Paper ST-02

Time Series Analysis: Separating Overlapping Events

M. Scott Elliott, FedEx Express, Memphis, TN

ABSTRACT

Analysts working with time series data often need to apply different events to their time series. An example of this problem is in dealing with global demand data, non-related holidays and events that need to be separated, especially those that overlap. SAS® Forecast Server and the SAS® High Performance Forecasting (HPF) language use 'dummy variables' to identify holiday events using one HPFEvents file for a project. A simple solution is to sort the data BY origin country and then use the same BY statement in the PROC TIMESERIES, but what if there are numerous locations per country? This paper describes one solution being implemented at FedEx that uses a macro to generate the HPFEvents file that will be used by the SAS® HPF program, handling each country separately. This method can also be used for local holidays and weather events that do not affect the entire time series.

INTRODUCTION

This paper explores how to separate overlapping events using SAS® Forecast Server and SAS® HPF programming language. The first section describes the problem of how holiday events can overlap and their impact of overlapping events. The second section provides a brief background of the SAS® Forecast Server products including the outputs and inputs of the product. The HPFEvents inputs are then explored further including examples. The next section explains overlapping events. That is followed by a section that provides a solution of how to separate non-related overlapping events including a macro example. The advantages and disadvantages of using that macro are then presented. The last section is the conclusion showing the results and recommendations for future enhancements.

FORECASTING PROBLEM

When all of the time series data is combined in a project, unrelated events can overlap and cause anomalies in the projected forecasts. The example that is presented below in Figures 1 to 3 are used to illustrate how two unrelated holidays in one time series can affect each other negatively. For proprietary reasons, the values are fictitious including the relative relationships of one week to the next week. The values are estimated for a general case that gift business commerce in the USA around Valentine's Day peaks at different times based on the day of week of the event and the business commerce in China surrounding Chinese New Year are very low. When the Chinese New Year holiday event crosses February 14th, forecast data with China and the USA become very difficult to predict into the future. The problem is compounded since the overlapping holiday events do not occur every year.

FIGURE 1. FEBRUARY AVERAGE DAILY BY WEEK FOR FOUR YEARS

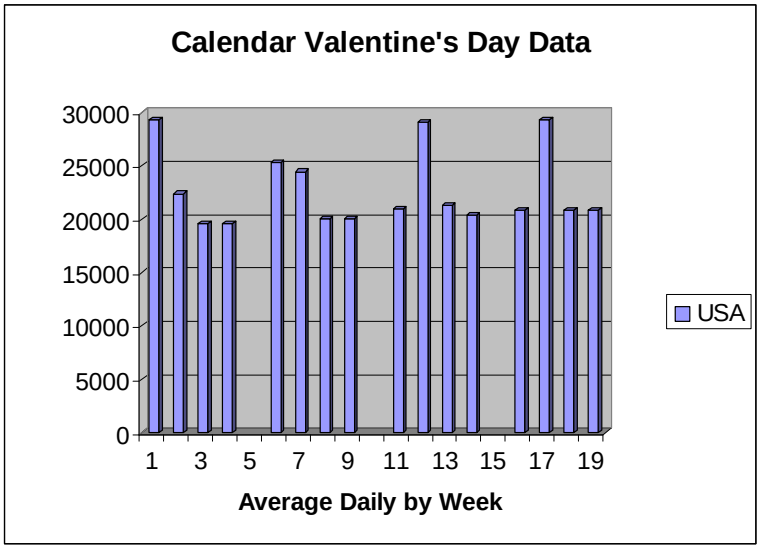

VALENTINE'S DAY HOLIDAY EXAMPLE

The data in Figure 1 shows four years of average daily values by week. Valentine's Day of February 14th is in the second week of each year, which corresponds to week 2, 7, 12, and 17. The three gaps separate the four years of

data. The peaking in weeks one and two are due to the day of week changing for February 14th. Based on the pattern in this graph, predicting a future forecast is not too difficult.

FIGURE 2. CHINESE NEW YEAR AVERAGE DAILY BY WEEK FOR FOUR YEARS

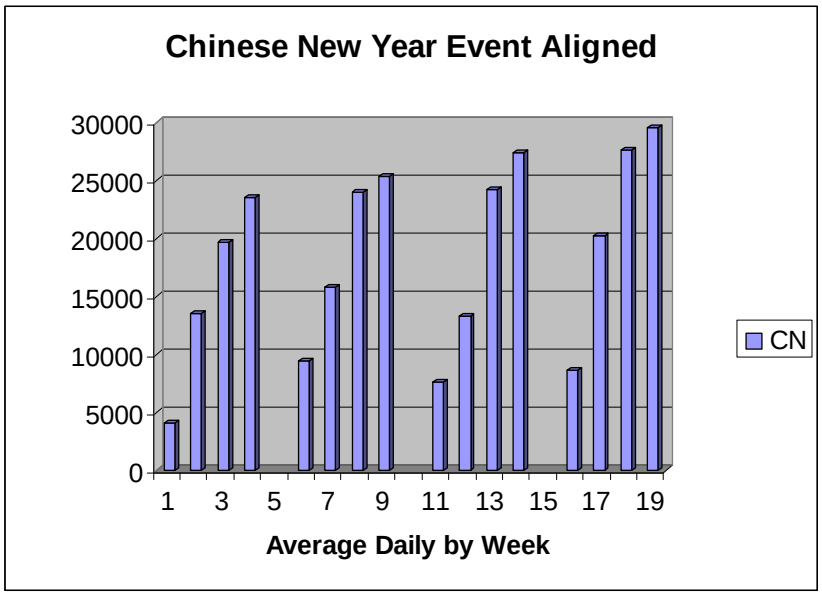

CHINESE NEW YEAR EXAMPLE

The data in Figure 2 shows four years of average daily values by calendar week for Chinese New Year. The data has been aligned by the starting week of Chinese New Year. In some years the first week starts in January while in other years the holiday starts in February and lasts one to three weeks depending on the location. SAS® Forecast Server software has tools to identify these types of holidays so that they can be aligned correctly. The tools will be discussed in length later in this paper. Like the Valentine's Day data in Figure 1, based on the pattern in this graph, predicting a future forecast is not too difficult.

FIGURE 3. COMBINED AVERAGE DAILY BY WEEK FOR FOUR YEARS

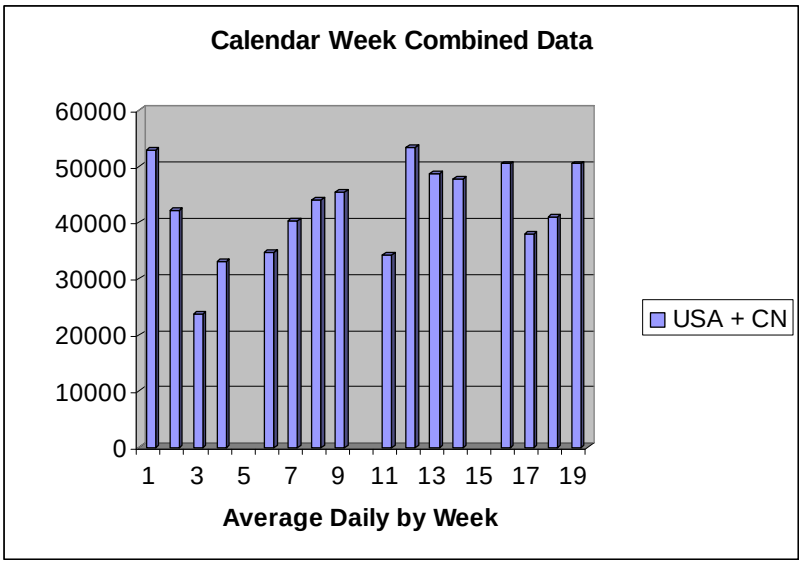

EXAMPLE OF COMBINING CHINESE NEW YEAR AND VALENTINE'S DAY HOLIDAYS

The data in Figure 3 shows four years of average daily values by week for Chinese New Year and the Valentine's Day holiday. Since the start of Chinese New Year varies, the effect of the holiday on the second week in February when Valentine's Day occurs fluctuates considerably. Based on the lack of a pattern in this graph, predicting a future forecast would be quite difficult. This paper examines one solution using SAS® Forecast Server software.

BACKGROUND OF SAS® FORECAST SERVER

The goal of SAS® Forecast Server software is to create multiple forecasts automatically for a large number of time series. The forecast accuracy is measured to identify which forecasts need analyst intervention to improve the results. The theory is that 80% of the forecasts can be created automatically, 10% cannot be accurately forecasted, and 10% needs extra attention by the analyst [Wells and Woodfield, 2008]. The idea is that the SAS® Forecast Server software will allow the analyst more time to focus on what needs to be enhanced rather than on production of the forecasts. The inputs to SAS® Forecast Server products are the time series historical data, the forecasting models, and one optional events file that is created using PROC HPFEVENTS. PROC HPFDIAGNOSE is then used to determine which model is most appropriate, and PROC HPFENGINE creates the forecast from the model selected [Truxillo, Wells, and Woodfield, 2007].

FORECASTING MODELS

SAS® Forecast Server generates SAS® HPF code for the analyst using a GUI interface called SAS® Forecast Studio that organizes the time series data as well as numerous parameters controlled by the analyst. For each time series, multiple forecasting models are fitted to the data. Numerous models are available in the SAS® Forecast Server in the model repository. Table 1 lists the model types in order of average run-time performance.

TABLE 1. FORECASTING MODELS IN SAS® FORECAST SERVER MODEL REPOSITORY

Model Type	Model Description
ESM	Exponential Smoothing Models (Winter's models)
IDM	Intermittent Demand Models for time series with many zero values
ARIMAX	AutoRegressive Integrated Moving Average with eXogenous variables (Box-Jenkins models)
UCM	Unobserved Components Models also known as structural decomposition models
OTHER	Simple models: the mean, random walk, and random walk with drift (used if other models fail) User-generated models can also be included in the model list

The above order of performance of the four most common models is useful to know which ones generate forecast outputs the quickest especially when dealing with very large time series data sets. However, the most accurate methods, on average, are ARMIAX, UCM, and ESM for continuous demand time series. IDM is the best method for time series with lots of gaps [Truxillo, Wells, and Woodfield, 2007]. The output forecast chosen by the SAS® Forecast Server software is the best model based on the accuracy measure chosen by the analyst.

ACCURACY MEASURES

The default accuracy measure in SAS® Forecast Server is MAPE (Mean Absolute Percent Error). MAPE is usually used by business analysts while RMSE (Root Mean Squared Error) is often used by academic statisticians. The model chosen using either accuracy measure is the smallest value [Wells and Woodfield, 2008].

SAS® FORECAST STUDIO EXAMPLE

The following figure is a snapshot of the output from SAS® Forecast Studio. The input data came from the SASHELP directory installed with the software. The SAS® data set name is AIR.sas7bdat. In the Modeling View, which is not displayed, the MAPE for the best ARIMA model is 2.82. Since the MAPE for the best smoothing model is 1.97, it is the model selected. A brief data analysis follows.

ANALYSIS OF THE PASSENGER AIRLINE FORECAST

Figure 4 on page 4 shows the forecasting view using SAS® Forecast Studio. This view displays the passenger airline historical data from January 1956 to January 1961 followed by a vertical black line separating the historical actual data from the forecast for the next year to January 1962. The light blue shade around the forecast line displays the confidence interval visually so the analyst can see potential problems. For example, the peak forecast pattern for July and August 1961 follows the similar pattern from 1960 with July higher than August. The override is available to the analyst to change the forecast due outside knowledge of the data. An event may have occurred that decreased the actual history in August 1961, such as a hurricane or civil unrest. However, it may be no event at all. Instead the decrease may be due to a seasonal pattern where each August with five Saturdays is higher than July, and each August with four Saturdays is lower than July, including the projected August 1961. It is up to the analyst to determine if a known event is non-repeating or will occur again the following year and let the SAS® Forecast Server software search for unknown events and determine if they are outliers.

FIGURE 4. SAS® FORECAST STUDIO OUTPUT FOR PASSENGER AIRLINE DATA

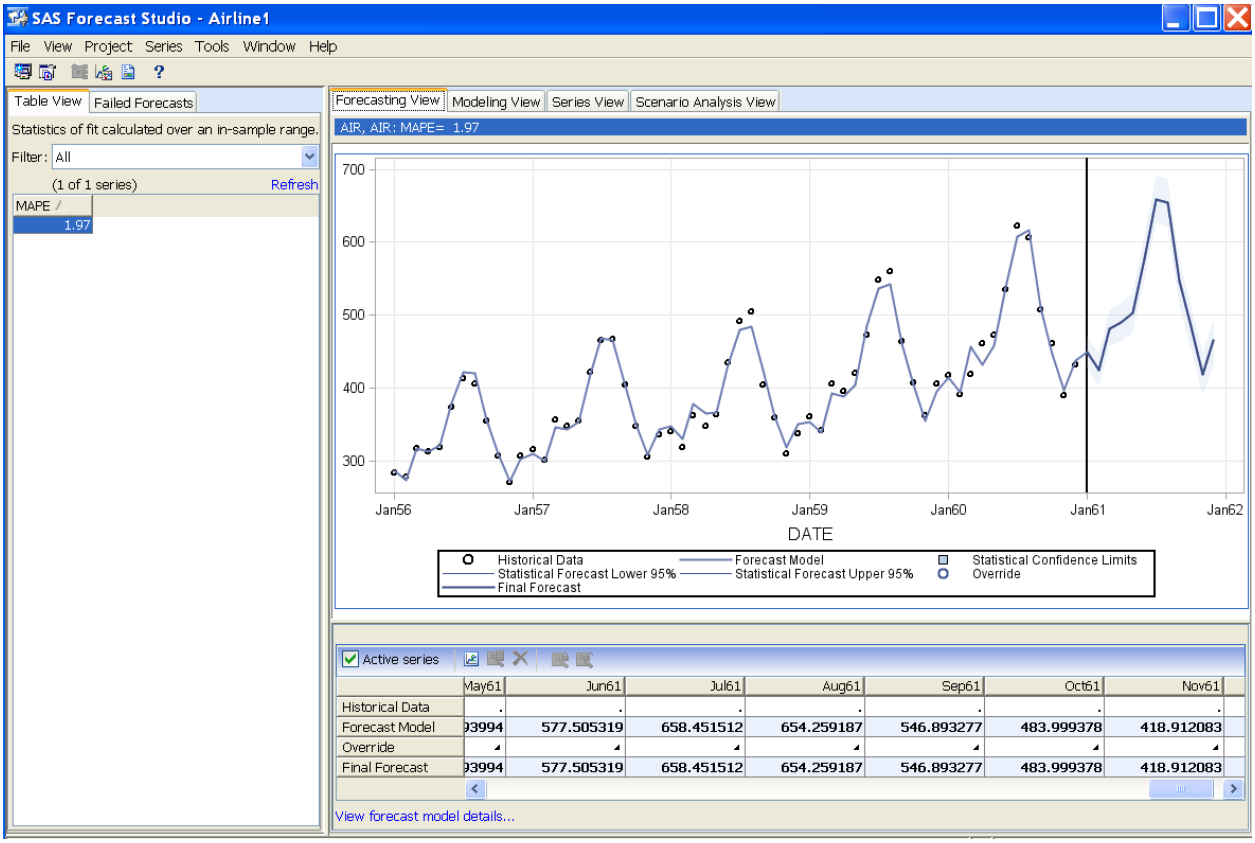

DESCRIPTION OF EVENTS

An event is a change in the time series data. Some of the most common events involve a level shift up or down, a change in the trend, or a single observation point also called a pulse. Some of the events repeat and some do not. In using the SAS® Forecast Server software, the event has a name, date and duration as well as several other options such as different types of growth and/or decay. The duration may be a single day like for a holiday or the rest of the time series for a permanent level shift change. PROC HPFEVENTS creates data sets that are inputs for the SAS® Forecast Server software. The EVENTDATA data set defines the event over the time series horizon, and the EVENTDUMMY creates a data set full of dummy variables with values of zero and one. SAS® HPF checks the event significance for the values greater than zero. The most common option for SAS® HPF is the 'use if significant' option, which is controlled by the analyst [SAS Institute Inc. 2006].

ADVANTAGES OF USING PROC HPFEVENTS

There are many advantages of using PROC HPFEVENTS. The data can be easily managed in the SAS® software. Once set up, the data can be easily maintained. Also, the dummy variables automatically adjust to the horizon of the time series and can be applied to different intervals such as weekly and monthly data. Another advantage is that the output data set can be used as input for other SAS® procedures outside of the SAS® HPF software like PROC REG [Truxillo, Wells, and Woodfield, 2007].

EVENT TYPES

Using PROC HPFEVENTS, there are many TYPES of events that can be used to estimate the shape of an event in the time series. A sample of the different types is listed in Table 2 on page 5. The shapes were downloaded from support.sas.com. The default shape refers to a permanent change, and the shape with finite duration refers to a temporary change where the duration after the event is greater than zero. For those types with SLOPE, the options are increasing slope referred to as GROWTH or decreasing slope referred to as DECAY [SAS Institute Inc. 2006]. An example of an overlapping event is where one event is defined as a permanent change level shift and another event is defined as a holiday pulse event in the same time period. This will be explained further in the next section.

TABLE 2. EVENT TYPES AND SHAPES

TYPE	DESCRIPTION	Default shape	Shape with finite duration
POINT	Point or pulse		
LS	Level shift		
RAMP	Ramp with SLOPE		
TR	Temp Ramp with SLOPE		
TC	Temp Change with SLOPE		
LIN	Linear trend		
QUAD	Quadratic trend		
CUBIC	Cubic trend		
INV	Inverse trend		
LOG	Logarithmic trend	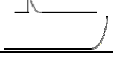	

EVENT GROUPS

Events can be grouped into man-made, natural, and holidays. Examples of man-made events are advertising, promotions, mergers and acquisitions, government and business policy changes, union strikes, tabloid scandals, and man-made disasters [Wells and Woodfield, 2008]. These events often start on a fixed date and have a duration that is short, long, or permanent. Man-made events can be level shift events, changes in trend events, and point events. For example, in business a common level shift event is due to a change in market share which causes a level shift followed by the normal data pattern at a different level. An example of this is a strike by one company that decreases their output and increases the output of their competitors. When the strike is settled, another level shift can occur excluding any customers who do not return to the previous company. A change in trend example could be due to successful advertising of a new product. This would be depicted as a ramp type event with slope=growth.

NATURAL EVENTS

Many of the natural events are to model the effect of a natural disaster [Wells and Woodfield, 2008]. Examples of natural events are floods, hurricanes, tornados, earthquakes, tsunamis, wildfires, and blizzards. Like the man-made events, natural events often start on a fixed date and have a duration that is short, long, or permanent. Natural events can be level shifts, changes in trend events, and point events. For example, Hurricane Katrina had a devastating impact on New Orleans. Many of the people affected moved to Baton Rouge and have not returned to New Orleans. In a time series model, this would be a level shift in both locations. New Orleans would be a shift downward while Baton Rouge would be a shift upward.

HOLIDAY EVENTS

The holiday events can be grouped into major holidays that disrupt the time series data and minor holidays that may or may not disrupt the data. Since overloading SAS® Forecast Server software with events can affect performance and the output as some models will fail, eliminating minor holidays is a prudent decision. The hard decision is determining which holidays are minor and should be ignored. Like man-made events and natural events, holiday events often start on a fixed date and have a duration that is short or long that may occur before the event as well as after the holiday event. Holiday events are rarely permanent. They are usually depicted as point or pulse events. The days before the holiday may be a normal day, a peak day anticipating the holiday, or a low day as the holiday celebration begins before the holiday event. Likewise, the days after the holiday may return to a normal day, peak as a recover day, or slowly recover over one or more days back to normal. The impact of the holiday may also vary based on the day of week of the event. Not all events need to be defined since the SAS® HPF program includes predefined holidays for Canada and the United States. Table 3 below lists the SAS® predefined US holidays events [Truxillo, Wells, and Woodfield, 2007]. An event description for each holiday has been added to compare the holiday data. Since the day of week of the event of the fixed date holidays varies, they are usually harder to forecast.

TABLE 3. SAS® PREDEFINED US HOLIDAY EVENTS

NAME OF SAS® PREDEFINED EVENT	DEFINITION	EVENT DESCRIPTION
NEWYEAR	January 1st	Fixed date
VALENTINES	February 14th	Fixed date
EASTER	Easter Sunday	Floats between March and April
MOTHERS	2nd Sunday in May	Fixed day of week
MEMORIAL	Last Monday in May	Fixed day of week
USINDEPENDENCE	July 4th	Fixed date
LABOR	1st Monday in September	Fixed day of week
HALLOWEEN	October 31st	Fixed date
THANKSGVING	4th Thursday in November	Fixed day of week
CHRISTMAS	December 25th	Fixed date

SAS® FORECAST STUDIO EXAMPLES

The following two figures are snapshots of the output from SAS® Forecast Studio comparing the output without using PROC HPFEvents and the output with PROC HPFEvents. The underlying SAS® source code is available but is not normally displayed in SAS® Forecast Studio. This example provides a visual representation of how using PROC HPFEvents can improve the forecast output as well as demonstrate an example of a time series that has become difficult to forecast. The input data came from the SASHELP directory installed with the software. The SAS® data set name is GULFOIL.sas7bdat. In the Modeling View, which is not displayed, the MAPE for the best ARIMA model is 6.28. Since the MAPE for the best smoothing model is 5.95, it is the initial model selected. When the events are defined, the new MAPE is 5.55. A brief data analysis and comparison follows.

FIGURE 5. SAS® FORECAST STUDIO OUTPUT FOR GULF OIL DATA WITHOUT PROC HPFEVENTS

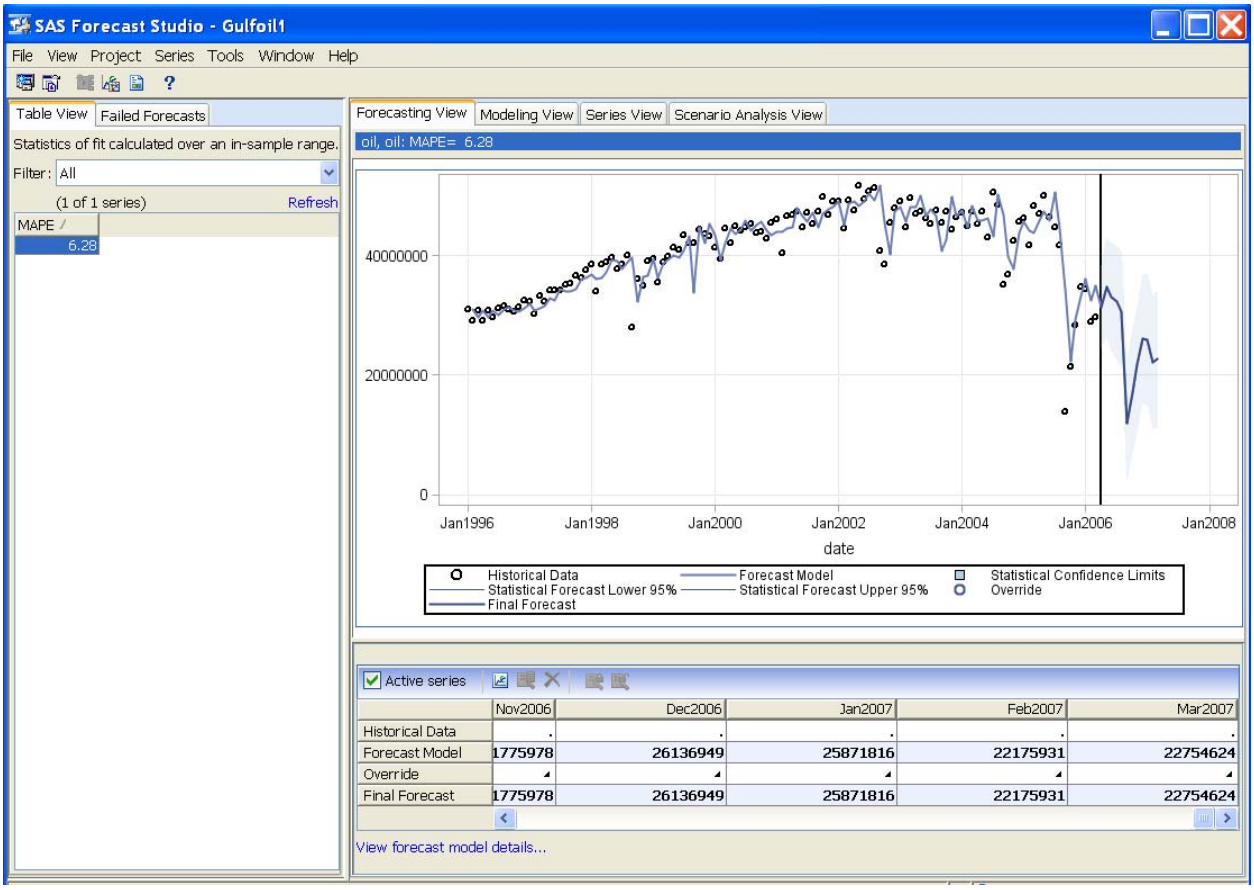

ANALYSIS OF THE GULF OIL FORECAST WITHOUT PROC HPFEVENTS

Figure 5 shows the forecasting view using SAS® Forecast Studio. This view displays the oil historical data of Gulf Oil from January 1966 to March 2006 and extending the forecast to March 2007. The light blue shade around the forecast line displays the confidence interval visually so the analyst can see potential problems. It is much wider than the Airline example in Figure 4. In this example, a significant event occurred in September 2005; however, in this example, the event is not defined. By using the outlier detection option, the September 2005 data was not included in selecting the model. It is visually obvious that more work needs to be done by the analyst to improve this forecast. Figure 6 shows the same data with PROC HPFEvents included.

ANALYSIS OF THE GULF OIL FORECAST WITH PROC HPFEVENTS

Figure 6 on page 7 shows two windows. One window is the updated forecasting view using SAS® Forecast Studio. The other window shows the Event Repository. The first event has been named DROP1. It is a pulse, single observation event for October 2005 and the usage in SAS® Forecast Studio is defined as 'use if significant'. The second event has been named NEWSTATE1. It is a level shift starting at November 2005 and will also be used if significant. Since almost every observation that is treated as an outlier is over the September to November time frame, the projected forecast maintains that pattern. Also, the override feature has been used to increase the September 2006 forecast by 12%. The analyst would be expected to explain the reason for the override and the override value chosen.

FIGURE 6. SAS® FORECAST STUDIO OUTPUT FOR GULF OIL DATA WITH PROC HPFEVENTS

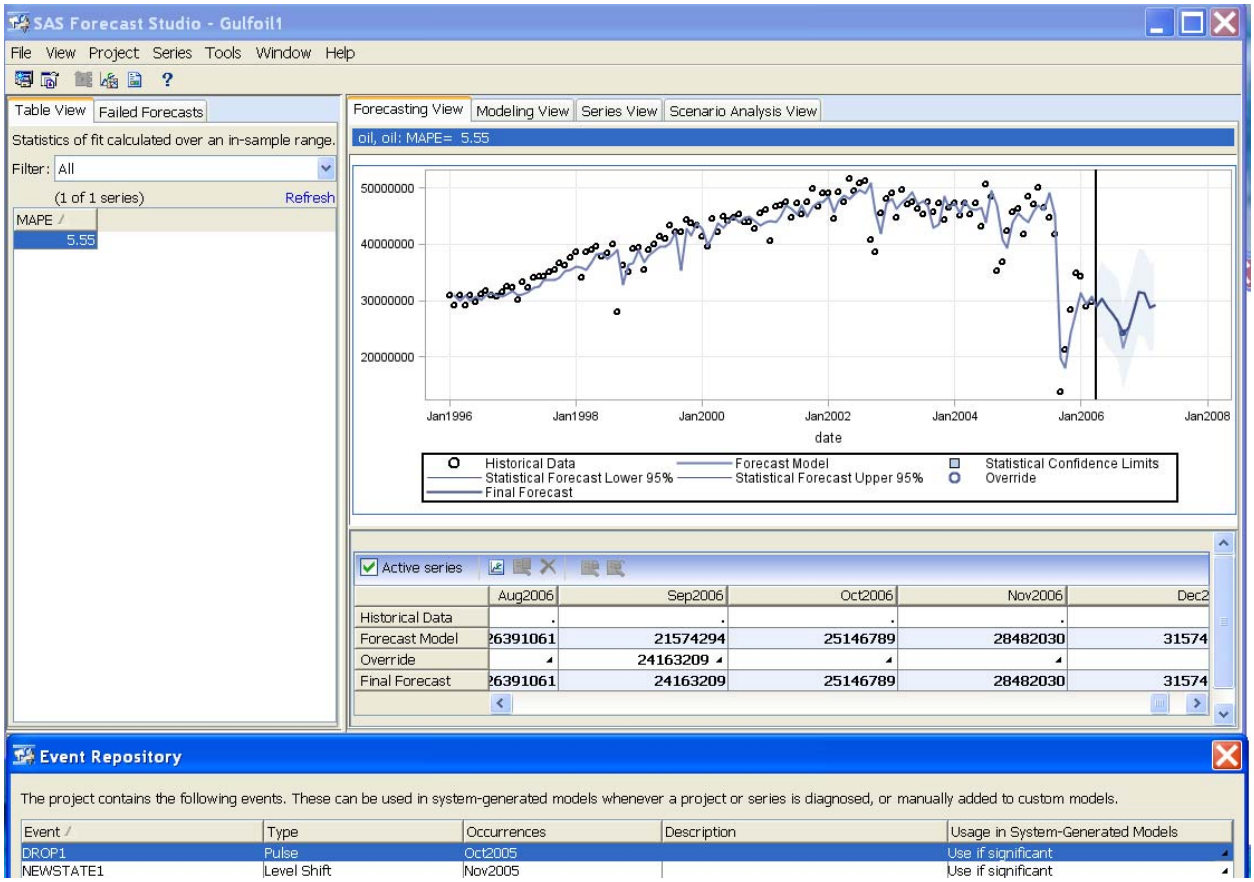

OVERLAPPING EVENTS

Overlapping events occur frequently especially in global data. Many countries have common holidays that they celebrate, but most have their own specific holidays that are not celebrated by other countries. Also, natural disasters sometimes overlap a holiday so that the normal holiday pattern is significantly changed. The Gulf Oil time series in Figure 6 is an example. Any events beyond the Drop1 event starting in October 2005 will overlap with the level shift event. When an event is defined in SAS® HPF software, one option is called RULE which becomes activated when events overlap. The values for this option are listed in the Table 4 below. The default is ADD which would work well with the Gulf Oil time series [SAS Institute Inc. 2006].

TABLE 4. DEFINITION OF RULE VALUES FOR OVERLAPPING EVENTS

RULE	NAME	DESCRIPTION
ADD	Add	Add the values
MAX	Maximum	Use the maximum value
MIN	Minimum	Use the minimum value
MINNZ	2nd Sunday in May	Use the minimum nonzero value
MINMAG	Last Monday in May	Use value with least magnitude
MULT	1st Monday in September	Multiply

RELATED OVERLAPPING EVENTS EXAMPLES

There are instances where overlapping events is preferred, especially if the events are related. An example would be holidays that occur over a weekend. In Dubai, United Arab Emirates, Friday is a very low commerce day because it is the start of their weekend. So, a holiday event that includes Friday will usually be lower than if the holiday occurred on a non-Friday. An option that can be used to solve this problem is that when the event is defined, the RULE=MIN could be included, which specifies that the minimum value of the overlapping data will be used [SAS Institute Inc. 2006]. Another example is a fixed holiday date occurring on a weekend. In some countries, an observed weekday holiday is used. In 2011, December 25th will occur on Sunday. Some countries will observe Monday, December 26th as the Christmas holiday; however, several European countries will not move the observed holiday to Monday, so

Sunday, December 25th should have very low business commerce due to it being Sunday while Monday, December 26th will only be slightly reduced or a normal business day instead of the Christmas holiday impact depending on the country. Again if RULE=MIN was used, then Sunday December 25th would be the lower of Christmas and regular Sunday impact. Handling this event could be quite difficult for countries observing Christmas on Monday.

UNRELATED OVERLAPPING EVENTS EXAMPLES

The Chinese New Year holiday frequently overlaps with the Valentine's Day holiday event as shown in Figures 1-3. Other examples would be Islamic holidays which follow a lunar cycle rather than a 365 day cycle. Due to the cycle, the overlap often occurs in just one year. This means that it may be significant the following two years, but the impact will decrease in most models over time. It is up to the analyst to determine the amount of time dedicated to handling overlapping events.

EFFECT OF OVERLAPPING EVENTS

When overlapping events occur in the data, SAS® Forecast Server may link the events even if they are not related. This potentially could cause an anomaly in the forecast that the analyst or customer may or may not discover. The above example of Chinese New Year overlapping with Valentine's Day may only slightly decrease future February 14th forecasts in China and slightly alter future Chinese New Year weeks in the USA. Not correcting for the overlap may have a negative impact on the accuracy of the forecast in both countries as well as the other countries in Asia that are affected by Chinese New Year.

SEPARATING OVERLAPPING EVENTS

There are several approaches to this problem. A simple solution is to sort the time series data BY origin country and then use the same BY statement in the PROC TIMESERIES. This approach will work if the input time series data has only one location per country otherwise an error is generated. Another solution would be to create a macro to generate multiple HPFEvents files that will then be used by the SAS® HPF programs. Each country would be handled separately and the BY variable will be origin location with the appropriate holiday effects applied.

MACRO EXAMPLE

In this macro the raw data time series, the output HFPEvents file and the event definitions are kept in separate files to be included in the SAS® source code. The &OCNTRY macro variable is the controlling variable that controls the inputs and outputs of the macro. In this example it has been set to 'FR' for France. Here is the macro source code:

```
/*=================================================================*/
/* MACRO variable, library names, and file definitions.            */
/*=================================================================*/
%let OCNTRY='FR'; /* origin country is controlling variable */
%let ENDFCST='31DEC2011'd;   /* forecast horizon */
%let libdir="/Forecast_Studio/Raw_Data";
%let infile=loc_horizon_ocntry_data;
%let inevents=/Forecast_Studio/hpfevents/&OCEVENT; /* origin country events */
%let woevents=/Forecast_Studio/hpfevents/WO_hpfevents; /* worldwide events */
%let outlib="/Forecast_Studio/hpfevents";
libname history &libdir;
libname events &outlib;
/** Print the values associated with all the global macro variables     */
%PUT _GLOBAL_;
  run;
%MACRO FCSTEVENTS;
/*=================================================================*/
/* Read in time series data and keep only the origin country history.   */
/* Sort the smaller data set by origin country, forecast metric, origin */
/* location, grouping hierarchy, product & date.                   */
/*=================================================================*/
data olocfcst;
  set history.&infile;
  if ocntry = "&OCNTRY1";
proc sort in=olocfcst;
 by ocntry metric oloc hierarchy_type group prodname date1;
/*=================================================================*/
/* Build events data set for the special days by origin country.        */
/*=================================================================*/
```

```
proc hpfevents data=olocfcst;
  by ocntry metric oloc hierarchy_type group prodname;
  var horizon_sum;
  id date1 end=&ENDFCST; /* endfcst would be the last day of forecast horizon */
/*== Include Origin Country and Worldwide Holidays ==*/
%INCLUDE "&INEVENTS";
%INCLUDE "&WOEVENTS";
eventdata out=events.horizon_events_intl_rmp_&OCNTRY1;
eventdummy out=events.horizon_event_dummies_intl_rmp_&OCNTRY1;
run;
%MEND;
%FCSTEVENTS
```

EXPLANATION OF MACRO

The purpose of the macro is to create input data sets for SAS® Forecast Server software so that the impact of the holiday events will be accommodated by SAS® HPF. The above example is built to match the forecast hierarchy which is controlled by the analyst. SAS® Forecast Studio provides an interface for creating the hierarchy. Since the time series data has been organized for processing based on this hierarchy, then the EVENTDATA and EVENTDUMMY data sets also should be organized by the same hierarchy. The input macro variable, &OCNTRY, is used to control the input and output streams. A sample of the two include files is provided below.

MACRO EXAMPLE SAMPLE INPUT SOURCE CODE

The event definitions are kept in a separate file to be included in the SAS® source code to decrease the size of the macro file. Otherwise, most of the macro file would be lines and lines of event definitions and their corresponding event dates. The worldwide events are defined and maintained in one small file instead of being replicated for every origin country. The origin country code has been added to the event definition name for ease of programming. Many holiday events are significant at a local level but are not significant at a global level. For example, Bastille Day is a major holiday but only for France. The first three event definitions would come from the "&INEVENTS" file, and the last worldwide event definition would come from the "&WOEVENTS" file. France was chosen to show how some holidays change observed days while others do not. For example, Bastille Day is not observed on Monday in 2002. Likewise May Day is not observed on Monday in 2005. France also has a holiday that only occurs on Monday for Whit Monday. Here is an example of that source code that is included in the macro in this format:

```
eventdef FR_Bastille_Day = ( '14JUL2003'd, '14JUL2008'd,
   '14JUL2002'd) / type=point; /* holiday remains Sunday, 14JUL2002 */
eventdef FR_May_Day = ( '01MAY2003'd, '01MAY2006'd,
   '01MAY2005'd) / type=point; /* no observed holiday on Monday, 02MAY2005 */
eventdef FR_Whit_Monday = ( '01JUN2009'd, '09JUN2003'd,
   '31MAY2004'd) / type=point; /* observed holiday always on Monday */
eventdef WO_Good_Friday = EASTER / SHIFT=-2 type=point;
   /* EASTER is a SAS® reserved word for Easter Sunday -2 days is Good Friday */
```

MACRO ADVANTAGES

The advantage of using this method is that since the subset data is smaller in size, the individual parts will run faster and provide quicker feedback for model error detection. Also, the SAS® HPF programs have more customization features that the analysts can use specifically for the ten to twenty percent of the time series that needs more attention by the analyst. For example, the analyst can control the HPFEvents file data by origin country.

MACRO DISADVANTAGES

The disadvantage of using this approach is that it subsets the data so it must be run using SAS® HPF instead of in the SAS® Forecast Studio GUI. At this time, there is not a way to import the SAS® HPF output into SAS® Forecast Studio so the output graphs and interface is not available. Another disadvantage of this approach is that it splits the project into multiple pieces.

CONCLUSION

The result of using the macro is that the analyst is able to prevent non-related holidays from affecting countries whose holidays overlap. This solution could be expanded to accommodate changes in the time series at a local level caused by natural disasters like floods, hurricanes, tornados, earthquakes, blizzards and wildfires or to estimate local holiday impacts of the Olympics in the host city, Mardi Gras in New Orleans, Sweetest Day in Chicago, St. Patrick's Day in Boston, and the Cherry Blossom Festival and inaugurations in Washington D.C. The SAS® Forecast Server products are powerful tools to handle the problem of forecasting events.

REFERENCES

http://support.sas.com/documentation/cdl/en/hpfug/63959/HTML/default/viewer.htm#hpfug_hpfevents_sect016.htm

Truxillo, C., Wells, C. and Woodfield, T. *Using SAS® High-Performance Forecast Software Course Notes*, Copyright © 2007 By SAS Institute Inc., Cary, NC 27513, USA.

Wells, C. and Woodfield, T. *Forecasting Using SAS® Forecast Server Software Course Notes*, Copyright © 2008 By SAS Institute Inc., Cary, NC 27513, USA.

SAS Institute Inc. *SAS® High Performance Forecasting 2.2: User's Guide, Volume 1 and 2*, Copyright © 2002-2006 By SAS Institute Inc., Cary, NC 27513, USA.

ACKNOWLEDGEMENTS

I would like to acknowledge the contributions of the following individuals who helped in their expertise of using the SAS Forecast Server® products. They are Pete Dillman and Chip Wells of SAS®, who helped train us on using the SAS® forecasting products and Deborah Smith of FedEx Express, who provided feedback and support in writing the macro and this paper. Finally, I would like to thank my family for their love and patience.

RECOMMENDED READING

For more information on the forecasting models that are built into SAS® Forecast Server, please read Chapter 5, *Using SAS® High-Performance Forecast Software Course Notes.* Most of the forecasting models have several different methods that are included with the SAS® Forecast Server product.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

M. Scott Elliott
FedEx Express
3680 Hacks Cross Road
Building H, Second Floor
Memphis, TN 38125
Work phone: 901-434-9092
Fax: 901-434-9634
E-mail: mselliott@fedex.com