Development of a Mining Fill Inventory from Multi-date Elevation Data[1]

Michael Shank[2]

Abstract. This paper describes the development of a spatial inventory of mining fills derived from multi-date elevation data. An IFSAR Digital Elevation Model (DEM) for approximately nine counties of Southern West Virginia was acquired in September 2003. This data was compared with a second elevation grid derived from 7.5-minute hypsographic (contour) data for the same area. A simple grid subtraction of the two data sources produced a crude representation of topographic changes occurring during the interim time period. However, error in both datasets complicated the process of isolating, identifying, and delineating fills, and limited the minimum size and depth of fills that could be resolved. In addition, the limited ability of IFSAR data to penetrate tree canopy required the creation of a simple forest/non-forest landcover classification, which was used to mask forested areas from further analysis. Several metrics were evaluated in an attempt to automatically discriminate between actual fills and remaining error artifacts, including variance, distance from cut areas, and drainage characteristics. These metrics were valuable for identifying the most obvious fills and discarding the obvious errors. However, the majority of 'gray zone' fill candidates were classified visually using a variety of image resources. All identified fill outlines were adjusted to approximate a 10-meter depth contour, and any shape anomalies due to misclassification in the forest mask, forested fill faces, and DEM error were removed.

After creating the initial fill inventory, GIS analysis tools were utilized to calculate fill characteristics such as area, volume, maximum depth, and length of buried stream. Preliminary results indicate the presence of over 500 fills that were not represented in an existing fill inventory digitized from permit maps. These previously unmapped fills cover an aggregate area of approximately 9700 acres, and over 138 miles of stream channels, based on a 15-acre minimum drainage area.

Additional Key Words: valley fill, change detection, IFSAR, DEM, Landsat, landcover classification.

_____

[2]Michael Shank is a GIS Database Administrator with West Virginia Environmental Protection, Charleston, WV 25303

I. Introduction

Explosive growth in the availability of geospatial data, along with continuing advances in computing capacity, has facilitated new kinds of analysis that heretofore have been impractical. New technologies for collecting elevation data, notably radar (IFSAR) and laser (LIDAR) based systems, along with diligent conversion of paper-based map assets, has produced a data-rich environment in which it is possible to characterize topographic change over time for a relatively large area, and at a reasonably detailed scale. Concurrent with database developments, computer performance now permits the visualization and analysis of datasets far larger than could be contemplated previously.

Figure 1. Example of a valley fill.

The need to characterize topographic change over time is particularly acute in the coalfields of Southern Appalachia, where the scale of surface mining operations has increased dramatically in the last decade. This expansion has not occurred without controversy, becoming the subject of national media attention, federal lawsuits, and even presidential politics. An environmental impact statement, released as a draft in 2003, brought together more than 30 scientific and technical studies to address this issue. At the center of the controversy is a practice commonly referred to as mountaintop removal, in which entire mountains are removed down to a base coal seam. All

intervening coal seams are recovered and "the additional volume of rock that is often generated…but cannot be returned to the locations from which it was removed…is typically placed in valleys adjacent to the surface mine, resulting in valley fills" (DEIS, 2003). An example of a valley fill is shown in figure 1.

Significant effort has been expended to develop a comprehensive spatial inventory of mining features, including valley fill locations, based on maps of individual permits issued by the West Virginia Department of Environmental Protection (WVDEP). This database has been extremely valuable for characterizing permitting activity. However, for the purpose of developing a comprehensive inventory of fills as they exist on the ground, remote sensing technology has definite advantages. For example, suitable permit maps sometimes are not available for closed mines where regulatory responsibility has ended. In other cases, maps may not reflect the true eventual extent of a fill. Finally, it is difficult to obtain a high level of confidence about which fills have actually been constructed without canvassing, in detail, the entire cadre of inspectors on every permit issued.

This study relates the development of a valley fill inventory derived from two elevation data sources, one that predates most valley fill activity, and a second that was collected recently. The inventory was created by processing the differences between the two datasets. The study builds on an earlier investigation that attempted to detect mining fills using LIDAR data for a single county (Shank, 2002). Based on the results of that study, the investigation was expanded to comprise ten counties of Southern West Virginia that encompass over 86% of known permitted valley fills (figure 2). A slight expansion of this area, to include an additional 8 USGS quadrangles, would capture nearly 95% of known permitted fills, and may be contemplated in the future.
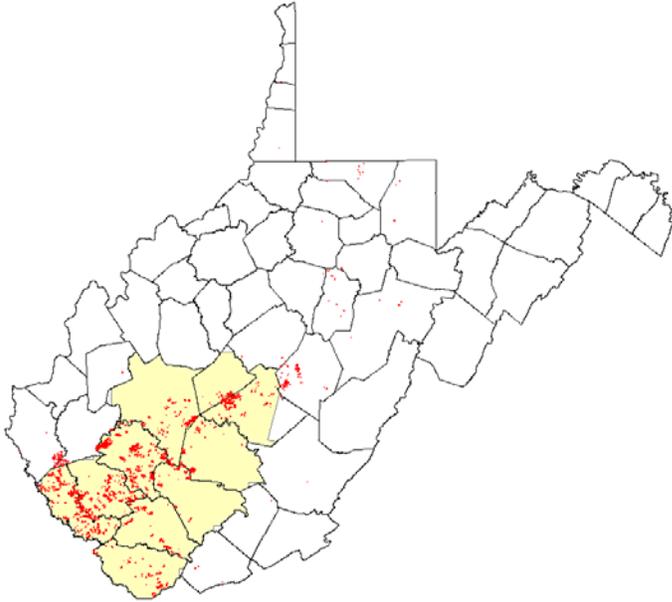
Figure 2. Study area, shown in yellow. Documented valley
fills are shown in red.

II. Data Sources

IFSAR elevation

TAGIS contracted with Intermap Technologies to acquire IFSAR elevation data for the entire study area. Intermap carries a STAR-3i IFSAR sensor on a modified Learjet Model 36. The aircraft typically flies at over 450mph at elevations exceeding 30,000 feet, and the sensor can image a 10km swath of ground in a single pass. These characteristics allow large areas to be collected at comparably low cost. The delivered elevation product is a regular grid of elevations at 5-meter intervals with a root mean square error (RMSE) error of 1.0 meter vertical. An additional image product, resembling a black & white photo, is also produced showing the radar returns. This product has a 1.25-meter pixel size.

Despite large advantages in cost and relatively high accuracy, IFSAR can exhibit characteristics that can affect product quality. These characteristics include layover, shadow, and saturation. *Layover* occurs in mountainous terrain because returns from the tops of mountains

arrive at the sensor ahead of those at the foot (because they are physically closer to the sensor). This creates an effect where mountains appear to lean toward the sensor and obscure detail along nearside slopes. This affect can be corrected, though it can result in data gaps. *Shadow* results from areas on the ground that are not illuminated by the radar pulse, *e.g.* a cliff facing away from the sensor. This represents a data gap, which can sometimes be avoided if another look angle is available from an adjacent pass. *Saturation* occurs when an arriving signal has a higher amplitude than the receiver can record. This is analogous to overexposing an image when using a film or digital camera, causing a loss of highlight detail.

In addition to Digital Surface Models (DSM), which map the elevations of first returns, Intermap delivers Digital Terrain Models (DTM), which attempt to remove buildings, vegetation, and other features in order to produce a "bare earth" model. This accomplished though a proprietary algorithm that works reasonably well when there is at least some return from the earth's surface. However, IFSAR has difficulty penetrating dense forest, and DTM products cannot be relied upon as an accurate representation of the earth's surface in these cases. For example, for an entire quadrangle in a relatively heavily forested area of Wyoming County, the average elevation difference between a LIDAR DTM and an IFSAR DTM was 59 feet.

Because of the tree canopy problem, the analysis was restricted to non-forested landcover types, such as rock, soil, or grassland, for which the radar product was relatively accurate. Reforested cut and fill areas initially could not be accurately recovered, and in several cases fills were not accurately delineated due to reforested faces. However, many times the fills were partially resolved, or identified through adjacent cut features, and were delineated manually.

Hypsography-based elevation

The pre-fill elevation grid was constructed from Digital Line Graph (DLG) hypsography. The DLGs were scanned and digitized from USGS 7.5 minute quadrangles. An examination of source maps indicates that the photography on which the hypsography was based was acquired between 1955 and 1969, with exception of two quadrangles bordering Kentucky (figure 3). The DLGs were converted to 10-meter elevation grids using ESRI's TOPOGRID algorithm and mosaiced into a single dataset. National map accuracy standards specify a vertical accuracy of no

more than one-half contour interval used in the source map. For the study area, where the contour interval is usually 20 or 40 feet, this means an expected error of about 3-6 meters (10-20 feet).

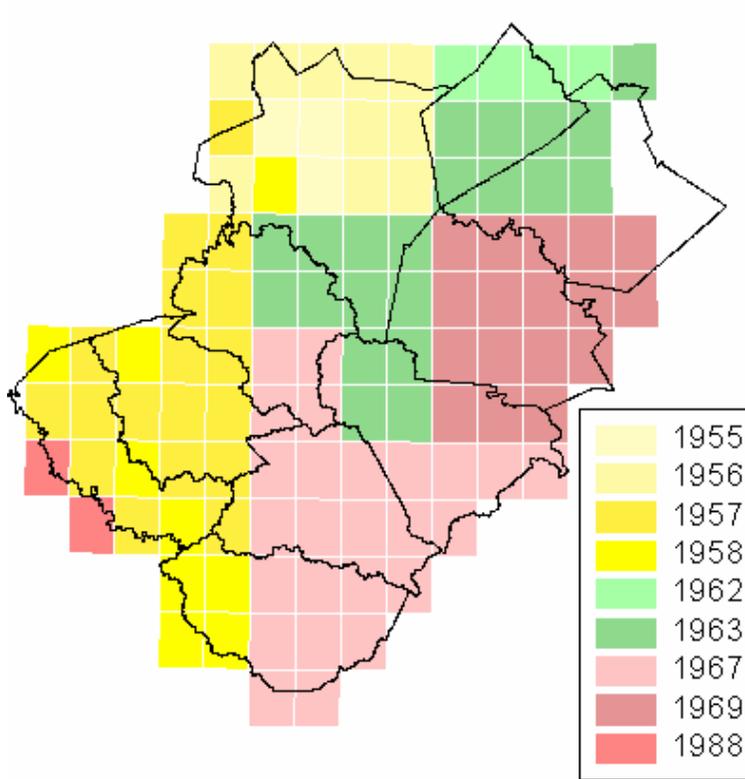|	
	1955
	1956
	1957
	1958
	1962
	1963
	1967
	1969
	1988

Figure 3. Image acquisition dates for the USGS contour data.

Forest mask

The final dataset used for the analysis was created from a pair of Landsat TM images that were separately classified into a forest/nonforest mask (figure 4). The two images had the following characteristics:

Path/Row	Acquisition Date	Area
18/33	08/10/2002	Northern
18/34	06/02/2003	Southern

The mask was used to restrict the analysis to cleared areas such as grass, soil or bare rock. The acquisition date of the northern image, more than one year prior to the acquisition of the

IFSAR dataset, would create a classification that lagged mining operations that were actively conducting clearing operations in the intervening period. Special attention was paid to this fact when finalizing fill polygon outlines, to make certain that the forest mask did not erroneously limit fills detected in areas covered by the source image.
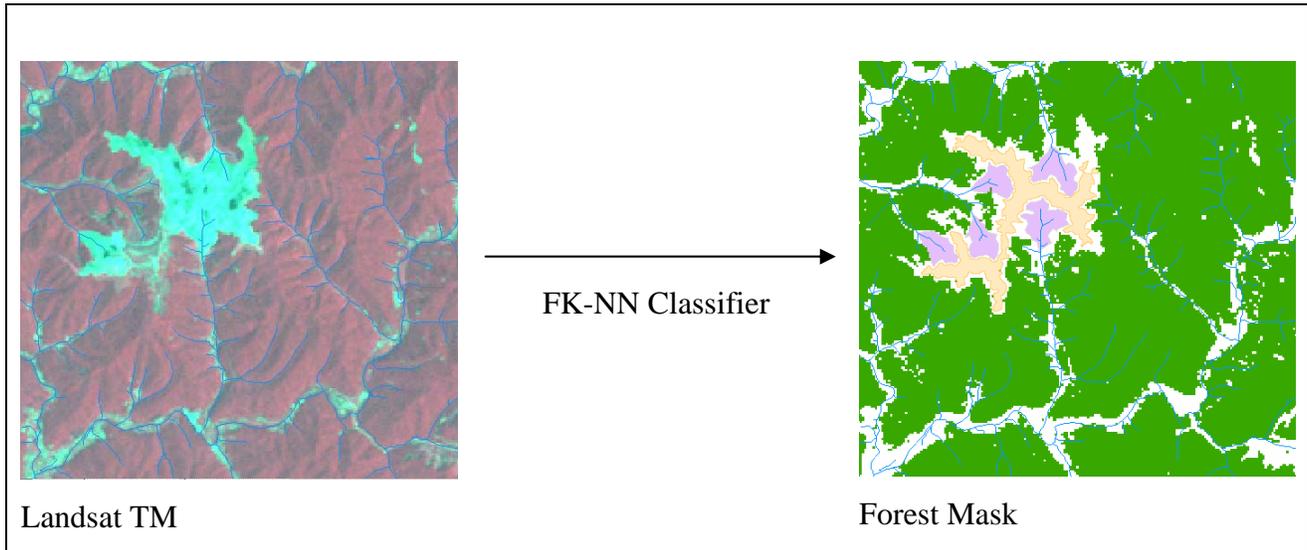

Figure 4. Landsat TM image, left, was classified into a forest/nonforest mask, shown at right, with cut/fill areas added.

The forest mask was produced using a landcover classification method based on the Fuzzy K-Nearest Neighbor (FK-NN) algorithm (Keller, 1985). This method assigns a landcover class to each image pixel based on the distance to the K nearest samples in a pattern space. The samples are selected from the source image to be representative of the landcover classes to be identified, and the pattern space is formed by assigning n spectral bands from the image to n orthogonal axes of the pattern space (figure 5). The distances to the K neighbors are inversely weighted and summed for each possible landcover class, with the sum of all classes equaling 1. Generally, the class with the highest value is assigned to the image pixel.

The limiting factor in spectral-based classification approaches such as FK-NN is the amount of overlap in the pattern space between two or more landcover classes. The FK-NN algorithm accounts for overlap by allowing partial, or fuzzy, class membership in multiple landcover classes, based on ideas drawn from fuzzy set theory originally presented in Zadeh (1965). Fuzzy class

membership does not always result from class overlap in the pattern space; it may in fact represent a real phenomenon on the ground, *i.e.* the occurrence of an intermediate gradation between two classes. For the purposes of this analysis, however, the central problem resulted from spectral overlap between prototype samples of forest and grass, which can be seen in figure 5. In an effort to enhance the results of the classification, pixels that had significant membership in more than one class were subjected to a neighborhood analysis, whereby the class totals for all immediate neighbors were summed, and the majority class assigned to the pixel in question. This approach attempted to extend the purely spectral-based approach, supplementing it with information on the spatial context in which the uncertain pixel resided.
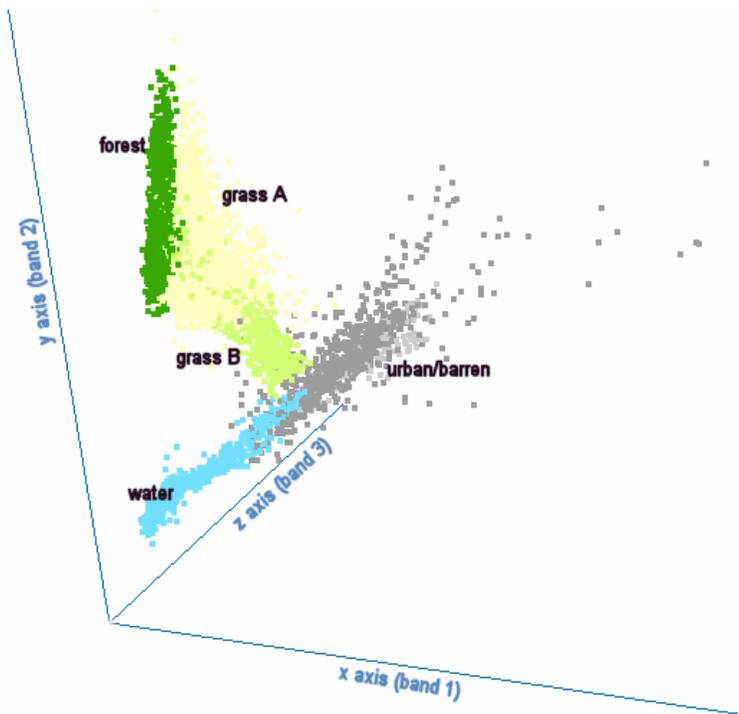

Figure 5. Example of a 3-dimensional pattern space used for landcover classification. The varying response of different landcover types in several spectral bands produces separation in the pattern space, which can be exploited by a variety of analysis methods.

III. Analysis Procedure

The logical flow of the analysis is shown in figure 6. The core of the entire process is simply to process a difference of two elevation grids representing distinct time intervals, $T_2 - T_1$, so that positive values reflect a net increase in elevation, and negative values a net decrease. The rest of the analysis is concerned with minimizing the effect of errors on the source data.

For this analysis, additional processing was required to account for the IFSAR sensor's inability to penetrate tree canopy. This required limiting the analysis to areas devoid of trees. This was accomplished through the use of a forest/non-forest mask, produced using the classification procedure described in the previous section. Areas where tree cover was detected were simply set to zero and were not processed further.

In a hypothetical circumstance where both elevation datasets were perfectly accurate, all areas of no elevation change in the difference grid would have a value of 0. Of course, real data is never prefect, so a zone needed to be expanded about 0 that eliminated most of the error deviations in the source data, while retaining real variations. This was accomplished through a reclassification of the difference grid into three categories—cut, fill, and null—with the null class representing the expanded zone of no change.
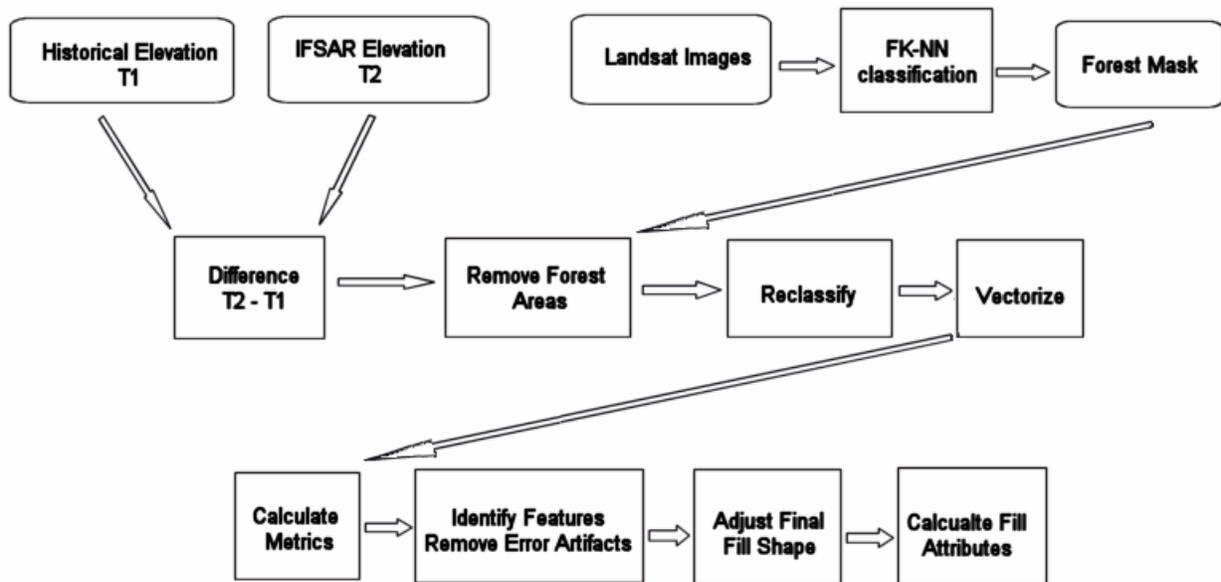

Figure 6. Analysis procedure.

Deciding where to set the thresholds that define the no change class depends on the characteristics of the source data and the nature of the features to be identified. As the positive threshold is set higher, noise due to error in the source data begins to drop out and real changes become increasingly isolated and easy to identify. At the same time, however, subtle changes are lost. During this study, it was apparent that some overlap existed between the largest errors and the shallowest fills, so that some compromise was needed. The positive threshold was initially set at 18 meters, with the knowledge that this would create error artifacts that would need to be removed by subsequent analysis. For cut areas, a 5-meter threshold effectively removed the vast majority of errors. Figure 7 shows the initial difference grid, and figure 8 shows the effect of isolating potential features by performing the reclassification.
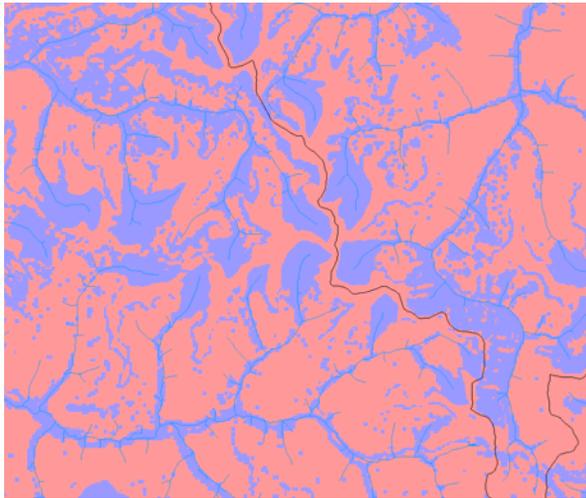

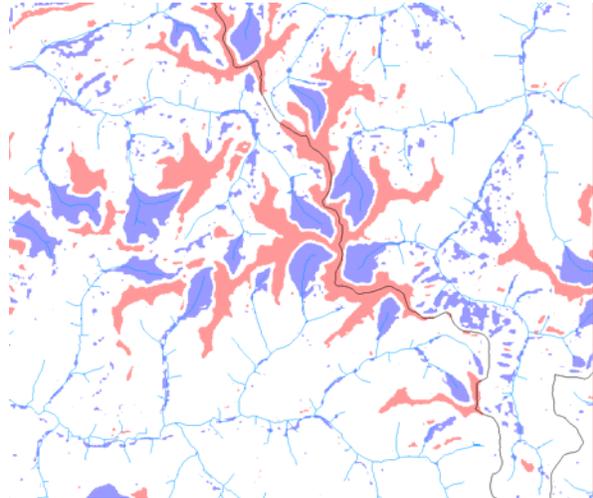Figure 7. Initial difference grid. Cut areas are red, fills blue.

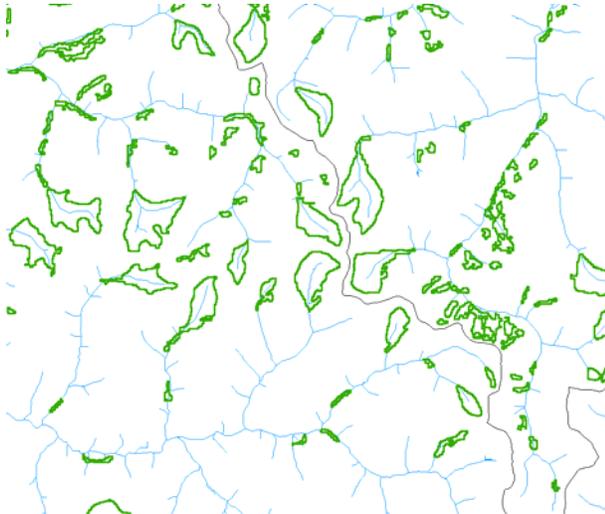Figure 8. results of initial reclassification.
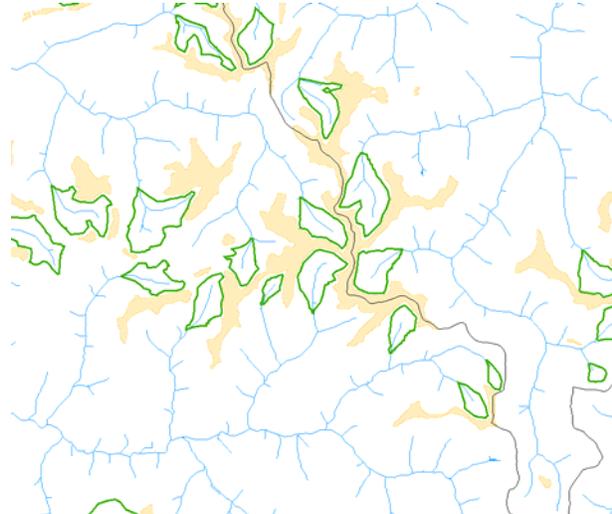
Figure 9. Initial vector polygon fill inventory.
Figure 10. Final inventory, shown with cut areas added.

Following reclassification, isolated cuts and fills were converted to vector polygons, and polygons less than 3 acres in size were eliminated (figure 9). This resulted in 9,661 fill candidates.

Several metrics were calculated for each fill candidate, including distance from a cut feature, standard deviation of depths within the polygon boundary, minimum and maximum depth, and the percentage of area that drained to a single location. Numerous attempts were made to develop a set of rules that would separate real features from error artifacts, including the use of a fuzzy pattern space classifier. However, the process resisted automation. Figure 11 plots fills against non-fills for two of the metrics used—average distance from cut, and standard deviation of depth within the polygon. As expected, fills tended to cluster nearer to cut areas and often exhibited higher standard deviation. However, it was not possible to separate the two classes without leading to errors of omission. The relatively large area of overlap between the two classes required extensive manual examination and verification using a variety of supplemental image and geographical data.
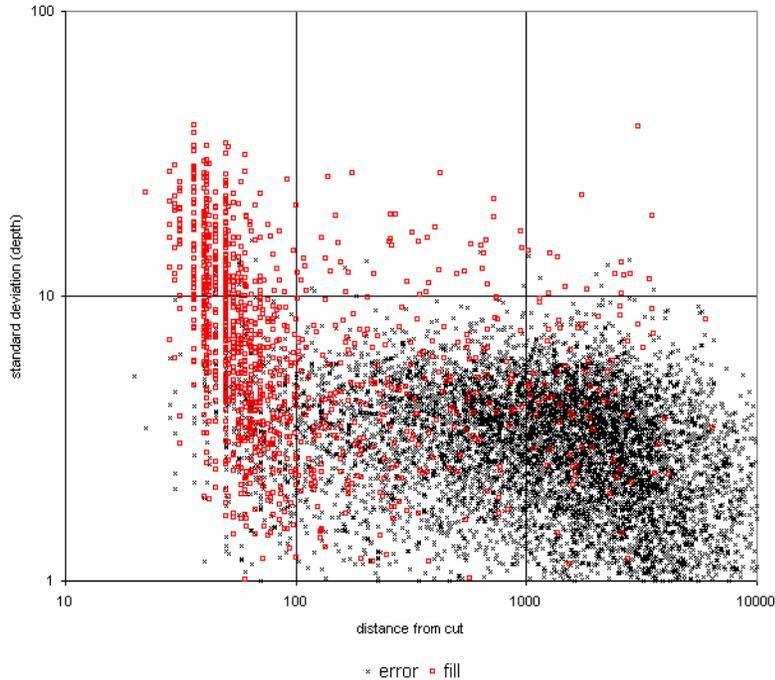
Figure 11. Plot of fills (red) and error artifacts (black) for two
metrics. Overlap always occurred, which prohibited
automatic identification of fill features.

After the identification process, final fill outlines were edited to approximate a 10-meter depth contour, and errors due to misclassification in the forest mask were corrected. At this time, additional small fills missed by the analysis were added as they were identified. During this phase, cut area polygons were particularly useful for drawing attention to the location of potential fills that were too shallow, small, or vegetated to have been detected. The final dataset exhibited characteristics shown in figure 10.

IV. Results

The final inventory comprised 1,329 fill polygons covering an area of 38,633 acres (60.4 square miles). To calculate length of stream under each fill, a vector stream channel network was derived from the hypsography-based elevation grid, after 1:24,000-scale National Hydrology Dataset streams were embedded and sinks removed.   Flow accumulations greater than 15 acres were reclassified as stream channels and converted to vector line data. The total length of 15-acre stream channels under fills was 479.9 miles. Volumes were calculated for each fill as well. However, the reliability of these calculations has not been evaluated.

The elevation-derived fill inventory was compared with fills digitized from permit maps submitted to WVDEP. The permit-based inventory was compiled as part of a program to create an integrated GIS database of active mining operations based on existing permit maps. The program also captured information for closed permits when maps were available. The permit-based fill inventory contained 1800 fill polygons for the study area, though the same valley was sometimes permitted twice, and a significant number of the permitted fills had not been constructed. The comparison indicated that 513 (38.6%) of the fills identified in this study did not overlap any fill captured from permit map sources. These fills represented 25.3% of the total surface area under fill, and accounted for 138.6 miles (28.9%) of stream channels. In several instances, the overlap between permitted and discovered fills seemed purely coincidental. For example, the largest (577 acres) and third largest (330 acres) fills discovered in the study intersected permitted fills that were only a few acres in size, and probably were not associated with the discovered fill.

In addition to locating previously undocumented fills, the discovered fills often deviated from permitted fills in significant ways. Some of the disagreement could be traced to fills that were under construction at the time of data collection, while others seem to reflect limitations in the permit-based inventory itself. For example, sometimes only the face of a fill was shown on permit maps, rather than the full extent. On other occasions, permitted fills were significantly smaller than what was indicated by the elevation study. Figure 12 depicts several examples of mismatches between elevation-derived and permit-based fill outlines.
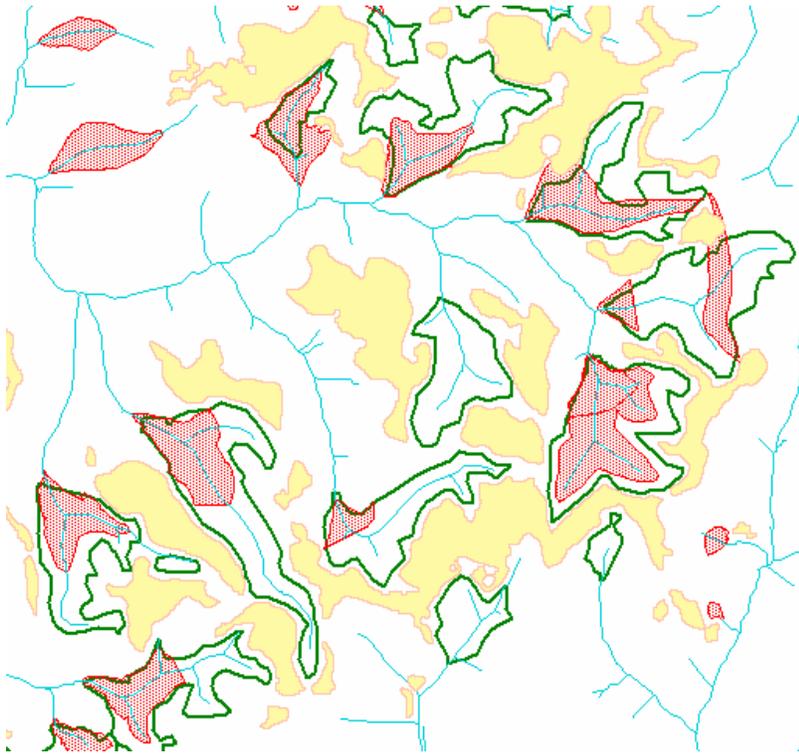
Figure 12. Example of mismatch between permit-derived valley fills and fills derived from the analysis. Fills digitized from permits are shown in red, fills extracted from elevation data are outlined in green, and cut areas are depicted as yellow.

Comparison with LIDAR Study

Results of the study also were compared with previous work that sought to identify valley fills using LIDAR, in place of IFSAR, as an elevation source. The LIDAR was collected as part of a floodplain mapping project in Wyoming County, which makes up about 9.28% of the total study area. The LIDAR study was expected to produce a superior result, and potentially could be used to gage the effectiveness of the larger study. This was due largely to the LIDAR sensor's ability to record at least some returns from the forest floor during leaf-off conditions, which could be used to construct a better bare-earth model of elevations.

Out of 117 fills confirmed by the LIDAR study of Wyoming County, 14 (11.9%) were missed when using the IFSAR dataset. Of the fills that were missed, 9 did not produce a fill polygon candidate, and 5 were erroneously discounted as a noise artifact. Of the 9 fills that did not

produce a fill candidate, they were either too shallow to be resolved, or partially covered by the forest mask, or both. Assuming a foolproof evaluation of fill candidate polygons could be made, the radar data exhibited a miss rate of 7.6% for Wyoming County.

The missed fills averaged 4.5 acres in size, compared to the overall average of 18.3 acres for the entire county. This supports the idea that the missed fills were relatively small and shallow. If Wyoming county comparison is extrapolated to the rest of the study area, taking into account the relative density of fills, it could be expected that as many as 160 additional small fills, covering 720 acres, could exist throughout the study area that have not been identified. This aggregate total is somewhat larger than the single largest fill (577 acres) but probably represents a point of diminishing returns, in terms of analyst time and data acquisition costs for LIDAR data.

V. Conclusion

An IFSAR elevation dataset was combined with historical elevation data to successfully develop an inventory of mining cut and fill features over a large area. The analysis significantly expanded the list of known fills over an existing inventory created from permit records, and provided verification of which permitted fills were actually constructed, or under construction, throughout the entire region. When compared with LIDAR data, which was used to perform a similar study, the IFSAR product was somewhat less adept at detecting small, shallow fills, particularly when partially revegetated. IFSAR's inability to penetrate tree canopy effectively required the creation of a landcover classification, which was used to mask forested areas from the analysis. In contrast to LIDAR, however, IFSAR-based elevation data can be acquired at a dramatically reduced cost, facilitating analysis over relatively large areas that would be difficult or impossible to fund otherwise.

References

Keller, J., M. Gray and J. Givens, "A Fuzzy K-Nearest Neighbor Algorithm", *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-15 no. 4, pp. 580-585, July/August 1985.

Shank, Michael, "Automated Detection of Mining Valley Fills From Multi-Date Elevation Data" *http://gis.wvdep.org/tagis/projects/fill_paper.doc*, 2003

Zadeh L., "Fuzzy Sets", *Journal of Information and Control*, Vol.8, pp. 338-353, 1965.