Energy Analytics – Opportunities for Energy Monitoring and Prediction with Smart Meters *

Karoline Ingebrigtsen[1], Volker Hoffmann[2], and Arne J. Berre[2]

[1] SINTEF Energy Research, Sem Sælands Vei 11, 7034 Trondheim, Norway
karoline.ingebrigtsen@sintef.no

[2] SINTEF Digital, Forskningsveien 1, 0373 Oslo, Norway
volker.hoffmann@sintef.no, arne.j.berre@sintef.no

Abstract. By 2019, Norway will complete the national rollout of advanced metering systems (AMS) for all customers. Beyond near-time monitoring of voltage quality and frictionless billing of customers, such a rollout opens a host of possibilities. However, a full-scale rollout is not without challenges. For instance, throughput limitations of radio-mesh networks, privacy considerations, and bounds on compute and storage infrastructure limit the cardinality of metering data to levels below that of which established techniques (for example non-intrusive load disaggregation) require. Pilot projects are now exploring how to mitigate these challenges as well as seeking novel opportunities that open up through data fusion and recent advances in machine learning. In this contribution, we outline the capabilities of the Norwegian AMS system and describe established use-cases and non-intrusive load monitoring. We then discuss a pilot on detection of electric vehicles. Based on preliminary findings, we map the path forward.

Keywords: Smart Meters, AMS, Energy, Analytics, Energytics, Norway

1 Introduction & The Norwegian Smart Meters

The Norwegian government has decided that by 2019, all electricity customers must have smart meters installed. Figure 1 shows an example. There is no requirement on the type and make of meter, but there is a number of regulations which the meters must fulfill [1]. These include the following.

- Report power consumption in intervals of at least 15 and at most 60 minutes.
- Follow a standard interface to allow communication with external equipment.
- Ensuring that no data is lost even during an outage.
- Providing security to avoid data and load control misuse.
- Metering both active and reactive power, both in- and outbound.

Although the core focus is on acquisition of active and reactive power, smart meters can also acquire voltage and current information on one or three phases.

Fig. 1. A smart meter produced by Aidon [4]. The two major players on the Norwegian market are Aidon and Kamstrup. Meters are usually delivered to the grid operators which then contract out installation to local companies.

In addition to enabling more accurate billing of customers, the installation of smart meters will open up new opportunities in energy monitoring and prediction. Examples include enabling of demand response opportunities by detection of high-load appliances, detection of electric vehicles for use as roaming batteries, frictionless accounting for prosumers (customers that feed power back into the grid), and improving safety and reliability by remote fault detection. Predictive maintenance of grid components also offers large potential savings.

However, there are a number of challenges and limitations. The meters communicate measurements via radio-mesh networks of limited bandwidth. Although meters can acquire and store measurements of voltage, current, and power at rates on the order of seconds to minutes, this amount of data cannot be moved out continuously. In practice, measurements are aggregated (summed, averaged, min/max, histograms) over durations of 15 to 60 minutes. Even at 15 minute intervals, data may be too coarse to support the outlined opportunities. There are also privacy considerations, as the data recorded by the smart meters is personal and sensitive.

As the number of electric vehicles (EVs) and other high consumption appliances increases, the capacity of the grid can be strained in periods of peak demand. The Norwegian Water Resources and Energy Directorate (NVE) has produced a report examining the impact of increased numbers of EVs in line with projections (1.5 million EVs by 2030) [2]. NVE concludes that the grid can handle the demand, but that problems can occur if a large number of EVs are charged simultaneously in an area, especially if the grid is already strained. The simple solution is costly infrastructure upgrades, but such upgrades can be deferred if the existing infrastructure can be managed more intelligently through, for example, better understanding of prosumers, identification of roaming storage, or peak shaving through demand response. Here, customers shift their time of energy consumption either manually or through automatic controllers.

The structure of this contribution is as follows. Section 2 introduces the ENERGYTICS ("Energy Analytics") project, Section 3 introduces Non-Intrusive Load Monitoring and its relevance for the smart meters installed in Norway, and Section 4 discusses a pilot demonstrator within ENERGYTICS – the automated detection of charging electric vehicles. Finally, Section 5 concludes the contribution and outlines future work.

2 The Energytics Project

To explore the possibilities opening up by national deployment of smart meters, SINTEF is collaborating with a number of Norwegian grid operators on the ENERGYTICS project. The project serves as a platform to coordinate the exploitation of smart meter data for the following four areas of focus.

1. Operation of AMS and additional services.
2. Real-time analysis of faults and events.
3. Analysing voltage quality and power consumption.
4. Maintenance and reinvestment decisions.

For each area, different demonstrators will be developed. Initially, focus is given to the third area with two demonstrators running that aim to (i) detect charging electric vehicles, and (ii) detect prosumers. At the time of writing, this effort is still ramping up. As such, this contribution is limited to reporting on the first demonstrator, which is the automated detection of electric vehicles from AMS data. This is a form of non-intrusive load monitoring.

Exploratory data analysis is carried out locally using the Python Data Analysis stack (Jupyter, Numpy, Pandas, Scipy, Matplotlib, Scikit-Learn) on data exported from the grid operators' internal systems. As the project matures, data provisioning, storage, and processing will scale to an Azure based cloud solution – both for batch, interactive, and stream processing.

3 Non-Intrusive Load Monitoring

Non-intrusive load monitoring (NILM) is a set of methods to perform load disaggregation based only on power readings at the mains connection. It contrasts intrusive load monitoring, which requires installation of monitoring devices at the circuit or appliance level. Broadly speaking, load disaggregation seeks to identify which appliances are active within a household at any given time [9].

For utilities, NILM is useful to better estimate demand over time and is a prerequisite for demand shifting schemes such as demand response. The enabling of demand response, especially if automated, can lead to savings for customers [6].

Appliances can be disaggregated by identifying characteristic patterns in, for example, their load profiles, spectral envelopes, transient features (e.g., motors spinning up), variations in real and reactive power, or simply the total power drawn. Figure 2 shows an example of the individual load profiles of two appliances as well as their aggregate load profile. If load profiles for individual appliances are available (empirically determined or derived from operating principles), practical load disaggregation is performed algorithmically – either by simple thresholding [8], signal processing [7], or statistical methods [5].

Algorithm performance is limited by the uniqueness of a given appliance signature as well as the sampling rate at which the power signal is acquired. For instance, a hair dryer and a heating iron may exhibit very strong similarities

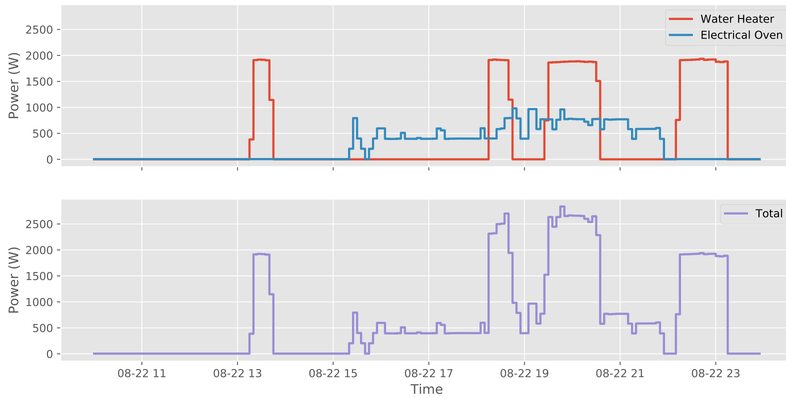

Fig. 2. Illustration of how two power signals stacks. *Top*: Instantaneous power consumption for a water heater (red) and an electrical oven (blue). Data is sampled every minute, and averaged over a five minute sliding window. *Bottom*: Summed instantaneous power consumption of the water heater and electrical oven. The task of load disaggregation algorithms is to infer the decomposition in the upper panel given only measurements corresponding to the bottom panel. *We gratefully acknowledge support from Boye Annfelt Høverstad who supplied the measurements from his private home.*

in usage durations and total power drawn [9]. However, devices such as water heaters and electric ovens differ significantly, cf. Figure 2. Obtaining reliable device signatures in the spectral domain requires sampling rates at the Hz level, while signatures visible in load profiles only require a sufficient number of samples to be drawn during the duty cycle – the duration of which can vary wildly between, for example, a washing machine and a hair dryer. In practice, load profile based disaggregation can be performed on data sampled at the level of a few minutes.

Sampling rates of the AMS infrastructure in Norway are on the order of one hour due to regulatory and technical limitations. Therefore, use of established load disaggregation techniques are unlikely to be immediately useful for loads that vary on timescales shorter than this. Future work must therefore focus on overcoming some of the limitations of the AMS infrastructure, and identification of alternative device signatures from available aggregate AMS data ("feature engineering"). Nevertheless, appliances with sufficiently long duty cycles may be possible to extract from aggregate load profiles. Electric vehicles, for example, have charge cycles on the order of hours, and are explored in the next section.

4 Detection of Electric Vehicles in Energytics

If grid operators can detect location, time, and duration of charging electric vehicles (EVs), they can quantify some of the flexibility in their power consumption and suggest charging to shift to other times of the day. They can also determine the potential for using EVs as an energy storage solution to shift power from off-peak to on-peak periods.

Assuming that the charging signatures of batteries are known, detection of charging EVs does not require complete load disaggregation. Instead, it is sufficient to search the signature of a charging EV within the load profile, which is typically done via a matched filter. Here, a template (the load signature of the charging battery) is slid across a signal (the power consumption time-series) and the cross-correlation calculated along the way. If the cross-correlation peaks and exceeds a given threshold, the template is located. This demonstrator attempts to determine whether this procedure is feasible even if the time-series data is degraded to samples every 10 or 60 minutes. As outlined, the procedure requires a time-series of aggregate load, and load profiles of charging EV batteries.

Since AMS data from the grid operators in ENERGYTICS have not been received yet, a dataset called UK-DALE (UK Domestic Appliance-Level Electricity) is used (with the caveat that these load patterns may not reflect those in Norway) [3]. Among others, the UK-DALE set contains time-series data of instantaneous power from the mains connection sampled at six second intervals. Here, the data is further downsampled by averaging over 10 and 60 minutes.

Lacking empirically determined signatures of charging batteries, a simple trapezoidal model is constructed instead. Here, demand rises linearly from 0 to about 3 kW over 25 minutes, remains stable for about five hours, and then linearly drops to 0 kW over the course of one hour and 25 minutes. At this stage and time-resolution, this is a sufficient approximation of empirically and analytically derived profiles [7].

The charging profile is then embedded in a three-day slice of the UK-DALE time-series, and the normalised cross-correlation computed at both 10 and 60 minutes time-resolution. The procedure is illustrated in Figure 3. Moving from top to bottom, it shows the following. In the top panel, the baseline power consumption from the mains connection of the first house in the UK-DALE dataset is shown with a 10 minute resolution. Underneath, the demand profile resulting from the charging battery described above is shown with its trapezoidal shape. In the third panel, the load profile of the charging battery is embedded in the consumption profile of the mains connection. The fourth panel again shows the resulting load profile, but also shows a version of the demand profile that is downsampled (by averaging) to a 60 minute resolution. Finally, the bottom panel shows the cross correlation of the battery charging template with the load profiles.

The performance of the matched filter is assessed by three criteria:

1. Is normalized cross-correlation largest at the time of charging? If not, the battery template cannot be detected at all.
2. Is it close to unity? Any automated implementation will operate on a threshold value which needs to be selected either analytically or empirically.
3. How unique is the largest peak in the normalized cross-correlation? Similar peak values indicate confusion and suggest that false positives will be found.

Fig. 3. A synthetically embedded and subsequently extracted signature of a charging battery in a three day time-series of power. Power is initially obtained as the average over 10 minute intervals. From top to bottom: (i) Power from the mains connection of a house in the UK-DALE dataset, (ii) a synthetic battery charge signature, (iii) the battery signature embedded in the above time-series, (iv) a representation of the same time-series downsampled to 60 minutes (as well as the original 10 minute data), and (v) the normalized cross-correlation of the template from the second panel with the time-series in the fourth panel.

From the bottom panel of Figure 3, the maximum correlation for 10 minute data (blue) is 0.96 and is at the point where the charging profile was imposed, as it should be. It can be seen that the other peaks in this correlation coefficient are substantially smaller than the maximum one. This implies that the charging profile is to a large degree recognized in the consumption data, and that other appliances are not easily confused with it.

For the 60 minute data (red), the maximum correlation coefficient is at the same point, but has a value of 0.84, 12 per cent below the maximum correlation for the 10 minute data. The other peaks in the correlation coefficient are closer to the maximum value than in the case for 10 minute data. This means that while it is possible to distinguish the EV, there is a higher chance of mistaking a load profile from other appliances to be a charging EV if hourly data is used (compared to 10 minute data). This holds especially true for appliances with a similar load profile as charging EVs, such as electric heaters without thermostats. Here, confusion can arise. The higher the power consumption by other appliances, the harder they are to distinguish from charging EVs. Additionally, when using hourly data, it is likely that only charging cycles with a length of minimum two to three hours can be detected.

Although not investigated at this stage, detecting prosumers will likely turn out to be even more difficult than detecting EVs because their load profiles may not be straightforward to synthesize.

5 Conclusion & Future Directions

Beyond summarizing the framework, opportunities, and limitation of the Norwegian national smart meter rollout, this contribution introduced the ENERGYTICS project with its four focus areas. Special focus was given to a pilot demonstrator for on-line detection of charging electric vehicles.

Direct application of established load disaggregation techniques is hampered by the long sampling intervals of the deployed smart meters, although even hourly sampling seems to support some degree of EV detection. Work will continue by giving more focus to inferring whether charge signatures can be detected in other (lower resolution, or even statistical) representations of AMS data. A similar line of inquiry will be followed for detection of prosumers.

Assuming success, there are many scenarios where on-line detection EVs and prosumers can support predictive use cases. Consider the following examples.

1. Mapping out deployment and charge patterns of EVs yields a description of their power demand as a function of time and location, and their capacity for short-term energy storage. Given this, a predictive model can be constructed.
2. Tracking the relation between energy ingested by prosumers in time and space allows for more accurate forecasting of energy production capacity. This is especially true when combined with auxiliary data sources such as numerical weather prediction.

Although not the focus of this work, promising future demonstrators in the remaining three focus areas of ENERGYTICS include the following three:

1. Automated monitoring of events (e.g. outages) in the electricity grid. Once patterns are identified, predictive models can be constructed.
2. Generating a price signal to the smart meters so that customers can adjust their consumption to the electricity price; either automatically or manually.
3. From historical data, determining whether there exists a relation between weather and outages. Based on this, a predictive model for outages based on weather forecasts can be built.

In conclusion, many of the opportunities presented by the rollout of a national AMS infrastructure are in a nascent stage. Grid operators, regulatory bodies, and research institutions are still exploring this brave new world. Once the opportunities are realized, the advanced metering system will no doubt contribute to more efficient and sustainable grid operations.

References

1. NVE, "FOR-1999-03-11-301: Forskrift om måling, avregning, fakturering av nettjenester og elektrisk energi, nettselskapets nøytralitet mv.", (Online Resource), `https://lovdata.no/dokument/SF/forskrift/1999-03-11-301`, Accessed 07-11-2017
2. NVE, "Hva betyr elbiler for strmnettet?", 2016, `http://publikasjoner.nve.no/rapport/2016/rapport2016_74.pdf`
3. J. Kelly, W. Knottenbelt, 2015, "The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes", Scientific Data, `http://dx.doi.org/10.1038/sdata.2015.7`
4. Midinvest, "Aidon Oy", (Online), `http://www.midinvest.fi/en/sijoitus/aidon-oy-2/`, Accessed 15-11-2017
5. J. Kelly, W. Knottenbelt, 2015, "Neural NILM: Deep Neural Networks Applied to Energy Disaggregation." In Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments (BuildSys '15). ACM, New York, NY, USA, 55-64. `http://dx.doi.org/10.1145/2821650.2821672`
6. H. Sæle, O. S. Grande, "Demand Response From Household Customers: Experiences From a Pilot Study in Norway", in IEEE Transactions on Smart Grid, vol. 2, no. 1, pp. 102-109, March 2011. `http://dx.doi.org/10.1109/TSG.2010.2104165`
7. P. Zhang, C. Zhou, B. G. Stewart, D. M. Hepburn, W. Zhou, J. Yu, "An Improved Non-Intrusive Load Monitoring Method for Recognition of Electric Vehicle Battery Charging Load", In Energy Procedia, Volume 12, 2011, Pages 104-112, ISSN 1876-6102, `https://doi.org/10.1016/j.egypro.2011.10.015`
8. Z Zhang, J. H. Son, Y. Li, M. Trayer, Z. Pi, D. Y. Hwang, J. K. Moon, "Training-free non-intrusive load monitoring of electric vehicle charging with low sampling rate", In Proceedings IECON 2014 - 40th Annual Conference of the IEEE Industrial Electronics Society, Dallas, TX, 2014, pp. 5419-5425, `http://dx.doi.org/10.1109/IECON.2014.7049328`
9. G. W. Hart, "Nonintrusive Appliance Load Monitoring", 1992, In Proceedings of the IEEE, vol. 80, no. 12, pp. 1870-1889