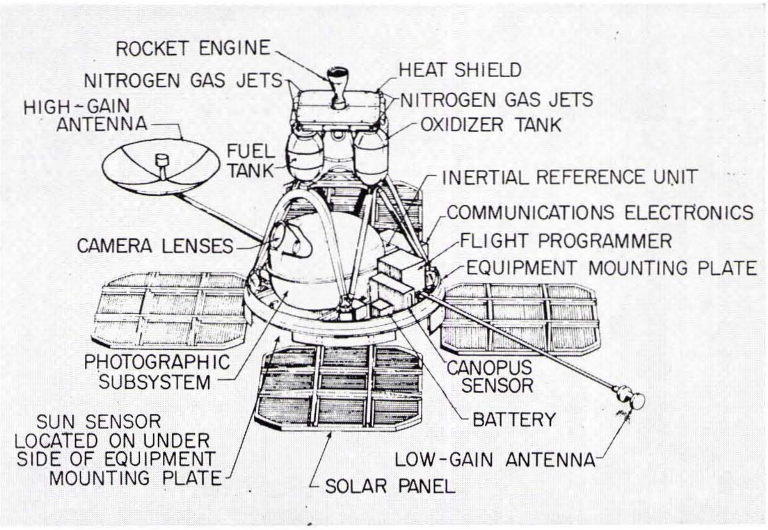

FIG. 1. Spacecraft configuration.

LEON J. KOSOFSKY*
National Aeronautics & Space Admin.
Washington, D. C.

Topography from Lunar Orbiter Photos

Landing sites to be studied from photos on film transmitted to earth by video from an unmanned spacecraft orbiting the moon.

(Abstract on next page)

INTRODUCTION

MISSION OBJECTIVES

NASA's LUNAR ORBITER PROGRAM is scheduled to put a series of unmanned spacecraft into orbit around the moon, beginning in 1966. The primary objective of the program is to produce information about the topography of selected portions of the lunar surface. The principal payload is therefore a photographic system.

Topographic information can serve many purposes. The specific needs which the Lunar Orbiter was designed to help fill are the selection of the safest possible landing sites for the Apollo manned landing missions and the determination and location of all surface hazards within the selected areas. Another unmanned NASA program—the Surveyor Program—is designed to soft-land instruments which will measure critical surface characteristics, such as bearing strength, at the point of landing. The Lunar Orbiter photography must provide the basis for extrapolating this point information to the surrounding area and to apparently similar areas.

TYPICAL MISSION

The spacecraft, weighing about 840 pounds (Figure 1), will be launched toward the moon

---

* Presented at the Semiannual Convention of the American Society of Photogrammetry, Dayton, Ohio, September 1965.

Abstract: *NASA's Lunar Orbiter program will put a series of unmanned photographic spacecraft into orbit around the moon beginning in 1966. The primary objective is topographic information to permit the selection and examination of lunar landing sites. A dual camera will produce simultaneous high-resolution monoscopic and medium-resolution stereoscopic exposures on film which will be processed on board and read out for transmission to earth. A combination of photometric and photogrammetric techniques will be required to extract the desired data from the photographs.*

from Cape Kennedy by an Atlas-Agena launch vehicle, and will place itself into lunar orbit by a retro-firing of its velocity control engine. After about three days in its initial orbit, it will transfer to the final orbit, which has a nominal altitude of 46 kilometers at perilune and 1850 kilometers at apolune. For the initial missions, the orbit will be at a low inclination to the equatorial plane. In order to achieve the desired ground resolution, photography must be accomplished at or near perilune. It is also necessary that the sun be between 10° and 40° above the horizon during photography.

The spacecraft ordinarily flies with its solar panels and sun sensor facing the sun, and with its star sensor locked on the star Canopus. This keeps the parabolic high-gain antenna pointed at the earth. It keeps this orientation even when it is in the moon's shadow, thanks to its Inertial Reference Unit. In order to take a sequence of photographs during any orbital pass, the spacecraft is reoriented so that the camera will be vertical at the middle of the picture sequence. After each picture sequence, the spacecraft returns to its sun-Canopus orientation.

As the moon's rotation brings the desired target areas under the orbit, additional picture sequences can be taken. The film is processed within the photographic package shortly after exposure, and the processed film is stored on the takeup reel until it can be read out for transmission to the Deep Space Network receiving stations on earth.

TOPOGRAPHIC INFORMATION REQUIREMENTS

The Apollo Program needs some very special kinds of topographic information from the Lunar Orbiter photography. The Apollo Lunar Excursion Module must land where the local slopes will not cause it to topple over, and where local obstructions will not damage its structure or equipment. Since the LEM's four feet are only about seven meters apart, and provide less than a meter of effective ground clearance, the focus of interest is on minute topographic features. There is no vital interest in absolute elevations above a datum. Fairly large tilts or warpages in the assumed datum can be tolerated, since they add only a small gravity vector increment to the local slopes produced by fine-scale roughness of the ground.

It will require a unique combination of photogrammetric techniques to extract topography at that level of detail from the Lunar Orbiter photography. A brief exposition of the nature of the photography is necessary to an understanding of these methods. The following information has been extracted from a more complete treatment of the Lunar Orbiter photographic system. (Reference 1)

Nature of the Photography

PHOTO COVERAGE PATTERNS

To begin with, the Lunar Orbiter's photographic system carries a dual framing camera. Each exposure puts two images on different portions of the single roll of film. At the nominal altitude, with the camera vertical, the wide-angle image covers a rectangular area about 32×37 km, at an effective ground resolution of 8 meters. The central part of this coverage, a narrow rectangle about 4×16 km, is also covered by a high-resolution image at a 1 meter ground resolution (Figure 2).

The camera builds up coverage in the line of flight by taking overlapping exposures. There are two overlap modes available. In the Low Repetition mode, the exposure interval produces an overlap of a little over 50 per cent on the wide-angle exposures. This leaves gaps between successive high-resolution exposures. In the High Repetition mode, there is a 5 per cent forward lap between successive high-resolution exposures—just enough to prevent gaps in this monoscopic coverage. Since there is always a simultaneous wide-angle exposure, these are overlapped some 87 per cent. In either mode, we have a choice of 4, 8, or 16 successive exposures in a strip. By photographing on successive orbital passes,

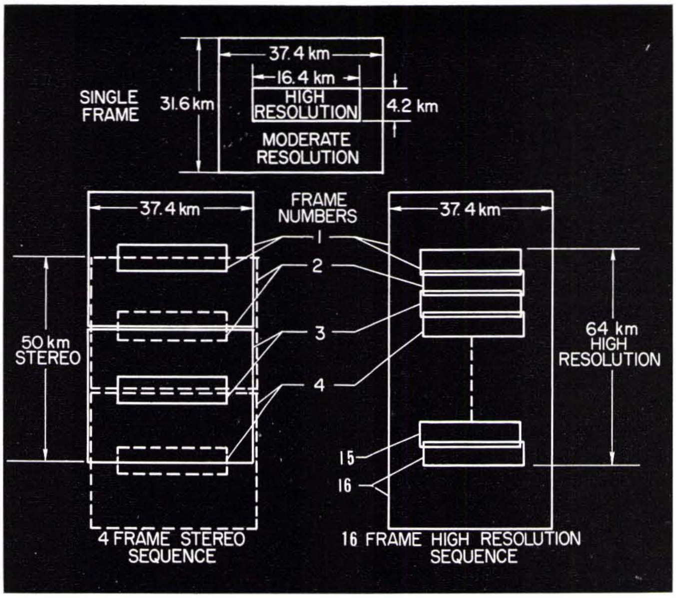

FIG. 2. Frame format and typical sequences.

we can build up blocks of sidelapping strips, if we wish to.

CAMERA DESCRIPTION

The entire photographic system, weighing about 135 pounds, is contained in a pressure shell. The camera makes its exposure through two quartz windows in the shell. (Figure 3 is an artist's rendition of the camera portion of the system.)

The high-resolution lens is a 24-inch *f*/5.6 lens especially built by the Pacific Optical Company. The optical path is folded by the mirror, placing the image on the platen. There is a focal plane shutter, with speeds of 1/25, 1/50, and 1/100 of a second. The platen moves during exposure to provide image motion compensation.

The other lens is an 80 mm Schneider Xenotar which has been stopped down to *f*/5.6. Its between-the-lens shutter has the same speed selection, and the platen also has IMC. This is called the wide-angle lens of the system, although it only has to cover a half-angle of 27° at the corners.

With the very slow film being used (Kodak

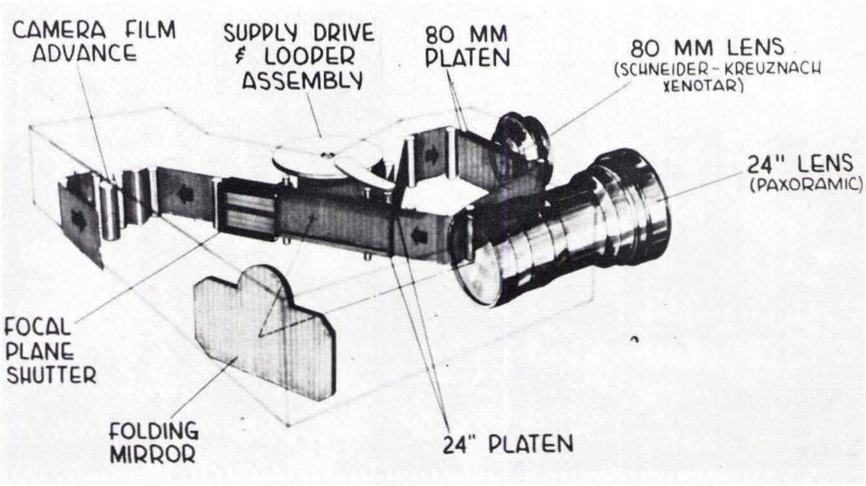

FIG. 3. Spacecraft camera (artist's rendition).

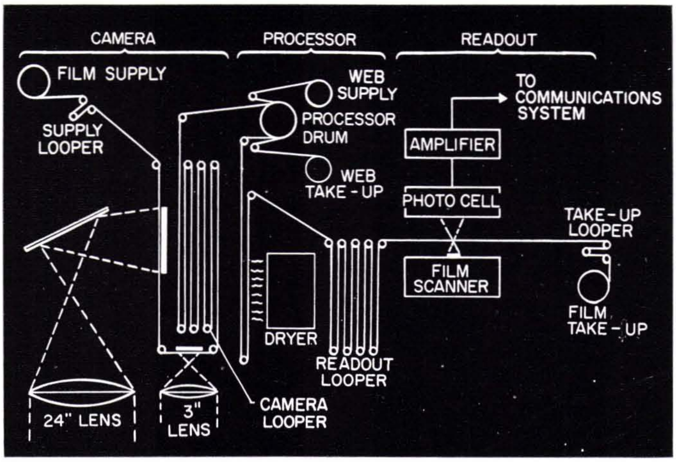

FIG. 4. Spacecraft photographic system schematic drawing.

S.O. 243), the only way to get the required ground resolution is to provide an accurate compensation for image motion. The *V/H* sensor (not shown) steals an unused part of the 24-inch lens' image to measure the image motion and control the motion of the platens. The rate information also goes to the exposure interval controller.

The film supply is a roll of unperforated 70 mm film, sufficient for 194 exposures (i.e., dual frames).

The film is processed within the photo system package (Figure 4) using the Bimat process. This is a diffusion-transfer process basically similar to what goes on in a Polaroid Land camera. (Reference 2)

READOUT

From the processor, the film is wound on the takeup spool. During readout, the film is pulled slowly backwards through the readout gate.

The readout schematic diagram (Figure 5) bears closer examination, since the characteristics of the readout system have a strong influence on the kind of photogrammetric data reduction that is feasible.

The function of the readout system is to scan the film to convert a sequence of image densities into electrical signals. The Lunar Orbiter film has a high density of information —the required ground resolution corresponds to 76 line pairs per millimeter on the film. The way the designers chose to read out at high resolution was to look at very small segments of the photo with a greatly demagnified flying spot.

In this system, the flying spot comes from a Line Scan Tube. The electron beam moves in only one direction, producing a line which is about $2\frac{1}{2}$ inches long at the phosphor. The scanner lens demagnifies this line to a length of 1/10 of an inch and focuses it on the film. The scanner lens is driven mechanically in the transverse direction (vertically on the schematic drawing), moving the scanning line across the film. The width of the film is covered by nearly 17,000 horizontal scans of the beam. The light transmitted by the film goes through the collecting lens to the photomultiplier tube, providing a voltage signal to the communications system.

When the scanner lens comes to the end of its travel, the film is advanced 1/10 of an inch and the lens then scans the next segment in the reverse direction. This scanning procedure breaks up each dual frame exposure, which occupies about a foot of film, into nearly 120 segments. We call these segments framelets.

It takes 40 minutes to read out each exposure. The rate is set by the video band-

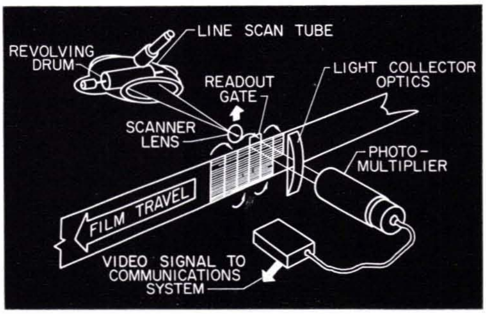

FIG. 5. Photographic system readout schematic drawing.

width available for transmission from the moon, which is 230 kilocycles.

There will be three ground stations equipped to record the Lunar Orbiter video transmissions. Each station has an 85-foot antenna and two sets of photo Ground Reproduction Equipment. After the video signal is separated from the carrier wave, it is fed to a kinescope tube. The electron beam in the kinescope moves only in the horizontal direction, putting a line on the 35-millimeter recording film. Continuous motion of the recording film builds up the framelet along the length of the film.

The film coming out of the Ground Reproduction Equipment is the primary replica of the spacecraft film. Its format is the succession of readout framelets, magnified about seven times (Figure 6). Alternate framelets are reversed end-for-end, due to the reciprocating travel of the scanner lens in the readout mechanism.

At this point, a faint-hearted photogrammetrist might just throw up his hands. It's hard enough to be told that you're not going to recover the camera film, without having to learn that the replica you will get only comes in tiny pieces. However, these difficulties should properly be considered as challenges to our ingenuity, rather than as insurmountable barriers. The measures available for recovering the geometry of the original exposure will be discussed further in this paper.

In addition to photogrammetrists, most users would like to look at larger pieces of the picture than just single framelets. There is an additional piece of ground equipment, called a reassembly printer, inherited from an Air Force project. It prints a mosaic of 14 framelets onto one piece of film. It automatically reverses the alternate framelets, and it prints adjacent framelets in very good register, at least to the unaided eye.

We call this composite of 14 framelets a subframe. In printing it, the reassembly printer reduces the scale slightly, so that the subframe is $6\frac{1}{2}$ times the scale of the spacecraft film.

The subframe is probably the most convenient format for general-purpose use. The scale permits veiwing without further magnification, and the size is manageable. However, a subframe by itself may not be very useful to the photogrammetrist, because it is only about half of a medium resolution frame. One has to mosaic two or three of them to make a frame. In addition, the matching of detail from framelet to framelet may be good enough for the eyeball, but not for accurate stereoplotting, unless some provision is made for correcting the position of each framelet.

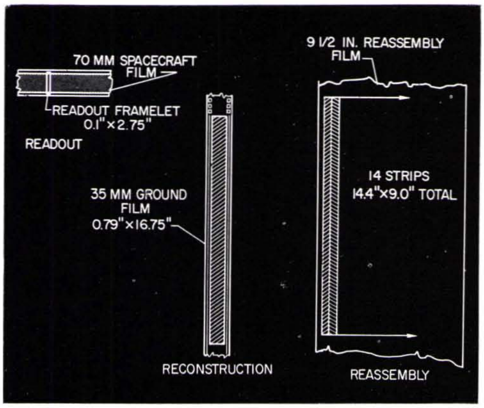

FIG. 6. Geometry of ground reconstruction and reassembly.

METRICAL CHARACTERISTICS OF THE MEDIUM-RESOLUTION PHOTOGRAPHS

Since we do intend to do stereo-triangulation and stereoplotting with the medium-resolution photographs, it would be well to look at the camera's basic metrical characteristics.

The camera is rigidly constructed, with image-motion compensation, a between-the-lens shutter and very high-performance lens, so that it does have the essentials of a good metrical camera. The angular coverage provides a base/height ratio of about one-third for pictures taken in the Low Repetition mode. This can be increased somewhat for pictures taken in the High Repetition mode, if the user is willing to work with segments of stereopairs. There is, then, a respectable stereo base.

The lens is not corrected for radial distortion. However, as far as we know, the distortion curve is regular, and we expect to obtain a good distortion calibration for each wide-angle camera. In summary, then, the negative roll in the spacecraft should have quite good metrical characteristics. All that is necessary is to correct for the distortions introduced by the readout process.

The straightforward answer to a problem of this kind is a camera reseau. One wants to put a set of precisely located marks on the picture before it gets cut up, so that he will know how the little pieces should be put together afterward.

The reseau in the Lunar Orbiter's camera (Figure 7) is a compromise between what a reseau really should be and the overwhelming

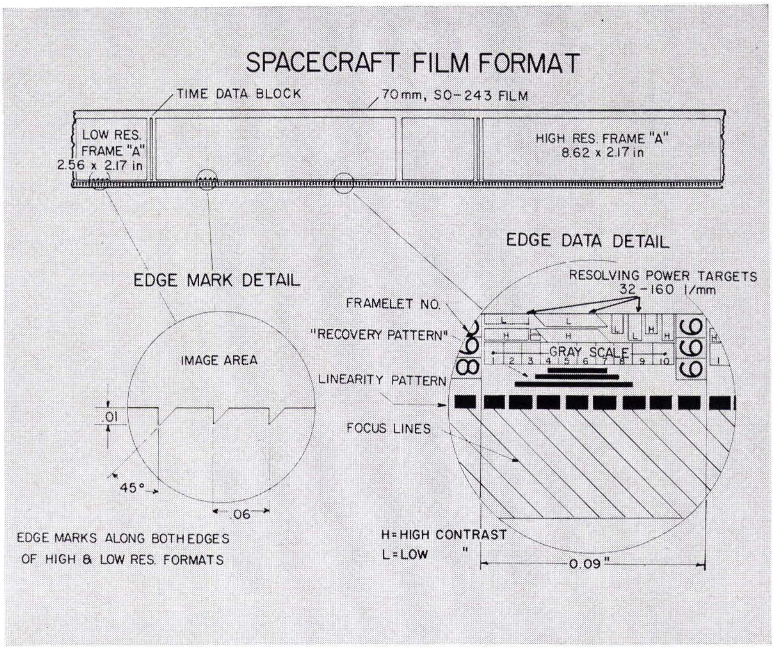

FIG. 7. Spacecraft film format and details.

need that a spacecraft on a lunar mission has for reliability. In this case, reliability considerations precluded the use of flashing light sources on a moving platen. The reseau is therefore limited to edge indentations, illuminated only by the natural light of the exposure. What we have is a sawtooth pattern along two edges of the picture. The pitch is fine enough so that after readout every framelet will have at least one reseau mark at each end. The reseau marks will appear on the reconstructed framelet and also on the reassembled subframe.

The preflight calibration of the cameras will include exposure of the reseau pattern on stable-base film, so that the location of each mark in the original format ought to be known within five microns.

We can assume, then, that we have tied down both ends of each framelet. If the measured length of a reconstructed framelet is different from the calculated value, we can apply a linear scale factor. This raises the question of possible nonlinearities in the relationship between the scanner lens travel in the spacecraft and the film drive on the ground. The best way to tie these down would be reseau marks inside the format. The second best way would probably be a set of reseau marks along the other two sides of the frame. Since both of these were ruled out by reliability considerations, the Lunar Orbiter photo system has a third means.

Figure 7 shows a series of time data blocks on the film. Spacecraft time is printed on the film with each exposure in binary code, by means of a set of 20 lights. The positions of these lights can be precisely located during the camera calibration exposures. The time blocks will of course be read out along with everything else on the film. This means that one framelet in every 120 is a record of known distances closely spaced along the framelet.

That should take care of the nonlinearity problem. It still leaves the possibility of random variations in the two drives. Our once-per-120-framelet sampling will pick up long-term drifts in the rates, but not haphazard changes. Here it will be necessary to rely heavily on the quality of the two mechanical drives, backed up by preflight testing. They are precision drives, and the variations are expected to be very small.

The discussion so far applies mainly to the accuracy of distances across the film, which is the direction of the stereoscopic parallax. Larger errors can be tolerated in the other direction, provided they aren't cumulative. It is clear that these errors will not build up from one framelet to the next, because of the reseau marks. Errors will exist within the width of a framelet, because the electron beam scans in the Line Scan Tube and in the kinescope are subject to nonlinearity and to drift. To correct for these, we will rely on the pre-exposed data running along one edge of the spacecraft film. This edge strip contains sensitometric step exposures, resolution patterns and linearity markings, and it is read out on every framelet, along with the picture. Comparison of the pre-exposed distances between selected markings with the distances measured on the reassembled subframes will provide the basis for corrections in the direction of the electronic scans.

By all these means, we hope to minimize the losses in geometric precision that arise from our inability to bring back the original film. Since we haven't done any testing yet, it is premature to say just what the loss in precision will be. It is the writer's estimate that we will be able to supply enough information to permit the photogrammetrist to recover the original films' geometry within 10 microns in the stereo direction, and somewhat more in the transverse direction. Of course, when the photogrammetrist knows what corrections to apply, he still has the laborious task of actually applying them. The final assessment will probably be that the scanning procedure costs a little bit in accuracy and a good deal more in labor and inconvenience.

Stereoscopic Data Reduction

Now that we have our reassembled stereo pictures and the data for correcting their geometry, we are ready to begin the triangulation step. We would like to bring all the photos in a strip to a common scale and adjust the strip to the control. If some of the Lunar Orbiter missions photograph blocks, we would like to do a block triangulation. The current plan for the first mission, however, is to photograph individual strips.

In the opinion of the writer, the nature of the reassembled pictures makes the analytical methods of triangulation inherently more convenient than the instrumental methods. It appears so much simpler to correct the plate coordinates mathematically than it would be to try to make precisely corrected plates.

To what kind of control will these strips be adjusted? The only existing control on the moon comes from measurements of astronomical camera plates. The Army Map Service has been at work for several years on the adjustment of a net containing some 250 measured points. Additional work has been going on at the Aeronautical Chart and Information Center, which should bring some additional points into a single, unified net. At that density, it is unlikely that many of the Lunar Orbiter's strips will cover existing control points. NASA is negotiating arrangements with the mapping agencies to put in new control points to order, as soon as the Lunar Orbiter photography is available. These points would be measured on the astronomical plates, and adjusted to the existing control points which surround them.

This control has obvious limitations. There is first the problem of identifying a unique point on a picture with 8-meter resolution when the point itself was originally located on pictures with 800-meter resolution. Then there is the related problem of the accuracy of the primary control. This is estimated to vary from ±500 meters near the center of the visible face to considerably larger errors at higher longitudes and latitudes.

The Lunar Orbiter has some information of its own to contribute to the triangulation adjustment. To begin with, it has camera station information, since the spacecraft will be continuously tracked by means of its transponder. The reduction of this tracking information should locate the spacecraft in a selenocentric coordinate system to about a hundred meters at any instant. Since there is a record on the film of spacecraft time at each exposure, the camera station can be calculated. The least reading of time on the time block is 1/10 second, so that the spacecraft's position in its orbital track at exposure time cannot be fixed more closely than about 160 meters. This limitation in the basic information becomes less significant in fixing the scale of a strip of pictures than it is for any single picture.

The other information that the Lunar Orbiter can contribute is data concerning camera attitude. For a number of reasons, the absolute attitude information will not be very accurate. One estimate gives a standard error of about one-half degree in roll and in pitch. On the other hand, the relative attitude information can be quite useful. The spacecraft maintains essentially the same inertial attitude throughout a photographic sequence. Its rotation rates are specified to be less than 1/100 of a degree per second. This is set by the size of its control jet orifices in relation to moments of inertia, and it will be tested.

When all this information is stirred into the triangulation adjustment, the usual product ought to come out: the pass point locations necessary for controlling the individual stereomodels. For the Lunar Orbiter data reduction we also wish to extract the adjusted camera station and the adjusted camera orientation. This information is necessary to the exploitation of the high-resolution

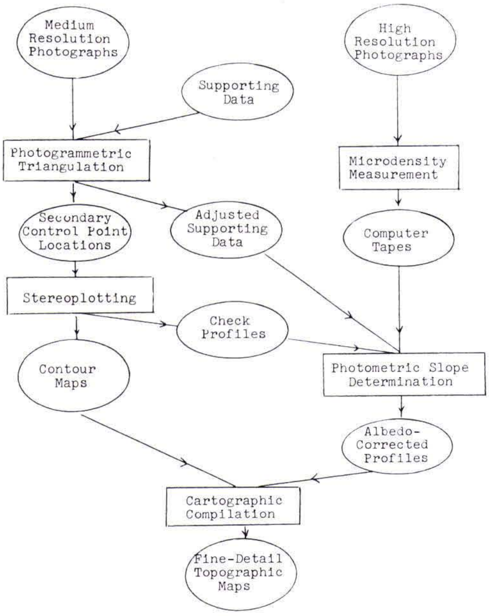

FIG. 8. Generalized data flow.

photos, as indicated in the data flow diagram (Figure 8).

As the diagram indicates, the stereoplotting process starts with stereoscopic photos and secondary control, and ends with contour maps.

This paper will not concern itself with the very interesting question of what constitutes planimetric detail on a large-scale moon map. Let us say that we are only interested in the relief of the surface. What range of contour intervals will give meaningful information from these stereomodels? (Remember that we are not interested in absolute elevations, but we are very much concerned about local slopes.) We know the ground resolution and the base/height ratio, and we can make an estimate of the residual optical and geometric distortions. It is the writer's estimate that contouring at a C-factor somewhere near 1,000 (excluding errors in the secondary control) is about what we can expect to be meaningful. This would be a 40-meter or 50-meter contour interval. If we can calibrate the system completely, and if the map-makers are industrious enough to make all possible corrections, perhaps 25-meter contouring is possible.

Since the Apollo Project must worry about obstructions less than 1 meter high, we have a measure of the extent to which we are forced to depend on the monoscopic high-resolution photos for all of our fine-detail information.

PHOTOMETRIC DATA REDUCTION

To anybody who is unfamiliar with the properties of the moon as a reflector of light, it might seem ridiculous to propose that we should use monoscopic photographs to fill in the relief detail that we couldn't get stereoscopically. There is, however, a valid foundation for this new kind of photogrammetry, based on the lunar photometric properties. The theoretical basis of the method has been explained elsewhere (References 3 and 4). This paper will deal with it only to the extent necessary to show what is being attempted with the high resolution pictures.

Let us begin by contrasting earth photography with lunar photography. In ordinary panchromatic aerial photographs, we see contrasts based on the total reflectivity of the surface materials. The direction of illumination and the direction of emission have negligible effects on the relative tones within a photograph. A dirt road looks lighter than the fields it runs through regardless of the slope of the ground. A notorious exception is the case of open water areas, because the specular reflection of the water surface is highly directional.

On lunar photographs, the effect of the *direction* of reflection on the intensity of light reaching the camera is very strong. The curve expressing the direction effect is called the photometric function. The light reaching the camera from any ground point depends on both the reflectivity of the surface material (which we define as albedo) and the photometric function. A spot on the lunar photo that is brighter than the surrounding area represents either an inherently more reflective material or a local slope toward the sun.

When the Ranger Program produced the first close-up pictures of the moon, several people made topographic measurements based on the directional effect. In order to do it, they made two assumptions: that all of the material had the same albedo, and that all of it exhibited the same photometric function as had been determined by telescope observations. Having made those assumptions, they were able to compute profiles of the lunar surface in the planes which included the sun, the camera, and the surface.

In the case of the Lunar Orbiter's high-resolution photos, there is no need to make these assumptions. We can get independent measurements of slope from the stereoplotter and then reconcile the two types of slope information by adjusting the albedo or the photometric function.

Let us examine the processes represented on the right side of the flow chart (Figure 8). The high-resolution photographs are distorted records of surface brightness. By means of extensive preflight photometric calibration and the transmitted sensitometric step wedges, we can remove most of the photometric distortions in the computer, recognizing the precision limitations of photographic methods of photometry. We first read off the pattern of film densities in the sun-camera-ground plane, using a scanning microdensitometer* and feed that into the computer. To do the profiling, for any photo, we have to know the camera's selenocentric position and attitude as accurately as possible. This information, of course, is available as a byproduct of the photogrammetric triangulation.

The additional data we need to do the profiling without making assumptions about the local albedos are a number of profiles, run in the stereoplotter, which coincide with profiles run in the microdensitometer.

These check profiles should be run as profiles directly, rather than interpolated from the contour map, because there is very much more relief information in a stereomodel than there is in the contour map extracted from that stereomodel.

An important point about these check profiles is that they cannot be chosen indiscriminately. They must be chosen where there is geologic and photographic evidence that the surface material is the same along the length of a profile. Also, since we are comparing information from two different levels of resolution, we should look for places where very fine-scale relief appears relatively negligible.

When the check profiles are thrown along with the other material into the box marked Photometric Slope Determination, the product coming out is a bundle of albedo-corrected profiles carrying very fine topographic detail. In order to tie the profiles together to form a full surface, we should also run some *transverse* profiles in the stereoplotter. It then becomes possible to add the fine-detail information to the original contour maps and to make plots (not necessarily form-line maps, although that is one obvious type of product) which will tell the Apollo Project all it needs to know about the topography of candidate landing sites.

It must be pointed out that nobody really expects to go through this entire process for all of the areas covered by a single Lunar Orbiter mission. We expect to do much visual screening in the first few weeks after the pictures come back, to select the limited areas which require the full data-reduction treatment and the larger areas which require partial work-up.

ORGANIZATIONAL ARRANGEMENTS

A final word remains to be said about the mechanics of getting all this work done. Throughout this paper the word "we" has been used to point out the things that can or should be done. The "we" does not refer only to the Lunar Orbiter Project or even to NASA by itself. The expectation is that most of the topographic work will be done by existing mapping organizations, principally the U. S. Geological Survey, the Army Map Service, and the Aeronautical Chart and Information Center. The arrangements for this have not yet been completed, but the various organizations within NASA that are concerned with lunar programs are currently supporting study efforts at all three agencies directed toward building the theoretical basis and the actual capability for doing the job.

REFERENCES

1. Kosofsky, L. J. and Broome, G. C., "Lunar Orbiter: A Photographic Satellite," *Journal of the Society of Motion Picture and Television Engineers*, September 1965, p. 773.
2. Tarkington, R. G., "The Kodak Bimat Process," PHOTOGRAMMETRIC ENGINEERING, Vol. 23, No. 1, January 1965, p. 126.
3. Rindfleisch, T., "A Photometric Method for Lunar Topography" in this issue, page 262.
4. Spradley, L. H., "Photometric Reduction of Lunar Orbital Data," presented at Semiannual Convention of American Society of Photogrammetry, September 22–24, 1965.

---

* In practice, microdensitometer reading of the reconstructed pictures may be replaced by the conversion of video tape records of the photo transmissions into digital form. The principle remains the same.