Evaluation of Flexibility Interoperability Platform

June 17, 2019

Uptime

EcoGrid 2.0 is a research and demonstration project funded by EUDP (Energiteknologisk Udviklings- og Demonstrationsprogram). The 9 partners in the project are:

Main Authors:

Name/Partner	Email
Troels Brødsgaard	troels@uptime.dk

TABLE OF CONTENTS

1 Introduction

Executive summary

To provide interoperability for incompatible home energy management systems in EcoGrid 2.0, a Flexibility Interoperability Platform was built. The evaluation of the platforms performance found that the FIP could perform activations within acceptable time frames and with success rates higher than those of the downstream HEMSs, but also showed that using a centralized service for this task carries the risk of introducing new failure scenarios unrelated to the HEMSs themselves.

Throughout heating season 3, the Flexibility Interoperability Platform (FIP) was to be used by the aggregators to facilitate their communication with home energy management systems (HEMSs) in the EcoGrid 2.0 households.

The need for the FIP was identified previously in the project. It acts as a translation layer between the EcoGrid 2.0 interoperability protocol (described in deliverable D4.5.2) and the native protocols of the HEMSs. Such functionality will always be required as long as the HEMSs does not implement the interoperability protocol themselves, and while the scope of EcoGrid 2.0 is limited to control of HEMSs, the functionality can be extended to the broader spectrum of distributed energy resources (DERs).

The point of this evaluation is to answer research questions 1 and 2 (RQ1 and RQ2):

1. What are the principles and pitfalls when building an interoperability platform to enable DER/Aggregator interoperability without protocol adoption at the DER level?
2. How can such an interoperability platform be implemented in the EcoGrid 2.0 project, and what are the control characteristics that aggregators experience when activating their portfolio?

From the period January 15th to April 5th the entire EcoGrid 2.0 setup was running in autonomous mode. During this period, the FIP was a central facilitator in all demonstrations involving HEMSs. The evaluation of RQ1 is based on experience and data gathered during autonomous operation in heating season 3, rather than specific demonstrations. In-depth description of this is contained in deliverable D5.4, while some points will be summarized below.

In addition to autonomous operations of the FIP, specific tests were conducted against the FIP to gather data to evaluate the portfolio control characteristics (RQ2). These tests are described below.

2 Demonstrations

Autonomous operation

As described in D5.4, an important principle in the FIP architecture is the Virtual DER, and isolation between Virtual DERs. While this was achieved in the FIP by implementing the Virtual DERs as lightweight processes with preemptive scheduling and isolated memory space, common error cases were caused by reliance on shared, external systems. The most prevalent examples of these have been:

- Log file rotations inadvertently destroying the logging pipe, rendering the FIP unable to log and filling up the log buffer until running out of memory.
- Queries accessing database tables with transaction auditing and data collection slowed to a crawl due to improper utilization of indexes or indexes growing too large to fit in memory.
- Downstream HEMS vendor cloud systems being overwhelmed due to sudden and large spikes in the number of requests.

Portfolio control characteristics

To discern portfolio control characteristics, the FIP is tested by a load controller, which plays the part of an aggregator. It makes a selection of households, which it then subjects to a sequence of control signals through the FIP. The control signals throttle the heating equipment of the selected households, checks that the change has been effectuated and then unthrottles the heating equipment. The message flow is recorded in order to calculate response times and error rates. The tests were carried out against varied portfolio sizes.

The response times from two such tests are plotted in Figure 1 and Figure 2. As can be seen from the plots, the response times vary dramatically depending on the number of households being activated simultaneously – the median response time for successful control is 11 seconds for 249 households and 52 seconds for 602 households, and the maximum response times are 15 and 66 seconds respectively.

The main reason for the large difference in response times is that the downstream HEMS vendor cloud systems struggle to handle the additional load. This means response times from the downstream systems increase, and fail relatively more often, meaning the FIP must perform more retries.

This difference in response times should not have a large impact on the aggregators ability to activate flexibility due to the time constants involved in household heating. But it should influence the aggregators control algorithm, as they must be prepared to wait for up to a minute before concluding on the success of an activation.

It is notable that the median response times in case of errors is lower than for successful control, and that the minimum and maximum error response times are respectively lower and higher than those for successful control.

The reason for lower median response times in case of errors is that errors are usually discovered early in the communication chain. Either the FIP or the HEMS cloud control system may error out early if it is aware of communication issues with a HEMS that cannot be solved by retries. In contrast, successful control attempts always have to reach the HEMS, and possibly peripheral equipment connected by low quality connections with high latency. Thus, if the FIP knows that a given HEMS is offline it can return an error with sub-second response times, while the fastest successful control actions will always take a few seconds.

The reason why the slowest errors have higher response times than the slowest successful control attempts is that some error conditions can only be determined when HEMS interactions time out, which obviously takes longer than successful HEMS interactions.

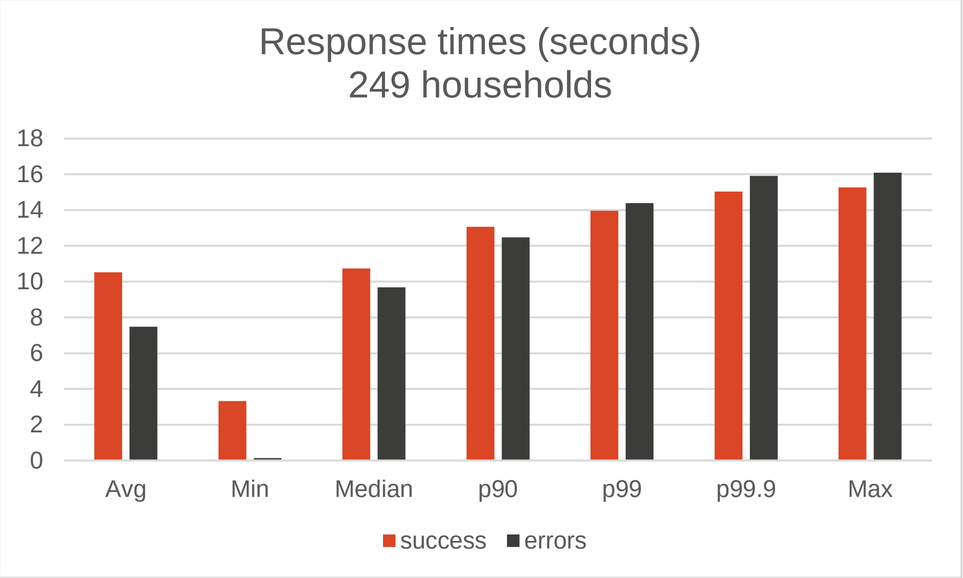

Figure 1 - Response times in test with 249 households

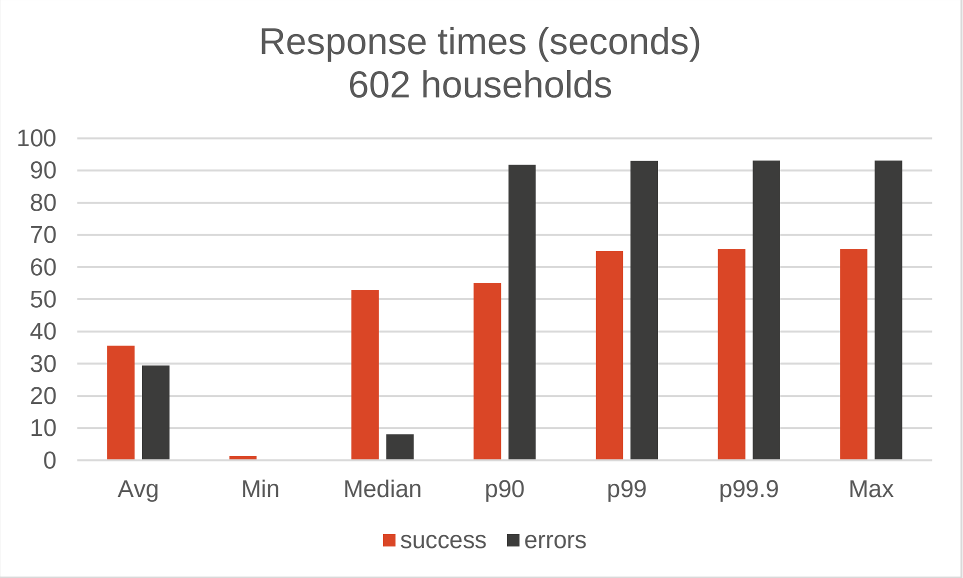

Figure 2 - Response times in test with 602 households

Another important aspect of the control characteristics is the number of successful HEMS responses during activation. For tests including more than 50% of the portfolio, the success rate was consistently between 70 – 80% of the activated portfolio. Individual HEMSs were much more error prone than others, ostensibly due to connectivity or configuration issues, which skewered the results from tests on smaller parts of the portfolio.

The success rates during autonomous operation were higher than those seen during synthetic load tests. During the period of autonomous operation, the aggregators performed more than 100.000 control actions through the FIP, with a success rate of 90%. These control actions led to almost 300.000 interactions with the HEMSs, which had a success rate of 67%. The reason for the greater success rate of FIP control actions was that many HEMS interactions would succeed when retried.

The success rate for activations has much greater impact on the aggregators ability to utilize the flexibility of its portfolio than the response times. With success rates of about 90%, the aggregators will have to account for this in extra uncertainty when calculating the flexibility of their portfolio. As some HEMSs failed much more than others, aggregators may want to exclude such HEMSs from their portfolio.

3 Conclusion

The FIP was able to perform activations within acceptable time frames and with success rates higher than those of the downstream HEMS systems. That said, while it is possible to provide a centralized interoperability platform to include non-compatible DERs in aggregator portfolios, the act of doing so carries the risk of introducing new failure scenarios. This means a DER may be fully functioning, but the aggregator is unable to control it due to issues within the interoperability platform. Principles such as isolation and fault tolerance lowers this risk, but do not remove it. Especially reliance on shared, external systems must be treated with care, to avoid that these become single-point of failures. Ultimately, scalability concerns can be only addressed by building it as a distributed systems, but building distributed systems which act as a single entity is not easily done.

Read more at www.ecogrid.dk