Article

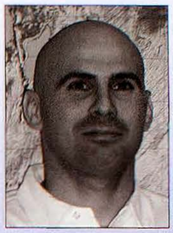

Mosaicing Tool for Aerial Imagery from a Lidar Bathymetry Survey

By Shachak Pe'eri and Yuri Rzhanov, Center for Coastal and Ocean Mapping, University of New Hampshire, Durham (NH, USA).

Abstract

Aerial imagery collected during lidar bathymetry surveying provides an independent reference dataset for ground truth. Mosaicing of aerial imagery requires some manual involvement by the operator, which is time consuming. This paper presents an automatic mosaicing procedure that creates a continuous and visually consistent photographic map of the imaged area. This study aimed to use only the frames from the aerial camera without additional information. A comparison between the features in the resultant mosaic and a reference chart shows that the mosaic is visually consistent and there is good spatial-geometric correlation of features.

Résumé

L'imagerie aérienne effectuée pendant les levés bathymétriques lidar constitue un ensemble de données de référence indépendant, pour la réalité de terrain. Le mosaïquage de l'imagerie aérienne requiert une intervention manuelle de l'opérateur, laquelle prend beaucoup de temps. Cet article présente une procédure de mosaïquage automatique qui permet d'obtenir une carte photographique continue et visuellement cohérente de la zone couverte. L'objectif de cette étude consiste à utiliser seulement les images de la caméra aérienne sans informations supplémentaires. Une comparaison entre les éléments dans la mosaïque résultante et une carte de référence montre que la mosaïque est visuellement cohérente et qu'il existe une bonne corrélation géométrique-spatiale des éléments.

Resumen

Las imágenes aéreas recogidas durante los levantamientos batimétricos efectuados con el lídar proporcionan una colección de datos de referencia independientes para la validación en el terreno. La composición de las imágenes aéreas en forma de mosaico requiere una cierta implicación manual por parte del operador, lo que toma mucho tiempo. Este artículo presenta un procedimiento para la composición automática en forma de mosaico, que crea un mapa fotográfico continuo y visualmente coherente de la zona representada en la imagen. El objetivo de este estudio es utilizar sólo los marcos de la cámara aérea sin información adicional. Una comparación entre las características del mosaico resultante y una carta de referencia muestra que el mosaico es visualmente coherente y que hay una buena correlación geométrico-espacial de las características.

For all types of hydrographic surveying measurements, it is important to have an independent reference dataset to provide control. The reference dataset supplies additional information and helps in troubleshooting problems that are encountered during the surveying process. Airborne lidar bathymetry (ALB) surveying is no exception to this concept. Aerial imagery typically is simultaneously recorded with a lidar survey as a part of the data collection. To date, the most common usage of the aerial imagery from lidar surveys is to provide ground truth to elevation measurements.

The term "mosaic" (also called photomosaic) refers to a composite of photographs created by stitching together a series of adjacent pictures of a scene (ASPRS, 2005). The mosaicing process is not new and there are several commercial off-the-shelf products that allow manual production of a mosaic. Although some commercial tools are automatic, they are not robust enough for all types of imagery and require initial manual involvement of the operator to allow other procedures to be automated. In this paper, an automatic procedure is described to create a continuous and visually consistent mosaic that consists of individual frames collected during a lidar survey.

In concept, a visually consistent mosaic is one that exhibits all the features observed in the frames without any artifacts due to image processing. The procedure uses only frames from the aerial camera without any navigation or attitude information from the aircraft's systems. The mosaic procedure should work on conventional image formats (e.g., bitmap format) and be efficient enough to include hundreds to thousands of individual frames in the process. The data source for the mosaic in the examples presented here is colour (RGB) imagery from a DuncanTech4000 digital camera onboard the U.S. Army Corps of Engineers' (USACE) lidar CHARTS system (Wozencraft and Lillycrop, 2006). The algorithms for this procedure were originally developed for mosaicing underwater imagery (Rzhanov et al., 2006). Similarly to the situation with underwater imagery, no navigation (i.e., positioning) or attitude (i.e., yaw, pitch, and roll) information is being used (e.g., positioning of the camera underwater is not available with the accuracy required for mosaicing) and only the frames are used.

Analog versus digital

In the past, an analog video system was used with the previous USACE ALB systems such as SHOALS-200 and SHOALS-400 (Figure 1a). Today, the current USACE SHOALS systems (e.g., SHOALS-1000 and SHOALS-3000) operate with a digital still camera (Figure 1b). The benefit in using the legacy analog video is that the relatively high speed of recording and high sequential frame overlap gives the operator the option to select the highest quality frames from the dataset. This benefit is not substantial when the legacy analog video frames used in the lidar survey are of low quality compared to the digital still-camera frames. Typically, the spatial resolution of the analog video is low. These systems have relatively slow radiometric adjustments to the varying light en-

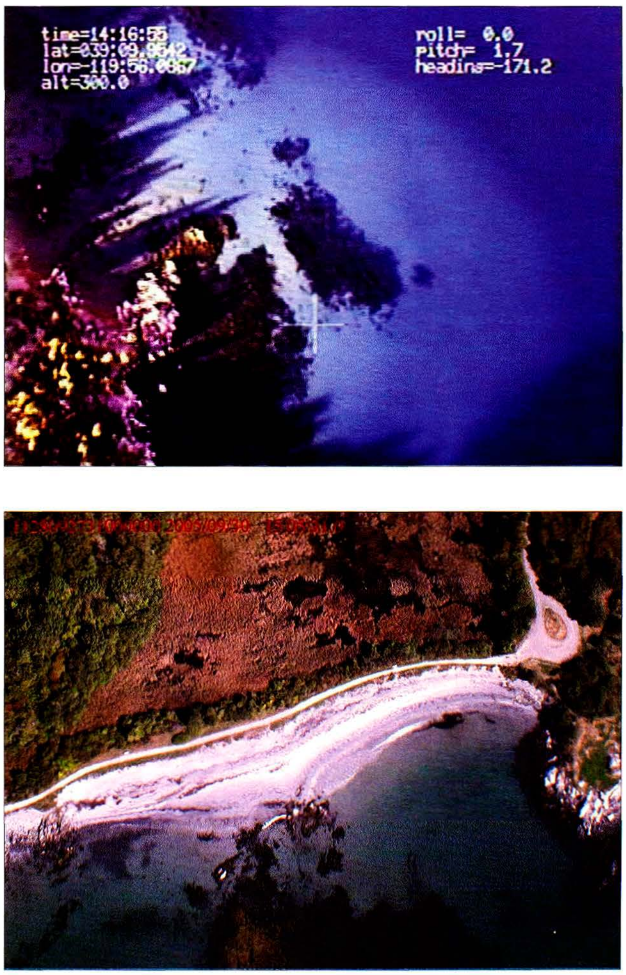

Figure 1: Sample frames of (a) the legacy analog imagery (SHOALS-400), Lake Tahoe, CA; and (b) the current Digital imagery (CHARTS: SHOALS-3000), Gerrish Island, ME.

tering through the aperture that degrades the qualities of the frames because of the optics. The analog data are re-sampled when converted to a digital format, further degrading the quality. In analog video, the overlay of the auxiliary information (e.g., time, navigation, and attitude) is imprinted on the image and substantially complicates the mosaicing procedure. The information overlay is detected as seafloor features in the mosaicing process and can be regarded as control points that need to be matched. The digital imagery has a smaller frame capture rate (~1 Hz), but has radiometric and geometric qualities that are higher than the legacy analog video images. Each digital frame is collected in three channels (i.e., RGB) and recorded to a calibrated colour scale. This is not the case with analog video frames that are re-sampled using channel-dependent spectral resolution. The spatial resolution may also be under-sampled in analog frames.

Because of the low quality of the analog video data, the mosaicing procedure developed here focuses on still digital imagery. The digital camera used in the study is a Duncan Tech4000 (DT4000) digital camera that has a 1600 (H) x 1200 (V) CCD imaging sensor (11.8mm x 8.9mm). Each frame is a 24-bit image compressed to a JPEG format by a Matrox frame-grabber. The JPEG compression is beneficial for the reduction of data through-put, but causes some loss of acquired data (ITU, 1993). All the frames used in the study were collected on an airborne platform moving at a speed of 70m/s at altitude of 400m.

Basic assumptions

Several assumptions are made for the mosaic process. The first assumption is a flat-earth assumption whereby the distance ratio between features in a frame stays constant. The flat-earth assumption is considered to be feasible if the ratio of variation in land elevation to aircraft altitude is less that 20 (Trucco and Verri, 1998). In a standard ALB survey, the altitude of the aircraft is usually 300 to 400m, whereas the land elevation varies on the scale of a few meters. Consequently, the land-to-altitude ratio in standard lidar survey is considered small enough that the land can be assumed flat. However, cases where the imaged area contains high topographic relief (e.g., coastal cliffs), errors will be introduced in the mosaic procedure.

The second assumption is that the type of motion in which the camera is advancing is known and can be modeled. This assumption is essential for the registration step that is discussed below. There are several camera-motion models that are commonly used to transform the image projection to the mosaic's plane. The choice of a camera-motion model depends on the tradeoff between the accuracy of registration results and required processing time. The models vary from the simplest 2-parameter translational model that allows the tracking of X and Y shifts. An 8-parameter perspective model allows relating overlapping views of the same planar surface (Hartley and Zisserman, 2004).

The third basic assumption is that the camera has been calibrated and the lens distortion is known. Some cameras have significant distortions associated with their optical systems. Correction for these distortions may drastically improve results of the mosaicing procedure. Additional pre-processing steps, such as removal of illumination pattern using contrast enhancement (Zuiderveld, 1994) and detrending (Rzhanov et al., 2000) can be applied to the imagery for enhancing the mosaic result.

Determination of Angles

Ideally, the frames received for photogrammetric processing should be taken by a camera with a vertical optical axis. However, photographs are often taken at an oblique camera angle. This can be due to an unstable platform or for special surveying considerations, as in the case of ALB surveying where the camera should simultaneously capture the same area surveyed by the lidar. In these cases, the frames first need to undergo a special procedure known as rectification and is defined as a photographic procedure by which a tilted aerial photograph is converted into one having no tilt (ASPRS, 2005). The tilts that occur in the aircraft, although kept to a minimum by the leveling of the camera system; do affect the position of objects on the photograph.

There are three tilt angles that correspond to the axis of motion of the aircraft (Figure 2). The swing angle (also called *kappa* (κ)) relates to the yaw motion that is a rotation in a horizontal plane about the normally vertical axis. The resulting image is the rotation of the photograph on its own plane about the photograph perpendicular. The x-tilt angle (also

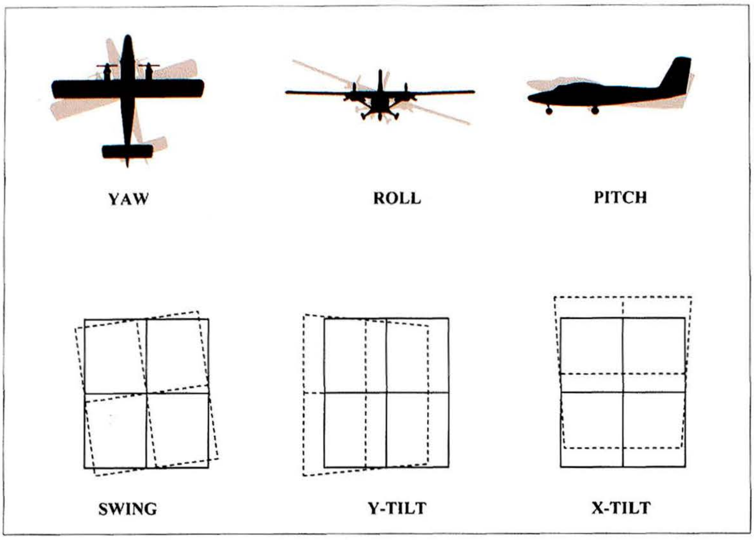

Figure 2: The axis of motion and the effect of corresponding tilts on the image.

called lateral tilt or *omega* (ω)) relates to the roll motion that is defined when the aircraft rotates about an axis that is aligned with the direction in which the aircraft is flying. The y-tilt angle (also called longitudinal tilt or *phi* (φ)) relates to the pitch motion and measures the degree to which an aircraft's nose tilts up or down. The resulting image from either tilt angles (ω or φ) results in a trapezoid-shaped area of ground coverage.

Since the frames used for mosaicing are from a planned airborne lidar survey where the plane's motion and path are monitored and controlled constantly, it is safe to assume that the axes of motion are approximately constant, although the result may not be zero. The aircraft is controlled to fly a planned flight course at a constant altitude and azimuth. The axes of motion fluctuate, but they remain around a constant bias. This assumption is also supported by long segment observations (e.g., more than 50 frames). Following the observation's results, it was decided that both the yaw and roll axes of motion can be considered as zero values. Only the pitch value has a non-zero bias that needs to be estimated for accurate image rectification. However, this is not typically done for most aerial surveys, but is used with the CHARTS system.

Mosaicing Procedure

To the knowledge of the authors, there is no commercial-off-the-shelf (COTS) software robust enough to process only conventional image formats without additional information to define initial conditions of registration. Navigation and attitude are usually required by COTS software, but there are cases where these data is missing or logged in a proprietary format and cannot be used.

Pre-processing

All still-digital frames from an ALB survey are encapsulated in a single movie file (e.g., Microsoft AVI format) that allows group processing of all the images. The overlay imprinted on the frames (e.g., time, navigation, or attitude) and other video-related artifacts are cropped to minimize errors during the pair-wise registration step in the mosaic process. In order to reduce the processing time, down-scaling of video frames is done according to compilation constraints on a final product resolution. It is important to note that the frame-size reduction affects the image resolution and should be used only in cases where the frame-size reduction does not prevent the operator to process the data efficiently.

Attitude correction

It is common in SHOALS ALB surveys that the camera is positioned at an offset pitch angle (~10°). This camera configuration allows the surveyed area to be imaged at the time of the laser measurement. Unlike the attitude variation of the camera during the survey that is considered insignificant, the offset pitch (and/or roll) angle may pose problems. If no attitude correction is performed, then the images are not projected to the same reference plane and inaccurate registration occurs between the frames. If registration is tried without attitude correction (i.e., ortho-rectification), then scale distortions will be added to the image and will affect the accuracy and consistency of the final product. The resulting mosaic without correction is a distorted image that shrinks or grows depending to the camera's pitch angle (Figure 3). The tilt of the camera was not known in this study, but was determined iteratively, where

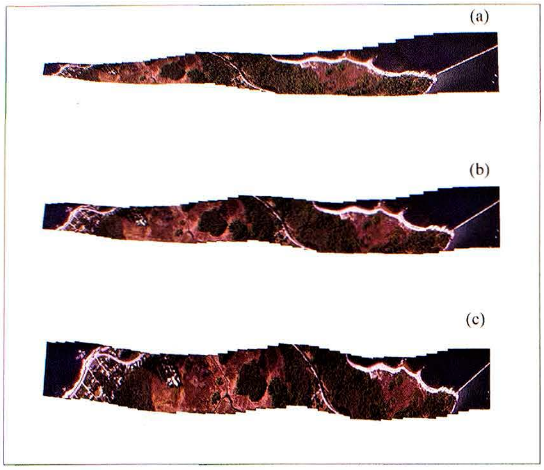

Figure 3: Mosaics produced at different pitch angle corrections: (a) no pitch correction, (b) 5° pitch correction, and (c) 10° pitch correction.

the value of the pitch angle was chosen according to the registration results that showed constant scale frames. The main assumption in this process is that camera attitude is changing slowly enough that the angle can be considered constant during the acquisition of two consecutive frames.

Pair-wise registration

The automated pair-wise registration step is the most CPU-intensive stage of the mosaicing procedure. The registration of one frame to the next overlapping frame depends on the presence of distinct features that appear in both frames. When features are present in both frames, it is possible to establish correspondence (i.e., matches) between points in both images. Locations of matching features allow the determination of a set of parameters for the chosen model of camera motion. In the case of featureless imagery, these parameters can be determined with the frequency-domain registration technique (Reddy and Chatterji, 1996). The harmonics in the two-dimensional Fourier spectrum of the images define the camera motion between the frames. Experience with underwater imagery indicates that the optimal camera-motion model is a 4-parameter rigid affine model (i.e., corresponds to a similarity transform) (Rzhanov et al., 2006). Although aerial imagery allows the use of an 8-parameter perspective transform, an additional step of ortho-rectification is required that adds more processing time and does not provide substantial benefits to the final mosaic product. The four parameters of the camera-motion model are X and Y shifts (lateral camera motion), scale (vertical camera shift perpendicular to the imaged surface) and rotation (about camera optical axis). At least two matches are required to solve for the parameters for this model. Larger numbers of matches require a least-squares solution and usually provide more accurate results. The pair-wise registration is done for all sequential frames for each ALB flight line and also for overlapping frames in the neighboring lines. A candidate for overlap between temporally non-sequential frames can be found iteratively or by using available geospatial information.

Global alignment

Quite often, several survey lines have overlap of frames. Pair-wise registration of overlapping frames introduces errors. The cumulative positioning error of the frames grows in the direction of the mosaic construction. The values of the errors can be spatially distributed using the technique known as global alignment (Sawhney et al., 1998). Corrections are applied so that the coordinate values of all points in the network will be consistent (ASPRS, 2005). The optimal alignment is the one with the minimum sum of squared distances between projections of the same features from different frames.

Blending

Once the location of each frame is determined for the mosaic, radiometric adjustments are required for a seamless transition between the frames. Simple stacking of the frames one on top of the other can generate noticeable artifacts. Improper calibration or variation in the illumination during the survey can cause some frames to have different colour values to the background. This problem is treated to some degree in the pre-processing stage (Section 4.1). Other problems such as small misalignments of features or discrepancies in background colour values are addressed by the blending procedure. There are several approaches used for blending (Boykov et al., 2001; Uyttendaele et al., 2001; Agarwala et al., 2004; Szeliski, 2004; Fattal et al., 2004). The graph-cut technique was used (Boykov et al., 2001) in this study. This approach finds a path with a minimum difference (in pixel values) between two overlapping images and stitches them along this path. The final result after completing the blending step is the mosaic.

Error sources

Error sources in the mosaic procedure can be encountered at any stage in the mosaic construction. Common errors can be grouped in three main types. The first type of error is caused by violation of one or more of the basic assumptions for the mosaicing procedure. Any deviation from flatness in the imaged scene causes parallax-related artifacts in the final mosaic. The horizontal location of features within frames with this violation will not be preserved and this will affect the registration accuracy. Another assumption that may be violated is modeling the camera motion. If the camera-motion model is too simple to correctly describe the camera's motion, then problems will occur at the registration step. It is important to note that if a complex model uses too many free parameters, it may significantly increase the processing time.

The second type of error is due to the amount of overlap and lack of details in the frames. The optimal overlap between two neighboring frames is in the range of 50% to 80%. The main problem with a smaller overlapping area between the frames is the number of available features to match frame-to-frame. The lack of sufficient overlap may affect the accuracy of the extracted registration parameters. Another possibility for error is when the images do not contain any features. This is a case where registration problems will occur using either the feature or the featureless registration technique.

The third type of error is radiometric distortion. An example of this would be a moving shadow error, typically found in areas that contain tall features. This error is similar to the violation of the flat-earth assumption. In the registration process, shadows may be treated as a feature. Registration of frames that were collected at different times of day may be problematic because of a shift in location and shape of a feature's shadows. A change in the sun's location or an overcast sky will change the shadows' dimensions and affect the registration process, accordingly. Another source of errors is artifacts in the frames. Images that contain either a permanent overlay (e.g., grid or permanent occlusions) or contain a shadow from an obstacle in the sky (e.g., airplane or cloud) also affect the registration. The last error in this group is radiometric differences in illumination due to a change in flight direction with respect to the sun, overcast sky, or camera calibration problems that will affect the registration and the blending stages.

Example Results

The mosaicing procedure was investigated on a case study in New Castle Island, New Hampshire. The frames were collected by USACE lidar survey on 3 October 2005 in the Portsmouth Harbour area using the DuncanTech DT4000 digital camera. New Castle Island contains a mixture of land and coastal features that includes both man-made (e.g., houses, roads, docks, and a fort) and natural features (e.g., trees, marsh vegetation, sandy and rock shores, ledges, and aquatic vegetation). The aerial imagery was processed following the mosaic procedures mentioned above and the final mosaic consists of 107 frames. The resolution that was required for the mosaic was defined by the goal of the study. It was decided that a 1:20,000 NOAA chart was adequate for evaluating the mosaic and the sensitivity of the produced mosaic to the camera attitude. Following the required resolution, a mosaic with a sub-meter pixel-resolution would more than suffice. The frames were down-scaled by a factor of 2 to 800 × 600 pixels (about 40cm × 40cm per pixel). The pitch angle at the time of the study was not known. An inspection of the constructed mosaic shows a strong spatial distortion indicating that no pitch correction was done on the frames (Figure 3). The average camera attitude was estimated from the imagery and a pitch-angle value of 10° was calculated. In order to evaluate the influence of the attitude correction on the mosaicing process, the pitch-angles values of 0°, 5°, 7°, 10°, 11°, and 15° were used to ortho-rectify the frames. Some errors in the registration using the pair-wise registration occurred mostly in the frames that consisted of small forests. These errors were corrected by manually identifying the matches in each pair of frames.

The resulting products are six mosaics: one mosaic for each pitch-angle value. Two of these mosaics are shown in Figure 4: 1) a mosaic without an attitude correction (pitch angle 0°), and 2) a mosaic with an attitude correction of 10°. All the mosaic

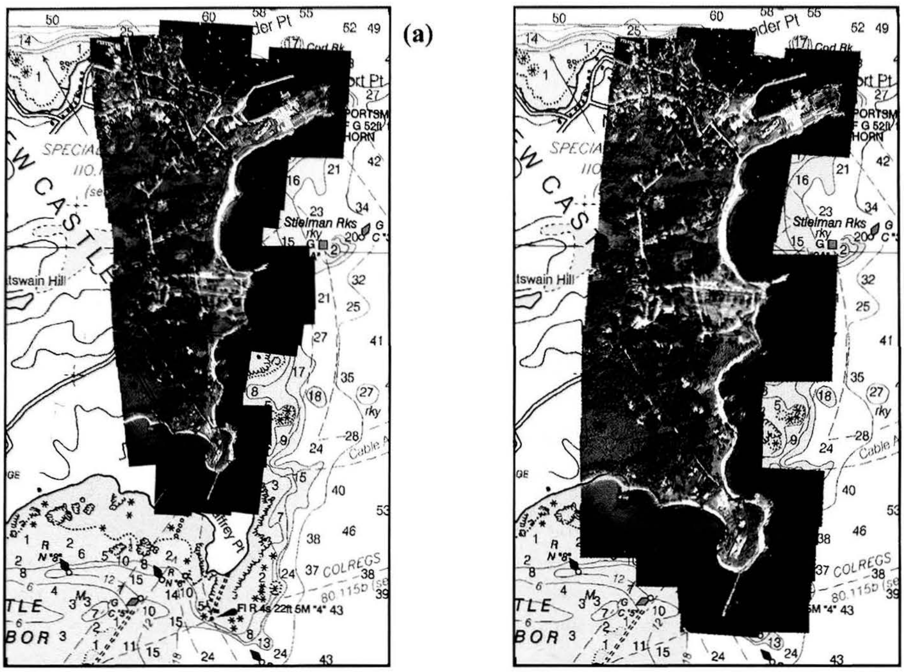

Figure 4: Final products from mosaic procedure: (a) without an attitude correction and (b) with a 10° attitude correction.

products were inspected visually and compared to a NOAA navigation chart (Figure 4). The mosaics were matched to the chart using a rigid affine model. The parameters extracted from the matching were X and Y shifts, scale, and rotation. The success of the mosaic procedure for each pitch angle was evaluated using two methods: match-point location and minimum-area difference. In the first method, a root-mean-square comparison of the match point from each mosaic was compared to the NOAA chart. In addition, a second comparison of the shorelines was done. The shoreline from each photomosaic product was digitised and compared to the digitised shoreline from the chart. The area between the two digitized shorelines was calculated and a comparison was made of defined shoreline features. The results of both comparisons were normalised to the largest value and are presented in Table 1.

The best-fit result from both comparisons was the 10° pitch-corrected mosaic. These results also correlated with the pitch-angle value in the attitude calculation step. Apart from the pitch angle, another error source that was noticed was the plane's shadow in the frames and sun-glint on the water surfaces that affected being able to obtain a visually consistent image in some areas.

Corrected pitch angle (degrees)	Matching point comparison	Area comparison
0	0.76	**0.82**
5	1.00	**1.00**
7	0.44	**0.61**
10	0.36	**0.43**
11	0.41	**0.48**
15	**0.42**	0.45

Table 1: Evaluation of the mosaic products as a function of pitch angle correction.

Discussion

The results of the study show that the described mosaic procedure is suitable for the aerial imagery received from the CHARTS' DuncanTech DT4000 digital camera, and can be applied to any standard still-airborne camera. It is important to know and correct the pitch angle of the digital camera. A schematic illustration of the suggested mosaic procedure is shown in Figure 5.

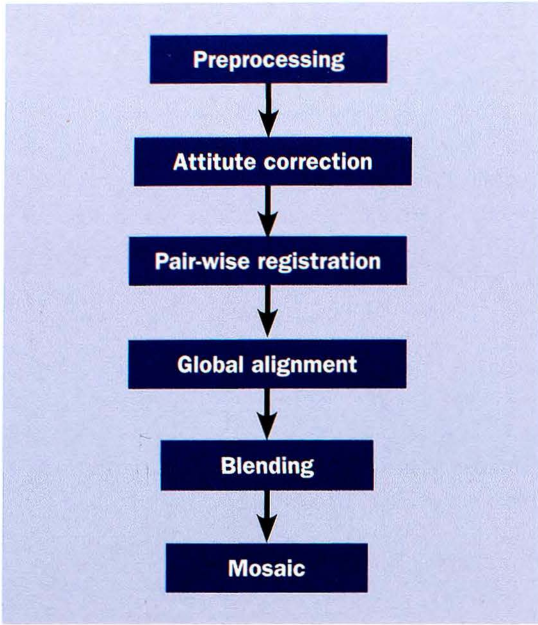

Figure 5: Schematic illustration of the mosaic procedure developed in this study.

An important aspect in the mosaicing procedure is the processing time that determines the feasibility for applied use. The processing-time calculations here were based on the use of the DuncanTech camera image frames. It was assumed that in a typical lidar survey about 100 frames will be used per site (2 km X 2 km). A PC with a 3GHz processor and 1GB of RAM is a typical computer workstation. It is assumed that the camera is calibrated and the lens distortion is corrected. In the first part of the procedure, the processing works on single frames and pairs. The first step is the attitude correction step that includes the re-projection of the frames and cropping the null pixels (no data areas). The time estimate for the attitude-correction step is about 0.3s/frame that results in a total of 0.5min of processing. The pair-wise registration step includes match-point extraction and calculations of the registration parameters. The time for this step is about 5.5s/pair that results to 9.5 min of processing. The second part of the procedure is the processing of the whole-image sequence. The global alignment includes iterative calculations of the registration coefficients between all the frames in the mosaic. The time estimate for the global alignment step is about 0.5min. The blending process includes merging of the frames and takes about 1min. To summarise, the estimated time to process 100 frames to an orthorectified mosaic takes about 11.5min (Table 2). The time of processing can be shortened. The main labour-intensive parts are quality control and search for non-sequential overlap. Both problems can be solved by using positioning information, even if they are not very accurate. The quality of the registration can be checked against the geo-information, so the automatic system will flag any inconsistency. This will create a more efficient procedure where non-sequential overlapping frames and some of the problematic sequential overlapping frames can be identified without operator interaction. It is important to mention that the failure of the automatic registration procedure will require manual processing. This will lengthen the processing time depending on the operator's skills.

Step	**Time (min)**
Reprojection	**0.5**
Pair-wise registration	**9.5**
Global alignment	**0.5**
Blending process	**1**
Total	**11.5**

Table 2: Time estimation of modules developed for the mosaic procedure. The estimate is based on 100 frames using a PC with 3 GHz processor and 1 GB of RAM.

Summary

We describe an automatic mosaicing procedure that uses ALB aerial photographs to create a continuous and visually consistent photographic mosaic of the lidar-surveyed area. The goal of this study was to use only the frames of an aerial RGB DuncanTech camera without any additional information. In order to assess if the mosaic produced by this procedure was successful, the following questions must be addressed:

1. Is the mosaic visually consistent?
2. How sensitive is the final product to the camera attitude?
3. Were the camera's tilt angles successfully estimated?
4. How well does the final product compare to a navigation chart?

Each frame was 1600 × 1200 pixels. The procedure that was developed includes: pre-processing, attitude correction, registration of two images, global alignment, blending and building the mosaic. The estimated time for this process is 11.5 minutes for 100 frames. A case study in New Castle Island, NH found the main error sources were from the 10° pitch correction that was required to all the frames and the plane's shadow in the frames. The final product was inspected visually and compared to a NOAA chart. A sensitivity test of the final mosaic product to the camera's pitch angle shows that knowledge of the tilt of the camera may simplify the procedure and shorten the processing time. The final mosaic is continuous and visually consistent and shows good correlation with features on the NOAA chart.

The final mosaic can be ortho-rectified and geo-referenced using ground control points. The main labour-intensive parts in this procedure are quality control and the search for non-sequential overlap. Both problems can be solved by using positioning information. The registration quality can be checked against the geo-information, so the automatic system will flag any inconsistency. Non-sequential overlapping images can be identified without operator interaction.

This procedure can provide visual background in a relatively short period of time using the aerial imagery that is collected simultaneously with ALB survey. The mosaic itself can provide a reliable reference tool for the lidar survey. Apart from ground truth, the image photomosaic can also aid in assessing the shoreline location and identifying the location and cause of unsuccessful lidar measurements in the survey.

References

Agarwala, A., M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen, 2004. Interactive digital photomontage, *ACM Transactions on Graphics*, v.23, n.3, pp. 292-300.

ASPRS, 2005. *Manual of Photogrammetry*, Eds., McGlone, C., E. Mikhail and J. Bethel, ASPRS Publications, 5th edition, pp. 1168.

Boykov, Y., O. Veksler, and R. Zabih, 2001. Fast approximate energy minimization via graph cuts, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, v.23, n.11, pp. 1222-1239.

Fattal, R., D. Lischinski, and M. Werman, 2002. Gradient domain high dynamic range compression, *ACM Transactions on Graphics*, v.21, n.3, pp. 249-256.

Hartley, R. and A. Zisserman, 2004. *Multiple view geometry in computer vision*, 2nd Ed., Cambridge University Press, pp. 655.

ITU, 1993. Terminal equipment and protocols for telemetric services: Information technology, digital compression and coding of continuous-tone still images – requirements and guidelines, *International Telecommunication Union - recommendation* T.81 (09/92), pp. 186.

Reddy, B.S. and B.N. Chatterji, 1996. An FFT-based technique for translation, rotation, and scale-invariant image registration, *IEEE Transactions of Image Processing*, v.5, pp. 1266-1271.

Rzhanov, Y., L.M. Linnett, and R. Forbes, 2000. Underwater mosaicing for seabed mapping. In *Proceedings IEEE International Conference on Image processing*, v.1, pp. 224-227.

Rzhanov, Y., L. Mayer, S. Beaulieu, T. Shank, S.A. Soule, and D.J. Fornari, 2006. Deep-sea geo-referenced video mosaics. *Oceans 2006 conference proceedings*: Abstract 060322-11.

Sawhney, H. S., S. Hsu, R. and Kumar, 1998. Robust video mosaicing through topology inference and local to global alignment, in *Proceedings of the 5th European Conference on Computer Vision*, Vol. II.

Szeliski, R., 2004. *Image alignment and stitching. A tutorial*. Technical report MSR-TR-2004-92, Microsoft Research,.

Trucco, E. and A. Verri, 1998. *Introductory Techniques for 3-D Computer Vision* Prentice Hall, Inc., pp. 343.

Uyttendaele, M., A. Eden, R. Szeliski, 2001. Eliminating ghosting and exposure artifacts in image mosaics, in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, v. 2, pp. 509-516.

Wozencraft, J. M. and J. W. Lillycrop. 2006. JALBTCX coastal mapping for USACE, *International Hydrographic Review*, v. 7, pp. 28-37.

Zuiderveld, K. 1994. Contrast Limited Adaptive Histogram Equalization, in *Graphics Gems IV*, Academic Press.

Acknowledgment:

This work was supported by Tyco Fellowship for Ocean Mapping and UNH/NOAA Joint Hydrographic Center grant NA05NOS4001153. The authors wish

to thank the Joint Airborne Lidar Bathymetry Technical Center of Expertise (JALBTCX) for providing the data for the study and J. V. Gardner from the University of New Hampshire for his numerous discussions and support.
The co-authors wish to acknowledge the thorough review and useful comments provided by Michael Casey and Jennifer Wozencraft on an earlier version of this paper.

Biography of the Authors

Shachak Pe'eri received the B.Sc, M.Sc, and Ph.D. degrees from the Tel Aviv University, Israel, Ramat Aviv, Israel, in 1996, 1997, and 2005, respectively, all in geophysics. He is currently working a research scientist in the Center of Coastal and Ocean Mapping, University of New Hampshire, Durham (NH, USA). Research activities have focused on experimental and theoretical studies of airborne lidar bathymetry.
Email: *shachak@ccom.unh.edu*

Yuri Rzhanov received the Ph.D. degree in Semiconductor Physics from the Russian Academy of Sciences, in 1983. Before joining the University of New Hampshire in 2000, he has been working at the Heriott-Watt University in Edinburgh, Scotland. Currently he is an Associate Professor at the Center for the Coastal and Ocean Mapping / Joint Hydrographic Center (CCOM/JHC) at the University of New Hampshire. His research interests include optical methods of seafloor mapping, blending techniques for construction of photomosaics from imagery acquired underwater, seabottom structure reconstruction from multiple views.
E-mail: *yuri@ccom.unh.edu*