Leveraging Advanced Metering Infrastructure for Distribution Grid Asset Management

Gary Atkinson
Bell Labs
Alcatel-Lucent
Murray Hill, New Jersey, USA
gary.atkinson@alcatel-lucent.com

Marina Thottan
Bell Labs
Alcatel-Lucent
Murray Hill, New Jersey, USA
marina.thottan@alcatel-lucent.com

***Abstract*—** We show how it is possible to develop a system to improve management of distribution grid assets using AMI. In particular, we demonstrate how distribution transformers can be monitored and therefore, managed, even when there are no direct communication channels between the transformers and the asset management system. As an example, using smart meter interval load measurements, load and aging of oil-immersed distribution transformers is estimated with high-temporal resolution, being limited only by the smart meter measurement interval.

Keywords—smart grid; distribution grid asset management; data analytics; oil-immersed transformer loading

I. Introduction

Increasingly, electric power utilities are deploying advanced metering infrastructure (AMI) in their systems with the goal to lower costs, improve operations, and enhance and expand services delivered to customers [1]. These infrastructures can make a large quantity of information available to the operator. It was indicated in [15] how utilities can benefit from employing analytic methods on data available from their system, including where and when to perform preventative maintenance. In this work we focus on an approach for how utilities can leverage their AMI investment to improve management of assets in their distribution grid.

In addition to the equipage at the distribution substations and the physical structures (utility poles, or underground vaults and conduit), there are a variety of other assets in the distribution grid, such as smart meters, distribution transformers, and relays. Assuming one per customer, typical power utilities can have 10^5 to 10^6 smart meters deployed. Therefore, assuming on average no more than 10 smart meters per transformer, distribution transformers can easily number in the tens to hundreds of thousands. By their nature, smart meters are relatively straightforward to monitor and manage. They typically have designed-in statusing and alarming capabilities and can use the AMI infrastructure to alert the meter management system of issues. While protection relays can number in the thousands, they are interconnected via teleprotection and thus can communicate their status. In contrast, there are often no direct communications channels to distribution transformers and typically they are not directly monitored.

Transformers, however, are subject to failures that can be hazardous (explosions) and/or give rise to unexpected power outages which in turn cause associated loss of revenue and customer dissatisfaction as well as increased operations costs. At another extreme, significantly under-loaded transformers can lead to unnecessary no-load losses and per-unit costs. Further, transformer replacement inventory must be maintained, resulting in cost overages if excessive inventory is carried. Owing to the difficulty of monitoring and maintaining transformers, as well as the possibility of an outage impacting multiple customers, we focus on developing an approach to improve the management of distribution transformers assets. Working with a public power utility, we developed a data analytics system that we believe accomplishes that goal.

The paper is organized as follows. After reviewing some key background ideas in Section II, we describe the analytics system methodology in Section III along with implementation and discuss some sample results. In Section IV, we discuss challenges we encountered and solutions for them. Section V concludes with observations and recommendations for utilities.

II. Background Concepts

A. Advanced Metering Infrastructure

An AMI system is a communications system capable of harvesting and analyzing interval data measurements from customer and other sites throughout a utility's service region and also providing communication from the utility to the measurement devices [2]. An AMI consists of the measurement devices (referred to here as *smart meters*), a meter data management system (MDMS), and a communications infrastructure to transport smart meter measurements to the MDMS as well as providing communications from the utility to the meters. A high-level view of the architecture is illustrated in Fig. 1. This architecture is generally applicable to any utility (electric, gas, or water), but our focus here will be on electric power utilities.

Smart meters are installed at customer locations and other load points in the grid and are capable of providing data measurements in intervals on the order of minutes, along with continuously available remote communications. For electric utilities, measurements can include data such as energy usage over the time intervals, instantaneous voltages and reactive power on each phase, power quality, and economic metrics. At the opposite end of the AMI is the utility's data and control center (DCC) where data is collected and stored, and possibly analyzed. These actions are performed by the MDMS, possibly with the assistance of a *head end*. When meter communication protocols are proprietary, the head end provides the interface or gateway between the meters and the MDMS. The head end also performs meter management duties such as monitoring and polling.

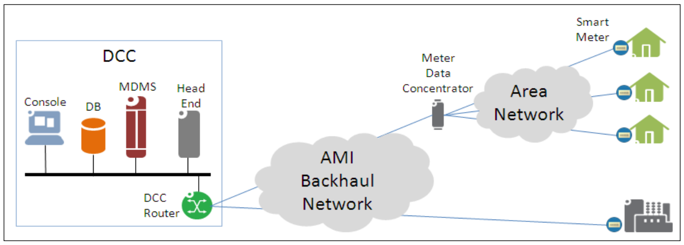

Fig. 1. High-level view of AMI architecture.

Communication between the meters and the DCC is provided by one or more networking technologies, such as copper- or optical fiber-based technologies, wireless (RF) mesh networks, and/or power line communications (PLC). Meters may also be homed to *meter data concentrators* for forwarding their data to the head end and on to the MDMS. It is not uncommon that a combination of networking solutions is deployed. For example, smart meters in neighborhoods may communicate via an RF mesh network which also includes meter data concentrators located in the local distribution substations. The data concentrators in turn communicate with the head end via a fiber-optic network.

B. *Distribution Grid Topology*

The majority of electrical power is generated in bulk in a 3-phase AC format at various types of power plants (e.g., fossil fuel, nuclear, hydroelectric) and then stepped up to high voltage (and low current) for long-distance conveyance over transmission lines to distribution systems for delivery to end users. From the transmission system, the power enters the distribution system at *distribution substations*, where the high voltage for transmission is stepped down via substation transformers to a level suitable for distribution to customers – see Fig. 2. From the substations, power is transported to customer premises via above- or below-ground distribution lines, also referred to as *feeders*. When the feeders reach the vicinity of a service location, the distribution line voltage is stepped down a final time via *distribution transformers*. Electrical service is then provided from the customer-facing or secondary side of the transformer to the (smart) *meter* monitoring the *customer premises*. For residential customers in the U.S., the service is typically one or two phases at 120 or 240 V. Commercial and industrial customers can receive all three phases of service at higher voltage levels.

Some useful observations about the distribution system are:

- Tree Topology – at any one time, the topology of the distribution system is a tree (i.e., it is *radial*). It turns out that feeders may be deployed in non-tree configurations (e.g., they can loop back to their substation or connect to another substation), but in such cases, there are also circuit breakers, reclosers, and/or relays deployed in which some are open to yield a tree.
- Type of distribution transformer assets – there can be a wide range of transformers according to nameplate ratings, cooling mechanisms, and winding insulation. For reference, commonly occurring examples are 25 or 37.5 kVA ONAN (oil-filled unit with natural convection flow in the tank and natural air convection cooling externally), which are often the familiar grey cylindrical canisters seen on utility poles [3].

For brevity, hereafter smart meters will be referred to as meters and distribution transformers as transformers.

III. INDIRECT TRANSFORMER MONITORING

Transformers in the distribution portion of electrical grids age with use due to electrical loading. Transformers regularly subjected to loading above their nameplate rating are more likely to fail sooner than those that are lightly loaded. An excessively aged transformer can fail catastrophically, resulting in unplanned outages to customers, often during periods of intense electrical usage. Grid operators prefer to replace aged transformers during planned outages (e.g., during periods of low usage) prior to catastrophic failure. However, operators do not want to replace transformers before they have reached their useful lifetime because premature replacement results in unnecessary asset costs and wasted effort.

These issues can be mitigated if utilities had a clear view of loading on their transformers. For that, quality day-to-day loading information per transformer would be needed. Obtaining such information directly can be difficult, often involving costly *in situ* measurements by field personnel, or requiring more-expensive transformers equipped with gas or thermal sensors and the associated telemetry infrastructure.

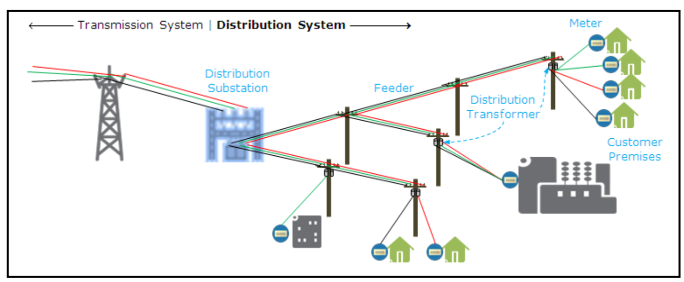

Fig. 2. Representative distribution system tree architecture.

However, now with the deployment of AMI infrastructure, meter measurements can be utilized to develop temporally high-resolution views of transformer loading *indirectly* to determine each operational transformer's aging acceleration. Having this information available will lead to multiple benefits, including enhancing the utility's ability to identify and prioritize candidate transformers for replacement prior to failure, minimizing service disruption time, reducing transformer stocks and associated inventory costs, and maintaining a more modern inventory.

A. *Key Assumptions*

To estimate a transformer's loading indirectly from measurements provided from a utility's AMI, it is assumed that each premises the transformer is providing service to is equipped with a (smart) meter. Also, it is assumed two types of dynamic high-frequency data are available: (i) meteorological data collected periodically (hourly or better) from sensors distributed throughout the operator's geographic footprint, and (ii) interval energy usage measurements from each premises equipped with a meter in small time intervals (5 to 60 minute intervals, with 15 minutes being typical).

Finally, it is assumed that reasonably reliable aging and thermal models are available for the transformers. For the purpose of illustrating how transformer load can be indirectly estimated from AMI data, we focus on the case of oil-immersed distribution transformers. For such transformers, there is a standards-based computational methodology based on accepted practices using models that apply broadly to distribution transformers of varying designs. Other models could be used and indeed many have been proposed, but they can be dependent on particular designs and it is not *a priori* clear how broadly they apply (see e.g., [4], [5], [6]). The standards-based models have known behavior, are computationally light, and can be applied broadly to the thousands of transformers in an operator's system.

B. Transformer Aging

A primary reason distribution transformers age and fail under operation is the degradation of winding insulation due to heating (see, e.g., [7]). For oil-immersed transformers, one of the most common types of transformers in the field, IEEE Standard C57.91-2011 [8] and IEC Standard 60076-7-2005 [10] give detailed procedures and formulae for calculating oil and winding temperatures in order to predict the thermal response of a transformer under load and, thus, the associated aging. Formulae are provided for both steady state and transient loading.

A transformer designed with thermally upgraded paper insulation and experiencing a load in a time interval $[t', t'']$ will undergo thermal aging at a rate ([8], [9], [10], [11])

$$A(t) = \text{aging rate} = A_0 \exp[-\alpha / (\theta_h(t)+273)] \qquad (1)$$

at a time $t\in[t', t'']$ where

- $A_0 \equiv \exp[\alpha/(T_0+273)]$
- $\alpha \equiv$ aging rate scaling parameter (K)
- $T_0 \equiv$ operational temperature target (°C)
- $\theta_h(t) \equiv$ transformer hottest spot temperature (°C).

Then the loss of transformer lifetime over the time interval is

$$L = \text{loss of lifetime} = \int_{t'}^{t''} A(t)\, dt \; . \qquad (2)$$

C. Transformer Thermal Modeling

Now (1) depends on knowing the transformer's hottest spot temperature $\theta_h(t)$. It can be related to the more easily measured top oil temperature $\theta_{top}(t)$, i.e., the temperature at or near the top of the transformer windings at the top of the oil tank.

$$\theta_h(t) = \theta_{top}(t) + \Delta\theta_{h,w}(t) - \Delta\theta_{h,o}(t) \qquad (3)$$

where $\Delta\theta_{h,w}(t)$ and $\Delta\theta_{h,o}(t)$ are the winding hottest spot-to-top oil temperature gradient over the top-oil temperature due to the winding response and oil viscosity, respectively. The top-oil temperature is based on the ANSI/IEEE Clause 7 top-oil rise model corrected for ambient temperature variation and in differential form ([8], [10]):

$$\tau_o\, d\theta_{top}/dt = -\theta_{top}(t) + \theta_{amb}(t) + \theta_{fl}\,[\,(I(t)^2R + 1) / (R+1)\,]^n \quad (4)$$

where

- $\theta_{top} \equiv$ top-oil temperature (°C)
- $\theta_{amb} \equiv$ ambient temperature (°C)
- $\theta_{fl} \equiv$ top-oil rise over ambient temp. at rated load (K)
- $\tau_o \equiv$ oil time constant (min)
- $I \equiv$ ratio of load to rated load
- $R \equiv$ ratio of load loss to no-load loss at rated load
- $n \equiv$ oil exponent, selected for each cooling mode (e.g., ONAN, ONAF, ODAF/WF, OFAF/WF) to approximate change in resistance with load

The hottest spot-to-top oil temperature gradients are each governed by an equation of the form ([8], [9], [10])

$$\tau_x\, d\Delta\theta_{h,x}/dt = a_x\,[\,b_x\,\Delta\theta_{hr}\, I(t)^{2m} - \Delta\theta_{h,x}(t)\,] \qquad (5)$$

where the "x" subscript is "w" for winding response or "o" for oil viscosity response and

- $\Delta\theta_{h,x} \equiv$ hottest spot to top oil temperature gradient over the top-oil temperature due to x (°C)
- $\Delta\theta_{hr} \equiv$ hottest-spot-to-top-oil gradient at rated current (°C)
- $\tau_x \equiv$ time constant for x (min)
- $a_x, b_x \equiv$ thermal constants for x
- $m \equiv$ load exponent for current versus winding temperature rise in steady state.

Given dynamic ambient temperature and loading inputs $\theta_{amb}(t)$ and $I(t)$, respectively, (4) and (5) are solved and combined with (3) to obtain the hottest spot temperature $\theta_h(t)$ which, in turn, is an input for the aging model in (1) and (2).

D. Implementation

Fig. 3 exhibits a high-level view of the Distribution Transformer Loading Analytics (DTLA) module flow as it processes transformers. The DTLA is designed to compute the loading and aging metrics for each in-service distribution transformer for a user-specified time period. The intention at this stage of development is that the transformer aging would be computed on a periodic basis, where any overloading or excessive aging could be flagged. Based on user inputs, the DTLA gathers information on the operational system elements (distribution transformers, meters, substations, connectivity) and the meteorological and energy usage time series data for the specified time period. For each transformer, the DTLA computes the aging a transformer's winding insulation undergoes over the period from the load variations sampled by its meters and from the varying ambient thermal environment it experiences over the period. The transformer's aging results for the period as well as other statistics are then stored in a database and are available for subsequent retrieval and analysis. This is repeated until all detected in-service transformers have been processed for the period.

The dynamic data that is provided to the algorithm is the ambient temperature and energy usage measurements (in W-hrs) over discrete time intervals by all meters. The ambient temperature can be used directly, but the energy usage must be converted to load measured in kW. A critical input for each transformer is which meters are measuring service usage provided by a transformer. Then for a set of meters $\{m\}$ receiving service from a transformer, each with a measured usage $\Delta E_m(t)$ over a time interval Δt, the average load in kW over the interval experienced by the transformer is given by

$$\text{Load}(t) = \textstyle\sum_{m\in\text{Transformer}} \Delta E_m(t) \,/(1000\, \Delta t\, pf)\,. \qquad (6)$$

Here, pf is the power factor for converting the real power to apparent power, since the apparent power is responsible for the heating in the transformer windings. When instantaneous real and reactive power measurements are provided by the meter, pf can be found from the real-to-apparent power ratio. Otherwise, an estimate of the power factor is made based on the type of user (residential, commercial, or industrial).

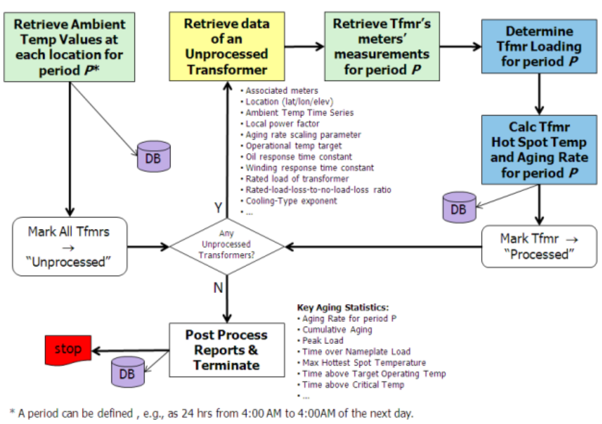

Fig. 3. High-level view of the DTLA flow.

Once the load on a transformer is known for the period, the system of equations (3) – (5) are readily solved using discrete time solution methods, in this case the Runge-Kutta family of algorithms, which readily incorporate time-dependent parameter functions (as well as non-linearities in the dependent variables). For integration of (1), we found the trapezoidal rule was sufficient and, together with fourth-order Runge-Kutta for the thermal model solves, provided less than a 0.05% combined relative error with 30-second time steps.

E. Sample Results

Fig. 4 illustrates a sample of the results available from the DTLA model based on actual data for an operator's 25 kVA-rated transformer in service during the 24-period between 4:00 AM July 1, 2012 and 4:00 AM July 2, 2012. Unknown transformer parameter values from [12] were used for this and the next example below. The estimated ambient temperature at the transformer location (blue) and the computed transformer hottest spot temperature (red) are measured against the left axis, while the load (green) and accumulated loss of lifetime (brown) series are normalized and measured against the right axis. Note that the hottest spot temperature rises above the transformer's target hottest spot temperature target of 110 °C (dotted line) for over 10 hours. Further, for about a half of an hour, it rises above 140 °C, a significant point for oil-immersed transformers because the oil can break down above that temperature and form gas bubbles, leading to numerous problems. The average load over the course of the 24-hour period is only 25.8 kVA which is barely above the rated load, but the bulk of the aging occurs when the load is significantly above the rated load (load normalized to 1 and referred to the right axis), between hours 10 and 20. As a result, the transformer experienced 2.8 times the aging (final value of LifeLoss) it normally would have under a rated load for the 24-hour duration. An aging calculation using [8] based on average load and ambient temperature for the same period gives the aging as 1.9 times the aging at rated load for 24 hours, and thus would have understated the amount of aging the transformer underwent by 50%, a large discrepancy. An operator might not have been aware of this without the detailed calculation of the aging based on the temporal resolution provided by high frequency data. This indicates that using, e.g., estimated transformer loading based on a daily average computed from monthly or quarterly meter readings can underestimate significantly the aging a transformer actually undergoes.

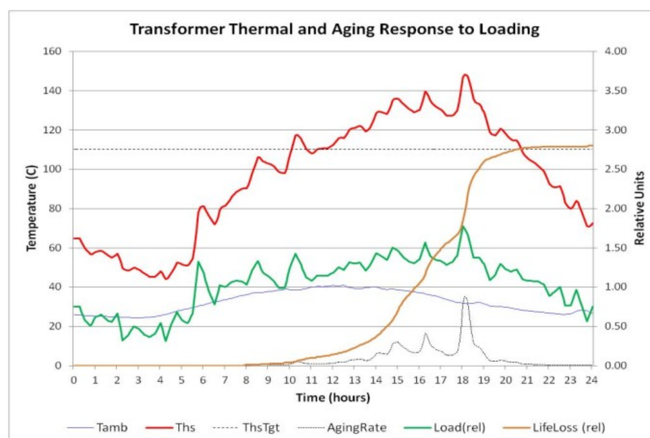

Fig. 4. Sample results for an in-service 25 kVA transformer.

Another interesting effect is shown in Fig. 5. There are two load and hottest spot temperature profiles shown for the same in-service 37.5 kVA transformer. The upper graph is for the 24-period between 4:00 AM July 1, 2012 and 4:00 AM July 2, 2012 ("summer") while the lower graph is for the 24-period between 4:00 AM February 1, 2012 and 4:00 AM February 2, 2012 ("winter"). The summer graph shows a prolonged period of loading above nameplate rating resulting in a hottest spot temperature above the target operating temperature of 110 °C (dotted line), and even rising briefly above the critical 140 °C level. The resulting aging for this transformer for this period is 1.25 times the nominal aging. By contrast, the winter graph shows many loading spikes above the nameplate rating, and resulting hottest spot temperature spikes above the operating temperature and even approaching the 140 °C mark, but the overall aging experienced by the transformer on this date was less than 0.4 of nominal aging. So, spikes don't result in aging as much as prolonged overages do. One can also see that it is possible for the transformer to exceed 140 °C on days that it is lightly loaded. Thus, catastrophic failures (explosions) could potentially occur when least expected. Thus, the high resolution enables possible flagging of transformers that might unexpectedly be at risk.

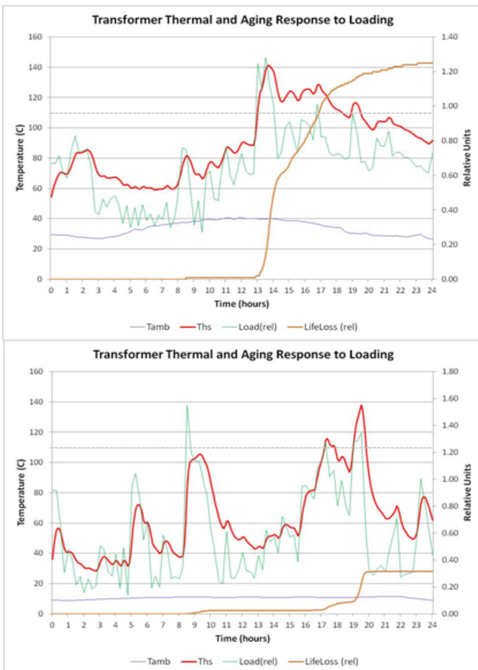

Fig. 5. Sample results for an in-service 37.5 kVA transformer. *upper*: summer, *lower*: winter.

IV. Challenges

During the course of the project, we encountered interesting challenges, two of which we mention here.

A. Topology

One significant challenge encountered early on was topology awareness by a utility. Utilities maintain the type of connectivity information about their system that supports their current operations. Unfortunately, this may not include the kind of information that would enable smart grid applications. This includes the application of interest herein, namely, load computation at a transformer based on usage measurements from smart-metered premises. For example, an operator may know what structure a transformer is mounted on and which structure a premises is receiving its service drop from. However, if multiple transformers share the same structure, what might not be known is precisely which transformer on the structure is providing service to particular premises. While before smart grid, this kind of information was not necessarily needed for utility operations, it is essential for estimating transformer load based on the meter energy measurements of the premises it provides service to. The authors in [13] developed an approach to reconstruct end-to-end connectivity in the distribution grid, but it assumes smart meters deployed at all topological nodes of the network, including at the transformers and feeder heads.

Due to space limitations, we summarize our approach to precision topology awareness. Upon studying the various databases maintained by the utility, certain patterns began to emerge. For example, one datum for the transformers that was tracked is what the phase connections were on the primary side. Further, in the customer database, which phases serviced the customer premises (corresponding to the transformer secondary side) was also tracked. So, if a customer premises receives 3-phase service (phases A, B, and C) from a structure (utility pole) with three single-phase transformers on it, where each transformer was connected line-to-line (A-B, B-C, and C-A), then the premises could be receiving its service from a set of single phase transformers connected in a delta configuration on the primary side. In addition to the phase connection patterns (Wye, Delta, open Delta, …) on the primary and secondary side, we were able to use other information, such as transformer age, status, and structure identifier to develop rules for pattern matching and data mining on a mash-up of multiple data sets. There were less than a hundred rules in total. The resulting topology discovery tool enabled a mapping between transformers and meters to be established based on operations data. Using this approach, we were able to map over 99.9%, all but a few, of the meters to transformers.

B. Centralized Architecture

A challenge that needs to be borne in mind of analytics platforms like the DTLA that rely on data harvested and stored in a centralized database is scaling. When performing the analytics, such platforms must establish the connection with the database server, communicate the queries, the queries must be performed by the DB, and then the results must be communicated back to the analytics platform. The actual time for the computational algorithms in the DTLA is very small compared to the database interactions. Thus, as the number of (transformer) assets in the system grows, at some point the number will exceed the system's capacity to provide results in a timely manner. This problem is compounded when there are other analytics applications utilizing the database concurrently.

One approach to improving the scaling performance is to parallelize the analytics platform architecture and distribute the database. It turns out that the transformer loading analytics is particularly amenable to this approach. The analytics machinery for transformer loading relies primarily on local information: local topology (meter-transformer connectivity), local meter measurements, local ambient temperature measurements. Furthermore, because the distribution grid has a tree topology, transformer and meter assets of different substations are essentially decoupled from each other. Therefore, instead of processing all the transformers serially through one platform instance on a central server that accesses a centralized database, the problem could be decomposed into subsets of transformer and meter assets, with each subproblem having a platform instance and database resources dedicated to those assets, including harvesting and maintenance of the meter data. Any coupled aspects of the problem could be centralized with presumably a much reduced computational load and demand on any central database. This should enable time and flexibility to coordinate centralized resource usage.

Employing distributed data management systems such as Hadoop [14], a possible distributed architecture for the DTLA is shown in Fig. 6. Each distribution substation hosts a node in the distributed architecture. Present at each node would be a meter data concentrator for harvesting and making available locally those measurements from all meters served by the substation's feeders. A local analytics platform instance would perform the loading analytics using only the local meter measurements and ambient temperatures from a weather station collocated at the substation. A node manager would perform local resource management and respond to other jobs (e.g., backup and data redundancy) from the main resource/application manager in the DCC. In this way the DTLA will be scalable since it will only need to process a substation's amount of transformers.

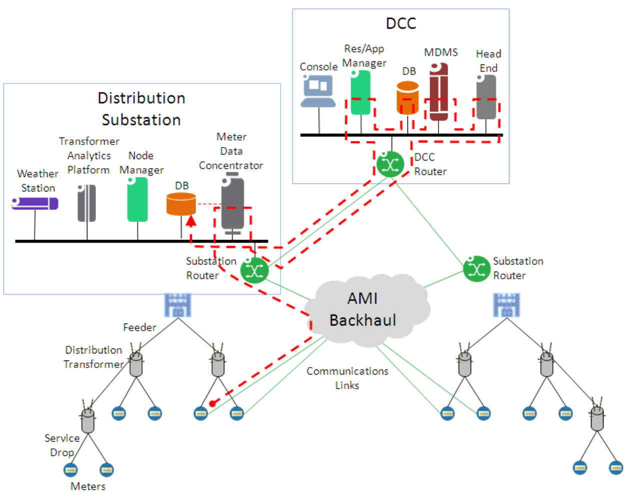

Fig. 6. Architecture of a distributed transformer management system. Local meter data flow shown (broken line) when a lone head end is located in DCC.

One design consideration is where to locate head ends. When meters are deployed that use proprietary protocols for communicating measurements, a head end is needed in the AMI system to make measurements available to the analytics and distributed database systems. Locating the head end in the DCC means that meter data will need to be sent to the DCC and then redistributed back down to the substation database (broken line in Fig. 6), thereby incurring bandwidth utilization and latency costs. In the current DTLA mode of running once daily, those costs are not significant, but in a real time application, an approach would be to deploy instances of the head end at each substation if tradeoff of the licensing costs against the bandwidth and latency costs justify it.

V. SUMMARY & REMARKS

We have shown how a utility can employ its AMI smart meter measurements to develop temporally high-resolution views of transformer loading and aging acceleration, even in the absence of direct communications channels with the transformers. A grid operator having such a detailed view of the loading and aging of its distribution transformer assets could reap multiple benefits, such as: (1) enhancing the utility's ability to identify and prioritize candidate transformers for replacement prior to failure; (2) minimizing service disruption time by planning needed outages instead of reacting when they occur; (3) reducing transformer stocks and associated inventory carrying costs; (4) maintaining a more modern transformer inventory; and (5) sharpening transformer asset management decisions such as when to repair, replace, or retire units, and replacement sizing, including less costly reduced-sized units. While the presented approach is based on oil-immersed standards-based transformer thermal and aging methodology, (reliable) models for other transformers types could be readily incorporated in the overall framework.

Finally, we also encountered challenges that suggest power grid operators should consider the following recommendations to help fully leverage their AMI investment: (1) update their grid topology database to provide complete end-to-end connectivity and details; (2) establish audit trail mechanisms (e.g., via bar- or QR-coding) for distribution transformers so that their loading history, and therefore aging, can be accurately estimated despite moving around the system; and (3) use market leverage to request transformer thermal models from manufacturers to improve accuracy of aging estimates.

REFERENCES

[1] "Assessment of Demand Response and Advanced Metering," Federal Energy Regulatory Commission, U.S. Dept. of Energy, December 2008.

[2] "Advanced Metering Infrastructure," Report v1.0, National Energy Technology Laboratory, U.S. Dept of Energy, February 2008.

[3] *Electric Power Engineering Handbook: Electric Power Transformer Engineering,* 2nd Ed., J. Harlow, ed., CRC Press, 2007.

[4] D. Susa, M. Lehtonen, H. Nordman,"Dynamic thermal modeling of distribution transformers," *IEEE Trans.On Power Delivery*, **20**(3), pp. 1919- 1929, July 2005.

[5] A. Mamizadeh, I. Iskender, "Analyzing and Comparing the Hot-spot Thermal Models of HV/LV Prefabricated and Outdoor Oil-Immersed Power Transformers," World Academy of Science, Engineering and Technology, **61**, pp. 34-39, 2012.

[6] L. Rivera, D. Tylavsky, "Acceptability of Four Transformer Top-Oil Thermal Models: Pt. 2: Comparing Metrics," *IEEE Transactions on Power Delivery*, **23**(2), pp. 866 – 872, 2008.

[7] M. Popescu, N. Mastorakis, L. Popescu-Perescu, "Modelling of Oil-filled Transformer," International Journal Of Mathematical Models and Methods In Applied Sciences, vol 3 (4), pp. 346-355, 2009.

[8] IEEE Std C57.91-2011, IEEE Guide for Loading Mineral-Oil-Immersed Transformers.

[9] "Permissible Loading of Oil-immersed Transformers and Regulators," U.S. Dept of the Interior, Bureau of Reclamation, Facilities Engineering Branch Denver Office, Denver, CO, Internet Version, 2000.

[10] IEC International Standard IEC 60076-7-2005, Power transformers – Part 7: Loading guide for oil-immersed power transformers, 2005.

[11] A. van Schijndel, J. Wetzer, and P. Wouters, "Forecasting Transformer Reliability," Proceedings, 2006 Annual Report, Conference on Electrical Insulation and Dielectric Phenomena, 2006.

[12] M. Kuss, A. Markel, and W. Kramer, "Application of Distribution Transformer Thermal Life Models to Electrified Vehicle Charging Loads Using Monte-Carlo Method," Proceedings, *25th World Battery, Hybrid and Fuel Cell Electric V,ehicle Symposium & Exhibition*, Shenzhen, China, 2011.

[13] V. Arya, T. Jayram, S. Pal, S. Kalyanaraman, "Inferring connectivity model from meter measurements in distribution networks," Proc. *Fourth Int'l Conference on Future Energy Systems*, pp. 173-182, 2013.

[14] http://hadoop.apache.org

[15] C. Rudin, et al, "Machine learning for the New York City power grid," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **34** (2), pp. 328–345, 20.