Robert Lorenz, Lars Dannecker, Philipp Rösch, Wolfgang Lehner, Gregor Hackenbroich, Benjamin Schlegel

Forecasting in hierarchical environments

Erstveröffentlichung in / First published in:

SSDBM: Proceedings of the 25th International Conference on Scientific and Statistical Database Management, Baltimore 29.-31.07.2013. ACM Digital Library, Art. Nr. 37. ISBN 978-1-4503-1921-8

DOI: https://doi.org/10.1145/2484838.2484849

Final edited form was published in "SSDBM '13: Conference on Scientific and Statistical Database Management. Baltimore 2013", Art. Nr. 37, ISBN 978-1-4503-1921-8
https://doi.org/10.1145/2484838.2484849

Forecasting in Hierarchical Environments

Robert Lorenz
SAP AG
robert.lorenz01@sap.com

Lars Dannecker
SAP AG
lars.dannecker@sap.com

Philipp Rösch
SAP AG
philipp.roesch@sap.com

Wolfgang Lehner
TU Dresden
Database Technology Group
wolfgang.lehner@tu-dresden.de

Gregor Hackenbroich
SAP AG
gregor.hackenbroich@sap.com

Benjamin Schlegel
TU Dresden
Database Technology Group
benjamin.schlegel@tu-dresden.de

ABSTRACT

Forecasting is an important data analysis technique and serves as the basis for business planning in many application areas such as energy, sales and traffic management. The currently employed statistical models already provide very accurate predictions, but the forecasting calculation process is very time consuming. This is especially true since many application domains deal with hierarchically organized data. Forecasting in these environments is especially challenging due to ensuring forecasting consistency between hierarchy levels, which leads to an increased data processing and communication effort. For this purpose, we introduce our novel hierarchical forecasting approach, where we propose to push forecast models to the entities on the lowest hierarch level and reuse these models to efficiently create forecast models on higher hierarchical levels. With that we avoid the time-consuming parameter estimation process and allow an almost instant calculation of forecasts.

Keywords

Forecasting, Hierarchy, Optimization

1. INTRODUCTION

Time series forecasting is the basis for manual and automatic planning in many application domains such as sales, traffic control and energy management [8]. The forecasting calculation is based on mathematical models called *forecast models* that use a parameterized relationship between past and future values to describe the behavior of a time series. To allow for accurate forecasts, the parameters are adapted to the specifics of the time series, by estimating them on a training data set using local and global numerical optimization algorithms with the goal of minimizing the forecast error measured in terms of some error metric [9]. This estimation is typically very time consuming, since it requires a large number of simulations. The reason is that the parameter search space increases exponentially with the number of forecast model parameters. In addition, many application domains exhibit a hierarchical data organization, where time series are aggregated along the hierarchy using some dimensional attributes [10, 4]. Forecasting in hierarchical environments is very complex, since it involves multiple entities and forecast models on different hierarchical levels. Thus, it is necessary to ensure forecasting consistency among all entities and across all levels of the hierarchy. This is especially an issue in application domains that require the availability of accurate forecasts at any point in time. One example is the energy domain, where accurate forecasts are the basis for the balancing of energy consumption and production on the electricity grid. There, with an increasing share of renewable energy sources, new requirements are posed on this balancing task. The reason is that renewables highly depend on exogenous influences (e.g., the weather), which means that their energy production cannot be planned like traditional energy sources. In addition, there is only insufficient storage capacity and thus, the energy production of renewables must be used when available. Research projects such as MIRABEL [12] and MeRegio [11] develop advanced technologies such as demand response systems and flexible energy requests [1, 2] to address the issue of real-time energy balancing and an increased utilization of RES. Most of these approaches have in common that they require an efficient and fast calculation of accurate predictions.

DOI: https://doi.org/10.1145/2484838.2484849

To allow constantly available accurate forecasts, we introduce our novel hierarchical forecasting approach that exploits the structure of hierarchically organized time series to increase the forecasting efficiency. In the following, we describe our approach using the energy domain as an example. However, our approach can be easily adapted to other application domains that deal with hierarchical time series. The European energy market is hierarchically organized with energy consumers and producers forming the lowest level of the hierarchy. They are are pooled into several balance groups, with the goal of balancing the energy consumption and production within these groups. Balance groups are controlled by balance responsible parties that form the second level of the hierarchy. The third level comprises transmission system operators (TSO) responsible for the operation of a certain segment of the transmission grid. Recently with an increasing deployment of smart grid technology, more and more entities on the lowest level are equipped with smart meters used to continuously monitor the amount of consumed or produced energy. The basic idea of our approach is to

Final edited form was published in "SSDBM '13: Conference on Scientific and Statistical Database Management. Baltimore 2013", Art. Nr. 37, ISBN 978-1-4503-1921-8
https://doi.org/10.1145/2484838.2484849

utilize the data processing capabilities provided by smart meters and to push the responsibility of creating and maintaining forecast models directly to the customer (consumers and producers). Thus, next to their task of recording current measurements, smart meters gain the capability to forecast future consumption and production using the provided forecast models. Each customer forecast model represents the individual energy consumption and production behavior of a single entity on the lowest hierarchical level. Entities on higher levels that are responsible for balancing energy demand and supply, can utilize the single customer forecast models to create a global forecast model representing the overall consumption or production behavior of all connected lower level entities. The used forecast model aggregation is much faster than the estimation or re-estimation of a global forecast model based on an aggregated time series. Thus, our approach allows balancing companies to almost instantly calculate very accurate forecasts as required for balancing energy demand and supply in real-time. Furthermore, customers can also benefit from the added forecasting capabilities, since they additionally enable advanced smart building and smart home systems.

2. HIERARCHICAL FORECASTING

Forecasting in a hierarchical system such as the European energy market, is currently organized in a way that higher level entities calculate forecasts based on the aggregated data of lower level entities. In this context, a balance responsible party (BRP) on the second hierarchy level, forecasts the energy consumption and production for its group of customers, using the data provided by the corresponding entities. Thus, after aggregating the data of all entities the BRP estimates a forecast model fitted to the aggregated data, which describes the energy consumption and production behavior of the entire balance group. Furthermore, typically the BRP maintains forecast models for large consumers and producers as well as private households in a certain granularity.

In the energy domain new measurements are constantly available, which means that entities on the lowest level continuously or in fixed intervals communicate changes to the responsible entity on the next hierarchy level. Since new measurements might change the time series behavior, it is necessary to adapt the forecast models the BRP is responsible for to the most recent time series updates. Besides high communication efforts, a forecast model adaptation in most cases means a re-estimation of the forecast model parameters, which renders the adaptation process almost as time consuming as the initial forecast model creation. Thus, the maintenance of all forecast models leads to large efforts for the BRP. Furthermore, the BRP might face the issue of not being able to fulfill possible execution deadlines posed by real-time balancing. To solve this issue, we propose a hierarchical forecasting approach that utilizes the hierarchical structure of the european energy market to substantially increase the forecasting calculation efficiency.

2.1 Forecast Model Aggregation

The core idea of our approach is a decentralization of the forecasting process. In particular, we push the responsibility to build and maintain a forecast model to the entities on the lowest level. A forecast model can be seen as a representation of a customer consumption or production time series, meaning that in an ideal case the values provided by the forecast model are identical to the actual measurements. While in the real world forecast models always describe an average behavior over a certain time frame and thus, in most cases exhibit a certain deviation between the model values and the real values, the forecast error is typically decreasing with an increasing aggregation level across the hierarchy. The reason is that in a group of entities fluctuations of a single entity are for the most part neutralized by other entities in the group. Following these facts, instead of providing and aggregating time series values we transmit the individual forecast models M_K of single customers and aggregate them to directly form a global forecast model M. No estimation of the global forecast model is required. Similar to the aggregated time series the global forecast model approximately describes the energy consumption and production of the entire group and can be used to predict future values. Figure 1 illustrates our hierarchical forecasting approach compared to the conventional model creation based on time series aggregation. For the conventional way on the left-hand side we aggregate the time series and subsequently use optimization algorithms to create the global forecast models. In contrast, using our approach on the right-hand side, we directly create the global forecast model based on the single customer models.

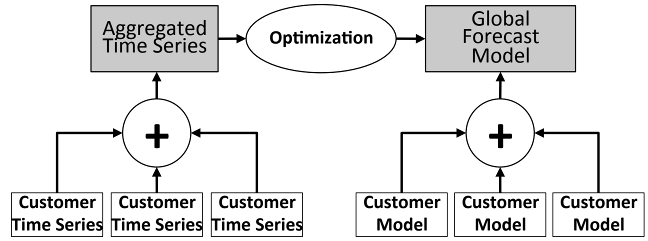

Figure 1: Estimation vs. Aggregation Approach

The aggregation of forecast models works as follows: A forecast model M consists of multiple forecast model components C, each comprising a time series coefficient x (i.e., the representative of the time series) and a corresponding parameter p. The linear model $a \cdot x + b$ for example comprise of the two forecast model components $a \cdot x$ and b, where a and b are the parameters and x and 1 (omitted) are the respective coefficients. The single customer forecast models are aggregated by creating a weighted linear combination for each of the forecast model components $\omega_k \cdot C_k + \omega_l \cdot C_l + \ldots$. Accordingly, to create a component for a specific global forecast model we involve all connected lower level entities into the linear combination. Since each entity consumes or produces a different amount of energy, they should likewise have a different influence on the global forecast model M. Therefore, we use the share of an entity on the global consumption (or production) ω as its weight. The share of an entity is estimated based on its average energy consumption in the recent past (e.g., in the last week) and the consumption of the group in the same time frame. Calculating a weighted linear combination is far less time-consuming than re-estimating the forecast model parameters using numerical optimization algorithms. Thus, we provide a very efficient way for entities on higher hierarchical level to calculate accurate forecasts.

Final edited form was published in "SSDBM '13: Conference on Scientific and Statistical Database Management. Baltimore 2013", Art. Nr. 37, ISBN 978-1-4503-1921-8
https://doi.org/10.1145/2484838.2484849

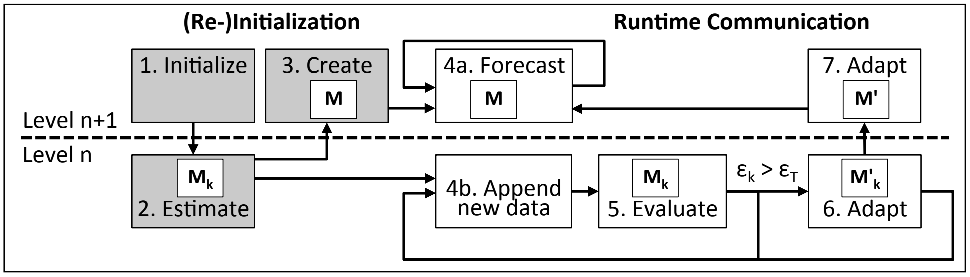

Figure 2: The communication protocol of the hierarchical forecasting system

A pre-requisite for our approach to work in the energy domain is that higher level entities have the possibility to determine the current global consumption or production for their group of lower level entities. Since smart grid technology is getting more and more deployed, transmission system operators may provide energy consumption and production information in different granularities. In addition, previously there was only a limited distribution of smart meters, which means that metering devices on the granularity level of balance groups were used to measure the global consumption and production data necessary for managing balance groups. Thus, it is safe to assume that higher level entities have the respective global measurement data available.

2.2 Hierarchical Communication

In addition to the forecast model aggregation, our hierarchical forecasting approach enables a more dynamic communication between the entities on different hierarchical levels. Instead of transmitting values in a fixed interval as typically done in the conventional approach, a lower level entity initiates an information exchange depending on the accuracy of its forecast model; as long as the base model is up-to-date there is no need to exchange information between the hierarchy levels. In general a forecast model transmission in the context of our communication protocol is a vector containing: the previous and current parameter combination of the forecast model, the average consumption/production, and some statistical information (e.g., date of last forecast model estimation, estimation error). Instead of just submitting the new parameter combination, we also added the old parameter combination to determine the direction vector of the change and the significance of the adaptation. The transmitted average consumption is used to determine the share of a single entity (compare Section 2.1).

The protocol is illustrated in Figure 2 and divides into two parts the initialization and the runtime part. The initialization is conducted when our hierarchical forecasting is used for the first time or when large organizational changes occurred. (1) It starts with requesting the individual forecast models from all lower level entities. (2) Each entity creates an up-to-date forecast model and transmits the model to the higher level entity. (3) There, all individual forecast models are aggregated to form the initial global forecast model. (4a) During runtime the higher level entity uses the initial global forecast model to calculate forecasts. (4b) At the same time, the entities on the lowest level continuously append new measurements to their time series and (5) use these information to evaluate the accuracy of the their forecast models. (6) As soon as the forecast error violates a given accuracy threshold, the respective entity adapts the forecast model by re-estimating its parameters and transmitting the changes to the higher level entity. (7) The higher level entity picks up the changes and adapts its forecast model accordingly, which is afterwards again used to calculate forecasts. During runtime, only entities that adapted their forecast model are transmitting information. Our communication protocol also works between multiple hierarchy levels in a cascading way. This means that entities on the lowest level communicate only with entities on the second level, which likewise inform entities on the third level and so on. With that we ensure that each level receives the provided information in the necessary granularity.

Thus, our communication protocol provides an efficient way for exchanging information within a hierarchy. In most cases it is sufficient to just transmit data when changes occurred at the lowest hierarchy level. However, for billing purposes, it is still necessary to collect the actual consumption or production data from the lower level entities at least once per accounting period. However, the transmission of these information can happen asynchronously to the actual hierarchical forecasting process and thus, under consideration of free transmission and processing resources. Thus, our communication protocol still increases the communication flexibility. Additionally, we can use these billing information to verify and adapt the average historic share.

2.3 Benefits of the Approach

With our hierarchical forecasting approach we provide an efficient way for creating global forecast models at higher hierarchy levels. The aggregation of forecast models is far less time consuming compared to the conventional parameter estimation and thus, enables a rapid provisioning of forecasting results. As a result, entities responsible for balancing energy demand and supply such as the balance responsible parties and transmission system operators are enabled to work in the new real-time balancing environment. In addition, customers also directly benefit from integrating forecast models into their smart meters. The forecasting capabilities enable them to better analyze their energy consumption behavior and predict their possible future demand. Furthermore forecasting at the customer site is the basis for enhanced functionalities in smart home and smart building systems, by determining energy saving potentials and more intelligent regulation of the energy usage in buildings. Even the usage of concepts such as the MIRABEL flex-offers [2] or other demand-response systems [7] can be better implemented with the availability of local forecasts.

Final edited form was published in "SSDBM '13: Conference on Scientific and Statistical Database Management. Baltimore 2013",
Art. Nr. 37, ISBN 978-1-4503-1921-8
https://doi.org/10.1145/2484838.2484849

3. RELATED WORK

Forecasting in hierarchical environments is an emerging issue that gains more and more attention in research and industry. The general forecasting in hierarchies is analyzed in multiple historic and recent studies [3, 6, 13] that typically examine the issue of the most beneficial aggregation. In general two approaches exist. First, the bottom-up aggregation calculates forecast on the lowest hierarchy level and aggregates them to higher levels. Second, the top-down approach conducts the forecast calculation on higher levels and the results are disaggregated. It is not ultimately decidable, which aggregation type is more beneficial. Additionally, all studies only consider the complete aggregation of results and do not deal with the aggregation of models.

Hyndman et al. [10] introduces a hierarchical forecasting approach that combines the top-down and buttom-up forecasting with the goal of providing better predictions than the single approaches. To do so they calculate independent forecasts on all hierarchy layers and use a regression model to create a combined forecast in accordance to the hierarchical structure. In contrast to our approach, Hyndman et al. aim at increasing the forecast accuracy instead of optimizing the forecasting efficiency. His approach is computationally more expensive, due to estimating multiple forecast models on the aggregated levels using numerical optimization methods.

Fischer et al [4, 5] published an sampling approach using only a sample of base forecast models for forecasting in hierarchies. Thus, forecasts on a specific hierarchy level may be based on a subset of optimized models on other levels. With their approach they achieved reliable forecast accuracy comparable to aggregating all base models. If they base the forecast on only 21% of the base models, the precision is as high as if all the models were considered. In addition, they proposed a hierarchical model configuration advisor, which can be used to estimate on which hierarchy level forecast models are required. In contrast to our approach, the concept of Fischer et al. is based on aggregating time series and forecast values, instead of aggregating forecast models to increase the forecasting efficiency. In addition, in our hierarchical forecasting approach we only communicate between hierarchy levels, when forecast models are adapted, while Fischer et al. transmit data for any forecasting calculation. Furthermore, for their approach the creation of forecast models on higher hierarchy levels still involves the estimation of forecast model parameters.

4. CONCLUSION AND FUTURE WORK

In this paper, we presented our novel hierarchical forecasting approach that exploits the structure of hierarchically organized time series to substantially increase the forecasting efficiency. The core idea of our approach is to decentralize the forecast process and to push the responsibility of building and maintaining forecast models directly to the lowest level of the hierarchy. This is motivated by the fact that smart meters are getting more and more deployed to customers, which means that we can exploit their data processing capacities by directly integrating forecasting capabilities. Thus, we allow to create global forecast models on higher hierarchical levels by simply aggregating the individual forecast models of single customers, instead of conducting a time-consuming parameter estimation using optimization algorithms. Furthermore, we also defined a communication protocol, that limits the communication between hierarchy levels to only important messages. Overall, we provide a very efficient way for calculating forecasts in hierarchical environments.

In the future, we want to further elaborate on the details of our approach and also want to enhance it by allowing heterogenous forecast models between hierarchy levels.

5. ACKNOWLEDGMENTS

The work presented in this paper has been carried out in the MIRABEL project funded by the EU under the grant agreement number 248195.

6. REFERENCES

[1] H. Berthold, M. Boehm, L. Dannecker, F.-J. Rumph, T. B. Pedersen, C. Nychtis, H. Frey, Z. Marinsek, B. Filipic, and S. Tselepis. Exploiting renewables by request-based balancing of energy demand and supply. In *11th IAEE European Conference*, 2010.

[2] M. Boehm, L. Dannecker, A. Doms, E. Dovgan, B. Filipic, U. Fischer, W. Lehner, T. B. Pedersen, Y. Pitarch, L. Siksnys, and T. Tusar. Data management in the mirabel smart grid system. In *1st EnDM (15th EDBT)*, 2012.

[3] D. M. Dunn, W. H. Williams, and T. L. DeChaine. Aggregate versus subaggregate models in local area forecasting. *Journal of the American Statistical Association*, 71:68–71, 1976.

[4] U. Fischer, F. Rosenthal, and W. Lehner. Sample-based forecasting exploiting hierarchical time series. *Proceedings of the 16th IDEAS*, pages 120–129, 2012.

[5] U. Fischer, C. Schildt, C. Hartmann, and W. Lehner. Forecasting the data cube: A model configuration advisor for multi-dimensional data sets. In *29th ICDE*, 2013.

[6] G. Fliedner. Hierarchical forecasting issues and use guidelines. *Industrial Management & Data Systems*, 101:5–12, 2001.

[7] B. Gnauk, L. Dannecker, and M. Hahmann. Leveraging gamification in demand dispatch systems. *Proceedings of the 1st EnDM workshop (EDBT 2012)*, pages 103–110, 2012.

[8] J. G. D. Gooijer and R. J. Hyndman. 25 years of time series forecasting. *International Journal of Forecasting*, 22(3):443–473, 2006.

[9] R. J. Hyndman. Another look at forecast-accuracy metrics for intermittent demand. *The International Journal of Applied Forecasting*, (4):43–46, 2006.

[10] R. J. Hyndman, R. A. Ahmed, G. Athanasopoulos, and H. L. Shang. Optimal combination forecasts for hierarchical time series. *Computational Statistics and Data Analysis*, 55(9):2579–2589, 2008.

[11] MeRegio Project. http://www.meregio.de/en/ edition, 2012.

[12] MIRABEL Project. http://www.mirabel-project.eu edition, 2013.

[13] A. Zellner and J. Tobias. A note on aggregation, disaggregation and forecasting performance. *Journal of Forecasting*, 19:457–469, 2000.