


EVLA Data Processing PDR

Scale of processing needs

Tim Cornwell, NRAO




Background



- EVLA correlator data rate will ~ 1000 times current correlator data rate
- Can 2009-era hardware handle the processing load?
- Can the software handle the processing load?



Scale of EVLA data processing



- Peak data rate out of correlator backend ~ 25 MB/s
- Total data volume for Peak 8-hr observation ~ 700 GB
- Floating point operations per float $\sim 100 - 10000$
- Peak compute rate ~ 5 Tflop
- Average/Peak computing load ~ 0.1
- Average compute rate ~ 0.5 Tflop
- Turnaround for 8-hr peak observation ~ 40 minutes
- Average/Peak data volume ~ 0.1
- Data for Average 8-hr observation ~ 70 GB
- Data for Average 1-yr ~ 80 TB





Scaling laws in computing



- “Rules of Thumb” by Gray and Shenoy
 - <http://www.research.microsoft.com/~jgray>
- Examples:
 - 1 Moore’s law: *Things get 4x better every 3 years*
 - 2 *You need an extra bit of addressing every 18 months*
 - 3 *Storage capacities increase 100x per decade*
 - 4 *Storage device throughput increases 10x per decade*
 - 7 *NearlineTape:OnlineDisk:RAM storage cost ratios are approximately 1:3:300*
 - 8 *In ten years RAM will cost what disk costs today*
 - 9 *A person can administer \$1M of disk storage*
 - 14 *Gilder’s law: Deployed bandwidth triples every year*
 - 15 *Link bandwidth increases 4x every 3 years*



Ops per second per \$






Detailed analysis



- Analyze processing in terms of FFT and Gridding costs
- Find scaling laws for various types of processing
- Express in terms of 450MHz Pentium III with Ultra-SCSI disk
- Use Moore's Law to scale to *e.g.* 2009
 - Performance/cost doubles every 18 months
- Many more details in EVLA Memo 24



Scale of EVLA data processing




- Typical cost equation

$$T^{mosaic} \sim 4 \cdot N_{mega-vis} \cdot t_{mega-grid} + 16 \cdot N_{cycles} \cdot N_{mega-pixel} \cdot t_{mega-FFT}$$

where units are in millions of visibilities or pixels



Detailed analysis



- Analyze processing in terms of FFT and Gridding costs
- Find scaling laws for various types of processing
- Express in terms of 450MHz Pentium III with Ultra-SCSI disk
- Use Moore's Law to scale to *e.g.* 2009
 - Performance/cost doubles every 18 months

| Observation | # pol | FOV arcsec | Cellsize arcsec | Pointings | Facets | Pixels | BW MHz | Freq res MHz | Vis chan | Image chan | IF's | T |
|--------------------------|-------|---------------|--------------------|-----------|--------|--------|-----------|-----------------|----------|------------|------|---|
| <i>primary beam (2D)</i> | 4 | 7200 | 0.3 | 1 | 256 | 24000 | 500 | 1.00 | 500 | 1 | 1 | |
| <i>primary beam (3D)</i> | 4 | 7200 | 0.3 | 1 | 1 | 24000 | 500 | 1.00 | 500 | 16 | 1 | |
| <i>ic of SGRA West</i> | 2 | 200 | 0.2 | 64 | 1 | 1000 | 70 | 0.5468 | 128 | 128 | 8 | |
| <i>nearby galaxy</i> | 2 | 600 | 0.5 | 1 | 1 | 1200 | 7 | 0.006 | 1166 | 1024 | 1 | |

| Observation | Data rate Mb/s | Total data GB | Image Mpixel | Visibilities Mvis | Minor cycles | single d | multiple d | mosaic d | Time d | # processors | rate TB/year |
|--------------------------|-------------------|------------------|-----------------|----------------------|--------------|-------------|---------------|-------------|-----------|--------------|-----------------|
| <i>primary beam (2D)</i> | 1.87 | 80.87 | 576 | 10108.80 | 10 | 28.50 | 35972.08 | 40.88 | 35972.08 | 71944.16 | 59.04 |
| <i>primary beam (3D)</i> | 1.87 | 80.87 | 9216 | 10108.80 | 10 | 130.48 | 194.08 | 232.88 | 130.48 | 260.96 | 59.04 |
| <i>ic of SGRA West</i> | 0.58 | 16.56 | 128 | 2070.28 | 100 | 19.97 | 296.85 | 34.20 | 34.20 | 102.59 | 18.14 |
| <i>nearby galaxy</i> | 0.65 | 56.58 | 1679.04 | 7072.12 | 10 | 38.30 | 122.53 | 56.96 | 38.30 | 38.30 | 20.65 |




For NRAO....



- Assume Moore's Law holds to 2009
 - Moore himself believes this.....
- Cost of computing for EVLA
 - ~ 10 – 20 processor parallel machine
 - ~ \$100K - \$200K (2009)
 - Archive ~ 50TB per year
 - ~ \$5K - \$10K (2009)
- Comparable to computing cost for ALMA
- Software costs
 - AIPS++ *as-is* can do much of the processing
 - Some development needed for high-end, pipelined processing
 - Some scientific/algorithmic work *e.g.* achieving full sensitivity, high dynamic range



For the observer...



- Moore's Law gives ~ 64 fold increase for a desktop
 - *I.e.* \$nK where $n \sim 1-3$
- Many projects do-able on (2009-era) desktop
 - *e.g.* 1000 km/s velocity range of HI for galaxy
 - *e.g.* Mosaic of SGRA West in all H recombination lines between 28 and 41 GHz
- Larger projects may require parallel machine or many days on a desktop
 - *e.g.* Full sensitivity continuum image of full resolution 20cm field
 - NRAO would provide access over the net



Current activities



- Installed SAN at the AOC
 - Increasing disk storage to 4TB
- Eight processor IBM Netfinity 370 server
 - Running RedHat Linux
 - 8GB memory
 - 1 TB disk

10Mbit

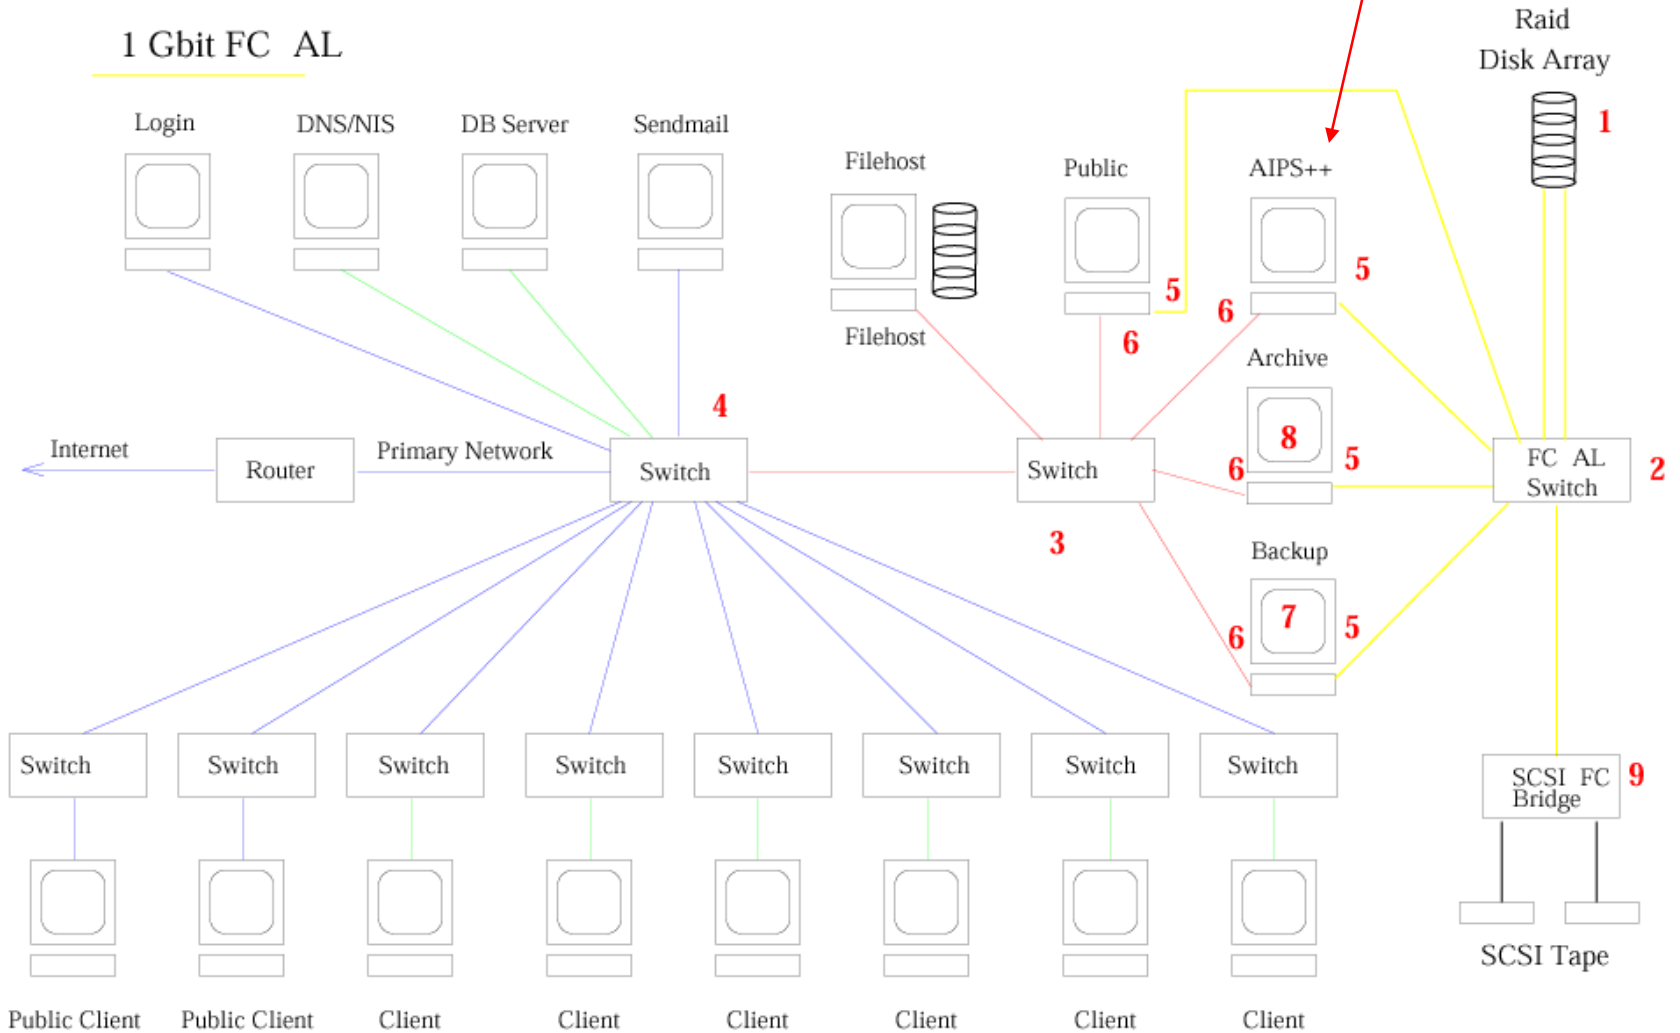
100Mbit

1 Gbit ether

1 Gbit FC AL


AOC Storage Area Network Diagram

Eight processor, 8GB, 1TB IBM NetFinity running RedHat Linux





Caveats



- Order of magnitude estimate
 - Development of computing uncertain
- Mix of science uncertain
- Processing requirements could escalate
 - Historically true
 - May be especially true for pipeline processing
- Moore's Law will throttle EVLA
 - As it did for the VLA