Important URLs:

- http://cass185.ucsd.edu/help/smei/pro_frames.html

  This is the main resource for all the UCSD/SMEI code that already exists.

- http://ips.ucsd.edu/

  Space Weather Forecast web page. Good to check up on periodically to make sure it is online and functioning correctly.

- http://smei.ucsd.edu/

  This is the SMEI main web page, like the Forecast page, should be checked periodically to make sure it hasn't dropped offline.

- http://smei.ucsd.edu/manager

  login: admin
  password: Mnt41eZ

  Location of the content management administration page for the SMEI page, this is where modifications can be made to the pages in case there are errors, or updates are needed.

- http://physics.nyu.edu/grierlab/idl_html_help/idl_con.html
- http://physics.nyu.edu/grierlab/idl_html_help/idl_alph.html

  IDL documentation, including information as well as built in functions.

Important Directories:

- $SMEI also /zshare/smei

  This is the location of most of the actual code, mounted from the smei server over NFS, exact path to files is included in the UCSD code documentation URL from above.

- $SMEIDB also /zshare/sswdb/smeidb

  Lots of stuff used by various apps, but rarely need to be accessed directly.

- $SMEIDBN also /smeidb
- $SMEIDCN also /smeidc

  These contain the SMEI satellite frames. SMEIDB is the primary, SMEIDC is a backup source that is spread over NFS.

  Other useful directories can be found by typing export at a console.

Important Programs:

- The Indexer (http://cass185.ucsd.edu/help/smei/for_s.html#smei_skyd and $SMEI/ucsd/camera/for/sky/smei_skyd.f)

  The indexer creates the skymaps from SMEI data frames. This program is usually only run in its entirety when drastic changes have been made to the way the data is processed, otherwise it only processes new data. Data is read from SMEIDC (the subnet backup of SMEIDB) and written to $SMEISKY0/sky/cX where X is the camera number.

  It is controlled by a program called skyd which spreads the load overas many machines on the subnet as possible using a background daemon and cron jobs. This daemon is usually run with zulu.smei.ucsd.edu as the master and as the user skyd. The directory with configuration files is $SKYD or /ztemp/skyd/skyd/cf and the configuration file is named by the host of the master, right now, skyd_zulu.cf . The main things that can be changed in this file are the grunts, which specify which systems on the subnet to send data, by default 2 (max_load) processes are executed, this can be lowered to 1 by specifying the grunt as hostname:1 .

  The rest is divided into sections like this:
  group_0:
  _min_orbit=400
  _avoidmoon=1
  _max_orbit=22126
  _level=3
  _destination=$SMEISKY0/sky/c1
  _source=SMEIDC?
  _catalogue=$SKYD/list/skyd_c1m2.txt
  _camera=1
  _done=0
  _mode=2
  _checkversion=1

  The section heading, group_X: where X is an integer number defines a section of data to be processed via round robin. Most of the other settings do not need to be modified.

  If changes are made to the configuration file, the daemon must be reloaded to recognize them via : skyd reload

  Similarly, skyd start and skyd stop start and stop the daemon. skyd coast will tell the daemon to stop when all currently running processes have completed, this is cleaner than just stopping skyd.

- Star Subtraction (Currently mishmashed into various systems as a standalone python script but should be uploaded to the python code tree $PYTHONPATH soon. The copy there is outdated. /home/soft/python/remove_stars.py)

  Star subtraction processes the indexed frames and removes the stars to clean up the final images. There are a lot of options here and they can be seen by checking the help text of remove_stars.py. Generally two processes can be run per machine (as with almost all of the processing routines). This python script, given certain variables will create a larger IDL script which is then passed to

IDL and run. Depending on the settings, it can take a very long time (about a week) to run (also... like most of the processing routines).

The Most Important Thing of All

- Backup DVDs

There are two kinds of backup DVDs. Regular DVDs and gzipped DVDs.

The process goes like this:
1 -> download frames from Air Force (this is done automatically via cron job on smei) (to $L1A/0_core or /hafb/l1a/1_get)
2 -> after sufficient data has been downloaded, it is moved to $L1A/2_dvd
3 -> this is the first step that we need to be involved in, Burning to DVD.

as the user soft, with a blank DVD in the drive, type l1a_dvd at the command line to begin DVD burning. It is useful to check the size of the $L1A tree to determine which step needs to be done. It looks like this:

```
zorro:/hafb/l1a# du -h
54G     ./0_core
5.1G    ./1_get
1.8G    ./2_dvd
11G     ./3_hold
80G     ./4_dvdgz
8.0K    ./5_attic
151G    .
zorro:/hafb/l1a#
```

Assuming that 2_dvd is larger than about 4 Gigs, l1a_dvd is the right choice to make. Once you do this it will download and burn the necessary buffer files, once complete, the program will give you a disc number, time and date and you will have to scroll back in the terminal to see which buffer files it burned (usually a series of files in day of year format, so you are interested only in the day of year, 100-105 for example)

Take this DVD to the printing system, edit the label to have the right time, frame numbers and disc number, put the DVD in the tray and then print it, this part is picky, sometimes it spits the disc out (the printer is on it's last leg) so you have to take the whole tray out, hit the middle button on the printer and then put it back in and hit the middle button again.

Once the label is printed, it's time for step 4.

After burning the disc, the files that were put on the disc were copied to $L1A/3_hold . Now you need to take the disc and put it into a different system than the system which it was burned in (as sort of data integrity test) and as the user soft run l1a_diff . This will compare the files on the disc to the files in the hold directory. Give the system a little bit of time to automatically mount the DVD before running 1la_diff or you will see that it claims not to have found anything.

Once this has run, scroll back and make sure that 'trouble' didn't show up anywhere in the comparison (it is usually OK if you get a read error on l1a_dvd.txt file) If there was trouble, the

easiest thing to do is to eject the disc and try it in another machine, if this still happens, find someone to walk you through the next steps.

Since l1a_diff worked perfectly, the verified files are compressed and move to $L1A/4_dvdgz

Now to cover a few little details, because this process is pretty linear, you can burn as many DVDs as needed at once, but each one will contribute to the NFS load of the server, typically 3-4 at a time is reasonable. The DVD burners in the 'S' machines, star, sid and sun are the most reliable, so use those to burn. Once you have labeled the disc, you should l1a_diff it in a different machine, typically the 'Z' machines, zelda, zulu, ziggy or zorro. It is ok to diff them in an 'S' machine as well, but preferably not the same 'S' machine it was burned on.

After completing a set of DVDs, and using the du -h command on the $L1A directory to confirm that the 3_hold directory is empty (it should be 8.0K if it has been cleaned up by l1a_diff) you can move to the next step, if you missed a 'trouble' error on one of the discs, you will see a stray .buf file in the directory, you will have to determine which DVD it came from (from the DOY) and re-run l1a_diff.

Next part:
4-> burn gzipped DVDs.

Once the 2_dvd directory is low in size (2-10 G) and 3_hold is empty, you can start burning gzipped DVDs, there is a time lag however, a certain amount of time (about a week) must pass before the gzipped files are able to be burned. Once there are ample files in 4_dvdgz, as soft user, run l1a_dvdgz and repeat a similar process to the previous DVD burning step. Print a label with the same information but make sure it is a DVDGZ label in the Epson CD labeler program.

Now just like before you need to put the disc into a different machine, this time running l1a_diffgz. This comparison completes the backups, deleting the gzipped file from 5_attic where it was moved when the file was burned to disc. If there is no trouble, just drop the disc into a case and that's it. These are also nice to do in bulk, 3-4 at a time for burning and differencing.

Section on IPS Tomography and NV3F files

Section on VOX files and VolumePro board

Anything Else

Contacts:
Bernard V. Jackson : bvjackson@ucsd.edu
Andrew Buffington : abuffington@ucsd.edu
P. Paul Hick : pphick@ucsd.edu
Mario M. Bisi: mmbisi@ucsd.edu
Elizabeth A. Jensen : ejensen@ucsd.edu
John Clover : jclover@ucsd.edu