FAT PIPE FUNDAMENTALS: MANAGING BANDWIDTH-HUNGRY APPLICATIONS

Wed, Oct 1, 2003, 3:00-4:00pm

Larry Bergman

Email: Larry.Bergman@jpl.nasa.gov

Building the Connection Between Government, Business, Education, and Technology

GBET

NASA

JPL
Jet Propulsion Laboratory
California Institute of Technology

A CENTRIC Event

Key Questions To Be Addressed

- What are “Fat Pipes”?
- What are the Internet technology trends that are enabling *fat pipes*?
- What are some examples of *fat pipe* applications: past, present, and future?
- What should I be doing to plan for *fat pipe* applications?

GBET

Outline

- Overview of Network Technology Directions
  - **How We Got here – A Brief History of the Internet & Computing**
  - **Next Bandwidth Enablers**
- Emerging Fat Pipe Applications
  - **Consumer**
  - **Business**
  - **Government**
  - **Aerospace and Science**
- Future Provisioning Strategies
- Wrapup
- Q & A

I.

INTERNET TECHNOLOGY TRENDS

Internet Growth Statistics

- IDC expects Internet growth to double each year for next five years
  - **Currently 180 petabits/day**
  - **By 2007, expectation is 5,175 petabits/day**

GBET

Internet Growth

- What is driving growth of Internet traffic, particularly in backbone networks?
- Will Internet traffic eventually reach steady state growth like the telephone system today?
  - **5% to 8% annual growth**
- Is Internet traffic growth sustainable doubling every year?

Source: Bill St Arnaud, CANARIE

Internet Hosts - Overall Trend

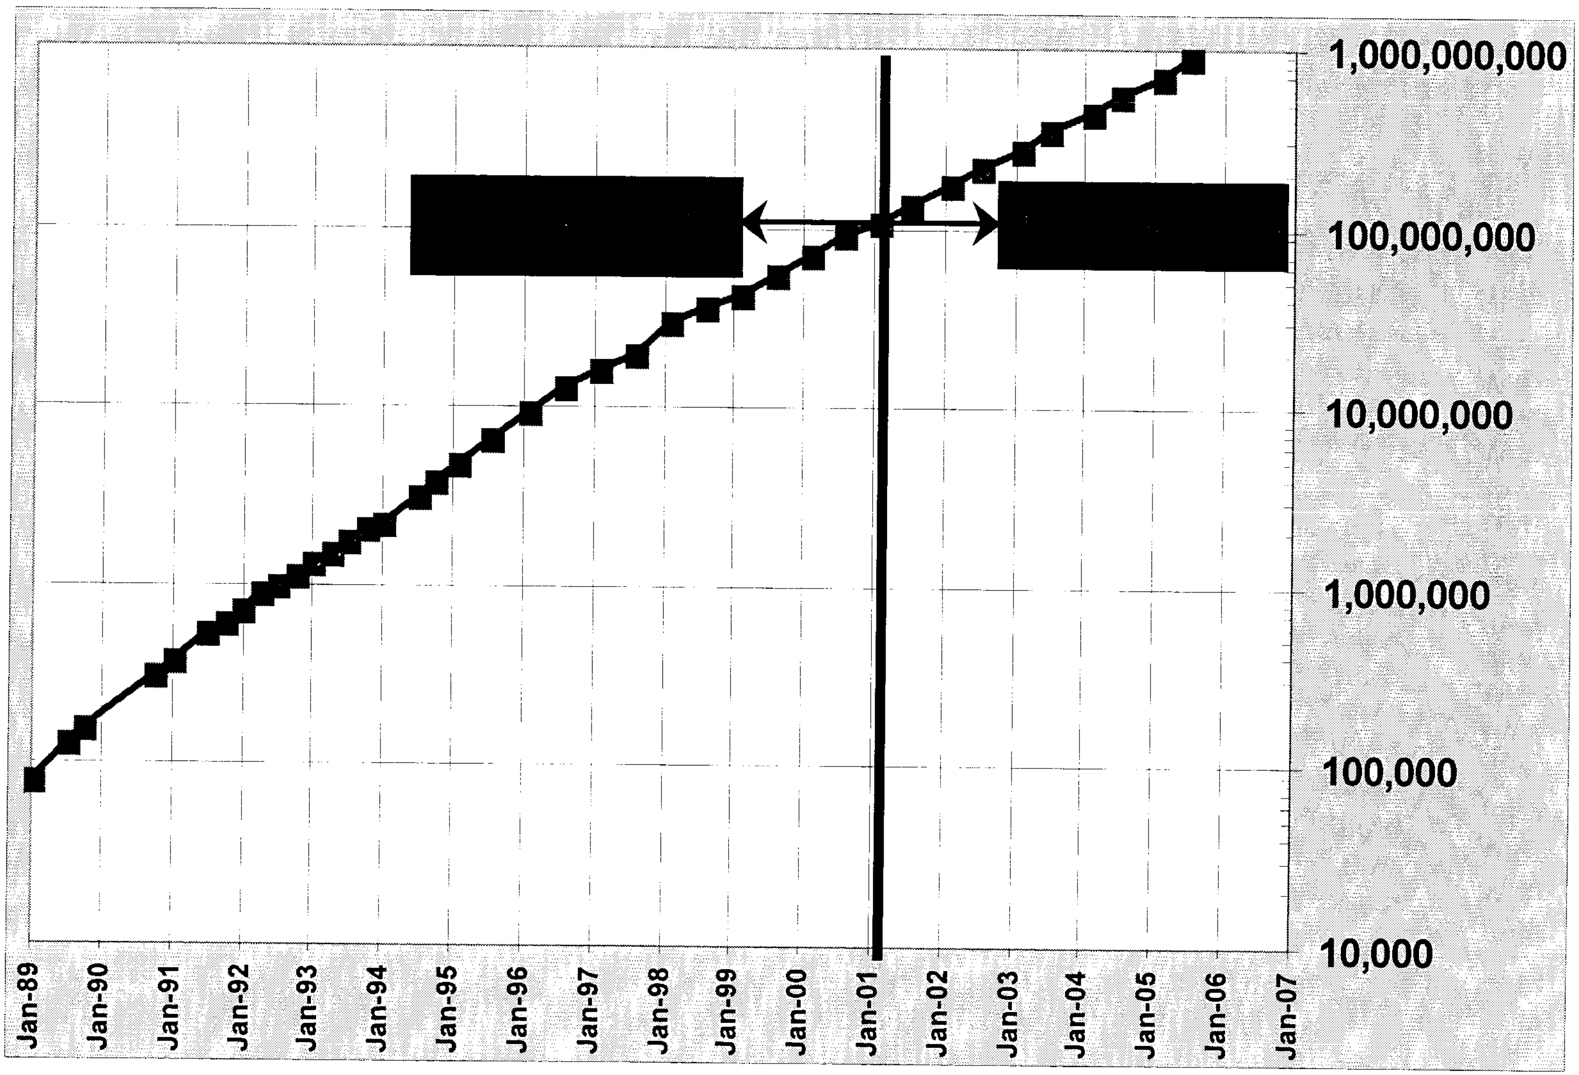

Source data:
M. Lottor,
Internet Software Consortium
<www.isc.org>

GBET

Factors Driving Internet Growth

- From backbone perspective what drives demand for bandwidth?
  - **New subscribers or more usage by each subscriber?**
- We know that new subscriber growth for telephone is ultimately limited by number of humans
- Internet growth is ultimately limited by:
  - **Not by number of humans**
  - **Not solely by number of connected computers**
  - **But by number of simultaneous TCP sessions (assuming no last mile bottleneck and CPU limitation) and pipe-size of individual connections (e.g., conversion to broadband)**
- Number of simultaneous TCP sessions is directly related to Moore's Law (at least capability) – and the types of applications
  - Internet growth (at a minimum) will be no less than Moore's Law

Source: Bill St Arnaud, CANARIE

To Date, Humans Have Been a Bottleneck To Internet Growth

- Humans can only do so much web surfing in a day
- Peer-To-Peer applications such as Napster, Grids and eScience allow computers to generate and receive traffic all the time
  - **These applications have much different traffic characteristics than human traffic**
    - Computers can maintain multiple connections at the same time
  - **Many simultaneous connections to different destinations**
  - **Each of which can burst to full capacity**

Source: Bill St Arnaud, CANARIE

What is a Fat Pipe?

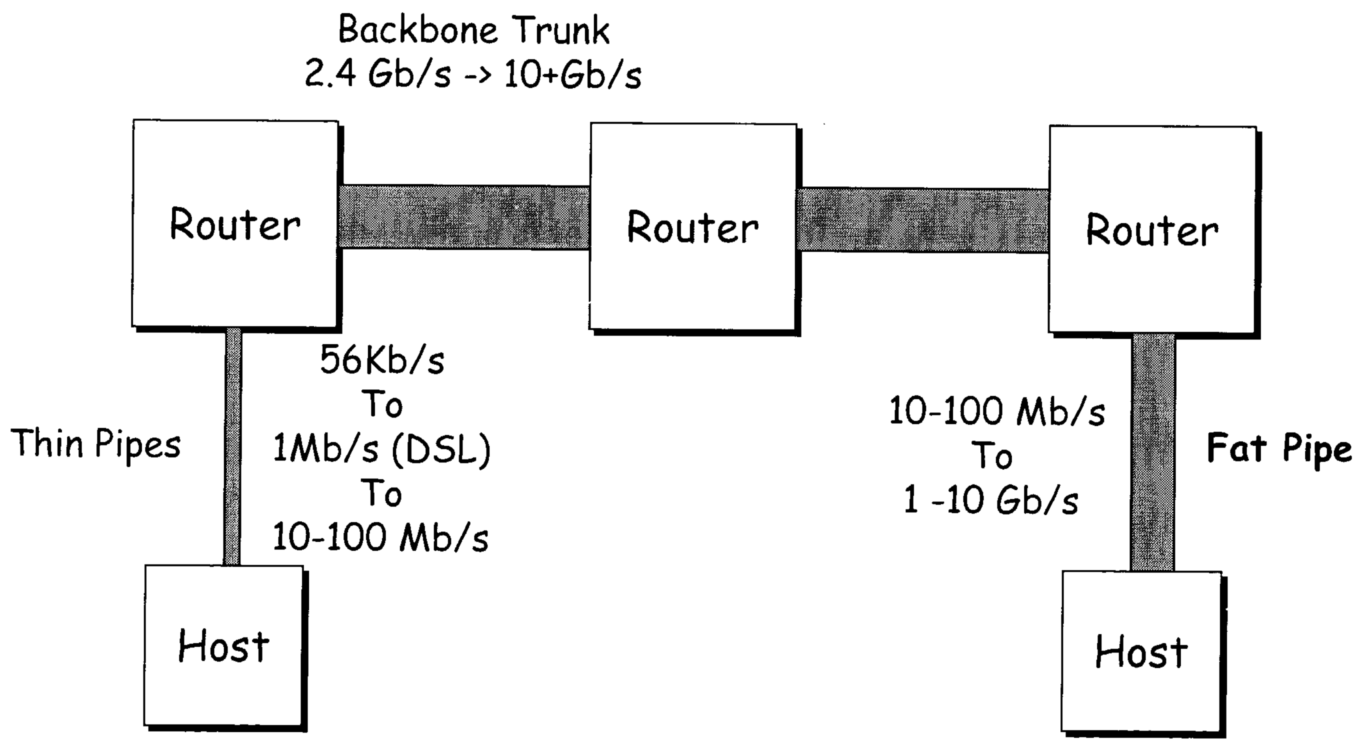

- And the goal post moves out each year for each tier!
- Example: TeraGRID uses 40Gbit/s!

Ingredients Feeding the Network Revolution

- Applications
- End Equipment
- Services
- Protocols
- Switching Fabric
- Transport Media

Computer Technology Drivers

- Very Large Scale Integrated (VLSI) Circuits
  - **Higher levels of integration**
  - **Faster clock speeds**
  - **Growing at Moore's Law**
  - **Approaching quantum limit**
- Micro-Processors
  - **Increasing speed (MHz to GHz)**
  - **Bus width and speed (64-bit emerging)**
  - **Multi-Instruction issue architectures**
- Memory
  - **KB to MB to GB**
  - **Still limits processor performance**
- Storage
  - **Magnetic disk, optical, tape**
  - **GB to TB**

GBET

Other Network Technology

- Protocols
  - **TCP/IP (v.6 is on the way … someday)**
- Network Interface Card Adaptors (NIC)
  - **1 Gbit/s Ethernet new de facto standard, 10 Gbit/s is on the way**
- Routers
  - **Custom processor IC's developed for router architecture**
  - **New functionality being added for security, striping**
- Fiber Optic Communications and Interconnects
  - **5 Gbit/s per lambda (1980) to 25 x 40 Gbit/s (2003)**
  - **Dispersion shifted fiber permits multiple wavelengths**
  - **Many advanced in III-V semiconductors and optoelectronics has resulted in compact O/E transceivers for 10 Gbit/s and beyond**
  - **All optical routing**

Are Computers Driving the Internet? or Vice Versa? (Pop Quiz)

- Which has more information capacity?
  - A) An Internet port (not trunk)
  - B) A state of the art computer
- If given the same parts, they should be about equal.
- However….
- Telecommunications system engineers often think serial at end points
- Computer system engineers often think parallel on several levels:
  - System buses
  - Multi processor architectures

II.

EMERGING FAT PIPE APPLICATIONS

Applications Evolving in Sophistication

- Early era (1970's and 1980's)
  - **Email**
  - **File transfer**
  - **Terminal emulation**

Mostly character oriented

- Web era (1990's)
  - **Hypertext**
  - **Object brokers (e.g., Corba)**
  - **Network savy applications and languages**

Mostly image oriented

- What's Next?
  - **Movies? HDTV? Or….**

Next Fat Pipe Applications?

- Caveat:
  - **This is one viewpoint (and I'm not that lucky!)**
  - **Predicting technology and applications that will be adopted by the public is always tricky!**

Next Fat Pipe Applications?

The possibilities:

- **High End Government, Aerospace, and Science Applications**
- **GRID**
- **Sensor Webs and Instruments (virtual instruments)**
- **Storage Area Networks (SANs)**
- **Home Entertainment Convergence**
  - Real Time Multimedia Services (HDTV, video conferences, games)
  - IP appliances
- **Enterprise**
  - Voice Over IP, video servers, storage networks,
- **E-Learning**
  - Video streaming and interactive services
- **Knowledge Management (KM)**

GBET

GRID

- An outgrowth of meta-computing from the early 1990's, it provides resource brokers for machines and people to locate data, cycles, storage, etc across the Internet automatically
- Takes care of all the mundane things people once had to do manually
  - **Earliest example: SETI @ Home**
- Attracted a lot of attention from industry for data and resource management within the enterprise
- GRID Forum formed
- Killer apps in this arena could be machine to machine!

CASA Gigabit Testbed

- One of first 5 gigabit network testbeds established by DARPA/NSF in 1989
- First to focus on meta-supercomputer applications in seismic image processing, global climate modeling, and chemical reaction dynamics
- Applications and network engineering team: Caltech, JPL, SDSC, LANL
- SONET OC-48 transport provided by: MCI, PacBell, and US West

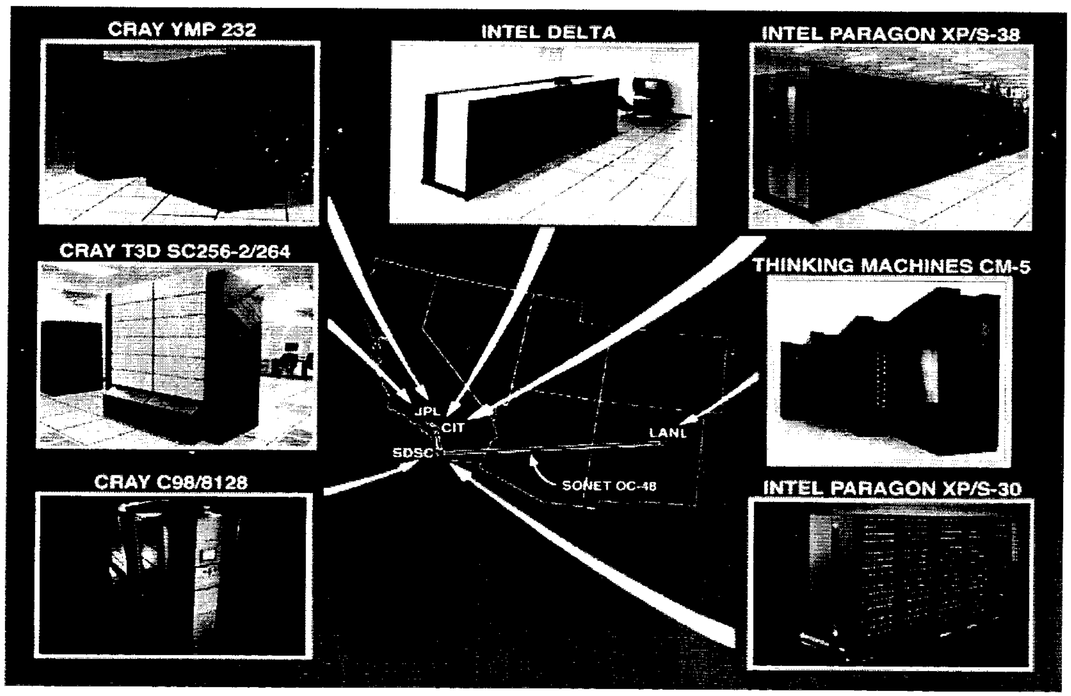

Salient Features

- HiPPI mapped onto SONET-48 frames (no ATM)
- HiPPI switches
- Firs t to demonstrate non-linear algorithm speedup with meta-supercomputer
- 300 GFLOPS
- 2000 km long
- 800 Mbit/s
- Achieved 500+ Mbit/s

NEED: 240,000 Mbit/s !!!!

CASA Application -- Modeling Landers Earthquake

Image Change Detection – Tracking Ground Movement

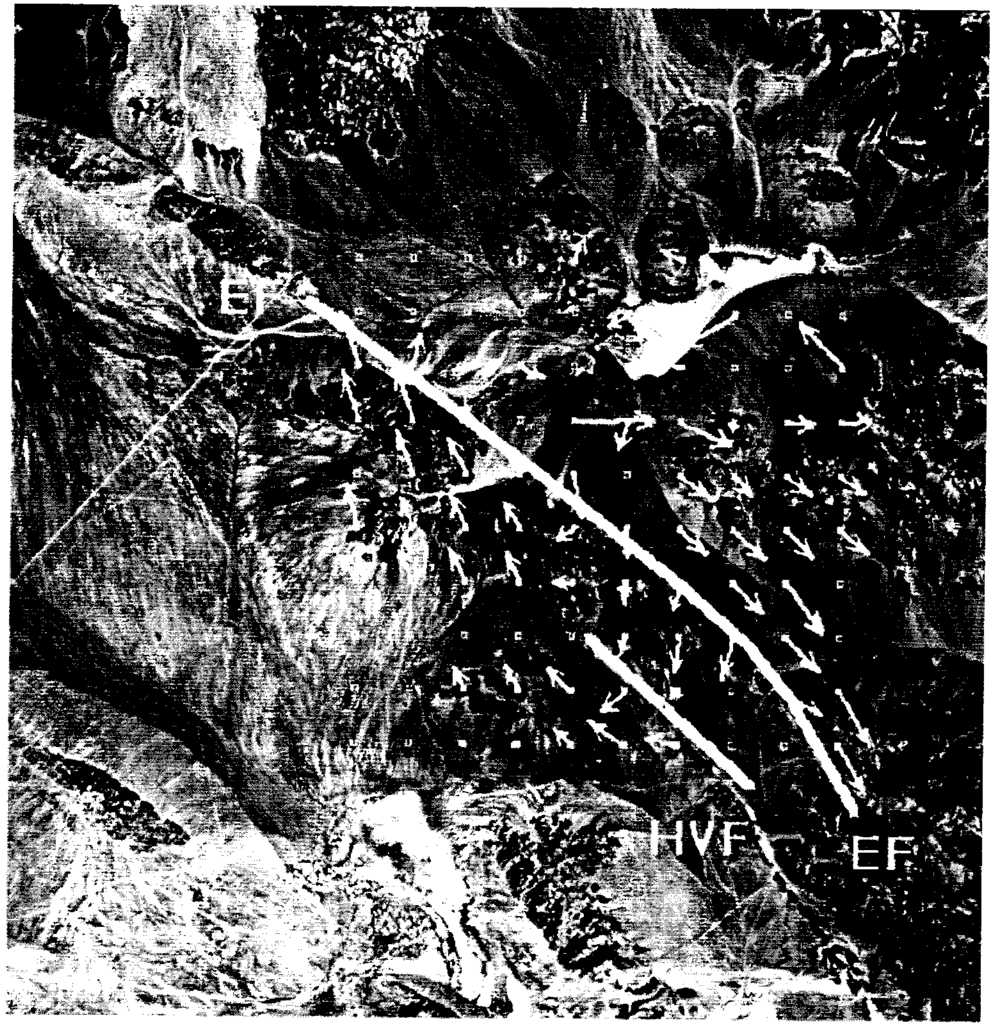

Each vector shows relative movement of a 100x100 pixel tile.

Source: Robert Crippen and Ron Blom / JPL

Pixel-Based Image Change Detection

Landers Earthquake, June 1992

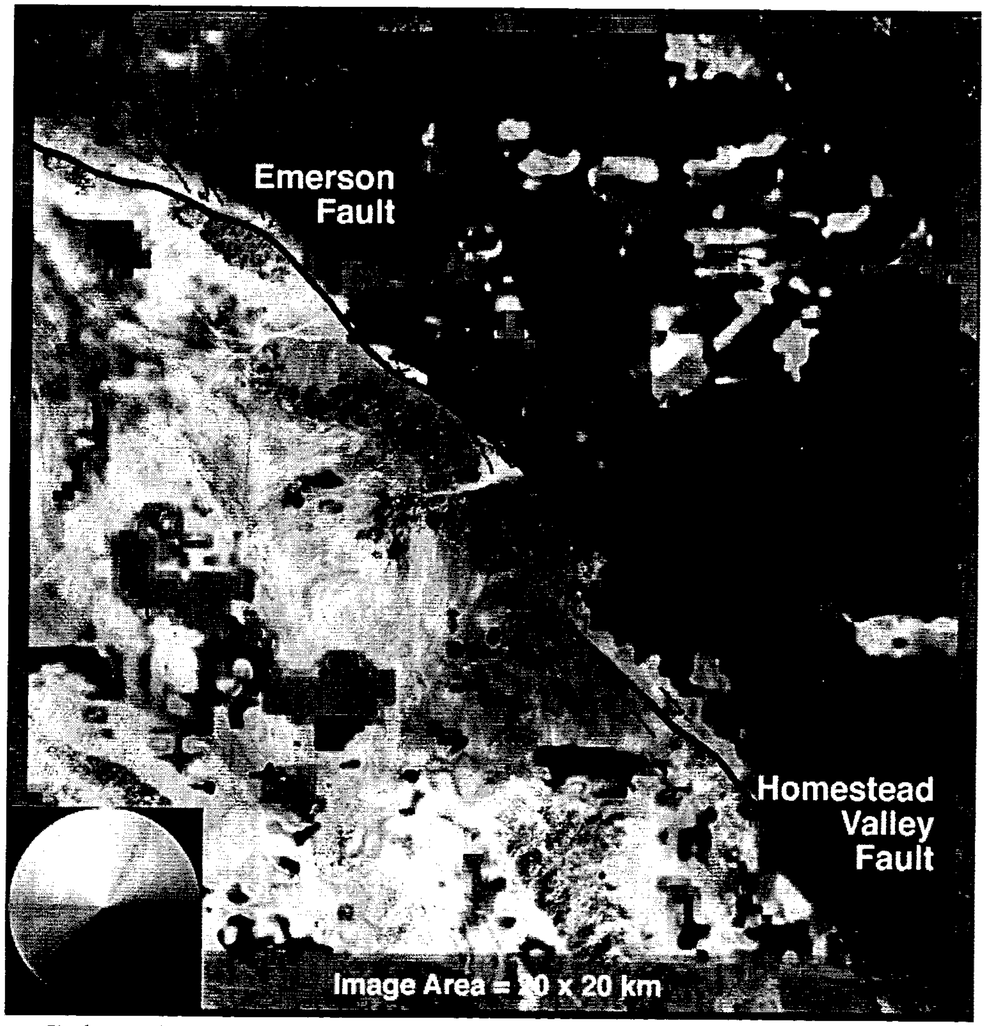

Color wheel encodes relative direction of ground motion for each pixel by comparison of similar tiles in two images

- Image shows displacement difference between two SPOT satellite images taken before and after the Landers earthquake
- 2050x2050 displacement map:
  - **10m pixel separation**
  - **1-2m precision**
- Hue discontinuity shows movement due to fault (black line).
- Unexpected activity appears in top right corner.
- Can be extended to monitor several images, and to general target detection.

Robert Crippen, Ron Blom, Paul Stolorz, Chris Dean / JPL

Desired CASA Data Movement

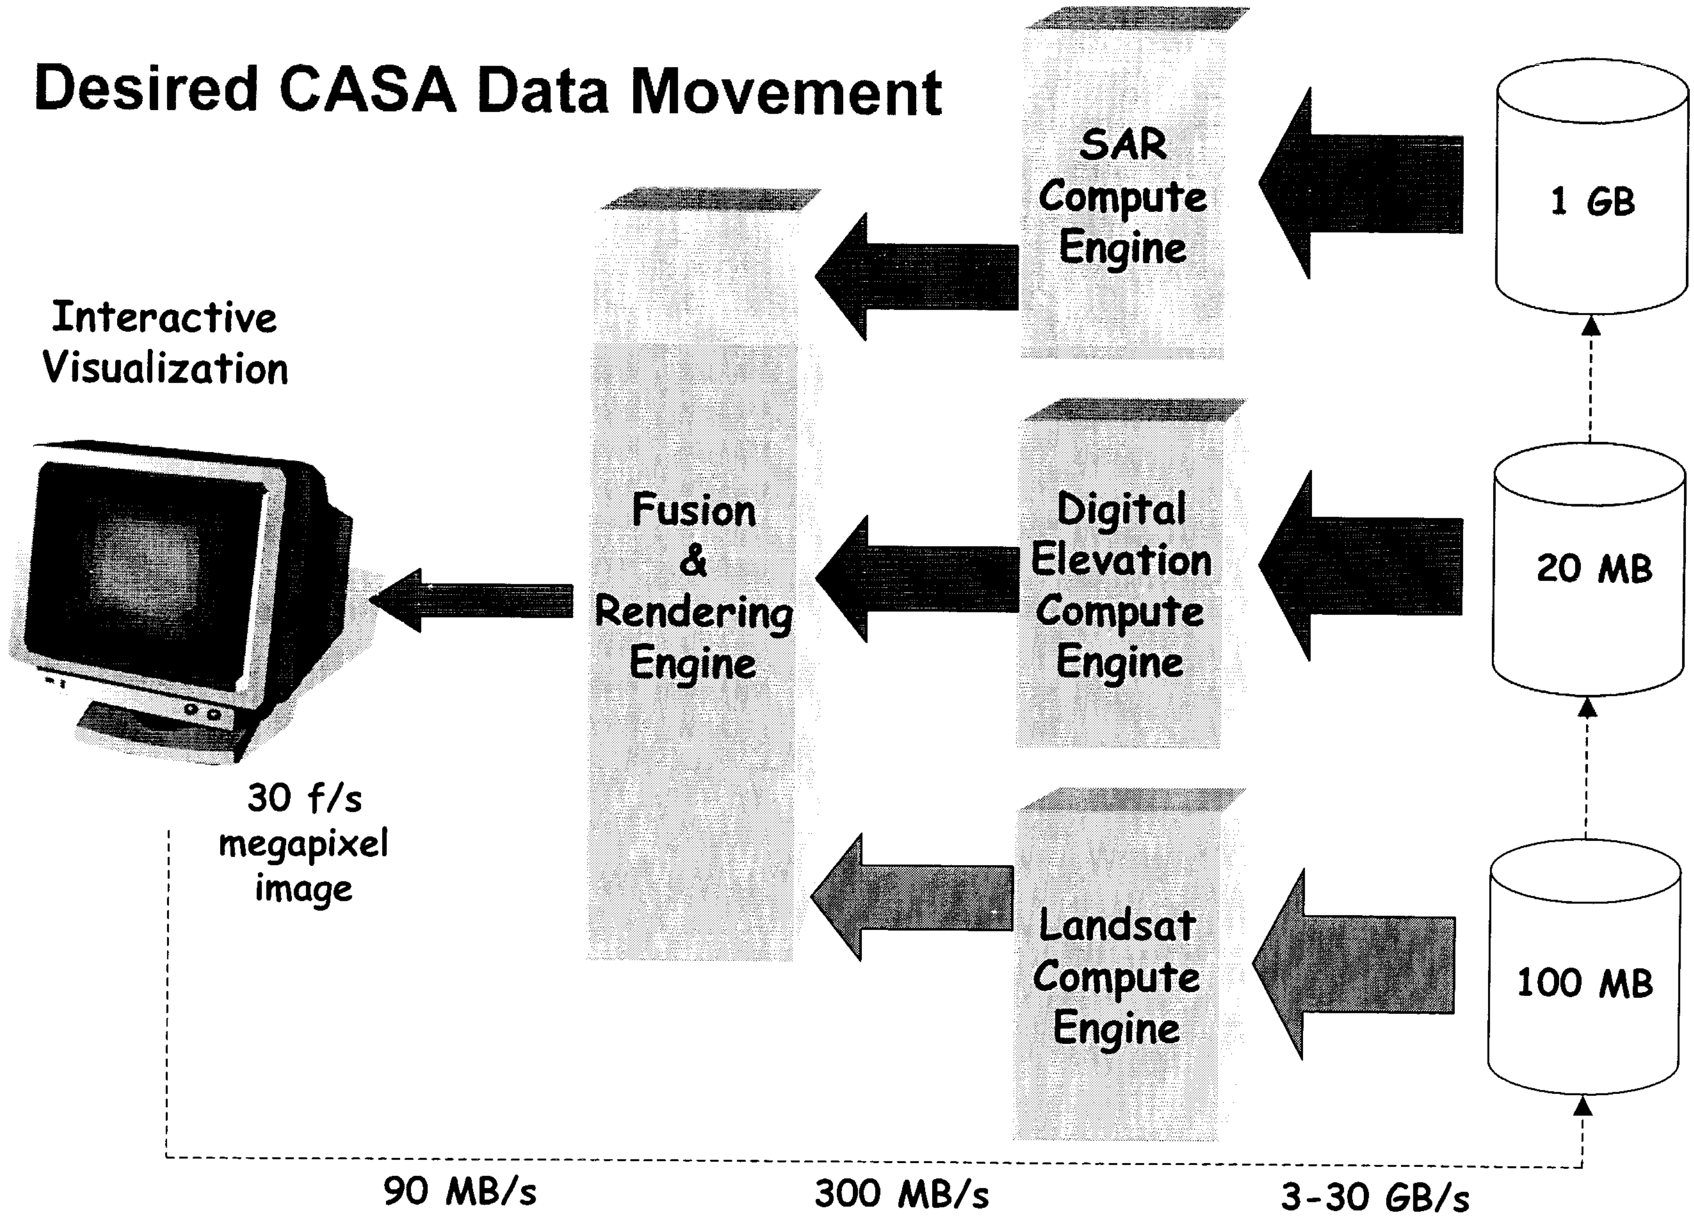

CASA Example: Important Provisioning Points

- Hosts
  - **Must be configured with appropriate Network Interface Controllers (NIC)**
  - **Internal I/O must support data movement commensurate with NIC (e.g., internal buses, CPU clock rate, DMA, ...)**
  - **Peripherals should also support target data rate (e.g. memory system, disk arrays, near-line storage)**
- Network
  - **Different parts of system require different bit rates**
    - Back end requires move data movement than front end
  - **Unless heavily buffered, visualization end may require dedicated bandwidth to insure low jitter**

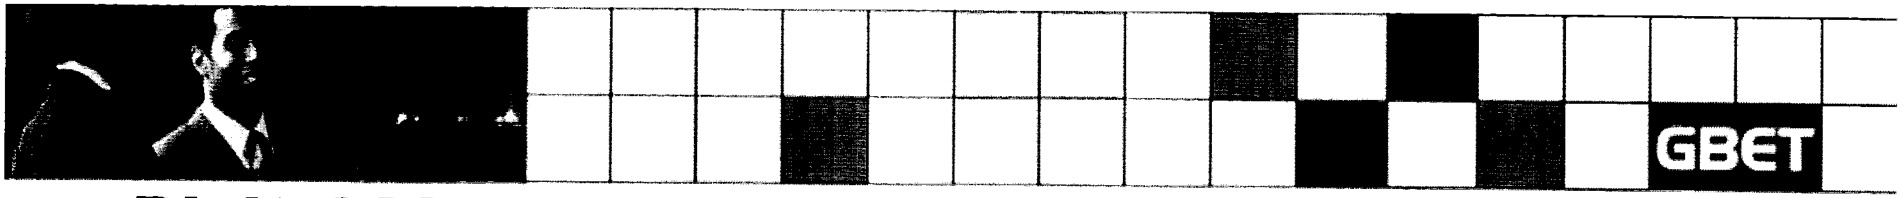

Digital Light Table
Landsat Mosaic - Western US

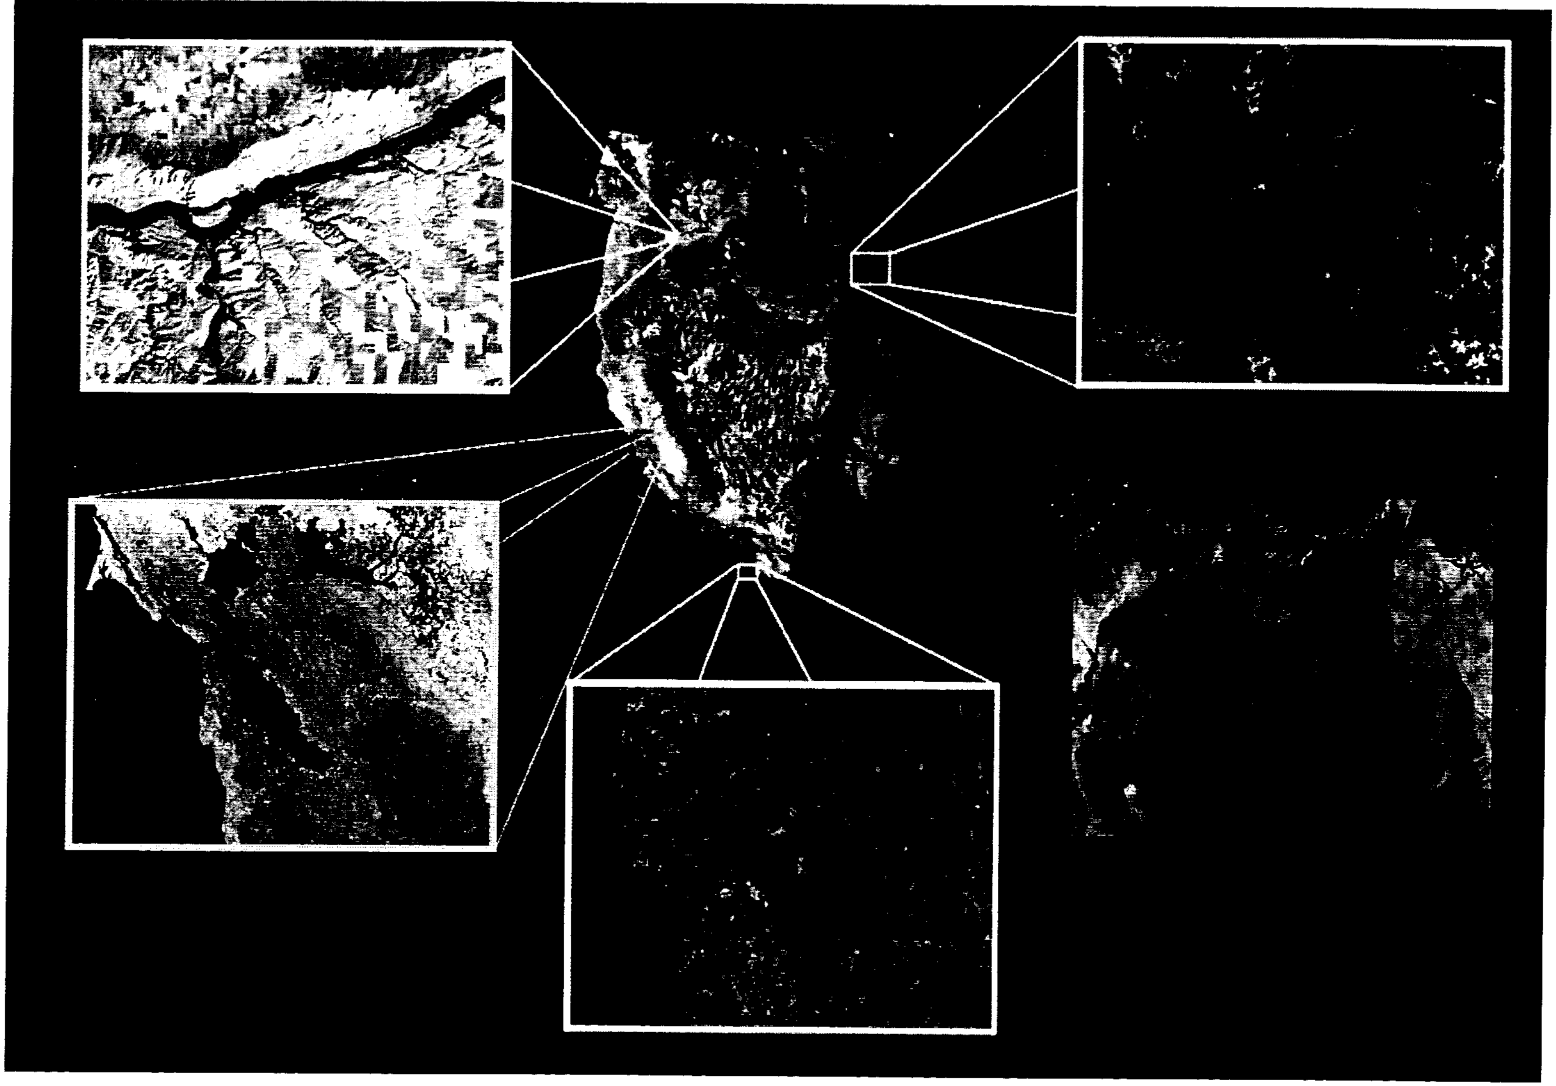

Source: Dave Curkendall / JPL

In Search of the Lost City of Ubar....

An Example in Data Fusion

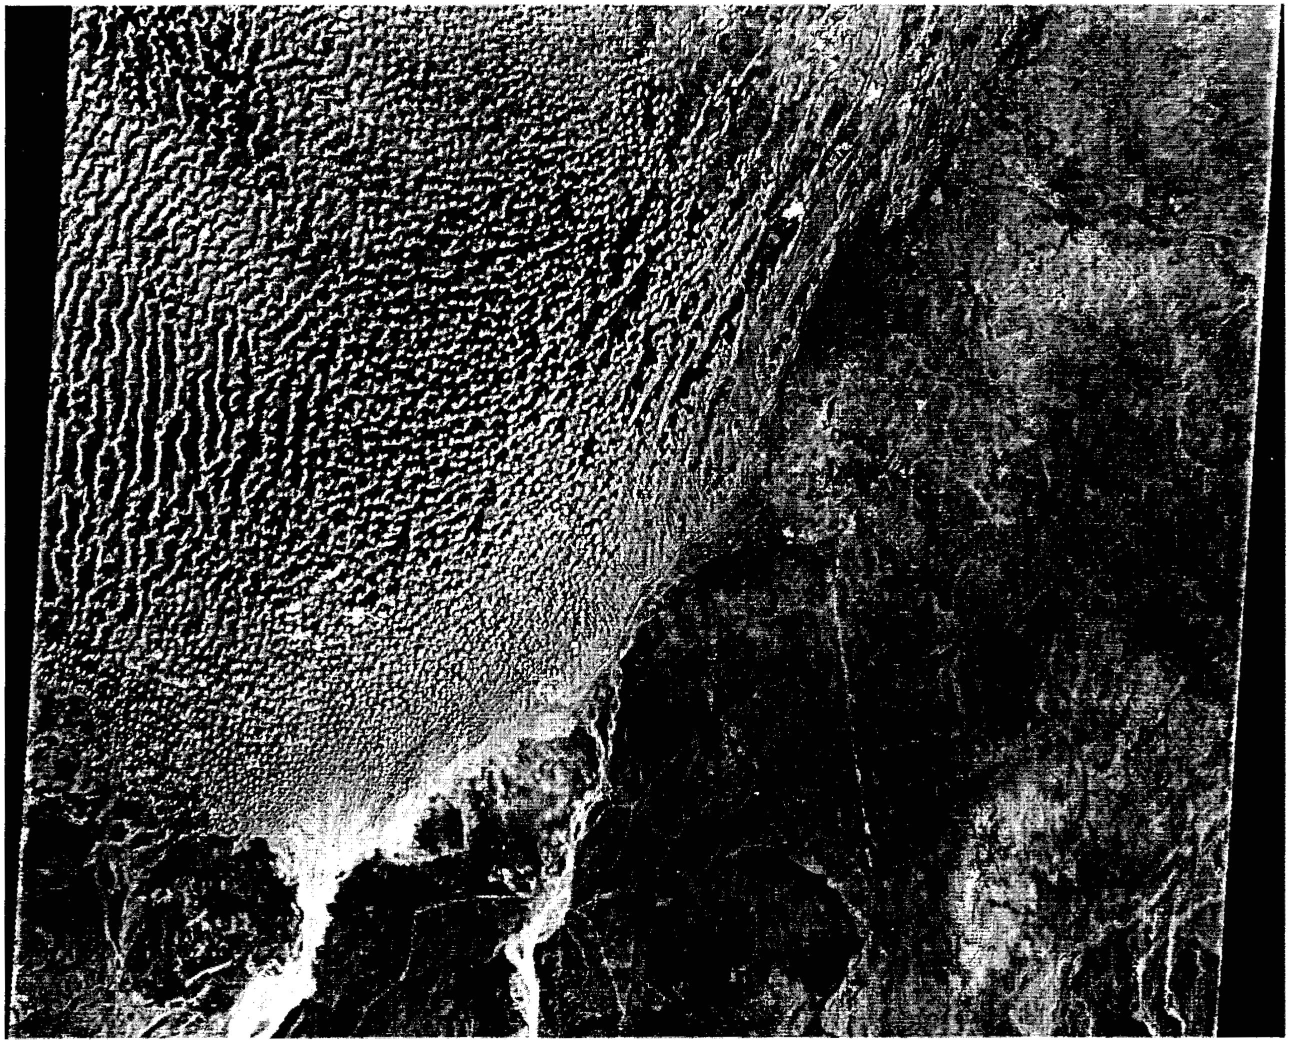

Source: Robert Crippen and Ron Blom / JPL

The Ubar Sinkhole 4X

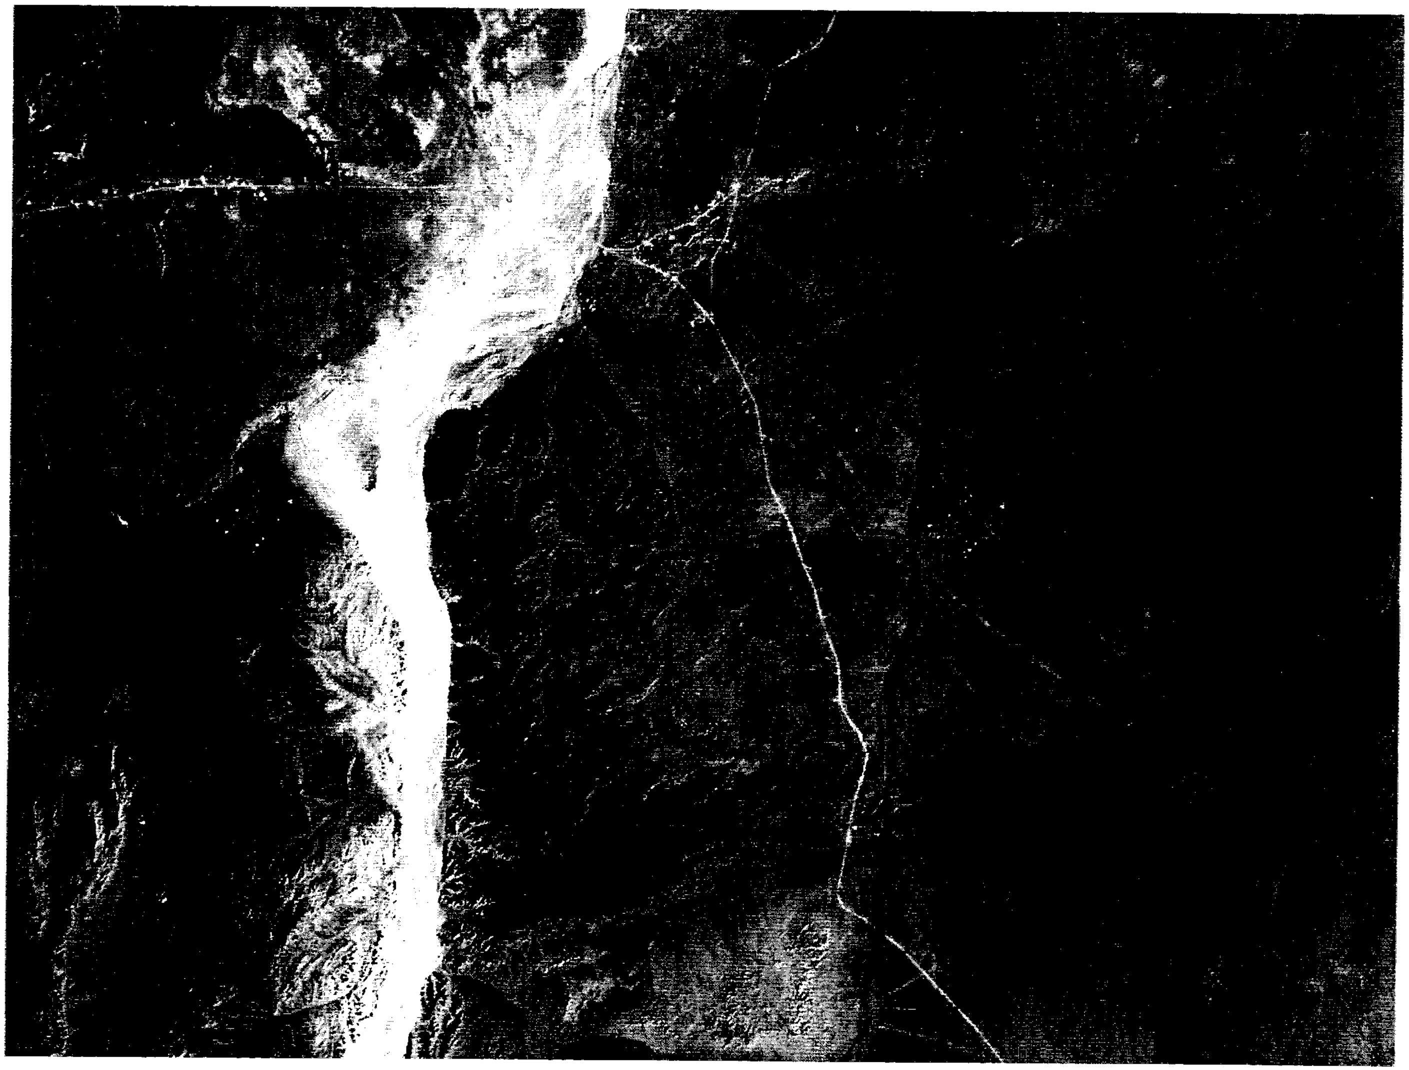

Source: Robert Crippen and Ron Blom / JPL

The Ubar Sinkhole 10X

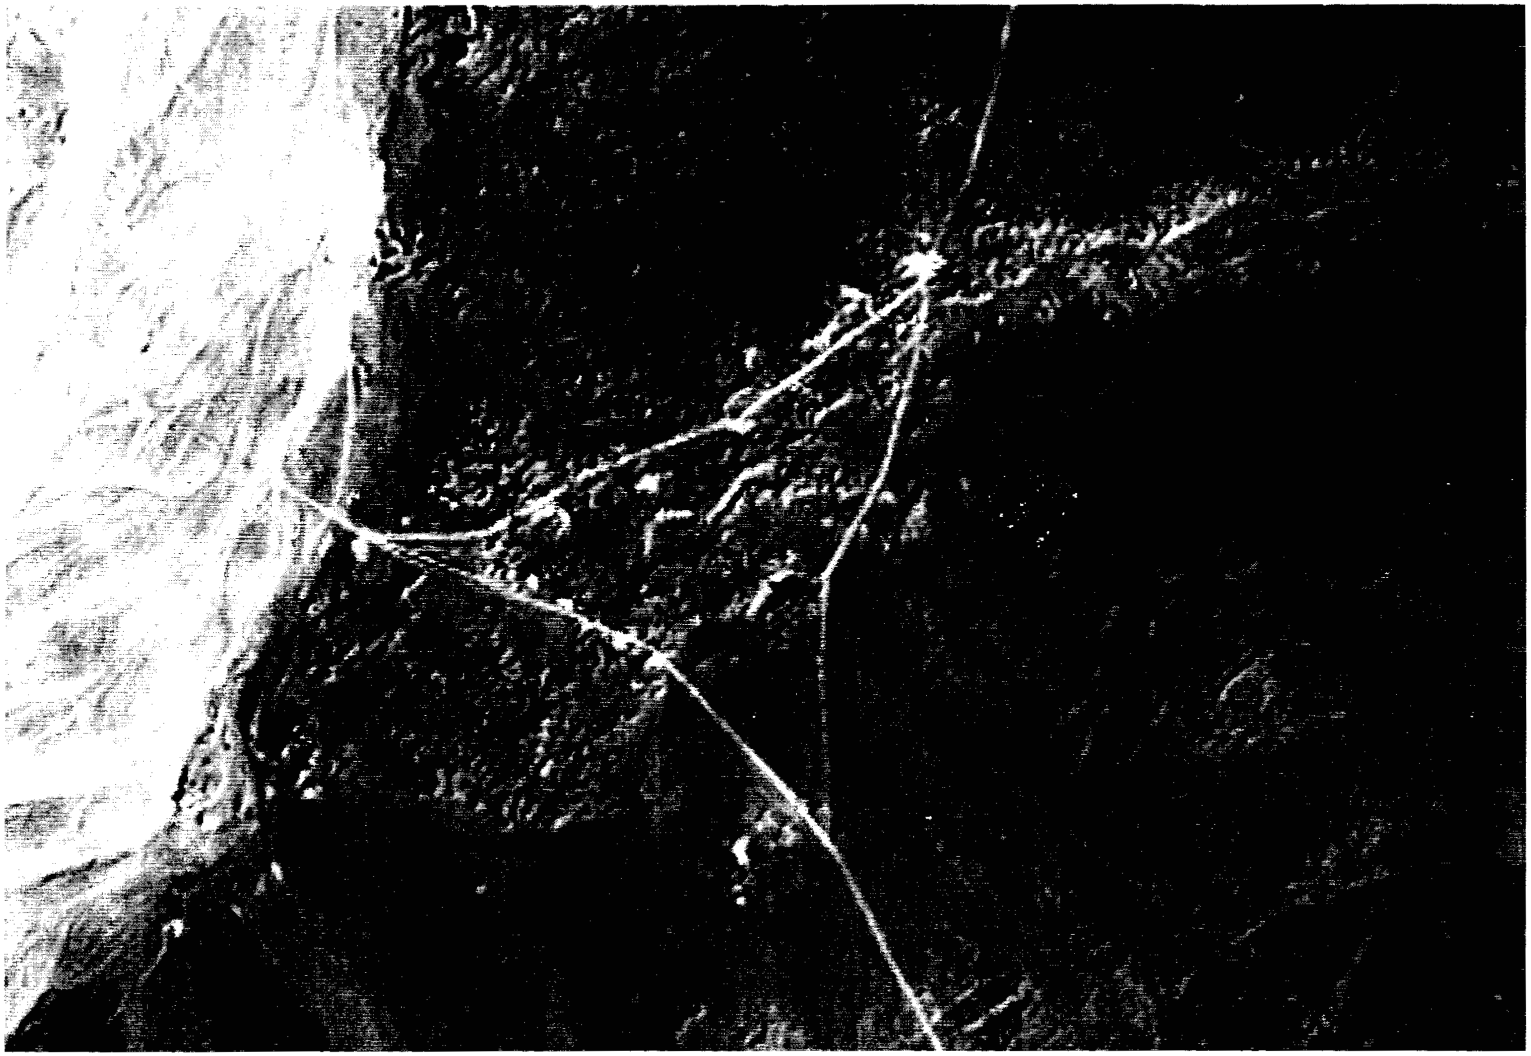

Source: Robert Crippen and Ron Blom / JPL

Sensor Networks

- Many examples in science
  - **Seismic, oceanic, astronomy, nuclear physics**
- Also, many emerging in commercial, civil, space, and military
  - **Weather**
  - **Traffic/road sensors**
  - **FAA aircraft control**
  - **Surveillance**
  - **Structural (aircraft, bridges, buildings, …)**
  - **Inter-Planetary Network (IPN)**
  - **Human**
  - **….and many others**

GBET

NEPTUNE

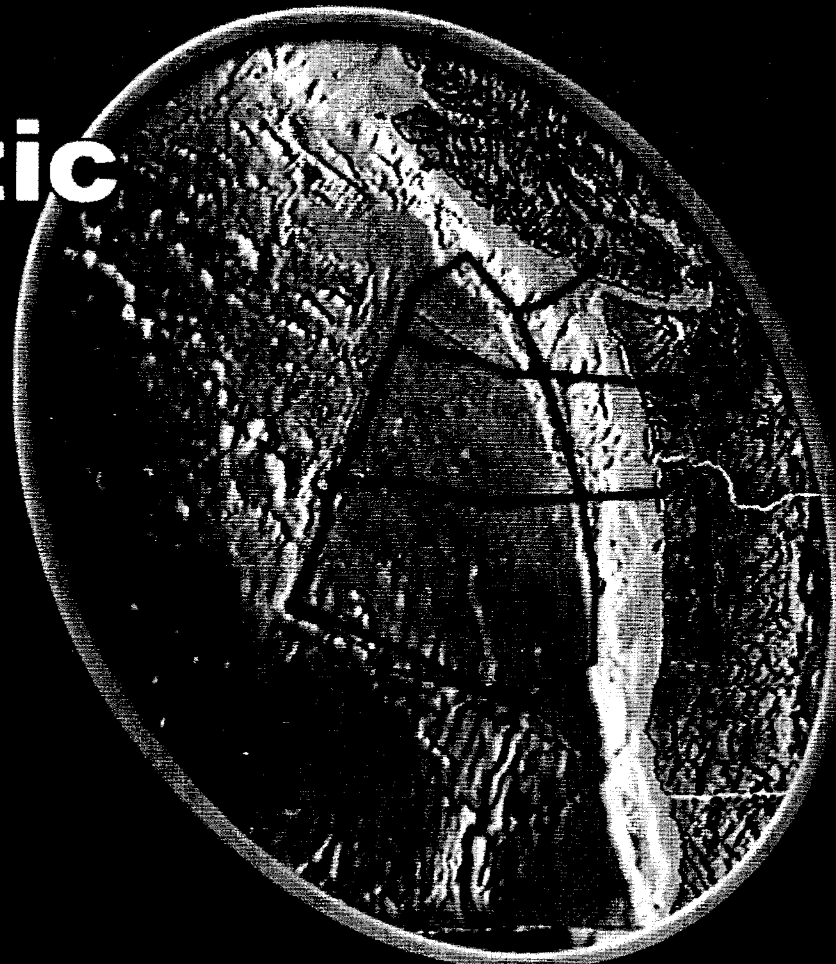

A Fiber Optic Telescope to Inner Space

www.ocean.washington.edu/neptune

Source: J. Delaney

CeV

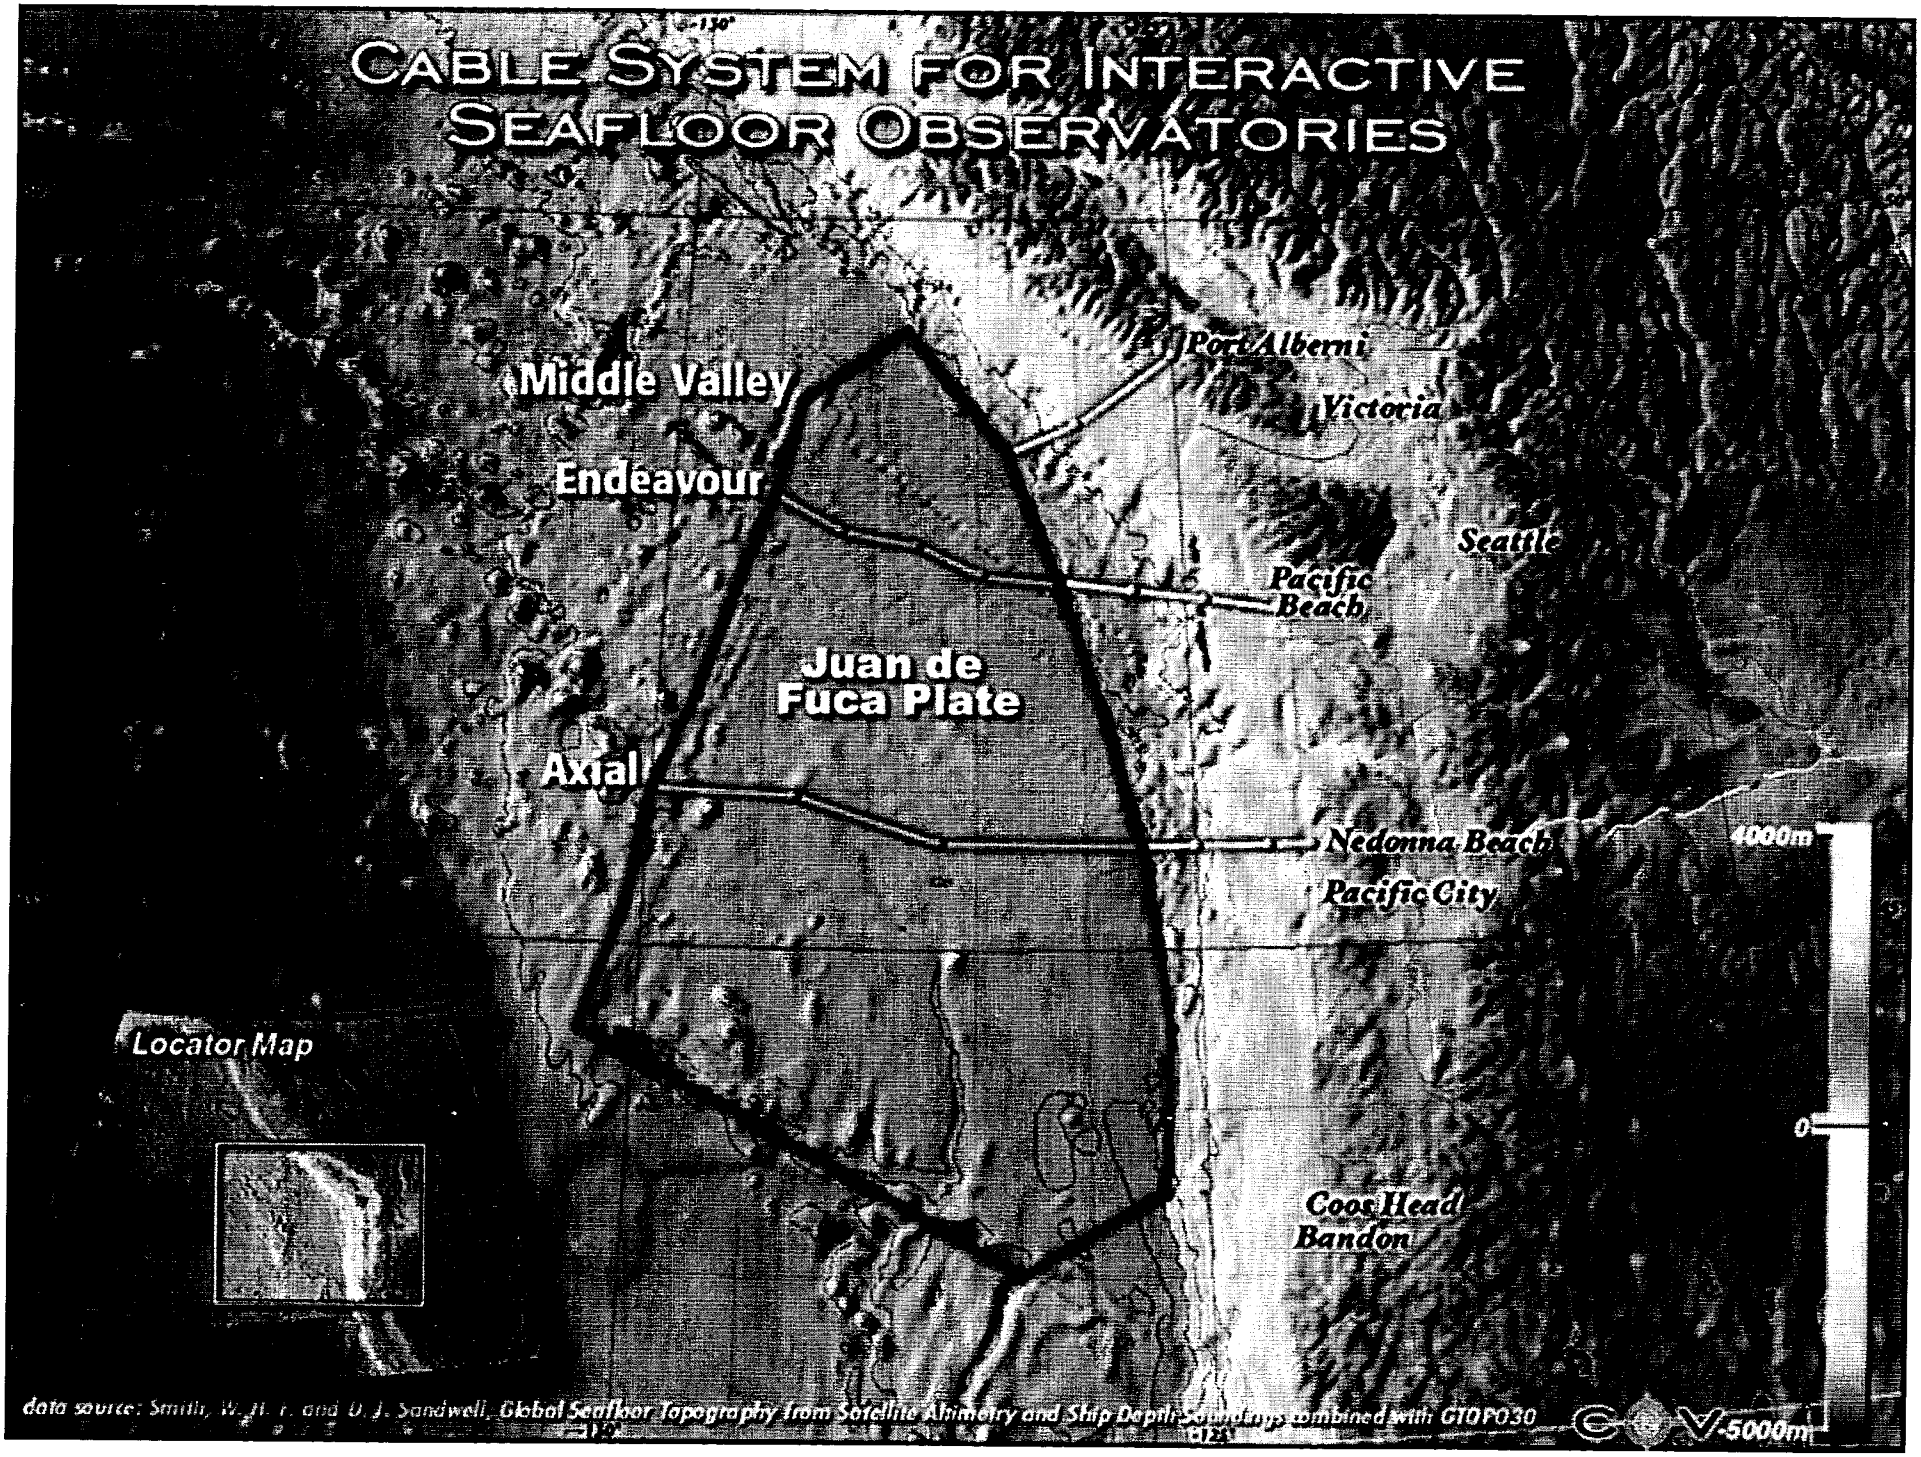

CABLE SYSTEM FOR INTERACTIVE SEAFLOOR OBSERVATORIES
Middle Valley
Endeavour
Juan de Fuca Plate
Axial
Port Alberni
Victoria
Seattle
Pacific Beach
Nedonna Beach
Pacific City
Coos Head
Bandon
Locator Map
4000m
0
-5000m
data source: Smith, W. H. F. and D. T. Sandwell, Global Seafloor Topography from Satellite Altimetry and Ship Depth Soundings combined with GTOPO30

NEPTUNE Overview

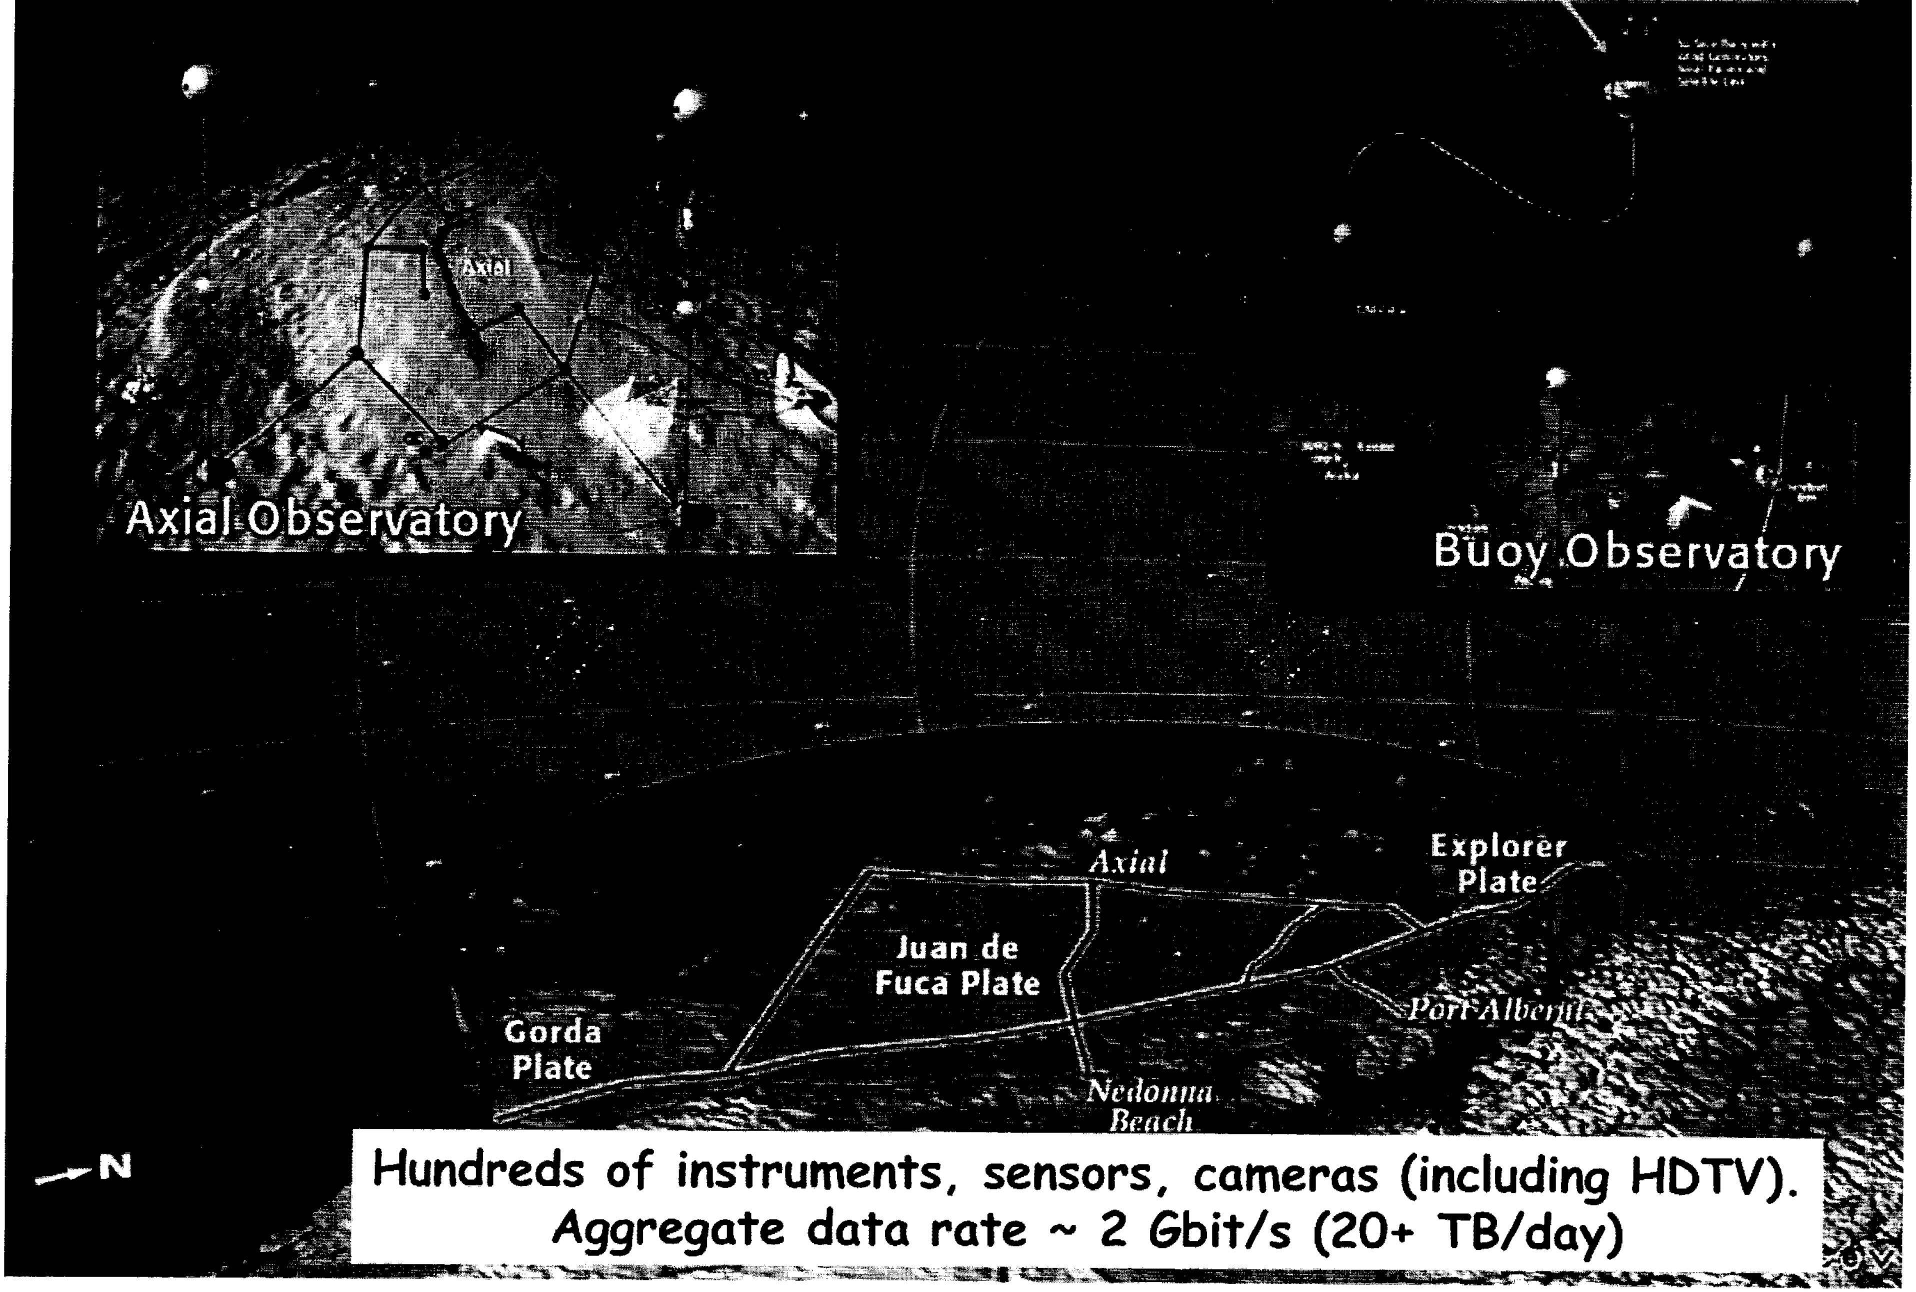

Hundreds of instruments, sensors, cameras (including HDTV).
Aggregate data rate ~ 2 Gbit/s (20+ TB/day)

National Virtual Observatory (NVO)

- The Astronomical Community is mounting a strong push with Code to deploy a Grid Computing approach to revolutionize astronomy.

- It is characterized by terascale distributed archives, linked by high speed networks and desktop access.

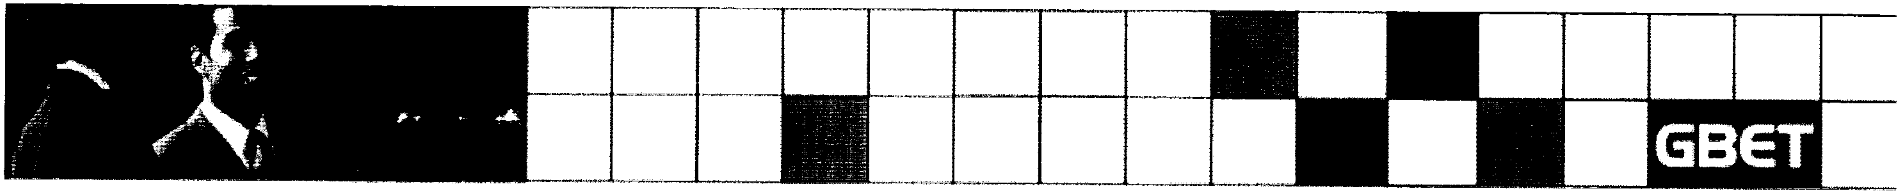

Remote Astronomy on Keck Telescope via ACTS

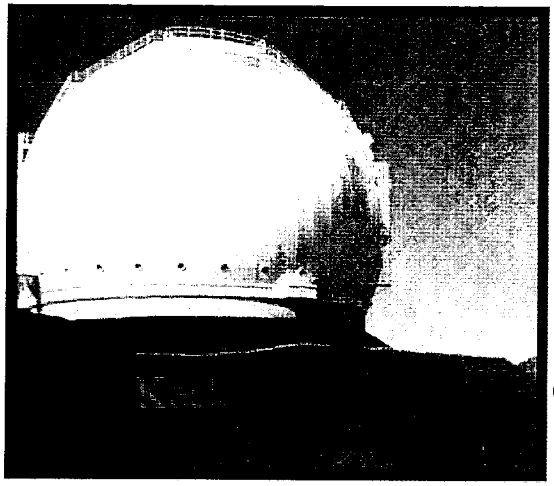

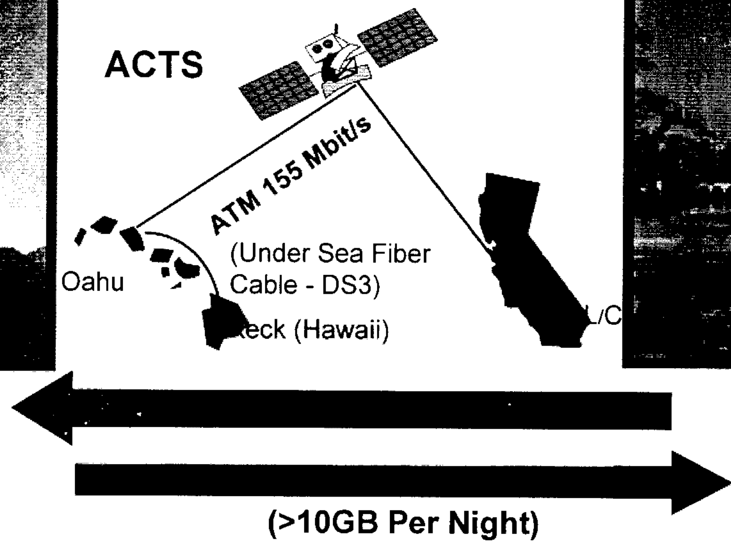

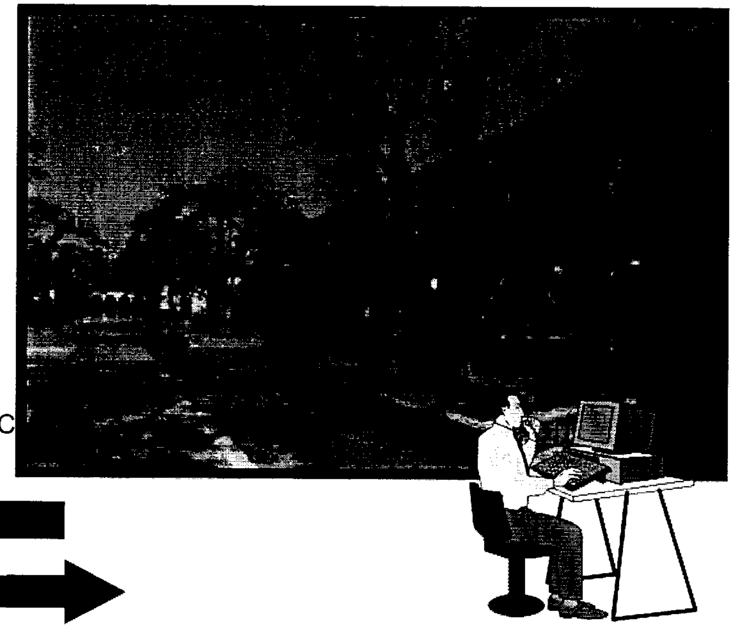

OBJECTIVES: Validate use of high latency gigabit satellites and network protocol for remote astronomy.

MOTIVATION: High altitude (14,000 ft), long travel time points the need for remote astronomy. High bandwidth satellite communications can reach remote locations where optical fiber is unavailable

DATA MOVEMENT: LRIS (4Kx4Kx16bit) instrument operated by Caltech Sun Workstation over 155 Mbit/s satellite and fiber optic link

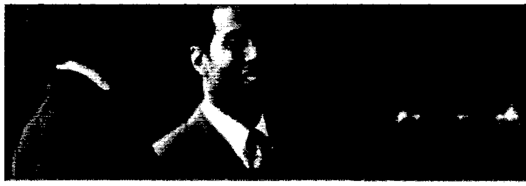

Digital Sky Virtual Observatory

- Multi-layer
- Multi-screen "Powerwall"
- Pan/Zoom; Image Processing
- Terabyte level Viewing Capacity
- Images and Catalog Data
- Plotting (2D, 3D)
- Annotations; Mark Objects
- Clipboard

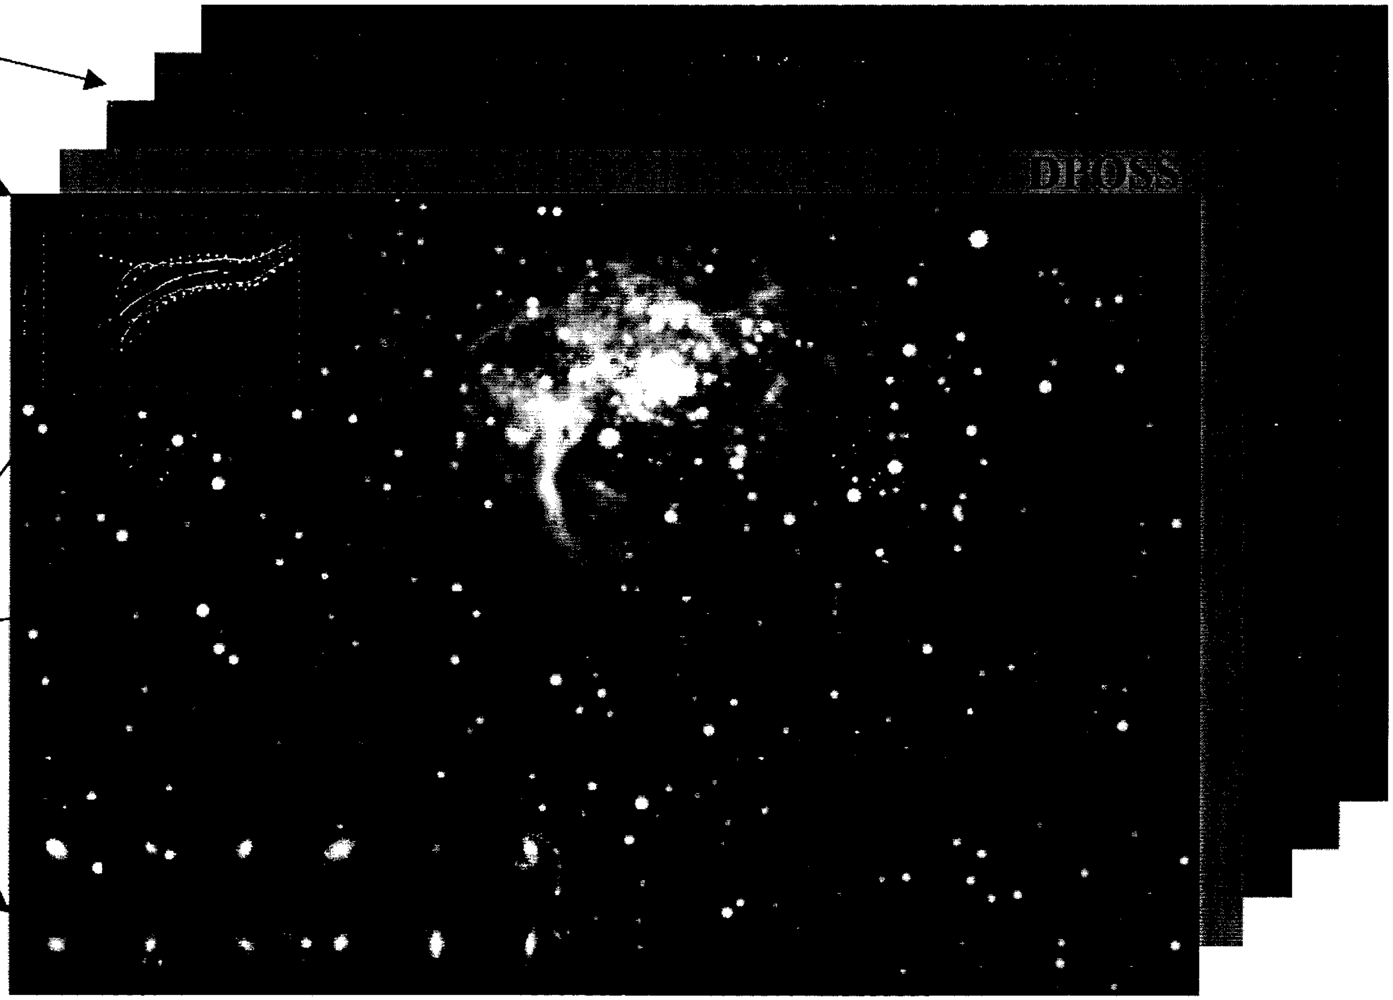

Collaborative Environment Vision

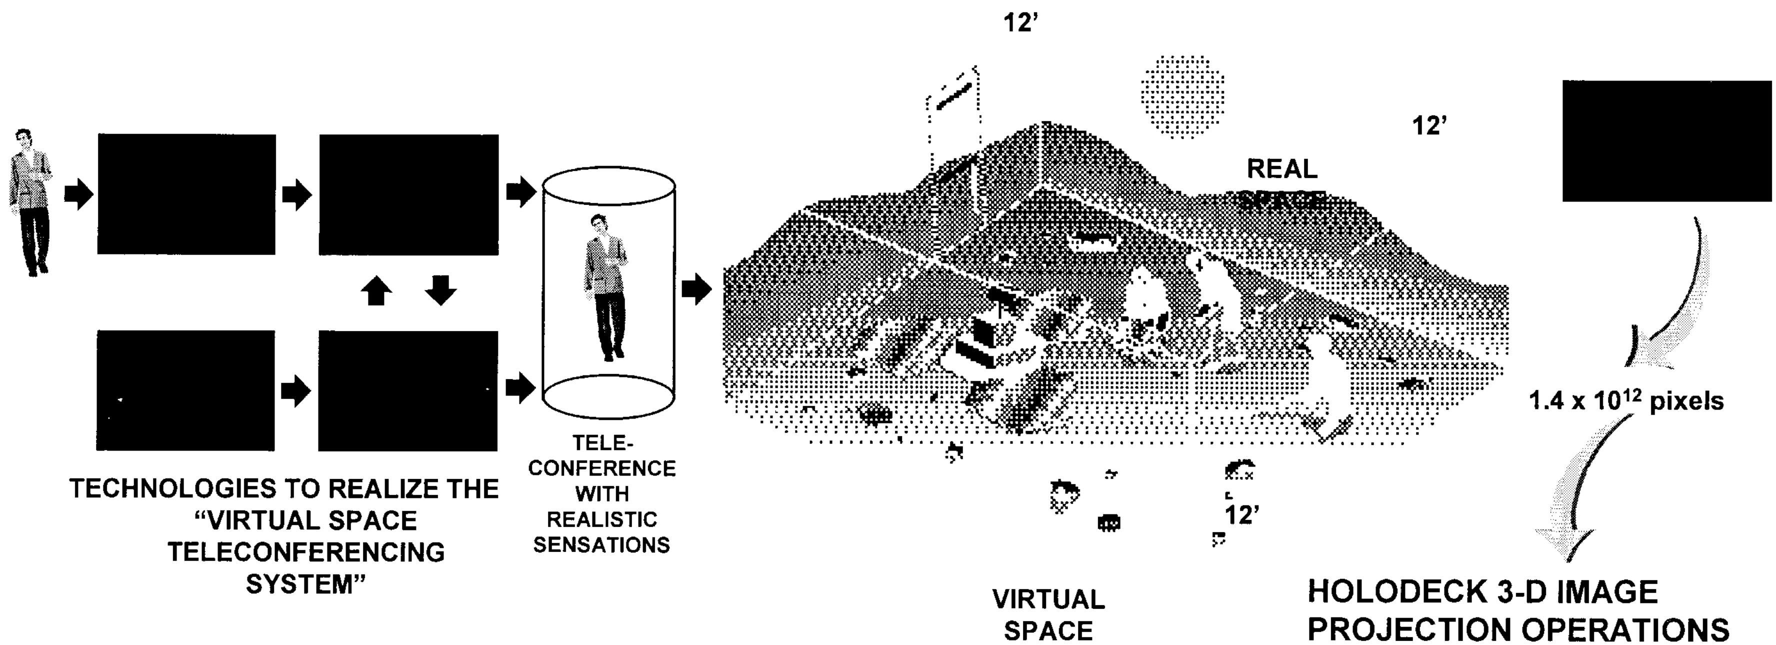

HOLODECK 3-D IMAGE PROJECTION OPERATIONS

= n log n x 100L x 10V x 60 frames/sec

$\cong 10^{18}$ operations/sec

- NO HEAD GEAR FOR 3-D VISUALIZATION/TRACKING
- PHOTOGRAPHIC PROJECTION QUALITY
- VIRTUAL SPACE TELECONFERENCE INTERFACE
- VIRTUAL PRESENCE FOR MANKIND ON OTHER PLANETS

Mission Vision Dome

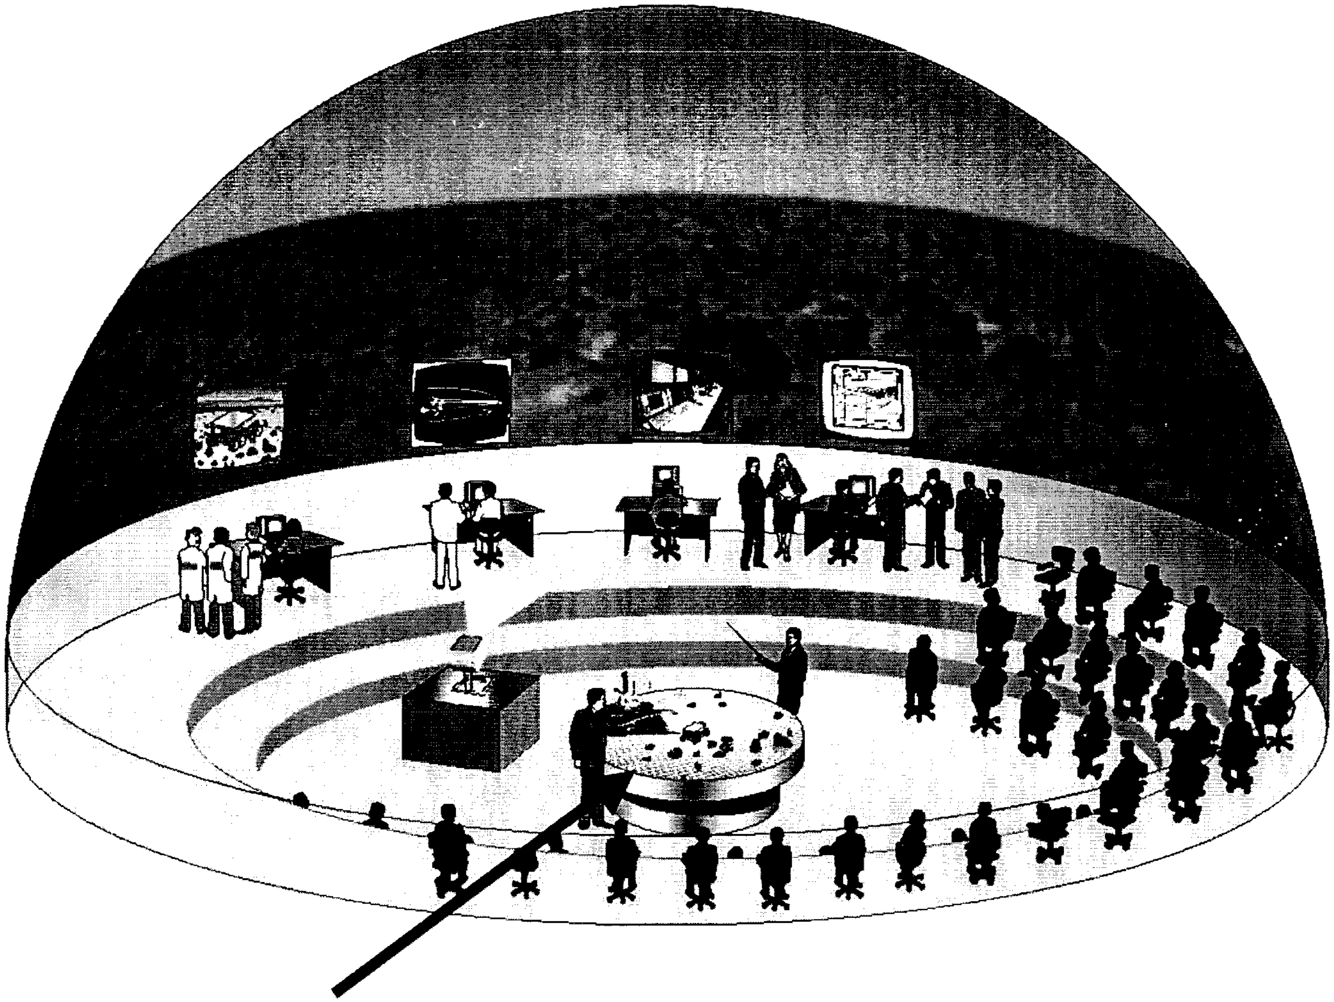

3-D Holographic Display Table

Source: Ed Chow / JPL

Multiplexed Holographic Projection Table

- Physical Optics Corporation (POC) is working with JPL in building a holographic display table based on the POC's holographic diffuser screen technology.
- Advantages include high resolution, multiple viewers/viewing angles, no mechanical parts, can improve with the projection technology, etc.

- In order to allow everyone to see 3-D objects at the same time, a special stereo projector is needed for each person standing in front of the table.

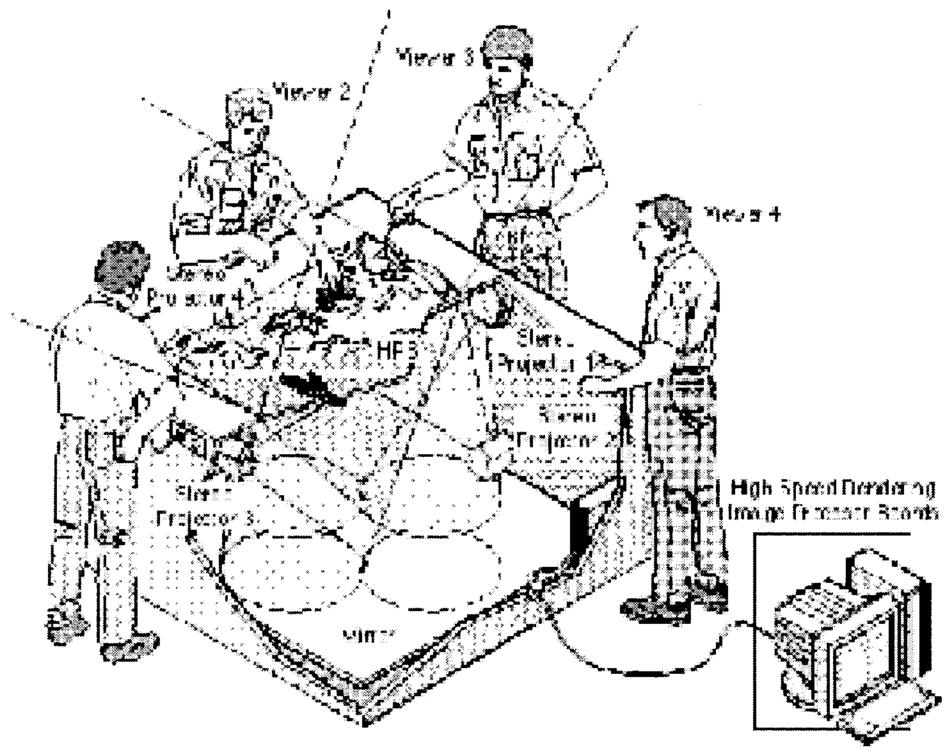

Ed Chow / JPL

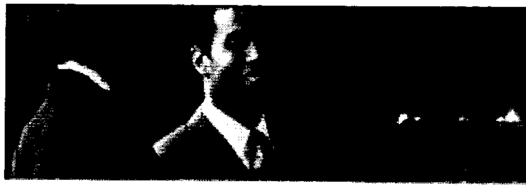

Virtual Environments Laboratory

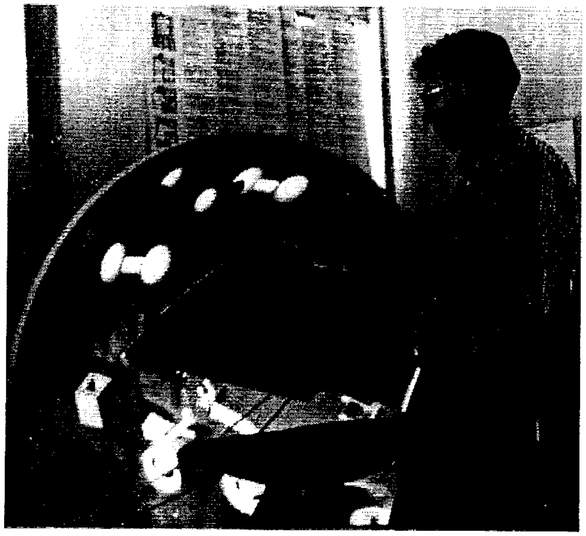

- *Gesture-based interaction*
- *Planning through simulation*

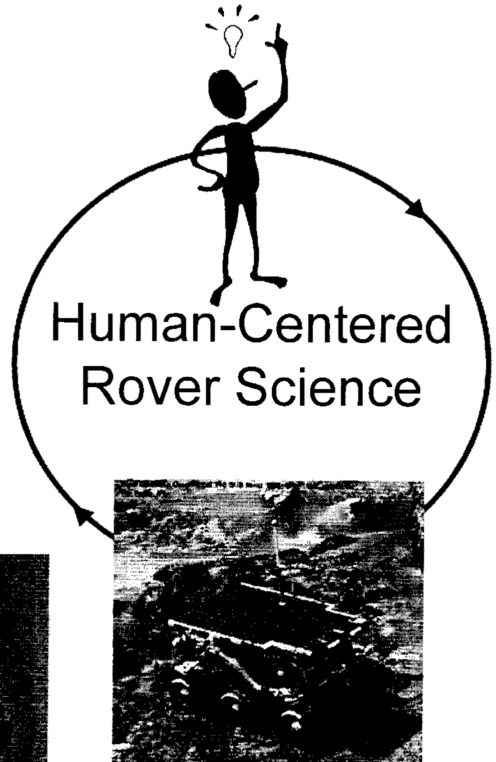

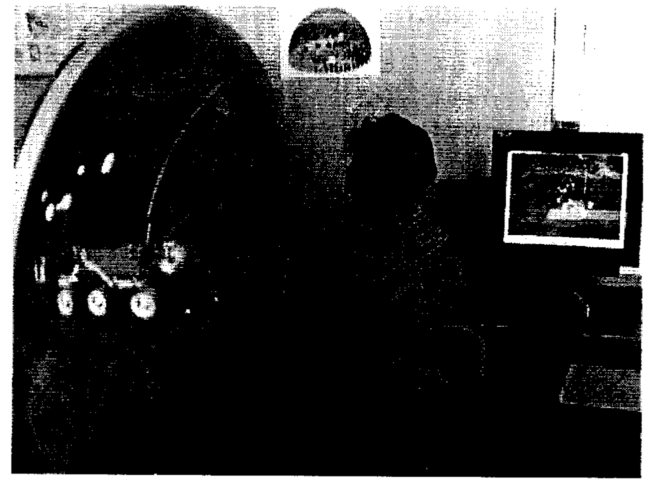

- *Supervised telemanipulation*
- *Haptic recording and display*

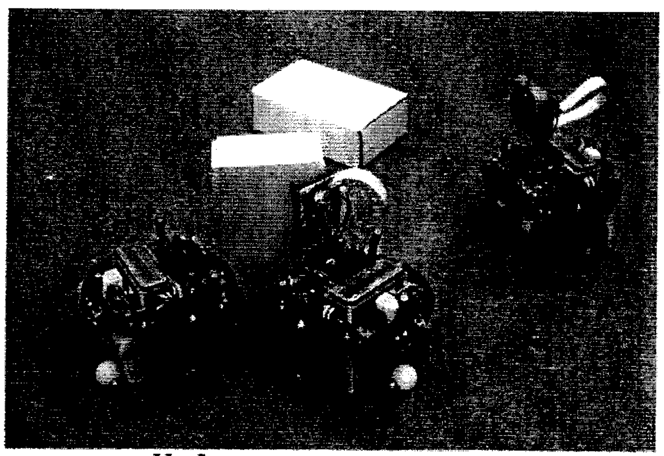

collaborative rovers

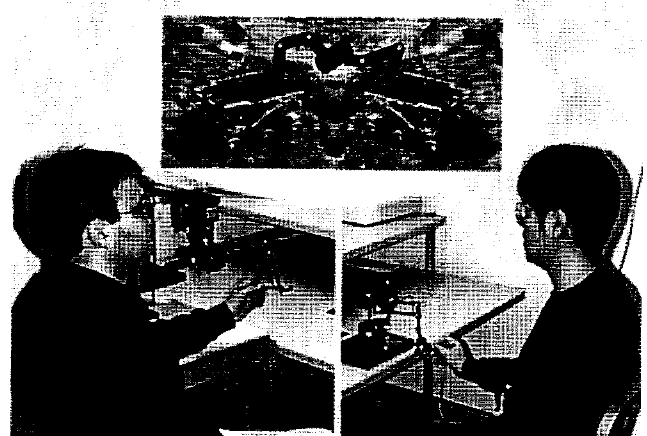

collaborative humans

3D Virtual Reality with Reduced Bandwidth

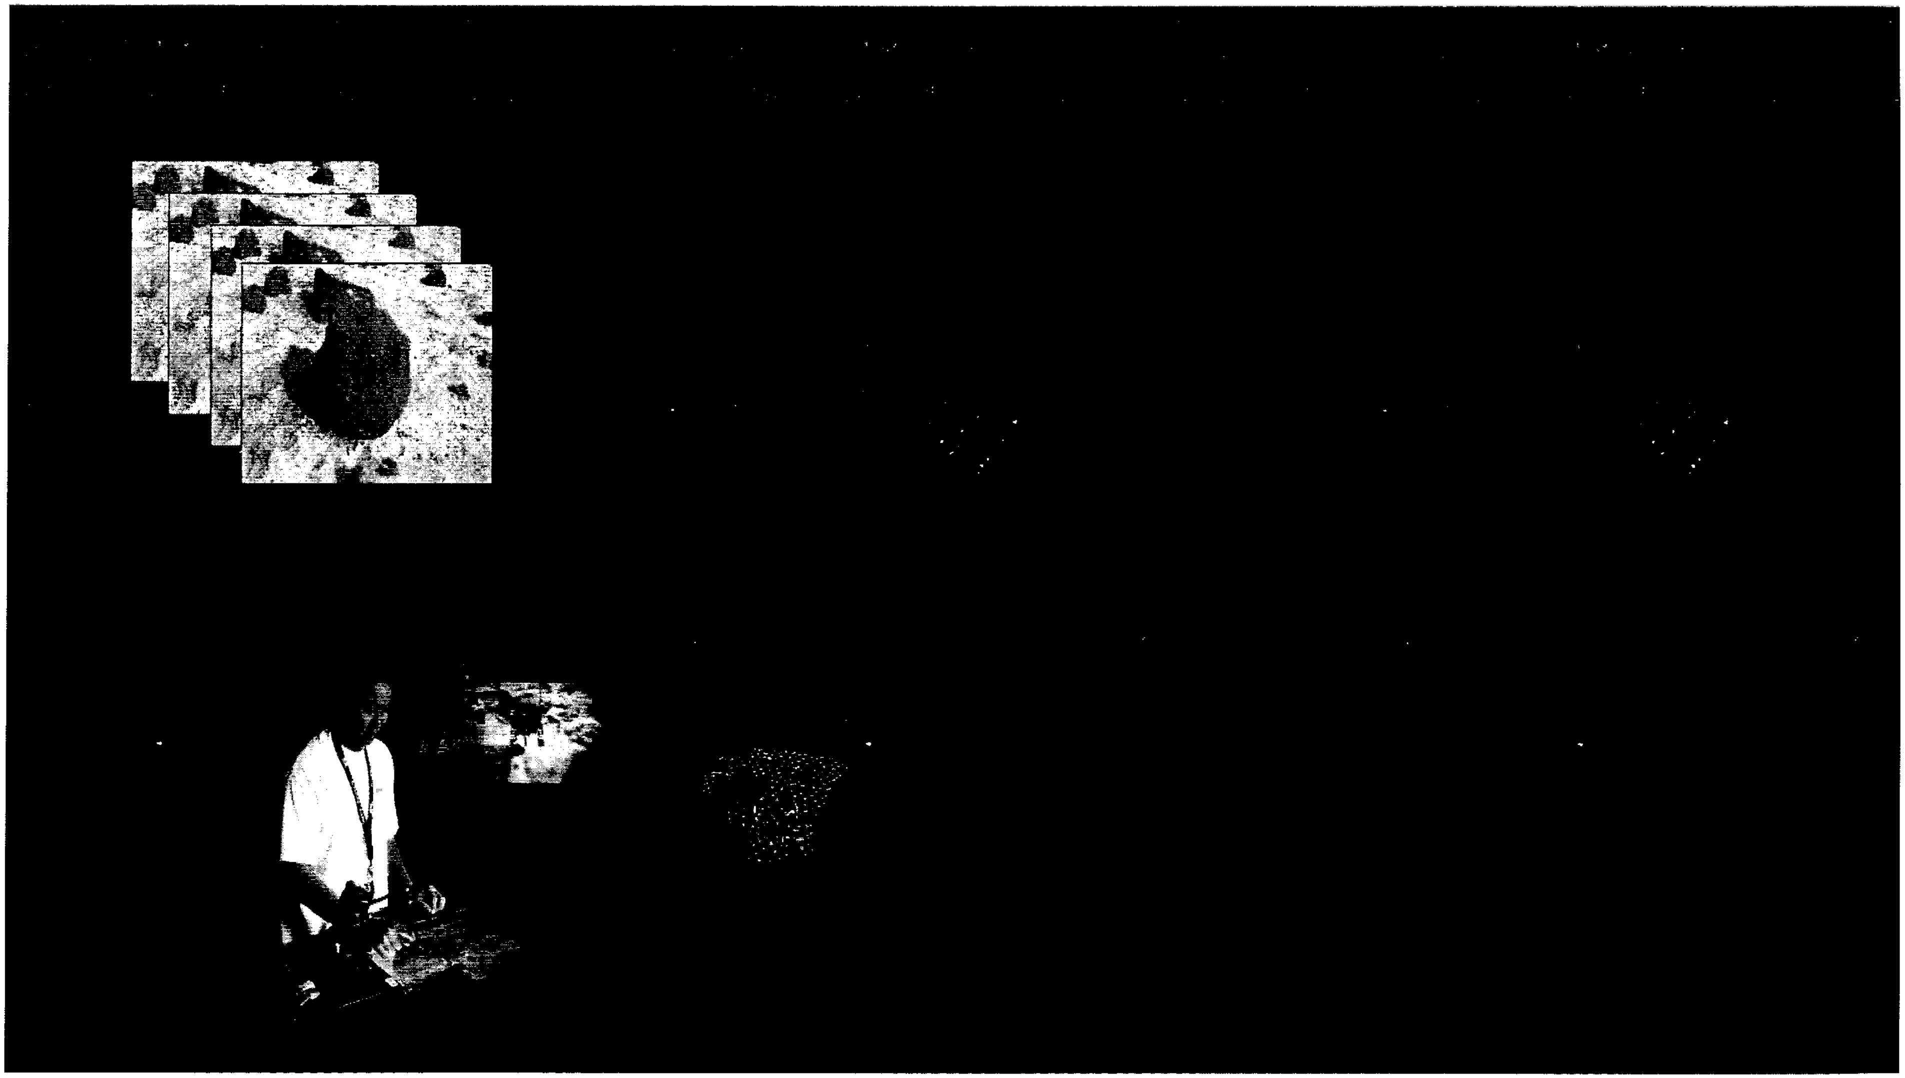

III.

INFRASTRUCTURE PROVISIONING STRATEGIES

Provisioning Strategies

- Streaming Applications
  - **Bandwidth reservation protocols (IPv.6) still coming**
  - **Most packet schemes will work by engineering sufficient over-capacity (3X and more peak to average load)**
  - **Trick is to estimate worst case peak loads on institutional trunks**
- Data Networks
  - **Keep in mind bandwidth asymmetry may exist between front end and back end of distributed applications**
  - **Can in some cases trade local processing (e.g., for additional compression) for network speed**
    - Depends on how expensive network resource is

Wrapup

- The next 20 years of the Internet will be much different from the last 20 years
- Expect more applications that use….
  - **three dimensional movie imagery, ultimately leading to VR**
  - **tapping immense sensor webs and data bases for information**
- Computer to computer communications will eventually dwarf human centric use of the web
- Real-time and non-real-time data services can co-exist with sufficient over provisioning of capacity (relatively inexpensive at the institutional level).

Q & A