Biometrics Overview of VDYP7 Components

James W. Flewelling, PhD, CF®

February 26, 2007

Prepared for the Forest Analysis and Inventory Branch, Ministry of Forests, British Columbia under contract1070-20/FLEW 07 032. James Flewelling is a consultant located in Seattle, Washington.

Introduction

This document provides an overview of the VDYP7 system from a functional biometrics perspective. From that perspective there are a number of major components to VDYP7; some are empirical and others are algorithmic. All the major components are described here. The purpose of the document is to provide the user with an understanding of what the major components of VDYP7 do, and how they interact. This will be useful in understanding how system outputs relate to the inputs, and why the system behavior for VDYP7 may sometimes be quite different from that of VDYP6.

The VDYP7 software embodies all the components described here and more. However, the user interfaces for the software do not always refer directly to the model components. This is because a single user action may invoke several different components, and the choice of components may depend on the attributes of the stands being processed. Based on this introduction, the user should be able to determine which model components are actually being invoked for the calculation of current or future yield for any given stand.

This document is not intended as an ongoing technical reference to the software. In a few instances, specific values or parameters are stated; such values could be modified in the release code or in various parameter files. Exact specifications should be sought in the VDYP7 overview and user's guides.

Conceptual Basis of VDYP7

The major external factor that lead to the decision to develop VDYP7 as a replacement for VDYP6 was the expectation that more on-the-ground inventories would be conducted, either for individual stands, or for populations of stands on a sampling basis. The key difference between VDYP6 and VDYP7 is that VDYP6 is primarily a yield model, and VDYP7 is primarily a growth model. This is not a black and white distinction in that the software allows VDYP6 to sometimes appears to be a growth model, and VDYP7 includes some components which are in fact yield functions. The distinction is that the projections within VDYP7 are from growth equations, and not from the reapplication of a yield function to an incremented age.

The VDYP6 software uses a modified version of the "air" projection described by Smith (1992). The essential features of that software are the yield equations which predict various yield statistics as functions of age, crown closure, species, and attributes related to geographic location and site index. The primary purpose of those equations is to predict volume yields in several merchantability categories as function of stand attributes which are generally derived from photo interpretation. The empirical functions were chosen and fitted in such a way that future yield predictions for a stand would be greater than current yield predictions by plausible amounts. However the functions do not take into account the likely changes over time in species mixture or crown closure. Furthermore, there is no clean mechanism to introduce current measured yields into the system. An earlier version of VDYP6 did have a “ground” projection mechanism which used basal area as an input and did not use crown closure.

To the extent that the input stand descriptions to VDYP7 are similar to those provided to VDYP6, the two systems must have similar capabilities for predicting current stand yield. In that situation, the two systems are functionally equivalent. The equation forms differ somewhat, and VDYP7 has somewhat more complex methods of dealing with species mixtures, and provides more detail for each predicted stand.

The major difference in VDYP7 is that once the current stand conditions are predicted, there is considerable room for the current predictions to be modified before moving on to the prediction of increment. In the extreme case, there is the option of replacing the predicted initial yield conditions with cruise summaries for individual stands. The more common case is where the vegetation resource inventory (VRI) sampling process provides a sample of all the stands within a certain population, usually an administrative unit. Through various adjustment processes such a sample can be used to modify the predicted initial yields of all the stands. Depending on exactly how this is done, the resultant estimated yields, summed over all the stands, can be unbiased estimates of the true yields in the population. That statement can be correct or approximately correct for several yield statistics such as volumes to certain utilization standards. Estimates for individual stands will not be unbiased.

The growth predictions themselves are made with growth equations which rely in part on yield equations and a slow approach towards the predictions from the yield equations. The precision with which growth is predicted is lower than would be expected for growth models applicable to single-species single-age stands growing in a limited range of site conditions. The high variability in observed growth rates, the large number of geographic regions, and the large number of species combinations are not the material from which precise growth models are made. Furthermore, a decision to force yield predictions at older ages to approach limits similar to those in VDYP6 occasionally limited the impact of some of the available growth data. Both VDYP6 and VDYP7 are expected to be extrapolated beyond the age ranges of the growth data; both use ad hoc limiting procedures at old ages; neither model attempts to deal with the ultimate breakup and replacement of stands.

VDYP7 Components

The various VDYP7 components all work towards creating or modifying VDYP7 stand descriptions. These stand descriptions consist of a specific set of variables, most of which are required to be known. In aggregate they comprise everything that is predicted about a stand at any one point in time. Yield statistics are on a per hectare basis; stand size is tracked by the VDYP7 software, but it is not relevant to the functioning of most of the VDYP7 components. The stand descriptions are most relevant for stands where the dominant trees have DBH’s in excess of 7.5 cm; this limitation arises from the lower DBH measurement limit is most of the permanent plot data.

The FIPSTART process accepts photo interpreted stand descriptions from the legacy FIP inventory standards. VRISTART accepts photo interpreted stand descriptions from the VRI standards. The primary difference between these standards is that the FIP standards include crown closure percentage, and the volume distribution between species, whereas the VRI provides basal area and basal area by species. VRIYOUNG is a subcomponent of VRISTART; it accepts FIP or VRI stand descriptions at young ages, and advances them to a condition where the stand meets some minimal criteria; those criteria usually include a site height of at least 6.0 m, breast height age of 5 years, and a basal area of 2.0 m^2/ha.

A sampling adjustment process, together with VRIADJST, provide a mechanism which allows the VDYP7 stand descriptions for all the stands in an administrative unit to be adjusted to better reflect the results of a current VRI sample within the administrative unit.

The VDYP7 growth process projects a VDYP7 stand description from the current year to some future year. It does this through a series of annual growth increments. The primary driving equations are for increments in site height, basal area and quadratic mean diameter (DQ). A secondary series of equations provides for updated estimates of basal area and TPH by species and utilization class. Similarly, all of the volume statistics are updated to reflect the increments in basal area, DQ and height. Species percentages of basal area do not change during a projection using the recommended default parameter settings.

VDYPBACK is a program that accepts a VDYP7 stand description, estimates the condition of the primary layer at some earlier age, and interpolates at intermediate ages. The earlier age is chosen such that the site curves will indicate a site height of about 9 m. No estimates are made for the history of the overstory layer should there be one. Long term planning processes used by the Ministry with VDYP6 had required yield curves for each stand. VDYPBACK is the mechanism within VDYP7 designed to create plausible yield curves which include the current stand condition. These curves allow the present planning mechanism to be continued with minimal change. However, the curves derived in this way are far from perfect; there may be better ways to estimate yield curves for future rotations.

VDYP7 Stand Descriptions

Every VDYP7 stand description consists of a fixed set of predefined descriptors. These are at a stand level, layer by species level, and utilization level. At the stand level, the relevant variables are YEAR, biogeoclimatic zones (BEC) and forest inventory zone (FIZ), and percent forest land (PCTFLAND). YEAR is an integer; in model development YEAR refers to the most recently completed growing season; thus 2006 would refer to a stand's condition after the 2006 growing season and before the 2007 growing season. BEC refers to one of the fourteen biogeoclimatic zones of British Columbia. Several newer BEC's are defined in Volume 1 - VDYP7 Overview (Table 7.1.10); stands with new BEC's generally have a corresponding original BEC used as a substitute for all computations. FIZ codes are single-character codes from a zoning system that predates BEC. PCTFLAND is a variable that allows for conversion between per-hectare yields on entire stands, and the per hectare yields on the forested portion of the same stands.

Layers. Stands may have multiple layers. The significant layer for which growth is modeled is the primary (P) layer. Additionally there may be an overstory (V) layer, which can include "veteran" layers in the FIP standard. The overstory layer is estimated, but is not grown. A layer must have associated with it one or more species groups (SP0). The species groupings and two-letter codes taken from earlier inventory documents are shown in Vol. 1 VDYP7 Overview, section 6. At least one of the SP0's must have a site height (HD), site index, total age, breast-height age, and years to breast height. The SP0 with the greatest basal area is considered the primary SP0. Each SP0 can be associated with one or more actual species codes, and percentages of those species. The percentages are primarily a book-keeping feature; their only significant influence comes about when the choice of site curve number (SCN) can vary with species within the primary SP0 group.

Utilization level. There are yield statistics for various utilization levels within each SP0 by layer combination. The utilization levels are based on DBH recorded in centimeters. The classes are defined as:

-1	4.0-7.49
0	≥ 7.5
1	7.5 -12.49
2	12.5 - 17.49
3	17.5-22.49
4	≥ 22.5

Hence utilization class 0 is the sum of classes 1 through 4. In this document, references to variables are for utilization class 0 unless otherwise stated. Class -1 is predicted, but the growth models make no use of yields in this class; most plot data used in fitting the growth equations were collected with protocols which excluded trees having DBH < 7.5 cm.

The statistics associated with each utilization level, for each layer and SP0, are basal area per hectare (BA), trees per hectare (TPH), and five cubic volume statistics. The latter are named in Vol. 3 VDYP7 console interface guide as whole stem volume (Vws), close utilization volume (Vcu), close utilization volume net decay (Vd), close utilization volume net of decay and waste (Vdw), and close utilization volume net of decay, waste and breakage (Vdwb). Additionally, Lorey height (HL) is computed for utilization classes -1 and 0. Lorey heights are weighted means of total tree heights, using tree basal area as the weight.

FIPSTART

This component generates a full VDYP7 stand description from the photo interpretations made using the FIP standard. That standard allows for an overstory layer and a primary layer. Each layer has a height, total age, site index, identified site species, crown closure, and percent volume by species group (SP0). The principal empirical function within FIPSTART is one that predicts basal area as a function of crown closure (CC), site height (HDOM), and inventory type group (ITG). The latter is a classification scheme applicable to the FIP standard; it defines some 42 groupings of primary and secondary species. Other empirical relationships predict the relative heights and quadratic mean DBH's of the species groups. Static relationships by species group relate volume to basal area, quadratic mean diameter (DQ), and Lorey height. Similarly there is a defined relationship between site height and Lorey height for the dominant species.

The primary computation task of FIPSTART is to allocate the basal area by species group such that the volume percentages on input are retained. For the primary layer, this is a root finding procedure which iteratively modifies the basal area percentages. For each trial set of percentages, volumes by species are calculated using the relationships cited above. The basal area percentages are adjusted until the resultant volume estimates agree with the percentage volume estimates from the photo interpreted inventory. In most cases, exact agreement is obtained.

In a limited number of situations, exact agreement with the volume percentages on input is not achieved. This situation occurs most frequently when the principal species groups have identical volume percentages on input. The cause of the problem is that several of the empirical relationships have been fit in groups based on lead species. Lead species can be different between the FIP standard and the VDYP7 standard. In the later, the lead species group is the one with the most basal area. Thus during the root finding process, the lead species may change for some of the relationships; the resultant discontinuities in the relationships pose numeric difficulties which are not always resolved. The photo interpreters assignments of species percentages should not be expected to have high precision; hence an occasional failure to exactly replicate those percentages is not viewed as a serious flaw in the prediction process.

The overstory layer, if any, is estimated by a simpler procedure. The species percentages from the input files are applied directly to the estimated basal area. Because volume to basal area ratios are usually predicted to differ between species groups, it is unlikely that the resultant volume predictions will be in accord with the input percentages. The recorded overstory components are usually of a single species; hence this potential problem occurs infrequently. The overstory computations are performed prior to those for the primary layer. The basal area of the overstory layer has a small negative affect upon that predicted for the primary layer.

VRISTART

VRISTART accepts the information collected as part of the VRI photo-interpretation process, and creates a VDYP7 stand description. The major difference between FIP input and VRI input is that the latter usually includes estimates for basal area and TPH. Furthermore, the species percentages on input apply to basal area, not volume. Hence the full complexity of the root finding procedure described for FIPSTART is not required.

A critical feature of this component is that the input basal area is assumed to be for all live trees with DBH > 7.5 cm. This interpretation differs from the VRI specifications, which do not use a specified diameter limit in defining basal area or TPH. This disconnect between the interpretation and use of basal area is addressed to some extent in the sampling adjustment process for primary statistics as described in a subsequent section of this document.

Species and species groups are handled in much the same way as in FIPSTART. For most purposes, species are grouped into the SP0 classes; actual species are tracked, but the only substantive use of species is in the selection of the site curve number (SCN) for SiteTools in cases where this is not already present on the input files.

VRISTART has several modes of operation; for a given stand only one mode is selected to be applied. The selection is based on the values of the input variables. All modes require that site index be available. The standard input mode requires basal area, TPH, site height and AGE; crown closure is not used. The VRIYOUNG input mode is described in a separate section of this document; it is applicable to young stands, and uses age or site height as its principal quantitative input; basal area, TPH and crown closure are not used. A special “MINIMAL” mode is available for stands with specified site height, and unspecified basal area and crown closure. Additionally there is a “crown closure” mode which requires site height and crown closure on input, but not basal area. It is the standard input mode which is being described in the following paragraphs.

The principal computational processes for the primary layer are:

1. Lorey height (HL) of the lead species is estimated as an empirical function of site height.
2. Estimate Lorey height of other species as functions of site height or as functions of Lorey height of the dominant species.
3. Assign DQ by species. This requires a root finding routine to allocate the total TPH to the various SP0. Empirical relationships between DQ of individual species and the lead species are invoked in the process. Limits on DQ by species are enforced during the root finding process.
4. The small utilization class, for DBH’s from 4 to 7.5 cm., have estimates made by SP0 for Lorey height, basal area, DQ and volume. These are estimated as empirical functions of the predicted yields for utilization class 0 (DBH $\geq$ 7.5 cm).
5. The basal area and TPH’s in utilization class 0 are apportioned between utilization classes 1 through 4.
6. Whole stem volumes are calculated by SP0 and utilization class. The four other types of volumes in the VDYP7 stand descriptions are predicted as empirical functions of the whole stem volumes and other relevant variables, the most important of which is usually the DQ of the utilization class.

A similar process is followed for the overstory layer if one is present. All trees in the overstory layer are assumed to have DBH $\geq$ 22.5 cm.

VRIYOUNG

This is a sub-component of VRISTART. The reason for having this startup mode is that the VDYP7 growth model depends on having basal area and TPH as known input parameters. These statistics are defined for utilization class 0, which includes all trees with DBH $\geq$ 7.5 cm. For young stands, the yield statistics are often zero, and give no useful indication of stocking or mean tree size. Hence, if a full VDYP7 stand description is made at some young age, the values in it will be mainly zeros, and it will not provide a useful starting point for the growth components. Accordingly, VRIYOUNG does not attempt to make a VDYP7 stand description for the current age. Instead it advances the age however many years are required so that a predicted stand would have trees large enough to make a VDYP7 stand description meaningful.

VRISTART makes a determination as to whether VRIYOUNG should be invoked. Generally young stands without a specified basal area qualify. Age, site index, and site curve number are the primary inputs. These allow site height to be determined at any age. The current age and a series of annual increments are to be examined. At each age, site height is determined from the site curve, and basal area is predicted with an empirical equation. At each age, the predicted HDOM and BA are compared against pre-set minimum targets. The first age where the targets are met becomes the age at which a full VDYP7 description is made. The targets are breast height age 6, site height 6 m, and basal area of about 2.0 m^2/ha.

At extremely low values of site indices it is possible that the target conditions will not be met, even if the age is advanced for lengthy periods. The VDYP7 software will stop seeking a solution after advancing for a fixed number of years, currently eighty. No VDYP7 descriptions will be available for these stands.

For stands which do meet the target conditions, the following computations are performed:

1. YEAR and ages are updated.
2. Site height of the lead species is updated using the site curves.
3. Input heights of non-lead species are discarded.
4. Total basal area is estimated as an empirical function of age, and site height. Coefficients of the relationship vary with BEC and inventory type group (ITG). The latter is determined by the species mixture indicated on the input description for the stand.
5. DQ is predicted from an empirical function of age and site height. Coefficients depend on BEC and ITG. TPH is calculated from BA and DQ; all three of these statistics refer to trees with DBH ≥ 7.5 cm.

At this point, all the input attributes required for a VRISTART's stand input mode are present. Accordingly, VDYP7 executes the standard VRISTART procedure. The species percentages supplied on the input file are used.

As an aside, the MINIMAL startup mode for VRISTART is an abbreviated version of the above, starting at step 4.

Static Volume Functions

Cubic volumes, defined in several different ways, are calculated as functions of several variables, primarily basal area, Lorey height and DQ. Coefficients of the relationships may vary with utilization class and BEC. Age is not a factor for whole stem volume, though it does affect some calculations for decay and breakage. The static volume equations, particularly those for whole-stem volume are robust and generally precise. In cases where the user questions the veracity of the predicted volumes, one or more of the input variables to the function will usually be suspect.

The series of volume functions starts with utilization class 0 for a species group. An empirical equation predicts the mean volume per tree as a function of DQ and HL. Coefficients vary with BEC and SP0. There are 74 sets of coefficients; these correspond to the 74 different BEC-based taper equations fit by Kozak(1994). The whole-stem volumes utilization classes 1 through 4 are then estimated as products of class basal areas and empirical estimates of the volume to basal area ratio; the sum of the volumes for class 1 through 4 are reconciled to force their agreement with the volumes calculated for class 0. Close utilization volumes by class are estimated as the product of empirical ratio functions and the estimated whole stem volumes. The same 74 groupings of BEC-SP0 are used for all the cited functions.

Similarly, empirical functions are used to predict the amount of decay and waste in each utilization class. The fraction of the close utilization volume which is predicted to be decay is a function of DQ and breast height age of the primary species. Coefficients of the relationship vary by decay group (DGRP). The latter are groupings of SP0 and BEC for which individual tree decay predictions equations were available. A waste prediction function depends mainly on the predicted decay; coefficients vary by SP0, and differ slightly between utilization class 4 and the three small-diameter classes. The waste function is the same for all BEC's.

Breakage is estimated as a function of DQ; coefficients vary by breakage groups (BGRP). The latter are a set of 34 combinations of SP0 and BEC. The breakage predictions are a close approximation of the results which would be obtained by applying tree level breakage tables by BEC as shown in tables prepared by the Ministry. The BEC groupings simply separate the coastal BEC's (CDF, CWH and MH) from the others.

The VDYP7 programs have mechanisms which allow for modifications of the predictions for decay, waste and breakage based on user-provided parameters.

Sampling Adjustment Process for Primary Statistics

The vegetation resource Inventory (VRI) was designed to allow for the unbiased estimation of many different quantities at some aggregate levels, such as that of a Tree Farm License (TFL). Typically this would start by using any available method to estimate quantities such as basal area and volume for each stand. VDYP6 has been the principal model used for that purpose for natural stands; VDYP7 is now available as a replacement. Regardless of where the starting estimates come from, a sampling estimator can be used to calculate an unbiased estimate of the totals for a population; typically the sampling estimator would be the product of the original estimate for the population total and a ratio of means, the later being derived from data on a set of randomly selected plots.

The sampling estimator is for totals for the population. Its use does not dictate how estimates should be made for the individual stands in the population. If a ratio estimator has been used for the total, it is common practice to derive a revised estimator for all the stands as the product of the original estimator and the sample ratio. This ensures that the revised stand estimates sum to the revised total estimate. However, other methods of distributing the revised totals back to the stands might be better in particular situations.

Regardless of the method of assigning revised yields to the individual stands, there is the potential for implausible or impossible predictions at the stand level. For example, consider that ratio estimators are calculated for basal area and TPH, and that the same ratios are applied to all stands. It could happen that the ratio for basal area is less than one, and the ratio for TPH is greater than one. Recall that VDYP7 defines basal area and TPH for the portion of the stand with DBH ≥ 7.5 cm. For a stand where the starting estimator for DQ is only fractionally greater than 7.5 cm, the adjustment procedure described here would be inadvisable. If the basal area is decreased and TPH is increased through the application of sample ratios, the resultant DQ may be less than 7.5 cm. The VDYP7 growth program will detect that the stand description is impossible, and it will refuse to grow the stand. This potential problem can be overcome with some fairly simple techniques such as eliminating the stand-level modifier for TPH, and instead applying a multiplier to the quantity (DQ – 7.5).

If all of the elements of a VDYP7 stand description were adjusted using ratio estimators, implausible or impossible stand descriptions would be expected to be a frequent occurrences. Accordingly, a multipart adjustment strategy is undertaken. First, several critical inputs for the stands to be processed by VRISTART are modified to be in accord with the VRI sample results. The input parameters being modified may include age, site height, site index, and basal area and TPH for utilization class 0. If sample-based changes in species composition are allowed, they occur at this stage. When the VRI ground sampling occurs in a growth year subsequent to that of the photo interpretation, multiple processes must be forced to occur to update the stand from the date of photo interpretation to the sampling year. Those processes are described in Volume 4 – VDYP7Batch Interface Guide. At the year of sampling, a set of VRISTART inputs are created for each stand; those inputs will have been adjusted for the sampling results. VRISTART is then run for all of the stands, creating VDYP7 stand descriptions.

Additional adjustment processes are still required. The reason for the additional processes is that without them, the only yield statistics which can be considered unbiased for the administrative unit are those which were directly adjusted. They include basal area at utilization zero, but do not include basal area at any other utilization level or any volume statistics. Accordingly, two more stages of the overall adjustment process are required. One of these deals with a set of seven important yield statistics, as shown in Section 3.9.1 of Volume 4 – VDYP7Batch Interface Guide:

1: HL (Lorey height, all trees with DBH > 7.5 cm).
2: WSV7.5, Whole-stem volume, all trees with DBH > 7.5 cm
3: BA12.5, Basal area, all trees with DBH > 12.5 cm
4: WSV12.5, Whole stem volume, all trees with DBH > 12.5 cm
5: V_CU12.5: Close utilization volume, all trees with DBH > 12.5 cm
6: V_D12.5: Close U volume, less decay, trees with DBH > 12.5 cm
7: V_DW12.5: Close U volume, less decay and waste, DBH > 12.5 cm

These are all-species aggregates, with no opportunity for separate values by SP0. Process documentation indicates whether or not separate adjustments by layer are currently allowed.

The final stage of the adjustment process is a VDYP7 process named VRIADJST. It accepts as input the values of the above seven variables, one stand at a time, together with the VDYP7 stand descriptions from VRISTART. For each stand, the full set of variables within the VDYP7 stand descriptions are modified to be in accord with the seven selected yield variables. The choice of these seven variables was subjective. The VRIADJST could have been written to deal with a different set of input variables. The variables chosen were those which Ministry personnel thought were of high importance both biologically and economically, and amenable to estimation in a robust fashion from VRI sample data.

The overall adjustment process that has been programmed is one out of many possible such processes. For some of the yield variables it relies upon regression models; the resultant estimates of population totals may differ from what a sample-theory ratio estimator would predict.

VRIADJST

The VRIADJST process is an algorithm designed to accept as input a VDYP7 stand description and a specific set of revisions to the all-species aggregated yields; the output is a revised VDYP7 stand description that reflects the specified revisions, and has all other attributes modified to be in accord with those revisions. The revised statistics are the specific set of seven variables shown in the previous section. The revised statistics may be for the primary layer only, for both the primary layer and overstory layer separately, or for the combined yields of both layers.

VRIADJST respects the integrity of some of the input variables on the VDYP7 stand description which it processes. These not-to-be-altered variables include age, site index, yields for utilization class -1, and BA and TPH for utilization class 0 (all trees with DBH ≥ 7.5 cm). The latter two variables are respected overall, and by SP0. The processes to adjust all the statistics for all the utilization layers involve an ad hoc series of adjustment equations, and some root finding routines. On occasion, the VRIADJST program will fail to meet all of its design objectives. Failure, when it does occur, is most often that TPH's by species had to be altered slightly. Additionally, there may be occasional small departures, less than one percent, from some of the specified revisions to the yields. Typically, the error in BA12.5 is less than 0.01 percent.

VDYP7 Growth

The VDYP7 growth model advances a VDYP7 stand description from one age to another through a series of annual increments. The primary equations in the increment process are for increments in site height, basal area and DQ. Other equations deals with Lorey height and the division of increments between species and utilization classes. Overstory layers, if any, are not incremented for variables other than age. The configuration files play an important part in the process. They force the selection of the recommended models, force model behavior at old ages, establish limits on yield and mean tree size, and control a critical interaction between a growth equation and a yield equation. They also control how several “compatibility variables” change over time. These variables allow the observed relationships between variables on the input VDYP7 stand descriptions to be reflected on the output VDYP7 stand descriptions. This mechanism allows individual stands to have non-normal relationships between variables; for example the proportion of a stand’s basal area due to trees in the highest utilization class (class 4: DBH ≥ 22.5 cm) may differ from that predicted by the static equations internal to the VDYP7 programs. The following discussion is a brief overview of the annual growth computations implemented with standard configuration files.

(1) The site species, current site height and site index are determined; the site curves are used to determine the expected annual increment at the current height. The grown height is calculated as the starting height plus the expected increment. Limits in the configuration file can affect this and many other processes. For example, every site curve may have a maximum age assigned to it; in that case, the site curve is effectively modified to become a flat line at the indicated age.

(2) Basal area growth, aggregated over all species, is predicted as a weighted average of two different growth estimates. The first growth estimate, the “fiat” estimate, is obtained as a difference between yield model values at two points, with a small adjustment which would eventually force the predicted yield to approach that of the yield model. The yield model is a function of site height, age and overstory basal area. The coefficients of the yield model vary with BEC and the percentages of basal area by SP0. The second growth estimate is an empirical function of site height, breast height age, and the unadjusted increment inferred from the yield model. The coefficients of the yield model and the growth model vary with BEC and the percentages of basal area by SP0. The combined estimate of annual increment is some fraction of the empirical estimate, plus the product one minus that fraction and the fiat growth estimate. The assignment of the fraction is age dependent, with controls referenced in the configuration files. Prior to a specified age, for example 80 years, growth is entirely from the empirical model. There is a transition period of about sixty years where the controlling fraction continually changes. Following the transition, the controlling fraction is zero, forcing the growth prediction to come entirely from the fiat model. The reason for this mixed approach to growth is that the empirical model is not always credible at the older ages, where there was often little or no growth data.

(3) Growth in DQ is predicted analogously to the growth of basal area. There is a fiat yield model, and empirical growth model and a weighting function which varies with age. Together the growth in BA and DQ determine the change in TPH.

(4) Update Lorey height (HL) for each species using static relationships and compatibility variables. The static relationships are simply the empirical equations which predict some yield variables as a functions of others. For example, the Lorey height of the lead species is a function of site height and TPH of the primary species. The associated compatibility variable at the start of the one-year increment period is a measure of the difference between the input HL and that predicted by the static equation; here the compatibility variable is defined as the difference in HL divided by the static prediction. The predicted compatibility variable at the end of the one year

period is forced to be slightly smaller in absolute value than at the start; the rate of decrease is controlled by coefficients indicated in the configuration file. The ending value for the primary species HL is derived from the static equation and the compatibility variable; the computational equation is the inverse of that used to define the compatibility variable. Increment in HL for the nonprimary species follows a similar course. The static equations to predict those HL's are functions of either site height or the Lorey height of the lead species. Which of the latter is used as the independent variable is dependent upon the species combination. The coefficients in the static relationships vary with species combination and an interior/coastal split.

(5) Update BA and TPH by species. The overall change in TPH is known from BA and DQ. The allocation to species being described here may be referred to as partial stand dynamics. The percentage distributions of basal area by SP0 are not allowed to change. The percentage distributions of TPH do change. The changes in the percentage of TPH allocated to each species are not modeled directly, but instead are determined through a prediction process for DQ by species. Species whose Lorey heights tend to be higher than that of other species in the same stand also tend to have the greater DQ's; this tendency is modeled with static prediction equations for each species. The coefficients of the static relationship vary by subject species, and the mixture of the other species in the stand. A reconciliation process adjusts the static predictions so as to achieve the correct total TPH. Another process focuses on how the relationship between actual DQ and the static DQ prediction should change during the one year growth period. This involves setting a change parameter for each the species present. No change is the preferred solution. However, various constraints usually invalidate that solution. Other constraints are based on the observed upper and lower limits on the DQ to HL ratios by species, with separate bounds for the coast and interior. In summary, this step involves the empirical static relationships between HL and DQ by species, then follows algorithmic processes to find a solution that violates no boundary conditions, or violates them to the least extent possible. Overall, reasonable relationships between DQ and HL tend to be maintained over time. Trends in change in TPH by species are an outcome of the process, and are not always reasonable.

(6) Compatibility variables are calculated for basal area and DQ of all the utilization classes for all species. Compatibility variables are also calculated for all of the volume variables. The compatibility variables are decreased in absolute terms using coefficients referenced in the configuration files. The revised compatibility variables and the static relationships for utilization class and volumes are combined to estimate all the yield components for all utilization classes for all species. Hence over time the relationships between the utilization classes tends to approach the empirically fitted relationships.

In fitting the growth model to permanent sample plot (PSP) data, there was a fair amount of grouping the data by BEC; also species are often combined. The BEC groupings are (CWH, CDF, MS), ICH, (PP, BG, IDF), (ESSF, MS), SBS, (AT, SWB, BWSS). The species groupings are (AC, D, MB), B, C, (E, AT), F, H, L, S, (PW, PY), (PA, PL). The data used for fitting were from 6,860 growth periods on individual plots. These included only the plots with the highest quality of data. Some 257 growth periods had been excluded due to high mortality criteria developed by the Ministry. The data are assumed to represent fully stocked conditions; temporary plots (TSP's) which are more representative of the inventory tend to have about fifteen percent less basal area than the PSP's. Overall, the model explains some 65% of the variation in growth rates observed in the PSP's.

VDYPBACK

The main purpose of this component is to make plausible estimates of merchantable yields at ages younger than the current age. VDYPBACK does not "grow" backwards. Instead the site curves and other functions are used to find the youngest possible "convergence age". That is a young age at which a reasonable estimate of stand conditions can be made, and at which all compatibility variables can be assumed to be zero. The convergence age is determined by procedures similar to those described for VRIYOUNG. The principal difference is that site height is required to be at least 9 m. Site species and site index are the main determinants of the convergence age.

At all ages from the convergence age to the current age, the yield functions alluded to within the VDYP7 growth function are applied to estimate a model BA and a model DQ. At the current age, the actual BA and DQ are of course known. These are used to compute a pair of factors which represent the ratio of achieved growth to that predicted by the yield models. Growth in this context is the difference between yields at the current age and the convergence age. The two derived factors are known to be valid at the current year and are assumed to be valid at all intermediate years. Species percents in basal area are assumed to be constant. Other assumptions allow the full complement of variables within the VDYP7 stand description to be predicted for all ages.

The fact that the VDYPBACK process does not use the growth equations which are within the VDYP7 growth model can produce some anomalies. One of these is that the inferred growth in the year prior to the current age may on occasion be considerably different from the inferred growth rate in the following year.

Discussion

Current yields can be estimated directly from photo interpreted attributes collected to the VRI standard. However directly inferred attributes such as site height may be biased. Furthermore the directly inferred attributes of basal area and TPH are not field verifiable. Data collected with the ground-based sampling protocols of the VRI are necessary to calibrate the yield predictions for any administrative unit or other population. The adjustment mechanisms developed as part of the VDYP7 system can provide the calibration mechanism. The adjustments are critical. Without adjustments, the VRISTART model component can not be assumed to be correct. In particular, the VRISTART component does not have any empirical equations to predict basal area to some DBH standard as a function of the photo interpreted basal area. That task must be handled as a sample-based adjustment.

Growth predictions are from heavily constrained growth equations. The use of an explicit growth equation has two major advantages over the VDYP6 approach of inferring growth from yield curves. First, for short term projections of individual stands, the relationships between various species, between utilization classes, and between the various measures of volume are maintained. For example, if a stand has an unusually high number of trees in the largest utilization class (DBH $\geq$ 22.5 cm), that condition is extrapolated into the future. The mechanism by which this is accomplished, through compatibility variables, is course, and will lack some of the internal consistencies which would be expected of an individual tree growth model. The second advantage of growth equations is that they take advantage of the starting distribution of species through more than a simple averaging process.

A weakness of the growth model is that it does not fully capture species dynamics. In common species mixtures such as spruce - aspen, the spruce component will typically increase from a very low percentage to a very high percentage. Though the data may have been adequate to model such species change in a few situations, the data or the attempted models were inadequate in other situations. Hence the default model parameters are set so as to force the species distribution by basal area to remain constant during the course of a projection.

The growth model does not deal with stand breakup and replacement, nor with significant insect or disease events. Long term planning requires that these be dealt with outside the VDYP7 growth modeling framework.

Growth predictions for individual stands contain error; this fact is known and accepted by everyone who uses growth models. However, there is sometimes a presumption that the growth models are correct on average. That assumption is not supported by data. Within the natural forests of British Columbia there has not been a comprehensive growth monitoring program; hence the average growth rates for any fixed population of stands is unknown. The growth model is put forward as being a reasonable way to project stands, subject to the limitations discussed here. Differences between the PSP characteristics and those of the overall inventory have been accounted for, but there has not and can not be a true validation in the absence of growth data that is known to be representative.

References

Kozak, A. 1994. Development of taper equations by BEC zones and species. B.C. Ministry of Forests. 25p.

Pojar, J. ; Meidinger, D. 1991. British Columbia: the environmental setting. p 39-68 in: Ecosystems of British Columbia. Special Report Series 6. ed: Del Meidinger and J. Pojar. B.C. Ministry of Forests, 330 p.

Smith, S., 1992. New Models for Yield Estimation and Projection. Interim Revision. Report to MOF. Sterling Wood Group, Victoria.