

Appendix A – Camera System Operations Manual

VT Calibration Camera Instructions	2
Launching the Software	3
Overview of major panels	4
Camera Control Box	8
Taking pictures.....	9
Setting Camera Parameters	11
Setting the appropriate index of refraction	13
Filling: comparing actual images with calculated images	14
“Tweak”-ing.....	21
Calibrating a Camera	23
Predicting Source Location.....	25
Locating the Source	27
Camera Location Map.....	30

VT Calibration Camera Instructions

(located at c:\calib\instructions\instructions.doc on bxvtcalib.lngs.infn.it in CR4 in Hall C) Please read these instructions before using.

In general, the software is designed to prevent any physical damage to the camera system. As long as one does not try to bypass the user interface there is little one can do which can not be fixed, even though you might well get confused. In general, it is good to leave the main camera power on (equivalent to having the batteries in) but the cameras 'off'. If there are any questions or doubts, please contact Bruce Vogelaar (vogelaar@vt.edu US cell: 540 239 5963) or Steve Hardy. Frequently, the associated black 'DOS box' will provide valuable insight into what is happening. This is a substantial software package which most users will only partially use. The complete code is located at c:\calib on bxvtcalib.lngs.infn.it The code is designed to 'preserve state' as much as possible, so that if it needs to be restarted, almost everything is remembered.

While the software is essentially done, this instruction document is very incomplete at the moment, so please check regularly for an updated version. First are several brief descriptions of the major windows, followed by sequences designed to accomplish specific tasks.

There are seven cameras mounted in the SSS. By calibrating the images, each pixel from each camera defines a unique and known ray within the SSS. Thus, examining which pixel a certain feature shows up in in multiple images, one can reconstruct where that feature was within the detector. Due to slight shifts of the camera lenses between images, four reference light 'points' in the camera housing are included in each image to be analyzed. Using these, one can 'tweak' the image to match the calibration image. The calibration images are key to mapping pixels to rays using the known PMT positions. This was done when the detector was only filled with air. However, after the vessels were installed and inflated, the cameras were moved. Since one can not see through the multiple vessel films, the current 'air' calibration is not as precise as it should be (only ~5cm). When the SSS is filled with water or PC, new reference images will be taken (the films should become much clearer since their index of refraction will closely match the liquid) and calibrations for 'water' and 'PC' done again to about 2cm. The spherical dome of the camera housing will act as a lens (even though the 'pin-hole' of the camera is at the center of the dome), resulting in the need of three calibrations.

Launching the Software

Double-click on VTCalib icon which brings up main control panel. This is used to launch various subwindows. (Normally as 'CameraUser' pwd: calib)



Overview of major panels

“Camera Parameters” is used to set the camera parameters. The method is to copy a startup script to the cameras which then runs when the cameras are turned on again. The parameters are strictly limited, since any request from such a script requiring user input at the camera itself will disable that camera forever.

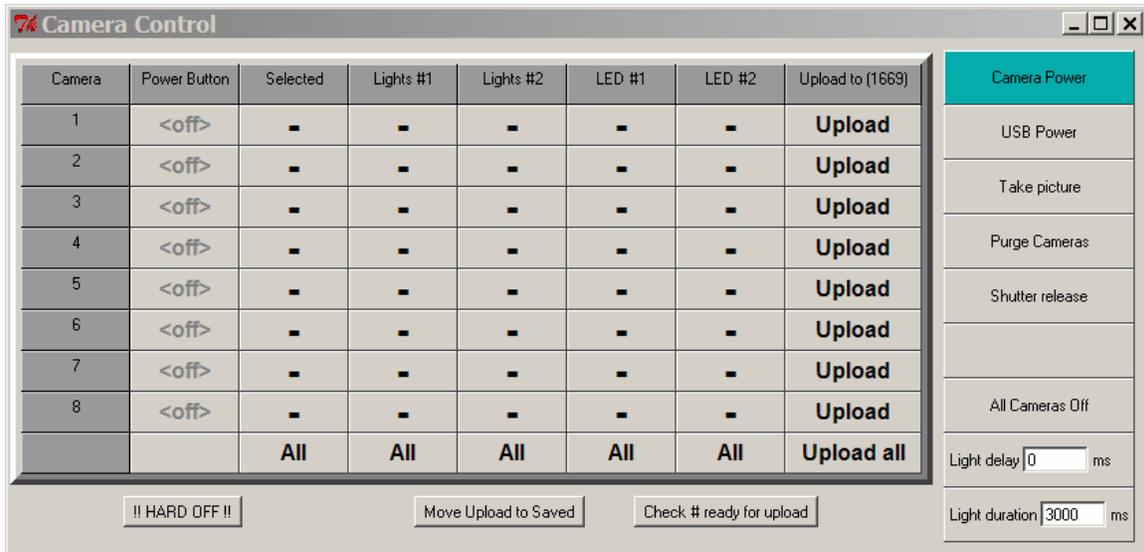
VT Calibration - Camera parameters

Set the camera parameters

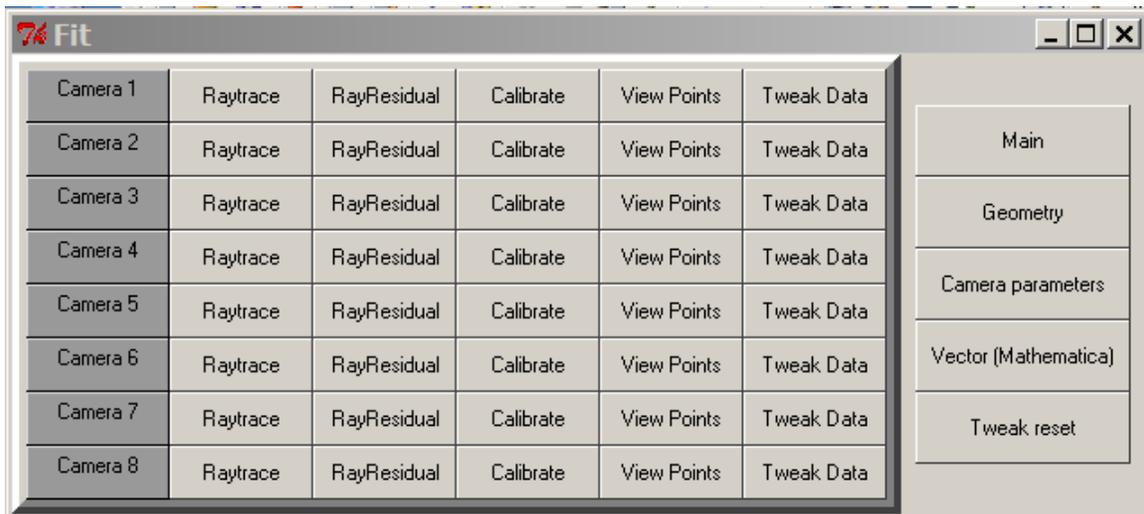
Still compression	<input type="radio"/> High <input type="radio"/> Normal <input type="radio"/> Low <input checked="" type="radio"/> None
Exposure mode	<input type="radio"/> Programmed <input checked="" type="radio"/> Longtime
Shutter speed (3076 - 16000000 microseconds)	<input type="text" value="2000000"/>
Aperture size (300 - 1600, 0.01 step size, 560=f/5.6 default)	<input type="text" value="560"/>
Zoom position (100 - 300, 150 normal)	<input type="text" value="150"/>
Focal distance (50 - 65535 cm) typ: 50 - PC; 1500 air	<input type="text" value="1500"/>
White balance mode	<input checked="" type="radio"/> Auto <input type="radio"/> Daylight <input type="radio"/> Fluorescent <input type="radio"/> Tungsten <input type="radio"/> Off
Focus mode	<input checked="" type="radio"/> Manual
Focus method	<input checked="" type="radio"/> Multispot
Still capture size	<input checked="" type="radio"/> High
Set auto rotate image	<input checked="" type="radio"/> Off
Strobe mode	<input checked="" type="radio"/> Off
Shutter sound	<input checked="" type="radio"/> Off
Zoom mode	<input checked="" type="radio"/> Manual
Image sharpness (-25 - 100, 50 default)	<input type="text" value="50"/>
Image sharpness radius (6553-655360, 6553 default)	<input type="text" value="6553"/>
Image sharpness threshold (0-25, 4 default)	<input type="text" value="4"/>
White balance - blue (16384-262144, 65536 default)	<input type="text" value="65536"/>
White balance - green (16384-262144, 65536 default)	<input type="text" value="65536"/>
White balance - red (16384-262144, 65536 default)	<input type="text" value="65536"/>
Exposure compensation (-200 - 200, 0 default)	<input type="text" value="0"/>

Copy Startup.csm to camera (takes effect on next camera power-up)

“Camera Control” is used to turn on/off the cameras and select which lights are on when the cameras take a picture. Everything is well interlocked, so one can not have lights on when they should not be on. This screen is also used to upload images from the cameras themselves (do NOT try and use Windows Explorer for this).



“Fit” is used to calibrate the camera images. Besides simply taking nice pictures, the cameras can be used as an analytical tool to locate sources, evaluate vessel shapes, and assist during the filling process.



“Locate LED” is the main analytical page. Using uploaded pictures, it allows reconstruction of source positions, inventory and visualization of insertion rods, ability to predict what the vessels and PMTs should look like during partial filling stages, and will ‘tweak’ images to correct for small shifts which occur from picture to picture.

VT Calibration - Locate LED

Various commands

Get Saved | Save | Move Upload to Saved | Project lines from sel. cams | Find (w Tweak) @ calc. xyz | Reinitialize

View Camera:

- Camera 1
- Camera 2
- Camera 3
- Camera 4
- Camera 5
- Camera 6
- Camera 7
- Camera 8

Selected Image Actions

Save as Background

Reset Image

Find LED

Load calculated x,y into cursor

Load/Save to Tweak Master

View Peak Info
(only available at Zoom 32)

Tweak Master Reset

Tweak Reset

Tweak All

Tweak View

Overlay

Options

Subtract Background

Select Image

Normalize Image

Find Pseudo Centroid

Selected Cameras:

- Camera 1
- Camera 2
- Camera 3
- Camera 4
- Camera 5
- Camera 6
- Camera 7
- Camera 8

Color	Mouse Position	Cursor Position	Calculated Position
R 0	x 0.0	x 0.0	x 1.0
G 0	y 0.0	y 0.0	y 1.0
B 0			

Color Option

Color Red Green Blue

Zoom

1/8 1/4 1/2 1 2 4 8 16 32

Enter LED Coordinates

Polar Rectangular

Enter Parameters for centroid

Centroid Range Sel. Thresh.

Insertion Rod & Tube Inventory

Number of rods above hinge	<input type="text" value="9"/>	Straight Tether Length (cm)	1180.0	LED r (pred)	227.9
Number below hinge (n-no hinge)	<input type="text" value="2"/>	Tether Short by (cm)	<input type="text" value="63.4"/>	LED t (pred)	108.3
Extension in Glove Box (cm)	<input type="text" value="40.0"/>	Arm Angle (degrees)	<input type="text" value="53.3"/>	<input type="button" value="Toggle View Geometry"/>	

Reconstruct Locations

LED			Source (depends upon rod geometry)		
x	0.0	r	0.0	x	0.0
y	0.0	theta	0.0	y	0.0
z	0.0	phi	0.0	z	0.0
arm angle 0.0					

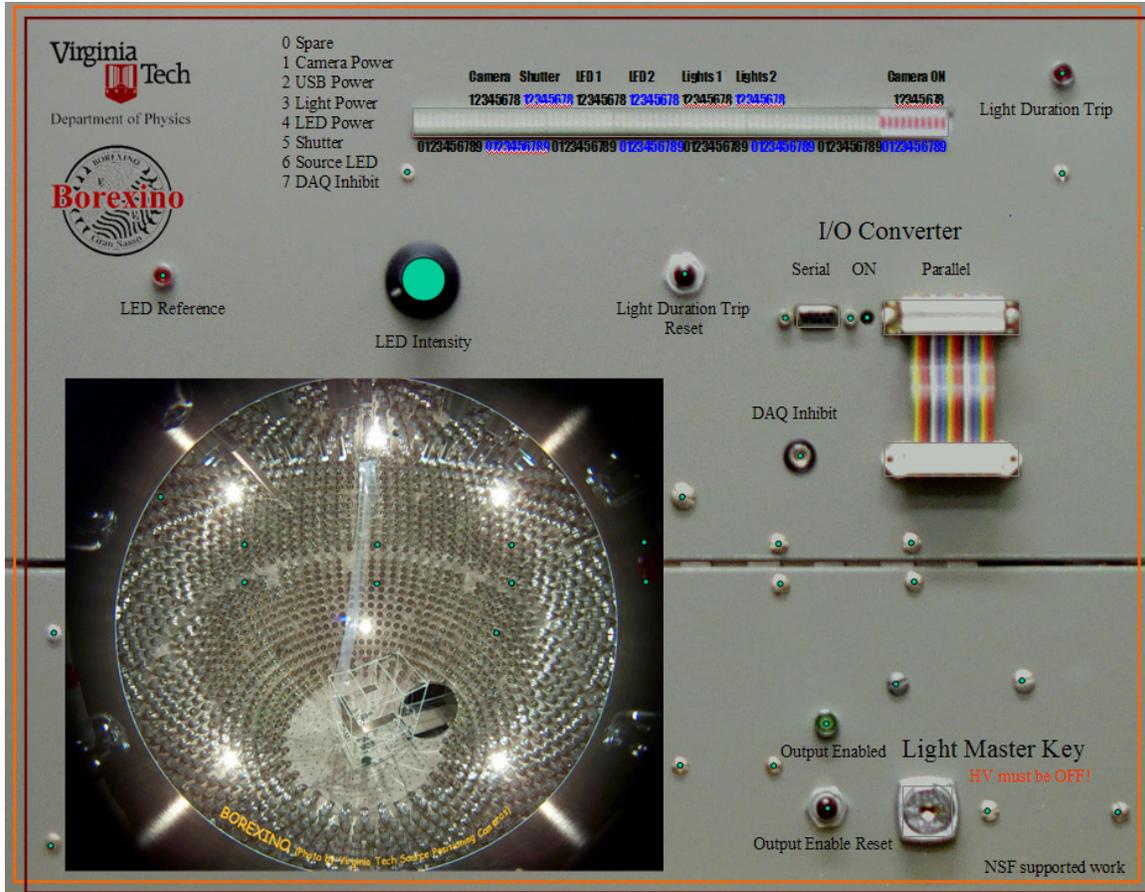
Current Index of Refraction: 1.0

“Gas Control” launches a stand-alone program used to control gas flow for the internal source calibration system. Its functions are described in a separate document, but a screen shot is included here for completeness.

V1	closed	G1	0 to 0	FM1	0 to 0	P1	off	LM
V2	closed	G2	0 to 0	FM2	0 to 0	P2	off	V1377
V3	closed	G3	unset	FM3	0 to 0	NV1	unset	V1320
V4	closed	G4	0 to 0	FM4	0 to 0	NV2	unset	OBPump
V5	closed	G5	0 to 0	CV1	unset	NV3	unset	IV
V6	closed	G6	0 to 0	CV2	unset	NV4	unset	IB
V7	closed	OM	0 to 0	CV3	unset	R1	unset	Covers
V8	closed	SG	0 to 0	CV4	unset	R2	unset	TM
						R3	unset	

Camera Control Box

The camera control box runs numerous relays and allows for visual confirmation of how the cameras are operating. It looks like:



Due to safety concerns, it is mounted behind a window, outside of CR4. The ‘Output Enable Reset’, ‘Light Duration Trip Reset’, and ‘Light Master Key’ have now been remoted to be inside CR4. (This whole box is scheduled to be replaced by commercial relay boards sometime in the future.)

The bank of LEDs indicate which cameras and lights are selected. The far right indicates if current is being drawn by the cameras (ie: they are ‘on’). If the main power to this box is turned off, then when it is powered up again, and the computer is running, then the ‘Output Enable Reset’ button needs to be pressed. Above this box, one can see the USB extender boxes, and their three lights each indicate power, remote power, and connection from left to right. The PMT HV and ‘Light Master Key’ are on the same ring to prevent the quartz lights from being turned on at the same time as the HV on the PMTs. (The system is on the UPS supply, so power should never be off.)

Taking pictures

Assuming the camera settings are correct (if not, see below for how to change them), one selects:

VTcalib:camera control:camera power (if not already on):USB Power: appropriate power button

The camera(s) should turn on one at a time, and connect as a Windows assigned drive letter which replaces the <off> label. If they do not connect, it could be because the camera power was off, in which case the camera will need to be turned off again and turned on again (sets default date, lost when power failed since the internal batteries have died). If they still do not connect, the computer may need to be rebooted. They have been known to be finicky from time to time, so it might take a couple of tries... (If a little warning window pops up saying a drive could not be found, simply close it.)

Once the cameras are on, select which camera will be taking the picture using the 'Selected' button. Then set which lights you will be using. One should only have to use 'Lights #1', as 'Lights #2' is only a backup. You can turn on/off all the lights by pushing the 'All' button at the bottom of the appropriate column. Note that you can not turn on the lights for any of the 'selected' cameras, since there would be too much local light at that camera. Next, select which LEDs you want. LED#2 is the two bright LEDs in each camera housing, and LED#1 is the four pin-point LEDs in each camera housing used to 'tweak' the image (due to slight shifts between successive images).

Next, adjust the 'light delay' which is how long after the shutter is pressed that the lights will turn on (normally 0ms). Then set the 'light duration' which determines how long the lights will stay on (normally 3000ms). Ensure that the PMTs are locked out if using any of the light banks! (Keyed system should prohibit any accidents.)

Next, click 'Purge Cameras' which will delete any images left on any camera turned on. Finally, select 'Take Picture' which will begin the picture taking sequence: DAQ inhibit, 0.5 s delay, shutter, light delay, lights&LEDs&Source LED on, light duration, lights&LEDs&Source LED off, 2 s delay, shutter off, 10 second delay, DAQ re-enabled.

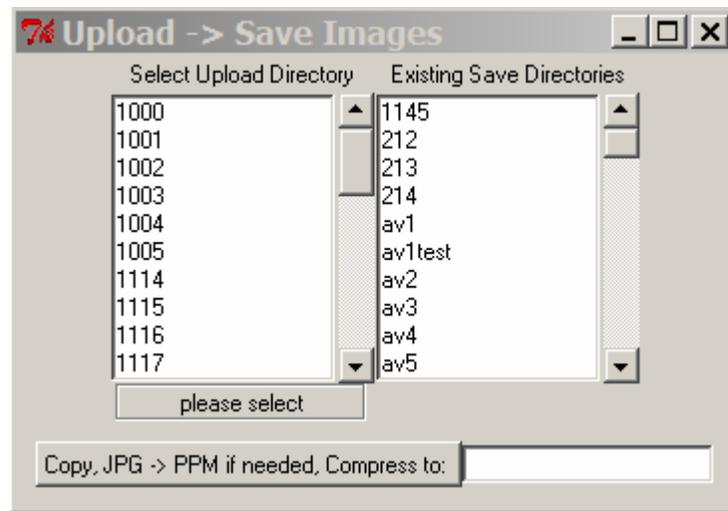
The 'selected' cameras will have taken pictures. You must wait for the cameras to finish processing before trying to upload the images. You can check for this by clicking on 'Check # ready for upload' to see if they are all there (result shows up in DOS box). If you want pictures from more than one camera and are using the main light banks, then you can reset which camera is selected, adjust the light map, and hit 'Take Picture' again.

When all the desired cameras have pictures, you can upload them all at once into the 'upload' directory (whose number is appears at the top of the 'upload' column. You can also upload single cameras at a time if desired, but each will show up in a separate 'upload' directory.

NOTE: for some reason, when taking uncompressed images (ie TIF), the shutter has to be pressed once between each 'real' picture. Thus, between each set of pictures, you must hit 'Shutter Release' which will simply press the shutters on the selected cameras to provide the extra step. Unfortunately cameras can get out of sync, depending upon the history of picture taking. To re-sync them, 'select' the cameras which need an extra shutter push, and hit 'Shutter Release', which will simply press the shutter button on the selected camera.

When finished, click 'All Cameras Off', which will cycle through all the 'on' cameras and turn them off. (You can also turn them off one at a time by clicking on their drive-letter in the 'Power Button' column.) Then turn off the 'USB power'. Leave the 'Camera Power' on (ie: blue), unless the cameras will not be used for a substantial time (over a week or so).

If you are taking regular 'monitoring' pictures, use Windows Explorer to copy the uploaded images to whatever directory you wish. If you intend to actually analyze the uploaded pictures, click 'Move Upload to Saved' which will bring up this screen:



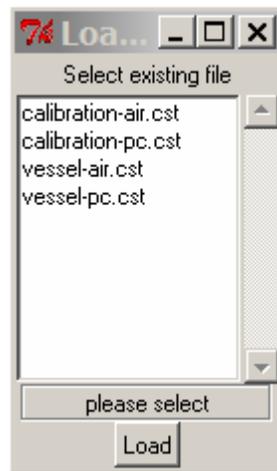
This will convert (if necessary) and compress images in your 'upload' directory to images in the named 'saved' directory. Only images in a 'saved' directory can be loaded into the later analysis steps. (This operation can also be done later from the 'Locate LED' screen.)

Setting Camera Parameters

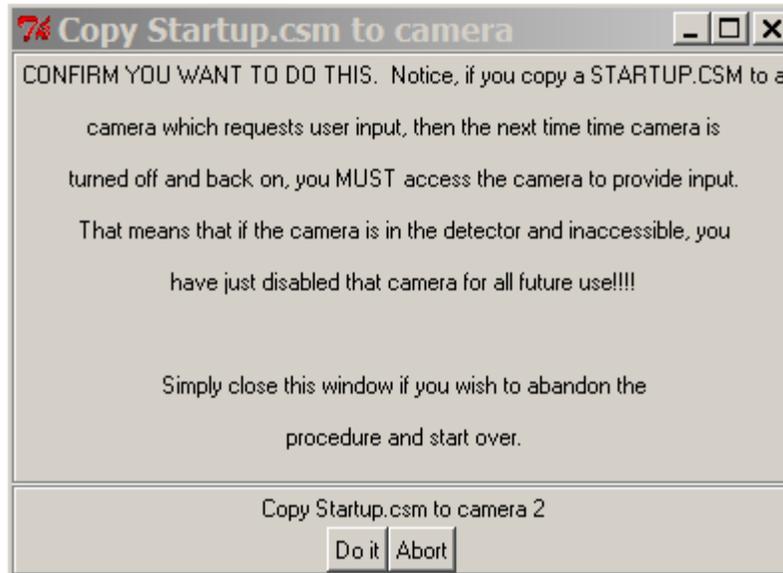
Turn on the cameras (see above) which you wish to set the parameters on. Since Kodak does not support these cameras any more, we can only adjust the parameters by writing a 'startup.csm' file onto the cameras which runs when the cameras are turned on. (This is sort of like the old 'autoexec.bat' files on a PC.)

After the cameras are turned on, select 'Camera Parameters' from the main launch screen. Here one can set a carefully selected number of options, save them to 'startup.csm' and then copy them onto the cameras. These setting will only become active once the cameras are turned off and then on again.

There are really only four sets of options that one needs: PC filled (monitor), PC filled (locate source), Air filled (monitor) and Air filled (locate source). To facilitate setting these, one can load them from a '.cst' file. Select 'Load from file (.cst)' which brings up



Select the appropriate one, and click 'Load'. Next, save this by clicking 'Save to Startup.csm'. Next, select the appropriate option on the bottom row to copy this file to the appropriate camera. A warning screen will appear, asking you to confirm this action.



If you are sure, then select 'Do it'. A confirmation window will appear, letting you know if it was successful. Close the window and the 'Camera Parameters' window.

All that remains is to turn the cameras off again. The next time they are turned on, the new settings will be active. They will remain the default settings until you next copy a 'startup.csm' file to the cameras.

Setting the appropriate index of refraction

When reconstructing the source location, it is assumed that the vessel is filled with only one kind of liquid. The index of refraction for this liquid is set in `c:\calib\fortran\geom.dat` and this file must be edited to reflect the correct index.

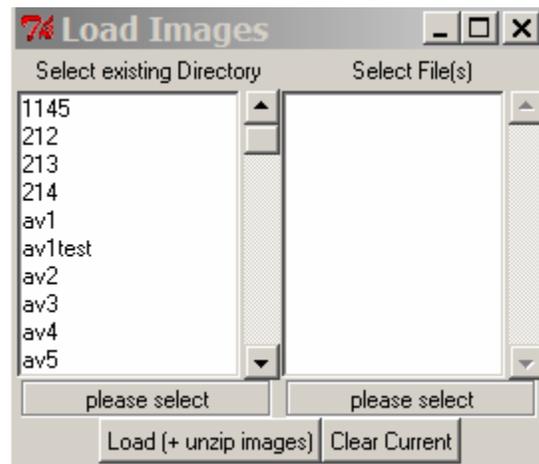
When predicting what each camera should see when the vessel is partly filled, both the index of refraction for the upper region and the lower region must be set. These are set in the file `c:\calib\raytrace\raygeo.dat` and this file must be edited to reflect the correct indexes.

The data files are self-explanatory, but their overall format must not be changed.

Filling: comparing actual images with calculated images

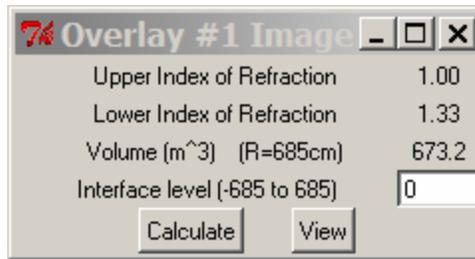
Without having fixed, known locations on the vessels to image, one can not reconstruct the vessel positions (shapes). The best one can do is to predict what each camera would see given various liquid levels within the SSS for spherical vessels. (A sort of pseudo 'rendering') This is accomplished by taking pictures and moving the results into a 'saved' directory as described above. Once there, ideally using the appropriate calibration (slightly different for air, water, pc) for the liquid covering the camera dome, one can calculate what the image should look like and visual comparison will indicate if the vessels are distorted, etc.

To do this, click VTcalib:Locate LED. If you have not already done so, move the images to be analyzed from 'upload' to 'saved' by clicking on 'Move Upload to Saved'. Next, select the appropriate 'saved' images to analyze by clicking on 'Get Saved' which brings up



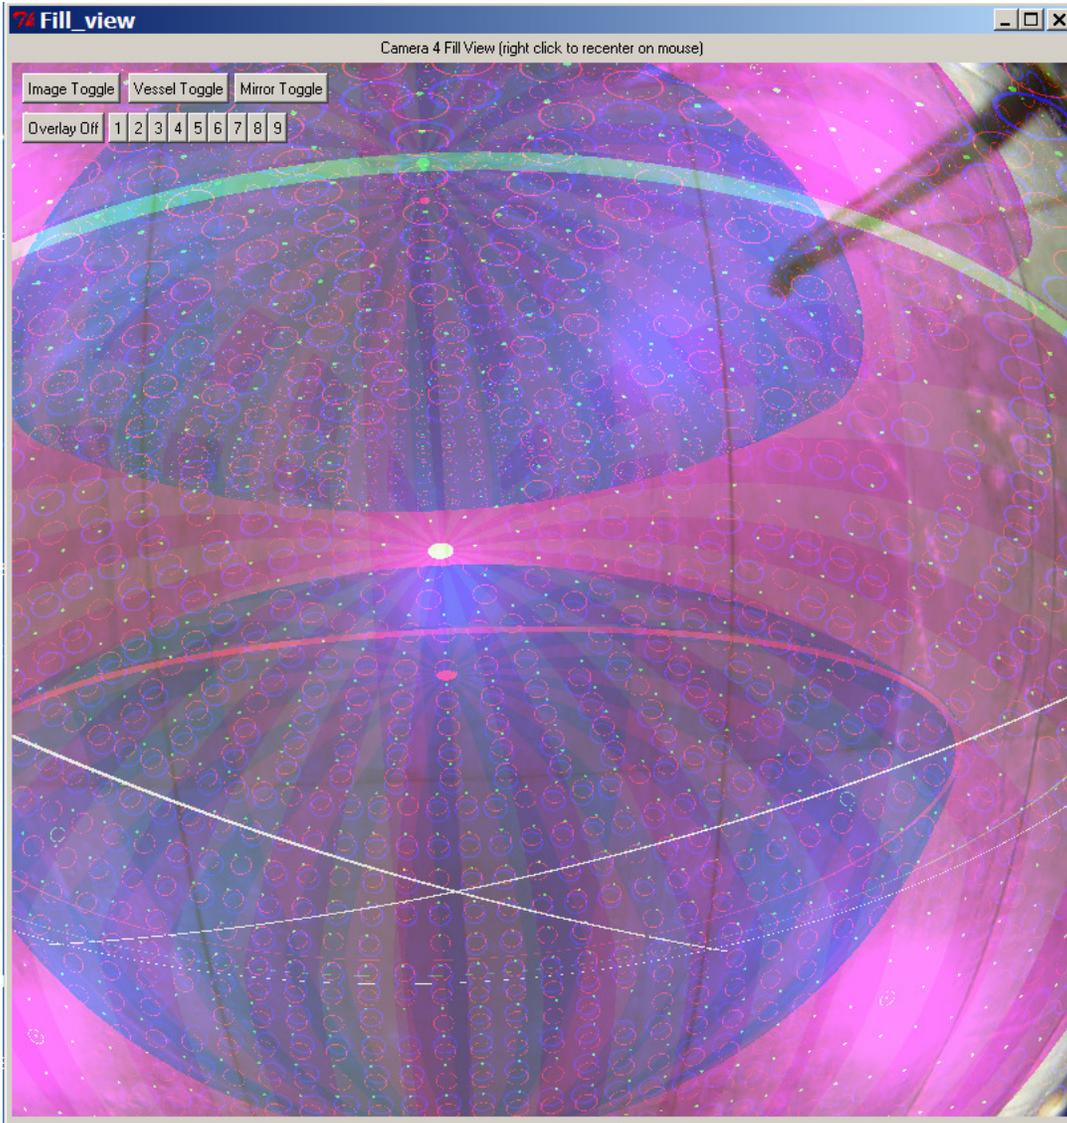
Unless you are sure, it is best to click 'Clear Current' first. Then select from the left screen the 'saved' directory to look in. Select the option/file from the right panel and click on 'Load (+unzip images)'. These pictures are copied into another directory, and kept there for future reference until you either clear or overwrite them with a subsequent 'load'. Either click 'Clear Tweak' or 'Tweak All' as desired.

Clicking 'Reinitialize' at this point isn't a bad idea, especially if you're new at this. Then select which camera you wish to analyze using the left radio-buttons. Then click on 'Overlay' (right side of Locate LED window) which will bring up

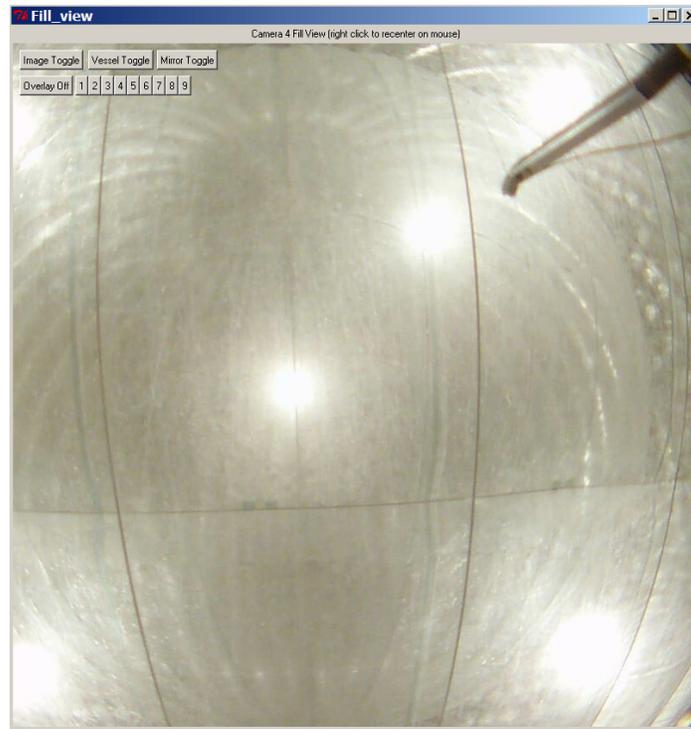


Check that the window title is indeed the camera you want, and that the indexes of refraction are the correct ones (see above to change these). Next, indicate the 'interface' level, where 0 is the center of the detector, and the units are in centimeters. (It calculates the total volume, assuming a 685cm radius sphere for reference.) Finally, click on 'Calculate'. This will take ~2 minutes, and you can watch the progress in the DOS box. It calculates the image twice, once with refraction at the interface (and total internal reflection if appropriate), and a second time where all the refracted rays are replaced with reflected ones. You only need to do this step if changing camera or liquid interface level.

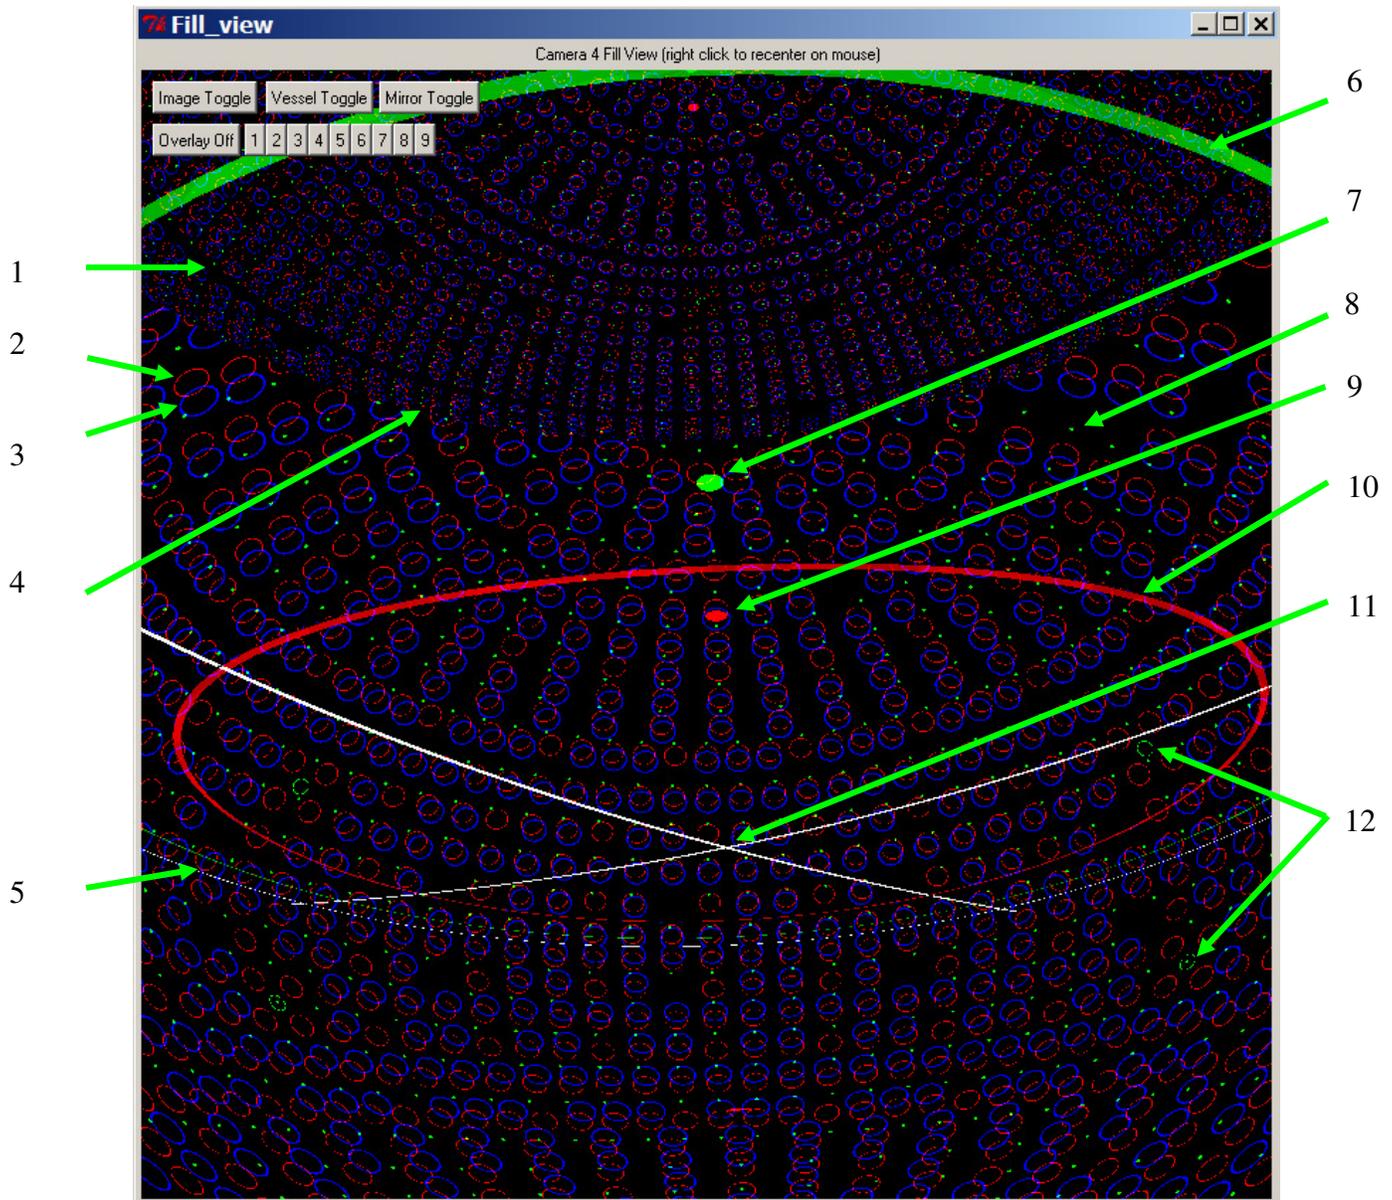
Otherwise, or after running 'calculate', click on 'View' which will bring up (for camera 4 and air/water interface at -250 cm.



This is hard to interpret at first glance, so instead let's turn off all overlay items by setting 'Overlay Off'. That gives you:



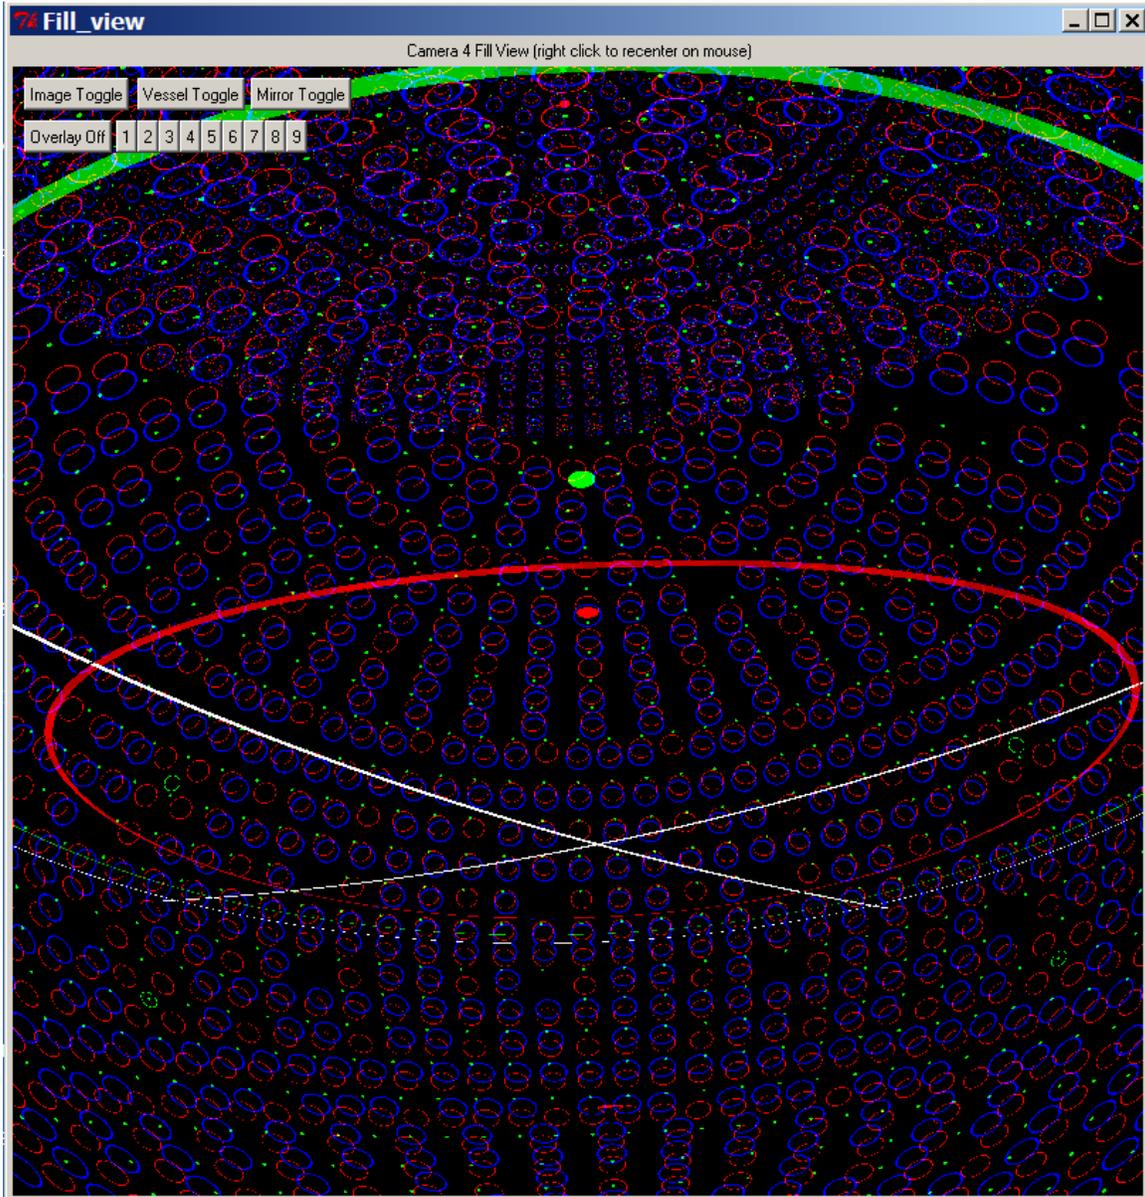
which is simply the image from camera 4. You can re-center the window by placing the mouse in the image and right-clicking. Next, let's turn off the 'Image', 'Vessels', and 'Mirror', and set 'Overlay' to 5. This, with a slight recentering, gives you:



You will notice several items:

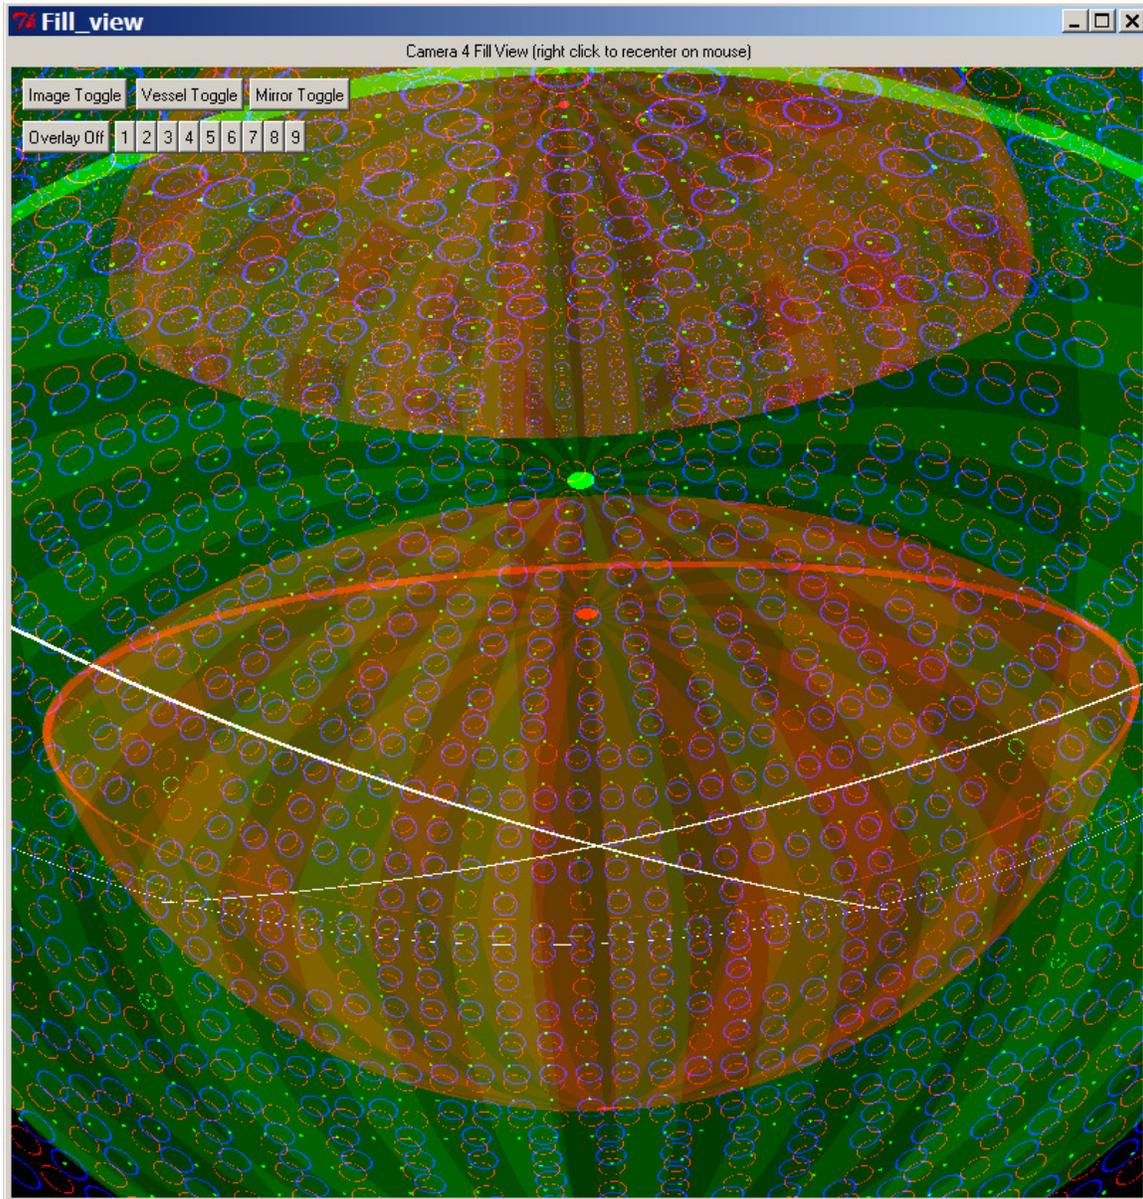
- 1) This band is the SSS equatorial ring.
- 2) The small red circles are the leading edges of the mu-metal shields
- 3) The small blue circles are the leading edges of the light concentrators
- 4) Notice the transition between total-internal reflection from the interface, and refraction through interface.
- 5) The large white circle shows the intersection of the interface with the sphere defined by the leading edges of the light concentrators. Notice the total reflection just above this.
- 6) The large green circle shows the intersection of the interface with the OV
- 7) The small solid green circles are the location of the OV poles.
- 8) The small green dots are the locations of the PMT feed-throughs in the SSS
- 9) The small solid red circles are the locations of the IV poles.
- 10) The large red circle shows the intersection of the interface with the IV.
- 11) The white cross is simply painted onto the interface plane and the lines extend out to the SSS radius.
- 12) Small green circles: location of other cameras. In this case, the lower one is a direct view, while the upper one is the *same* camera, but repeated due to total internal reflection from the interface.

Now turn on 'Mirror'ing and the image changes to:



Notice that superimposed onto the refracted light coming through the interface, one also has light as though the interface was a perfect mirror.

Finally, add the 'Vessels' back in. This paints transparent 'orange peels' indicating what a spherical vessel would look like. In practice, the vessels are either too opaque (in air) or likely to be too transparent (in PC) to be easily viewed. Thus one may need to toggle the 'Vessel' option on and off several times (when the image is also present) to be able to notice any differences. Below, the 'Vessel' option is toggled back on:



Notice how peculiar the vessels look when the detector is partially filled. Without matching images to calculations, it would be virtually impossible to interpret what was going on.

What remains, is to overlay such calculations onto actual images in such a way as to allow reasonable interpretation of the results. To accomplish this, the relative ratio of calculation to image is set by the 'Overlay 0-9' buttons. One may need to toggle on/off things like 'Image', 'Vessel', and 'Mirror' to help figure things out. With a little practice, one can actually get quite good results. Deviations between calculations and images are cause for concern, but it is hard to quantify vessel deformation without an analytical expression for their geometry (they are simply assumed spherical for now).

“Tweak”-ing

Since commercial cameras were used, the lens retracts and returns between each image taken. As a result there can be a slight (several pixel) shift between images. To correct for this, four ‘pin-hole’ light sources in each camera housing can be turned on while taking pictures. Doing this for the ‘calibration’ image provides fixed reference points for future images. These are then shifted and scaled so the four reference points line up with those in the calibration image. The process is fairly automated.

First, let’s describe how the reference points are defined. Images with ‘lights on’ are taken (with LED #1 bank also on). These images are moved to a ‘saved’ directory, and then loaded into the ‘Locate LED’ window as described above (these are the same images used to calibrate the optics). Select the camera to work with, and then click on ‘Reinitialize’ in the upper right, then ‘Load/Save to Tweak Master’ which brings up (assuming work with Camera 1):

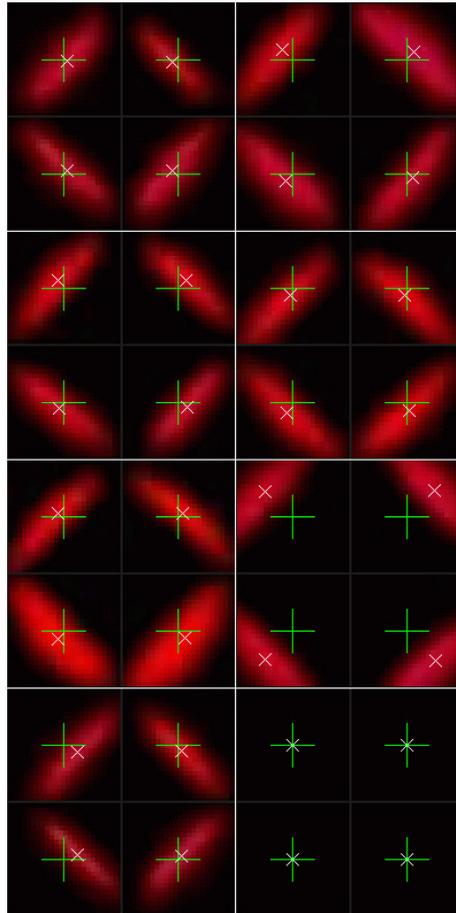


Next, click on Load from ‘Upper Left’. The image in the ‘Locate LED’ window will be redrawn close to where the upper left reference light is. It probably will not be exactly where desired, but by using the right mouse to drag the image around the reference light can be approximately centered. (One might also need to change scale to help get oriented.) The oblong shape of the ‘pin-hole’ is due to the fish-eye effect of the lens. Next click on ‘Load Calculated x,y into Cursor’ and repeat until the white ‘x’ lines up with the cross-hairs. At that point, click Save to ‘Upper Left’. Repeat for the other three lights. Next repeat for all seven cameras.

At this point ‘Tweak Master’ will have been defined. (It is located at: c:\calib\calibrate\tweak_master.dat. You might want to save a copy in that directory for later reference.) Clicking ‘Tweak Master Reset’ will copy this file to c:\calib\fortran\tweak_reset.dat. Clicking ‘Tweak Reset’ will copy this latter file to c:\calib\fortran\tweak.dat. (A bit redundant for historical reasons...) It is this final file which will be written to during the ‘tweak’ search and used for subsequent analysis.

Subsequent images which are ‘saved’ and then analyzed (assuming that LED #1 bank was turned on) can now be corrected for any offset/scaling required. Clicking ‘Tweak All’ will try to locate all the reference points in the new images, and write the

results to the second data set in 'tweak.dat' as well as scale and offset factors. To examine how well this has occurred, click on 'Tweak View' which will bring up:



This is a composite picture, examining pixels around each of the four 'pin-hole lights' for each of the cameras (1,2,3,4,5,6,7,8 left to right then down) The green 'plus' signs indicate the pixels where the pin-hole lights showed up in the 'calibration' image, while the white 'x's indicate where they are found in the current image.

notes:

- if one of the images looks really wrong, simply do not use that one for locating the source. Otherwise, a more hands-on approach can be used, but is not documented here.
- take care to make sure to 'Tweak Reset' before calibrating an image (to avoid double corrections).
- you can see the tweak scale and offset values actually being used in the 'DOS box' whenever an image is 'calculated'.

Calibrating a Camera

Having done this once, there is no longer a need to start from ‘scratch’. The idea is to take an image of the detector with the quartz lights on (as well as LED#1 for future ‘tweaking’). Since we know the physical location of all the light concentrators, we can use the central pixels for these, coupled with the known location of the cameras, to determine well defined rays in space. Since there are over 2000 PMTs, a simple functional form to correct the optics is over-determined, and a simple fitting routine is used to determine the best parameters for the function. (A pin-hole camera would be a simple direct calculation, but then wouldn’t have any light collection efficiency either.)

The process is to *calculate* an image, using the above function and *known* PMT positions, and compare it to the actual image. If the two match, the calibration is good. An uncertainty of one pixel in a camera corresponds to about a 1 cm uncertainty at the center of the detector. The process is iterative, starting with a few identifiable mating points in both the calculated and true images. This is enough to set the x,y offset of the CCD image, as well as an overall scale and rotation (other parameters found from earlier work should be left alone at this point). The calculated image is then regenerated, and a few more identifiable mating points will become apparent. With these added, the function is again fit (minimizing the pixel shifts between the calculated and true-image calibration points), perhaps with a few more parameters active. Repeating this process, one will wind up with perhaps one to two hundred calibration points. A good fit will result in less than a few pixels difference between any calculated and true image calibration point.

To define a calibration point, the calculated image is overlaid onto the true image, and can be drug around until a calculated light concentrator overlies a true-image light-concentrator. When satisfied, the true image pixel and x,y,z position of this light-concentrator (from a table generated by raytrace.f for experts) are saved together as a calibration point. One then goes to another region of the image, and repeats the process. If such a ‘match’ is later seen to be wrong (can’t be fit), it can be modified or deleted. The chi-squared used in the fit is the sum of the pixel distances squared for all calibration points between the true-image pixel and the calculated pixel given by the fit function operating on the associated x,y,z position (normalized to the number of points). This approach avoids the problem that the amount of ‘drag’ will depend upon the function parameters *which change* between iterations of adding new calibration points.

While the camera’s ‘pin-hole’ is located at the center of the camera housing glass dome, there is likely to be some lensing effect when the detector is filled with either water or PC. This means that there will be in fact three calibrations: air, water, and pc. Each of these requires a clear picture where all of the light concentrators are visible. Once calibrated, the quartz lights will never need to be turned on for source work, and thus the PMT HV never needs to be turned off (although the DAQ needs to

be inhibited while the LED lights are on). (This has been verified during tests – the dark rate recovers in a few seconds.)

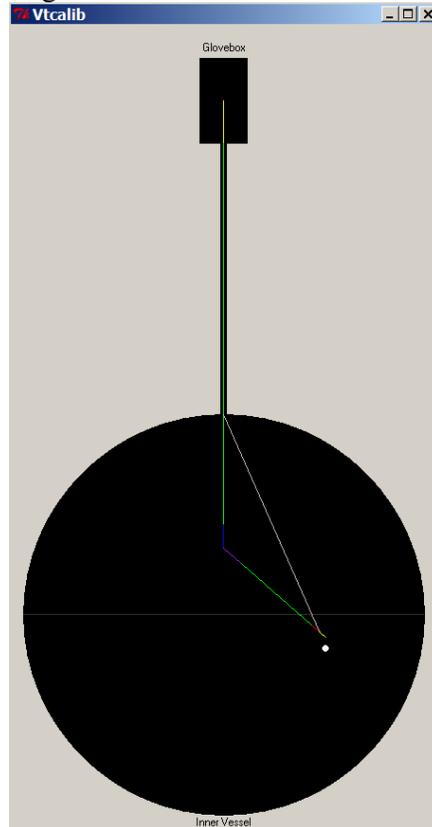
This iterative process is best learned by individual training.

Predicting Source Location

The location of the source can be predicted knowing the number of insertion rods, the tether length, and the azimuthal orientation of the insertion rod. Certain geometry parameters are hard-coded, such as the length of the rods, the distance from the glove-box to the center of the detector, the radius of the IV, etc. These only need to be set once in the file `c:\calib\fortran\rod.dat`. Open the 'Locate LED' control window, and locate the following section:

Insertion Rod & Tube Inventory					
Number of rods above hinge	<input type="text" value="9"/>	Straight Tether Length (cm)	1180.0	LED r (pred)	223.1
Number below hinge (n-no hinge)	<input type="text" value="2"/>	Tether Short by (cm)	<input type="text" value="56.0"/>	LED t (pred)	112.0
Extension in Glove Box (cm)	<input type="text" value="40.0"/>	Arm Angle (degrees)	<input type="text" value="50.0"/>	<input type="button" value="Toggle View Geometry"/>	

Enter the parameters in the left column, and *either* 'Tether Short by' or 'Arm Angle'. Ensure 'entry' by either typing a carriage return, or by clicking in another entry box. The predicted r (from detector center) and theta (from the North pole) location of the LED is displayed as 'LED r' and 'LED t'. Note that the source is actually slightly beyond this, along the direction from the hinge to the LED. To view the geometry, you may select 'Toggle View Geometry' which brings up a graphic display of the situation, updated every time a parameter is changed.



It is well worth playing around with this a little bit to get a feel for how changing parameters will impact the end of the rod location. No harm can be done by simply trying various combinations.

For example, if the hinge is at the center of the detector and has a bend of 90 degrees, raising the rod up (by say increasing the 'extension into glove box') will move the hinge up, but also increase the hinge angle past 90 degrees, since the tether will be held against the vertical insertion rod for a greater length. (note: the software will allow the source to be predicted anywhere, even through the vessel, or with the hinge past 90 degrees, even though these are not physically possible.)

Using the predicted location of the LED will facilitate locating the LED with the cameras. Note that the azimuthal angle is simply read off-of the source rod orientation.

Locating the Source

If not already done, take pictures as instructed above and move into 'saved' directory. Load these 'saved' images into the 'Locate LED' control panel using 'Get Saved'. Click on 'Reinitialize' for good measure. Enter the predicted LED position based on the rod configuration as found above into this panel:



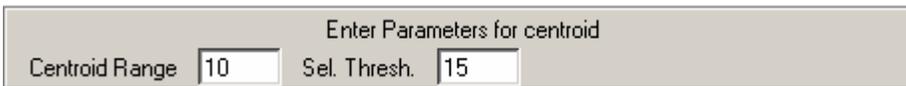
Enter LED Coordinates

270.0 135.0 0.0 Polar Rectangular Calculate

and click on 'Calculate'. In the top of the 'Locate LED' control panel, click on 'Find (w tweak) @ calc x,y,z'. If all goes well, the location of the LED and source are returned at the bottom left of the 'Locate LED' control panel.

[If there are no real images to look at yet, you can mimic the above process by simply entering the LED coordinates (as found by rod geometry or otherwise) and clicking 'Calculate'. Then, rather than doing an automated 'Find', click through each of the cameras (on the left side of 'Locate LED') and for each camera (other than Camera 8 – which is a dummy camera station) click on 'Load calculated x,y into cursor' (right side of 'Locate LED'). Then click on 'Reconstruct' at the lower right of the 'Locate LED' window.]

Sometimes the automated 'Find' will have problems. For example, the source LED might be blocked for a given camera by the insertion rod itself. The software will try to identify such situations by checking that there is indeed an LED seen in the picture. The sensitivity for finding an LED is set in:



Enter Parameters for centroid

Centroid Range 10 Sel. Thresh. 15

If an LED is not found for a given camera, that camera will not be used in reconstruction. The cameras which are being used will have a check mark in 'Selected Cameras' on the left side of the 'Locate LED' control panel.

For a careful determination of position, one should utilize the 'Tweak View' feature to make sure that the automatic 'tweaking' looks good, and also click through each of the cameras to make sure that the LED image is properly centered in the display window. Any camera which is suspect can be removed from the reconstruction process by unchecking it under 'Selected Cameras'. When satisfied, click 'Reconstruct' to get the reported LED and source position.

One can also manually locate the LED in an image using the image manipulation tools (whose parameters are set automatically when using the 'Find' feature):

Color	Mouse Position	Cursor Position	Calculated Position
R 0	x 45.5	x 42.1	x 42.1
G 0	y 117.9	y 121.7	y 121.7
B 0			

Once the LED is centered (by dragging the image around while holding the left mouse button down), click several times on 'Load calculated x,y into cursor'. If still not satisfied, simply manually place the centroid at the cursor crosshairs. Then select this camera, repeat as needed for other cameras, and then 'reconstruct'.

The use of background subtracted images will hopefully not be needed, but is available. Additional diagnostics on the reconstructed position will show up in the 'dos-box':

```

Range :          8.938
x0,y0,z0:    189.886    -1.130   -189.117
x,y,z:       189.886    -1.130   -189.117
Camera 2 dist =    4.73233032
Camera 3 dist =    3.80692625
Camera 4 dist =    2.38101673
Camera 5 dist =    2.11745572
Camera 6 dist =    1.38304567
Camera 7 dist =    0.411929756

```

This example says that cameras 2-7 were used in reconstruction, and the ray from camera 2 passed within 4.7 pixels of the reconstructed position, etc. If any one is abnormally

large, it would be good to visually check the image in the 'Locate LED' panel, and perhaps either fix it, or else remove it from the reconstruction process.

Camera Location Map

