


Project Deliverable

Project Number: 325275	Project Acronym: SAPPHIRE	Project Title: System Automation of PEMFCs with Prognostics and Health management for Improved Reliability and Economy
Instrument: Collaborative Project (CP)		Thematic Priority FUEL CELL AND HYDROGEN JOINT UNDERTAKING
Title: Report on the advantages of PHM-aware control		
Contractual Delivery Date: April 30, 2016		Actual Delivery Date: May 5, 2016
Start date of project: May 1, 2013		Duration: 36 months
Organisation name of lead contractor for this deliverable: SINTEF		Document version: Version 1.1
Organisation notes:		

Dissemination level (Project co-funded by the European Commission within the Seventh Framework Programme)		
PU	Public	PU
PP	Restricted to other programme participants (including the Commission)	
RE	Restricted to a group defined by the consortium (including the Commission)	
CO	Confidential, only for members of the consortium (including the Commission)	
		Public



Authors (organizations):

Johannes Tjønnås (SINTEF), Federico Zenith (SINTEF), Ivar Halvorsen (SINTEF)

Abstract :

This deliverable summarises the potential and validation of the proposed, implemented and tested control strategies within the project.

Keywords :

Control; Monitoring, Prognostics; humidity; air bleed; CO poisoning;

Revision History

Rev.	Date	Description	Author (Organisation)
1.0	April 12, 2016	First draft	Johannes Tjønnås (SINTEF)
1.1	May 3, 2016	Revision & QA	Federico Zenith (SINTEF)




Table of Contents

1	SCOPE OF THE DELIVERABLE	4
2	CONTROLLER MECHANISMS IMPLEMENTATIONS AND VALIDATION	4
2.1	Main controller and recorded performance	4
2.2	Dry-Out controller and recorded performance	5
2.3	CO-mitigation controller and recorded adaptive air bleed performance	6
2.4	In-situ polarization curves and recorded data	9
3	DANTHERM POWER STACK DATA ANALYSIS AND POTENTIAL	9
3.1	Monitoring of current voltage stack properties	10
3.2	Restart performance effect	15
3.3	Potential of Rejuvenation Techniques	17
4	CONCLUSIONS AND FURTHER WORK	20

List of Appendices

NO TABLE OF CONTENTS ENTRIES FOUND.

1 Scope of the Deliverable

The objective of this deliverable is to evaluate the project control strategies, specified, developed and implemented in the confidential deliverables D7.2 and D7.4. This is done by inspecting the micro-CHP system data made available from Dantherm Power's experiments. The middle level control and observer implementations were validated in ZSW's lab and presented in D7.4, and are not repeated here due to confidentiality. Plots showing the main mechanisms are however shown together with short comments.

No clear system long-term degradation was shown in the first system experiments, i.e. no clear harmful operating conditions was detected. However, a connection between the number of restarts and system performance was indicated. It was therefore decided to restart the system regularly in the second system experiment in order to provide more data for identifying the restart-performance relation.

2 Controller mechanisms implementations and validation

The overview of the controller architecture is presented in Figure 1. The focus of this chapter is highlighting the mechanisms of the middle level controls, acting as a layer between the diagnostics and prognostics and the regulatory layer

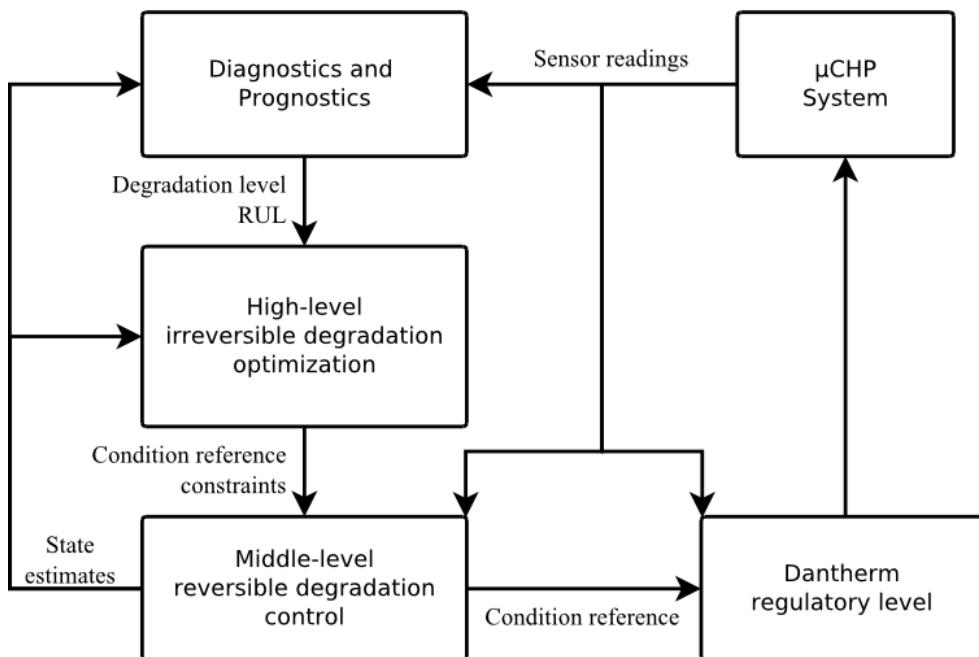



Figure 1 Controller structure

2.1 Main controller and recorded performance

The task of the main controller is to handle and prioritize the middle-level controller actions. The controller actions are prioritized according to the following order: in-situ polarization



curve, CO controller and drying controller, where only one middle-level controller action is available at a given time.

The main controller should also maintain knowledge of the state information from the regulatory level, mainly in order not to allow the controllers to interfere with each other's function, or take into account data while the regulatory level changes its operating states. In some cases, however, it was found that the CO compensator is activated in a wrong state and it erroneously increased air bleed during current down-ramping, as in Figure 2.





Figure 2: Case of mistaken air bleed increase. Data from December 30, 2015.

In this case, the reduction of current caused a corresponding increase in voltage, and the CO compensator, which was coincidentally started immediately before the current ramping began, attributed the voltage increase to removal of CO poisoning, and proceeded to increase air bleed from 0.4 % to 0.75 %, likely without the actual necessity of it.

2.2 Dry-Out controller and recorded performance

Another medium-term controller designed by SINTEF was the anti-drying controller. Initially a controller to avoid both flooding and dry-out conditions, it was downsized after ZSW showed in their laboratory that no flooding was possible in Dantherm Power's stack within the feasible range of operations. Since the anti-flooding controller would only have been one more item that could malfunction and had no practical purpose any more, it was removed.

The anti-drying controller monitors the flow resistance through the cathode side of the stack. Since flow in the small cathode channels is laminar, the flow resistance is supposed to be approximately constant, and the pressure drop should be about linear with gas flow. However, in actual operation, the humidity in the cathodic side must be kept slightly above



100 % in order to maintain the membrane fully humidified, and some droplets are present in the channels: this increases the flow resistance somewhat.

The flow resistance is therefore monitored with an additional pressure sensor over the cathode (in the demonstration, two sensor were used at inlet and outlet; a differential pressure sensor would have been sufficient), which Dantherm Power acquired for the nominal price of 68 €. Whenever the measured flow resistance drops below 1.6 mbar/(L/min), it is assumed that drying is occurring, and the controller requests a reduction of stack temperature to increase condensation.

In the demonstration, data analysis indicates that the controller detected drying conditions on a regular basis, as shown in Figure 3.




Figure 3: Operation of the anti-drying controller. Data from October 30 to November 3, 2015.

In particular, the controller was activated regularly when the system current was brought to its lowest level, 15 A. The reason for this is that the reduced reaction rate in the system generated less water on the cathode, which then induced drying. The controller could however not affect the situation, since in this condition its manipulated variable, the coolant inlet temperature, was already at its saturated minimum level.

Whereas the controller did not influence the performance of the system, it was able to diagnose drying conditions consistently through the demonstration.

2.3 CO-mitigation controller and recorded adaptive air bleed performance

A novel element of the control system was an adaptive air bleed strategy that exploited the fast voltage response upon increase of air bleed in the presence of CO poisoning. The system



worked mostly as intended, but some lessons were learned in terms of its interaction with other processes affecting voltage in the system.




Figure 4: Typical pattern of bleed increase. Data from December 4, 2015.

In Figure 4, a typical bleed increase pattern is illustrated. The system starts a bleed increase at 18:20, starting from 0.4 % (a very low value for industry standards). As the bleed is gradually increased in steps, voltage increases correspondingly, until no further response is detected at 1 %; the system then settles at 0.85 %, going down one and a half times the step up, and waits six hours for the next attempt at increasing air bleed.

The purpose of the last downward step is to allow the controller to gradually reduce air bleed if no CO contamination is detected. Since this contamination is slow at the ppm concentrations that are typical of μ CHP reformat, it is necessary to wait several hours between tests.

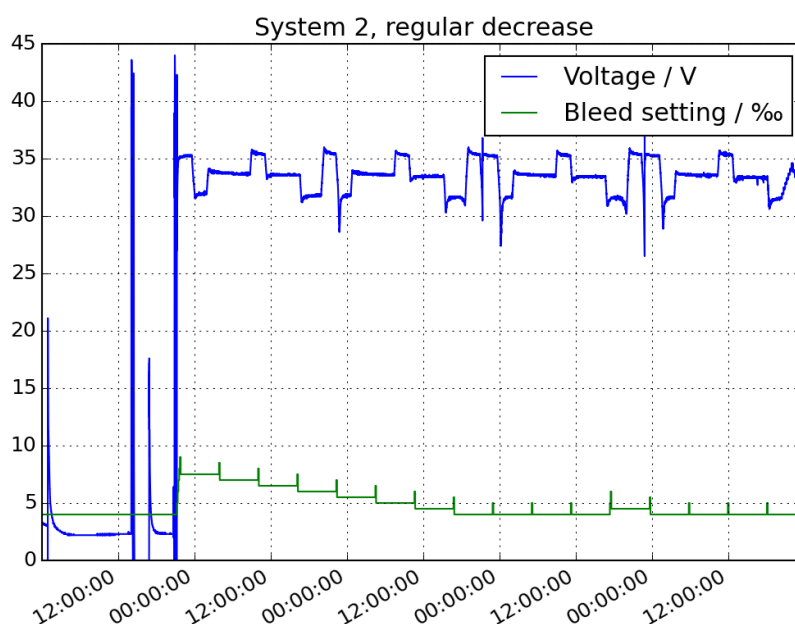



Figure 5: Gradual reduction of air bleed. Data from January 28 through February 1, 2016.

Over several days, the air bleed may be gradually reduced, as seen in Figure 5. As system #2 is started up in the evening of January 28, the CO compensator increases the air bleed rapidly up to 0.8 ‰, to compensate for the high CO concentration at system start-up, due to transients in the fuel processor. The air bleed is then gradually reduced until January 30, when it stabilises around the minimum value of 0.4 ‰. Because of the application of a variable load pattern, it would have been very difficult to detect any changes in voltage over a long time, as clear from the figure.

As previously indicated in Figure 2, a scheduled air-bleed test was influenced by a simultaneous change in load and produced an unnecessary increase of air bleed. The damage done by these interferences is minor, both because they are unlikely and have little consequences, and can in any case be removed by simply deactivating the CO controller during ramping (such a feature is actually already available, but was not employed). In other applications, such as automotive, this might be a more serious problem since the load is changing continuously, which causes variable CO poisoning as well as it blocks the execution of the CO compensator. A more sophisticated controller would be required, possibly based on model-based estimation. Such a controller has also been synthesized and tested in a simulation environment, but was not implemented in this application since constant loads are more common in μ CHP applications.

Finally, it was reported by the operators at Dantherm Power that the compensator sometimes failed to act when significant CO poisoning was present. This was typically the case after large steps in current, which caused equally large transients in the fuel processor: CO concentration levels were once manually measured to be over 50 ppm, even if only temporarily. In these cases, the controller could be forced to start a test immediately after a



transient in the fuel processor is detected (this functionality is also already available in the controller).

2.4 In-situ polarization curves and recorded data

Both demonstrated systems were outfitted with a routine to sample automatically a partial polarisation curve every week, but due to operational difficulties System #1 did not develop a complete set. The polarisation curves measured for System #2 were more in number and evenly spaced, and are presented in Figure 6.

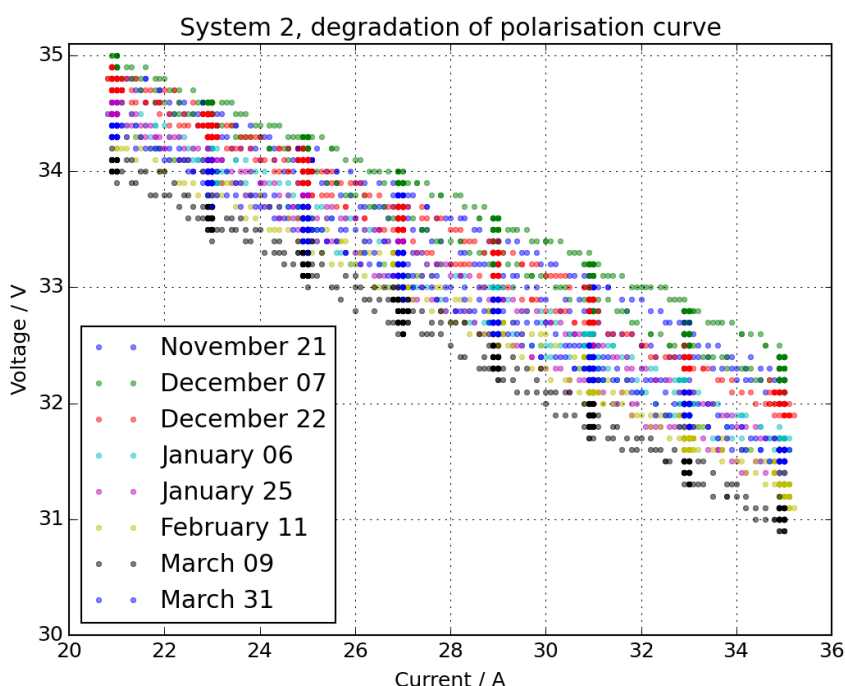


Figure 6: Polarisation curves over the demonstration.

The evolution of these polarisation curves indicates that there was not much performance degradation initially, with even a performance improvement in the beginning of December 2015; this corresponds with the period in which regular start-stops were performed to boost the system's performance. From January, performance degraded consistently until March, when a number of unplanned shutdowns resulted in a new regeneration of the stack. System performance is in the following chapter discussed in more detail.

3 Dantherm Power stack data analysis and potential

For each of the tests, the Dantherm Power systems were ran for 3000 hours. In order to track the stack performance at different loads, it was proposed to identify the voltage-current relation. This allows a parameterized monitoring of the stack, and also provides a model for comparing and evaluating the different systems at different times. In the following an overview of the findings is presented.



3.1 Monitoring of current voltage stack properties

The Sapphire operating load profile, repeating itself every day, render voltage-current relation monitoring possible. A simple least squares method is implemented in order to parameterize and keep track of the stack performance over time. Figure 7 show data recorded during two days for the given load profile.

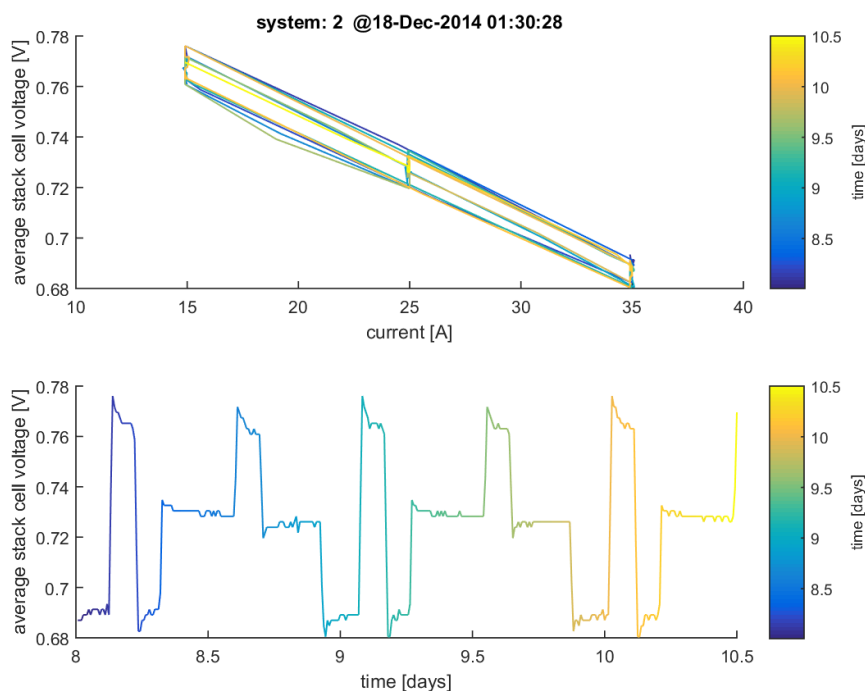


Figure 7 Voltage-current relation

In order to validate this approach, automatic in-situ polarization curve experiments were implemented in the μ CHP system. In Figure 8 two examples of the implemented experiments are shown. The dotted lines are based on the best-fit linear least square parameters identified for the data. The same method is used for extracting the parameters from regular operation. In this case the method uses data from a 2-day sliding window. The main difference between regular operation and the polarization experiment is that the experiment has higher resolution, but a smaller region of operation. This implies a locally more precise parameter identification for the experiment, while the solution from the continuous operations are valid over a larger region.




Figure 8 In-situ polarization curve examples

Typical parameters over a longer period are shown in Figure 9. These parameters are the slope of the polarisation curve in the region of interest, and the *projected* voltage at zero current; this is *not* the open-circuit voltage, but the *intercept* on the voltage axis that would correspond to a linear voltage model. This *intercept voltage* is not directly measurable, but is inferred from the data. In this figure the green plot marks the border of the 95% confidence on the parameters. System restarts involve dynamic behaviour such that a linear relation between voltage and current are typically not maintained, and the identified parameters are not valid. In order to cope with this situation only high-correlation results are considered. In Figure 10 the voltage-current correlation is shown for the whole period, but only data points with correlation less than -0.95 are considered valid. The voltage dynamics (hysteresis) introduced by the load profile will affect the parameter estimate, and a slightly underestimated stack performance is recorded. This does however not cause any concern as long as the load profile stays the same; also, monitoring of the relative performance is of the main interest.

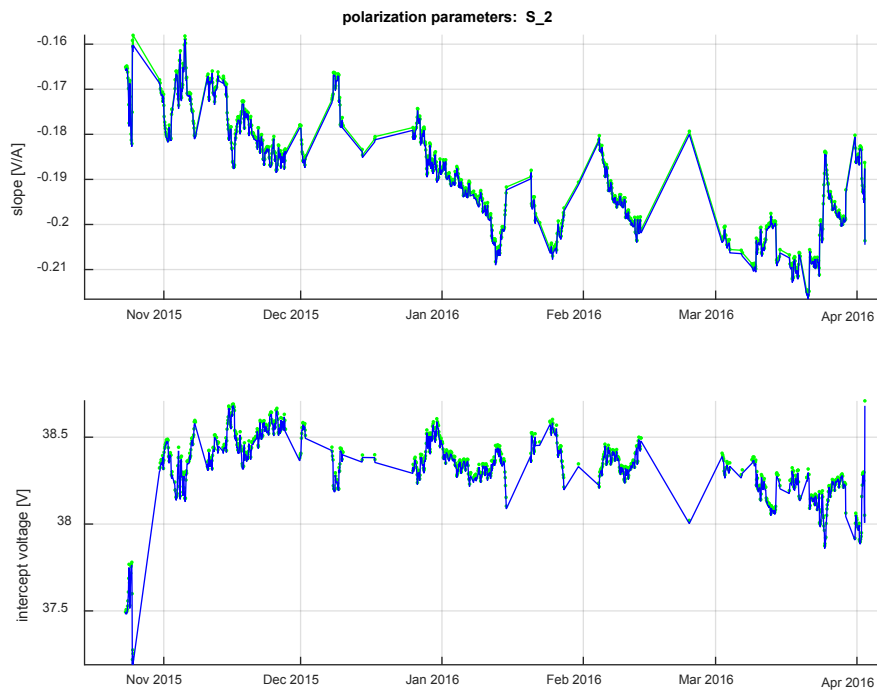



Figure 9 Voltage-current parameterized by least-squares parameters

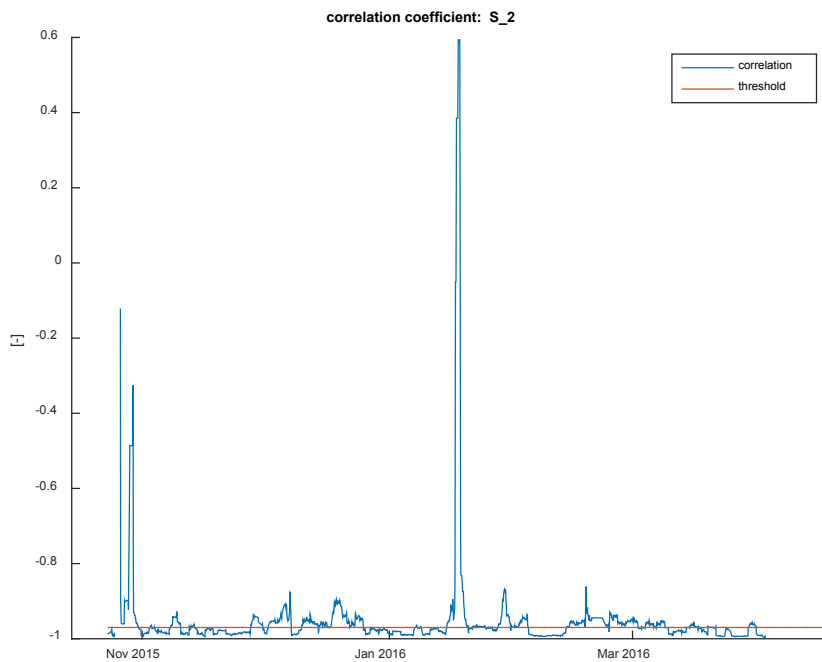



Figure 10 Voltage-current correlation, only high negative correlation is considered



In Figure 11, the parameters are shown for the main systems and the main testing periods of the Sapphire project. The blue and red stars mark the times and parameters calculated based on the in-situ polarization tests (only done for the last test batch). For both systems the continuous estimates develop consistently with the in-situ tests. But for System 2 there are offsets by increased slope steepness and intercept voltage. As the In-situ tests are carried out at higher currents, this indicates that the system may be operated closer to its mass transport loss boundaries (more negative double partial derivative of the polarization curve).

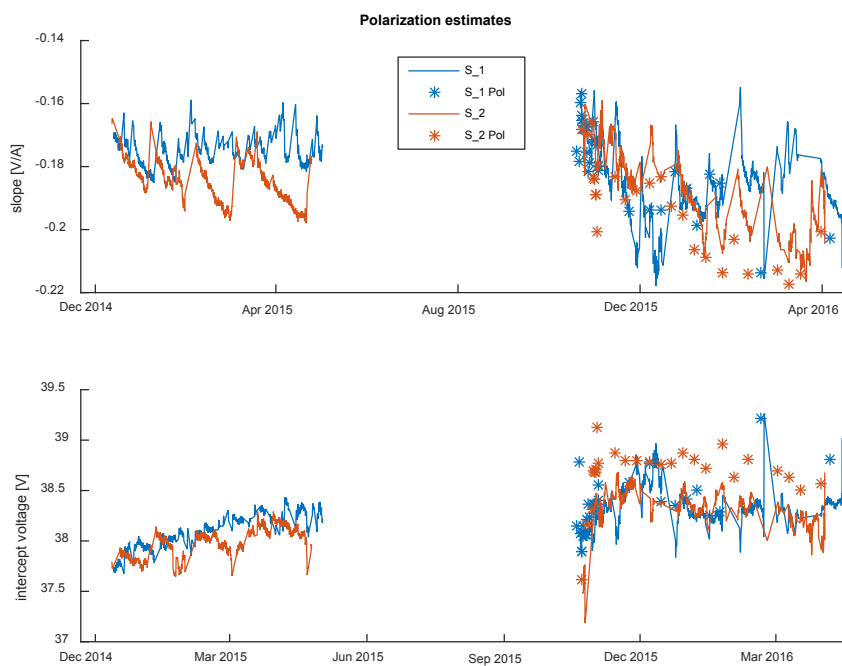


Figure 11 Voltage-current relation summary for all Dantherm Power tests





Figure 12 Residual plot for Dantherm Power system 2, first testing session

In order to monitor the stack performance, a scalar indicator is chosen and calculated for each valid time step. This value is the estimated voltage at 35 A, which is the normal operating current of the system; the performance of both systems is shown in Figure 13. In the first testing batch, the middle-level controllers were not implemented. The performance of system 1 was improving in the first test and second half of the second test period. System 2 had a stable performance in the first test, while it was decaying in the second one. Based on the first test performances it was hypothesized that repeated restarts of the system may improve importance.





Figure 13 Estimated system performances shown for the Dantherm Power tests.

3.2 Restart performance effect

During the project, restart frequency has been discussed as a possibility for increased stack performance. In order to evaluate this hypothesis, a frequency measure was defined: the number of restarts within a period of four weeks was considered a good balance between measure resolution and system response time. In Figure 14, the frequency estimates for each system and test are shown.




Figure 14 Identified start/start stop frequency over a four week period, starts/stops happening within 1 day of operation are not considered.

The system restart effect on the stack performance can be tested by correlating the performance and the frequency measures. The results, shown in Figure 15, indicate that the hypothesis holds, with the exception of results from system 2 in the first batch; their lower correlation may be due to the lower frequency of restarts. According to the chosen metric, the correlation sensitivity is reduced and the results will be dominated by the on-time decay slope. The on-time decay should therefore be considered together with the general system decay in a complete analysis. This is left for the discussion in the next chapter and future projects.





Figure 15 Start/stop frequency correlated with the performance estimate.

3.3 Potential of Rejuvenation Techniques

In this section the potential for estimating the system Remaining Useful Life (RUL) based on restart mechanisms are discussed.

In the previous section, a correlation between overall development of cell performance and the preceding start/stop frequency was observed. This observation was based on identifying the voltage-current relation within two cycles of the load profile. This allowed a wider and normalized set of stack data to correlate with restart frequency. In the following a slightly simpler approach is used by considering the stack voltage performance directly at 25 A. The results can be generalized by using the model found in the previous section. The 25 ampère level is selected due to symmetry nullifying the transients and the amount of data at this operating condition.

During the tests there are several stop/starts. Some are scheduled, but some are not, and in some period with some operational problems the stops are more frequent. The detailed nature of the starts and stops, and also the procedure of handling these states are confidential and not known to the authors.

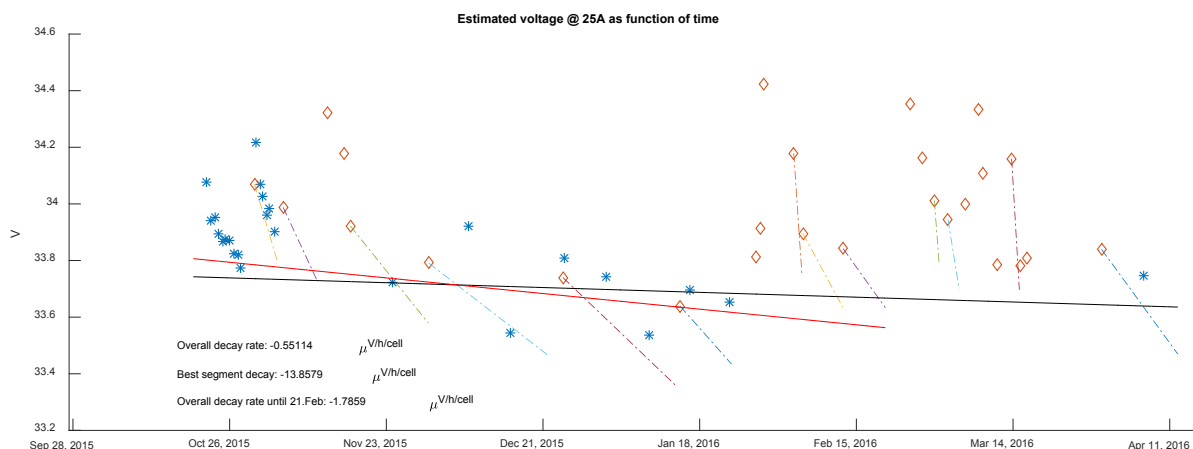



Figure 16 Development of the voltage for each continuous operating period and overall for Dantherm Power system 1.

Figure 16 shows how the voltage at 25 A develops during the test. In this figure every diamond marks the start of a continuous operating period. The dashed lines mark the fitted performance, within continuous operation, based on 25 A load current.

It can be observed that there is a certain decay of performance during the continuous operation periods, but then there is a voltage increase at the next startup, here denoted as a *rejuvenation*. The local decay in each period is the slope of each segment. In some periods the slopes are quite steep: this is typically due short up-time and transient issues. In some of the starts above, the operation has only lasted for a very short time, and therefore a decay line is not fitted. However, each diamond still represent a full startup, and a voltage has been registered for a short period.

A detail of the measured voltage in the continuous operating period starting February 12th up to the 21th is shown in Figure 17: it can be observed that the fitting of the straight line through the voltages at 25 A describes the decay as function of time quite accurately, limiting the transient effects.

In Figure 16, the asterisks (*) mark the estimated 25 ampere results from the polarization curve experiments. Ideally, these would be found on the segment in that operation period. It fits reasonably well, but with some deviation for a few points.

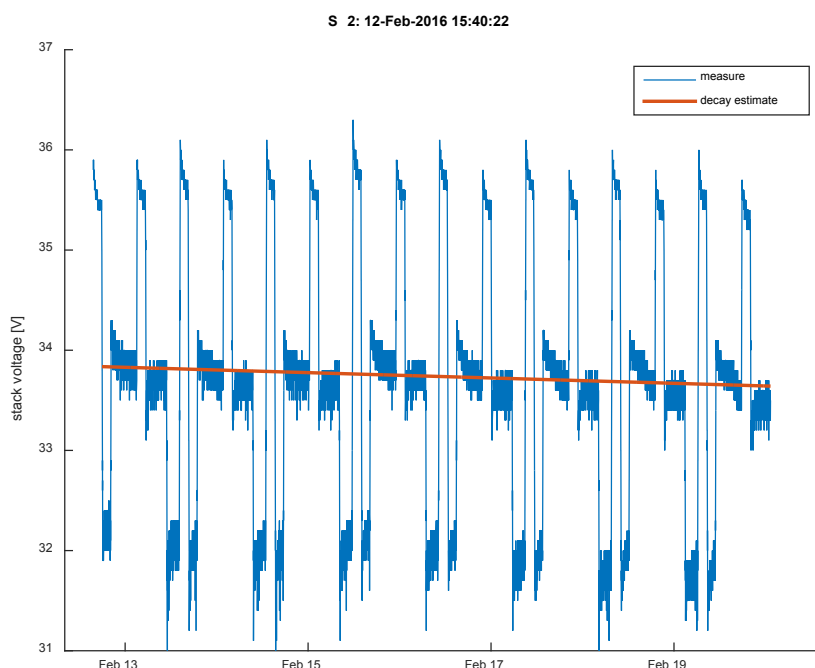


Figure 17. Detail of the data and a single fitted decay line from 12-21. February.

The overall decay calculated at all periods in the data set where the current is 25 A shows a decay rate of $-0.55 \mu\text{V/h/cell}$. (Solid black line in Figure 16). However, it is observed that there is a period from end of February with many stops and short on-time periods, and then the last long operation period in April is actually back to the November level. By analysing data up to February 21, the overall slope is found to be $-1.79 \mu\text{V/h/cell}$. The least step decay in continuous operation is in early December with a decay rate of $-13.86 \mu\text{V/h/cell}$. Note that there is significant less overall decay than the decay that can be observed within a continuous operation period.

Thus, for a better prediction of remaining useful life, the rejuvenation after each stop/start must be taken into consideration.

The relation between the stop/start frequency and overall decay is illustrated conceptually in Figure 18. The red curve indicates continuous operation until end of life. The blue and green curve have a certain stop/start frequency where the voltage increases just after a start-up compared to the voltage just before the previous stop.




Figure 18. Indicated voltage development with varying start/stop frequency

The overall decay is reduced with increasing higher frequency. Thus, in theory, by selecting a suitable frequency (blue and green), the expected remaining useful life (RUL 2 and 3) should be extended compared to the continuous operation (red, RUL 1).

The implementation may be flexible, e.g. a constant frequency imposed on the system, or some kind of adaptive procedure that depends on the actual performance development and "natural" start/stops of the system.

It is reasonable that there is some maximum frequency when no more improvement can be expected. This has not been concluded from data. It should also be considered whether too frequent start/stops may have other negative effects on the fuel cell or other components in the rest of the system like the reformer.

However, an improvement from RUL 1 for the long-term continuous operation prediction towards RUL 2 when there are a number of stop/starts seems to be supported by all the experimental data, and may be exploited by implementing an appropriate control procedure.

Further studies should focus on the physical mechanisms that led to these observations. A better understanding of the phenomena may be used to develop better and more precise PHM control algorithms.

4 Conclusions and further work

In this report a qualitative evaluation of the developed control structure implemented on the Dantherm Power system was done. The quantitative verification and validation of these implementations was done through dedicated experiments at ZSW.

The operation and actions of the controllers were according to expectations, except for some incidents where the air-bleed controller was started during system load change. This can either be handled at the main controller level, or by implanting an air-bleed controller that does not assume constant load.

No clear long-term performance decay was shown in the first system tests, but a dependence on system restarts was inferred. For the second tests it was therefore decided to run system 2 with regular restarts for an initial period: correlation between system restarts and performance was shown for three of the four tests. Higher correlation is



expected by accounting for local up-time decay, especially for the cases with low restart frequency; furthermore, the normalization gained by removing frequency and uptime dependency will provide a more detailed measure of the system decay.

Concepts and ideas for controlled system restarts are proposed, but they need further study. In order to be validated, datasets from systems running for the full life need to be analysed.