paul-lez-Durance, France; victor.moncada@cea.fr (V.M.);
     chakib.belafdil@cea.fr (C.B.); raphael.mitteau@cea.fr (R.M.); marie-helene.aumeunier@cea.fr (M.-H.A.)
3   Department of Electrical and Electronic Engineering, University of Cagliari (UniCa), Piazza d'Armi,
     09126 Cagliari, Italy; fabio.pisano@unica.it (F.B.); barbara.cannas@unica.it (B.C.)
4   Department of Signal Theory and Communications, Universitat Politècnica de Catalunya (UPC), Jordi Girona
     1-3, 08034 Barcelona, Spain; josep.ramon.casas@upc.edu (J.R.C.); philippe.salembier@upc.edu (P.S.);
     rocco.clemente@upc.edu (R.C.)
*   Correspondence: aleix.puig.sitjes@ipp.mpg.de
†   Membership of the Team Name is provided in Acknowledgments.

Abstract: Wendelstein 7-X (W7-X) is the leading experiment on the path of demonstrating that stellarators are a feasible concept for a future power plant. One of its major goals is to prove quasi-steady-state operation in a reactor-relevant parameter regime. The surveillance and protection of the water-cooled plasma-facing components (PFCs) against overheating is fundamental to guarantee a safe steady-state high-heat-flux operation. The system has to detect thermal events in real-time and timely interrupt operation if it detects a critical event. The fast reaction times required to prevent damage to the device make it imperative to automate fully the image analysis algorithms. During the past operational phases, W7-X was equipped with inertially cooled test divertor units and the system still required manual supervision. With the experience gained, we have designed a new real-time PFC protection system based on image processing techniques. It uses a precise registration of the entire field of view against the CAD model to determine the temperature limits and thermal properties of the different PFCs. Instead of reacting when the temperature limits are breached in certain regions of interest, the system predicts when an overload will occur based on a heat flux estimation, triggering the interlock system in advance to compensate for the system delay. To conclude, we present our research roadmap towards a feedback control system of thermal loads to prevent unnecessary plasma interruptions in long high-performance plasmas.

Keywords: Wendelstein 7-X; plasma facing components protection; imaging diagnostics; thermography; hot spot detection; thermal events detection

1. Introduction

Wendelstein 7-X (W7-X) is the largest drift optimized stellarator. It aims at demonstrating reactor-relevant plasma performance of the stellarator line with nearly steady-state plasmas of up to 30 min [1]. With a heating power of 10 MW, protecting the plasma-facing components (PFCs) is of vital importance in preventing damage to the device in steady-state.

W7-X experimental campaigns define a step-by-step road map with increasing performance to reach steady-state operation. The operation phase named OP1.2 [2] was designed to test the island divertor concept in which the target plates intersect the helical field lines inside the magnetic islands at the edge for power and particle exhaust [3]. W7-X has a three-dimensional helical shape with a five-fold modular symmetry (see Figure 1). The divertor is composed of 10 units with several target modules (see Figure 2). During this experimental phase, we used an inertially cooled test divertor made of fine-grain graphite tiles. This allowed us to test the image analysis system designed to protect the PFCs in steady-state without danger of damaging the high-heat flux divertor.
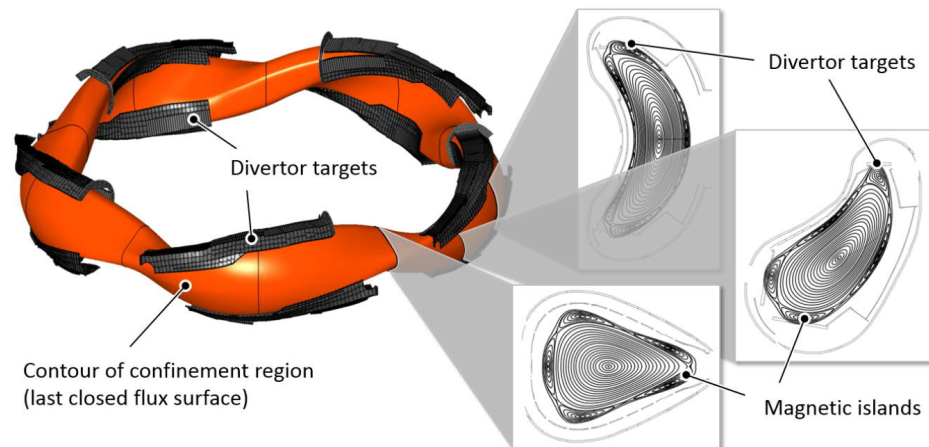
Figure 1. W7-X five-fold modules with its 10 divertor units, 5 upper and 5 lower divertor units. The divertor targets intersect the magnetic islands at the edge for power and particle exhaust.
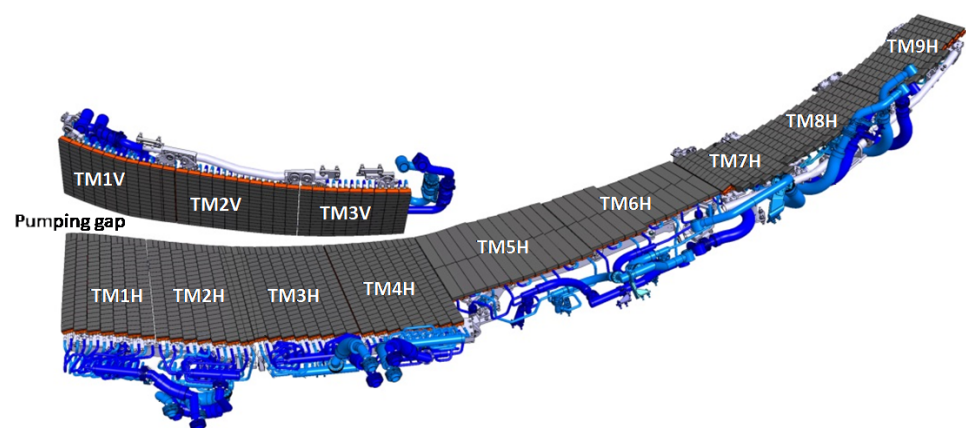
Figure 2. The divertor units are 5 m long and 1 m wide with several target modules: vertical targets (TM1V-TM3V), low-iota targets (TM1H-TM4H), low-load target modules (TM5H-TM6H) and high-iota targets (TM7H-TM9H). The divertor water-cooled tiles are made of CFC (Carbon Fibre Composite) except the two low-load central target modules, which are made of fine-grain graphite.

The steady-state operation phase (OP2) is planned to start in 2022. By then, all the PFCs will be water-cooled and the new high-heat flux water-cooled divertor will sustain heat fluxes of up to 10 MW/m^2. The divertor tiles are made of a Carbon Fibre Composite (CFC) layer joined to a CuCrZr cooling structure. The divertor's maximum operational temperature is limited by a Cu interlayer, which should not exceed a sustained temperature of 475 °C. This temperature is reached at 10 MW/m^2 when the surface temperature is 1200 °C [4,5]. The other water-cooled PFCs are the baffles, the heat shields, the wall and pumping gap panels, and the poloidal closures. The baffles and the heat shields are made

of fine-grain graphite tiles, which are limited to 400 °C because of the brazed joint between the cooling pipe and the heat sink. The wall and pumping gap panels and the poloidal closure are made of water-cooled stainless steel, which cannot exceed 200 °C. See Figure 3 for an overview of the PFCs and their maximum operational temperatures.
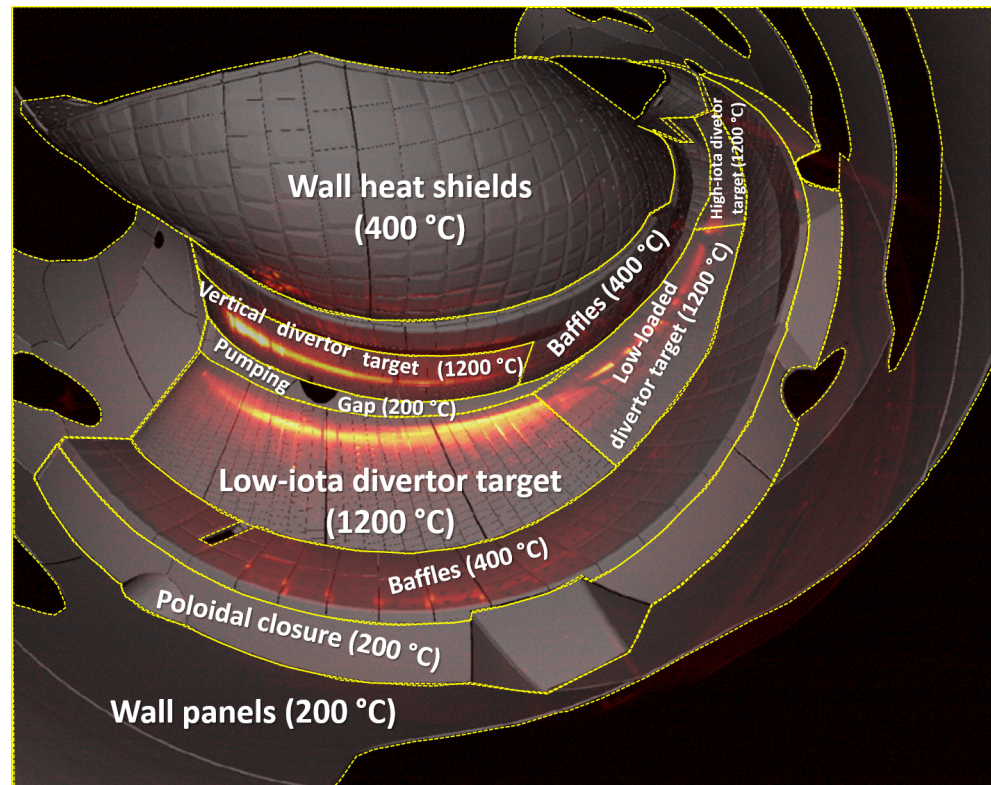
Figure 3. Infrared image overlaid on the CAD of the different plasma-facing components with their maximum operational temperatures.

The imaging system will have to guarantee a safe steady-state operation of the device by monitoring the PFCs in real-time and prevent overheating by triggering the interlock system if their temperature limits are compromised. The interlock is part of the safety system of the machine and stops the plasma heating systems when it is triggered. Several similar safety systems have been developed for other fusion devices with the same goal [6–9]. These systems typically rely on pre-defined regions of interest and it triggers the interlock when the temperature limits of these regions are breached.

In Section 2 the system to protect the PFCs of W7-X in OP2 is described. The thermography diagnostic, used for the detection of thermal events, is introduced together with the architecture of the real-time image analysis system. Section 3 goes into detail on the image analysis algorithms, the scene models concept (Section 3.1), the processing pipeline (Section 3.2), and the overload detection algorithm (Section 3.3). Section 4 illustrates the results obtained with the analysis of the data from the OP1.2 campaign. Section 5 explains the ongoing research towards a full feedback control system of thermal loads and Section 6 summarizes the major contributions and results.

2. Real-Time Imaging System

The thermal loads on the PFCs of W7-X are detected with the thermography diagnostic that monitors the 10 divertors and the surrounding baffles. W7-X is a stellarator with five-fold symmetry, but as far as safety is concerned, toroidal symmetry cannot be assumed. Heat loads in one module cannot be used to infer heat loads in other modules. We must provide 100% observation of divertors.

In OP1.2, nine immersion tubes (Figure 4) and 1 endoscope (Figure 5) were used [10]. At the beginning of OP2, with plasma energy still limited to 1 GJ, the diagnostic will comprise 8 water-cooled immersion tubes and 2 steady-state endoscopes. The immersion tubes, however, are not designed for steady-state operation and steady-state endoscopes will replace them when the plasma energy will progressively increase up to 18 GJ.
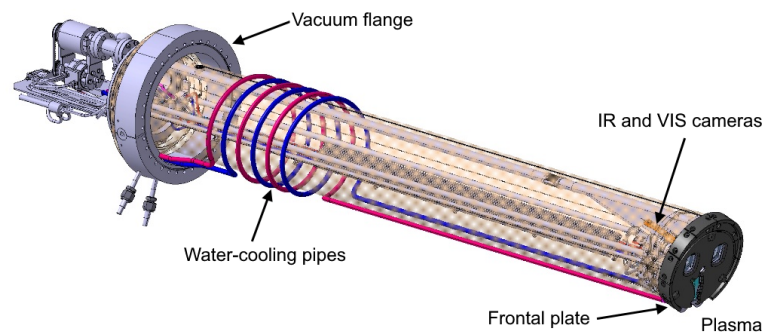
Figure 4. The immersion tubes provide an optical resolution of the order of 5 to 20 mm/pixel on the divertor, depending on the distance and the viewing angle of the divertor. In the high-iota region, the lower resolution may lead to measuring lower temperatures from hot spots smaller than the resolution limit. The immersion tubes are equipped with microbolometer cameras IRCam Caleo 768k L covering the spectral range from 8 to 10 μm at a frame rate of 100 Hz.
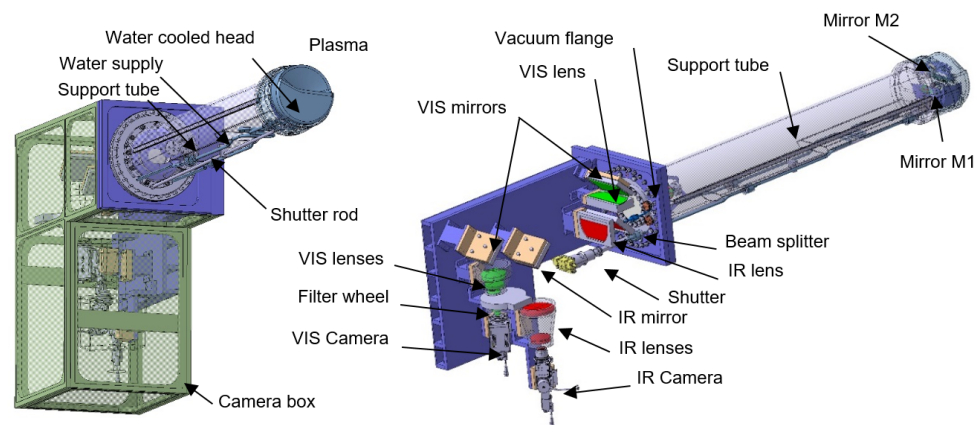
Figure 5. The steady-state endoscopes consists of an off-axis Cassegrain optical system with a pinhole aperture of 6 mm. A set of mirrors transmits the light and a dichroic splitter divides it into visible and infrared beams, which are detected by the cameras after being corrected by a group of lenses. They have an optical resolution of 8 mm/pixel and they are equipped with an Indium Antimonide (InSb) infrared camera covering the spectral range from 2 to 5.7 μm (filtered at 4.1 μm) at a frame rate of 100 Hz.

The image analysis of the thermographic images is a fundamental part of the machine safety system of W7-X (see Figure 6). The algorithm is based on real-time image processing techniques. The overload detection algorithm detects when a PFC is overheated and interrupts operation immediately through the interlock system to guarantee that the operational limits of the PFCs are not exceeded.
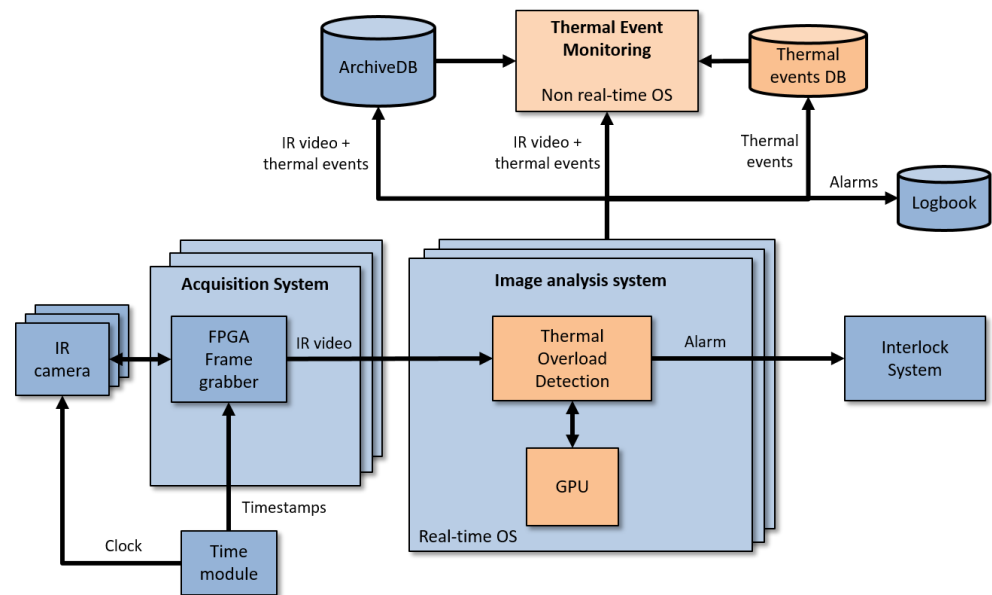
Figure 6. An MTCA-based (Micro Telecommunication Computing Architecture) frame grabber acquire the infrared images and synchronizes them with the W7-X universal time. The raw images are transferred to the Image Analysis System, a set of real-time computers, one for each camera, that calibrates and analyses the infrared data. The video streams are analyzed through computer vision techniques to detect overloads. When the integrity of the PFCs is compromised, the system triggers the Interlock. We will implement some of these algorithms on a GPU to speed up computation. The detected overloads and alarms are sent to the Thermal Event Monitoring system in the control room and registered into a database, the Archive, and the Logbook for future reference.

3. Image Analysis

3.1. The Scene Model

During the image analysis, the system checks the temperatures against the operational limits of the different PFCs, considering the thermal properties of each component. In many fusion devices, this is performed using regions of interest (ROIs). In the proposed system, we use the concept of scene model instead (see Figure 7). The scene model provides a pixel-wise mapping of each camera field of view with the CAD model of the in-vessel components. The CAD model is distorted to compensate for the strong lens distortion of the optical system and to match the field of view [11]. A code script builds the scene models using ray-tracing techniques, one for each camera. When in-vessel components are changed in between campaigns, the scene models are updated automatically. In this way, we improve maintainability and we guarantee the models match the installed components.

3.2. Image Processing Pipeline

Prior to the analysis of the thermographic images, the acquired images have to be calibrated. This includes applying a non-uniformity correction (NUC), bad pixel interpolation, and temperature conversion, using Planck's Law to convert from photon flux to temperature. This conversion requires the true emissivity value of the imaged surface. The scene model provides for each pixel the nominal emissivity value of the target PFC. The nominal emissivity, however, has to be corrected during operation because of the erosion and re-deposition of the target material or deposition of other impurities. An emissivity correction map provides this correction factor. Initially, it will be updated every week, but we may adjust this periodicity depending on how fast the emissivity changes significantly. We can perform the measurement of the new emissivities when the machine is in thermal equilibrium by observing the differences in the apparent temperature of the surfaces. This correction procedure can also be triggered automatically if these differences in the apparent temperature are above a certain threshold.
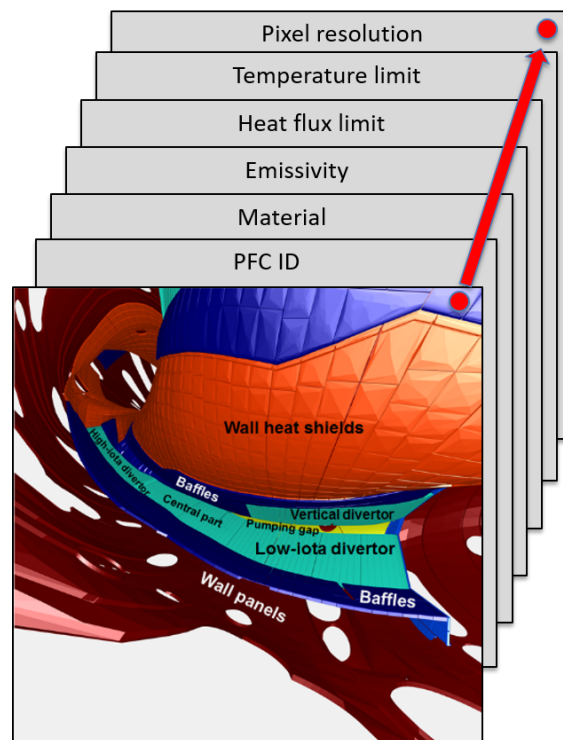
Figure 7. The scene model is a multi-channel image matching the camera field of view, each channel providing different PFCs properties for each pixel: identifier, material, emissivity, heat flux limit, temperature limit and pixel resolution.

The overload detection algorithm then analyzes the calibrated images (see Section 3.3 for in-depth details). This algorithm evaluates the risk of overload for each PFC considering the current heat flux, the maximum operational temperature of each component (also provided by the scene model), and the presence of surface layers. If there is a risk of overload, the interlock is triggered.

Because of carbon erosion, thin surface layers can develop on the PFCs because of re-deposition. They show as high-temperature hot spots because of their low thermal capacity and their poor thermal connectivity to the bulk of the tile and the cooling system. They do not pose a risk to the PFCs integrity but they may be a source of false alarms [12]. We update periodically a surface layer map to aid in the thermal event risk evaluation. We construct the surface layer map from dedicated discharges with special modulated heating. A train of pulses can be used to measure the time evolution of the surface temperature. The different hot spots respond differently because of their different thermal capacities and conductivities. The normalized temperature decay time is a discriminative feature that it is used to detect surface layers and delaminations [13].

We summarize the image analysis pipeline from the image acquisition to the interlock activation in Figure 8.
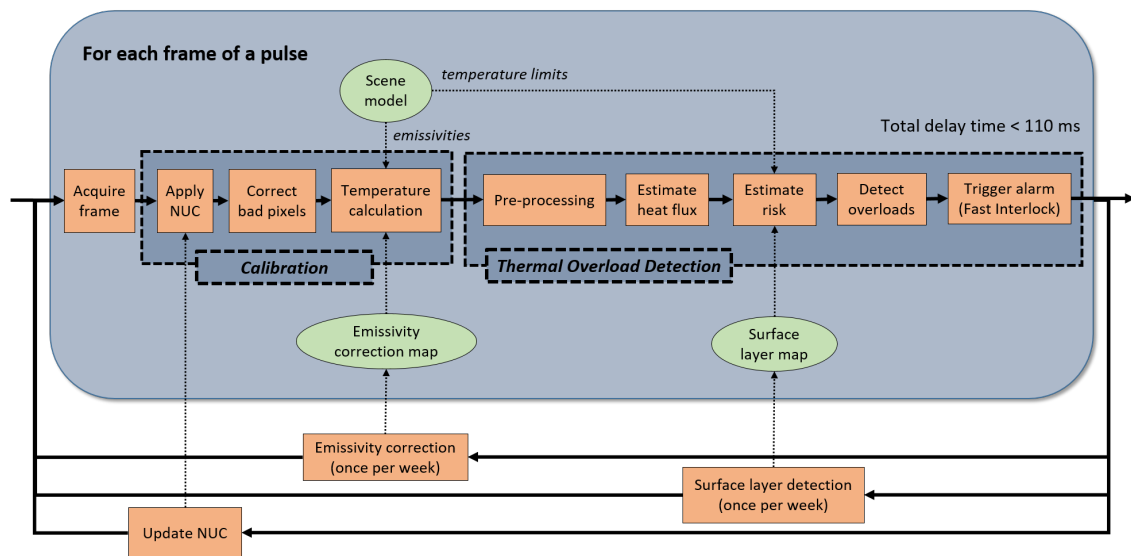
Figure 8. The system calibrates and analyses each frame of a pulse. The calibrated images are analysed by the Thermal Overload Detection algorithm. If there is risk of overload of the PFCs, the system sends an alarm to the interlock to stop operation. The total delay of the pipeline must be less than 110 ms, including the stopping of the heating systems. The scene model and the emissivity correction map provide the thermal properties of the imaged components and their temperature limits. A surface layer map is used to adjust the risk in presence of surface layers. The non-uniformity correction (NUC), emissivity correction map, and surface layer map are updated periodically with dedicated procedures.

3.3. Thermal Overload Detection Algorithm

The Thermal Overload Detection algorithm must trigger a stop operation (i.e., stop of heating power) before the maximum operational surface temperature of the PFCs is exceeded. The system has a delay time because of the processing and the time needed to stop the heating systems. To compensate for this delay and to anticipate when an overload would occur, it predicts the surface temperature using a transient model (Equation (1)) [12]. Equation (1) approximates the temperature increase ΔT given the heat flux q [kW/m^2] and a time interval Δt [s]. C_m is a material-dependent constant ($C_m = \pi \cdot \lambda \cdot \rho \cdot C_p/4 \times 10^{-6}$ [s(kW)^2/(m^4 K^2)], where λ is the diffusion coefficient, ρ the material density and C_p the specific heat capacity).

$$\Delta T = q\sqrt{\frac{\Delta t}{C_m}} \tag{1}$$

This model is valid for uncooled as well as cooled components before reaching the steady-state temperature or during fast transient heat loads when the heat propagates into the material down to the actively cooled heat sink. Close to the steady-state, however, this model overestimates the increase in the surface temperature, resulting in a conservative assessment. In any case, no overload can occur in steady-state conditions.

The system triggers an alarm when the temperature reaches the threshold T_{th} defined by Equation (2). This threshold depends on the temperature limit T_{limit}, the current heat flux q and a given reaction time ($t_r = 0.16$ s; see Figure 9 for a justification of this value).

$$T_{th} = T_{limit} - q\sqrt{\frac{t_r}{C_m}} \tag{2}$$

The system estimates the current heat flux and predicts the temperature evolution in the next 0.16 s according to the transient model and sets the temperature threshold for each frame. It is required, however, to process one frame and stop the heating systems in less time ($t_{max\_delay} = 0.11$ s) allowing a 50 ms of safety margin. The safety margin is subtracted to allow for the uncertainty of the prediction. The temperature evolution is predicted using a model which is valid for 1-dimensional semi-infinite solids. This is, of

course, an approximated model. Furthermore, the measurement of the heat flux is also very noisy, resulting in a prediction with high uncertainty. The safety margin value of 50 ms is a parameter of the system and it has been derived experimentally from the current data and it may need further adjustment in future campaigns.
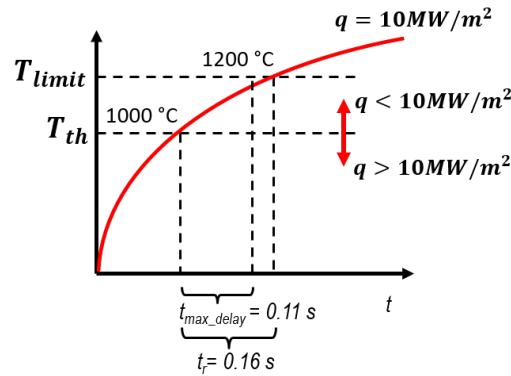
Figure 9. On the divertor targets, if the algorithm triggers the interlock when it detects that the temperature breaches the limit ($T_{limit} = 1200\ °C$), the temperature will exceeded this limit before the heating is interrupted due to delay of the system. To prevent this, it is required to stop at a lower threshold (T_{th}). At 1000 °C it takes 0.16 s to reach the temperature limit when the heat flux is 10 MW/m² (the maximum heat flux allowed). This provides us with a reference reaction time $t_r = 0.16$ s. The system, however, must stop at a lower threshold if the detected heat flux is higher than 10 MW/m² or it can stop at a higher threshold if the detected heat flux is lower.

According to this, the Thermal Overload Detection algorithm is designed as follows:

(1) After the acquisition and calibration of the thermal image T, the image is filtered to avoid false alarms due to remaining uncorrected bad pixels and noise coming from neutrons. It is clear that hot pixels can trigger false alarms, but also cold pixels must be removed as they can suddenly become hot, creating an apparent (false) peak of heat flux. This may lower the threshold in excess, triggering a false alarm. We remove hot pixels with a spatial filter, a morphological opening with a structuring element of 3×3 pixels and 8-connectivity. Cold pixels are removed by a complementary morphological closing followed by a reconstruction [14]. The uncertainty of the temperature calibration is then added to the thermal image resulting in the corrected image T'.

(2) Following, the corrected thermal image is averaged over time $\overline{T'}$ with a moving average to reduce its noise. The average window size n has to be kept small (2 to 5 frames) to avoid adding too much delay to the system.

(3) The temperature threshold T_{xy}^{th} [K] is computed with Equation (3) for each pixel (x, y) in real-time using the model in Equation (2), the reaction time $t_r = 0.16$ s, the temperature limit T_{xy}^{limit} [K] of the PFC, the current heat flux q_{xy} [kW/m²], and the material constant C_{xy}^m [s(kW)²/(m⁴ K²)]. The scene model provides the temperature limit and material constant for each pixel.

$$T_{xy}^{th}(t) = T_{xy}^{limit} - q_{xy}(t)\left(\frac{t_r}{C_{xy}^m}\right)^l \tag{3}$$

where l is a parameter that controls the linearity of the model. In theory, l should be 0.5 (the temperature increases as a square root), but this only holds if the heat-flux is constant. In some situations, when the heat-flux is increasing quickly, the temperature increases faster and the system may react late. We have found that a linear prediction ($l = 1$) works better in practice for safety purposes in these cases. The divertor overload case in Section 4 is an example of this: when the heating is increasing in

power at the beginning of the discharge, the heat flux is also increasing quickly and the temperature increases faster than the square root model.

(4) Note that Equation (3) requires the computation of the heat-flux in real-time. Currently, 2-dimensional thermal calculations in real-time are not possible at W7-X, since we cannot assume toroidal symmetry [15]. The heat-flux can be estimated roughly according to Equation (1) using the increase of the averaged corrected temperature between two frames $\Delta \overline{T'_{xy}}$ and the frame time interval $t_f = 0.01$ s, resulting in Equation (4).

$$q_{xy}(t) = \frac{\Delta \overline{T'_{xy}}(t)}{\left(\frac{t_f}{C^m_{xy}}\right)^l} \tag{4}$$

(5) $\Delta \overline{T'_{xy}}$ is then time-filtered in Equation (5) with a learning rate α to avoid triggering false alarms due to fast transients.

$$\Delta \overline{T'_{xy}}(t) = (1 - \alpha)\Delta \overline{T'_{xy}}(t - t_f) + \alpha \left(\overline{T'_{xy}}(t) - \overline{T'_{xy}}(t - t_f) \right) \tag{5}$$

Fast transients refer to fast physical events with high heat flux that last for a short period of time, not enough to actually heat and damage the PFCs. However, the temperature prediction is based on the measured heat flux, and it reacts assuming that the current heat flux is sustained for 0.16 s. A short high increase of heat flux can then trigger a false alarm, even lasting just few ms. To prevent this, the heat flux has to be averaged over time to react only to sustained heat flux changes.

(6) In case the heat-flux calculation under-estimates the real heat flux, we follow a conservative approach to stay on the safe side. The temperature threshold finally used (Equation (6)) is the minimum between the computed according to Equation (3) and an upper-bound $T^{th\_max}_{xy}$:

$$T^{th}_{xy}(t) = \min \left\{ T^{limit}_{xy} - q_{xy}(t)\left(\frac{t_r}{C^m_{xy}}\right)^l, T^{th\_max}_{xy} \right\} \tag{6}$$

The upper-bound values are:

- Divertor targets: 1034 °C
- Baffles and heat shields: 387 °C
- Pumping gap: 172 °C

These values are computed according to Equation (2) considering a $t_r = t_{max\_delay} = 0.11$ s (without safety margin) and assuming a heat-flux of 10 MW/m^2 for the divertors, 0.5 MW/m^2 for the baffle and heat shields and 0.2 MW/m^2 for the pumping gap. Note that we do not require here the safety margin of the reaction time because the calculation of the threshold is deterministic. The safety margin was added to account for the uncertainty of the heat-flux estimation. Also note that setting a temperature threshold upper-bound is equivalent to setting a heat flux lower-bound or a minimum value of the heat flux (see Equation (2)).

(7) An estimated risk $\hat{r}_{xy}$ is then computed in Equation (7) for each pixel with the averaged corrected temperature $\overline{T'_{xy}}$, the temperature threshold T^{th}_{xy} and a correction factor to consider the presence of surface layers S_{xy} [13].

$$\hat{r}_{xy}(t) = \frac{\overline{T'_{xy}}(t)}{T^{th}_{xy}(t)}(1 - S_{xy}) \tag{7}$$

Note that the estimated risk image is normalized for the different temperature limits and heat fluxes on the different PFCs, allowing to detect overloaded regions in the entire field of view naturally, avoiding the use of regions of interest.

(8) The overloaded regions are detected by clustering all pixels that have an estimated risk equal or greater than 1. If the cluster's area (in physical units) is bigger than a minimum area A_{min} an alarm is triggered. The scene model also provides the physical area covered by each pixel in the image.

4. Results from OP1.2 Campaign

During the operation phase OP1.2, inertially-cooled test divertor units were used to test the island divertor concept of W7-X. This allowed us to test the imaging software for PFCs protection without danger of damaging the water-cooled high-heat-flux divertor. Note that a water-cooled divertor has similar dynamics to an inertially-cooled divertor in non-steady-state conditions when transitioning from one steady-state to another because of a change of the heat flux. Hence, the validation of the overload detection algorithm with an inertially-cooled divertor can be translated to a water-cooled one in non-steady-state conditions. In steady-state the temperature remains constant, no overload can occur, and no safety system is needed.

During OP1.2, the system was not yet automated and it required manual inspection after the plasma discharges to assist the experiment leaders. We observed significant thermal events, and we reported them in [16]. These experiment programs make up a dataset of 25 discharges in different magnetic configurations. It includes high power discharges, overload experiments and unexpected events (hot spots, divertor overloads, strong leading edges, baffle overloads, and shine-through hot spots). This dataset has been used to test the Thermal Overload Detection algorithm.

Note that this dataset is not valid for inferring the true alarm rate versus false alarm rate, since it consists of many extreme cases. It is not representative of a normal operation. Also note that the false alarm rate will depend on how the machine is operated. If the machine is operated close to the limits the false alarm rate will be higher than operating at low power with larger safety margins. The dataset is valid, however, for qualitative analysis.
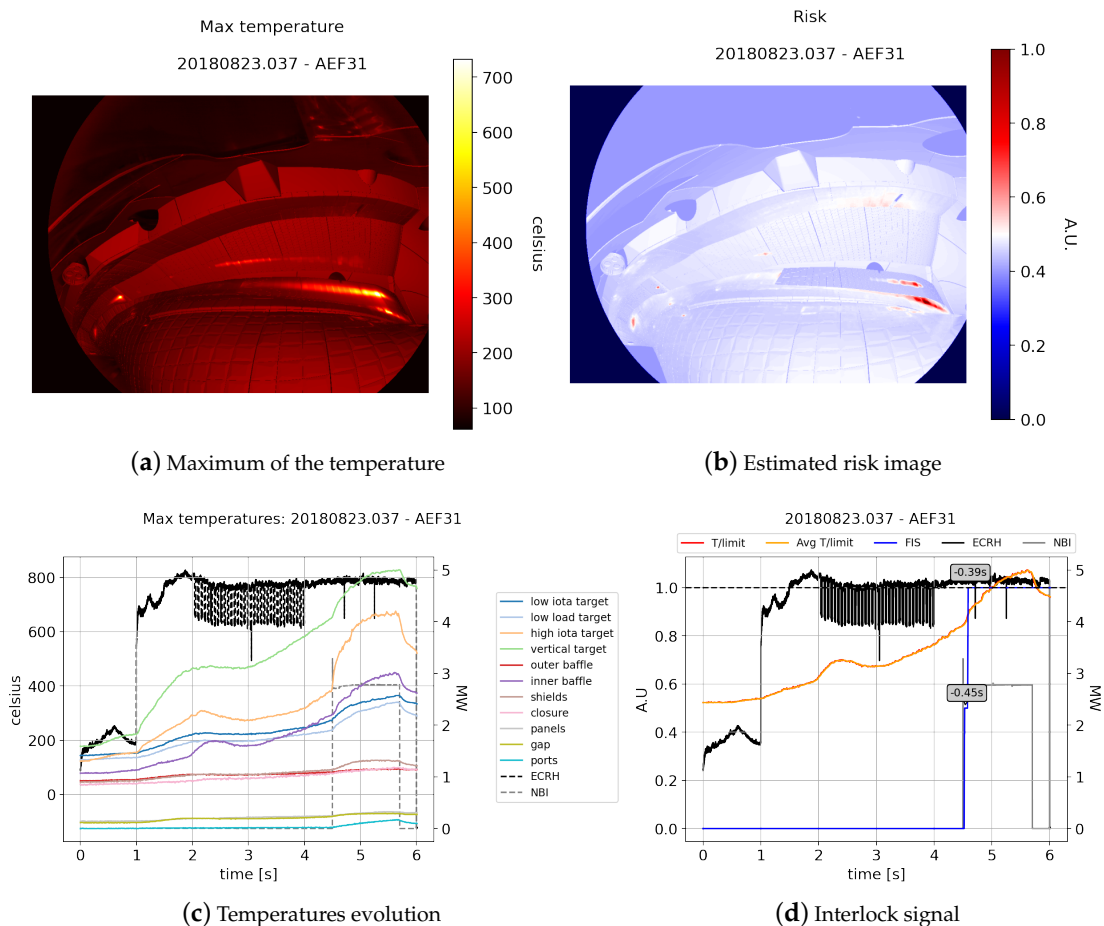
Since the system was not connected to the interlock during OP1.2 and the experiments were not interrupted, it is possible to measure the time difference between the time when the algorithm would have triggered the alarm (when the estimated risk reaches 1) and the time when the temperature reaches the limit. We define this as the anticipation time $t_{anticipation} = t_{T=T_{limit}} - t_{risk=1}$, and we can measure it for each pulse. The algorithm tries to anticipate the overload and compensate for the system delay t_{max_delay} (0.11 s). The goal is to have an anticipation time greater than the system delay. In this context, a late alarm occurs when the alarm is triggered after the temperature reaches the limit, when $t_{anticipation} < t_{max\_delay}$. We consider a false alarm when the alarm is triggered and no overload occurs within (arbitrarily) 1 s window, when $t_{anticipation} > 1$ s. Since the alarm is triggered before the event occurs (it is predictive) it is not possible to deterministically relate an alarm to an overload. The longer the time exists between the alarm and the overload the less likely they are related. Hence the 1 s threshold used here is arbitrary.

The reported results are obtained with the following parameter configuration: $n = 3$, $\alpha = 0.15$, $l = 1.0$, $A_{min} = 0$. Note that, for this test, we do not consider a surface layer factor and that we assume that the measured surface temperature is the true bulk temperature of the tiles ($S_{xy} = 0$ in Equation (7)). Figure 10 shows the anticipation times distribution on the validation dataset comprising 238 infrared sequences (25 discharges $\times$ 10 cameras, with some unavailable cameras for some shots). The results show no late alarms. The algorithm was also tested with 198 video sequence from randomly selected discharges. In these sequences no event occurred and no alarm was triggered.

Figures 11 and 12 show two examples of events: an unexpected baffle overload and a divertor overload that, in this case, was part of the experiment design.
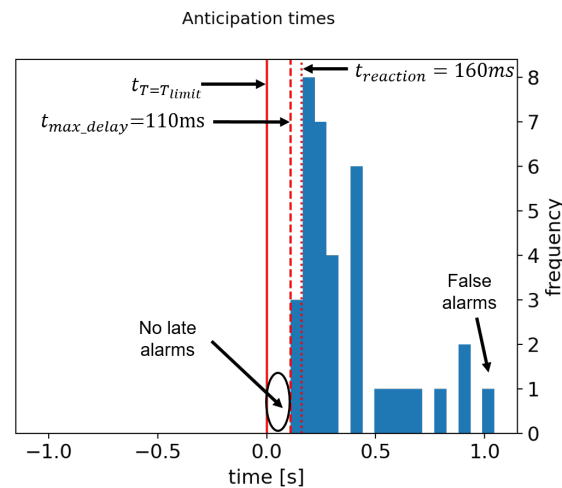
Figure 10. Distribution of the anticipation times on the validation dataset relative to the time when the temperature reaches the limit $t_{T=T_{limit}}$. All alarms are triggered in time (with $t_{anticipation} > t_{max\_delay}$). Only one false alarm is triggered, considering a false alarm when $t_{anticipation} > 1$ s.

(**a**) Maximum of the temperature

(**b**) Estimated risk image

(**c**) Temperatures evolution

(**d**) Interlock signal

Figure 11. A baffle overload example. (**a**) Maximum of the temperature over time. (**b**) Maximum of the estimated risk over time; the strike-lines which are within the temperature limits of the divertor are barely visible; the overload of the baffle stands out. (**c**) Maximum temperature for each PFCs over time; the inner baffle reaches 400 °C at 5 s. (**d**) Time evolution of the maximum temperature normalized by the temperature limit over the image, time evolution of the maximum averaged temperature normalized by the temperature limit over the image, and Fast Interlock Signal (0.5 corresponding to a warning and 1.0 to an alarm); We can observe that an alarm is triggered 0.39 s before the maximum temperature reaches the limit (when the baffle temperature reaches 400 °C) well in advance to compensate for the 0.11 s of the system delay.

(**a**) Maximum of the temperature

(**b**) Estimated risk image

(**c**) Temperatures evolution

(**d**) Interlock signal
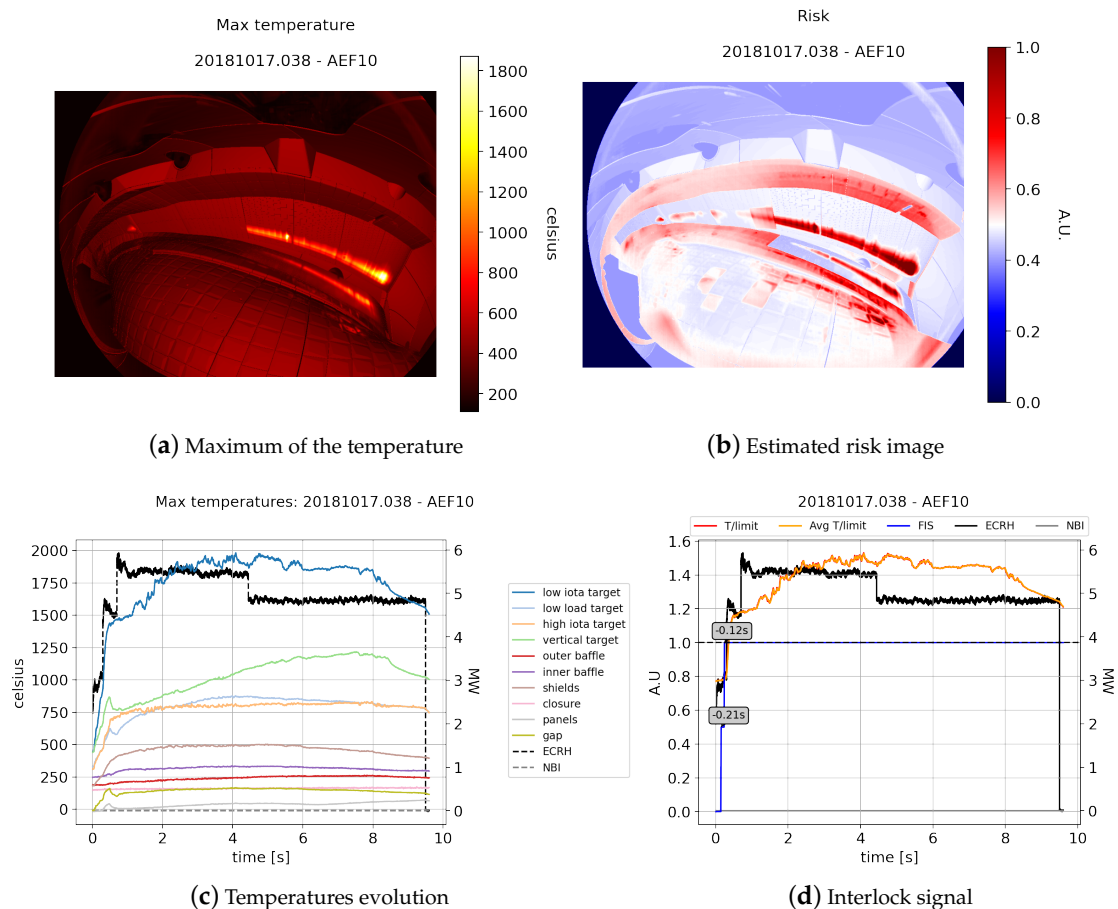
Figure 12. Divertor overload example. (**a**) Maximum of the temperature over time. (**b**) Maximum of the estimated risk over time. (**c**) Maximum temperature for each PFCs over time; the low-iota divertor target reaches 1200 °C at 0.4 s. (**d**) Time evolution of the maximum temperature normalized by the temperature limit over the image, time evolution of the maximum averaged temperature normalized by the temperature limit over the image, and Fast Interlock Signal (0.5 corresponding to a warning and 1.0 to an alarm); We can observe that an alarm is triggered 0.12 s before the maximum temperature reaches the limit (when the divertor temperature reaches 1200 °C) just in time to compensate the 0.11 s of the system delay.

5. Towards Feedback Control of Thermal Loads

With increasing plasma performance in subsequent operational phases, W7-X will require a feedback control system able to prevent plasma interruptions due to overloading to achieve high-performance long-plasma operation. The feedback control will have to trigger countermeasures when the risk of overheating reaches a threshold limit [17]. It has to be informed of the current development of the thermal events, their type, cause, and risk. This requires a high-level scene understanding which can only be achieved through advanced computer vision and machine learning techniques [18].

The Thermal Overload Detection algorithm can protect the machine through the interlock, but it cannot provide enough information on the thermal events for effective feedback control. We foresee a dual detection scheme with two algorithms running in parallel (see Figure 13). The Thermal Overload Detection algorithm guarantees that the PFCs operational limits are not exceeded and the Thermal Event Detection algorithm (still under research), detects promptly the thermal events, tracks them over time, classifies them and evaluates their risk. It informs the feedback control system of the ongoing events through a real-time network. The advantage of this dual detection scheme is that the Thermal Event Detection algorithm does not need to comply with the strict real-time and high-reliability requirements of a safety system and it can use more advanced state-of-the-art machine learning techniques.
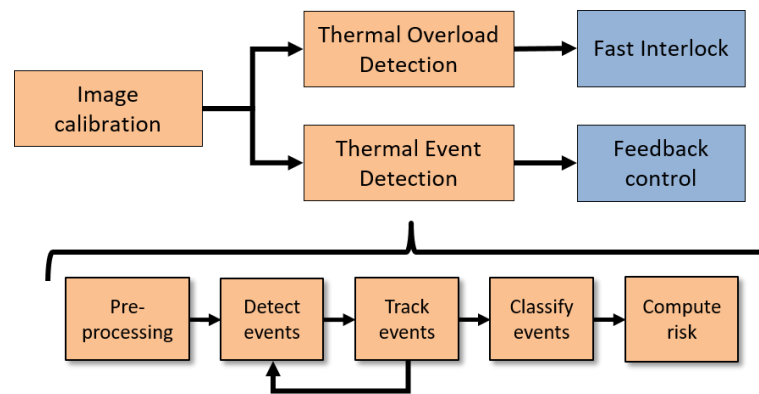
Figure 13. The machine is protected with a dual detection scheme. The Thermal Overload Detection algorithm detects when a PFC has risk of overloading and triggers the interlock to stop operation. The Thermal Event Detection algorithm promptly detects the thermal events, tracks them over time, classifies them and computes their risk. It sends all this information to the feedback control system. When an event reaches a certain risk, the feedback control has to take action to prevent the overloading of the PFCs and the plasma interruption.

The Thermal Event Detection algorithm can evaluate the estimated risk $\hat{r}_i$ of each event i using Equation (8).

$$\hat{r}_i = \max_{(x,y)\epsilon A_i} \hat{r}_{xy} \tag{8}$$

where $\hat{r}_{xy}$ is the estimated risk computed with Equation (7) for each pixel within the event segmented region $(x, y)\epsilon A_i$. The temperature threshold T_{xy}^{th} in Equation (3) is computed with a larger reaction time to allow the feedback control to take action ($t_r = 0.5$ s). When the estimated risk $\hat{r}_i$ reaches 1 for an event, the feedback control has to take the proper action according to the type of event to avoid the PFCs overloading and the interruption of the plasma by the Thermal Overload Detection algorithm.

In a first development version of the feedback control, the system will act on the heating systems, ramping down the ECRH (Electron Cyclotron Resonance Heating) power until a safe operation is reached or stopping NBI (Neutral Beam Injection) operation if fast particle losses are detected. In later stages, we will introduce more advanced feedback control of thermal loads by plasma detachment or by controlling the strike-line shape and positioning employing the control coils below each divertor unit.

The strike-line control is of special importance at W7-X. It is well known that the developing bootstrap current will move the strike-line towards the pumping gap in long discharges [19]. Also, beta effects can put some components at risk and this will require advanced feedback control. We can control the strike-line shape and position by employing the control coils located underneath the divertor targets of W7-X. By changing their currents, we can modify the topology of the magnetic islands at the plasma edge to reposition the strike-line on the divertor targets. Also, we can use the Electron Cyclotron Current Drive (ECCD) as an actuator to control the toroidal current. Ongoing research in this direction can be found in [20–22].

6. Conclusions

The protection of W7-X during steady-state plasmas requires checking in real-time the temperature of the PFCs against their limits, and triggering an alarm to the interlock and interrupt operation when these limits are compromised. Rather than using a fixed temperature threshold for each PFC and react when these thresholds are breached, a predictive algorithm sets a dynamic threshold in real-time for each pixel based on the different temperature limits and current heat-flux distribution. This dynamic threshold determines the estimated risk of overload during the processing time and it allows to trigger the alarm in advance to compensate for the delay in the system.

In this design, we avoid the use of regions of interest to monitor the different in-vessel components with different temperature limits. Instead, a scene model provides a pixel-wise correspondence of the entire field of view with the CAD model of the PFCs and their properties, such as emissivity and temperature limits. The scene model is constructed automatically by code and the operator does not need to know the position of each component and their limits. When the in-vessel components are changed, the scene models are updated automatically. The risk images encode the temperature limits of the different components, and the detection of overloads is performed naturally instead of treating each region of interest separately. This is useful when an event affects several components with different operational limits. An example of this is when a strike-line extends from the divertor targets towards the baffles.

We could validate the Thermal Overload Detection algorithm with the data acquired during the OP1.2 campaign when we used an uncooled test divertor. The results show the algorithm can detect all the safety-relevant thermal events that we manually observed during the campaign with enough anticipation time to compensate for the system delay. This proofs that the system would have promptly interrupted operation had it been connected to the interlock. The test of the algorithm with an uncooled divertor can be translated to a water-cooled divertor in non-steady-state conditions when an overload can occur. Therefore, the system is ready for commissioning and to provide a safe machine operation in the next campaign, when water-cooled PFCs will be used.

To achieve high-performance steady-state operation, however, it is also required to avoid most of the plasma interruptions through the feedback control. We envisage a PFCs protection system based on a dual detection scheme. The Thermal Overload Detection algorithm will protect the device using real-time image-processing techniques, while the Thermal Event Detection algorithm will use more advanced machine learning to detect, track and classify the different thermal events. The first provides high reliability aimed at machine protection and the latter allows a higher level of scene understanding, paving the way for a future feedback control of the thermal loads.

References

1. Endler, M.; Baldzuhn, J.; Beidler, C.D.; Bosch, H.S.; Bozhenkov, S.; Buttenschön, B.; Dinklage, A.; Fellinger, J.; Feng, Y.; Fuchert, G.; et al. Wendelstein 7-X on the path to long-pulse high-performance operation. *Fusion Eng. Des.* **2021**, *167*, 112381. [CrossRef]
2. Klinger, T.; Andreeva, T.; Bozhenkov, S.; Brandt, C.; Burhenn, R.; Buttenschön, B.; Fuchert, G.; Geiger, B.; Grulke, O.; Laqua, H.P.; et al. Overview of first Wendelstein 7-X high- performance operation. *Nucl. Fusion* **2019**, *59*, 112004. [CrossRef]

3. Jakubowski, M.; Endler, M.; Feng, Y.; Gao, Y.; Killer, C.; König, R.; Krychowiak, M.; Perseo, V.; Reimold, F.; Schmitz, O.; et al. Overview of the results from divertor experiments with attached and detached plasmas at Wendelstein 7-X and their implications for steady-state operation. *Nucl. Fusion* **2021**, *61*, 106003. [CrossRef]

4. Boscary, J.; Peacock, A.; Stadler, R.; Mendelevitch, B.; Tittes, H.; Tretter, J.; Smirnow, M.; Li, C. Actively Water-Cooled Plasma Facing Components of the Wendelstein 7-X Stellarator. *Fusion Sci. Technol.* **2017**, *64*, 263–268. [CrossRef]

5. Stadler, R.; Vorköper, A.; Boscary, J.; Cardella, A.; Hurd, F.; Li, C.; Mendelevitch, B.; Peacock, A.; Pirsch, H. The in-vessel components of the experiment W7-X. *Fusion Eng. Des.* **2009**, *84*, 305–308. [CrossRef]

6. Huber, A.; Kinna, D.; Huber, V.; Arnoux, G.; Balboa, I.; Balorin, C.; Carman, P.; Carvalho, P.; Collins, S.; Conway, N.; et al. The near infrared imaging system for the real-time protection of the JET ITER-like wall. *Phys. Scr.* **2017**, *T170*, 014027. [CrossRef]

7. Huber, V.; Huber, A.; Kinna, D.; Matthews, G.F.; Sergienko, G.; Balboa, I.; Brezinsek, S.; Lomas, P.J.; Mailloux, J.; McCullen, P.; et al. The Software and Hardware Architecture of the Real-Time Protection of In-Vessel Components in JET-ILW. *Nucl. Fusion* **2019**, *59*, 076016. [CrossRef]

8. Herrmann, A.; Drube, R.; Lunt, T.; De Marné, P. Real-time protection of in-vessel components in ASDEX Upgrade. *Fusion Eng. Des.* **2011**, *86*, 530–534. [CrossRef]

9. Sieglin, B.; Faitsch, M.; Herrmann, A.; Brucker, B.; Eich, T.; Kammerloher, L.; Martinov, S. Real-Time capable Infrared Thermography at ASDEX Upgrade. *Rev. Sci. Instrum.* **2015**, *86*, 113502. [CrossRef] [PubMed]

10. Jakubowski, M.; Drewelow, P.; Fellinger, J.; Puig Sitjes, A.; Wurden, G.; Ali, A.; Biedermann, C.; Cannas, B.; Chauvin, D.; Gamradt, M.; et al. Infrared imaging systems for wall protection in the W7-X stellarator. *Rev. Sci. Instrum.* **2018**, *89*, 10E116. [CrossRef]

11. Pisano, F.; Cannas, B.; Fanni, A.; Sias, G.; Jakubowski, M.W.; Drewelow, P.; Niemann, H.; Puig Sitjes, A.; Gao, Y.; Moncada, V.; et al. Tools for Image Analysis and First Wall Protection at W7-X. *Fusion Sci. Technol.* **2020**, *76*, 933–941. [CrossRef]

12. Herrmann, A.; Sieglin, B.; Faitsch, M.; Team, A.U. Surface Temperature Measurement of In-Vessel Components and Its Real-Time Validation. *Fusion Sci. Technol.* **2016**, *69*, 569–579. [CrossRef]

13. Ali, A.; Niemann, H.; Jakubowski, M.; Pedersen, T.S.; Neu, R.; Corre, Y.; Drewelow, P.; Sitjes, A.P.; Wurden, G.; Pisano, F.; et al. Initial results from the hotspot detection scheme for protection of plasma facing components in Wendelstein 7-X. *Nucl. Mater. Energy* **2019**, *19*, 335–339. [CrossRef]

14. Vincent, L. Morphological Grayscale Reconstruction in Image Analysis: Applications and Efficient Algorithms. *IEEE Trans. Image Process.* **1993**, *2*, 176–201. [CrossRef]

15. Sieglin, B.; Faitsch, M.; Herrmann, A.; Martinov, S.; Eich, T.; Team, A.U. Real-Time infrared thermography at ASDEX Upgrade. *Fusion Sci. Technol.* **2016**, *69*, 580–585. [CrossRef]

16. Sitjes, A.P.; Gao, Y.; Jakubowski, M.; Drewelow, P.; Niemann, H.; Ali, A.; Moncada, V.; Pisano, F.; Ngo, T.T.; Cannas, B.; et al. Observation of thermal events on the plasma facing components of Wendelstein 7-X. *J. Instrum.* **2019**, *14*, C11002. [CrossRef]

17. Mitteau, R.; Belafdil, C.; Balorin, C.; Courtois, X.; Moncada, V.; Nouailletas, R.; Santraine, B. WEST operation with real time feed back control based on wall component temperature toward machine protection in a steady state tungsten environment. *Fusion Eng. Des.* **2021**, *165*, 112223. [CrossRef]

18. Martin, V.; Moncada, V.; Travere, J.M.; Loarer, T.; Brémond, F.; Charpiat, G.; Thonnat, M. A Cognitive Vision System for Nuclear Fusion Device Monitoring. In *International Conference on Computer Vision Systems*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 163–172.

19. Gao, Y.; Geiger, J.; Jakubowski, M.W.; Drewelow, P.; Endler, M.; Rahbarnia, K.; Bozhenkov, S.; Otte, M.; Suzuki, Y.; Feng, Y.; et al. Effects of toroidal plasma current on divertor power depositions on Wendelstein 7-X. *Nucl. Fusion* **2019**, *59*, 106015. [CrossRef]

20. Pisano, F.; Cannas, B.; Fanni, A.; Sias, G.; Gao, Y.; Jakubowski, M.; Niemann, H.; Sitjes, A.P. Learning control coil currents from heat-flux images using convolutional neural networks at Wendelstein 7-X. *Plasma Phys. Control. Fusion* **2020**, *63*, 025009. [CrossRef]

21. Böckenhoff, D.; Blatzheim, M. Application of improved analysis of convective heat loads on plasma facing components to Wendelstein 7-X. *Nucl. Fusion* **2019**, *59*, 086031. [CrossRef]

22. Böckenhoff, D.; Blatzheim, M.; Hölbe, H.; Niemann, H.; Pisano, F.; Labahn, R.; Pedersen, T.S. Reconstruction of Magnetic Configurations in W7-X using Artifical Neural Networks. *Nucl. Fusion* **2018**, *58*, 056009. [CrossRef]