Combine Harvest Area Determination by Vector Processing of GPS Position Data

S. T. Drummond, C. W. Fraisse, K. A. Sudduth

Abstract. *The measurement of actual harvested area per unit time is an important component in the creation of accurate crop yield maps. For row crops, such as corn, these measurements can be made manually on most conventional yield monitors. However, in drilled or broadcast crops a more accurate and automated method is required. In this study, a vector method was developed to determine actual combine harvest area at each time step of the harvest process from a global positioning system (GPS) trajectory. The algorithm was coded into a geographic information system (GIS), and modifications were made to increase computational efficiency. The method was compared with a previously reported raster method of harvest area determination on data collected during the 1997 drilled soybean harvest, using high accuracy real time kinematic GPS data. The vector method improved harvest area estimates by an average of 11% over the assumption of a constant, full swath width and provided a number of distinct advantages over the raster method.*
Keywords. *Precision farming, Yield mapping, Swath width, GPS, GIS.*

The equipment and techniques necessary for crop yield measurement have progressed significantly in the last decade. Pierce et al. (1997) reviewed a variety of grain flow sensors, several of which are commercially available today. Through the inclusion of position determination techniques, most common of which is the differential global positioning system (DGPS), the ability to create grain yield maps has become fairly commonplace, both among researchers and producers.

The accuracy and quality of these yield maps are important in the development of effective precision management strategies. However, the creation of accurate yield maps is a process made difficult by a number of possible error sources. Blackmore and Marshall (1996) identified six main groups of errors affecting yield map accuracy and intuitively ranked them with regard to their effect on grain yield maps. They ranked the error sources in this way:

1. Unknown crop width entering the header during harvest.
2. Time lag of grain through the threshing mechanism.
3. The inherent "wandering" error from the GPS.
4. Surging grain through the combine transport system.
5. Grain losses from the combine.
6. Sensor accuracy and calibration.

Article was submitted for publication in November 1998; reviewed and approved for publication by the Power & Machinery Division of ASAE in July 1999.

Mention of trade names or commercial products is solely for the purpose of providing specific information and does not imply recommendation or endorsement by the U.S. Department of Agriculture or the University of Missouri.

The authors are **Scott T. Drummond**, Computer Specialist, **Clyde W. Fraisse**, ***ASAE Member Engineer***, Agricultural Engineer, and **Kenneth A. Sudduth**, ***ASAE Member Engineer***, Agricultural Engineer, USDA-ARS Cropping Systems and Water Quality Unit, Columbia, Mo. **Corresponding author:** Scott T. Drummond, USDA-ARS, 269 Agricultural Engineering Bldg., Columbia, MO 65211; voice: (573) 882-1146, fax: (573) 882-1115, e-mail: drummonds@missouri.edu.

It is interesting to note that these authors felt swath width determination was the largest error source affecting grain yield map accuracy. Though the magnitude of these errors has not been well defined, Stafford et al. (1997) suggested that errors averaging up to 10% could be experienced in normal combining operations. This is an extremely high level of non-random error to contend with in yield mapping and demonstrates the importance of swath width determination for crops not grown or harvested in defined rows.

Our primary goal in this study was to develop and test a vector-based method of determining combine cutting width and harvested area using combine position data. Additionally, we wanted to apply the vector-based method to a set of typical field data to help us better understand the nature and magnitude of the errors caused by swath width variation. Finally, we wanted to investigate the feasibility of using this method compared to other, previously suggested methods.

Literature Review

The problem of swath width determination has been widely studied from a variety of approaches. These approaches range from methods to minimize the effects of swath width on yield measurement to direct measurement techniques, such as the use of ultrasonic sensors. What follows is a survey of the techniques which have been used (or might possibly be used) to measure swath width or harvested area or to minimize its effect on yield map accuracy.

Operator Estimation

Reitz and Kutzbach (1996) pointed out the fact that some commercial yield monitoring systems offer the possibility to manually vary the recorded cutting width. This approach is an improvement over the assumption of constant swath width. However, it is extremely difficult to implement in the field for at least two reasons. First, the combine operator seldom has the time and concentration available to make the right corrections precisely at the right times. Secondly, even when the combine operator can make the adjustments properly, or has a second person available to do so, the effective resolution (and resulting accuracy) of these systems would likely be measured in feet.

Effect Minimization

Methods have been suggested which could minimize the effect of varying swath width on yield maps. Birrell et al. (1996) proposed a raster accumulation method by which the geographically harvested area corresponding to the header travel during each time interval was determined from the GPS position information. The grain mass harvested during that time interval was then distributed to each grid cell traversed, according to the ratio of the harvested area in any particular grid cell to the total area harvested for the time interval. After all sample points had been processed, each cell would contain the total grain mass harvested in that cell. Similar techniques have also been suggested (i.e., Potential Mapping, Blackmore and Marshall, 1996; Equal Area Grids, Pierce et al., 1997).

If the cell size is large in relation to the combine header width and the GPS positioning error, this method could be quite effective in removing the effects of combine swath width on yield maps. However, this method is obviously quite susceptible to missing data, or even data that has been shifted improperly (i.e., improper time lag, position errors, etc.). Additionally, these methods by definition provide a raster format for the final yield data, and this may or may not be desired for a particular application.

Ultrasonic Sensors

Several groups have investigated the use of ultrasonic sensors for automatic swath width measurement (Vansichen and De Baerdemaker, 1991; Reitz and Kutzbach, 1996; Stafford et al., 1997; Sudduth et al., 1998). In general, the results from these studies have been quite promising, particularly when the distance between crop edge and sensor is relatively small. However, in practice these systems may be difficult to implement for a variety of reasons. First, when the crop edge is dense and “crisp”, the sensor measurements are usually quite accurate (various authors have reported errors in the range of 2 to 10 cm). However, under certain relatively common conditions (i.e., sparse crops, crops bent over from a previous pass, crops that have fallen or matted, extremely windy conditions, etc.), ultrasonic swath width sensing may be problematical or even impossible.

Machine Vision

Hoffman (1996) described the Demeter project, an automated vehicle guidance system for a forage harvester based upon high speed (5-30 Hz) crop edge detection from camera images. While the current system does not employ any yield measurement device, the machine vision approach might be useful in either minimizing the effect of swath width variation through vehicle guidance, or as a means to measure cutting width outright.

While it may be possible to achieve quite acceptable accuracies through this approach, the implementation of such a system for swath width detection could be quite difficult. Without a doubt, the implementation of a machine vision system for swath width measurement would require a significant investment in terms of additional hardware costs.

Combine Position

The improvements in GPS equipment accuracy over the last few years have prompted researchers to consider GPS position data as a means of computing effective combine harvest area. It is important to consider the effect of GPS accuracy on the suitability of such methods. For example, conventional DGPS methods, with accuracies measured in the 1- to 3-m (rms) range, are unsuitable for computing harvest area, as these errors represent 20 to 60% of the header width on a 5-m grain platform. By comparison, real time kinematic (RTK) GPS provides horizontal accuracies in the 2- to 4-cm (rms) range, or less than 1% of header width on a 5-m platform, and would be a sufficiently accurate data source for determining harvest area. For systems with accuracies between these two extremes (10 to 100 cm), computational methods using combine position data may provide some benefit, however the effects of GPS error under these conditions should be carefully considered.

Han et al. (1997) proposed a method for determining effective combine harvest area from the sequential GPS positions of the combine in the field (fig. 1). The method consists of initializing a high resolution bitmap representing the pre-harvest crop conditions of the field (zero = no crop, one = crop). The bitmap is then progressively updated at each step of the harvest process, by turning all of the cells whose centers fall within the area traversed by the combine header to zeros. The summed area of the cells which changed from ones to zeros during each time interval represents the area harvested for that time interval.

The major advantages of this approach are that it requires no additional equipment, beyond that required for yield mapping, and no additional user input. Another important consideration is that the accuracy of methods based upon combine position should significantly improve as the accuracy of the positioning equipment improves. Additionally, though the method might require significant processing power, it could likely be implemented in real-time. As long as cell size is very small in relation to header width, this method provides a simple, efficient means of estimating harvest area. However, the memory necessary to store a high resolution bitmap for a large field is considerable. For example, for a 50-ha square field with a 5-cm cell size, about 25 MB of memory is required. As the accuracy of GPS equipment steadily and significantly improves, cell sizes this small, or possibly even smaller, may be required to minimize harvest area errors using this approach.

Materials and Methods

Algorithm Development

In an attempt to build upon the advantages of the raster method suggested by Han et al. (1997), we developed a vector method which similarly determines the effective harvest area from combine position data. The polygon representing the area covered by the header (based upon

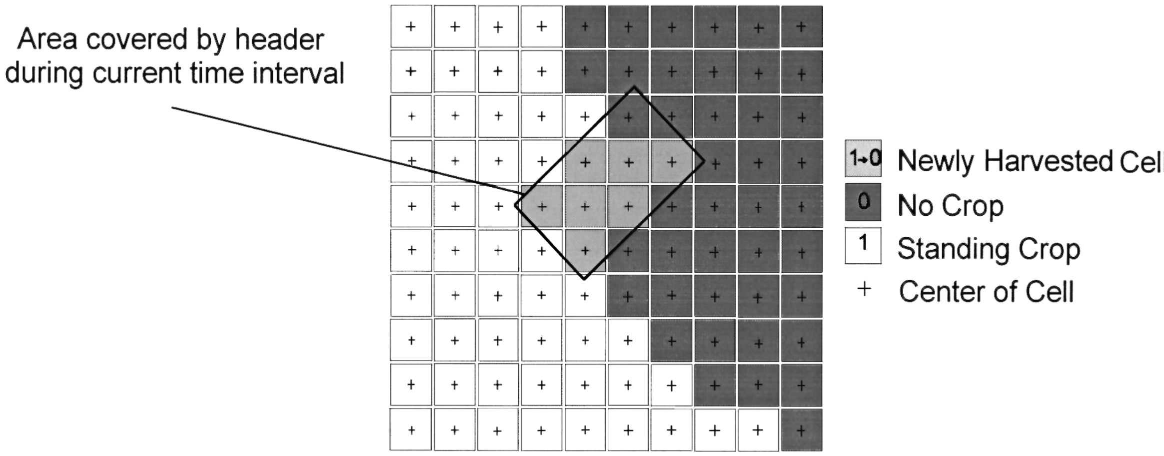

Figure 1–Graphical representation of the bitmap method suggested by Han et al. (1997). Area harvested is estimated by multiplying the number of newly harvested cells times the area of each cell. Cell size in the figure is greatly exaggerated for illustrative purposes.

position and trajectory information) during each time interval is determined. Each polygon is then processed in chronological harvest order by subtracting, in the boolean sense, all previously processed polygons from the current one. In a more formal sense, let the area traversed by the combine header for time interval i of the harvest process be represented by a_i. In set notation, the actual harvested area for each time interval, A_i, is given by:

$$A_i = a_i - \left(a_1 \cup a_2 \cup \ldots \cup a_{i-1}\right) \tag{1}$$

The resulting polygons represent the actual harvested area during each time step. Figure 2 shows a graphical representation of this concept for three time intervals during which the combine header areas overlap. Intuitively, it is obvious that this technique should be able to extract the actual harvested areas to the level of positioning accuracy attained for each time interval of the harvest process.

However, this approach is not without drawbacks. Boundary conditions at the edges of fields, or at islands within the fields, whether real or due to missing data, may cause overestimation of actual harvested areas along these boundaries. If the true locations of these boundary conditions are known, judicious use of "boundary" polygons could be used to correct these problems.

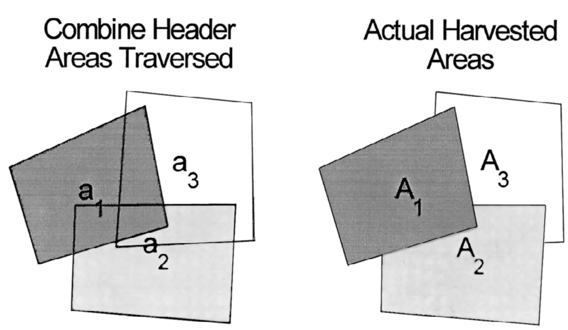

Figure 2–Example of determining actual harvested areas (A_i) from the combine header areas covered (a_i) on three ordered time intervals.

A more serious drawback of this simplistic approach is the computational complexity of the algorithm. Given that we have N discrete time intervals and associated polygons which require processing, then clipping may be required on $1/2\ N^2$ polygon pairs. Thus, with respect to the polygon clipping routine, the complexity of the algorithm is order N^2. In other words, as we double the size of our dataset, we quadruple the number of polygon clipping operations that must take place.

There are techniques we can use to reduce the processing time of this problem. The most obvious of these is a bounding window check of the two polygons upon which we are about to perform the polygon clipping algorithm. If the bounding windows of the two polygons do not overlap, the clipping algorithm need not be performed for that particular pair of polygons. While this modification to the algorithm significantly reduces processing time, it does not reduce the computational complexity of the problem below order N^2.

A technique which could be used to reduce the complexity of the algorithm would be to group adjacent, consecutive time interval polygons together into individual "passes" or "coverages". Coverages might be distinguished from one another by a number of possible techniques, i.e., header raised/lowered, combine heading or speed changed, break in time sequence, etc. Once an individual coverage has been processed, and the resulting individual harvest area polygons (A_i) have been output, the entire coverage could be represented by a single "boundary" polygon (fig. 3).

This technique should significantly reduce the computational complexity of the problem. For example, consider the best case of a square field with $\sqrt{N}$ parallel swaths with each coverage containing $\sqrt{N}$ individual polygons. Now, each new polygon a_i must be clipped with less than $\sqrt{N}$ other coverages and less than $\sqrt{N}$ individual polygons within its own coverage. This yields an overall computational complexity with respect to the clipping algorithm of order $N^{3/2}$—a considerable improvement over N^2. For example, for a given field with 10,000 polygons, the order N^2 algorithm might require as many as 100,000,000 polygon clipping operations, while the $N^{3/2}$ algorithm would have an

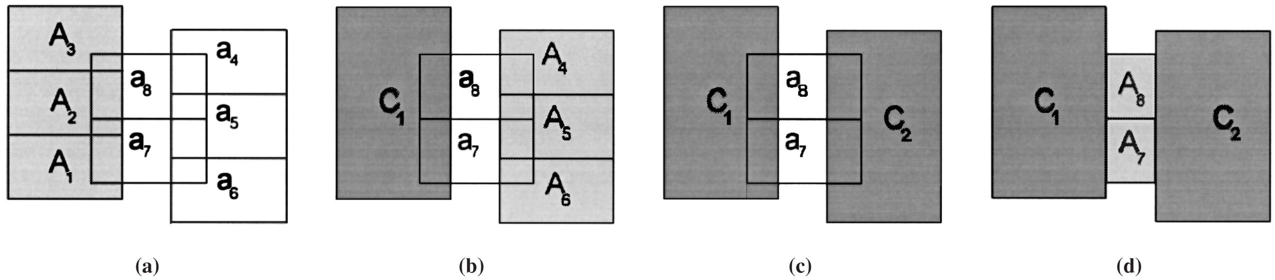

Figure 3–Example of algorithm with coverage implementation: (a) actual areas $A_{1,2,3}$ in coverage C_1 are processed and recorded; (b) actual areas in coverage C_2 are processed and recorded; (c) header areas a7,8 are processed against C_1 and C_2 and (d) all actual harvest areas have been processed and output.

upper limit of 1,000,000 such operations. Furthermore, it is clear that the algorithm is inherently parallel in nature and could easily be implemented in a multiprocessor environment, further reducing the total amount of time necessary to achieve a solution.

The complexity and efficiency of the clipping algorithm used is another important issue that should be considered when implementing such a procedure. Many such algorithms exist (i.e., Vatti, 1992), and vary in complexity depending upon such factors as whether the input polygons might be described as: convex; concave; self-intersecting; self-touching; containing holes; etc. To be certain, one can envision unique field situations in which polygons of each of these classes might be created, and for the overall algorithm development to be robust, a polygon clipping routine which can handle all of these cases would be required. For the sake of efficiency, it may be necessary to implement an algorithm which can handle most, but not all of these cases. For example, Weiler and Atherton (1977) described an attractive, efficient edge labeling algorithm which can be slightly modified so that it will handle all polygons except those of the self-intersecting variety. For a practical implementation of the vector approach to harvest area measurement, such a polygon clipping algorithm might be an excellent selection, as the occurrence of a self-intersecting polygon would almost certainly be due to a data collection error, and should likely be removed from analysis.

Analysis Procedures

The algorithm outlined in the previous section was implemented using ARC/INFO version 7 Arc Macro Language (AML) routines (code available from authors upon request). While it was understood that this approach would incur significant overhead and thus affect processing speed, we felt that the reduction in development time and the portability of the finished product were acceptable trade-offs for these initial investigations.

The polygon clipping operation was performed using the ARC/INFO IDENTITY command. This command takes two arguments, an input polygon or polygons (in our case, the current, unprocessed polygon) and a second "identity" coverage (in our case, a previously processed coverage). The output coverage contains all of the input polygons and only those portions of identity coverage that overlap the input polygons. In addition, the feature attribute table preserves ID information from both the input and identity features. These intersecting polygon(s) have a non-zero feature identification code with respect to the identity coverage, and in this way can be selectively removed from the original input polygon(s). The remaining polygons represent the actual harvested area. Several other ARC/INFO procedures were considered and rejected for use in the polygon clipping operation. For example, the ERASE command appears to be precisely the right procedure to perform the clipping operation we require. However, there are certain cases where ERASE creates extraneous external polygons, and since feature attributes are not preserved with this command, there is no easy, automated way to select and remove these extraneous polygons.

Our trial dataset consisted of an approximately 10-ha soybean field located in central Missouri. The crop was harvested with an R42 Gleaner combine with a 4.6-m-width grain platform, and yield data were collected with an AgLeader 2000 yield monitoring system. Position information was collected using a pair of Ashtech Z-Surveyor GPS receivers operating in real time kinematic (RTK) mode. The manufacturer's stated accuracy of these receivers was 3 cm (rms) as operated for this study. On the 4.6 m header used, this corresponds to a positioning error of 0.7% of the header width. This level of positioning accuracy seemed adequate for studying variations in combine harvest area.

A significant fraction of the area under study required replanting due to water ponding problems in the early season, and the replanted soybeans were harvested approximately two weeks after the soybeans from the first planting. As a result, the harvest pattern for the western two-thirds of this field was quite complex, including a number of curved trajectories, and it was difficult for the driver to keep a constant swath width entering the header at all times. Position and heading data were used, along with knowledge of the geometry of the header and its position in relation to the GPS antenna, to create a map of the areas covered by the header during each second of the harvest process (fig. 4). These polygons were created by connecting the positions of the header ends at the beginning and end of each of these second time intervals.

Each of the polygons was identified by a time tag, such that grain flow information collected by the yield monitor could be recombined with the actual harvested areas after processing was complete. New coverages were initiated whenever (1) a time-break occurred in the polygon list, (2) the heading of the combine changed rapidly and significantly, and (3) shortly after the combine header sensor indicated a raising of the combine header. These steps were taken to insure that each coverage would consist

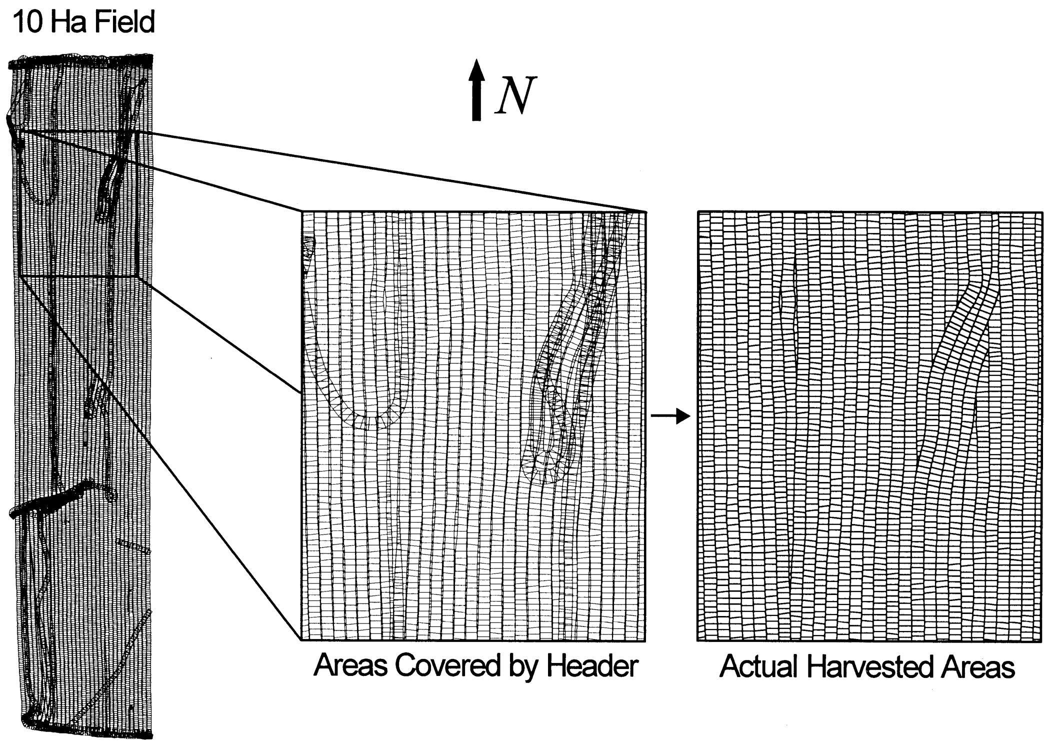

Figure 4–Polygons representing the areas covered by the header, and the actual harvested areas determined by the vector-based method for a small region of the field.

of contiguous but likely non-overlapping polygons. (It is clear that additional algorithmic efficiency would result if this could be guaranteed.) Next, our algorithm was applied to the polygon dataset, and the actual harvested polygons and their associated areas were determined (fig. 4).

For comparison, the bitmap method was implemented on this same initial polygon set, with both 10-cm and 30-cm square cell sizes. These sizes were selected since they spanned the range suggested by Han et al. (1997) for high precision positioning systems. Additionally, to evaluate errors with respect to the conventional assumption of a constantly full swath width, the total area covered by the header during each time interval was calculated and the results were compared to both the raster and vector methods.

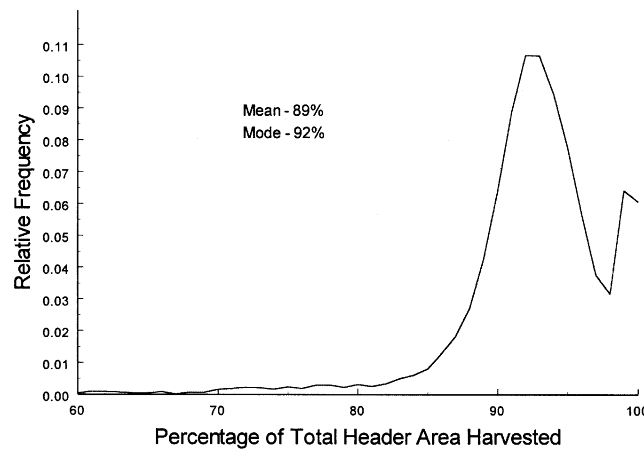

Figure 5–Distribution of the percentage of total area covered by the header which was actually harvested during each time interval.

Results and Discussion

One objective of this study was to better understand the magnitude of errors that might be introduced in yield maps by assuming a constant swath width when harvesting was actually accomplished with an unknown and varying swath width. We observed a significant amount of variability between actual harvested area and total area covered by the header on the test field. Figure 5 shows the distribution of these observations, with two distinct peaks. The smaller one occurs at 100% of header width and represents the harvesting of full-swath lanes within the field with standing crop on both sides of the header. The second occurs at about 92% of total header area and represents the more common harvest situation where the operator is trying to keep the header as full as possible while making certain no crop is left standing. Point rows and relatively small cleanup strips would likely account for those observations lower than 80%, and the distribution in this region appears to be quite random. Over all time intervals, the average harvested area was approximately 89% of the total header area covered. If we assume positioning error is negligible, then for this dataset we could expect an average of 11% error in harvest area estimation on a point-by-point basis, just from the assumption of a constant, full header width of crop entering the combine. If all observations where the header was operating less than half full were removed from the data set, this error was reduced to approximately a 7% error in harvested area. This is still a very significant error source in the calculation of instantaneous crop yield. Similar results were achieved with the bitmap method, with average errors for both the 10 cm and 30 cm cell sizes of approximately 11 to 12%, compared to a constant, full swath of crop.

Next, we considered whether the complexity of the harvest pattern would affect these results. As previously mentioned, the western two-thirds of the field required a more complex harvest pattern, due to the large replant area which included multiple end-rows and numerous curved trajectories. We compared the results within this area to the results from the eastern one-third of the field, which was harvested more conventionally, in long, approximately parallel swaths. The distributions for these two regions were very similar, with both showing bimodal patterns (fig. 6). However, the easier-to-harvest eastern region of the field had a significantly smaller average error (8%, as compared to the assumption of constant, full swath width) than did the western region of the field (12%). It seems reasonable to expect that, as we investigate fields with progressively more complex harvest patterns, this general trend would continue. More research on larger fields containing highly variable landscapes (i.e., terraced fields) seems warranted.

A comparison of the results from the vector method to those from the bitmap method of Han et al. (1997) was also performed. Figure 7 shows the distribution of the differences in calculated harvest areas for each individual harvest polygon between the bitmap method and the vector method, for both bitmap cell sizes. For the 10 cm cell size, the difference between the vector method and the bitmap method was seldom more than ±0.2 m^2, with a standard deviation of 0.157 m^2. This would correspond to less than a 2% difference on the average-sized polygon of 8.24 m^2. For the 30 cm cell size implementation, a significant number of differences greater than ±1 m^2 were noted, with a standard deviation of 0.402 m^2, which represents a nearly 5% difference on the average-sized polygon. These results emphasize the fact that there is the potential for significant differences between the results obtained by the vector method and the bitmap method, particularly as the selected cell size for the bitmap method increases. However, it is important to realize that the vector method should compute actual harvest areas accurately only when GPS positioning error is very small. In situations where GPS positioning error is relatively large in relation to header width, the vector method introduces unnecessary computational complexity with no real accuracy improvement over the simpler raster method. Furthermore, there is a point where GPS accuracy levels are poor enough that any means of computing harvest area from GPS position data would provide little or no advantage, or possibly a disadvantage compared to the assumption of constant swath width.

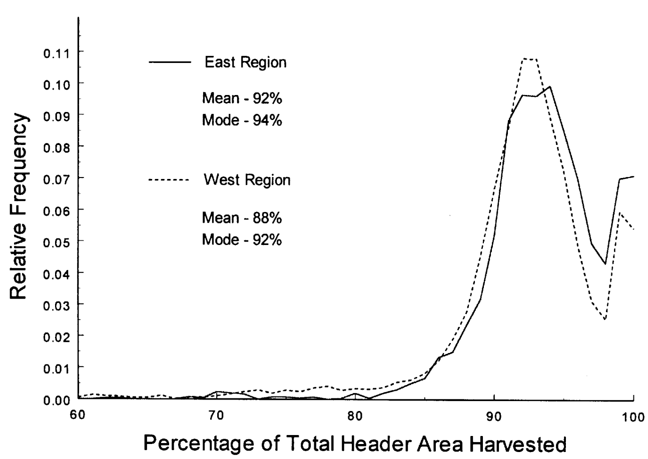

Figure 6–Distribution of the percentage of total area covered by the header which was actually harvested during each time interval for two dissimilar regions of the field.

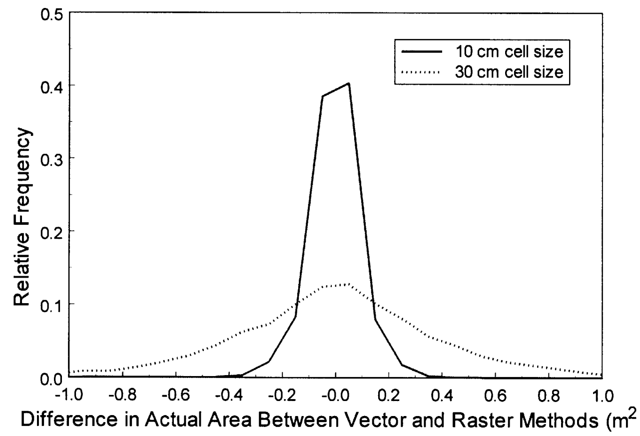

Figure 7–Distribution of differences in calculated harvest areas between vector and bitmap methods for both 10- and 30-cm bitmap cell sizes.

Beyond the potential for more accurate harvested area and yield estimates, a useful byproduct of the vector method is that it provides us with the actual polygons representing the harvested areas. With these polygons, it is possible to create a yield map of classed polygons, as opposed to the more conventional classed point maps (fig. 8). Maps of these classed polygons are created by placing each polygon observation into a particular class, based upon a given classification variable (i.e., yield), and "coloring" each class differently, in just the same manner as is done for a common classed point map.

Classed polygons provide significant advantages in the interpretation phase of yield mapping. First, they allow the user to intuitively see how and in what order the crop was harvested, helping the user to locate areas of the yield map that may provide unreliable results. Due to the effects of combine dynamics (ramping, starting and stopping in the crop, etc.), knowledge of the exact harvest pattern can dramatically affect how a yield map is interpreted. Secondly, with classed point maps, points which represent extremely small areas are given precisely the same symbol size as points which represent larger, more representative areas. This can mislead the user in the evaluation of yield maps, particularly where harvest patterns are unusual. Classed polygon maps allow the user to more easily ignore small areas in the interpretation, since these areas will be represented by a proportionally small area.

For example, consider the series of maps in figure 8. The first four shaded yield classes all represent reasonable yield ranges for this particular field, with the last class (> 3.4 Mg/ha) being at the maximum of the reasonable yield values we could expect to obtain. Notice that for the full, constant swath width map (fig. 8a), we are likely to consistently underestimate the actual crop yield across the entire area, yet we have very few points with unreasonably high yield levels. For the classed point map using calculated harvest area (fig. 8b), we tend to more

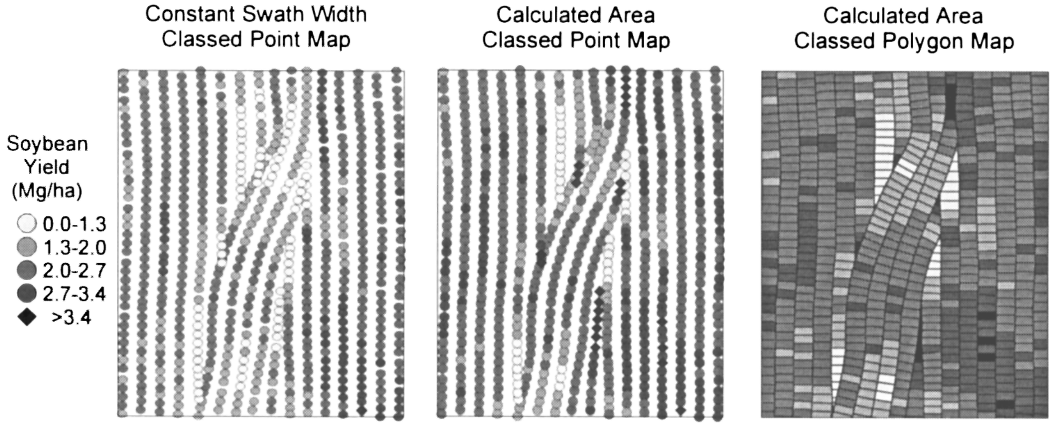

Figure 8–Comparison of classed point and classed polygon for mapping crop yields from harvesting with varying combine cutting width.

accurately estimate yield over the majority of the map, however, there are far more points where (due to small areas and moderate grain flows) we tend to receive many possibly unreasonably large results. These points may significantly affect the visual interpretation of the yield map, though in reality they represent very small areas of the field. In the classed polygon map (fig. 8c), this problem is resolved since those yield observations which represent very small areas are given proportionally small areas in the visual representation of the data.

Conclusions

The measurement of actual harvested area is an important component in the creation of accurate crop yield maps. In this study, a vector method was developed and implemented to determine the actual combine harvest area at each time step of the harvest process using combine position information. The method provides several advantages over a previously reported raster method. It can provide significant error reduction over the raster method, particularly as raster cell size increases in relationship to positioning accuracy. Although the vector algorithm is computationally more complex, it would not require any additional computational resources as positioning technologies become more accurate. The raster method would require steadily smaller cell sizes and larger memory requirements to achieve comparable accuracies.

The vector method was used to estimate errors in calculating instantaneous crop yield due to varying swath widths. Results from a 10 ha drilled soybean field showed that average errors of 8 to 12% in instantaneous area estimation are possible, if a full, constant swath width is assumed. Additionally, the bimodal error distribution indicated that any constant swath estimation method would perform poorly on this dataset. More investigation of swath width distributions over a variety of fields could provide additional useful information.

A useful by-product of the vector method is that it provides the actual polygons representing the harvested areas. With these polygons, it is possible to create a yield map of classed polygons, as opposed to the more conventional classed point maps. These classed polygon maps provide distinct advantages over classed point maps in visual yield map interpretation.

References

Birrell, S. J., K. A. Sudduth, and S. C. Borgelt. 1996. Comparison of sensors and techniques for crop yield mapping. *Comput. & Electronics in Agric*. 14: 215-233.

Blackmore, B. S., and C. J. Marshall. 1996. Yield mapping; errors and algorithms. In *Proc. 3rd Int. Conf. on Precision Agric*., 403-415. Minneapolis, Minn.: ASA, CSSA, SSSA.

Han, S., S. M. Schneider, S. L. Rawlins, and R. G. Evans. 1997. A bitmap method for determining effective combine cut width in yield mapping. *Transactions of the ASAE* 40(2): 485-490.

Hoffman, R. 1996. A perception architecture for autonomous harvesting. ASAE Paper No. 96-3068. St. Joseph, Mich.: ASAE.

Pierce, F. J., N. W. Anderson, T. S. Colvin, J. K. Schueller, D. S. Humburg, and N. B. McLaughlin. 1997. Yield mapping. In *The State of Site-specific Management for Agric*., eds. F. J. Pierce, and E. J Sadler, ch. 11: 211-243. Madison, Wis.: ASA, CSSA, SSSA.

Reitz, P., and H. D. Kutzbach. 1996. Investigations on a particular yield mapping system for combine harvesters. *Comput. & Electronics in Agric*. 14: 137-150.

Stafford, J. V., B. Ambler, and H. C. Bolam. 1997. Cut width sensors to improve the accuracy of yield mapping systems. In *Proc. Precision Agric. '97,* 2: 519-527. London, UK: SCI.

Sudduth, K. A., S. T. Drummond, W. Wang, M. J. Krumpelman, and C. W. Fraisse. 1998. Ultrasonic and GPS measurement of combine swath width. ASAE Paper No. 98-3096. St. Joseph, Mich.: ASAE.

Vansichen, R., and J. De Baerdemaker. 1991. Continuous wheat yield measurement on a combine. In *Proc. ASAE Symp. on Automated Agric. for the 21st Century*, 346-355, Chicago, Ill. St. Joseph, Mich.: ASAE.

Vatti, B. R. 1992. A generic solution to polygon clipping. *Commun. ACM* 35(7): 56-63.

Weiler, K., and P. Atherton. 1977. Hidden surface removal using polygon area sorting. *Comput. Graphics* 11(2): 214-222.