LA-UR-11-02817

Approved for public release;
distribution is unlimited.

Title: MCNP Monte Carlo & Parallel Computing

Author(s): Forrest B. Brown

Intended for: UNM Workshop on Monte Carlo for Particle Therapy Treatment Planning, University of New Mexico, Albuquerque NM, May 16-18, 2011

Los Alamos
NATIONAL LABORATORY
EST. 1943

University of New Mexico Workshop on Monte Carlo for
Particle Therapy Treatment Planning, 16-18 May 2011

LA-UR-11-02817

MCNP Monte Carlo & Parallel Computing

Forrest Brown

Monte Carlo Codes, XCP-3
Los Alamos National Laboratory

MCNP Monte Carlo & Parallel Computing

Forrest Brown, Monte Carlo Codes, LANL

MCNP is a general purpose Monte Carlo particle transport code developed at Los Alamos National Laboratory over the past 30+ years. The most recent production versions, MCNP5 and MCNPX, have been merged into MCNP6. MCNP6 provides very general capabilities for modeling geometry, defining particle sources, tallying a wide variety of physical phenomena, high fidelity representation of collision physics, variance reduction techniques, and criticality calculations. MCNP6 will track 32 different types of particles over a wide range of energies, including neutrons, photons, electrons, protons, muons, etc., plus heavy ions.

MCNP has a wide range of capabilities which make it useful for medical physics calculations. These abilities span its geometry representation, physics models, and source, tally and variance reduction capabilities. This talk reviews the history and capabilities of MCNP, and provides numerous examples of MCNP applications to medical physics and proton radiography experiments. Because all applications of Monte Carlo methods are limited by computer speeds, present and planned MCNP capabilities for parallel computation are also reviewed.

- **MCNP**
  - History & Overview
  - Applications – General
  - Applications – Medical Physics
  - MCNP6 – Capabilities & Status
  - Applications – Proton Radiography

- **Parallel Computing**
  - Hierarchical Parallelism
  - Parallel Scaling – MPI & Threading
  - Future – New Approach for Exascale

MCNP
Monte Carlo Code

What is MCNP?

- **General purpose Monte Carlo N-Particle radiation transport code.**
  MCNP5 & MCNPX → MCNP6

- **Tracks 32 different kinds of particles**
  Neutrons, photons, electrons, protons, muons, etc., plus heavy ions.

- **Standard features that make MCNP versatile and easy to use include:**
  - A powerful general source, criticality source, and surface source
  - Both geometry and output tally plotters
  - Many variance reduction techniques
  - A flexible tally structure
  - An extensive collection of cross-section data

- 3D general geometry
- PC, Mac, Linux, Unix, Sun support
- Parallel (MPI + threads)
- 350K+ lines of code
- Extensive verification / validation

- 400+ person-years development
- 10,000+ users world wide
- 15,000+ reference citations
- Distributed by RSICC code center
- Export controlled

MCNP History

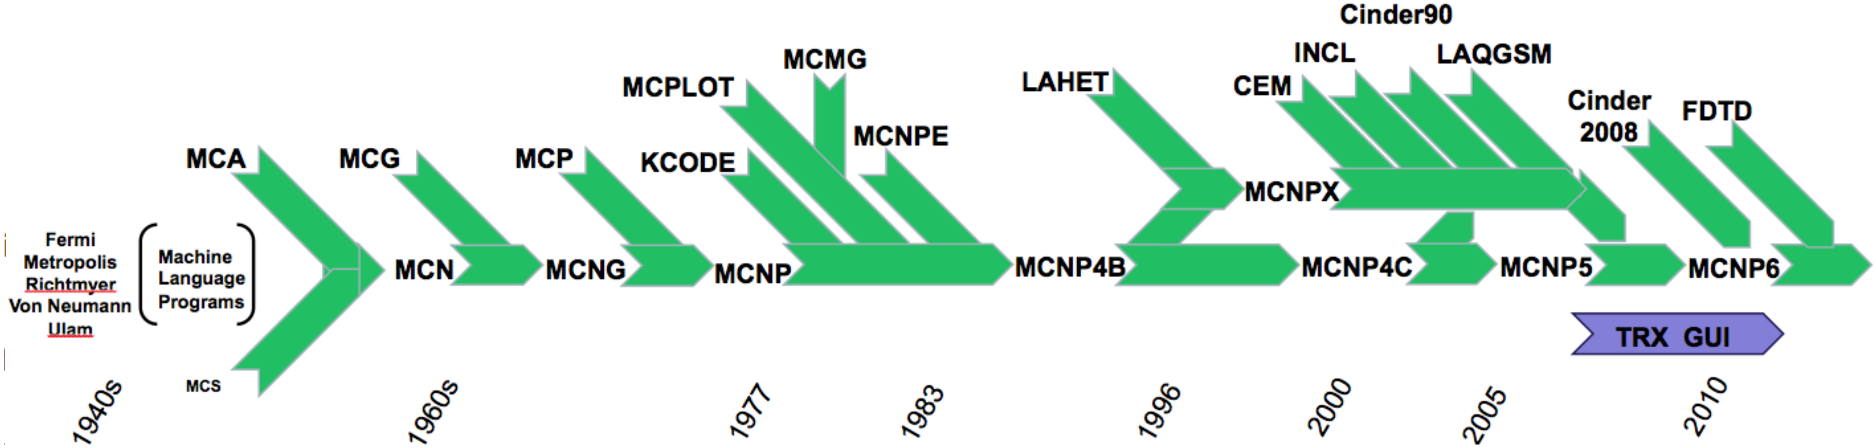

- **Monte Carlo transport of particles**
  - **MCNP5 - neutrons, photons, electrons**
  - **MCNPX - neutrons, photons, electrons + many more particles & ions**
  - **MCNP6 - merged code + more, 2011 - beta, 2012 - full release**
- **For 30+ years, MCNP & its data libraries have been supported by the Monte Carlo team at LANL**
  - **Roots of MCNP go directly back to von Neumann, et al.**
  - **Continuous development, support, R&D, V&V**

What Can MCNP Do?

Monte Carlo Codes
XCP-3, LANL

Detailed models of geometry & physics

- General 3D combinatorial geometry
- Repeated structures
- Lattice geometries
- Geometry, cross section, tally plotting
- ENDF/B-VII physics interaction data

Calculate nearly any physical quantity

- Flux & current
- Energy & charge deposition
- Heating & reaction rates
- Response functions
- Mesh tallies & radiography images
- K-effective, β_{eff}, η
- Fission distributions

Unique features for criticality calc's

- Shannon entropy of the fission source for assessing convergence
- Dominance ratio, k_1 / k_0
- Stochastic geometry
- Isotopic changes with burnup (mcnpx)
- Wielandt acceleration (soon)

> 10,000 users around the world

- Fission and fusion reactor design
- Nuclear criticality safety
- Radiation shielding
- Waste storage/disposal
- Detector design and analysis
- Nuclear well logging
- Health physics & dosimetry
- Medical physics and radiotherapy
- Transmutation, activation, & burnup
- Aerospace applications
- Decontamination & decommissioning
- Nuclear safeguards

Portable to any computer

- Windows, Linux, Mac, Unix
- Multicore, clusters, netbooks, ASC, …
- Parallel, scalable - MPI + threads
- Built-in plotting

Support

- Extensive V&V against experiments
- Web site, user groups, email forum
- Classes - 1 week, 6x / year

Many Mission Examples

- **Stockpile Stewardship**
  - Criticality Safety
  - Radiography
- **Nuclear regulation**
  - Verify requests from NRC & industry
- **Nuclear reactor design & analysis**
  - Reactor physics analysis
  - Verification/validation
- **Threat reduction**
  - Urban consequences
- **Non-proliferation**
  - Reactor actinide inventories
  - Portal monitors
  - Active interrogation
  - Detectable Quantities of materials
- **Medical & health physics**
  - Shielding design
  - Radiology, radiation therapy
  - Treatment planning
- **Proton radiography simulation, for beams in the GeV range**
  - Experiments
  - Simulation
- **Benchmarking & data testing**
  - ENDF/B-VII data testing,
- **Parallel calculations**
  - ASC teraflop systems
  - Linux clusters
- **Others…**
  - Fukushima reactor accident
  - Oil well logging tool design
  - Semiconductor radiation damage
  - Radiography for BP oil well damage

Nuclear Regulatory Commission Use of MCNP

Criticality Safety:

- To assess the criticality safety of licensed facilities that handle fissionable materials.

Radiation Shielding:

- To benchmark other shielding and dose calculation computer codes and methods used by NRC staff.
- To verify licensees’ shielding and dosimetry calculations.

Radiation Dosimetry:

- Assess planned and unplanned worker radiation exposures.
- Assess public exposure from planned licensing actions.

Medical:

- To understand the radiation safety implications of using radiation in medical diagnosis and treatments.

MCNP = Benchmark for Nuclear Reactor Design codes

MIT
Research Reactor

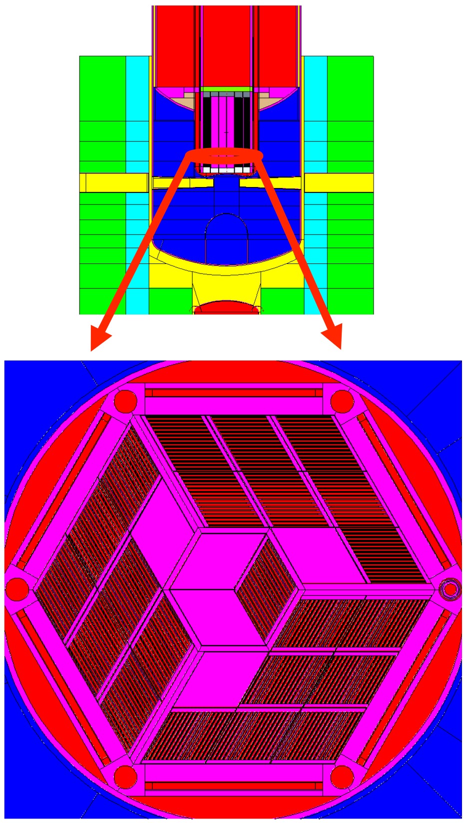

ATR
Advanced Test Reactor

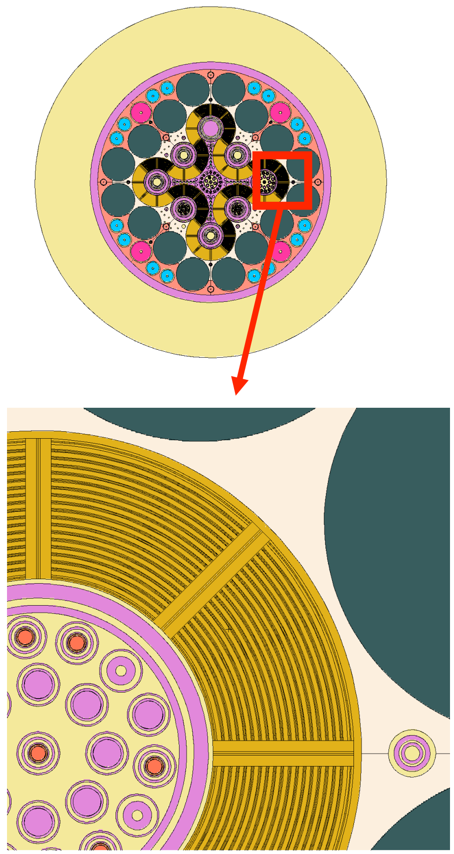

PWR
Pressurized Water Reactor

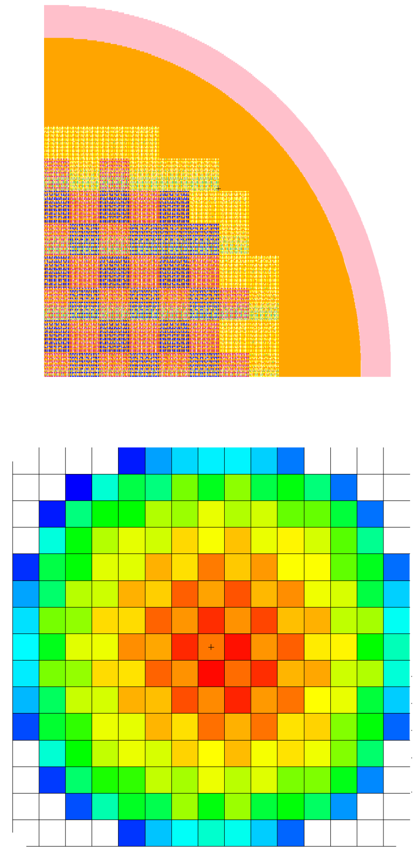

VHTR
Very High Temperature Gas-Cooled Reactor

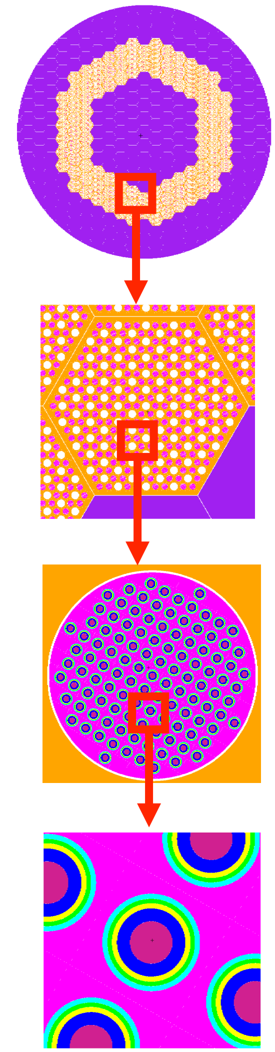

- Accurate & explicit modeling at multiple levels
- Accurate continuous-energy physics & data

Advanced Reactor Design - VHTR, HTGR, ...

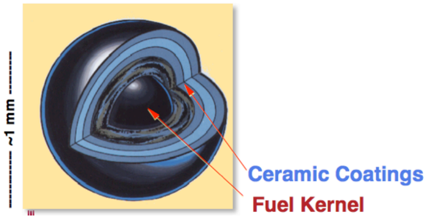

TRISO FUEL PARTICLES

FUEL COMPACTS

FUEL BLOCK

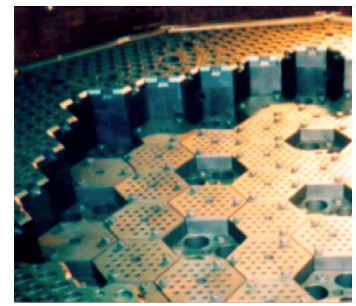

CORE

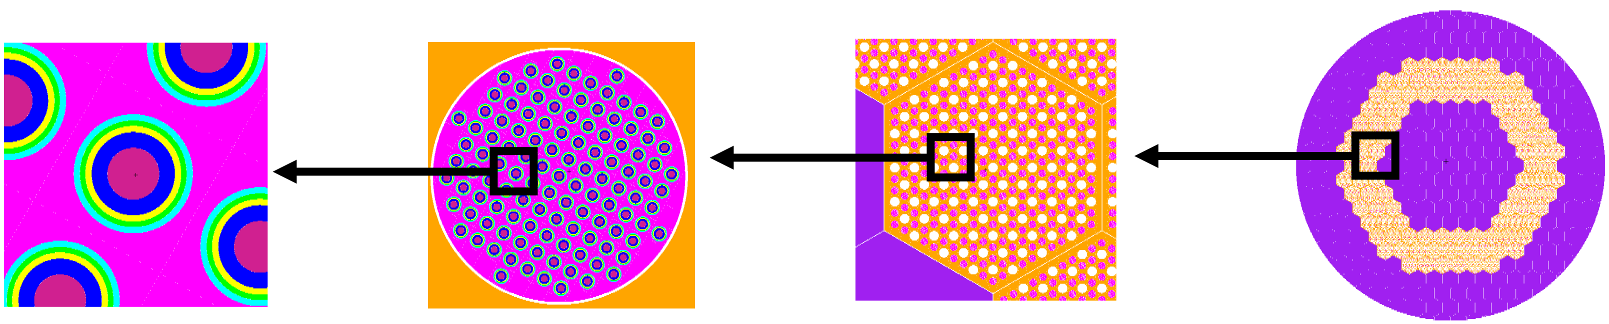

MCNP model - accurate & explicit at multiple levels

Analysis of IND Dose Effects

Google maps

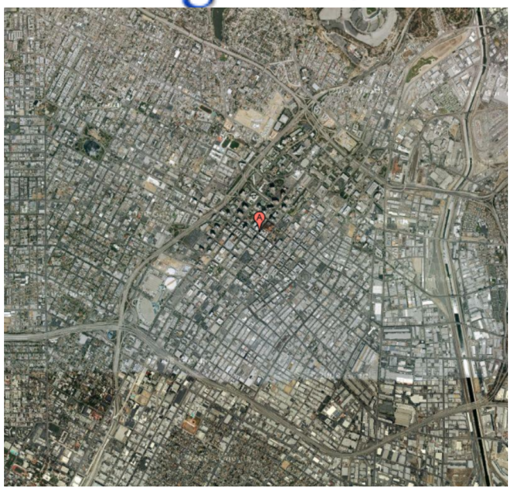

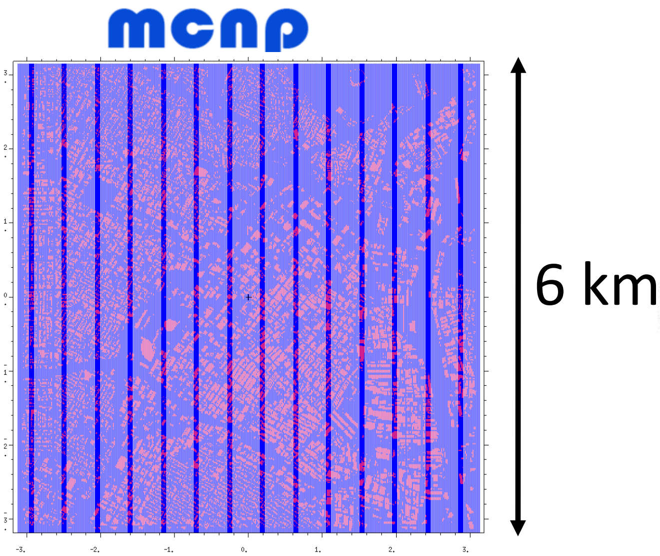

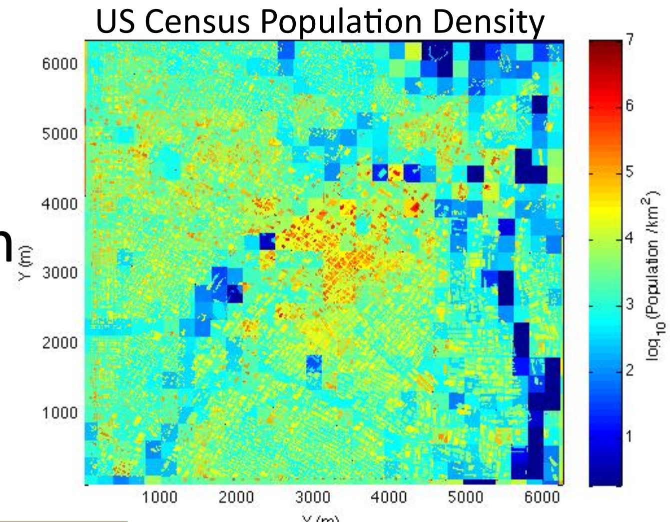

Prompt Radiation Effects - Dose

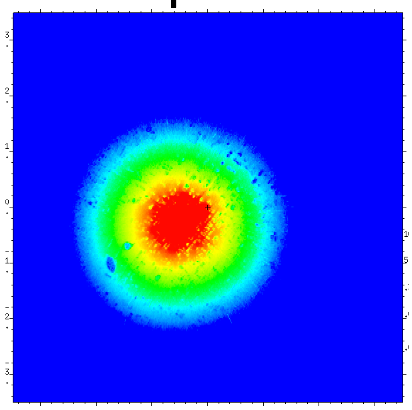

Neutron Dose (from neutron leakage)

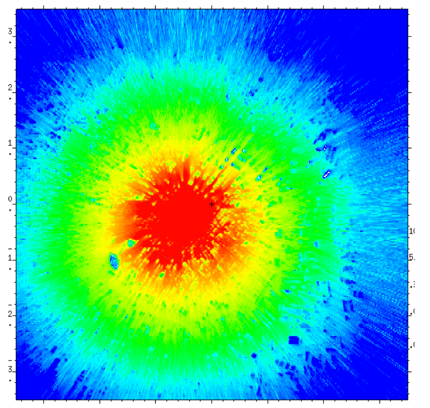

Gamma Dose (from neutron capture)

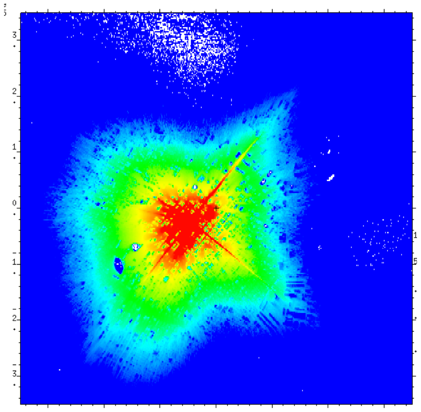

Gamma Dose (from gamma leakage)

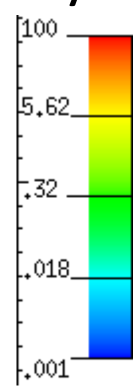

Fallout Dose

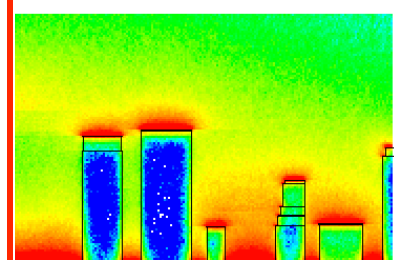

Gamma Dose

Dose contours from a 20 kT Little Boy device in downtown LA

Analysis of IND Circuit Effects

Google maps

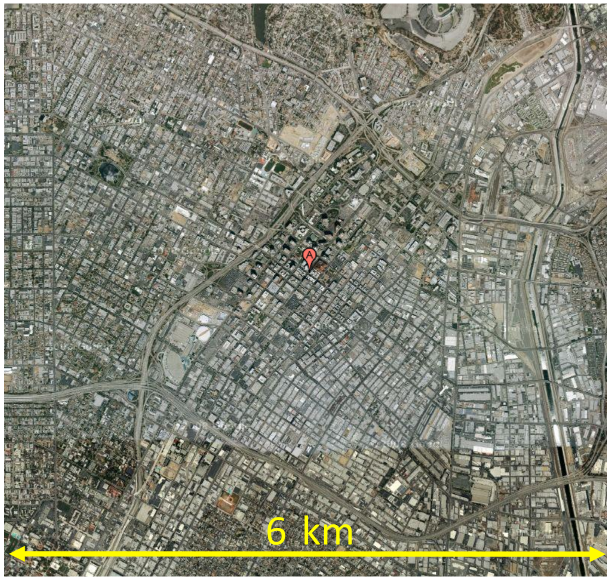

mcnp

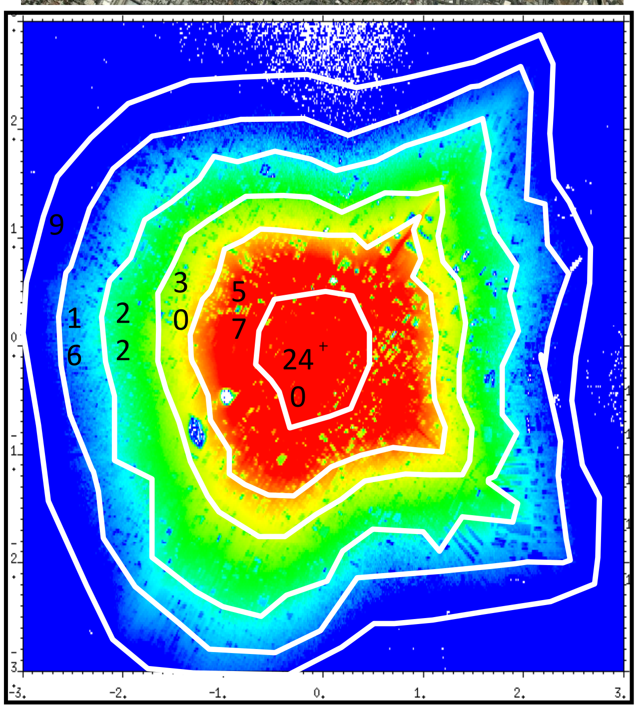

- Surface burst of Fat Man ~ 10 kT
- White contours: electric field strength in kV/m
- Circuit failure expected above 10 kV/m
- Color scale: photon flux
- Note 4 orders of magnitude decrease in underlying photon flux

$\frac{\gamma \text{ Flux}}{\gamma \text{ cm}^3}$

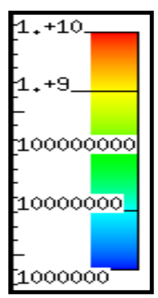

The results presented here are based on source region simulation levels from MCNP. This is part of the LANL EMP start-up project's goal of incorporating first physics principle source region calculations.

Uses of MCNP in Medical & Health Physics

Some real examples, not all-inclusive:

- **External Radiation Source Design**
  - Optimize filter materials thickness to get optimum energy, angle and spatial characteristics.
- **Shielding Design**
  - Analyze shielding to minimize staff radiation dose & determine construction costs.
  - Investigate skyshine & calculate build up factors
- **Radiation Detection**
  - Investigate scattering in design of CT machines
- **Dose reconstruction**
  - Dose received after event
  - Dose to various organs from internal radiation exposure
- **Treatment Planning**
  - Determine optimum radiation beams
- **Evaluate proposed forms of radiation therapy**

Radiography Calculations

- Radiography tallies

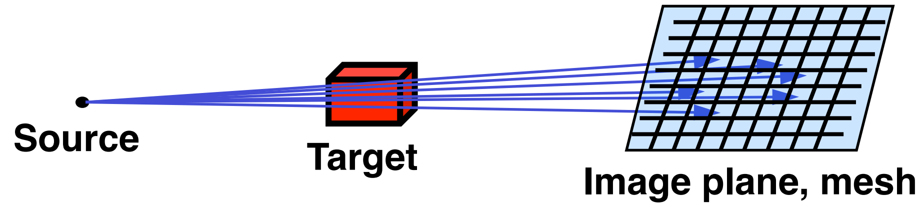

- Neutron and photon radiography uses a grid of point detectors (pixels)
- Each source and collision event contributes to all pixels

MCNP Model of Human Torso

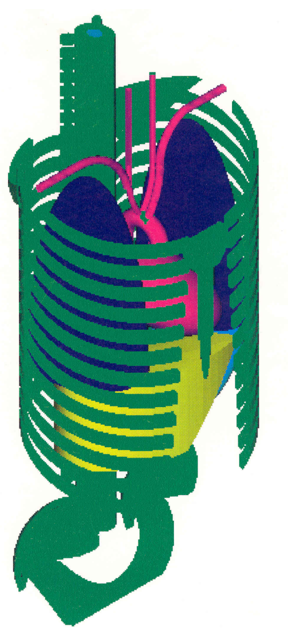

Simulated Radiograph - 1 M pixels

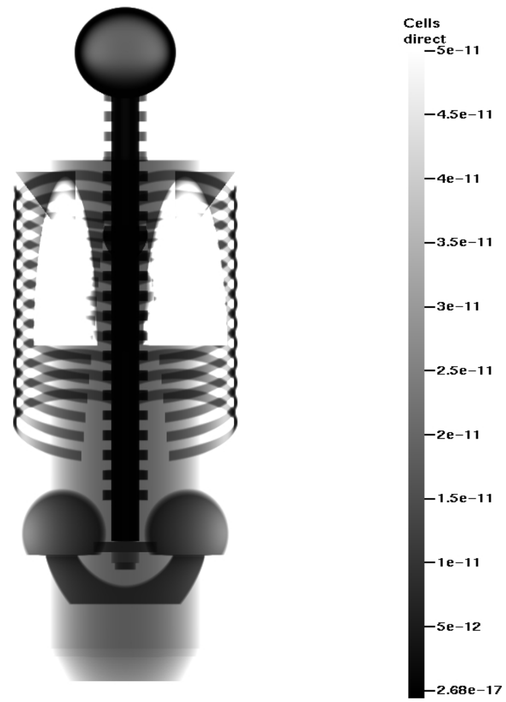

Treatment Planning with MCNP

- MCNP is used to calculate dose distributions throughout a target region from different radiation beam orientations.

- Post processing programs are then used to (perhaps) combine different beams to maximize dose to tumor while minimizing dose to surrounding healthy tissue. These programs usually overlay the dose contours over patient specific CT images.

- While the peak tumor and tissue dose are usually based on in-phantom dose rate measurements, the simulation is necessary to determine more advanced parameters, such as the dose volume histogram.

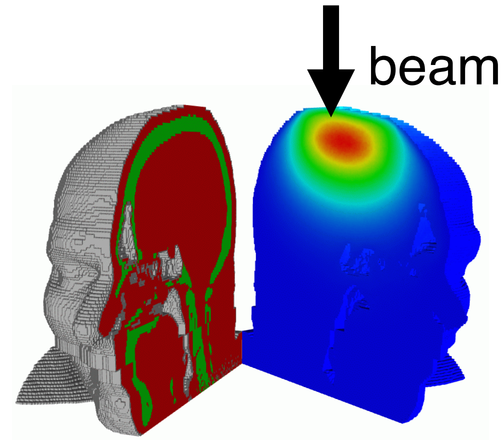

CT Based geometries are possible to represent in MCNP

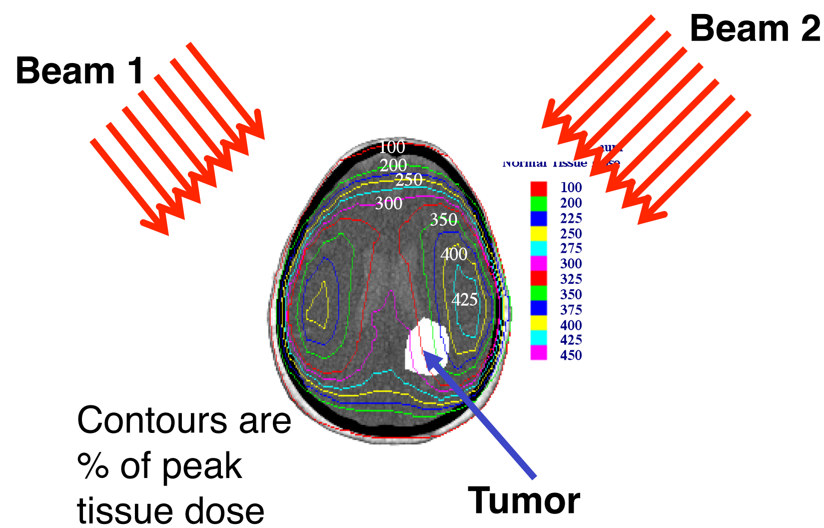

Medical Physics – Phantoms & Voxel Models

Monte Carlo Codes
XCP-3, LANL

Zubal phantom

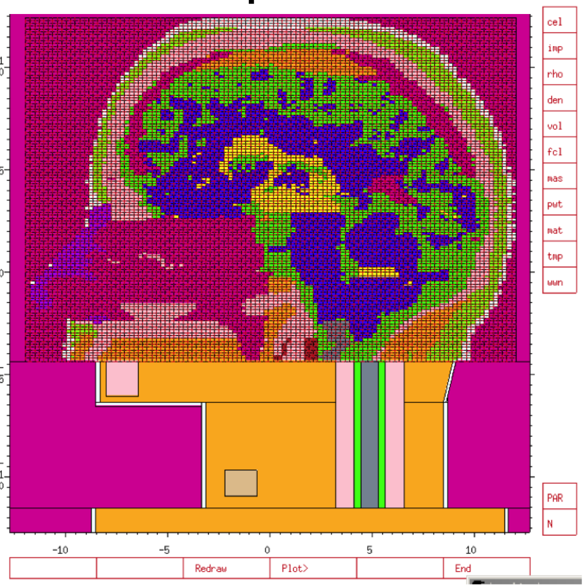

Yanch, MIT

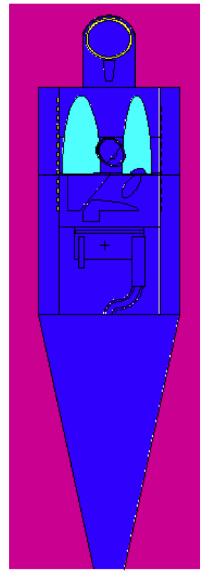

ORNL

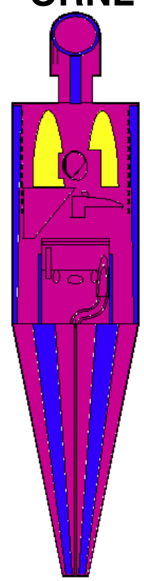

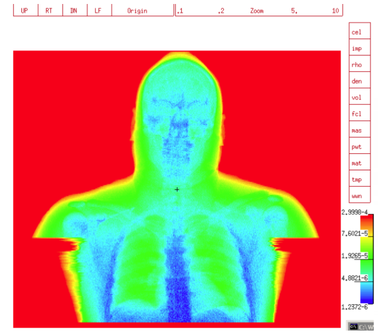

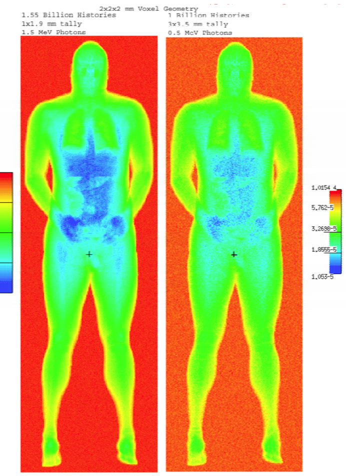

Radiographs of VIPMan model,
1x1x1 mm voxels (above),
2x2x2 mm voxels (right)
Images from MCNP5 plotter

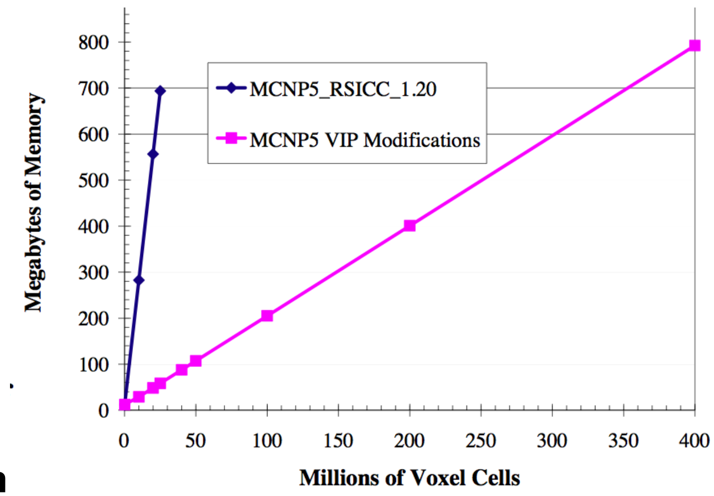

VIP Man

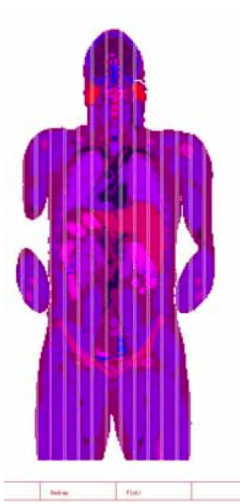

Snyder head phantom

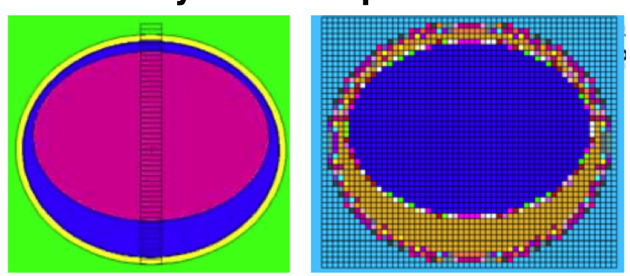

Medical Physics - Treatment Planning

- **MCNP6**
  - 3D unstructured mesh
  - Embedded in 3D MCNP geometry
  - Many applications
    - Radiation treatment planning
    - Linkage to Abaqus
  - Under development

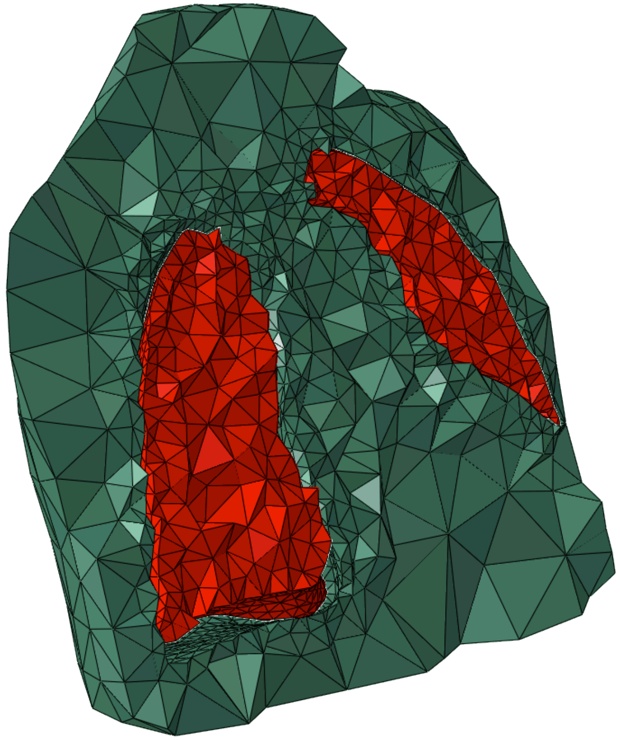

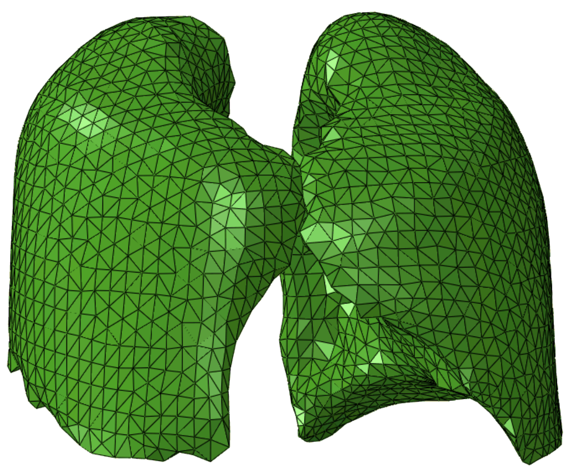

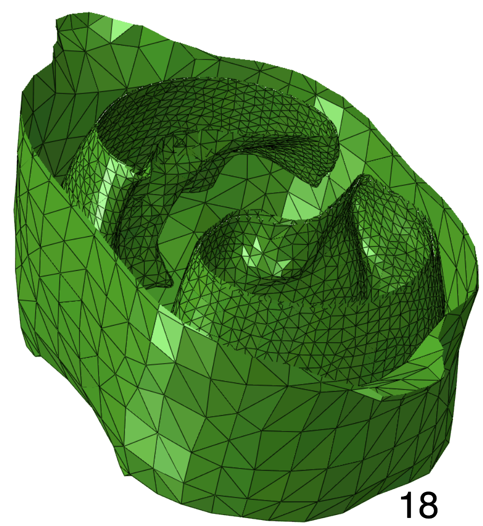

MCNP is widely used for radiation cancer therapy research

The code is ideally suited for use in medical applications because of the accuracy of its physics models, the unique set of clinically relevant features, and the responsive support provided by the developers and the user community.

We used MCNPX to verify the Mass General Hospital Proton Center, and this information has gone into the design of the MDACC proton center and others, which are used to treat > 5K people a year.

Wayne Newhauser, Ph. D.
Dept of Radiation Physics

THE UNIVERSITY OF TEXAS MD ANDERSON CANCER CENTER

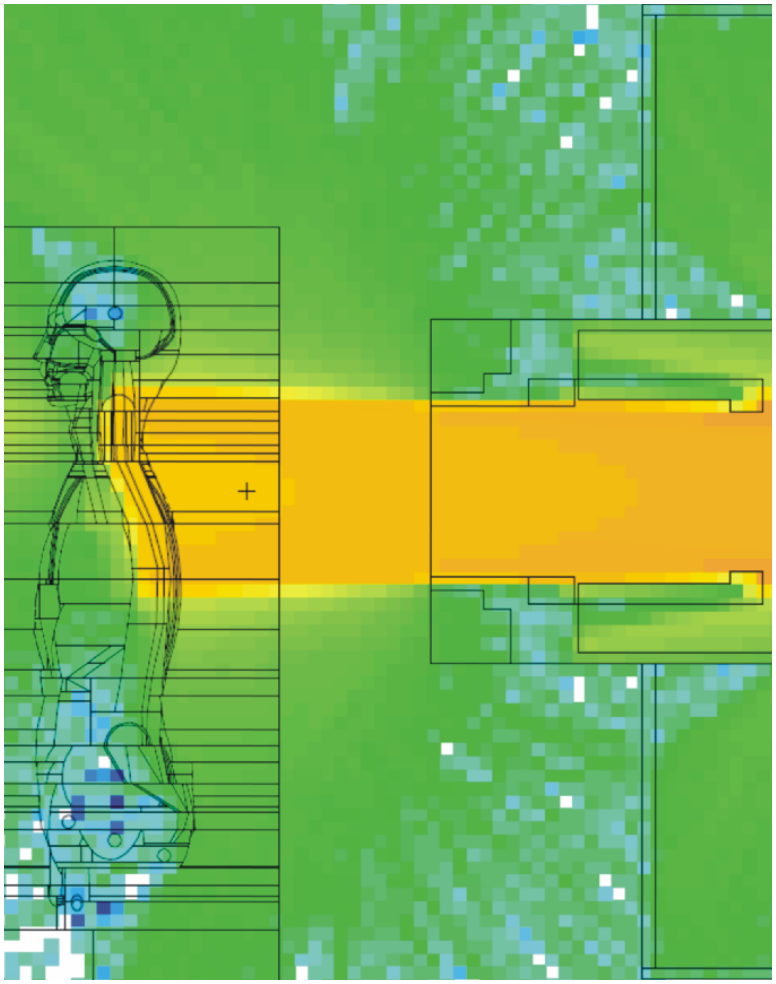

proton fluence and dose contours (arb units)

Medical Physics - Dose Calculations

- Patient-CT based model of **knee & end of accelerator**
- **Calculate dose throughout knee**
- Study impact of moderating/ shielding materials & B^{10} conc. in knee
- Need other code to determine neutron production in accelerator target

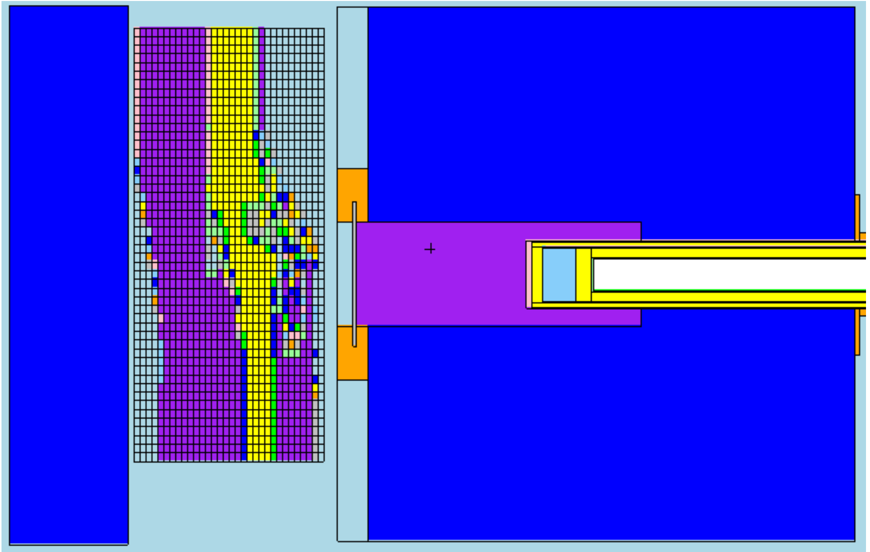

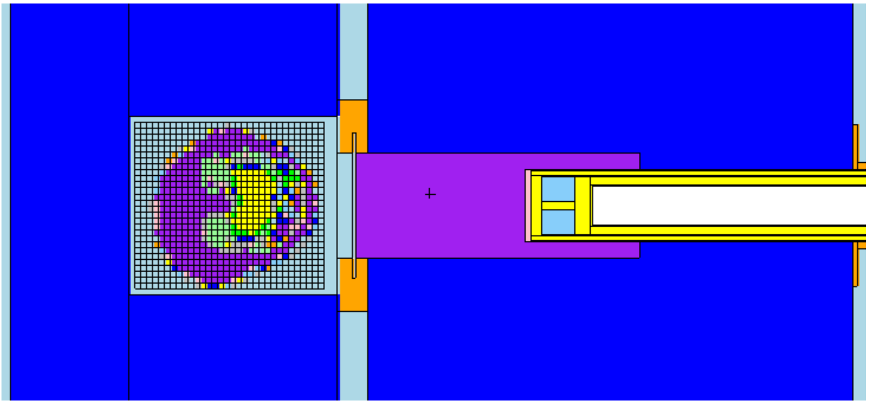

J. R. Albritton, ‘‘Analysis of the SERA treatment planning system and its use in boron neutron capture synovectomy,’’ M. S. thesis, Massachusetts Institute of Technology, 2001.

Gierga DP, Yanch JC, Shefer RE, “An investigation of the feasibility of gadolinium for neutron capture synovectomy”, Med Phys. 2000 Jul;27(7):1685-92.

Pictures from mcnp plotter

Proton & Carbon Therapy Applications

Monte Carlo Codes
XCP-3, LANL

Proc. ND2010 April 26-30, 2010, Jeju Island, Korea, in press:

Spectrum Measurement of Neutrons and Gamma-rays from Thick $H_2{}^{18}O$ Target Bombarded with 18 MeV Protons

M. Hagiwara [a], T. Sanami [a], Y. Iwamoto [b], N. Matsuda [b], Y. Sakamoto [b], Y. Nakane [c], H. Nakashima [c], K. Masumoto [a], Y. Uwamino [d] and H. Kaneko [e]

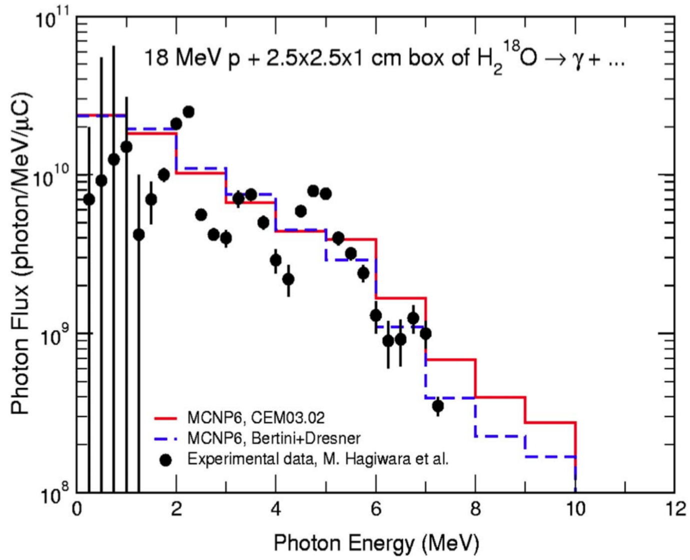

Experimental data are from: Y. Iwata et al., Phys. Rev. C64 (2001) 054609;
QMD, HIC, and LAQGSM03 results are from: H. Iwase et al., AIP 769 (2005) 1066

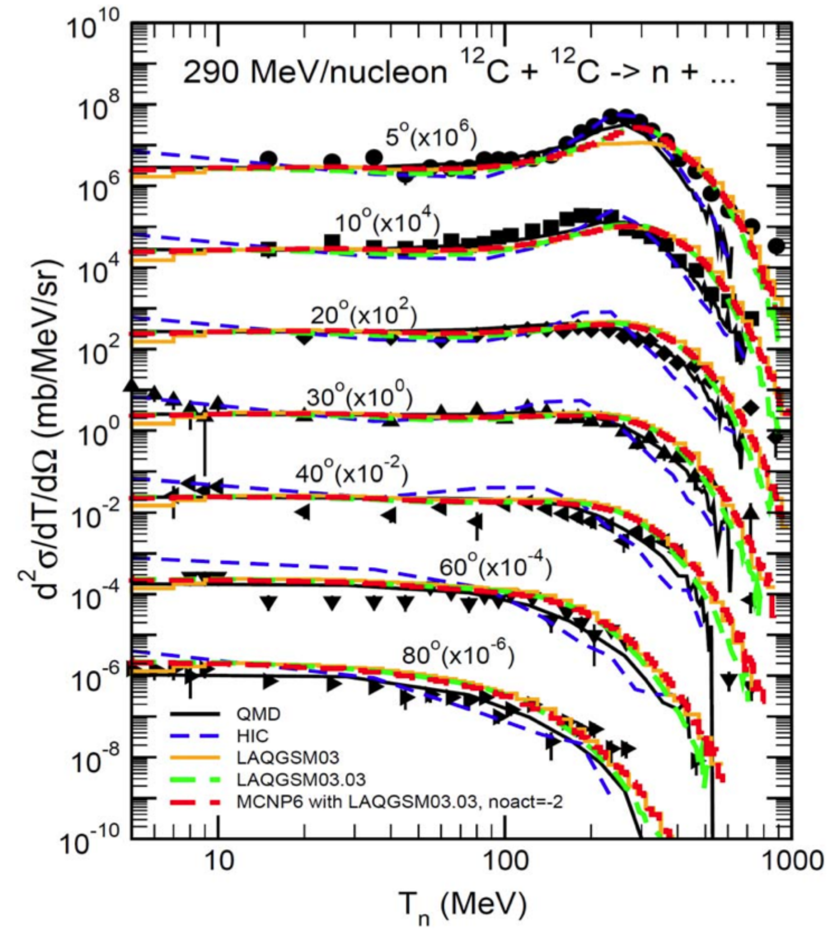

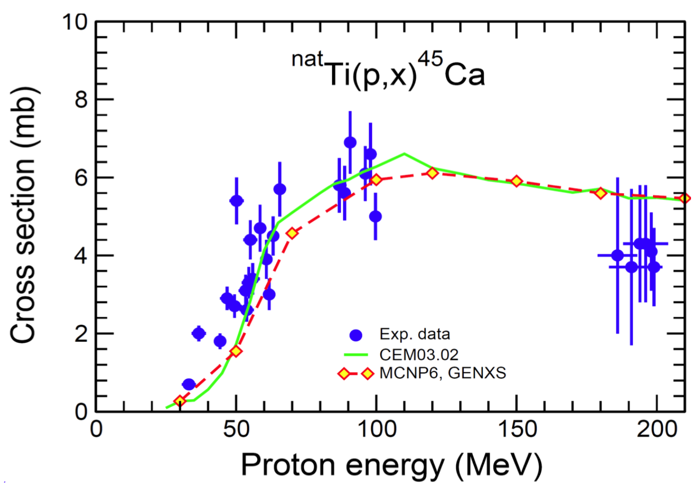

Radiochim. Acta **98**, 447–457 (2010)

Excitation functions of nuclear reactions leading to the soft-radiation emitting radionuclides ^{45}Ca, ^{49}V and ^{204}Tl in beam collimator materials used in proton therapy

By S. M. Qaim*, K. Kettern, Yu. N. Shubin+, S. Sudár# and H. H. Coenen

References

Monte Carlo Codes
XCP-3, LANL

High energy transport in MCNPX & MCNP6 have been validated & used in proton therapy for a variety of clinical and research applications, see, e.g.:

- M. R. James et al., NIM A562 (206) 819
- J. D. Fontenot et al., Med. Phys. 34 (2007) 489
- M. C. Harvey et al., Med. Phys. 35 (2008) 2243
- J. Herault et al., Med. Phys. 35 (2008) 2243
- W.D. Newhauser et al., Phys. Med. Biol. 52 (2007) 4569
- J. C. Polf et al., Med. Phys. 34 (2007) 4219
- P. J. Taddei et al., Phys. Med. Biol. 53 (2008) 2131
- Y. Zheng et al., J. Nucl. Mater. 361 (2007) 298
- T. Urban and J. Kluson, PNST10136 (2010)

Low-energy transport code MCNP5 has been validated & even more widely used in a variety of medical applications, see, e.g.:

- J. T. Goorley et al., LANL Report LA-UR-02-7205
- T. Goorley and D. Olsher, LA-UR-05-2755
- A. Lazarine and T. Goorley LANL Reports LA-UR-05-4598,LA-UR-05-6402, and LA-UR-06-4904
- A. L. Reed, LA-UR-10-4133
- I. Gerardy et al., J. Phys.: Conf. Ser. 102 (2008) 012012
- J. Zhang et al., Health Phys. 91 (2006) S59
- I. Gerardy, Appl. Rad. Isot. 68 (2010) 735
- Y.J. Huang and M. Blough, J. Appl. Clinic. Med. Phys.,11 (2010) 46
- H. G. Hughes and J. T. Goorley, LA-UR-10-05424 (2010)

Distributed with MCNP:

Medical Physics Calculations With MCNP: A Primer,

A.L. Reed, LA-UR-07-4133, Los Alamos (2007)

Resources: mostly code intercomparisons, but some benchmarks

- Computing Radiation Dosimetry – CRD 2002 (published by OECD)
  http://www.oecdbookshop.org
- QUADOS – European lead Code Intercomparison (2003)
  http://www.oecd-nea.org/download/quados/quados.html
- EURADOS & CONRAD (EU intercomparison) (Active & Ongoing)
  http://www.eurados.org
- American Nuclear Society - Computational Medical Physics Working Group (Active & Ongoing)
  http://cmpwg.ans.org/

MCNP6
Status

MCNP6 Beta Release

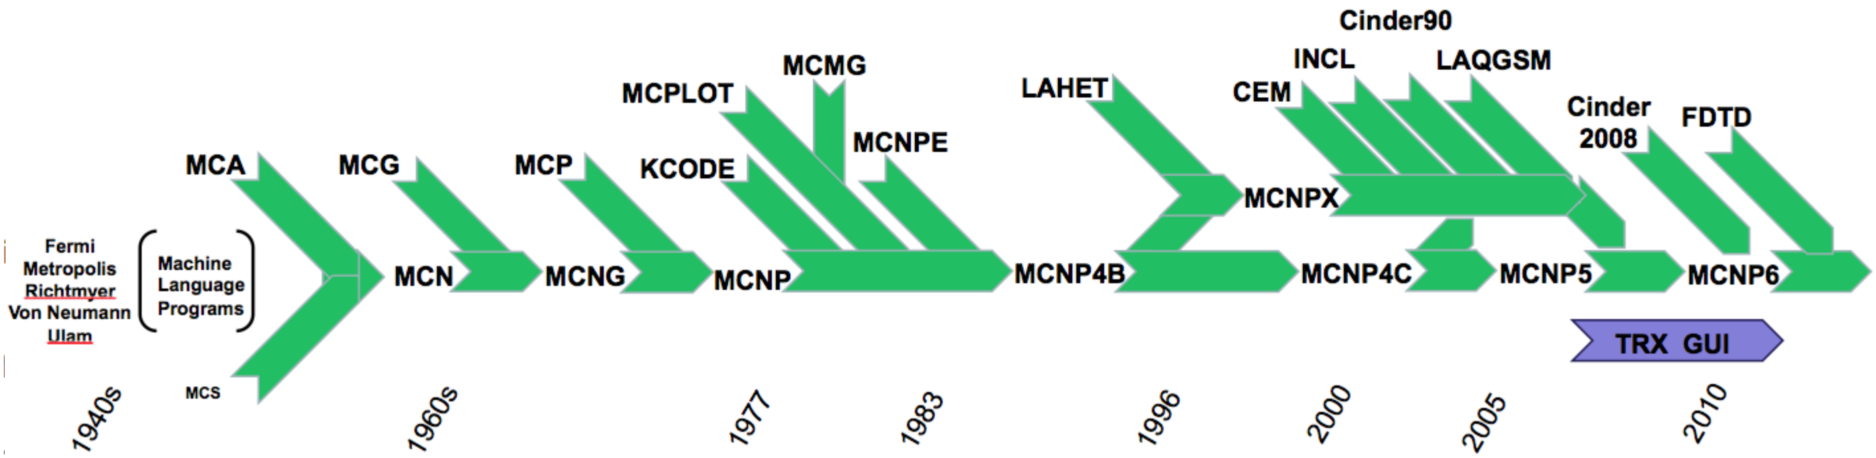

- MCNP6 – beta release sent to RSICC for a limited set of beta testers
- MCNP6 – full release by RSICC expected in 2012
- Culminates 5 years of effort combining all features of MCNPX-2.7.0 into MCNP5
- Both MCNP5 & MCNPX are now frozen - future development will occur in MCNP6

Support from DOE/NNSA, DOE, DoD,
DRTA, DHS/DNDO, NASA, & others

Current & Future MCNP6 Efforts

The LANL MCNP6 team has more than 12 full time and 5 part time staff working on the following:

- **Improved Physics**
  - Incorporate new INCL, add delta rays, improve stopping power , add Rutherford scattering, allow particle to pick up charge as they slow down
- **Improved Software parallelism**
  - to be able to utilize >10K processors w/ mpi, R&D into Cray Fortran
- **Improved Delayed Particle Emissions**
  - better energy and angle correlations, beta and alpha emissions
- **Efforts for EMP**
  - Adding Electric Fields, Improved magnetic fields, specialized tallies
- **Integration of Unstructured Mesh**
  - work with weight windows mesh, charged particle tracking
- **Optical Light**
  - refraction, reflection, Cherenkov radiation
- **Moving Objects**
  - Realistic simulation of moving vehicles
- **Sensitivity and Uncertainty**
- **Automatic Weight Windows Generation**
  - from SN calculations – LANL's PARTISN.

MCNP contains a lot of physics

- **Incorporates other codes as libraries:**

– LAHET	high energy transport	LANL
– CEM	high energy transport	LANL
– LAQGSM	high energy transport	LANL
– CINDER	unstable nuclei database	LANL
– ITS	electron transport	SNL
– MARS	high energy transport	FNAL
– HETC	high energy transport	ORNL

- **Utilizes Nuclear and Atomic Data**
  - LANL, LLNL, BNL, EU, Japan

- **Large energy range (eV – 100s of GeV)**

MCNP Physics

- **MCNP is physics rich – try to use best data, models, & theory**

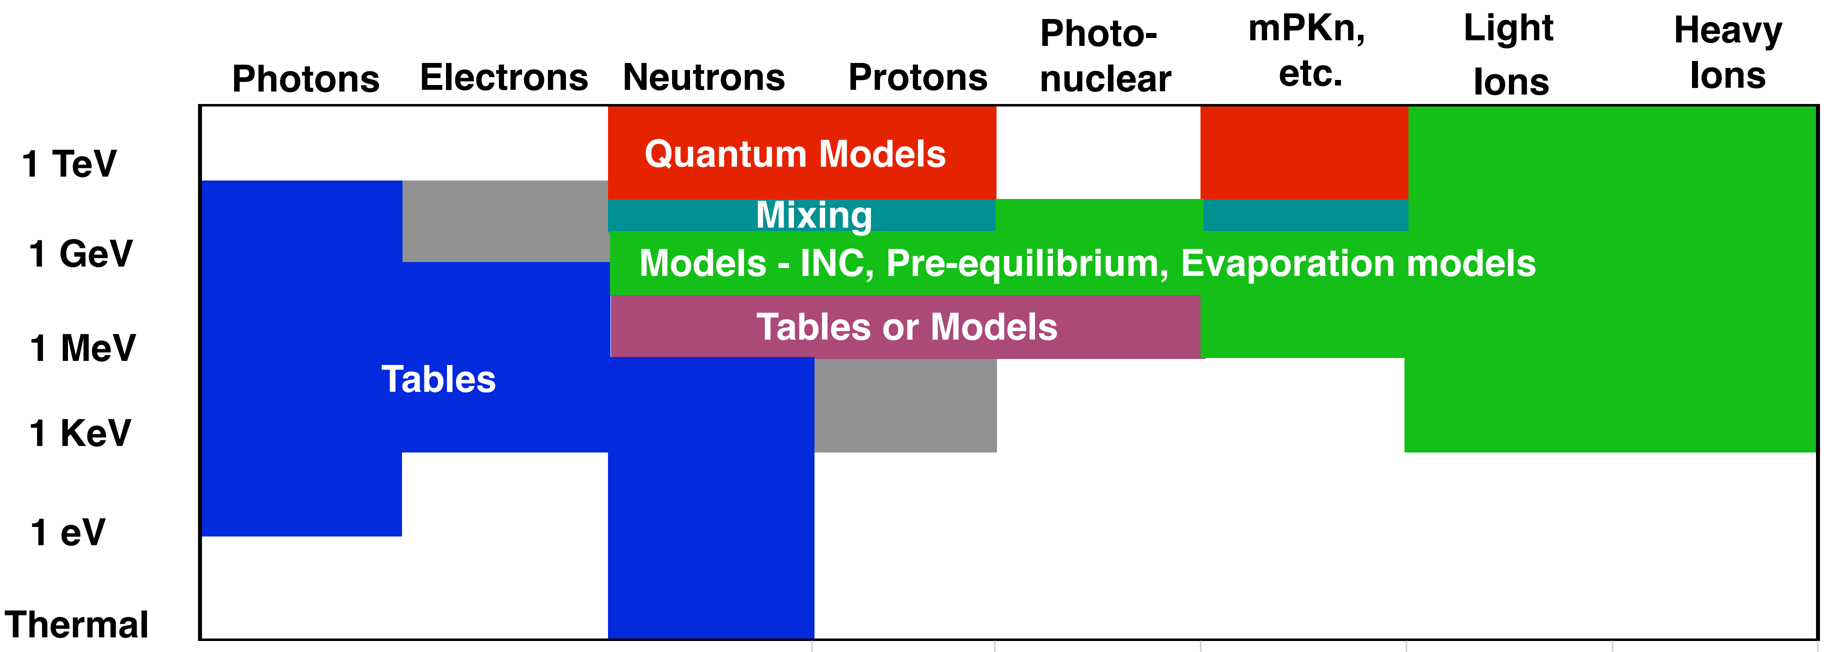

- **Recent physics improvements include:**
  - Photon induced fission multiplicity
  - Characteristic muonic X-rays
  - Exact delayed gamma emissions
  - Visible light
  - Improved photoatomic form factors
  - Upgrades to CEM & LAQGSM 3.03
  - GEF photofission yield

MCNP6 Status

- **MCNP6 contains:**
  - MCNP6 = development version of MCNP at LANL, since 2004
  - Includes:
    - All MCNP5-1.60 capabilities (mpi + threads)
    - High energy protons & magnetic fields, for proton radiography
    - All MCNPX 2.7.D capabilities (mpi)
    - CINDER 2010 decay & depletion
  - Unstructured mesh, for linking with ABAQUS
  - Structured mesh, for linking with PARTISN
  - Adjoint-weighted perturbation estimators

- **MCNP6 in (very) limited beta release to outside LANL**
  - Recipients are active collaborators and sponsors
  - Full beta access within LANL and LLNL

- **Active Validation Efforts**
  - Comparisons with experiments included in test suites
  - High energy proton, heavy ion interactions
  - Delayed photon and neutron spectra
  - Subcritical multiplication
  - Expanded criticality suite (119 problems)
  - Perturbation verification suite
  - Kobayashi benchmarks – streaming through ducts & voids
  - Reactor kinetics parameter benchmarks
  - Production / depletion (CINDER) soon

- **Nightly Regression Test suites**
  - 3 platforms (Linux 32, Linux 64, Windows 64)
  - 5 compilers (Intel 10+11, PGI 7, Pathscale 3, gfortran)
  - Serial, mpi, omp, mip+omp
  - Array bounds checking
  - 875 problem input files
  - Total: 10,000 runs each night

MCNP6 Status

- MCNP & MCNPX teams have adopted MCNP6 as the base for all future development
- To go from Beta release to Production release:
  - Assurance of reliability and accuracy for criticality
  - Assurance of reliability and accuracy for other apps
  - Comparable performance
  - Complete documentation
- Future Work
  - Cleanup coding style
  - Remove duplicate features
  - Extend parallel threading capability to new features
  - New Features
- General release through RSICC
  - 2012

Proton Radiography

Richard Prael, Grady Hughes, John Zumbro, John Sarracino,
Jeff Bull, Lon-Chang Liu, Stepan Mashnik, Arnold Sierk,
Forrest Brown, Tim Goorley, Jeremy Sweezy,
Robert Little, Morgan White, Elizabeth Selcow,
Nikolai Mokhov (FNAL), Sergei Striganov (FNAL),
Konstantin Gudima (Acad. Sci. Moldova)

Proton Radiography

- **For many experiments being conducted now at LANL & BNL, high-energy proton beams are directed at test objects to produce radiographic images**
  - LANL: 800 MeV proton beams
  - BNL: 24 GeV proton beams
  - Proposed: 50 GeV proton beams

- **Proton beams are collimated & focused by magnetic lenses**

- **Both the design of the experiments & analysis of results are carried out using MCNP6, the latest LANL development version of MCNP**
  - All MCNP5 features plus:
  - Continuous-energy proton physics up to 50 GeV
  - Models for multiple Coulomb scatter, nuclear elastic scatter, etc.
  - Direct tracking of protons through magnetic fields
  - COSY-map tracking of protons through magnetic fields
  - Many additional particle types being added to account for background

Proton Radiography

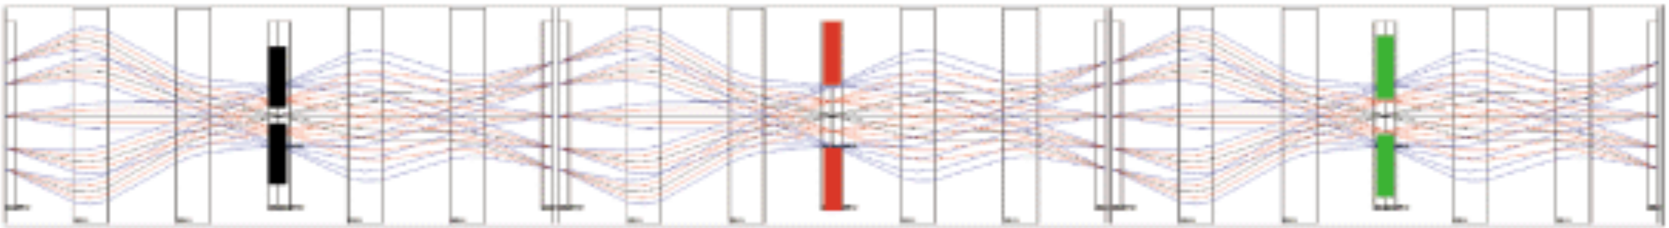

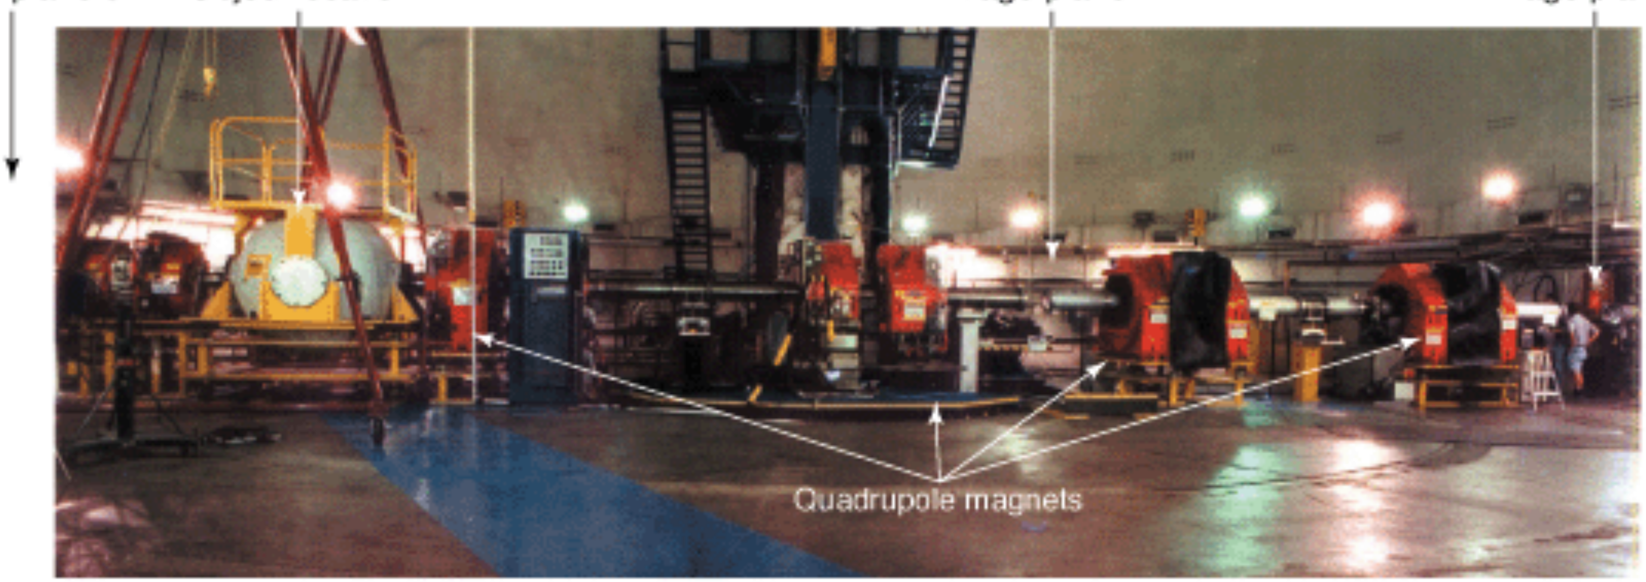

Explosive proton radiography experiments are conducted at the Los Alamos Neutron Science Center facility. In these experiments, a proton beam traveling inside a tube penetrates a target placed in a spherical vessel (left) to contain the explosion. Quadrupole magnets (orange) focus the scattered protons onto imaging detectors. This particular setup uses three imaging stations, including one installed in front of the target to examine the profile of the incoming proton beam. Collimators are located inside the beam tube.

Brookhaven Experiments

Nominal trajectories for BNL AGS E955 / E963 beamline

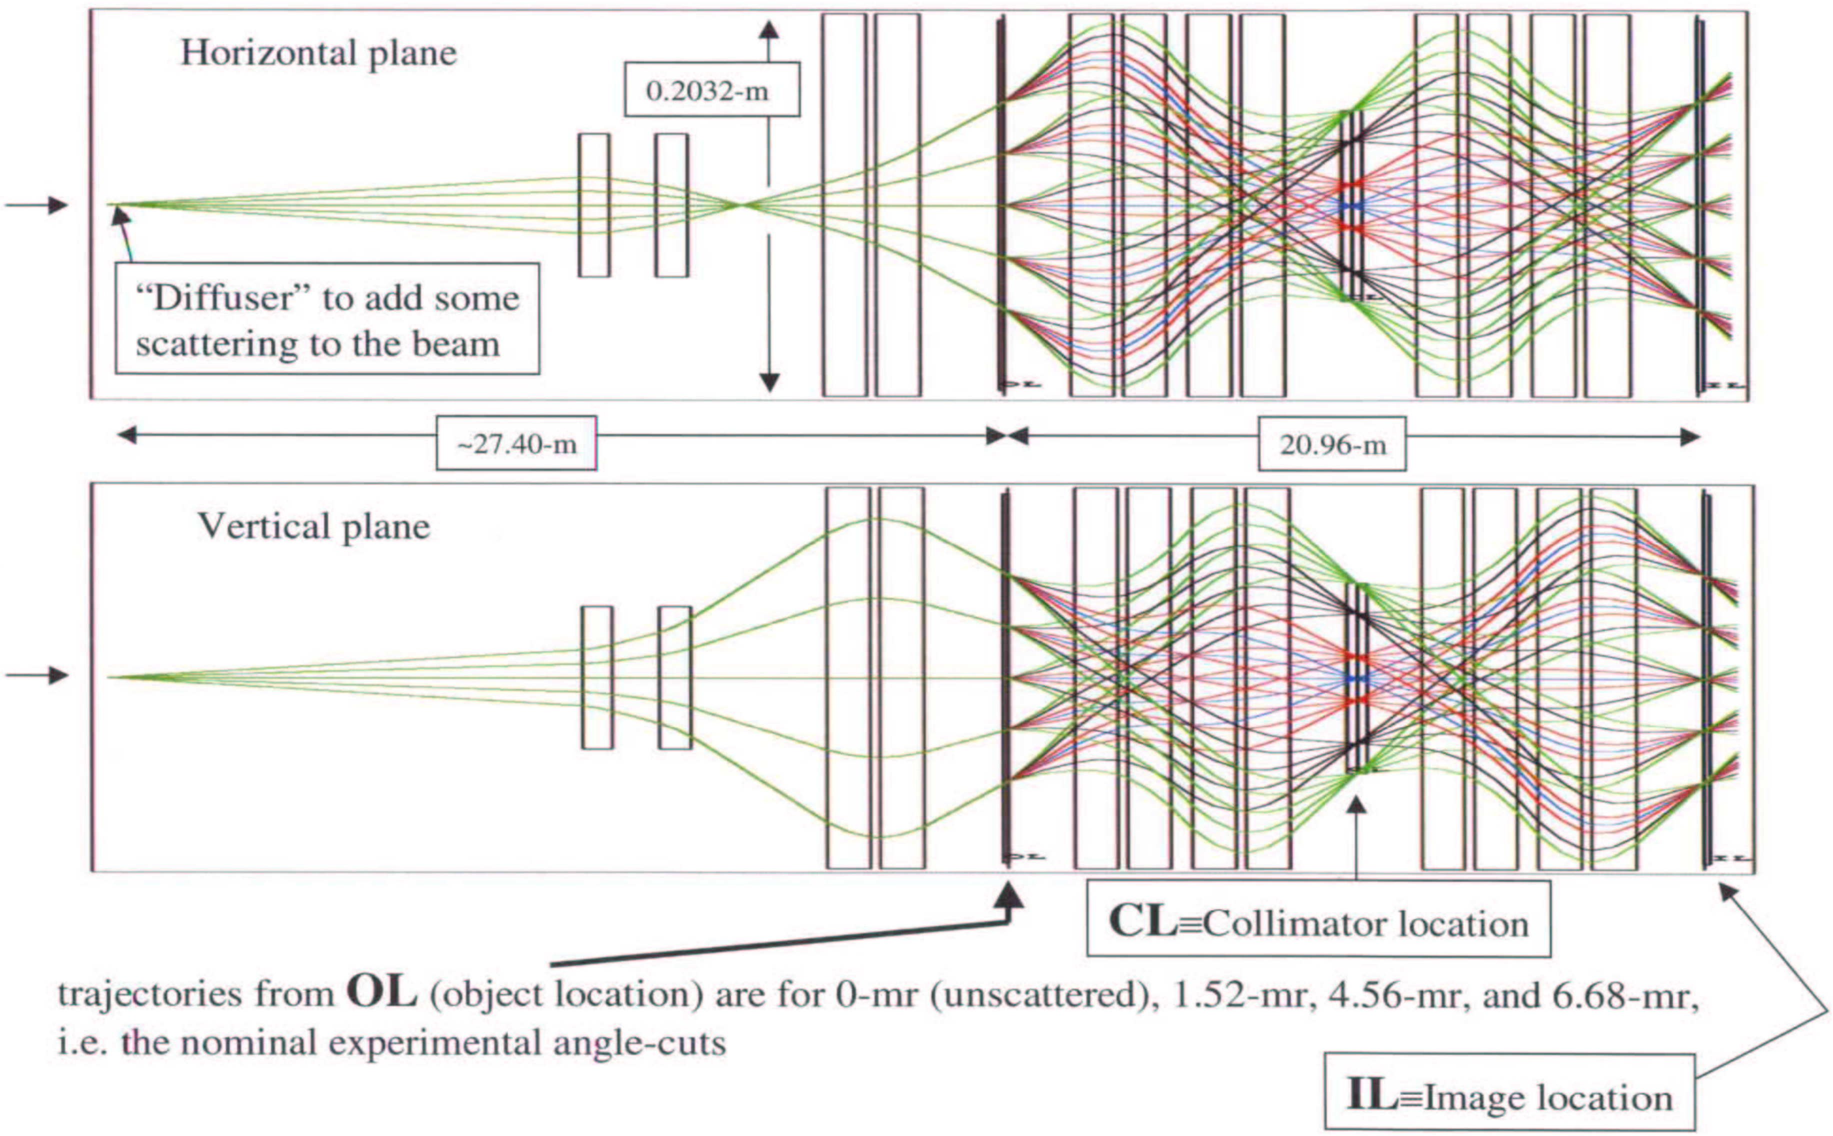

trajectories from **OL** (object location) are for 0-mr (unscattered), 1.52-mr, 4.56-mr, and 6.68-mr, i.e. the nominal experimental angle-cuts

MCNP6 Simulation

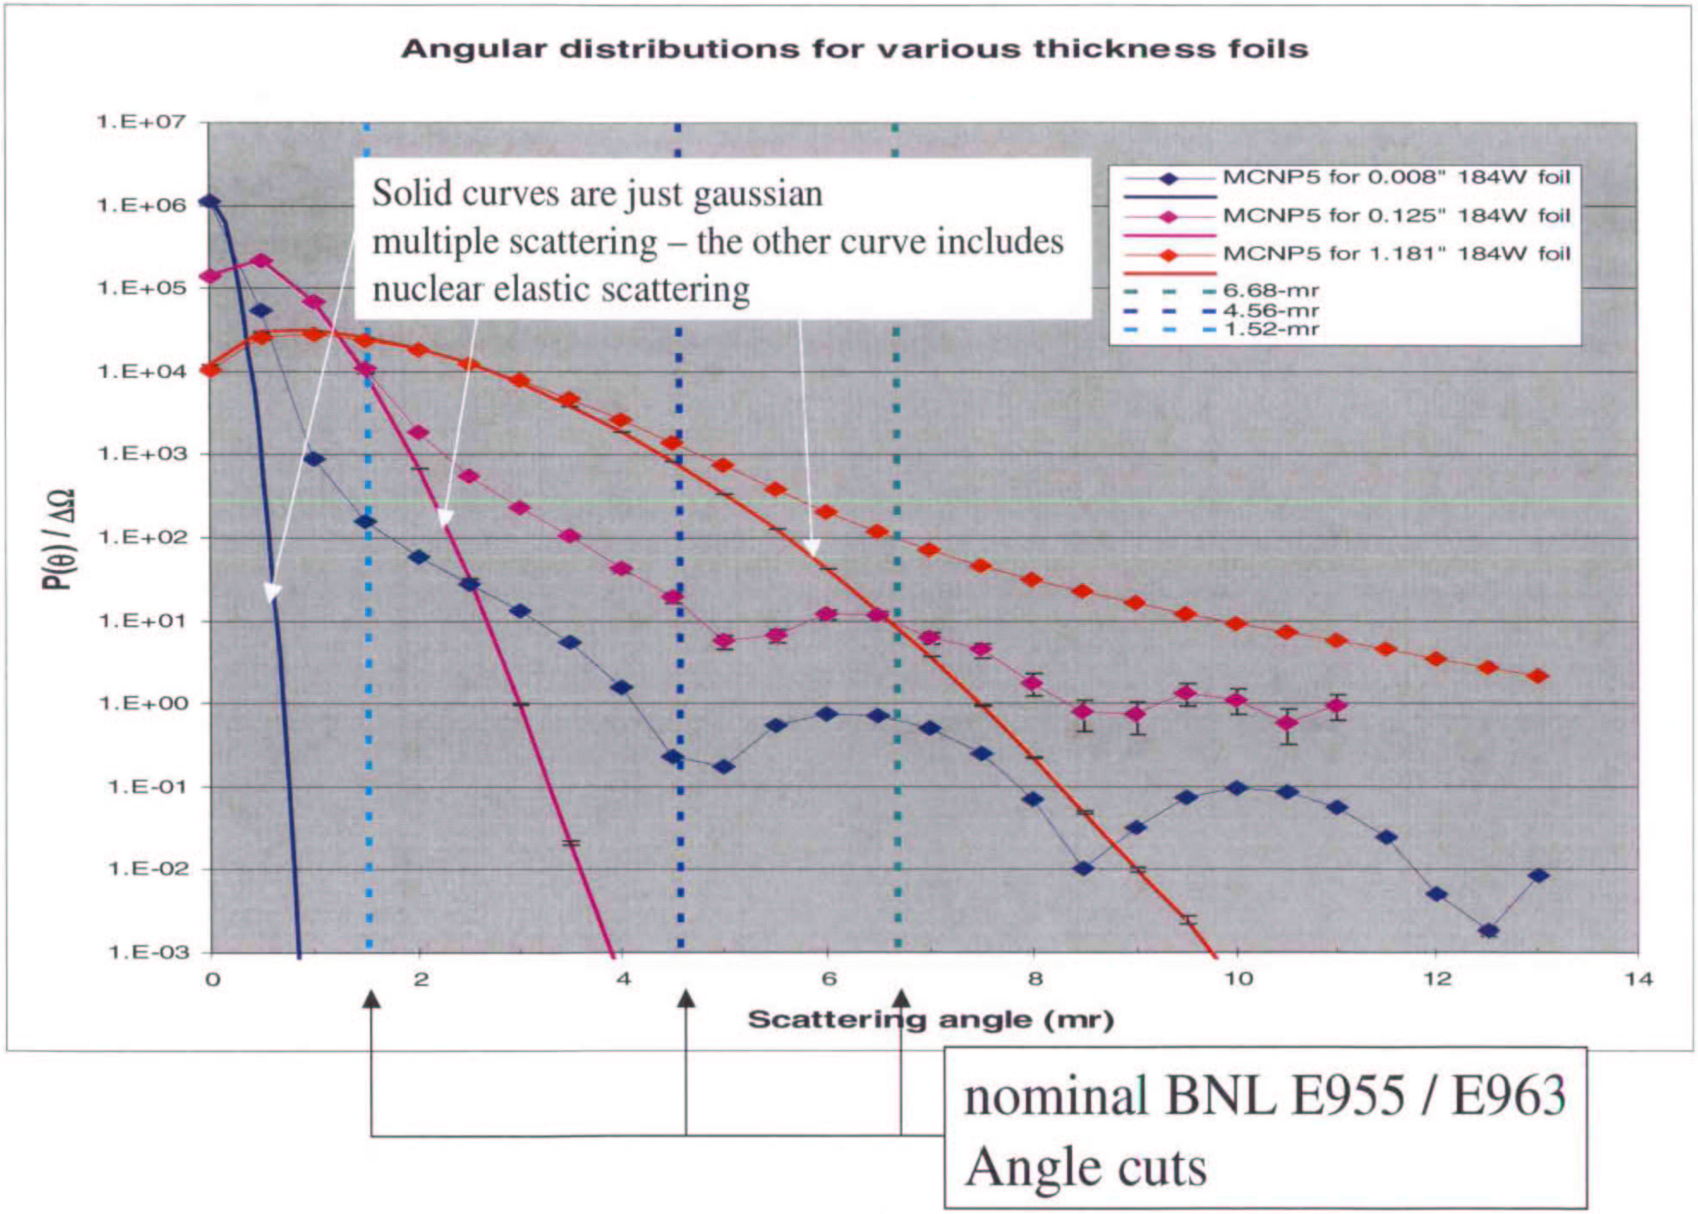

Brookhaven Experiments

Iron target: Blue = data, Yellow = MCNP6 simulation.

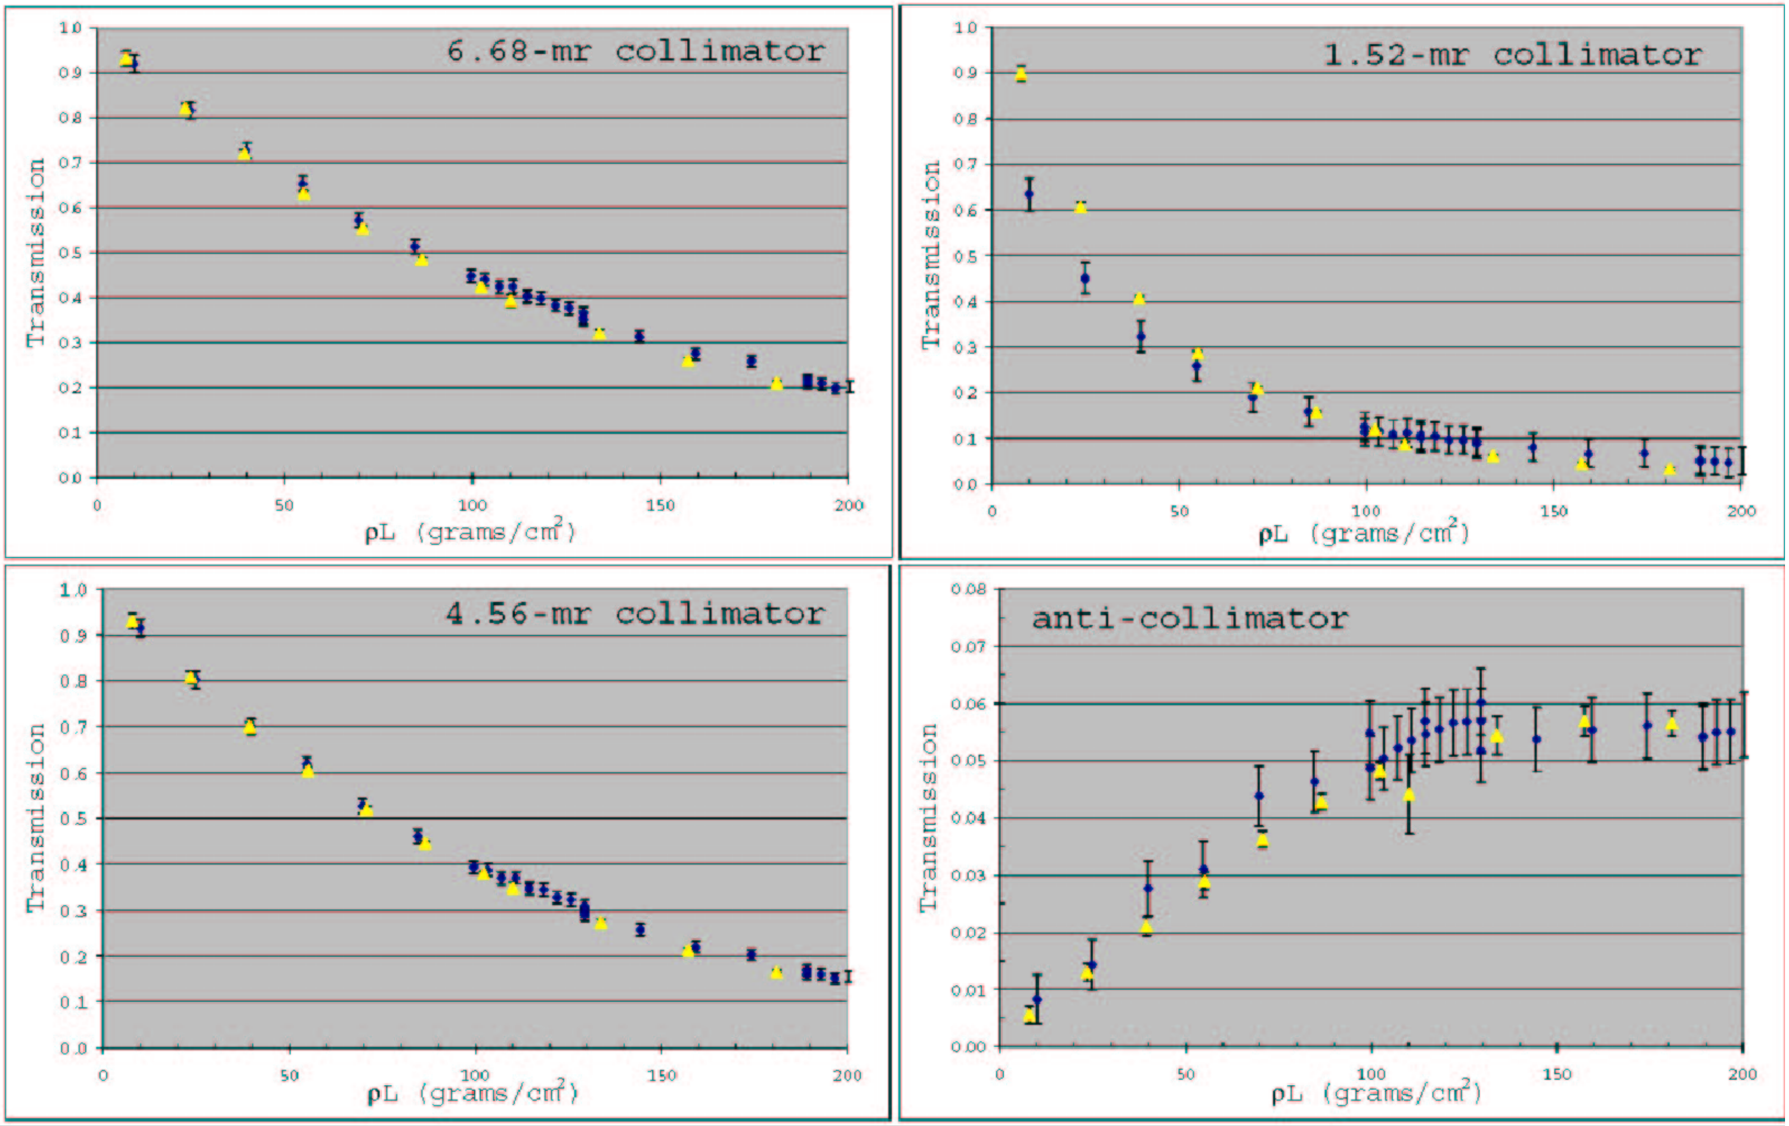

Brookhaven Experiments

Monte Carlo Codes
XCP-3, LANL

Beryllium target: Blue=data, Yellow=MCNP6, Magenta= w/o nuclear elastic

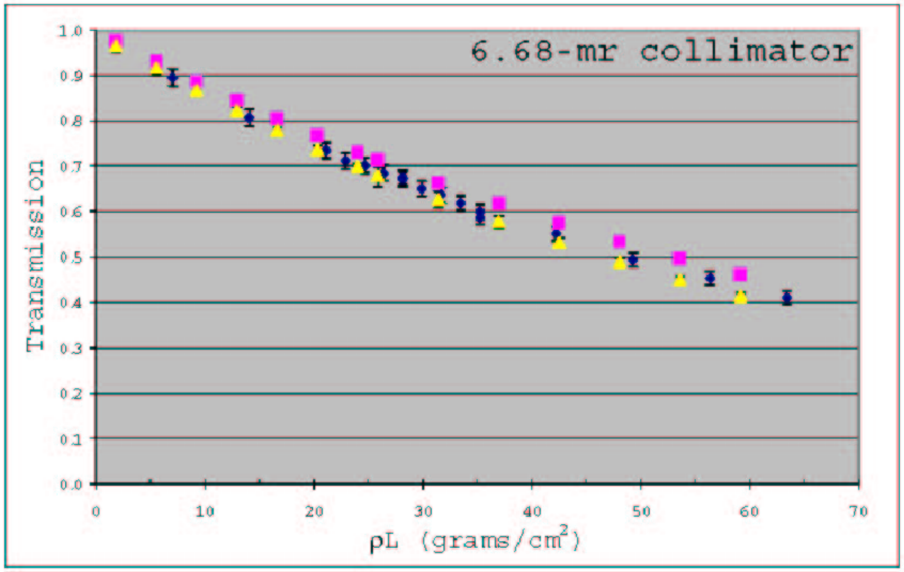

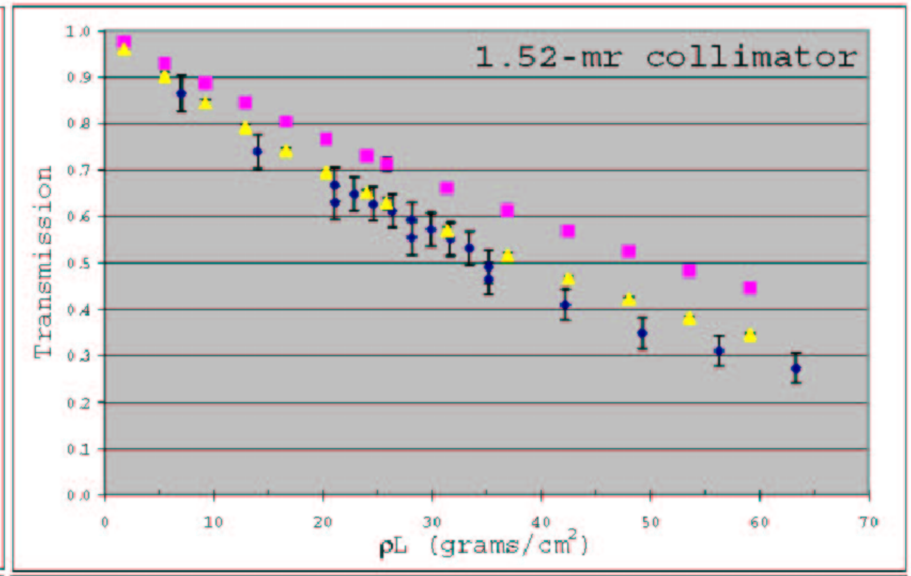

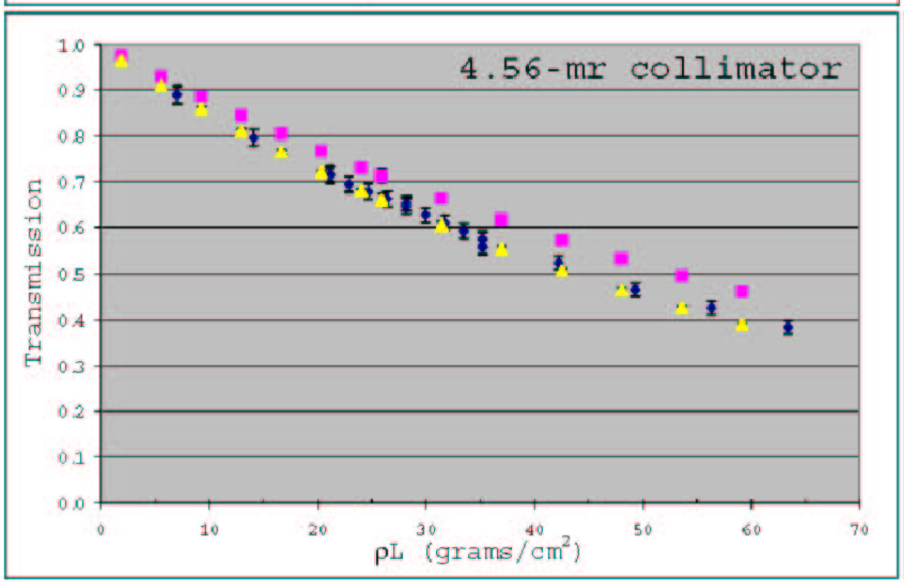

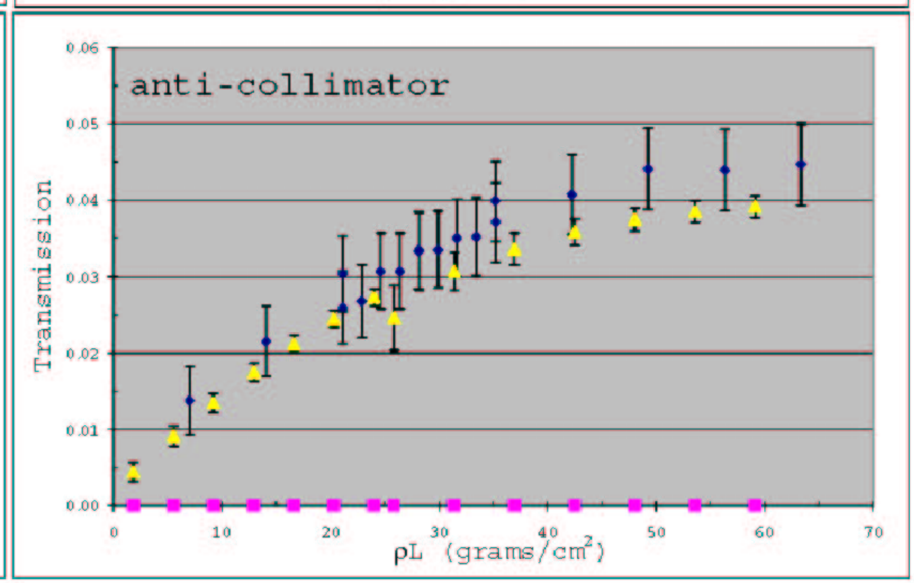

Proton in Air & Constant B Field

No Energy Straggling

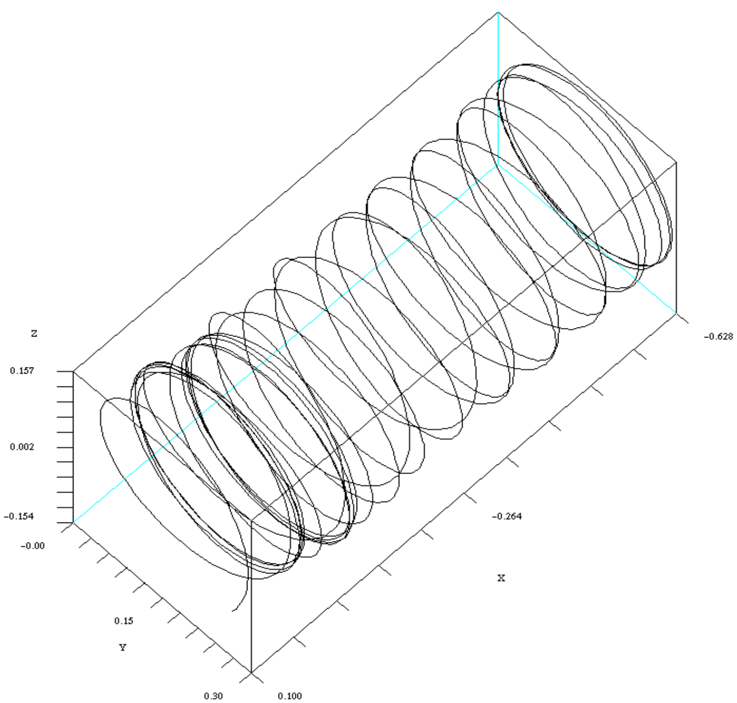

With Energy Straggling

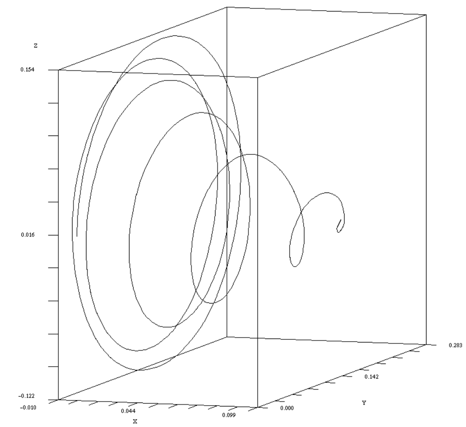

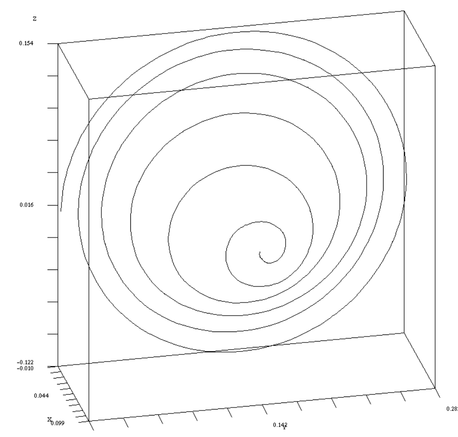

Parallel
Monte Carlo

Trends in Computing Technology

- **Commodity chips**
  - Microprocessor speed → ~2x gain / 18 months
  - Memory size → ~2x gain / 18 months
  - Memory latency → ~ no change (getting worse)
- **High-end scientific computing**
  - Key driver (or limit) → **economics:** mass production of desktop PCs & commercial servers
    power: reduced energy usage for very large parallel clusters
  - Architecture → **clusters**: with small/moderate number of commodity microprocessors on each node
    multicore: multiple CPUs per processor permits threading within each node processor
- **Operating systems**
  - Desktop & server → Windows, Linux
  - Supercomputers → Unix, Linux

CPU performance on supercomputer ➔ same as desktop PC

High-performance scientific computing ➔ parallel computing

Hierarchical Parallelism

- **For clustered SMPs,**
  - Use message-passing to distribute work among slaves ("boxes")
  - Use threading to distribute histories among individual processors on box

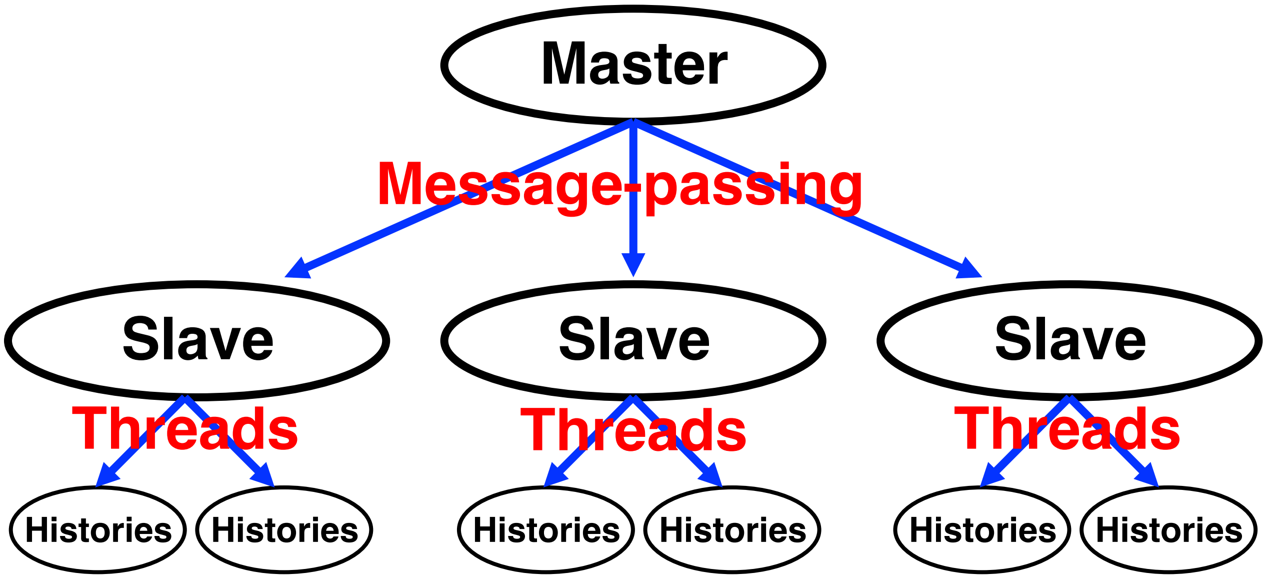

- **Only the master thread on each slave uses MPI send/recv's**
- **Threads on each slave share memory**

Simplified Master/Slave MPI Scheme

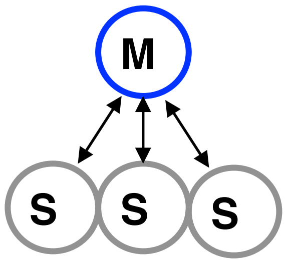

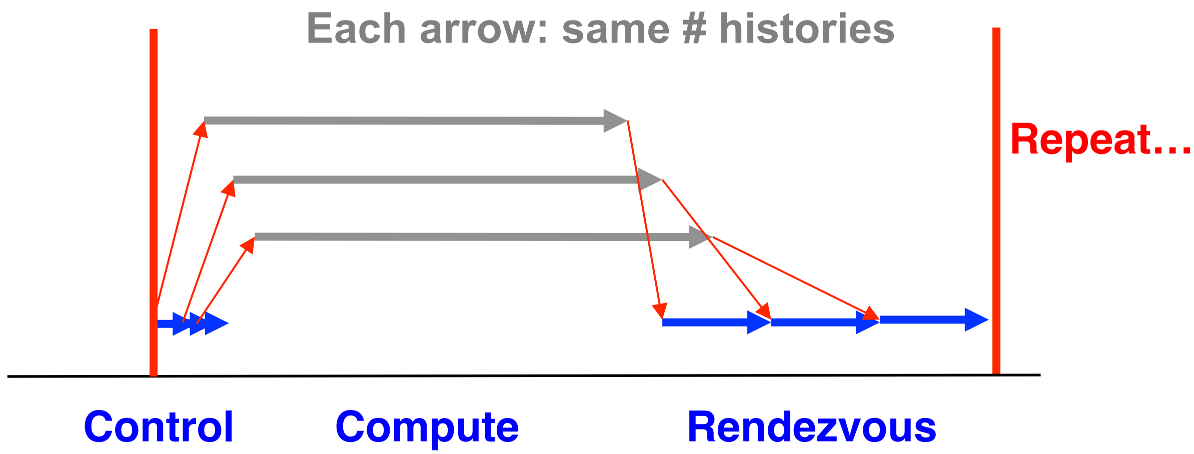

- **For efficiency, want (compute time) >> (rendezvous time)**
  - Compute time: Proportional to #histories/task
  - Rendezvous time: Depends on amount of tally data & latency+bandwidth for message-passing

MCNP – With Load Balancing & Fault Tolerance

Monte Carlo Codes
XCP-3, LANL

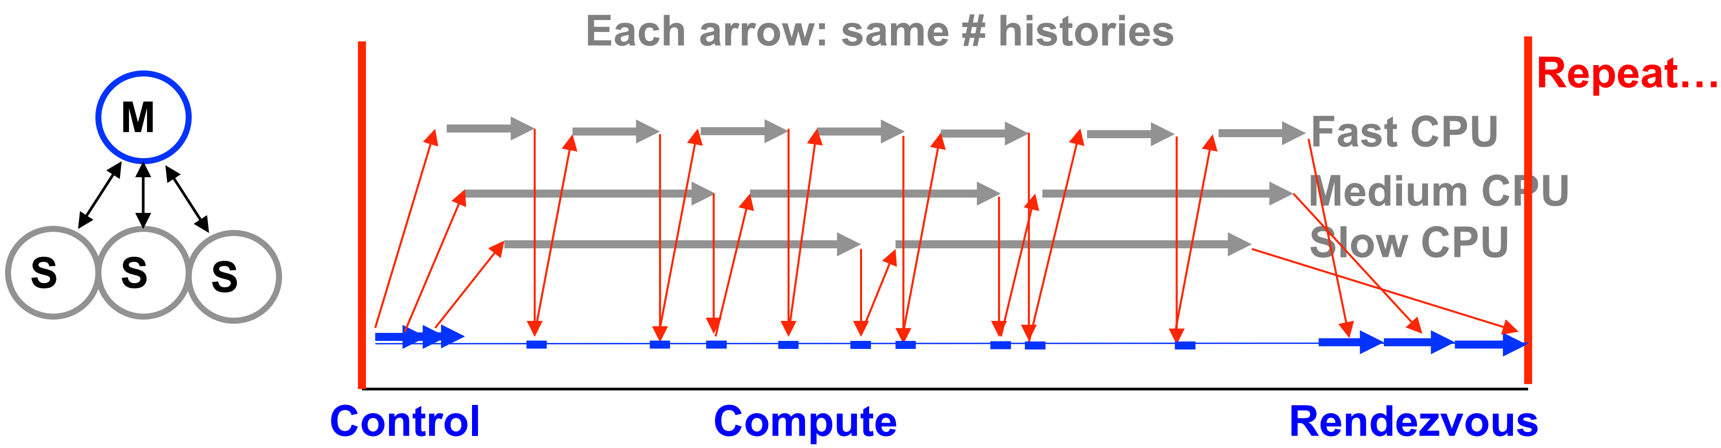

- **Load balancing:** Self-scheduling of histories on slaves
- **Fault tolerance:** Periodic rendezvous to save restart files
- **Parallel efficiency:** [compute time] / [compute + rendezvous time]

Parallel MC Performance Scaling

- **Scaling models, for master/slave with serial rendezvous**
  - "fixed" = constant number of histories/rendezvous, M (constant work)
  - "scaled" = M histories/slave per rendezvous, NM total (constant time)

Histories/rendezvous	**Speedup**
fixed	$S = N / (1 + cN^2)$
scaled	$S = N / (1 + cN)$

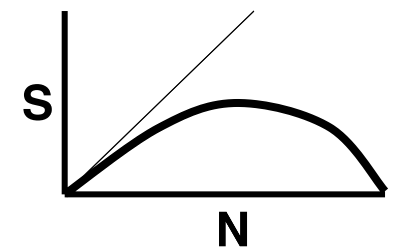

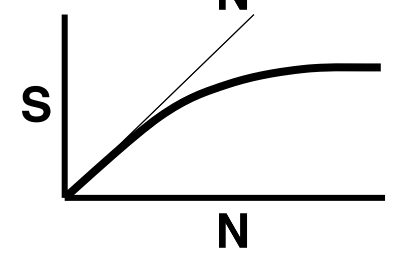

N = number of slaves

$c = (s + L/r) / T_1$

$T_1 \sim M$, more histories/rendezvous → larger T_1, smaller c

S+L/r, fixed, determined by number of tallies, ….

As M→infinity, c→0, S→N (limit for 1 rendezvous)

DOE Advanced Simulation & Computing – ASC

Blue Mountain – 3 TeraOps

(R.I.P.)

Lightning – 30 TeraOps

Q – 20 TeraOps

(R.I.P.)

Red Storm
Blue Gene/L
Hurricane
Turing
Cielo

Roadrunner – 1.3 PetaOps

[with Cell processors]

MCNP5 Parallel Scaled Speedup

mcnp Monte Carlo Codes XCP-3, LANL

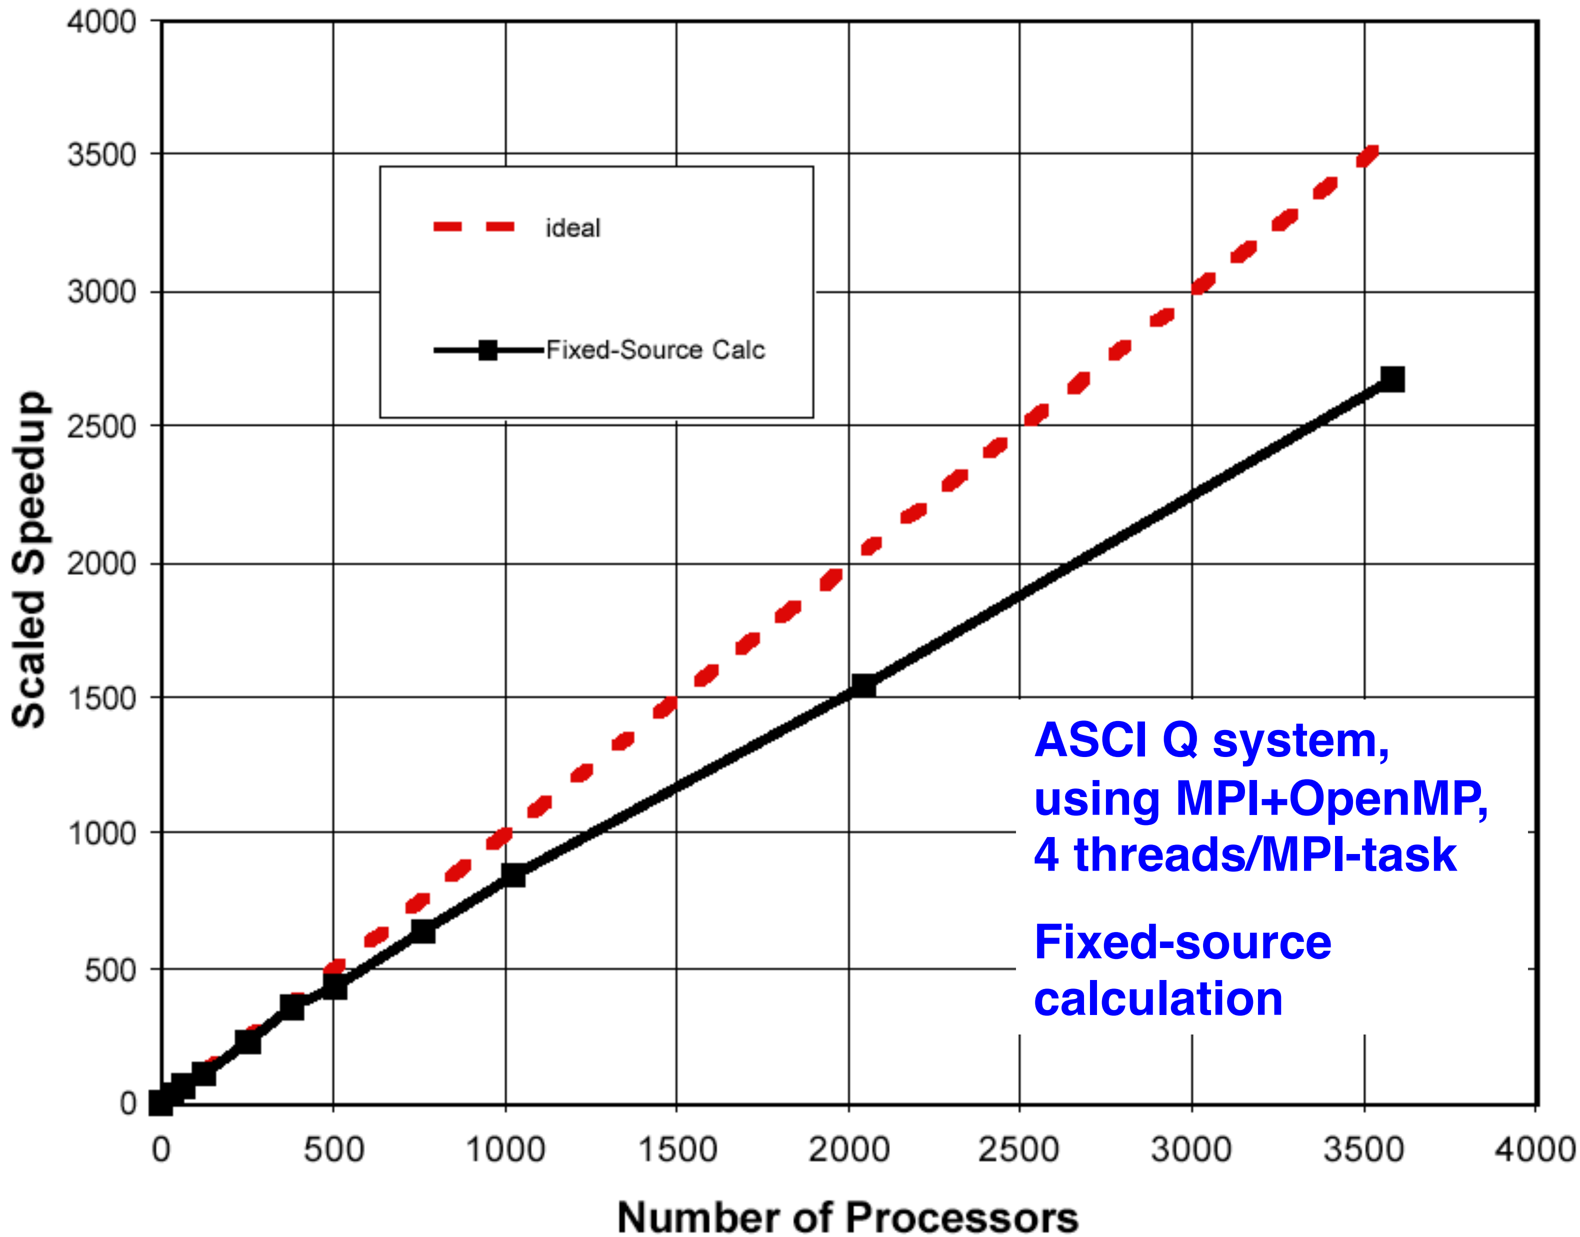

ASCI Q system, using MPI+OpenMP, 4 threads/MPI-task

Fixed-source calculation

MCNP5 Parallel Calculations

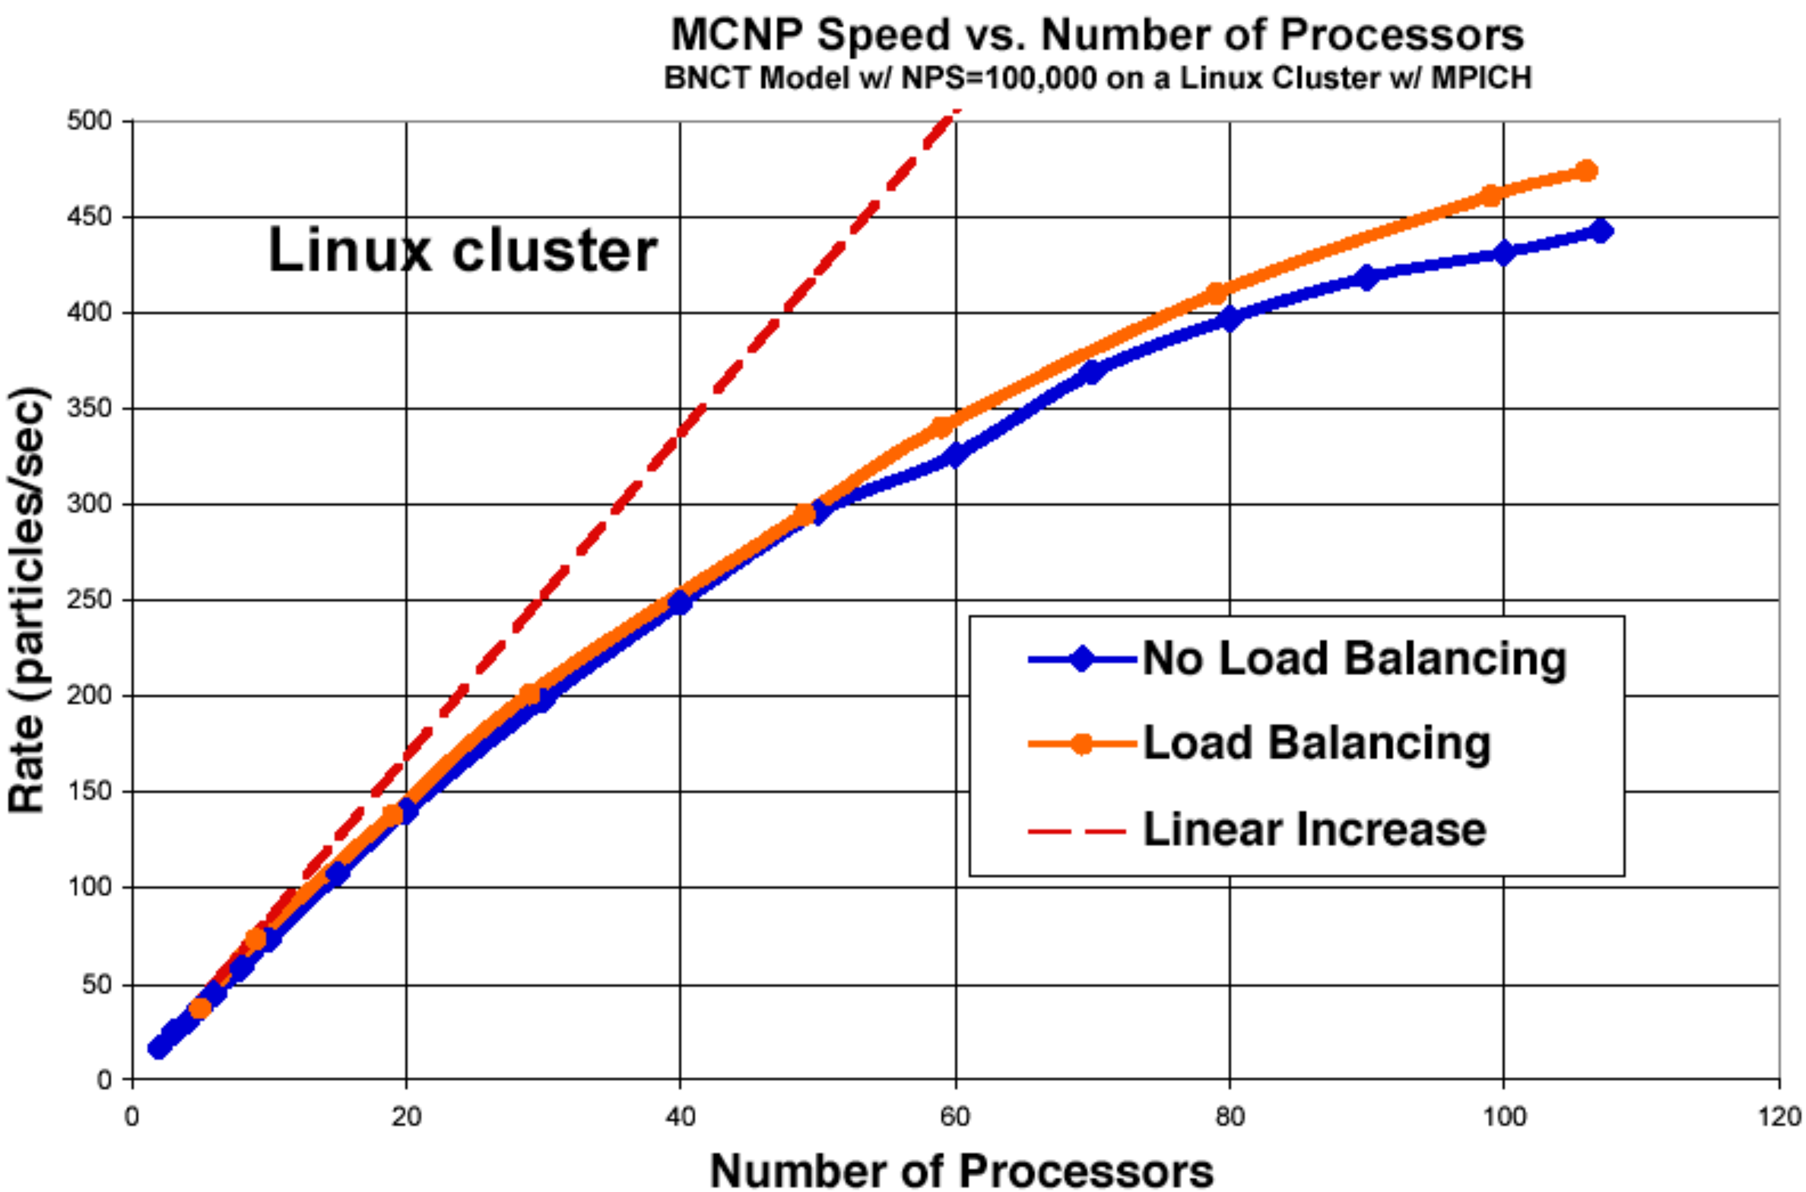

MCNP5 - Threading with OpenMP

- **MCNP5 performance - both serial & parallel - depends strongly on the Fortran-90 compiler & options used**
  - Runtime factors of 2-4x with <u>different compilers</u> on same hardware
  - Runtime factors of 2-4x with <u>different options</u> on same hardware & compiler

- **Parallel performance**
  - MCNP5 has always supported parallel calculations with **message-passing** (MPI) & **threading** (OpenMP)
  - Prior to mid-2006, Fortran compilers for Windows/Linux/Mac did a terrible job at threading. We recommended using only MPI.
  - Recently, using OpenMP threading with Intel compilers on Windows/Linux/Mac shows excellent speedups -- nearly 2x on dual-core, 3-4x on quad-core

MCNP5 - Threading on the Mac Pro

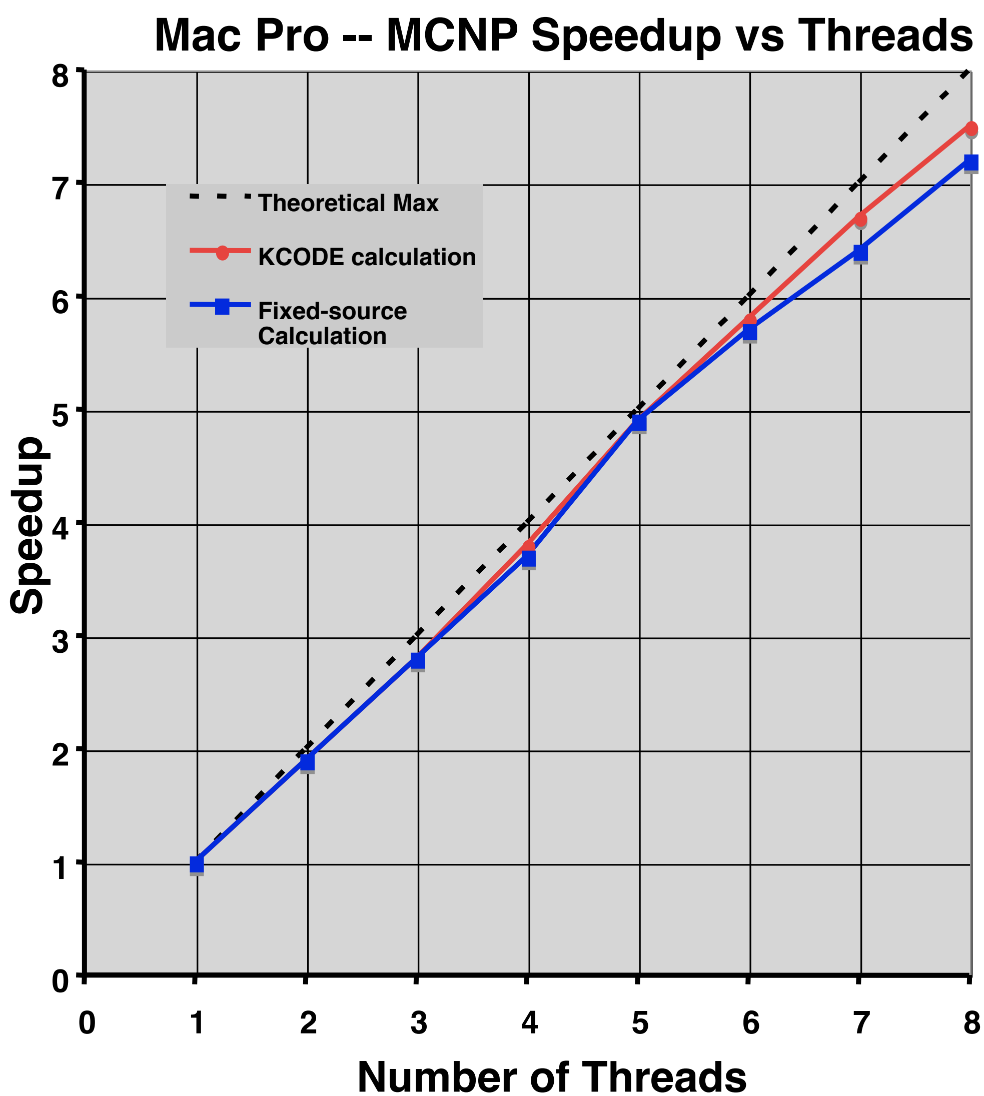

Hardware

- Mac Pro
- 2 x Quad-core Xeon
- 3GHz
- 8 GB memory

Software

- Mac OS X 10.4.11
- Intel F90, 10.0.017
  -O1 -openmp
- MCNP5 / 1.50

MCNP Calculations

- KCODE
  - BAWXI2 benchmark
  - kcode 5000 1 10 204
- Fixed-source
  - oil-well log, mode n
  - nps 500000

MCNP5 - Threading

Monte Carlo Codes
XCP-3, LANL

2009-03-03

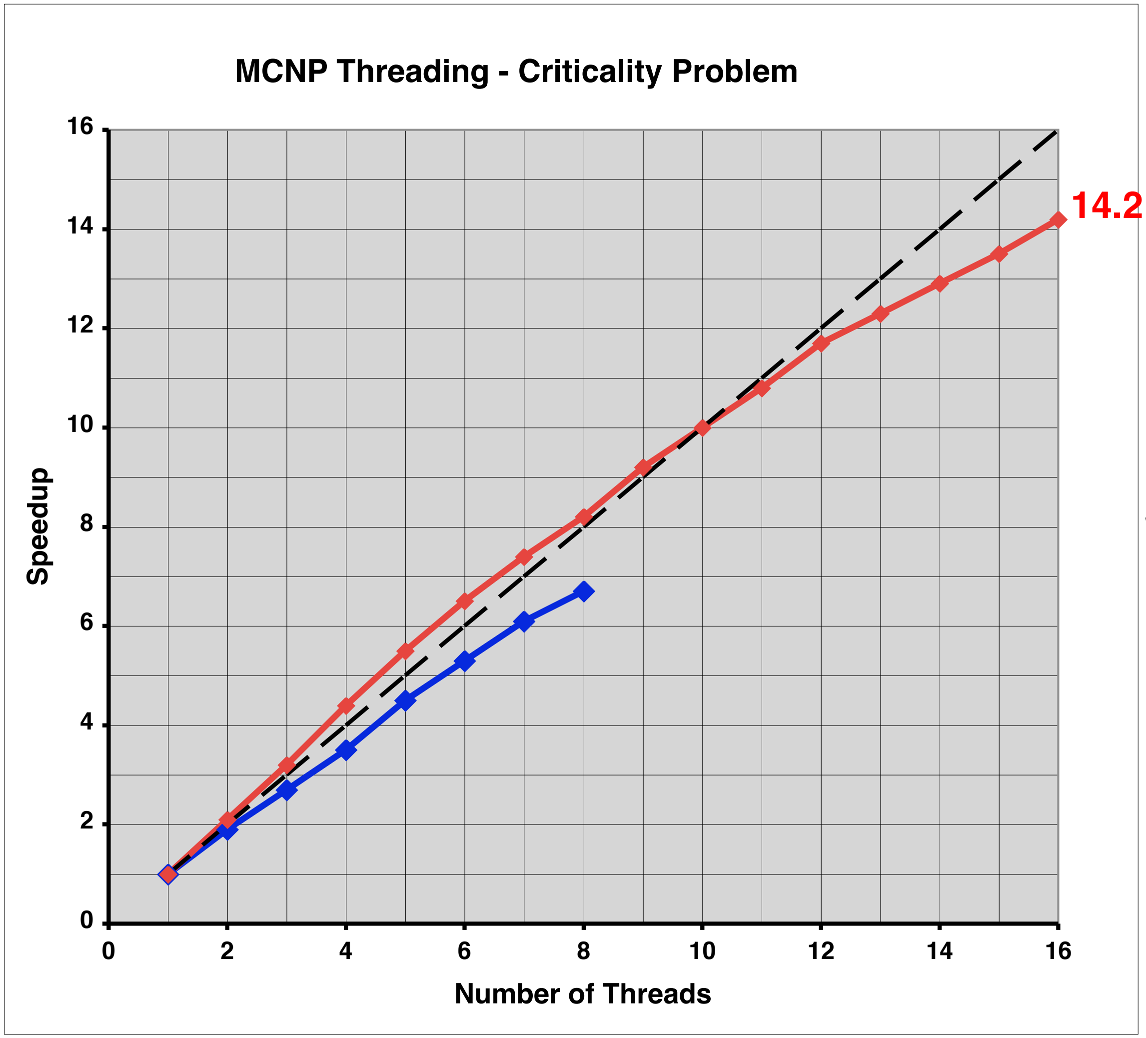

Hardware

- Lobo
  - 4 x Quad-core AMD Opteron
  - 2.2 GHz, 32 GB memory
- Mac Pro
  - 2 x Quad-core Intel Xeon
  - 3GHz, 8 GB memory

Software

- MCNP5-1.51
- Intel-10 F90, "-O1 -openmp"

MCNP Calculations

Criticality Calculation

BAWXI2 benchmark

kcode 25000 1 10 204

MCNP5 - Threading

2009-03-03

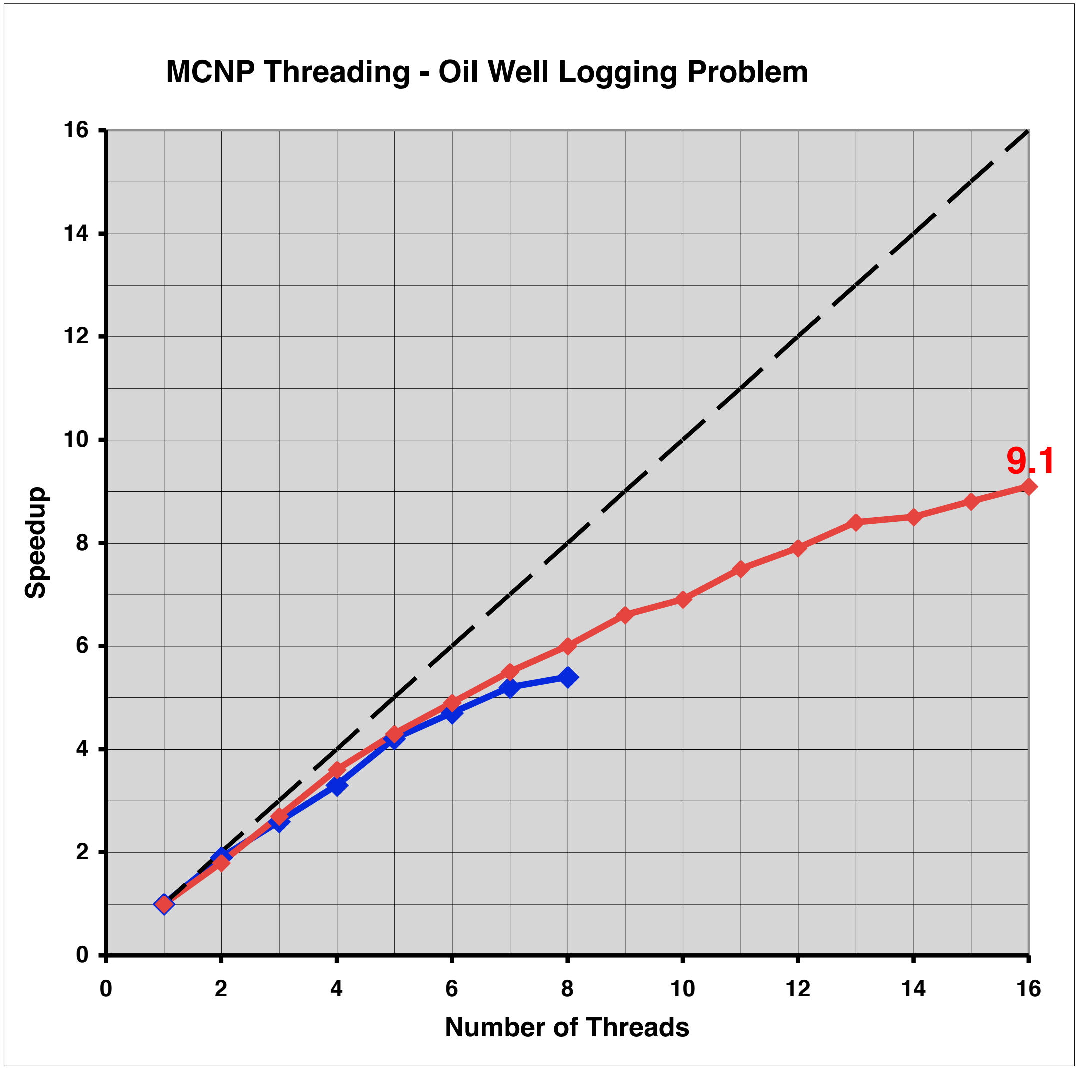

Hardware

- **Lobo**
  - 4 x Quad-core AMD Opteron
  - 2.2 GHz, 32 GB memory
- **Mac Pro**
  - 2 x Quad-core Intel Xeon
  - 3GHz, 8 GB memory

Software

- MCNP5-1.51
- Intel-10 F90, "-O1 -openmp"

MCNP Calculations

Oil Well Logging Calculation

inp12 benchmark

Nps 500000

Parallel MC Summary

Monte Carlo Codes
XCP-3, LANL

- **Master/slave algorithms work well**
  - Load-balancing: Self-scheduling
  - Fault-tolerance: Periodic rendezvous
  - Random numbers: Easy, with LCG & fast skip-ahead algorithm
  - Tallies: Use OpenMP "critical sections"
  - Scaling: Simple model, more histories/slave + fewer rendezvous
  - Hierarchical: Master/slave MPI, OpenMP threaded slaves
  - Portability: MPI/OpenMP, clusters of anything

- **Remaining difficulties**
  - Memory size: Entire problem must fit on each slave
    - Domain-decomposition has had limited success
      - Should be OK for reactor problems
      - May not scale well for shielding or time-dependent problems
      - For general 3D geometry, effective domain-decomposition is unsolved problem
    - Random access to memory distributed across nodes gives huge slowdown
      - May need functional parallelism with "data servers"

Parallel Processing For Large Monte Carlo Calculations

Domain Decomposition

If a Monte Carlo problem is too large to fit into memory of a single processor

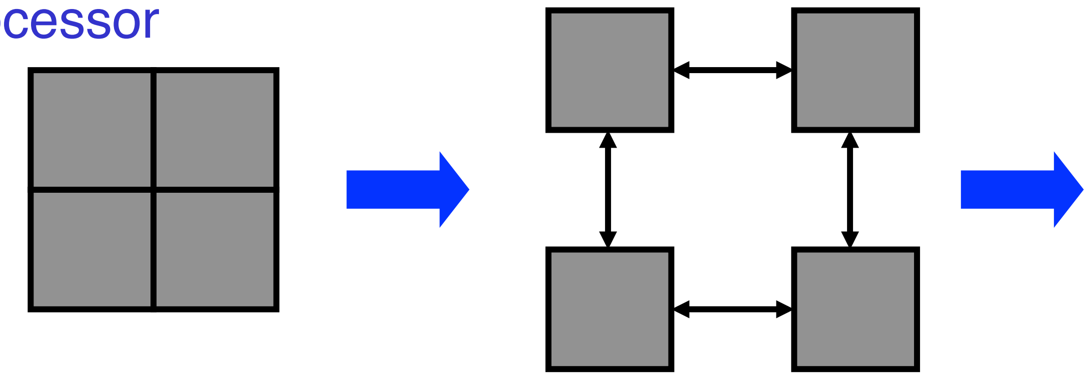

Decompose problem into spatial domains

Follow histories in each domain in parallel, move particles to new domains as needed

Collect Problem Results

– Need periodic synchronization to interchange particles among nodes
– Use message-passing (MPI) to interchange particles

➜ **Domain decomposition is often used when the entire problem will not fit in the memory of a single SMP node**

Parallel Monte Carlo

- **Inherent parallelism is on particles**
  - Scales well for all problems
- **Domain decomposition**
  - Spatial domains on different processors
  - Scales OK for Keff criticality calculations, where particle distribution among domains is roughly uniform
  - Does **not** scale for time-dependent problems due to severe load imbalances among domains
- **Domain decomposition - scaling with N processors**
  - **Best: performance ~ N (uniform distribution of particles)**
  - **Worst: performance ~ 1 (localized distribution of particles)**

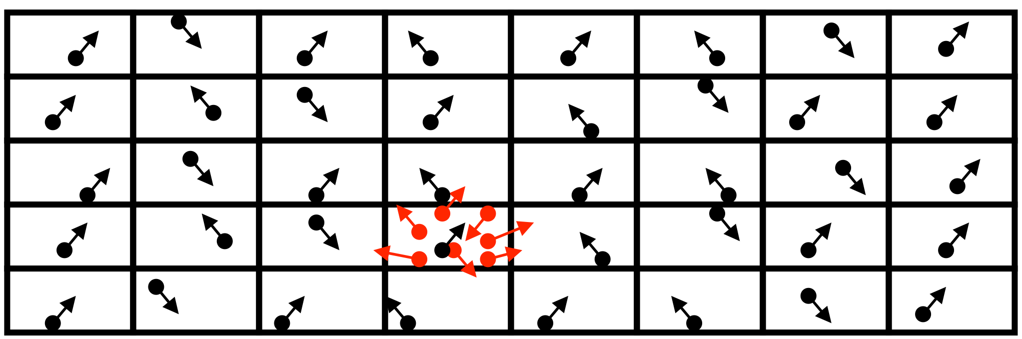

Parallel Monte Carlo

- Data is distributed by domain decomposition, but parallelism is on particles

- Solution ?

Parallel on particles + distributed data

- **Particle parallelism + Data Decomposition**
  - **Existing parallel algorithm for particles**
  - **Distribute data among processor nodes**
  - **Fetch the data to the particles as needed (dynamic)**
  - **Essentially same approach as used many years ago for CDC (LCM) or CRAY (SSD) machines**
  - **Scales well for all problems (but slower)**

Parallel Monte Carlo

- **Particle parallelism + data decomposition -- logical view:**

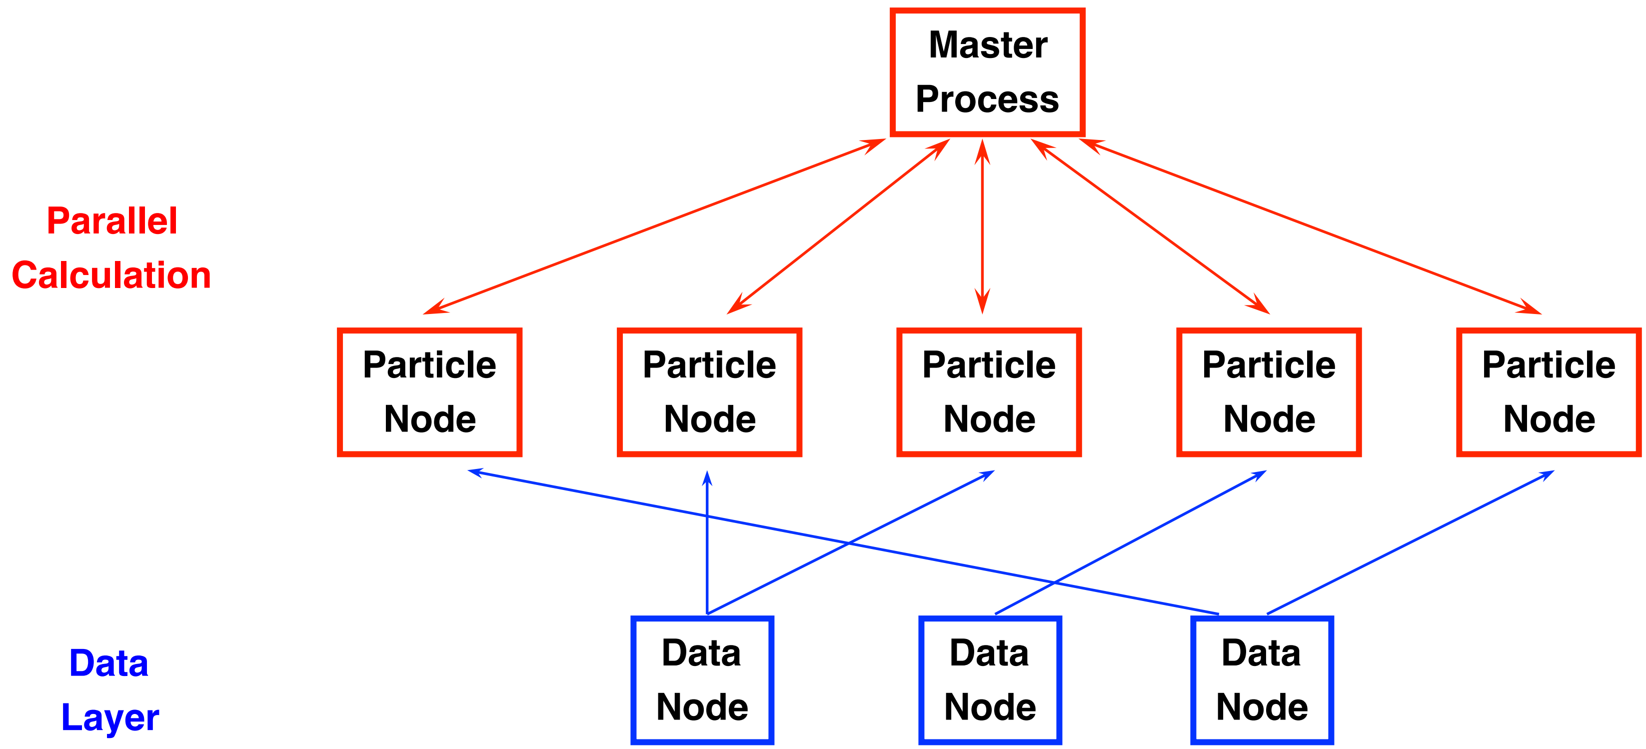

- **Mapping of logical processes onto compute nodes is flexible:**
  - Could map particle & data processes to **different** compute nodes
  - Could map particle & data processes to **same** compute nodes
- **Can replicate data nodes if contention arises**

Parallel Monte Carlo

- Particle parallelism + data decomposition

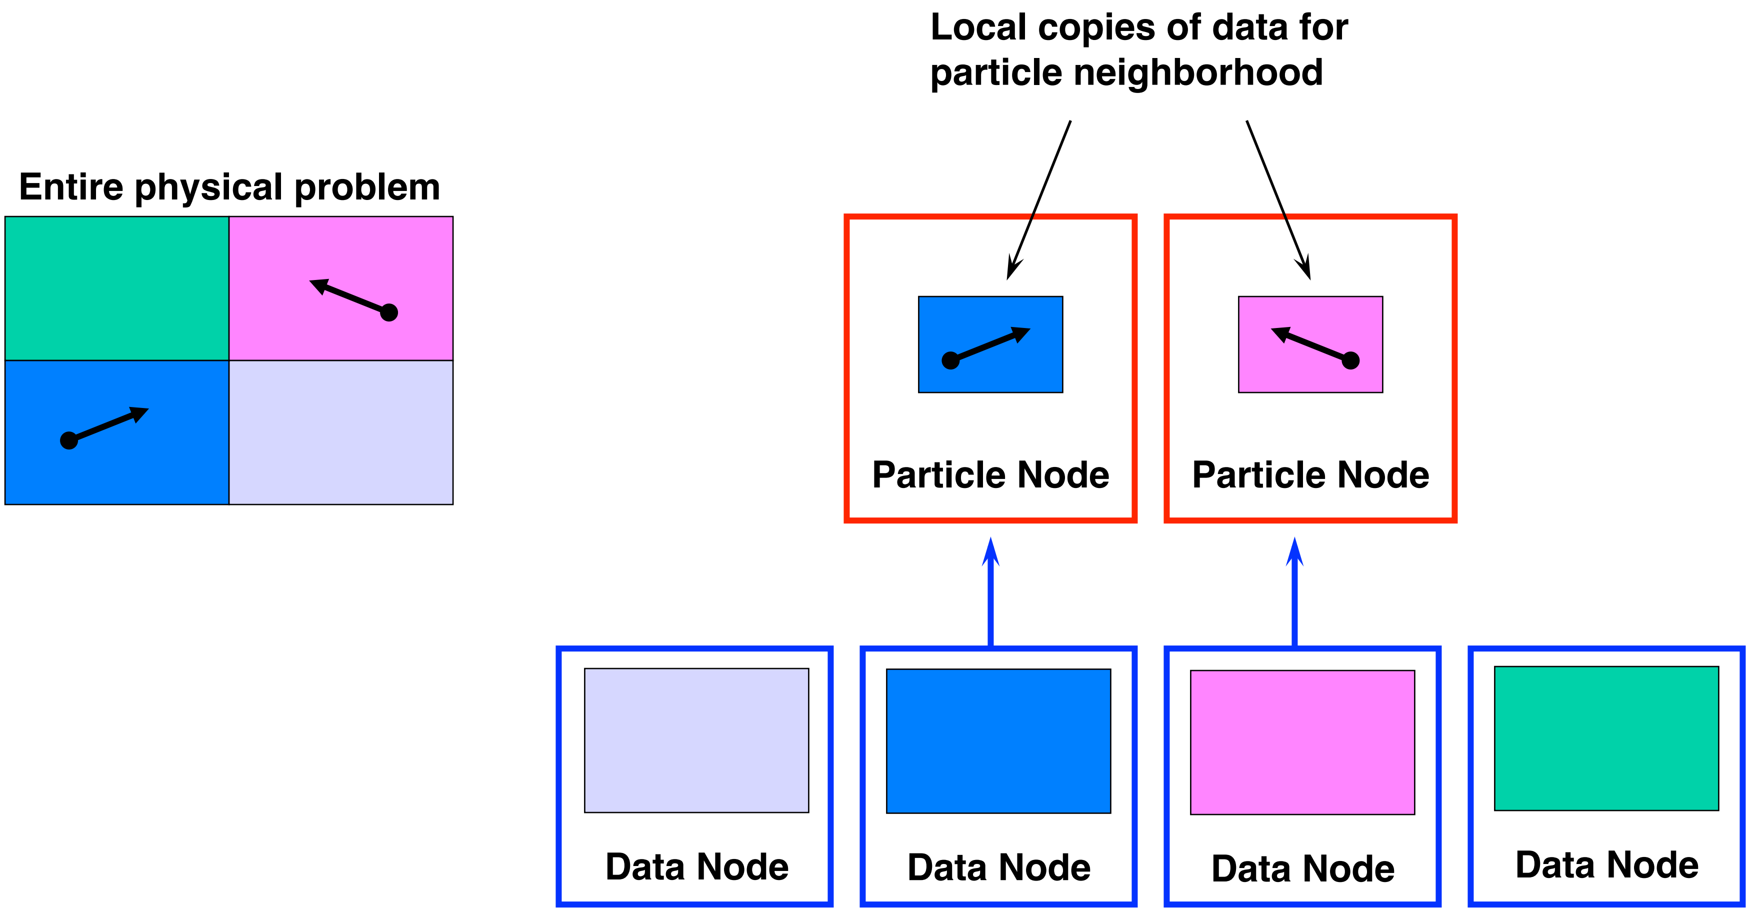

- **History modifications for data decomposition**

```
source

while   wgt > cutoff

.          compute distances & keep minimum:
.                    dist-to-boundary
.                    dist-to-time-cutoff
.                    dist-to-collision
.                    dist-to-data-domain-boundary

.          move particle
.          pathlength tallies

.          if   distance == dist-to-data-domain-boundary
.                    fetch new data

.          collision physics
.          roulette & split
. . .
```

Parallel Monte Carlo

- **Data windows & algorithm tuning**
  - Defining the "particle neighborhood" is an art
  - Anticipating the flight path can guide the pre-fetching of blocks of data
  - Tuning parameters:
    - How much data to fetch ?
    - Data extent vs. particle direction ?

Entire physical problem

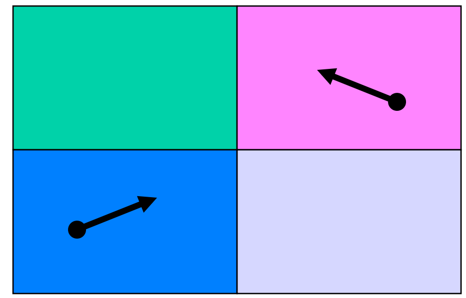

- **Examples**

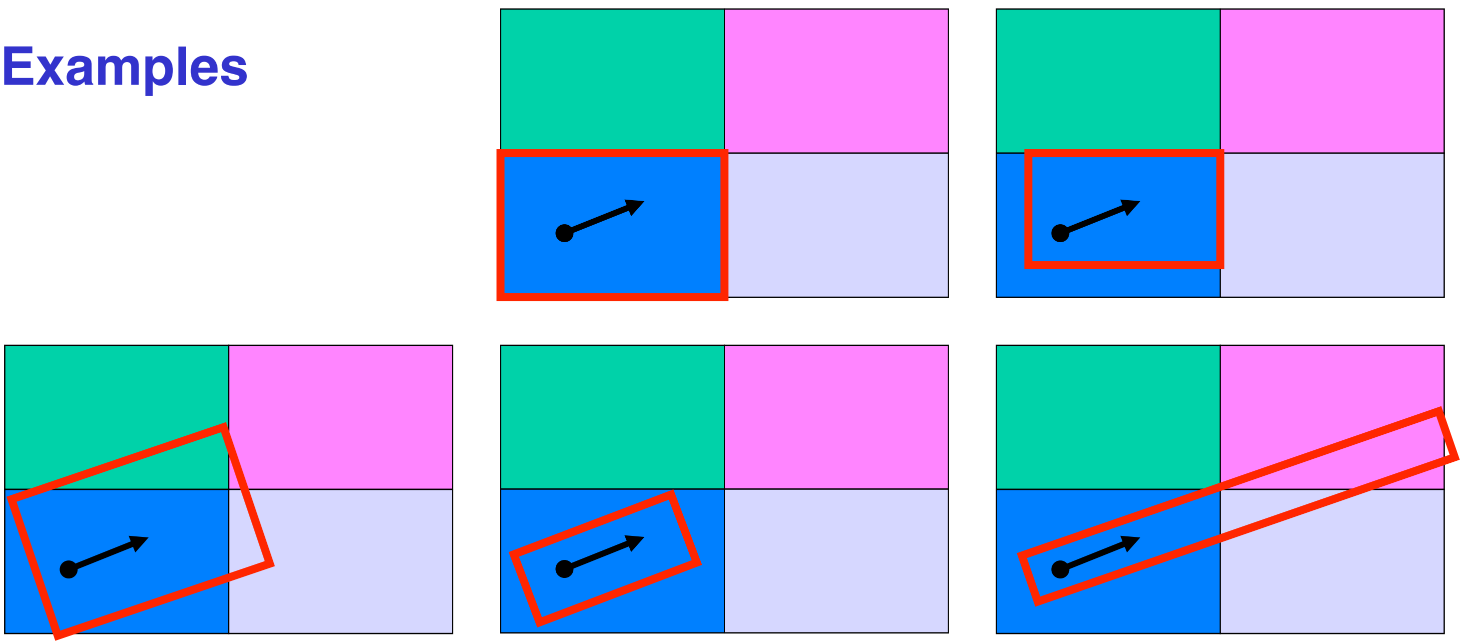

Conclusions

For Monte Carlo problems which can fit in memory:

- Concurrent scalar jobs - ideal for Linux clusters
- Master/slave parallel algorithm (replication) works well
  - Load-balancing: Self-scheduling
  - Fault-tolerance: Periodic rendezvous
  - Random numbers: Easy, with LCG & fast skip-ahead algorithm
  - Tallies: Use OpenMP "critical sections"
  - Scaling: Simple model, more histories/slave + fewer rendezvous
  - Hierarchical: Master/slave MPI, OpenMP threaded slaves
  - Portability: MPI/OpenMP, clusters of anything

For Monte Carlo problems too large to fit in memory:

- Spatial domain decomposition (with some replication) can work for some problems
- Particle parallelism + data decomposition is a promising approach which should scale for all problems