NASA High Performance Computing (HPC) Directions, Issues, and Concerns: A User's Perspective

Dr. Robert C. Singleterry Jr.
NASA Langley Research Center

HPC China

Oct 29th, 2010

Overview



- Current Computational Resources
- Directions from a User's Perspective
- Issues and Concerns
- Conclusion?
- Case Study – Space Radiation
- Summary

Current Computational Resources

- Ames
  - 115,000+ cores (Pleiades)
  - 1-2 GB/core
  - LUSTRE
- Langley
  - 3000+ cores (K)
  - 1GB/core
  - LUSTRE
- Goddard
  - 10,000+ Nehalem cores (1 year ago)
  - 3GB/core
  - GPFS
- Others at other centers

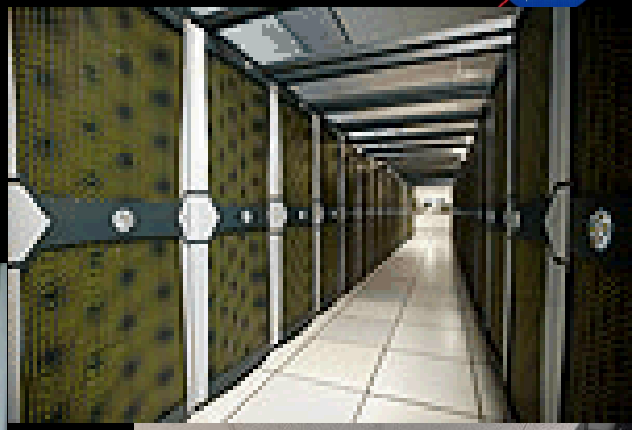





Current Computational Resources

- Science applications
  - Star and galaxy formation
  - Weather and climate modeling
- Engineering applications
  - CFD
    - Ares-I and Ares-V
    - Aircraft
    - Orion reentry
  - Space radiation
  - Structures
  - Materials
- Satellite operations, data analysis & storage

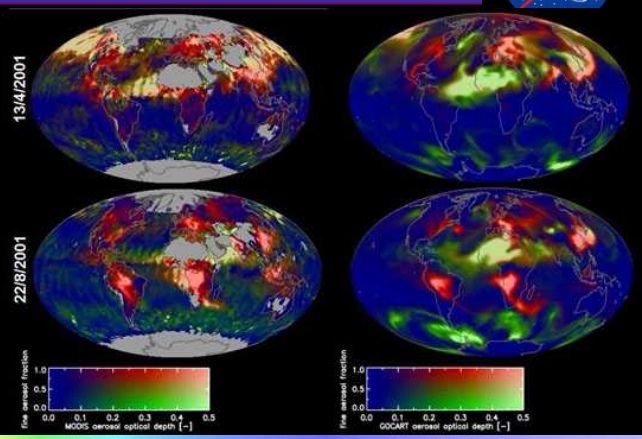

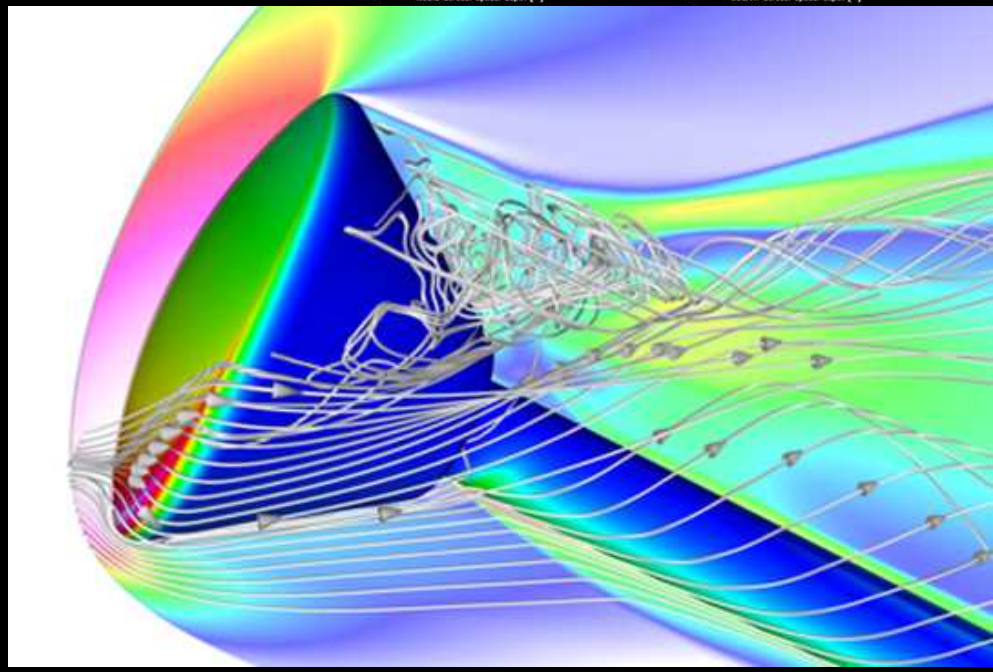

Directions from a User's Perspective

- 2004: Columbia
  - 10,240 cores
- 2008: Pleiades
  - 51,200 cores
- 2012 System
  - 256,000 cores
- 2016 System
  - 1,280,000 cores
- Extrapolation!!!
- Use at own risk

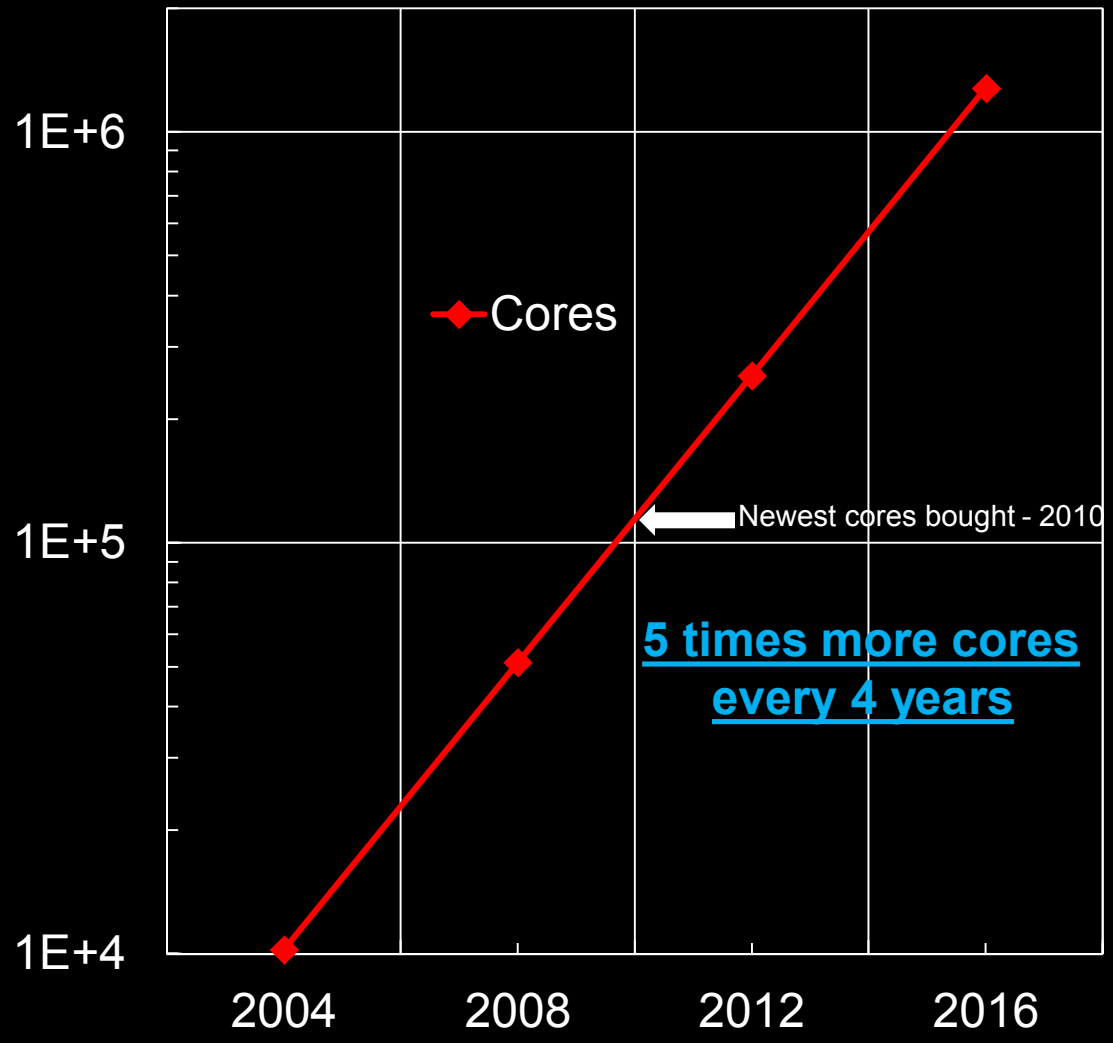

Issues and Concerns



- Assume power and cooling are not issues
  - Is this a valid assumption?
- What will a "core" be in the next 6 years?
  - "Nehalem"-like – powerful, fast, and "few"
  - "BlueGene"-like – minimal, slow, and "many"
  - "Cell"-like – not like CPU at all, fast, and many
  - "Unknown"-like – combination, hybrid, new, …
- In 2016, NASA should have a 1.28 million core machine tightly coupled together
- Everything seems to be fine Maybe???

Issues and Concerns?



- A few details about our systems
  - Each of the 4 NASA Mission Directorates "own" part of Pleiades
  - Each Center and Branch resource control their own machines in the manner they see fit
  - Queues limit the number of cores used per job per Directorate, Center, or Branch
  - Queues limit the time per job without special permissions from the Directorate, Center, or Branch
- This harkens of a time share machine of old

Issues and Concerns?



- As machines get bigger, 1.28 million cores in 2016, do the queues get bigger?
- Can the NASA research, engineer, and operation users utilize the bigger queues?
- Will NASA algorithms keep up with the 5 times scaling every 4 years?
  - 2008: 2000 core algorithms
  - 2016: 50,000 core algorithms
- Is NASA spending money on right issue?
  - Newer, bigger, better hardware
  - Newer, better, scalable algorithms

Conclusions?



- Is there a conclusion?
- There are issues and concerns!
  - Spend money on bigger and better hardware?
  - Spend money on more scalable algorithms?
- Do the NASA funders understand these issues from a researcher, engineer, and operations point of view?
- Do researchers and engineers understand the NASA funder point of view?
- At this point, there is no conclusion!

Case Study – Space Radiation

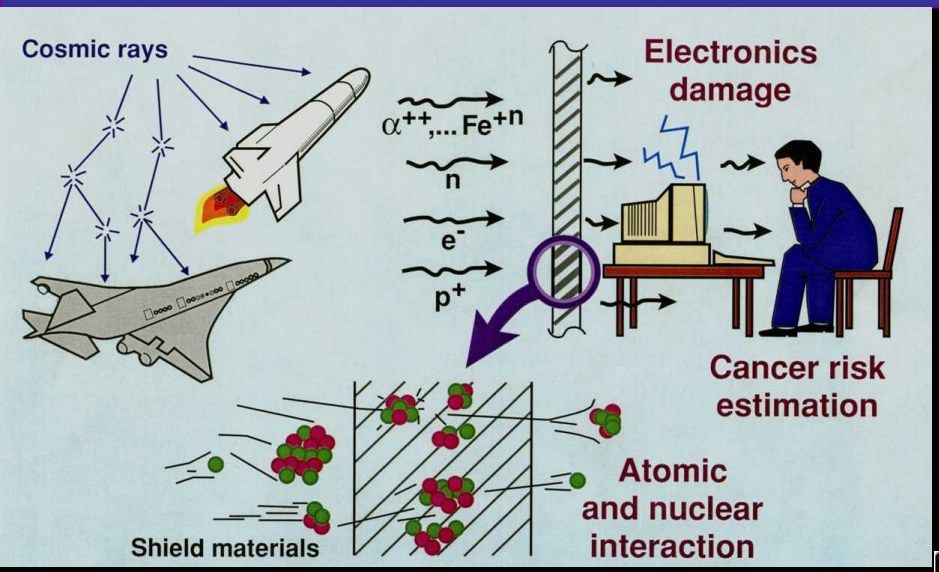

- Cosmic Rays and Solar Particle Events
- Nuclear interactions
- Human and electronic damage

- Dose Equivalent: damage caused by energy deposited along the particle's track

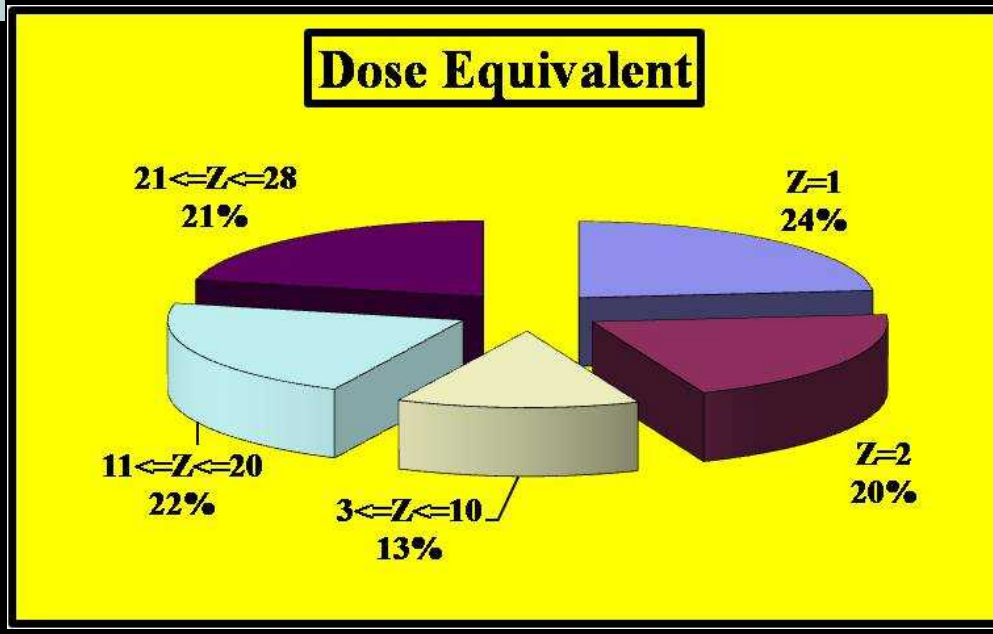

Previous Space Radiation Algorithm



- Design and start to build spacecraft
  - Mass limits and objectives have been reached
- Brought in radiation experts
- Analyzed spacecraft by hand (not parallel)
- Extra shielding needed for certain areas of the spacecraft or extra component capacity
- Reduced new mass to mass limits by lowering the objectives of the mission
  - Throwing off science experiments
  - Reducing mission capability

Previous Space Radiation Algorithm

- Major missions impacted in this manner
  - Viking
  - Gemini
  - Apollo
  - Mariner
  - Voyager



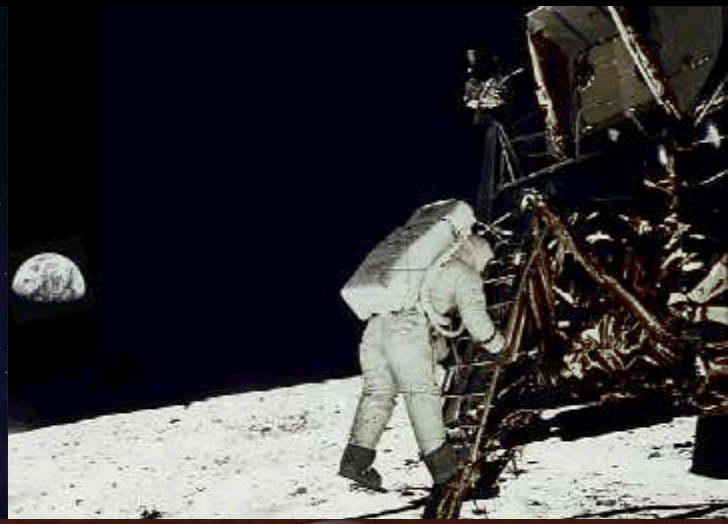





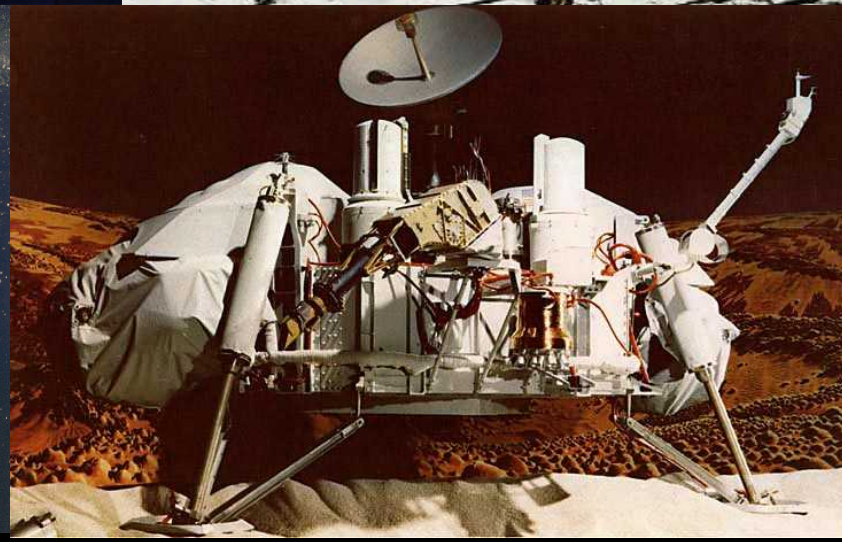

Previous Space Radiation Algorithm

SAGE III

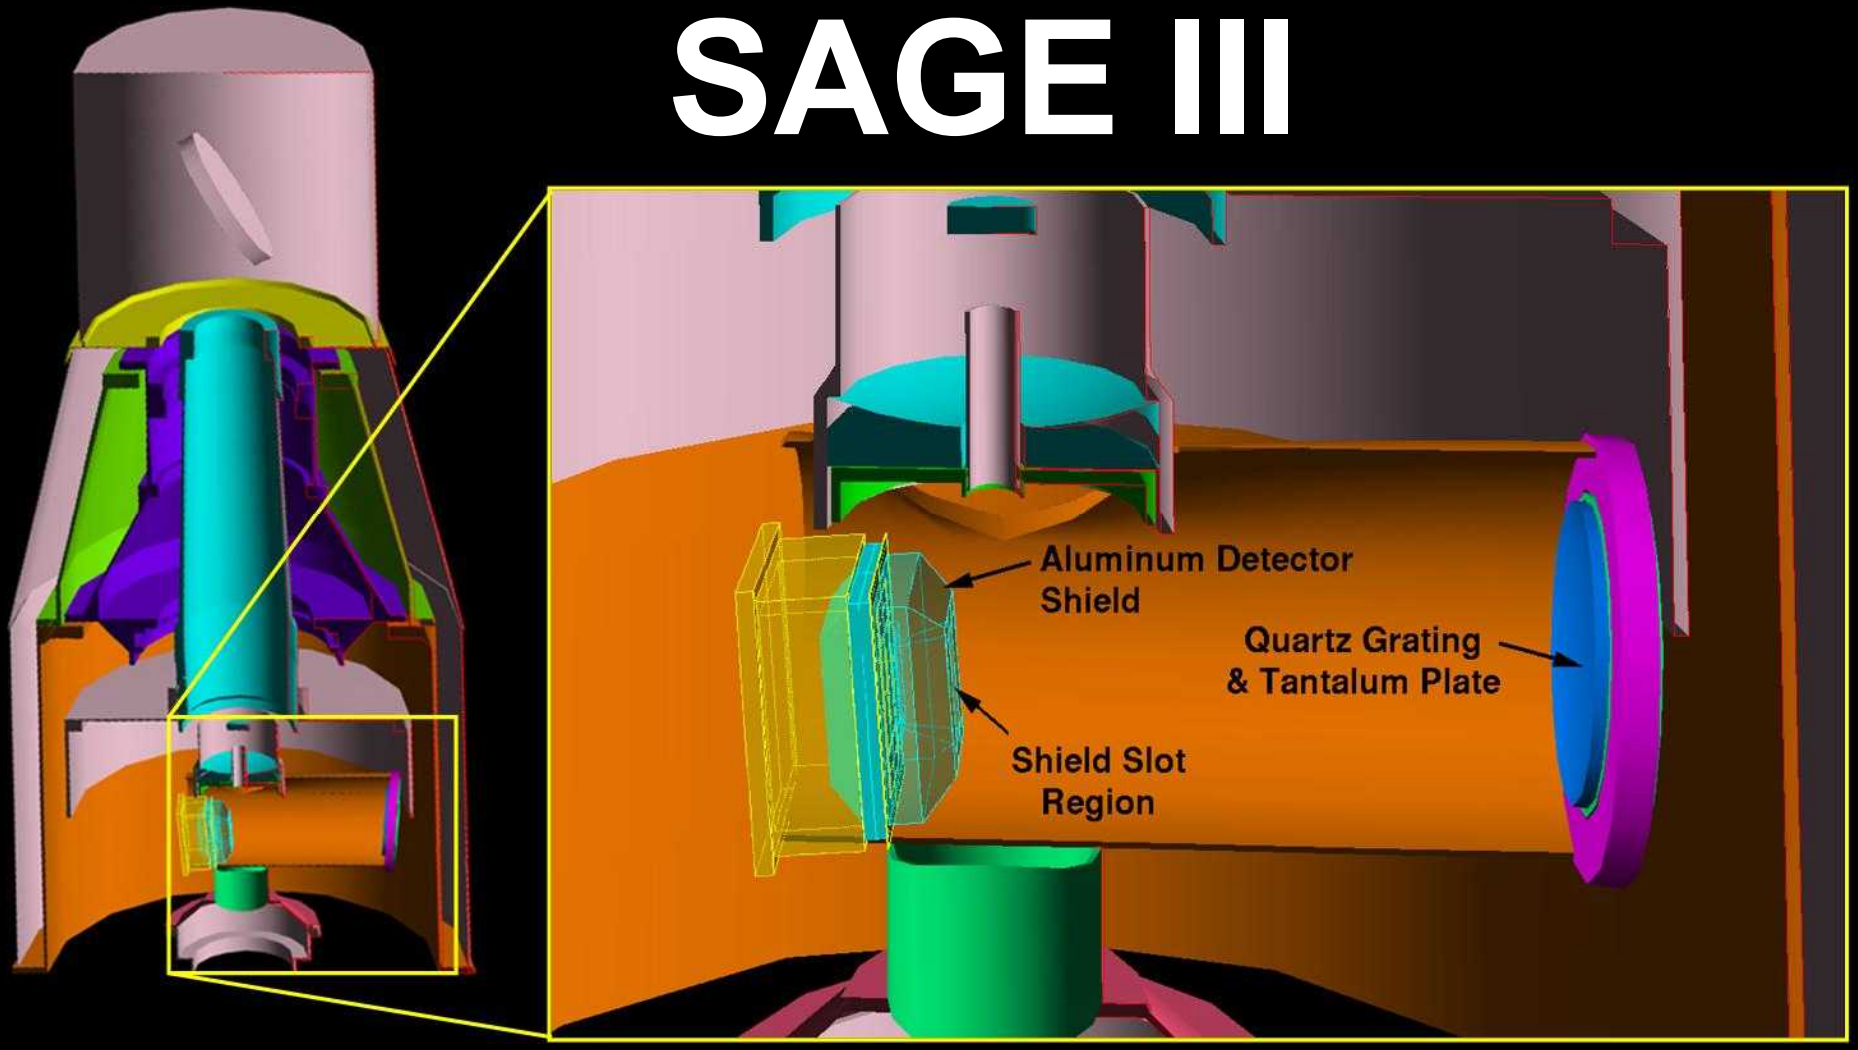

Primary Space Radiation Algorithm

- Ray trace of spacecraft/human geometry
- Reduction of ray trace materials to three ordered materials
  - Aluminum
  - Polyethylene
  - Tissue
- Transport database
- Interpolate each ray
- Integrate each point
- Do for all points in the body - weighted sum

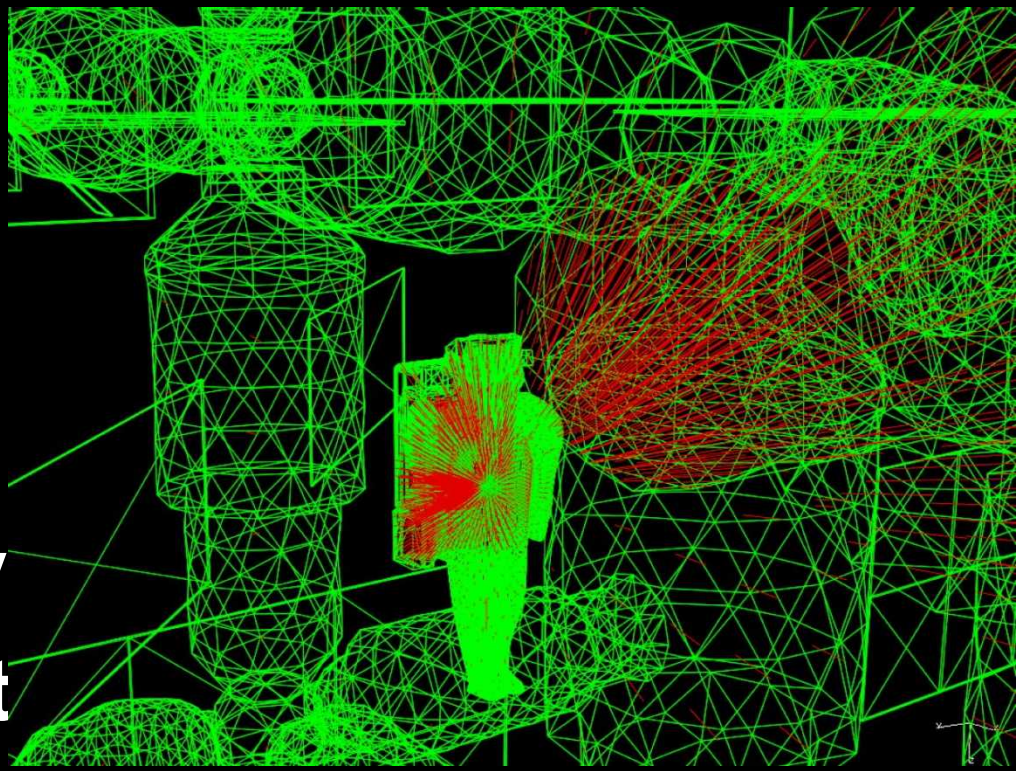

Primary Space Radiation Algorithm



- Transport database creation is mostly serial and not parallelizable in coarse grain
- 1,000 point interpolation over database is parallel in the coarse grain
- Integration of data at points is parallel if the right library routines are bought
- At most, a hundreds-of-core process over hours of computer time
- Not a good fit for the design cycle
- Not a good fit for the HPC of 2012 and 2016

Imminent Space Radiation Algorithm

- Ray trace of spacecraft/human geometry
- Run transport algorithm along each ray
  - No approximation on materials
- Integrate all rays
- Do for all points
- Weighted sum

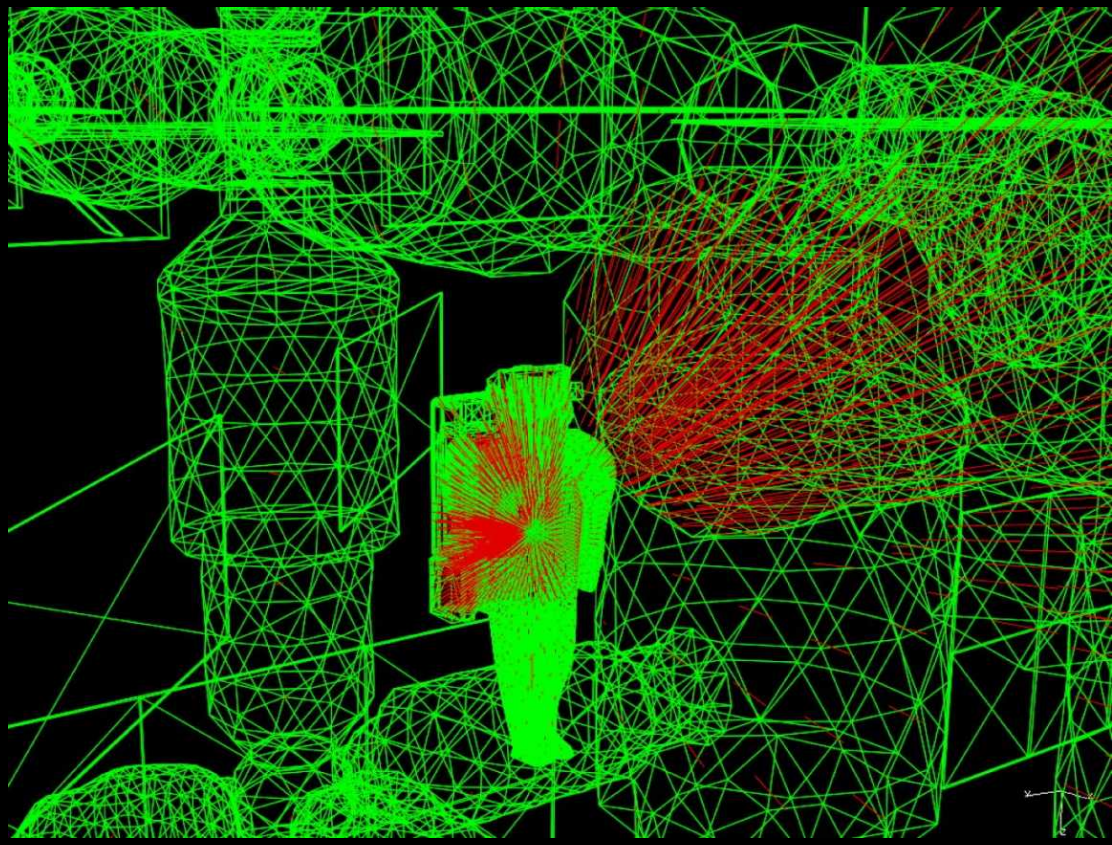

Imminent Space Radiation Algorithm



- 1,000 rays per point
- 1,000 points per body
- 1,000,000 transport runs @ 1 min to 10 hours per point (depends on rays)
- Integration of data at points is bottleneck
  - Data movement speed is key
  - Data size is small
- This process is inherently parallel if communication bottleneck is reasonable
- Better fit for HPC of 2012 and 2016

Future Space Radiation Algorithms



- Monte Carlo methods
  - Data communications is bottleneck
  - Each history is independent of other histories
- Forward/Adjoint finite element methods
  - Same problems as other finite element codes
  - Phase space decomposition is key
- Hybrid methods
  - Finite Element and Monte Carlo together
  - Best of both worlds (on paper anyway)
- Variational methods
  - Unknown at this time

Summary



- Present space radiation methods are not HPC friendly or scalable
  - Why care? Are the algorithms good enough?
  - Need scalability to
    - Keep up with design cycle wanted by users
    - Slower speeds of the many core chips
    - New bells & whistles wanted by funders
- Imminent method better but has problems
- Future methods show HPC scalability promise on paper but need resources for investigation and implementation

Summary



- NASA is committed to HPC for science, engineering, and operations
- Issues & concerns about where resources are spent & how they impact NASA's work
  - Will machines be bought that can benefit science, engineering, and operations?
  - Will resources be spent on algorithms that can utilize the machines bought?
- HPC help desk creation to inform and work with users to achieve better results for NASA work: HeCTOR Model