A POOR MAN'S ESP: STATISTICAL HYDROGRAPH TRACE ADJUSTMENT

T. C. Pagano* and D. C. Garen

ABSTRACT

Water managers increasingly demand forecasts of hydrograph characteristics beyond seasonal volumes. Currently, most volumetric forecasts of seasonal streamflow are produced using statistical regression techniques. One way to provide additional long-lead information about peak flows, low flows, and number of days to a particular threshold is to apply the Ensemble Prediction System, first developed by the National Weather Service. This Ensemble Prediction System involves the calibration of a hydrologic simulation model, model initialization using current watershed states, and forcing based on a number of observed historical meteorological traces. The output is a series of "possible future" daily hydrographs, from which the above mentioned characteristics can be derived. As a low-cost alternative, an ensemble hydrograph forecast can be obtained by rescaling historical flow traces by the existing statistical seasonal streamflow volume forecast. The volume of streamflow from past years is mapped into the distribution of the statistical error bound of the official seasonal volumetric forecast to obtain a multiplier factor for each year. This multiplier is then applied to each daily flow value of the historical year. While this simple system has many shortcomings and limitations, it still provides useful information and may serve as a credible "naïve forecaster" baseline against which to compare the performance of other ensemble forecast systems.

INTRODUCTION

The Natural Resources Conservation Service produces seasonal water supply forecasts monthly, January through June, in partnership with the National Weather Service (NWS) and local cooperating agencies, such as the Salt River Project in central Arizona. A typical water supply forecast, as published, includes the name of the forecast location (e.g., "White River near Meeker"), the forecast target season (e.g., "April-July"), the long term historical average flow volume, and the forecasted flow volumes corresponding to each of 10%, 30%, 50%, 70%, and 90% exceedance probabilities.

Forecasts are produced using the statistical principal components regression technique described by Garen (1992). The regression equations so derived are used to compute the median (50% exceedance) value of the seasonal water volume forecast distribution. Human expertise or other guidance may be used to shift this median value up or down, if the forecaster feels an adjustment is in order. A probabilistic error bound is then added to the forecast. This error bound typically narrows as the season progresses as there is less uncertainty about the forecast. The forecasts assume a normal (Gaussian) error distribution whose width is proportional to the root mean squared error of the forecast equation during jackknife calibration. At locations where the relationship between the seasonal streamflow volume and the predictor variables is not linear, nonlinear equations are used in which transformed (square root, cube root, natural logarithm) streamflow values are predicted by a linear equation. In these cases, the normal error bound is applied before the streamflow is un-transformed to obtain the final forecast value. For example, the White River near Meeker (Colorado) forecast is currently based on a natural logarithm transform of the flow. Its forecast error distribution is therefore log-normal.

The demand for hydrologic information in addition to the seasonal streamflow volume has led to the development of advanced forecasting tools and products, based on a hydrologic simulation model rather than statistical procedures. The use of a hydrologic simulation model makes it possible to generate long-lead information about hydrograph characteristics such as peak flows, low flows, and number of days to a particular flow threshold. The most commonly used procedure to accomplish this is the Ensemble Prediction System (ESP) of the NWS (Day, 1985; hereinafter referred to as NWS ESP), which is described in the next section.

Paper presented Western Snow Conference 2003

* US Department of Agriculture, Natural Resources Conservation Service, National Water and Climate Center, 101 SW Main Street, Suite 1600, Portland, OR 97204-3224; Ph (503) 414-3010; email tpagano@wcc.nrcs.usda.gov or dgaren@wcc.nrcs.usda.gov

Simulation modeling, however, is an extremely resource-intensive activity. It needs a great deal of input data requiring intensive screening and quality analysis on daily and sub-daily data values for both calibration and real-time forecasting. Calibration itself is a time-consuming mix of science and art. Additionally, any model is restricted by its simplified representations of hydrologic processes and will never be able to reproduce the observed streamflow in all situations and climates, requiring constant vigilance and adjustments by the forecaster to keep the model on track. These requirements are very demanding on personnel, which may or may not be available to a given agency. In water year 2003, for example, only four NRCS hydrologists operated and maintained statistical forecasting procedures for over 700 forecast locations across the West. Each forecaster typically has less than three working days to create, analyze, adjust, coordinate, and issue forecasts for close to 175 points simultaneously. In comparison, a typical simulation modeling hydrologist at the NWS's Colorado Basin River Forecast Center is responsible for less than 20 water supply forecast locations. It is clear that the implementation of a full ESP system would be very difficult for a staff the size of the NRCS forecasting group, making an alternative method requiring fewer personnel resources desirable.

This paper represents an attempt to develop an extremely low-cost method for generating simulation model-like daily ensemble streamflow forecasts. Although there have been past efforts to derive monthly volume forecasts from seasonal streamflow volume forecasts by statistical disaggregation (Hoshi and Burges, 1980; Pei et al., 1987; Reese and Krzysztofowicz, 1989), such a method has not been applied to daily flows and is probably not feasible. The method presented here for obtaining daily ensemble streamflow traces is very feasible.

ESP CONCEPTS AND TERMINOLOGY

The NWS ESP method involves the calibration of a hydrologic simulation model, model initialization using current watershed states, and forcing based on a number of observed historical meteorological traces. The output is a series of "possible future" daily hydrographs, which can be analyzed statistically to derive any desired hydrograph characteristic. Typically, the NWS River Forecast Centers employing this method operate the Sacramento Soil Moisture Accounting Model (Burnash et al., 1973) coupled with the HYDRO-17 snow model (Anderson, 1973). After the model is initialized with the current watershed state, it is forced with meteorological traces generally from the mid-1970's to the end of the 1990's (i.e., ~25 years). The procedure therefore assumes that the historical climate is a good analogue for the future climate. It essentially answers a series of "What If?" questions. That is, given the watershed state today (wet or dry, high or low snowpack), what would the future streamflow be if a meteorological sequence like what occurred in 1984 reoccurred during the remainder of the season? Or what if a meteorological trace such as 1986 reoccurred? Passing many historical meteorological sequences over today's watershed gives the forecaster an "ensemble" of possible futures.

The "observed" streamflow is that which is measured at a location, as realized by nature. A collection of past observed years is often referred to as "climatology". The "historical simulation" is a model's best attempt to reproduce the observed flow, by forcing it with the soil moisture and snow state of that year, as well as the observed meteorological trace of that year. For example, a model is initialized with April 1st 1983's soil moisture and snow states. The model is then run forward in time, forced with April 1st-July 31st 1983's observed precipitation and temperature. The resultant streamflow is the historical simulation of 1983's flows. A "conditional simulation" involves pairing basin states and meteorology from different years. For example, a model is initialized with April 1st 2003's basin states and forced with April 1st-July 31st 1983's observed precipitation and temperature. The resultant conditional hydrograph will differ from the historical simulation, as will its seasonal volume. At forecasting time, the NWS ESP initializes the model with the most recent basin states and develops as many conditional simulations as there are years of historical meteorological data available.

STATISTICAL HYDROGRAPH TRACE ADJUSTMENT

The alternate ensemble forecasting methodology proposed herein bypasses the simulation model and uses the observed streamflow data directly as a historical "simulation". Its conditional traces are rescaled versions of the observed streamflow so that the seasonal volumes are consistent with the error distribution of the official seasonal statistical water supply forecasts described above.

As described above, the seasonal statistical water supply forecast contains a 50% probability of exceedance ("median") value along with its probabilistic error bound for the total flow volume over a particular period (e.g., April-July). Knowing the shape and moments of this distribution, one can calculate the forecast streamflow volume at any probability of exceedance level. Conversely, one can specify a particular streamflow volume amount and calculate its probability of exceedance. These volumes and probabilities can be compared to the climatological distribution of observed seasonal flow volumes at this location. Generally, the forecast (conditional) distribution will be shifted and narrowed relative to the observed (unconditional) distribution.

The hydrograph adjustment procedure consists of the following steps (see also the appendix):

1) Determine the flow volume for the seasonal period of interest for each year in the historical streamflow record.
2) Develop a probability distribution of the historical observed flow volumes (the unconditional distribution).
3) Produce a streamflow forecast, typically from a statistical model, consisting of a median value and an error distribution (the conditional distribution).
4) For each year in the historical record, do the following:
    a) Determine the unconditional exceedance probability of the flow volume from the distribution obtained in step 2.
    b) Find the flow volume from the forecast (conditional) distribution that corresponds to the exceedance probability obtained in step 4a.
    c) Compute the ratio of the conditional flow volume from step 4b to the observed volume.
    d) Multiply each observed daily flow by the ratio computed in step 4c.

This procedure amounts to a mapping of the historical distribution to the forecast distribution. An example of these calculations is given in Table 1 and illustrated in Figure 1. The final column in Table 1 contains the ratio used to accomplish the mapping and is the factor multiplied with each daily flow in that year's hydrograph to obtain the adjusted hydrograph. Doing this for each year results in an ensemble of hydrographs whose seasonal volumes are consistent with the official water supply forecasts. The values shown here are conditioned on the Apr-July forecast for the White River near Meeker (Colorado), issued on February 1st 2003.

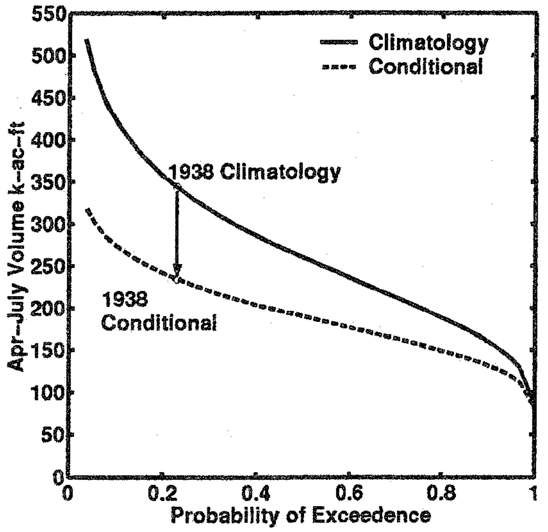

Figure 1. Probability of exceedance of climatological (solid) and conditional (dashed) April-July seasonal flow volumes for the White River nr Meeker 1936-2002. The conditional distribution is based on the February 1st 2003 official water supply outlook. The remapping of a single historical year's flow volume (1938) is shown.

Table 1. Example calculations used in rescaling seasonal flows, based on the February 1st 2003 April-July forecast for the White River nr Meeker. The Apr-July seasonal volumes from 1936-2002 have a log-normal climatological distribution with median 260 k-ac-ft, log median of 5.56, and log standard deviation of 0.379. The forecast distribution is log-normal with median 190 k-ac-ft, log median of 5.25 and log standard deviation of 0.286.

Year	Observed Apr-Jul Volume	Unconditional Probability of Exceedance	Corresponding Conditional Apr-Jul Volume	Conditional/Observed Ratio
1977	81	0.99	79	0.97
1992	166	0.88	136	0.82
1969	260	0.50	190	0.73
1938	344	0.23	235	0.68
1984	519	0.03	320	0.62

For snowmelt dominated basins, this linear rescaling is generally a fair approximation of behavior during the heart of snowmelt season, but it can produce physically unrealistic values during baseflow conditions before and after the snowmelt season. While daily snowmelt flow values easily span an order of magnitude, baseflows typically vary only in a small range. Therefore, a separate rescaling procedure is necessary to obtain realistic baseflow values outside the forecast target season.

To accomplish this, a linear regression equation is first developed to predict the post season baseflow volume (e.g., August-December total flow volume) using the forecast period flow volume (e.g., April-July). If real-time streamflow data are available, pre-season baseflow volumes (e.g., November flow at the start of the water year) can also be used as a predictor. If this second variable is used, the forecast will take advantage of the temporal autocorrelation in baseflows. The pre-season baseflow period should correspond to a time of year when the streamgage is not frozen.

While this equation remains fixed and universal across all years, when forecasting, the inputs to the equation are the conditional forecast period flow volume for the individual year, and, if available, the observed pre-season baseflow volume. The ratio of this predicted conditional post seasonal flow volume to the observed flow volume can be used to rescale the observed post seasonal daily flow values, as before. This procedure is repeated for the individual months during the pre-season baseflow, such as February and March. Table 2 shows the post-season conditional baseflow calculation procedure for the February 1st 2003 forecast.

Table 2. Calculation of the post-seasonal August-December flow volumes. Post season volumes are predicted using the equation Aug-Dec volume (k-ac-ft) = 24.6 + 0.0618* November average flow (cfs) + 0.2211* Apr-July volume (k-ac-ft). The November 2002 observed average flow is 260 cfs. All volumes listed below are in k-ac-ft.

Year	Conditional Apr-Jul Volume	Observed Aug-Dec Volume	Conditional Aug-Dec Volume	Conditional/Observed Post-Season Ratio
1977	79	61	58	0.95
1992	136	79	71	0.90
1969	190	110	83	0.75
1938	235	111	93	0.83
1984	320	190	112	0.59

Simple linear rescaling of the daily flow preserves the recession characteristics of the hydrograph if the recession has the form of $Q_{n+1} = kQ_n$. That is, in the absence of new water entering the stream, tomorrow's flow equals today's flow multiplied by a constant (k) that is less than 1.0. In nature, however, a better representation is that k is not a constant, being smaller (quickly recessing) during high flows and larger (slowly recessing) during lower flows. An exponential recession model of the form $Q_{n+1} = aQ_n^b$ is used by Martinec et al. (1983). The authors have developed an empirical iterative non-linear rescaling method that preserves an exponential recession, and this technique is currently being evaluated. Such a procedure may improve the performance of the technique during low flows and eliminate the need to develop separate pre- and post-season baseflow rescaling equations.

AN OPERATIONAL EXAMPLE

A daily forecast containing 67 ensemble members was developed for the White River nr Meeker (figures 2 and 3) using the procedure described above. The forecast was conditioned on the April-July water supply outlook issued February 1st 2003 as well as the realtime streamflow data available in November 2002. As 2002 was an exceptionally dry year, the observed streamflow for November 2002, at 72% of average, is the second driest November on record. The median of the seasonal forecast distribution, 190 k-ac-ft, is 73% of the climatological median of 260 k-ac-ft and 65% of the climatological mean of 290 k-ac-ft.

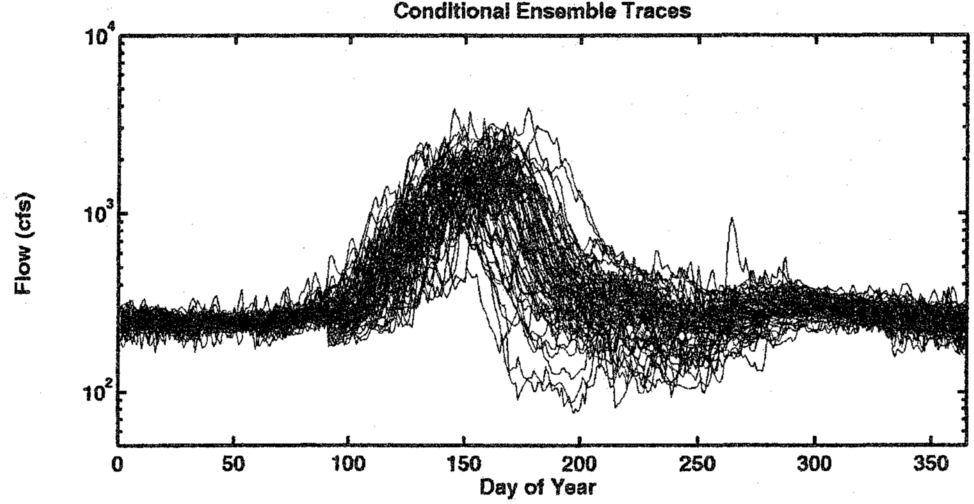

Figure 2. Daily conditional flow prediction. Each trace corresponds to an individual ensemble member. The discontinuity in flows on day 91 corresponds to the beginning of the Apr-July target season and the use of different rescaling parameters.

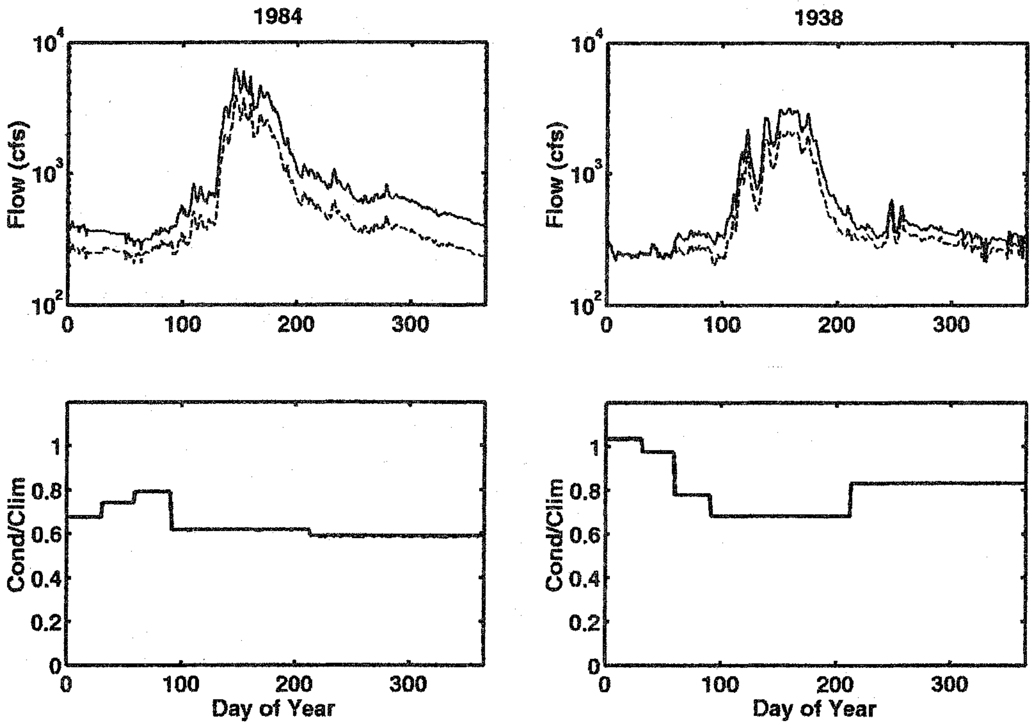

Figure 3. Individual ensemble members conditioned on the February 1st 2003 forecast (1984, left, 1938 right). The top subpanels display the daily flow values from the observed climatological flow (solid) and conditional flow (dashed) versus the day of the calendar year. In both examples, the forecast flow is less than the historical flows. The bottom panels display the daily conditional flow value divided by the climatological flow. The bottom plots represent the rescaling factors derived in the previous section. The stair-step behavior reflects separate rescaling parameters being used for the different seasons of the year (i.e. Jan, Feb, Mar, Apr-July, Aug-Dec).

A variety of hydrograph characteristics can be derived from the ensemble of traces. For example, a user may be interested in the expected peak flow value over the season. This information can be used for river rafting purposes or in planning supplemental releases for the environment. Figure 4a shows the probability of exceedance of the climatological annual maximum peak flow (solid), and the peak flow distribution derived from the ensemble forecast above (dashed). Not surprisingly, given the less than average seasonal forecast, the peak flow is expected to be low. The median of the forecast distribution is 2,150 cfs, compared to the climatological median peak of 3,000. For comparison, the Colorado Basin River Forecast Center issues a peak flow forecast for the White River near Meeker on the first of the month from March to June. Their March 1st 2003 forecast predicted a 50% chance of the peak flow exceeding 1,600 cfs. The observed peak flow was 3,820 cfs on June 2nd after unusually warm temperatures dramatically increased runoff efficiency during an otherwise low streamflow season.

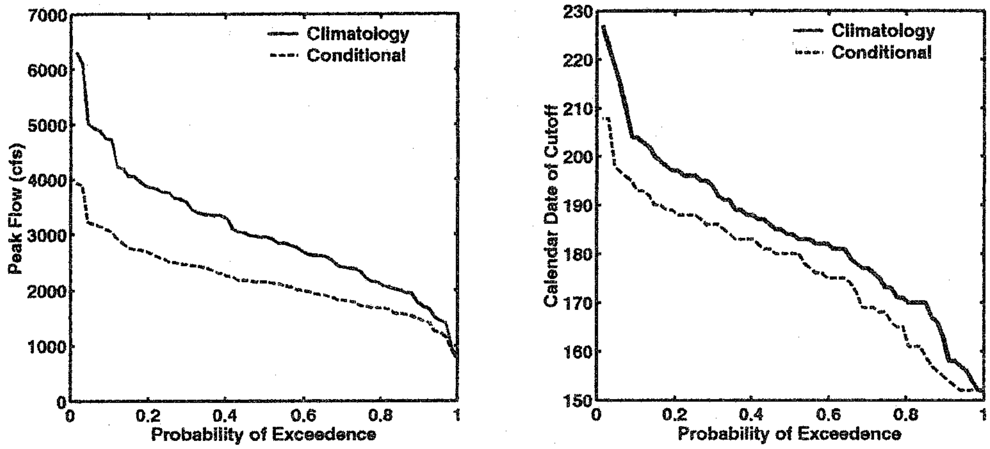

Figure 4a,b. Conditional and climatological distributions of peak flow (left) and calendar date of the first flow less than 750 cfs after June 1st (right). Conditional distributions are based on the rescaling of daily flows according to the February 1st 2003 seasonal water supply outlook. See text for discussion.

Users are also interested in the date that the streamflow may fall below a particular threshold, for example 750 cfs. Such a threshold may activate a cutoff of water rights for a junior irrigator. If this cutoff occurs earlier than usual, a farmer may determine that there will not be sufficient water to irrigate. While the technique described above does not explicitly shift the hydrographs of individual traces forwards or backwards in time or alter their overall shape, the downward rescaling of the hydrographs does suggest that the cutoff will occur earlier. Figure 4b shows the probability of exceedance of the first date after June 1st that flow at the White River near Meeker will fall below 750 cfs. The median climatological cutoff date is 184 (July 3rd), and the conditional median is 180 (June 29th). There is considerable uncertainty in this conditional distribution, but in this case, the observed cutoff date in 2003 (June 30th, day 181) was near the median.

STRENGTHS OF PROPOSED TECHNIQUE

This procedure is simple to use and has trivial data and computational requirements. Calibration of the various parameters is completely automatic, and no manual calibration or intervention is required. As the daily hydrographs evolve from direct rescaling of historical flows, the overall shape and character of the hydrograph bears a strong resemblance to actual streamflow behavior. In comparison, a simulation model is limited in its ability to reproduce all features of the hydrograph and may introduce undesirable biases.

Unless a weather generator is employed, simulation model ensemble members are limited by the availability and quality of historical daily meteorological data. Typically, a NWS ESP simulation will contain 20-30 ensemble members, from the 1970s onward. If the subset of years used to force the ESP model is exceptional compared to the entire period of record (e.g., many of the wettest years on record occurred in the last 20 years), the forecast conditional distribution will contain this bias. In contrast, the procedure described herein can create as many ensemble members as there are years of historical streamflow data. While this procedure will not work in ungaged basins, some locations have as many as 100 years of historical streamflow data. This tool faces the same difficulties as any simulation model in forecasting regulated (versus natural) flow.

Finally, the seasonal streamflow volume distribution produced by the NWS ESP model is not necessarily consistent with the distribution of the official water supply outlooks. Some water managers are required by law to operate reservoirs on the official outlooks but may want to take advantage of the NWS ESP daily flow values by forcing them into, for example, a reservoir operations model. Rectifying this difference may pose a problem. In comparison, by using the rescaling technique described here, the distribution of flow volumes is *by design* consistent with the official water supply outlooks.

LIMITATIONS OF PROPOSED TECHNIQUE

The limitations of this simple model are rife and manifest. As with all tools, the user must recognize its shortcomings and anticipate situations in which it will produce unrealistic results. This model should not be used in a careless manner if the fate of natural resources is at stake. This section will describe some of the primary limitations of this model.

This tool's primary purpose is to translate the seasonal water supply outlook forecasts into an ensemble of daily streamflow traces. These daily traces, if aggregated into seasonal volumes, will contain no more information about seasonal volumes than the original water supply forecast that was used to create it. This tool should not be used in a diagnostic sense to determine what the seasonal volume should be. It can only be used after that seasonal volume has been determined.

The skill of this ensemble forecasting procedure can only be as good as the seasonal water supply forecast used to create it. Further, the success of this tool hinges on the strong correlation between individual daily flows and the seasonal volume. This assumption is appropriate in snowmelt dominated basins of the interior West but can be less valid in arid or mixed rain-snow basins. This tool has not been tested on basins that have "peaky" streamflow from rainfall events during the winter, such as the Cascades of the Pacific Northwest. In determining if the model is appropriate for the basin, the user is advised to develop a time series plot of the moving window correlation of seasonal flows with 30-day flow volumes. The model is valid during times of the year when this correlation is high.

This procedure is not intended for short term (1 day - 1 month) forecasting. It does not guarantee continuity between the streamflow observed the day the model is initialized and the flow forecasted for the following day. In comparison, a simulation model can be explicitly updated with the most recent observed streamflow. NWS ESP traces will have a very narrow spread into the near future, eventually turning into "spaghetti" in the distant future, as the system becomes chaotic. The statistical ESP time series is uniformly "spaghetti" from the near future to the end of the season. Only a full simulation model would be able accurately to answer a question such as "Will the streamflow peak tomorrow and will that be the last peak for the season?"

Rescaling the daily flow values does not change the shape of the hydrograph or timing of its peaks. If the NWS ESP model is initialized with less snow than was present during the historical simulation, this snow will run out early and the bulk of the streamflow will occur at the beginning of the season. The tail of the season in the conditional simulation will be much drier than the historical simulation. The simple statistical technique described here instead distributes the reduction of water evenly across the season. A simulation model will also be able to create non-linear behavior in the hydrograph that a statistical simulation cannot. For example, in the extreme case of a very dry year with no snow, most NWS ESP simulation traces may simply display a flat hydrograph throughout the rest of the season. The statistical technique will create hydrographs shaped like a snowmelt pulse, but with unreasonably small flow values before and after the pulse. The former result is more physically realistic than the latter. The non-linear rescaling technique described earlier may alleviate some of the problems with hydrograph shape but would still not change runoff timing.

There is a critical difference between the simulation model ESP and the statistical ESP in the interpretation of their forecasts' individual ensemble members. The simulation model answers the question, "Given today's basin states, what would happen if the rest of the season's meteorology were identical to that of a given historical year?" In comparison, once mapped into the forecast distribution and rescaled, the statistical conditional traces are divorced in meaning from the historical flow. For example, consider if 1985 was, meteorologically, an extremely dry winter followed by an extremely wet spring, and 1986 was the opposite, with a wet winter and a dry spring, but their total seasonal flows were almost identical. If initialized with the soil moisture and snowpack at the end of the current winter and forced with these historical traces, the NWS ESP would produce a conditional streamflow trace for 1985 that would be wetter than the trace for 1986. The statistical ESP only rescales based on the total seasonal flow volume and as such would produce conditional traces for 1985 and 1986 that would be almost equal. This subtle difference is important but potentially difficult to explain to users not familiar with the model.

Finally, it is generally recommended that the seasonal water supply outlook error bound distribution take a non-linear form. Occasionally, the lower bound associated with a standard normal distribution can yield negative streamflows. While statistically valid, negative streamflow is physically impossible. When this happens operationally, new lower bounds are arbitrarily developed by the forecaster in an ad-hoc manner. It is difficult, but not impossible, to apply the statistical ensemble procedure to an irregularly shaped water supply outlook distribution. In contrast, a log-transformed seasonal water supply forecast will never produce a negative lower bound, and the problem of negative streamflow will not be encountered.

CONCLUSIONS

In summary, this technique is a simple approximation of streamflow behavior in snowmelt dominated basins. This tool is useful in translating small shifts in the seasonal forecast distribution into daily flows. It is a convenient statistical method for calculating on the fly any number of hydrograph characteristics, such as peak flow, date that the flow will fall below a particular threshold and so on. It is also useful for reservoir operators that seek daily flow guidance that is consistent with the published seasonal volumes of the official water supply outlooks. For example, these daily traces may be used as the input to a reservoir management decision support system. However, the ease and simplicity of this technique are also its faults in that it can fail when faced with unusual hydrologic situations or the shifts in the forecast distribution are extreme. This tool may find use in select locations, among knowledgeable users until the start of the snowmelt season. After the snowpack beings to ripen, the user is strongly encouraged to shift over to a physically based simulation model. Alternately, this tool may be a credible baseline against which to compare the performance and accuracy of more complex simulation models.

LITERATURE CITED

Anderson, E. A., 1973. National Weather Service River Forecast System -- snow accumulation and ablation model. NOAA Technical Memorandum NWS-HYDRO-17.

Burnash, R. J. C., R. L. Ferral, and R. A. McGuire, 1973. A generalized streamflow simulation system -- conceptual modeling for digital computers. Joint Federal-State River Forecast Center (Sacramento, California) and State of California Department of Water Resources.

Day, G. N., 1985. Extended streamflow forecasting using NWSRFS. Journal of Water Resources Planning and Management, 111(2), 157-170.

Garen, D. C., 1992. Improved techniques in regression-based streamflow volume forecasting. Journal of Water Resources Planning and Management, 118(6), 654-670.

Hoshi, K., and S. J. Burges, 1980. Seasonal runoff volumes conditioned on forecasted total runoff volumes. Water Resources Research, 16(6), 1079-1084.

Pei, D., S. J. Burges and J. R. Stedinger, 1987. Runoff volume forecasts conditioned on a total seasonal runoff forecast. Water Resources Research, 23(1), 9-14.

Martinec, J., A. Rango, and E Major, 1983. The Snowmelt Runoff Model (SRM) User's Manual, NASA Reference Publication 1100, Washington DC.

Reese, S, and R. Krzysztofowicz, 1989. Stochastic disaggregation of seasonal runoff into monthly runoffs. Report to USDA Soil Conservation Service West National Technical Center, Portland, OR, from Department of Systems Engineering, University of Virginia, Charlottesville. 153 pp.

Appendix: Flow diagram of the statistical hydrograph adjustment procedure.