Documentation

- More mundane details about CP-V on the Sigma 9 for Living Computer Museum users is overviewed in another set of slides.
- This document gives only the briefest overview of CP-V.
- For more technical details consult the following manuals:
- Xerox CP-V Operations Reference (901675G).
- Sys. Programming (903113A) and Sys. Management (901674G).
- Xerox CP-V Data Base (901995C).
- Xerox CP-V Batch Processing (901764F).
- There are also a myriad of other manuals, including my favorites the UTS Technical manuals.

User Privilege

- The default user privilege is X'40'.
- This is stored in byte 0 of X'8C43' in the user's `JIT`.
- Privilege X'80' allows a user to map for reading the running monitor.
- At AU pages containing the COC ring and input buffers were mapped without error onto page 0 (which any user can access).
- Privilege X'A0' allows a user to initiate a ghost job (from :SYS).
- Privilege X'B0' allows a user to map for writing the running monitor.
- Privilege X'C0' bypasses normal file security and is required for serious system programming work.

File Security

- Each user normally uses files in his login account.
- The only file hierarchy is the account and file directories.
- The account directory (`KEYM=8`) and file directories (`KEYM=31`) are processed much like regular keyed files.
- Files may be `READ NONE/ALL/CALKINS`, *etc.*
- Files may also be `WRITE NONE/ALL/CALKINS`, *etc.*
- Files may also be `EXECUTE NONE/ALL/UNDER`, *etc.*
- Up to 8 `RD/WR/EX` accounts (each) may be specified.
- `UNDER` will allow one processor to be named.
- `!COPY FILE OVER FILE(WR(UNDER),UN(APL),EX(ALL))`

Star/Temporary Files

- There are 7 star or temporary files kept track of in the Assign-Merge record.
- Example: if you `!FORT` with no arguments the binary output will go to file `$` through `M:GO`.
- You can then `!LYNX $` or `!LYNX` and create a temporary lmn.
- To list/copy these files with `!PCL` specify them as '.G' or X'0011C7'.
- X'C7' is the EBCDIC code for G and 11 is your user number.
- The temporary load module is X'0011D2' or '.L'.
- Consult the Assign-Merge documentation for more details.

Getting Privilege

- Getting privilege can be problematic.
- If logging in to the :SYS,LBE account (:USERS does not exist) you should receive X'C0' privilege.
- However, at AU we had `SYSPRIV` lines in `LOGON`.
- I set this to X'FFFF0000' on the CPCP PO tape, else it would drop you down to X'A0' privilege (and not log in :SYS).
- With front panel access, you can execute a spin loop (B $).
- Then you are likely the current/mapped user and you can set your privilege from the Sigma 9 front panel.
- AU also had `PRODPRIV` lines. Two letter account names could only log on to these lines.

SUPER not in AM account

- SUPER (sometimes known as WHO) has security checks.
- Even with X'C0' it may fake you out with: `A603 LOAD MODULE DOES NOT EXIST.`
- Run it under `DELTA` and check for load words early from .8C01 and .8C02 where it may check for various accounts.
- By-pass the failure (exit true) for both and you should be good.
- Similar checks are in GRAN for rewriting a granule (CPCP only).
- The CPCP boot tape was changed with the expected password.
- The CPCU boot tape version doesn't even log to the console.
- esc esc DELTA ;I= (tab) will get you to the code to bypass.

File Types

- There are three user file types: random, keyed, and consecutive.
- Random files are contiguous granules and depend on the user for structure. Use `FOPEN.D` to create.
- Keyed files can have a key max of up to 31.
- Edit files have a key max of 3.
- Consecutive files don't allow direct record access, only sequential.
- Keyed file records can be access directly or sequentially.
- CP-V keeps track of record ends (*i.e.* it isn't determined arbitrarily by CR/LF or LF record separators.

GENMD

- `GENMD` is a program to modify load modules.
- It is part of the boot process to patch system processors.
- ,+loc,value . 12/19/14 KGC would be a common entry.
- A segment name may appear before the comma and a symbol afterwards.
- For CPCU boot tape you will see many changes listed this way in the console output (which were done on-line).
- `LIST` is a very useful `GENMD` command.
- The changes are kept in keyed records starting at X'FF0000'
- However, the `DELETE` command removed this audit trail!

ANLZ

- `ANLZ` allows the systems programmer to examine/change the running monitor. Invoking `<DELTA` allows direct monitor access.
- `ANLZ` also assists in analyzing crashes.
- `GHOST1` launches `ANLZ` automatically after a crash to produce a half inch thick printout (depending on memory size & users).
- It knows the current crash number (`mod 8`) and opens the corresponding `MONDMPx.:SYS` file where $0 \leq x \leq 7$.
- Sometimes CP-V crashes so hard it asks: `ENTER I/O ADDRESS FOR TAPE DUMP`.
- `!GJOB ANLZ` and `<IN TA` will request and process the tape.
- `<RUN PP` will drill through the processor page chains.

CONTROL

- `CONTROL` allows the X'C0' privileged user to change system parameters.
- It allowed the X'80 privileged user to examine system parameters.
- `CONTROL` prompts with `CONTROL HERE` and - on the next line.
- `-OUM=107` (CPCP) or `-OUM=64` (CPCU) sets online maximum users.
- `-BUM=8` (CPCP) or `-BUM=6` (CPCU) sets batch maximum users.
- Thus these correspond with `!ON 107` and `!ONB 8` console keyins.
- A wide range of other system parameters can be examined/changed to tune the system for optimum performance.

DRSP – Dynamic Replacement of Shared Processors

- SYSMAK0/1 in GHOST1 initially writes the monitor, overlays and shared processors to the swap device on a tape boot.
- DRSP allows changes while CP-V is running.
- DRSP salutation is `DRSP HERE` and on the next line >
- Typical commands are `>LISTALL #10-12`, `>LIST #10-18`, `>DROP APL;` `>ENTER`, or `>REPLACE`. `>ENTER` and `>REPLACE` take an extended argument field which includes `FROM source`, and optional arguments following separated by commas: `PERM,J`, *etc.*
- `PERM` is short for permanent and will/not load it after the next reboot (after clean take down) or crash.
- J is for `JIT` access and there are other flags.

ELLA

- `!ELLA` is the error log lister. It wants .A0 privilege to awaken the `ERR:FIL` ghost and access to `ERRFILE.:SYS`
- Typical commands are `SU` for summary, `CL` for chronological listing, `DE,F0` to select device `F0`, and `TI,12:00` to start at a given time.
- Others include `TY,23` to select only type 23 entries (time stamps), `TY,0` turns that selection back off.
- Using ELLA to analyze hardware errors is an essential skill.

Fast Save; Fast Restore

- Telefile replaced `FSAVE` with `BACKUP` and `FRES` with `RESTORE`.
- `FSAVE/FRES` used "spin-loops" (`B $`) while waiting for I/O to finish and thus were not at all friendly to any other CP-V users!
- They also followed their own logic in accessing the file system.
- `BACKUP/RESTORE` used OS logic to access the file system.
- `BACKUP/RESTORE` introduced superblock—tape blocks $>$ 2K.
- George wrote code to implement superblock logic into `FSAVE/FRES` and eliminate the spin loops.
- Look for `JCL` in `:SYS/:SYSJOBS` to emulate.
- Non-superblocked save tapes are labelled tapes and hence "PCL-compatible."

FSAVE/FRES – 2

- Notes for PO tape `CPCU`:
- This PO tape has addresses `A80/1` identified as `BT`s.
- I checked that these are the Wangs (thus proper drivers).
- Thus to FRES from you will need to insert the following into `FRESJCL` after the `!LIMIT` and before the `!FRES` cards.
- `!ASSIGN M:EI,(DEVICE,BT).`
- You will also need to change `(9T,1)` to `(BT,1)` on the `!LIMIT`.
- This can be done online with `!EDIT FRESJCL.:SYS` and `!BATCH FRESJCL.:SYS` online or ghost.
- In ghost you will not need to specify the :SYS account.

SPY

- SPY allows privileged users to perform certain functions.
- Initiating ghost jobs via the GJOB n/l command is one.
- If an account is specified, expect the system to crash with a SC 7E-40.
- Commands are: DISK, GHOST, COC, IOQ, HELP, MEMORY, BATCH.
- These can be abbreviated to one letter.
- >M will output memory usage about every 12 seconds.
- P7 was written at AU to give better processor tracking.

SYSCON

- `SYSCON` or SYStem CONtrol was used primarily to partition/return private disk drives and tape drives for/from maintenance.
- Major commands are `PA`, `RE`, `LIST`, and `END` for partition, return, list, and exit.
- `PA` and `RE` can be spelled out too.
- `PA` and `RE` take a symbolic address (A80 for the tape drive at .80) as an argument.
- The `PA` and `RE` commands would automatically reduce/increase the total number of disk spindles or tape drives available on the system.
- `CONTROL` also allows you to change `TSP` and `T9T`.

Other Important Processors

- GRAN (absolute granule editor) CPCP write password=NOT4USE!
- RADMAP map of OS on disk for use with GRAN!
- The D account has many useful tools like `FP`, `KEYLIST`, `SIZELIST`, and `UKFE`.
- `GRAN.D` and `OC.D` often come into :SYS via the PO tape.
- `OC` allows the privileged user to enter console commands online.
- These are logged on the real console. It may be only `KEYIN` type commands work (*i.e.* no device retries allowed).
- Similarly, the X account has many useful programs, primarily for the priviledge X'80' user, however.

TIME, ERR, HELP, *etc.*

- `TIME` returns the system time and date.
- `ERR xxx` will expound on error code `xxx`.
- The `ERRMSG.:SYS` file must exist for that to happen.
- Software checks can be examined via `ERR 87E40` for a 7E-40.
- `HELP acc` will process the `HELPFILE` in a given acc, such as `:SYS`, `D`, `X`, `STATL`, `GAMEL`, `GAMEB`.
- If there is no `HELPFILE` in the current account, and no argument was given, it will prompt for an account.
- `OP` lists the options, which are the lines with characters in the first 4 columns.
- `ACC` allows you to change accounts; `END` exits.

DEVDMP

- DEVDMP or Device Dump is an off-line processor to COPY a physical disk (RAD or removeable) to tape as a logical unit.
- It writes a copy of itself on the tape and thus the tape can be used to restore the disk.
- Example command: `#7212,DC3F0,COPY[,MTxxx]`, where the `MTxxx` portion is optional.
- 7212 is a high speed RAD serial number of type DC.
- Use 7277 for an 86 Mb disk and device type DP.
- Just enter `REST` or `REST,MT080` to restore it.
- The Telefile version seems to have changed `>COPY` to `>SAVE` and `>CMPR` to `COMPARE`.

VOLINIT

- VOLINIT or VOLume INITializer is another off-line processor.
- However, later versions were able to be run on-line.
- Specifically it did a surface test, wrote headers, and assigned alternate tracks for bad spots by flawing them.
- We generally included flawing information on the label for later comparision since running it destroyed previous information.
- An on-line program, INITVOL.D was used to establish serial numbers and account information for private packs afterwards.
- When I started in 1978 a 6" card deck was used. I soon converted it to tape!

BOOTing

- Generally when you boot the Sigma 9 you `MEMORY CLEAR, SYS RESET, I/O RESET, LOAD`, and then move the `COMPUTE` switch to run.
- This assumes you have set the `UNIT ADDRESS` appropriately.
- X'80' and X'81' are currently the Wang tape drives.
- X'F0' is the current LCM boot disk address.
- In software these are generally referred to as A80 and AF0 or 9TA80 and DPAF0.
- An important exception is for an `Operator Recovery` on CP-V when you ensure `SENSE 3` is illuminated (on), DO NOT `MEMORY CLEAR`, but then boot from disk.

PO Tape Boot–1

- The Sigma boot tape is known as a PO (punched output) tape, from the days of paper tape!
- The tape is a standard Xerox labelled tape after the monitor and processors.
- As such you can space to the first file marker and back up two records to find the `:LBL and :ACN` records.
- The standard Sigma 6/7/9 IPL (initial program load) loads 88 bytes into locations .2A-.3F.
- On a tape boot a set of instructions are read in to read the monitor root and branch to its start address `INITIAL`.

PO Tape Boot—2

- .2A LI,1 76 /.2210004C/ number of pages to read (CPCP) value.
- .2B LI,2 .698 /.22200698/ special byte advance value for page 0.
- .2C LI,4 10 /.2240000A/ retries.
- .2D LI,0 DA(.3A) /.2200001D CDW address.
- .2E SIO,0 *.25 /.CC000025/ Start I/O.
- .2F TIO,3 *.25 /.CD300025/ Test I/O.
- .30 BCS,12 .2F /.69C0002F/ Check for I/O completion.
- .31 CW,3 .3F /.3130003F/ Check status.
- .32 BANZ,4 .37 /.69400037/ Need to retry?
- .33 AWM,2 .3A /.6620003A/ Advance destination.
- .34 LI,2 .800 /.22200800/ Advance amount for other pages.

PO Tape Boot—3

- `.35 BDR,1 .2C /.6410002C/ Keep reading.`
- `.36 B INITIAL /.6800770E/ Start OS (CPCP value).`
- `.37 LI,0 DA(.3C) /.2200001E/ Set CDW address.`
- `.38 BDR,4 .2E /.6440002E/ Decrement/retry.`
- `.39 B $ /.68000039/ Spin (if no more retries).`
- `.3A DATA X'02000168' Read to BA(.5A).`
- `.3B DATA X'00000800' Size of a page (max read).`
- `.3C DATA X'4B000000' Space Record Backward.`
- `.3D DATA X'20000000' Command Chaining.`
- `.3E DATA X'08000000' Transfer in Control.`
- `.3F DATA X'6EFE0000' Ignore Int pending, auto, reserved.`

PO Tape Boot—4

- At `INITIAL` in module `INITIAL` it checks if the `SYSGEN` matches the running configuration, *i.e.* Sigma 6/9, X560.
- The Sigma 9 traps to .4D on instruction exception if an Add Doubleword is done on an odd register.
- The external interrupts are set to `MTW,0 0`.
- The internal traps and interrupts are initialized.
- I/O and `CLOCK3` are armed and enabled.
- It then branches to `MONINIT` in module `BOOTSUBR`.
- When it returns it initializes write locks, memory map, and monitor `JIT`, starts up `ALLOCAT, GHOST1`, arms/enables the rest of the interrupts, and branches to the scheduler `T:SE`.

PO Tape Boot—5

- `MONINIT` checks for what type of boot this is.
- If it is a tape boot it prompts for device address changes for the card reader, line printer, and swapper.
- If the console address must be changed it spins.
- It determines how much memory is available (up to the genned amount anyway) and sets up the free page chain, taking care to allocate the monitor `JIT` and `UMOV`.
- It also requests the current date/time.
- If not a tape boot it asks if you want `DELTA`. This refers to `XDELTA`, the executive debugger. If not its pages are released.

Disk Boots—1

- .2A RD,0 0 /.6C000000/ Check for SS3 (Op. Rec.)
- .2B BCR,2 .2E /.6820002E/ Skip OR
- .2C XPSD,0 .3994 /.0F003994/ Do OR
- .2D DATA /.FFFF0000/ Screech Code for OR
- .2E LI,0 DA(.38) /.2200001C/ CDW address for I/O
- .2F SIO *.25 /.CC000025/ Start I/O
- .30 TIO *.25 /.CD000025/ Test I/O
- .31 LI,0 15000 /.22003A98/ Spin value
- .32 BDR,0 .32 /.64000032/ Spin
- .33 BCS,12 .30 /.69C00030/ Go back and check status
- .34 B *.5F /.E800005F/ B to `INITIAL`

Disk Boots—2

- .35 DATA /.00010C00/ cyl, head, sec for .5A–.3FFF
- .36 DATA /.00020000/ cyl, head, sec for .4000–.7FFF
- .37 DATA /.00020509/ cyl, head, sec for .8000–.8271
- .38 DATA /.030000D4/ Seek using BA(.35)
- .39 DATA /.2E000004/ BC=4
- .3A DATA /.02000168/ Read into WA(.5A)
- .3B DATA /.2E000000/ BC=0 (64Kb)
- .3C DATA /.08000523/ Transfer In Control to DA(.A46=BOOTIC)
- .3D DATA /.00000000/
- .3E DATA /.00000000/
- .3F DATA /.00000000/

M:MON

- M:MON is the internal name of the CP-V operating system.
- It is built by the LOADER as a paged and overlaid load module.
- Historically this was done via a LOCCT file built by PASS2 and passed to the LOADER. Telefile and Andrews used LYNX.
- There are two types of overlays: mapped and unmapped (UMOV).
- The mapped overlays can be associated with a user.
- UMOV is an extension of the operating system and contains non-user code such as scheduler, swapper, and device handlers.
- The monitor root is mapped one-to-one so it can run mapped or unmapped.

M:MON overlays

- The mapped monitor overlays start at .8000 and must end by .8BFF and are thus a maximum of 6 pages.
- UMOV starts at .8E00 and was rather large in CP-V F00 to contain the micro-processor controlled device code.
- The unmapped monitor overlay is optional, unless your system is too large (users, devices, CFUs, *etc.*)
- SUSPTERM had to be below .8000 as I recall.
- The monitor overlays have names like: STEPOVR, RMAOV, OPEN, MULOV, MISOV, LTAPE, LDLNK, KEYIN, OPENTP, ENQOV, ECBOV, DEBUG, and CLOSE. KEYIN runs as a ghost.
- Additional overlays like: TQOV1, TQOV2, RTOV, and UMOV (the unmapped overlay) are optional.

JIT

- Every user, including the Monitor (except KEYIN?) has a Job Information Table or JIT. Location .4F points to it.
- It is always located at mapped (and for the monitor unmapped) location .8C00 and extends to the end of the page: .8DFF.
- If the BIG option was sysgened or the user gets too ($>$ 19Kw?) large, he also had an AJIT (additional JIT) page at .8E00 to contain his swapping command lists and page chain.
- The first few bits indicate if the user is online (.80) or ghost (.40).
- The rest of that word contains his user number.
- The next words have his user name (2) and account (3).

JIT

- JB:PRIV is at .8C43.
- J:CCBUF was at .8CC8 for CP-V C00 but .8CCC for later releases.
- Thus CPCP has it at .8CCC to facilitate maintaining and testing software for Telefile, whereas the other boot tape (CPCU) has it at .8CC8 and thus some incompatibilities exist.
- .8C28 points to the TREE; .8C2B points to the DCB chain.
- M:UC is at .8C2C and is a system DCB (Data Control Block).
- TSTACK is at .8C4E and is a temporary stack used by the monitor and command processors.
- The rest of the monitor's JIT (to .8DFF) is allocated to the monitor's TSTACK.

SBUFx and Assign/Merge

- The monitor utilized the address space between the AJIT and user code (.A000—.1FFFF) for special buffers.
- Thus when tape or disk blocks are read they are processed in the pages .9000, .9200, ...9E00.
- These first two pages are known as `SBUF1` and `SBUF2` and any extra user pages are mapped there for processing.
- The monitor moves the information from `SBUFx` into user space as needed.
- At login all batch and online users are assigned a file storage granule for use as their assign-merge record.
- DCB assignment information is stored here and merged into running programs at each job step.

DCBs—1

- Data Control Blocks or DCBs are the linkage between user programs and file and print streams.
- `!SET` (`!ASSIGN` in batch) is used to establish the connection.
- Programs can modify these with appropriate `M:OPEN`, `M:READ`, and `M:CLOSE` monitor service calls.
- DCBs all start with either M: or F: where the M: are considered system DCBs and the F: user DCBs.
- CP-V maintains a set of system DCBs with various default settings.
- SI for source input and BO for binary output are assigned when a user specifies `!FORT fid over bo`, for example.
- Also: C, CI, CL, CO, DO, EI, EO, GO, LI, LL, LO, OC, PO, SL, UC.

DCBs—2

- FORTRAN unit numbers correspond with F:unit number.
- F:101 and F:102 are default terminal in/out.
- F:105 is card reader input; F:108 is line printer output.
- Thus these correspond with similar single digit IBM streams.
- Similarly, `COBOL` I/O streams have an `F:` prepended.

Disk Storage—1

- There are 4 types of disk storage.
- 1. The operating system uses up boot space when it is installed.
- For this we will include the monitor overlays and shared processors.
- Where the overlays and processors are can be found via DRSP.
- However, often a RADMAP was produced to document M:MON.
- 2. The second type is PSA (permanent swapping allocation) and includes #1.
- Since the size of processors (and number of slots) can vary, the system will check how many "cylinders" remain and reduce the number of allowed users accordingly.

Disk Storage—2

- The empty processor slots (`M:DUMLM`) can be filled with small processors `M:Jx` instead.
- 3. PER is allocated symbiont storage space. When this is exhausted, the system will use PFA.
- When the batch queue is lost, PER will be released.
- 4. PFA (Permanent File Allocation) is set aside for user files.
- Logging on requires 1 granule for your assign-merge record. This comes from PFA so you will enter state SQR if there are no granules.
- A patch to `LOGON` was enabled at times to steal a granule from `DUMPFILE` for privileged users. However, releasing this granule was likely not a good idea!

PASS

- PASS originated as a way for teachers in class accounts (nomenclature like COSC110) to change the protected password and look at password protected files.
- However, at times it also became a way for select users (system programmers) to obtain .C0 privilege on non-system lines.
- Thus PASS was loaded as a shared processor with JIT access.
- NEIN was a popular processor at some sites to allow the (privileged) user to change his log-in file directory.
- M:Jx, where $0 \leq x \leq 7$ was originally used at AU to make large/empty shared processor slots small.
- However, M:J0 eventually subsummed PASS's purpose when class accounts ceased.

User Space

- The virtual user space is from .A000 to .1BFFF. By issuing the `!EXTEND` command it can extend to .1FFFF.
- .1C000–.1FFFF is normally reserved for special processors like `TEL` and `DELTA`, and shared libraries (`:P11`, `:PCC`, `:PDD`).
- CP-V switches fairly seamlessly between these.
- The user's data area (protection type 00) is first.
- The user's DCBs (protection type 10) follows.
- The user's procedure (protection type 01) follows.
- User space above procedure can be obtained for various uses.
- LINK-built LMNs are limited to 2 pages of DCBs starting at .16C00 and thus may have virtual space available below that.