

AN INTRODUCTION TO THE MCNP CODE

by

J. Kenneth Shultis and Richard E. Faw

The MCNP Code, developed and maintained by Los Alamos National Laboratory, is the internationally recognized code for analyzing the transport of neutrons and gamma rays (hence NP for *neutral particles*) by the *Monte Carlo* method (hence MC). The code deals with transport of neutrons, gamma rays, and coupled transport, i.e., transport of secondary gamma rays resulting from neutron interactions. The MCNP code also deals with the transport of electrons, both primary source electrons and secondary electrons created in gamma-ray interactions.

As of this writing, the MCNP code is distributed as version 5. The instruction manual is distributed by the Radiation Safety Information Computational Center at Oak Ridge National Laboratory.

This tutorial document highlights certain aspects of the MCNP input code. Users are expected to have access to the MCNP instruction manuals. With version 5, the enormous MCNP manual has been split into 3 volumes. Volume I gives an overview (Ch. 1) and theory (Ch. 2) of the code. Volume II is the *User's Guide* that defines all the commands and options of the code (Ch. 3), gives many examples (Ch. 4), and describes the code's output (Ch. 5). Volume III is a *Developer's Guide* that gives many of the technical details of code and are needed by only the MCNP experts. Some of the notation used in the MCNP documentation uses historical terminology. For example, the term *card*, historically a punched card, should be interpreted as a line of the input file.

For the novice user, Ch. 1 of Vol. I of the manual presents an overview of MCNP that summarizes the preparation of input files, the execution of the code, and the interpretation of results. This is highly recommended reading. After gaining some experience with MCNP, the beginning user should periodically browse through the remainder of Vol. I to obtain a better understanding of the theory behind the many features of MCNP.

Volume II is essential for both the novice and expert user. This is the documentation that formally defines all the commands and options that make MCNP such a powerful radiation transport code. In this primer there are many margin notes indicating the pages in Vol. II that discuss in more detail the subject being presented in this primer.

The MCNP documentation is very comprehensive; thus, it is difficult for new users of the code to distinguish between information essential to learning how to use the code and information needed only under very specialized circumstances. For this reason, this tutorial document was prepared to introduce the novice with the more basic (and essential) aspects of the MCNP code.

1 Structure of the MCNP Input File

An input file has the following form:

```
Message Block {optional}
    blank line delimiter {optional}
One Line Problem Title Card
Cell Cards [Block 1]
    blank line delimiter
Surface Cards [Block 2]
    blank line delimiter
Data Cards [Block 3]
    blank line terminator {optional}
```

A more detailed listing of the elements of the MCNP input file is given in Ch. 3.V starting on page 3-154 of the MCNP5 manual.

1.1 Annotating the Input File

It is good practice to add liberally comments to an input MCNP file so that it is easier for you and others to understand what problem is addressed and the tricks used. A *comment line* begins with *C* or *c* followed by a space. Such a line is ignored by MCNP. Alternatively, anything following a *\$* sign on a line is ignored. See Fig. 4 on page 27 for a well-annotated MCNP input file.

2 Geometry Specifications

The specification of problem geometry is treated in several sections of the MCNP manual. In Vol. I, beginning on page 1-12, there is an introduction to geometric specification. Discussion continues in Ch. 2 Sec. II (page 2-7). Sections II and III of Ch. 3 provide detailed instructions on preparation of problem input cards and Finally, Ch. 4 Sec. I provides examples of geometry specifications.

MCNP treats problem geometry primarily in terms of regions bounded by first and second degree surfaces. Cells are defined by intersections, unions, and complements of the regions, and contained user defined materials. Intersections and unions are described in an accompanying illustration.

The union operation may be thought of as a logical OR, in that the union of A and B is a new region containing all space either in region A OR region B. The intersection operation may be thought of as a logical AND, in that the result is a region that contains only space common to both A AND B. The complement operator *#* plays the roll of a logical NOT. For example *#* (A:B) represents all space outside the union of A and B.




Figure 1. Left: the union $A:B$ or “A or B”. Right: the intersection $A B$ or “A and B”.

MCNP uses a 3-dimensional (x, y, z) Cartesian coordinate system. All dimensions are in centimeters (cm). All of the space is composed of contiguous volumes or *cells*. Each cell is bounded by a surface, multiple surfaces, or by infinity. For example, a cube is bounded by six planes. Every (x, y, z) point must belong to a cell (or be on the surface of a cell). There can be no “gaps” in the geometry, i.e., there can be no points that belong to no cell or surface. Every cell and surface is given by the user a unique numerical identifier.

2.1 Surfaces – Block 2

Table 3.1, taken from the MCNP instruction manual, lists the surfaces used by MCNP to create the geometry of a problem. All refer to a Cartesian coordinate system. A surface is represented functionally as $f(x, y, z) = 0$. For example, for a cylinder of radius R parallel to the z -axis is defined as $f(x, y, z) = (x - \bar{x})^2 + (y - \bar{y})^2 - R^2$, where the cylinder’s axis is parallel to the z -axis and passes through the point $(\bar{x}, \bar{y}, 0)$. The MCNP input line for such a surface uses the symbol *C/Z* (or *c/z*, since MCNP is case insensitive) is

```
1 C/Z 5 5 10 $ a cylindrical surface parallel to z-axis
```

defines surface 1 as an infinitely long cylindrical surface parallel to z -axis with radius 10 cm and whose axis passes through the point $(x = 5 \text{ cm}, y = 5 \text{ cm}, z = 0)$. Note that the length of the surface along the z axis is infinite. Note also the in-line comment, introduced by the *\$* symbol.

Table 1. MCNP Surface Cards

Mnemonic	Type	Description	Equation	Card Entries
P PX PY PZ	plane	general normal to x -axis normal to y -axis normal to z -axis	$Ax + By + Cz - D = 0$ $x - D = 0$ $y - D = 0$ $z - D = 0$	$A B C D$ D D D
SO S SX SY SZ	sphere	centered at origin general centered on x -axis centered on y -axis centered on z -axis	$x^2 + y^2 + z^2 - R^2 = 0$ $(x - \bar{x})^2 + (y - \bar{y})^2 + (z - \bar{z})^2 - R^2 = 0$ $(x - \bar{x})^2 + y^2 + z^2 - R^2 = 0$ $x^2 + (y - \bar{y})^2 + z^2 - R^2 = 0$ $x^2 + y^2 + (z - \bar{z})^2 - R^2 = 0$	R $\bar{x} \bar{y} \bar{z} R$ $\bar{x} R$ $\bar{y} R$ $\bar{z} R$
C/X C/Y C/Z CX CY CZ	cylinder	parallel to x -axis parallel to y -axis parallel to z -axis on x -axis on y -axis on z -axis	$(y - \bar{y})^2 + (z - \bar{z})^2 - R^2 = 0$ $(x - \bar{x})^2 + (z - \bar{z})^2 - R^2 = 0$ $(y - \bar{y})^2 + (y - \bar{y})^2 - R^2 = 0$ $y^2 + z^2 - R^2 = 0$ $x^2 + z^2 - R^2 = 0$ $x^2 + y^2 - R^2 = 0$	$\bar{y} \bar{z} R$ $\bar{x} \bar{z} R$ $\bar{x} \bar{y} R$ R R R
K/X K/Y K/Z KX KY KZ	cone	parallel to x -axis parallel to y -axis parallel to z -axis on x -axis on y -axis on z -axis	$\sqrt{(y - \bar{y})^2 + (z - \bar{z})^2} - t(x - \bar{x}) = 0$ $\sqrt{(x - \bar{x})^2 + (z - \bar{z})^2} - t(y - \bar{y}) = 0$ $\sqrt{(x - \bar{x})^2 + (y - \bar{y})^2} - t(z - \bar{z}) = 0$ $\sqrt{y^2 + z^2} - t(x - \bar{x}) = 0$ $\sqrt{x^2 + z^2} - t(y - \bar{y}) = 0$ $\sqrt{x^2 + y^2} - t(z - \bar{z}) = 0$	$\bar{x} \bar{y} \bar{z} t^2 \pm 1$ $\bar{x} \bar{y} \bar{z} t^2 \pm 1$ $\bar{x} \bar{y} \bar{z} t^2 \pm 1$ $\bar{x} t^2 \pm 1$ $\bar{y} t^2 \pm 1$ $\bar{z} t^2 \pm 1$ ± 1 used only for 1-sheet cone
SQ	ellipsoid hyperboloid paraboloid	axis parallel to x -, y -, or z -axis	$A(x - \bar{x})^2 + B(y - \bar{y})^2 + C(z - \bar{z})^2 + 2D(x - \bar{x}) + 2E(y - \bar{y}) + 2F(z - \bar{z}) + G = 0$	$A B C D E$ $F G \bar{x} \bar{y} \bar{z}$
GQ	cylinder, cone ellipsoid paraboloid hyperboloid	axis not parallel to x -, y -, or z -axis	$Ax^2 + By^2 + Cz^2 + Dxy + Eyz + Fzx + Gz + Hy + Jz + K = 0$	$A B C D E$ $F G H J K$
TX TY TZ	elliptical or circular torus. Axis is parallel to x -, y -, or z -axis		$(x - \bar{x})^2/B^2 + (\sqrt{(y - \bar{y})^2 + (z - \bar{z})^2} - A)^2/C^2 - 1 = 0$ $(y - \bar{y})^2/B^2 + (\sqrt{(x - \bar{x})^2 + (z - \bar{z})^2} - A)^2/C^2 - 1 = 0$ $(z - \bar{z})^2/B^2 + (\sqrt{(x - \bar{x})^2 + (y - \bar{y})^2} - A)^2/C^2 - 1 = 0$	$\bar{x} \bar{y} \bar{z} A B C$ $\bar{x} \bar{y} \bar{z} A B C$ $\bar{x} \bar{y} \bar{z} A B C$
XYZP		surfaces defined by points – see pages 3-15 to 3-17		

Every surface has a “positive” side and a “negative” side. These directional senses for a surface are defined formally as follows: any point at which $f(x, y, z) > 0$ is located in the positive sense (+) to the surface, and any point at which $f(x, y, z) < 0$ is located in the negative sense (-) to the surface. For example, a region within a cylindrical surface is negative with respect to the surface and a region outside the cylindrical surface is positive with respect to the surface.

2.2 Cells – Block 1

We illustrate how surfaces and Boolean logic are used to define cells by considering a simple example of a cylindrical storage cask whose wall and ends are composed of iron 1-cm thick. Inside and outside the cask are void regions. Suppose the outer cylindrical surface is that used in the illustration in the previous section. The geometry for this problem is shown in Fig. 2.




Figure 2. Geometry for a simply cask. Numbers in triangles are surface identification numbers and numbers in circles define the cell identification number.

To define the inside surface of the cask, we need another cylinder inside and concentric with the first cylinder but with a radius smaller by 1 cm. We shall call this smaller cylindrical surface number 4, so that the surface definition lines in the input file for these two cylinders are

```

1  C/Z  5 5 10    $ outer cylindrical surface
4  C/Z  5 5 9     $ inner cylindrical surface

```

To define the base and top of the cask, planes perpendicular to the z -axis are needed at locations $z = 40$ cm and $z = 60$ cm, respectively. Similarly, to define the base and top of the inner cavity of the cask two more planes perpendicular to the z -axis are needed at $z = 41$ cm and $z = 59$ cm. These four planes are defined by

```

2  PZ   40        $ base of cask
3  PZ   60        $ top of cask
5  PZ   41        $ base of inner cavity
6  PZ   59        $ top of inner cavity

```

These six surface definition cards (or input lines) can appear in any order in Block 2 of the input file.

With the problem surfaces defined, we now begin to define the columns or cells which must fill all (x, y, z) space. These cell definition cards comprise the content of Block 1 of the input file. First,

we define the inner void of the cask as cell 8. This volume is negative with respect to surface 4, positive with respect to the plane 5, and negative with respect to plane 6. Thus, cell 8 is defined as

```
8 0 -4 5 -6 IMP:N=0 IMP:P=1 $ inner cask void
```

The first number on a cell definition card is the cell number (arbitrarily picked by the user). Here the second entry 0 denotes that the cell is filled by a void, and -4 5 -6 indicate that all points in cell 8 are inside the cylinder 4 AND are above plane 5 AND are below plane 6. region. The last two IMP specifications define the *importance* of this region to neutrons (N) and (P). Neutrons in this cell have zero weight and photons have unit weight (e.g., for a photon transport problem). We'll discuss importances later. The order of surfaces in an intersection string is immaterial. Thus, we could have defined cell 8 by intersection of surfaces -6 -4 5.

Now consider the iron shell of the cask. Suppose this cell is number 7 in a sequence and consists of material 5, as yet to be defined, with density 7.86 g/cm³. Space within this cell is negative with respect to surface 1, positive with respect to surface 2 and negative with respect to surface 3 AND also cannot be inside the void or cell 8. This cell can thus be define as

```
7 5 -7.86 -1 2 -3 #8 IMP:N=0 IMP:P=1 $ iron cask shell
```

Although the *complement operator* # (for NOT) is often a convenient way to exclude an inner region, this operator often reduces the efficiency of MCNP. In fact, one theoretically never has to use #. The region outside cell 8 can be defined by the union string (4:6:-5) and the definition of cell 7 can equivalently defined as

```
7 5 -7.86 -1 2 -3 (4:6:-5) IMP:N=0 IMP:P=1 $ iron cask shell
```

Now suppose that cells 7 and 8 describe all space of interest for radiation transport. In other words, suppose that all photons passing outside the outer surface of the finite cylinder may be killed, i.e., may have path tracking ended. One needs to define this space and to require that within this space the photon importance is zero. This is done by specifying a cell, say cell 9, that is the union of all regions positive with respect to surfaces 1 and 3 and negative with respect to surface 2, that is,

```
9 0 1:3:-2 IMP:N=0 IMP:P=0 $ graveyard
```

This could also be defined using the complement operator, that is, by stating that the kill zone is all space outside the union of cells 7 and 8, namely

```
9 0 #(7:8) IMP:N=0 IMP:P=0 $ graveyard
```

Note that the second entry on this cell card is zero, indicating a vacuum and that the photon importance is set to zero.

3 Data Specifications – Block 3

This section describes how one describes the physical properties of materials in MCNP calculations, how one specifies the radiation source, how one specifies the scoring of results (tallies) and how one specifies the “physics of radiation transport calculations to be performed. In short, this third input block provides almost all problem specifications other than the geometry

An introduction to Block 3 commands is provided Ch. 1 Secs. I and II of the MCNP manual. More detail on the theory is provided in Ch. 2 Secs. III through V. Section IV of Ch. 3 provide detailed instructions on preparation of problem input “cards and Ch. 4 Secs. IV and V provides examples of source and tally treatments.

3.1 Materials Specification

Specification of materials filling the various cells in an MCNP calculation involves the following elements: (a) defining a unique material number, (b) the elemental (or isotopic) composition, and (c) the cross section compilations to be used.

Sec. 3.F
pp. 3-114 to
3.124

Note that density is not specified here. Instead, density is specified on the cell definition card. This permits one material to appear at different densities in different cells. Suppose that the first material to be identified in problem input is (light) water and that only gamma-ray transport is of interest. Comment cards (cards beginning with C or c) may be used for narrative description. In the following card images, the designation M1 refers to material 1. For a compound, unnormalized atomic fractions may be used. For example,

```

c -----
c WATER for gamma-ray transport (by atom fraction)
c -----
M1  1000    2    $ elemental H and atomic abundance
     8000    1    $ elemental O and atomic abundance

```

The designations 1000 and 8000 identify elemental hydrogen, atomic number $Z = 1$, and elemental oxygen ($Z = 8$). The three zeros in each designation are place holders for the atomic mass number, which would be required to identify specific isotopes of the element and which, generally, are required for neutron transport, as described later. For gamma ray and electron transport, one need only specify the atomic number. For compounds or mixtures, composition may be specified by mass fraction as indicated by minus sign, as follows

```

c -----
c WATER for gamma-ray transport (by mass fraction)
c -----
M1  1000  -0.11190  $ elemental H mass fraction
     8000  -0.88810  $ elemental O mass fraction

```

Error/warning messages can be avoided by assuring that mass/atomic fractions sum to unity.

For neutron transport problems, often a specific isotope of an element must be specified. The isotope Z Aid number ($Z A$ Identification) contains six digits ZZZAAA in which ZZZ is the atomic number Z and AAA is the atomic mass number A . Thus ^{235}U has a Z Aid number 092235 or simply 92235. If neutron cross sections for an element composed of its isotopes in their naturally occurring abundances is desired, then the Z Aid is specified as ZZZ000. Note, such elemental neutron cross section sets are not available for all elements. Often you must list all the important isotopes. As an example, light water for neutron problems could be defined as

```

c -----
c WATER for gamma-ray transport (by mass fraction)
c   (ignore H-2, H-3, O-17, and O-18)
c -----
M1  1001.50c -0.11190  $ H-1 and mass fraction
     8016.50c -0.88810  $ O-16 and mass fraction

```

Here 1001 and 8016 provide atomic number and atomic mass designations, in the form of the Z Aid numbers. The .50c designations identify particular cross section compilations and are described later.

When hydrogen is molecularly bound in water, either pure or as a constituent in some other material, the binding affects energy loss in collisions experienced by slow neutrons. For this reason, special cross-section data treatments are provided to take binding effects into account. To require this special treatment, an additional card of type MT is required,

```

c -----
c WATER for gamma-ray transport (by mass fraction)
c   (ignore H-2, H-3, O-17, and O-18)
c -----
M1  1001.50c -0.11190  $ H-1 and mass fraction
     8016.50c -0.88810  $ O-16 and mass fraction
MT1 lwtr.01

```

Without the MT card, hydrogen is treated as if it were a monatomic gas. Treatment of binding effects for other nuclides and other materials is described in Appendix G of Vol. I of the MCNP manual.

In Appendix A, many examples of material specifications are given.

3.2 Cross-Section Specification

Neutron reactions and cross-section data tables are described in Section III of Chapter 2 in the MCNP manual. A comprehensive list of data tables is provided in Table G2 of Appendix G (part of Vol. I). Specification of a particular cross-section data set depends somewhat on the nature of the problem to be solved and somewhat on the data available to the user. Not all data are available to all users. For users obtaining data through the Radiation Safety Information Computation Center (RSICC), a common choice would be the .50c cross sections available in the rmccs library and are taken from the ENDF/B-V evaluated nuclear data files.

In some instances, cross sections are available for elements with naturally occurring atomic abundances. For example, natural chromium would be identified by ZAID 24000.60c. However, cross sections for the isotope ^{50}Cr would require the identification 24050.60c and would require the endf60 library containing data from the ENDF/B-VI evaluated nuclear data file. The ENDF/B-VI cross sections are included with the MCNP-5 distribution package. In some instances, it is necessary for the user to define a natural element as a combination of isotopes. This is the case for many of the light elements such as helium and lithium.

3.3 Source Specifications

The source and type of radiation particles for an MCNP problem is defined by the SDEF command. The SDEF command has many variables or parameters that are used to define *all* the characteristics of all sources in the problem. The SDEF command with its many variables is one of the more complex MCNP commands and is capable of producing an incredible variety of sources – all with a single SDEF command. And only one SDEF card is allowed in an input file.

Sec. 3.D pp.
3-51 to 3-77

On the SDEF line values of the variables in Table 2 are entered, if other than the default values, that are needed to characterize the source. The = sign is optional, i.e., PAR=1 is equivalent to PAR 1. Values of variables can be specified at three levels: (1) explicitly (e.g., ERG=1.25), (2) with a *distribution* number (e.g., ERG=d5), and (3) as a function of another variable (e.g., ERG=Fpos). Specify variables at levels 2 and 3 requires the use of three other source cards: the SI (source information) card, the SP (source probabilities) card, and the SB (source bias) card.

Chapter 3 Sec. D gives a complete description of the SDEF command and the use of its variables. This is a very tersely written section and is difficult for the novice to understand all the features. As an MCNP user gains experience, this section should be periodically reread. Each reading almost always provides new insight and understanding of the SDEF command.

Chapter 4 of the MCNP manual has several examples of complicated sources. These are well worth studying. However, we often need fairly simple sources and such examples are not provided in the MCNP manual. It takes many readings of the few pages in Chapter 3 describing all the source commands and options to sometimes see how to do something fairly simple. Below are a few examples of fairly simple source definitions that may help you to understand better how to specify sources for MCNP.

When developing a new source definition, always check and recheck that source particles are truly being generated where you think they should be. HINT: Always use the VOID card and the PRINT 110 statement somewhere in block 3 of the input file. The PRINT 110 causes the starting locations, directions, and energies of the first 50 particles to be printed to the output file. Examine this output table to convince yourself that particles are being generated as you expect.

Table 2. Source variables for the SDEF command


Variable	Meaning	Default
CEL	cell	determined from XXX, YYY, ZZZ and possibly UUU, VVV, WWW
SUR	surface	0 (means cell source)
ERG	energy (MeV)	14 MeV
DIR	μ , the cosine of the angle between VEC and UUU, VVV, WWW. The azimuthal angle is always sampled uniformly in $[0, 2\pi]$	Volume case: μ is sampled uniformly in $[-1.1]$ (isotropic). Surface case: $p(\mu) = 2\mu$ for $\mu \in [0, 1]$ (cosine distribution).
VEC	reference vector for VEC	Volume case: required unless isotropic. Surface case: vector normal to the surface with sign determined by NRM.
NRM	sign of the surface normal	+1
POS	reference point for positioning sampling	0, 0, 0
RAD	radial distance of the position from POS or AXS	0
EXT	Cell case: distance from POS along AXS. Surface case: cosine of angle from AXS	0
AXS	reference vector for EXT and RAD	no direction
X	x-coordinate of position	no X
Y	y-coordinate of position	no Y
Z	z-coordinate of position	no Z
CCC	cookie-cutter cell	no cookie-cutter cell
ARA	area of surface (required only for direct contributions to point detectors from a plane surface source)	none
WGT	particle weight	1
EFF	reference efficiency criterion for position sampling	0.01
PAR	type of particle source emits	= 1 (neutron) if MODE N or P or N P E; = 2 (photon) if MODE P; = 3 (electron) if MODE E

3.3.1 Point Isotropic Sources

Two Point Isotropic Sources at Different Positions

```


c ----- Source: two point isotropic 1-MeV photon sources on x-axis
SDEF ERG=1.00 PAR=2 POS=d5 $ energy, particle type, location
SI5 L -10 0 0 10 0 0 $ (x,y,z) coords of the two pt sources
SP5 .75 .25 $ relative strengths of each source
  
```



Point Isotropic Source with Discrete Energy Photons

```


c ----- Source: point isotropic source with 4 discrete photon energies
SDEF POS 0 0 0 ERG=d1 PAR=2
SI1 L .3 .5 1. 2.5 $ the 4 discrete energies (MeV)
SP1 .2 .1 .3 .4 $ frequency of each energy
  
```



Point Isotropic Source with a Histogram of Energies

```


c ----- source: point isotropic with 4 histogram energy bins
SDEF POS 0 0 0 PAR=2 ERG=d1 $ position, particle type, energy
SI1 H .1 .3 .5 1. 2.5 $ histogram boundaries
SP1 D 0 .2 .4 .3 .2 $ probabilities for each bin
  
```



Point Isotropic Source with a Continuum of Energies

```


c ----- source: point isotropic with Maxwellian energy spectrum
SDEF POS 0 0 0 PAR=2 ERG=d1 $ position, particle type, energy
SP1 -2 0.5 $ Maxwellian spectrum (2) with temp a=0.5 MeV
  
```



Two Point Sources with Different Energy Distributions

```

c --- 2 pt iso sources: src 1 (4-bins) src 2 (4 discrete Ei)
SDEF PAR=2 POS=d1 ERG FPOS d2
SI1 L -10 0 0 10 0 0 $ coords of srcs on x-axis
SP1 .4 .6 $ rel strengths of sources
DS2 S 3 4 $ energy distributions
SI3 H .1 .3 .5 1. 2.5 $ E bin limits src 1
SP3 D 0 .2 .4 .3 .2 $ bin prob for src 1
SI4 L .3 .5 .9 1.25 $ discrete Ei for src 2
SP4 .20 .10 .30 .40 $ rel freq for src 2
  
```




3.3.2 Isotropic Volumetric Sources

Rectangular Parallelepiped Parallel to Axes

```

c --- volumetric monoenergetic source inside a rectangular parallelepiped
SDEF X=d1 Y=d2 Z=d3 ERG=1.25 PAR=2
SI1 -10. 10. $ x-range limits for source volume
SP1 0 1 $ uniform probability over x-range
SI2 -15. 15. $ y-range limits for source volume
SP2 0 1 $ uniform probability over y-range
SI3 -20. 20. $ z-range limits for source volume
SP3 0 1 $ uniform probability over z-range
  
```




Source in a Complex Cell: Enclosing Cube Rejection Method

```

c --- Cell 8 is some complex cell in which a monoenergetic isotropic
c volumetric source exists. A rectangular parallelepiped envelops
c this cell (MCNP does NOT check this!). Points, randomly picked
c in the rectangular parallelepiped, are accepted as source points
c only if they are inside cell 8.
c
SDEF X=d1 Y=d2 Z=d3 ERG=1.25 PAR=2 CEL=8
c NOTE: source limits are larger than the cube boundaries, and hence
c source positions picked outside cube are rejected.
SI1 -12. 12.      $ x-range limits for source volume
SP1 0 1          $ uniform probability over x-range
SI2 -11. 11.     $ y-range limits for source volume
SP2 0 1          $ uniform probability over y-range
SI3 -13. 13.     $ z-range limits for source volume
SP3 0 1          $ uniform probability over z-range

```




Source in a Complex Cell: Enclosing Sphere Rejection Method

```

c --- Cell 8 is some complex cell in which a monoenergetic isotropic
c volumetric source exists. A sphere envelops this cell {MCNP
c does NOT check this!). Points, randomly picked in the sphere,
c are accepted as source points only if they are inside cell 8.
c
SDEF POS=0 0 0 RAD=d1 CEL=8
SI1 0 20.        $ radial sampling range: 0 to Rmax (=20cm)
SP1 -21 2        $ weighting for radial sampling: here r^2

```




3.3.3 Line and Area Sources (Degenerate Volumetric Sources)

Line Source: Degenerate Rectangular Parallelepiped

```

c --- Line monoenergetic photon source lying along x-axis
c This uses a degenerate Cartesian volumetric source.
c
SDEF POS=0 0 0 X=d1 Y=0 Z=0 PAR=2 ERG=1.25
SI1 -10 10      $ Xmin to Xmax for line source
SP1 -21 0       $ uniform sampling on line Here x^0

```




Disk Source (Degenerate cylindrical Source)

```

c --- disk source in x-y plane centered at the origin.
c This is a degenerate cylindrical volume source.
c
SDEF POS=0 0 0 AXS=0 0 1 EXT=0 RAD=d1 PAR=2 ERG=1.25
SI1 0 11        $ radial sampling range: 0 to Rmax
SP1 -21 1       $ radial sampling weighting: r^1 for disk source

```




Plane Source (Degenerate Rectangular Parallelepiped)

```

c --- rectangular plane source centered on the origin and perpendicular
c     to the y-axis. This uses a degenerate Cartesian volumetric source.
c
SDEF POS=0 0 0 X=d1 Y=d2 Z=0 PAR=2 ERG=1.25
SI1 -10 10 $ sampling range Xmin to Xmax
SP1 0 1 $ weighting for x sampling: here constant
SI2 -15 15 $ sampling range Ymin to Ymax
SP2 0 1 $ weighting for y sampling: here constant

```




Line Source (Degenerate Cylindrical Source)

```

c --- line source (degenerate cylindrical volumetric source)
SDEF pos=0 0 0 axs=1 0 0 ext=d1 rad=0 par=2 erg=1.25
SI1 0 1 $ axial sampling range: -X to X
SP1 -21 0 $ weighting for axial sampling: here constant

```




3.3.4 Monodirectional and Collimated Sources

Monodirectional Disk Source

```

c --- Disk source perpendicular to z-axis uniformly emitting
c     1.2-MeV neutrons monodirectionally in the +ve z-direction.
c
SDEF POS=0 0 0 AXS=0 0 1 EXT=0 RAD=d1 PAR=1 ERG=1.2
      VEC=0 0 1 DIR=1
SI1 0 15 $ radial sampling range: 0 to Rmax (=15cm)
SP1 -21 1 $ radial sampling weighting: r^1 for disk

```




Point Source Collimated into a Cone of Directions

```

c --- Point isotropic 1.5-MeV photon source collimated into
c     an upward cone. Particles are confined to an upward
c     (+z axis) cone whose half-angle is acos(0.9) = 25.8
c     degrees about the z-axis. Angles are with respect to
c     the vector specified by VEC
c
SDEF POS=0 0 0 ERG=1.25 PAR=2 VEC=0 0 1 DIR=d1
SI1 -1 0.9 1 $ histogram for cosine bin limits
SP1 0 0.95 0.05 $ frac. solid angle for each bin
SB1 0. 0. 1. $ source bias for each bin

```



With this conical source, tally normalization is per source particle in 4π steradians. To normalize the tally per source particle in the cone, put `WGT=fsa2` on the SDEF card, where `fsa2` is the fraction solid angle of the cone (0.05 in the above example).

This conical collimation trick can also be used to preferentially bias the emission of particles in certain directions. The `SIn` entries are the upper bin cosine limits $\mu_i \equiv \cos\theta_i$ in ascending order. The first entry is -1 . Angles are with respect to the direction specified by `VEC`. The `SPn` entries give the fractional solid angle $\text{fsa}_i = [(1 - \mu_{i-1}) - (1 - \mu_i)]/2$ for the bin from μ_{i-1} to μ_i , and the `SBn` entries give the desired relative probabilities for emission in each angular bin. Note the first probability must be 0 for the unrealistic bin from $(-\infty, -1)$.


3.3.5 Multiple Volumetric Sources

Two Cylindrical Volumetric Sources

```

c --- 2 volumetric sources uniformly distributed in cells 8 & 9.
c     Both sources emit-1.25 MeV photons. Surround both source cells
c     by a large sampling cylinder defined by the POS RAD and EXT
c     parameters. The rejection technique is used to pick source
c     points with cells 8 and 9 with the specified frequency.
c
SDEF ERG=1.25 CEL d1 AXS=0 0 1 POS 0 0 0 RAD d2 EXT d5
SI1 L 8 9 $ source cells: src 1 =cell 8, src 2 =cell 9
SP1 0.8 0.2 $ 80% from src 1; 20% from src 2
SI2 0 50 $ radius of cyl. containing cells 8 & 9
SI5 -30 30 $ axial range of cyl. containing src cells

```




Two Cylindrical Sources with Different Energy Photons

```

c --- Two spatially different cylindrical monoenergetic sources.
c     The size and position of each cyl. source depends on the
c     source energy (FERG).
c
SDEF ERG=d1 POS=FERG d8 AXS=0 0 1 RAD=FERG d2 EXT=FERG d5
c
c -- set source energies: .667 MeV for region 1 and 1.25 MeV for region 2
SI1 L 0.667 1.25 $ fix energies: .667 MeV for region 1 and 1.25 MeV for region 2
SP1 0.4 0.6 $ 20% from src 1(Cs-137); 80% from src 2 (Co-60)
c -- set positions of the 2 source cylinders
DS8 S 9 10 $ based on source chosen, get position
SI9 L -30 0 0 $ center for spatially sampling of source 1
SP9 1 $ prob. distn for src 1 center
SI10 L 30 0 0 $ center for spatially sampling of source 2
SP10 1 $ prob. distn for src 2 center
c -- set radius and axial limits for each source cylinder
DS2 S 3 4 $ distn for sampling radially from each src axis
SI3 0 20 $ radial sampling limits for src1
SP3 -21 1 $ radial sampling weight for src1 r^1
SI4 0 10 $ radial sampling limits for src2
SP4 -21 1 $ radial sampling weight for src2 r^1
DS5 S 6 7 $ distns for sampling axially for each src
SI6 -10 10 $ axial sampling limits for src1
SP6 -21 0 $ axial sampling weight for src1 r^0
SI7 -30 30 $ axial sampling limits for src2
SP7 -21 0 $ axial sampling weight for src2 r^0

```




Two Arbitrary Volumetric Sources with Different Energy Photons

```

c --- 2 volumetric monoenergetic sources in complex-shaped cells 8 & 9
c     Spatial sampling uses the rejection technique by placing a finite
c     cylinder over each source cell. A random point inside a cylinder
c     is accepted as a source point only if it is inside the source
c     cell. Location and size of the sampling cylinders and source
c     photon energies are functions of the source cells (FCEL).
c
SDEF CEL=d1 POS=FCEL d2  AXS=0 0 1  RAD=FCEL d5  EXT=FCEL d8  ERG=FCEL d20
c
SI1 L 8 9                $ choose which cell source region to use for source
SP1  0.4 0.6             $ 40% from src 1; 60% from src 2
c -- set POS for each source
DS2 S 3 4                $ based on the cell chosen, set distribution for POS
SI3 L -30 0 0            $ center for spatially sampling of source 1
SP3  1                    $ prob. distn for src 1 center
SI4 L 30 0 0             $ center for spatially sampling of source 2
SP4  1                    $ prob. distn for src 2 center
c -- set RAD for each source (must completely include cells 8 or 9)
DS5 S 6 7                $ distns for sampling radially from each src axis
SI6  0 20                $ radial sampling limits for src1
SP6  -21 1               $ radial sampling weight for src1
SI7  0 10                $ radial sampling limits for src2
SP7  -21 1               $ radial sampling weight for src2
c -- set EXT for each source (must completely include cells 8 or 9)
DS8 S 9 10               $ distns for sampling axially for each src
SI9  -10 10              $ axial sampling limits for src1
SP9  -21 0               $ axial sampling weight for src1
SI10 -30 30              $ axial sampling limits for src2
SP10 -21 0               $ axial sampling weight for src2
c -- set energies of photons for each source
DS20 S 21 22
SI21 L 0.6938 1.1732 1.3325 $ Co-60 spectra for src 1
SP21 D 1.6312E-4 1 1       $ frequencies of gammas
SI22 L 0.667              $ Cs-137 spectrum for src 2
SP22 D 1

```



3.4 Tally Specifications

A technical description of the various types of tallies permitted in MCNP calculations is given in Ch. 2.V of the manual. Details of specifying tallies using tally cards and tally modification cards is given in Ch.3.IV.E.

Table 3. Types of tallies available in MCNP

Mnemonic	Tally Type	particles p
F1: p	surface current	N P N,P E
F2: p	average surface flux	N P N,P E
F4: p	average flux in a cell	N P N,P E
F5a: p	flux at a point or ring	N P
FIP5: p	pin-hole flux image	N P
FIR5: p	planar radiograph flux image	N P
FIC5: p	cylindrical radiograph flux image	N P
F6: p	energy deposition	N P N,P
F7: p	fission energy deposition in a cell	N
F8: p	pulse height distribution in a cell	P E P,E

The basic types of tallies are current at a surface, flux at a surface, flux at a point or ring, and flux averaged over a cell. Similar to flux tallies over a cell are various tallies of energy deposition. The accompanying table identifies mnemonics for specifying tally type. The designation F1:P, for example, would be changed to F1:N for neutrons or F1:E for electrons.

pp. 3-78 to
3-89

Unless otherwise specified, tallies are normalized to one source particle. Except for tallies F6 and F7, designating a tally as *F1:P, for example, multiplies the tally of each event by the particle energy. This results in tallies of energy flux or energy current. Tallies F6 and F7 are already in energy units.

Multiple tallies may be made at any one surface or any one cell. Similarly, the same type of tally may be performed at many surfaces or cells. If, for example, surface flux tallies were made at three surfaces, it would be necessary to identify them separately as F2:P, F12:P and F22:P.

The following sections describe the physical nature of several tallies. In the description, time dependence is suppressed, which is the normal case in MCNP calculations. The flux is integrated over time, and might better be called the fluence.

3.4.1 The Surface Current Tally (type 1)

Each time a particle crosses the specified surface, its weight is added to the tally, and the sum of the weights is reported as the F1 tally in the calculation output. Note that there is no division by surface area. Nor is there a distinction between direction of surface crossing. When used with problem geometry voided (zero density), the tally is useful for verifying conservation of energy and conservation of number of particles. Technically, if $J(\mathbf{r}, E, \Omega)$ were the energy and angular distribution of the flow (current) as a function of position, the F1 tallies would measure

$$F1 = \int_A d\mathbf{r} \int_E dE \int_{4\pi} d\Omega J(\mathbf{r}, E, \Omega)$$

$$*F1 = \int_A d\mathbf{r} \int_E dE \int_{4\pi} d\Omega E J(\mathbf{r}, E, \Omega)$$

3.4.2 The Average Surface Flux Tally (type 2)

Suppose a particle of weight W crosses a surface, making angle θ with a normal to the surface. This particle makes a contribution $W|\sec\theta|/A$ to the flux (fluence) at the surface. The sum of the contributions is reported as the F2 tally in the MCNP output.

Technically, if $\Phi(\mathbf{r}, E, \Omega)$ were the energy and angular distribution of the fluence as a function of position, the F2 tallies would measure

$$\begin{aligned} \text{F2} &= \frac{1}{A} \int_A d\mathbf{r} \int_E dE \int_{4\pi} d\Omega \Phi(\mathbf{r}, E, \Omega) \\ *F2 &= \frac{1}{A} \int_A d\mathbf{r} \int_E dE \int_{4\pi} d\Omega E \Phi(\mathbf{r}, E, \Omega) \end{aligned}$$

3.4.3 The Average Cell Flux Tally (type 4)

Suppose a particle of weight W and energy E makes a track-length (segment) T within a specified cell of volume V . This segment makes a contribution WT/V to the flux (fluence) in the cell. The sum of the contributions is reported as the F4 tally in the MCNP output. Technically, if $\Phi(\mathbf{r}, E, \Omega)$ were the energy and angular distribution of the fluence as a function of position, the F4 tallies would measure

$$\begin{aligned} \text{F4} &= \frac{1}{V} \int_V d\mathbf{r} \int_E dE \int_{4\pi} d\Omega \Phi(\mathbf{r}, E, \Omega) \\ *F4 &= \frac{1}{V} \int_V d\mathbf{r} \int_E dE \int_{4\pi} d\Omega E \Phi(\mathbf{r}, E, \Omega) \end{aligned}$$

3.4.4 Flux Tally at a Point or Ring (type 5)

This type of tally makes use of what some might call a variance reduction technique, namely, use of the “next event estimator. For each source particle and each collision event, a deterministic estimate is made of the fluence contribution at the detector point (or ring in an axisymmetric problem). To simplify description of this type of tally, assume that calculations are being performed in a uniform medium. Suppose a particle of energy E and weight W from an isotropic source is released at distance R from the detector point. Ray theory methodology as used in the point-kernel method dictates that the contribution $\delta\Phi$ to the fluence at the detector point is given by

$$\delta\Phi = \frac{W}{4\pi R^2} e^{-\mu(E)R},$$

in which $\mu(E)$ is the linear interaction coefficient for the particle of energy E . Note that $1/4\pi$ per steradian is the angular distribution of a point isotropic source. Now suppose that a collision takes place at distance R from the detector point and that, to reach the detector point, a scattering angle of θ_s would be required. Here E is the energy of the particle after the collision and W is its weight. If $\mu(E, \theta_s)$ is the linear interaction coefficient per steradian for scattering at angle θ_s , then $\mu(E, \theta_s)/\mu(E)$ is the probability per steradian for scattering at angle θ_s . Geometric attenuation remains as $1/R^2$, and the contribution $\delta\Phi$ to the fluence at the detector point is given by

$$\delta\Phi = \frac{W\mu(E, \theta_s)}{\mu(E)R^2} e^{-\mu(E)R}.$$

3.4.5 Tally Specification Cards

At least one tally card is required, with the first entry on the card being **Fn:p1**, in which **n** is the tally type, and **p1** stands for N (neutron tally), P (photon tally), N,P for joint neutron and photon tallies, and E for electron tallies. Following the tally type is a designation of the surfaces for the

tally (types 1 and 2), or the cells (tally 4). For the type 5 detector tally, there follows a designation of the position of the detector. Specification of energy deposition or pulse-height tallies or other specialized tallies is not discussed in this summary. Nor is discussion of the treatment of repeated structures.

3.4.6 Cards for Optional Tally Features

A table on p. 3-77 of the MCNP manual lists many optional features of tallies. The more important of these are tally energies (En card), tally multiplier (FMn), and fluence-to-dose conversion factors (DEn and DFn). These are addressed individually in the sections to follow, mostly by way of examples.

Sec. 3.E pp.
3-89 to
3-114

3.4.7 Cards for Surface and Cell Tallies

The card

F1:E 1 2 T

requires electron current tallies over surfaces 1 and 2, and the total over both surfaces. Recall that the current tally is not divided by surface area. The card

F2:P 1 (1 2) (2 3 4) T

requires photon surface fluence tallies for surface 1, the average over surfaces 1 and 2, the average over surfaces 2 through 4, and the average (T) over all surfaces 1 through 4. Similarly, the card

F4:N 1 (2 3 4)

requires cell neutron fluence tallies for cell 1 and the average over cells 2 through 4. No composite average is called for.

3.4.8 Cards for Point-Detector Tallies

In the sense of an experiment or Monte Carlo calculation, as the volume of a cell approaches zero, the path length segments in the cell and the number of particles intersecting the surface of the cell also approach zero, and the flux tally becomes indeterminate. However, there is a way of computing the flux at a point making use of the last-flight-estimator.⁴ The tally is called for by a card such as

F5:P X Y Z R

Again, 5 designates the tally number and P the type of particles being tallied. X, Y, and Z specify the point at which the tally is made, and R designates the radius of a spherical exclusion zone surrounding the point. The need for an exclusion zone is evident from the $1/R^2$ term in the flux contribution tallied, namely,

$$\delta\Phi = \frac{W}{4\pi R^2} e^{-\mu(E)R}.$$

If R approaches zero, the tally approaches infinity. Such close encounters make the detector flux tally much less stable than the cell or surface flux tally. This instability is minimized by establishing a minimum radius surrounding the point within which the flux contribution is tallied. For those events that would have occurred within the exclusion radius, an approximate flux contribution is computed using a method described on page 2-89 of the MCNP Manual. The Manual also describes the use of a ring detector useful for problems with symmetry about the axis of the ring. For a detector point in a void region, R should be set to zero. Several point detectors may be specified on one tally card, e.g.,

F5:P X1 Y1 Z1 R1 X2 Y2 Z2 R2

3.4.9 The Tally Energy Card

Suppose one wanted to subdivide the total flux or current tally number n into energy groups, say E1 to E2, E2 to E3, and E3 to E4. This might be useful, for example, to isolate an uncollided component of the flux. This may be accomplished by use of a tally energy card in the form

En E1 E2 E3 E4

p. 3-90

The result would be tallies for four energy groups, not three. Group (bin) upper limits are specified on the card. The lowest bin would extend down from E1 to zero or to a specified cutoff energy for the type of particle being tallied. If all tallies have the same energy group structure, a single card may be used, with n replaced by 0.

3.4.10 The Energy Multiplier Card

Optionally associated with the tally energy card is an energy multiplier card of the form

p. 3-98

```
EMn M1 M2 M3 M4
```

Here the multiplier M_k is applied to each contribution to the tally for the k th energy group. This card is useful, for example, to convert a fluence per source photon to a flux per curie source strength. For this example, one would add the following EM card for, say tally F64.

```
EM64 3.7E+10 $ photons per sec / curie (assuming 1 photon/decay)
```

The units of tally F64 would then be “photons $\text{cm}^{-2}\text{s}^{-1}$ per Ci.

3.4.11 Dose Energy and Function Cards

Suppose one wanted to compute a dose rate of some type associated with a flux or current tally, either total or by energy group. For example, suppose one wanted to compute

p. 3-97

$$*F4 = \frac{1}{V} \int_V d\mathbf{r} \int_E dE \int_{4\pi} d\Omega \mathfrak{R}(E) \Phi(\mathbf{r}, E, \Omega),$$

in which $\mathfrak{R}(E)$ is a fluence-to-dose conversion factor. MCNP will carry out this calculation, obtaining values of $\mathfrak{R}(E)$ by interpolation in a table placed in the input file. The form of the table is

```
DE4 A E1 E2 ... Ek
DF4 B F1 F2 ... Fk
```

Entries E1 through Ek are tabulated values of energy and F1 through Fk are corresponding tabulated values of $\mathfrak{R}(E)$. Entries A and B, either LOG or LIN, specify linear or logarithmic interpolation. If omitted, the default is logarithmic interpolation in both. If all tallies have the same dose conversion factors, a single table, designated by DE0 and DF0, may be used to avoid repeating the table.

3.4.12 The Tally Comment Card

If tallies are modified, it is good practice to explain the modification in a comment card that will be printed in the output file for the calculation. For example, an explanation of the nth tally could be entered in the card

p. 3-89

```
FCn This tally has units of Sieverts per source photon
```

Continuation lines may be added so long as there are blanks in columns 1 through 5.

3.4.13 Miscellaneous Data Specifications

The Mode Card This card is used to specify the type of problem, i.e., type of source particles to be tracked. Every input file must have a MODE card somewhere in block 3. In the line

p. 3-24

```
MODE x
```

the variable x may be N, P, E, or a combination such as N,P. When the mode is specified, the PAR entry may be omitted on the SDEF source definition card.

Time or History Cards The usual method for limiting how long MCNP should run is to specify either the maximum number of source particle histories or the maximum execution time. The maximum number of histories N is specified on the card

p. 3-133

```
NPS N
```

In addition, or as an option, the computing-time cutoff T, in minutes, may be specified by the card

p. 3-134

```
CTME T
```

The Print-and-Dump Cycle Card By default, an output file is created only at the conclusion of a calculation, a binary continuation file, RUNTPE, is written every 15 minutes, and no tally-plot file, MCTAL, is written. Options to control the dump cycle are provided the command PRDMP card

```
PRDMP NDP NDM MCT NDMP DMMP
```

in which NDP is the increment for printing tallies in the output file (positive number is number of histories negative the time in minutes 0 for no intermediate dump); NDM is the increment for writing a RUNTPE file (a positive number is the number of histories, a negative number is the time in minutes, and 0 suppresses all intermediate dumps); MCT is a flag to write tallies for plotting (1 yes, 0 no); NDMP is the maximum number of dumps written in the RUNTPE file (all by default); and DMMP is related to the use of multiple processors in the execution of MCNP. A typical card might read

```
PRDMP 0 60
```

With this card, at most, 60 minutes of computing time would be lost if a calculation were aborted.

4 Variance Reduction

The challenge in using MCNP is to minimize the computing expense needed to obtain a tally estimate with acceptable relative error (as well as satisfying nine other statistical criteria). For many deep-penetration problems, a direct simulation (analog MCNP) would require far too many histories to achieve acceptable results with the computer time available. For such cases, the analyst must employ “tricks” to reduce the relative error of a tally (or its variance) for a fixed computing time, or to reduce the computing time to achieve the same relative error.

Two basic approaches can be applied to reduce the computational effort for a particular problem: (1) simplify the MCNP model, and (2) use non-analog simulations. In the first approach, the model geometry and the physics used to simulate particle transport can often be simplified or *truncated*. For example, it is a waste of computing effort to use a detailed geometric model of a region that is far from the detector tally location and that has little influence on the radiation field near the detector. Similarly, it is a waste of computer time to track neutrons as they thermalize in a shield if only the fast neutron fluence in some structural component is sought. For such a problem, once a neutron leaves the fast energy region, it can be killed without affecting the tally.

The second basic approach to reduce the variance of a tally is to modify the simulation process itself by making certain events more or less probable than actually occur in nature. Such a modified simulation is referred to as *nonanalog* Monte Carlo. As discussed in this section, MCNP has many nonanalog options many of which an analyst can use in combination to make a difficult analog problem much more tractable. These nonanalog tricks can be categorized into three general methods: (1) population control, (2) modified sampling, and (3) partially-deterministic calculations. In population control, for example, the number of particles in regions of high/low importance can be artificially increased/reduced. In modified sampling methods, certain events can be altered from their natural frequencies. Finally, in the partially-deterministic methods, part of the random-walk simulation can be replaced by a deterministic point-kernel type of calculation.

4.1 Tally Variance

Before discussing the tricks used to reduce the variance of MCNP tallies, it is appropriate to examine exactly what it is that we are trying to reduce. When we run a random-walk simulation, the i th history contributes a score x_i to the tally. If the particle (or its daughters) never reaches the tally region, then $x_i = 0$, whereas, if it reaches the tally without interaction, the score x_i generally is very large. The probability any history will contribute a score between x and $x + dx$ is denoted by a probability distribution function (PD) $p(x) dx$. In an MCNP simulation, we seek the mean score (or expected value) of x , namely

$$\langle x \rangle \equiv \int_0^{\infty} x p(x) dx. \quad (1)$$

Unfortunately, we don't know $p(x)$ *a priori* (although MCNP will construct it and generate a plot of it – see examples p. 2-106). Instead, we estimate $\langle x \rangle$ by the mean \bar{x} of the scores of N particles, i.e.

$$\bar{x} \equiv \frac{1}{N} \sum_{i=1}^N x_i. \quad (2)$$

As $N \rightarrow \infty$, the *strong law of large numbers* guarantees that $\bar{x} \rightarrow \langle x \rangle$, provided $\langle x \rangle$ is finite.

The variation in the different scores x_i is measured by the *standard deviation* of the population (histories), which for large N

$$S^2 \equiv \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2 \simeq \overline{x^2} - \bar{x}^2, \quad (3)$$

where

$$\overline{x^2} \equiv \frac{1}{N} \sum_{i=1}^N x_i^2. \quad (4)$$

The estimated variance of the mean \bar{X} is then

$$S_{\bar{x}}^2 = \frac{1}{N} S^2. \quad (5)$$

The *central limit theorem* states that if we repeated the simulation a large number of times (each with N histories), the variation of the means \bar{x} from each simulation will be distributed normally about the true mean $\langle x \rangle$ and have a variance $S_{\bar{x}}^2$. It is this uncertainty or variance we are trying to reduce in our MCNP simulations, i.e., for a fixed number of particles, we seek an estimate \bar{x} which has the least uncertainty or minimum $S_{\bar{x}}$.

4.1.1 Relative Error and FOM

In any variance reduction method, we change the simulation and hence change the underlying distribution $p(x)$ so that it produces fewer zero-score histories and becomes more concentrated about its mean $\langle x \rangle$. By making $p(x)$ more concentrated about its mean (which remains the same as the mean of the analog PD), the variance of the mean $S_{\bar{x}}^2$ will be less than that of the analog PD, i.e., our estimate of the mean will be more precise.

For each tally, MCNP not only calculates the sample mean \bar{x} , but several other statistics. One of the most important is the *relative error* R defined as

$$R \equiv S_{\bar{x}}/\bar{x}. \quad (6)$$

Clearly, we want to make R as small as possible with as few histories as possible. As discussed in the manual, R generally must be less than 0.1 for meaningful results (and even smaller if point/ring detectors are used). From Eqs. (5) and (6), it is seen that $R \sim 1/\sqrt{N}$. Thus increasing the number of particle histories is generally a very poor way of reducing R . This property of the relative error is the great weakness of the Monte Carlo method, since generally many histories must be generated to obtain acceptable results.

Another important statistic generated by MCNP is the *figure of merit* (FOM). This is defined as

$$FOM \equiv \frac{1}{R^2 T}, \quad (7)$$

where T is the simulation time, which is proportional to N the number of histories run. Since $R^2 \sim 1/N$, we see that, except near the beginning of the simulation, the FOM should remain relatively constant. Also, for different simulations of the same problem, the simulation with the largest FOM is preferred since it requires the least time or produces the smallest error.

Now on to ways of how to perform nonanalog techniques with MCNP.

4.2 Truncation Techniques

The basic idea behind truncation methods is to reduce the time per particle history by either simplifying the geometry or the physics used to generate the random walk for each particle. Proper application of this approach for variance reduction requires considerable experience and intuition by the analyst, since any simplification in the geometry or physics introduces a *bias* in the tally. Although a very *precise* (i.e., low variance or relative error) can be achieved, the tally estimate may not be very *accurate*. Generally, multiple runs with different approximations must be made to assess the importance of any simplification. MCNP can give you no warning about errors caused by geometric simplifications. Even for physics simplifications, MCNP produces, at best, a warning in the output, but no indication of whether serious bias has been introduced.

4.2.1 Energy, Time and Weight Cutoff

The CUT command is used to specify a minimum energy, time, or particle weight below which the particle is killed. The values specified on the CUT card apply everywhere in the geometry. Here is an example:

```
CUT:p j 0.075 $ kill photons with E < 75 keV
```

In this example, whenever a photon falls below 75 keV, it is killed.

The ELPT is like the CUT card, but allows you to specify the cutoff on a cell-by-cell basis. For example,

```
ELPT:p 0.01 0.02 0.03 0.04 0.05 $ energy cutoffs
```

terminates photons in cell 1, 2, 3, 4, and 5 that have energies less than 10, 20, 30, 40, and 50 keV, respectively. Should both the ELPT and the global CUT commands be used, the higher limit prevails.

The CUT and ELPT commands are particularly useful for energy deposition tallies for which low energy particles make little contribution. However, for neutron problems, use CUT and ELPT carefully since low energy neutrons cause most of the fissions and produce most of the capture gamma photons.

4.2.2 Physics Simplification

The PHYS command is used to specify energy cutoffs and the physics treatments to be used for photons, neutrons and electrons. Each particle has different parameters which are specified with this command. An example:

```
PHYS:p 0.01 1 1 $ simple physics > 10 keV, no Brems, no coherent
```

The first parameter specifies that simple physics is to be used for photons above 10 keV. In simple physics, no fluorescence from photoelectric interactions is produced, no binding effects are used in Compton scattering, and no coherent scattering is included. The second parameter 1 specifies no Bremsstrahlung production, while the third parameter 1 sets the coherent scattering cross section to zero. Omitting this command is equivalent to `PHYS:p 100 0 0`, so that detailed physics with Bremsstrahlung and coherent scattering is the default.

An example for neutrons:

```
PHYS:n 5.0 0.1 $ max sigma table energy; analog capture below 100 keV
```

Here cross section data only below 5 MeV is retained (to save data storage memory). For neutrons below 0.1 MeV, *analog absorption* (direct simulation) will be used, while above 0.1 MeV, *implicit absorption* is used.

As an example of how different physics simplifications can affect the run time, consider a point isotropic source emitting 7-MeV photons into an infinite iron medium. The ambient dose equivalent 20 cm from the source is estimated by using an F2 spherical surface detector. In Table 4 the tally mean and the runtime are shown for different physics assumptions.

p. 3-131

p. 3-133

pp. 3-124 to
3-129

Table 4. Results for different physics models for a point 7-MeV photon source in an infinite iron medium. Tally is the ambient dose equivalent 20 cm from the source. Results are for simulations of 5×10^5 source photons.

Test File	MCNP Commands	Description	F2 Tally (Sv/photon)	Relative Error	FOM	Time (min)
PHYS0		default	1.232E-16*	0.0070	3000	6.84
PHYS1	PHYS 10 0 0	detailed physics + brem. + coh. scat.	1.231E-16*	0.0070	1720	11.77
PHYS2	PHYS 10 0 1	no coherent scattering	1.235E-16	0.0070	3560	5.79
PHYS3	PHYS 0.01 0 1	simple phys. > 10 keV + no coherent scat.	1.227E-16	0.0070	1804	11.38
PHYS4	PHYS 0.001 0 1	simple phys. > 1 keV + no coherent scat.	1.222E-16	0.0070	1963	10.57
PHYS5	PHYS 10 1 1	no bremsstrahlung or coherent scattering	1.100E-16	0.0075	8200	2.19
PHYS6	PHYS 0 1 1	simple physics + no brem.+ no coh. scat.	1.104E-16	0.0075	6840	2.63
PHYS7	PHYS 0 1 1 CUT:p j 0.1	detailed physics + no brem.+ no coh. scat.	1.103E-16	0.0075	7260	2.48
PHYS8	PHYS 0 0 1 CUT:p j 0.1	detailed physics + with brem.+ no coh. scat.	1.229E-16	0.0070	4880	4.17

* Did not meet all 10 statistical tests

4.2.3 Histories and Time Cutoffs

Normally an MCNP run is terminated when a certain number of particle histories have been run or a desired computing time has been exceeded. These cutoffs are specified by the NPS and CTME commands such as

```
NPS 1000000 $ stop after a million source particles have been run
CTME 20.0   $ stop run after twenty minutes
```

If both are specified, the first cutoff to occur causes program termination.

4.3 Nonanalog Simulation

In many problems, very few of the source particles reach the detector or region used for the tally, i.e., few particles ever *score*. The number of particles reaching the tally region can, however, often be dramatically increased by abandoning a strict analog simulation. Of course, the expected value of the tally must not be changed. How can the tally remain unchanged when we artificially force more particles to the scoring region? The key is to assign each particle a *weight*, and, as the particle is “forced” towards the scoring region, the particle weight is decreased in a manner such that the average of the particle weights reaching the detector is the same as the expected tally in a true analog simulation. Thus, if we make a certain event in a particle history m times more likely, we must multiply the particle’s weight by $1/m$ to avoid biasing the tally expectation.

MCNP has many nonanalog simulation options whose use, often in combination, can decrease the variance of a tally without increasing the computational expense.

4.3.1 Simple Examples

To understand the basic idea of nonanalog techniques, consider the simple slab transmission problem illustrated in Fig. 3. In this problem a point isotropic source is placed on one side of a slab shield, and the problem is to determine the fraction of source particles that reach the opposite face of the slab. A direct analog simulation is represented by Fig. 3(a).




Figure 3. Examples of analogue and nonanalogue Monte Carlo simulations.

Source Biasing A non-analog simulation can considerably reduce the computing effort compared to an analog simulation. For example, in the analog simulation half of the source particles are “wasted”, i.e., those emitted away from the slab cannot reach the scoring surface and computer time is wasted sampling source directions and tracking these particles to the problem boundary. It would be more efficient to start each source particle towards the slab, as shown in Fig. 3(b). However, by restricting or biasing the source emission directions to only those oriented toward the slab, twice as many particles will subsequently penetrate the slab in case (b) compared to analog case (a), for the same number of particle histories tracked. To avoid doubling the transmission tally (no. transmitted per real (analog) source particle), we adjust the source particle’s weight in the biased simulation to be 0.5. Thus the average of the *weights* of transmitted particles still equals the particle transmission fraction obtained with the analog simulation. Moreover, for the same number of source particles, twice as many reach the tally surface in case (b) compared to case (a), and thus the variance of the case (b) tally is less.

Splitting Another technique for increasing the number of particles reaching the scoring surface is illustrated in Fig. 3(c). Here the slab is conceptually divided into two sublayers. Whenever a particle crosses from the layer nearest the source to the layer nearest the tally surface, it is *split* into two particles, each with half of the original particle’s weight and both moving with the same velocity as the original particle. The random walk simulation is then performed independently for each new particle, beginning at the entry point into the second layer of the original particle. Twice as many particles will now reach the tally surface (thereby reducing the tally’s variance); but, since their weights have been reduced by one-half, the expected tally value remains unchanged.

Russian Roulette When a particle reaches a region of space far from the tally region it is unlikely, with further random walk simulation, to reach the tally region, and we can save computing effort by terminating its history. Thus, in the example case (c), when a particle in the right-hand sublayer returns to the first left sublayer, we may think it has a relatively poor chance of returning yet again the right sublayer and reaching the tally surface. When a particle reenters the first layer from the second, the particle is hence killed with a probability of 0.5. If the particle survives this winnowing process, its weight is increased by a factor of two to keep the simulation unbiased.

Implicit Absorption Those particles which are tracked through the slab but are absorbed before they reach the tally surface represent wasted computing effort. Another variance reduction trick is to replace *analog capture* with *implicit capture*. At a collision site, a particle is killed, in an analog simulation, with a probability σ_a/σ_t (analog capture). However, in implicit capture, the particle is allowed to continue on its trajectory as if no interaction had occurred but with the particle's weight changed to $1 - (\sigma_a/\sigma_t)$ times its original weight. In this way, no particles are lost due to absorption, but absorption effects are properly accounted for.

4.4 MCNP Variance Reduction Techniques

MCNP offers a variety of variance reduction techniques based on different nonanalog simulations. The art of using MCNP to solve difficult problems is to use these program features to obtain both precise and computationally efficient results. In this section, the use of several of the most useful variance reduction techniques are described.

The variance reduction techniques offered by MCNP can be categorized as follows:

1. *Population Control Methods*: These methods artificially increase/decrease the number of particles in spatial or energy regions that are important/unimportant to the tally score. Specific population control methods include
 - Geometry splitting and Russian roulette (IMP)
 - Energy splitting/roulette (ESPLT)
 - Weight cutoff (CUT, PWT)
 - Weight windows (WWE, WWN, WWP, WWG, WWGE)
2. *Modified Sampling Methods*: These methods artificially increase the likelihood of events that increase the probability a particle reaches the tally region. Included in MCNP are
 - Exponential transform (EXT, VECT)
 - Implicit capture (PHYS)
 - Forced collisions (FCL)
 - Bremsstrahlung biasing (BBREM)
 - source direction and energy biasing (SDEF, SP, SB, SI)
 - neutron-induced photon production biasing (PWT, 2-31)
3. *Partially Deterministic Methods*: These methods replace the random-walk process by a deterministic process (e.g., exponential attenuation) to move particles from one region to another. In MCNP the following are available:
 - Point and ring detectors (F5a)
 - DXTRAN spheres (DXT, DXC)
 - Correlated sampling (PD, 2-143)

The selection of effective variance reduction methods for a particular problem requires considerable experience and skill on the part of the analyst in interpreting the MCNP output. To gain experience in using these variance reduction techniques, the novice is encouraged to try using them on simple problems, sometimes separately and sometimes in various combinations. Through such experimentation, valuable experience and insight into variance reduction is gained. In the sections below, some of the simpler variance reduction techniques are discussed and illustrated.

4.4.1 Geometry Splitting

In geometry splitting, *importances* are assigned to each cell in the problem. Generally, cells near the tally region should have a greater importance than cells farther away. When a particle leaves a cell with importance I_1 and enters a cell of importance I_2 , the particle is split/rouletted according to the ratio I_2/I_1 . For example, if $I_2/I_1 = 2.75$ the entering particle is split into 3 particles with 75% probability and into 2 particles with 25% probability. If $I_2/I_1 = 0.6$, the entering particle is killed with 40% probability and allowed to survive with 60% probability. Of course, in each splitting or Russian roulette the weight of the remaining particles is adjusted to leave the tally unbiased. This technique of geometry splitting with Russian roulette is very reliable since, if no other biasing techniques are used, all the particles in a cell will have the same weight regardless of the paths taken to reach the cell.

p. 2-126,
3-33

The importance of a cell can be defined on the cell definition line, such as

```
c    Set cell importance on the cell definition line
20 1 -7.86 10 -20 IMP:p=7 $ cell 20; matl 1; density; defn; importance
```

or the importances of all cells can be set in Block 3 with the IMP command

```
c    Set cell importances in a geometric progression
IMP 1 2m 2m 2m 2m 2m 0 $ import. of cells 1--7 = 1 2 4 8 16 32 0
```

The importance of a cell is intimately related to the average adjoint fluence in the cell (a quantity generally not known *a priori*). As a practical matter, the cell importances should be adjusted so as to keep the population of particles in the cells relatively constant as one moves from the source region to the tally region. First, perform a short run with all importances set to unity, examine the “cell population” found in output print table 126, and estimate the cell importances by the ratio of cell populations P in adjacent cells, i.e. $I_n \simeq P_{n-1}/P_n$.

As a general rule, adjacent cells should not have importances differing by more than a factor of 2 to 4. Consequently, it is often necessary to subdivide cells into many cells to prevent adjacent cell importances from changing too rapidly. Finally, when Russian roulette is used to terminate some particles, information is lost; subsequent building up the particle population with large cell importances cannot regain the lost information.

A warning about large importances. In problems with large attenuation of particles between the source and tally region, importances of cells near the tally can reach many orders of magnitude. For these cases, if a VOID command is used to flood the geometry with particles in order to find geometry errors, the cell importances are still in effect, and a few source particles will be magnified into millions of particles, all of which MCNP must track. Instead of a short run, hours or days can be required!

4.4.2 Weight Windows

The weight-window variance reduction technique adjusts the weights of particles as they change energy and move through the various cells in the problem geometry. In each cell, a lower weight bound and an upper bound (defined as a multiple of the lower bound) are specified. If a particle entering a cell or a particle created in the cell has a weight above the upper bound, the particle is split such that all split particles are within the weight window. Similarly, if a particle has a weight below the lower bound, Russian roulette is used to increase the particle’s weight until it lies within the window or until it is killed.

(2-124)
p. 2-140,
3-36

In most problems weight-windows is preferred over importance biasing.

Advantages:

- The weight-windows variance reduction technique is a space and energy biasing scheme, whereas importance sampling is only a spatial biasing technique.

- Weight window discriminates on particle weight before taking appropriate action. Geometry splitting is done regardless of the particle's weight.
- Weight windows uses absolute bounds, whereas geometry splitting is based on ratios so that a particle's weight can grow or decrease without limit. This is particularly useful when using ring and point detectors with which large particle weights can cause large tally perturbations.
- Weight windows is applied at surfaces and collision sites whereas geometry splitting occurs only at surfaces.
- Weight windows is immune to weight fluctuations caused by other biasing techniques, whereas geometry splitting preserves such fluctuations.
- Weight windows can be turned off in large cells, in which no single importance applies, by setting the lower limit to zero.
- Weight windows can be generated automatically by MCNP whereas cell importances requires considerable insight by the user.
- Weight windows is more compatible with other variance reduction techniques such as the exponential transform.

Disadvantages:

- Weight windows are not as straight forward as geometry splitting. Without the automatic weight-window generator, weight windows would be very difficult to use since window limits of each cell are difficult to predict. By contrast, cell importances are much easier to guess.
- When the weight of source particles is changed, the weight-window limits have to be renormalized.

Generating Weight Windows The weight-windows biasing is specified with the WWE, WWN and WWP commands. The reader is referred to the manual for the command parameters. However, the non-expert rarely enters these command directly; rather, the *weight-windows generator* is usually used to automatically calculate these command parameters.

p. 3-46

To use the weight-window generator, the WWG (and, optionally, the WWGE) commands are placed in Block 3 of the input file. The WWG command is

```
WWG It Ic Wg xr yr zr J
```

where I_t is the tally number, I_c is a reference (usually source) cell, W_g is the value of generated lower weight-window bound (if 0, set to 0.5 of source particle weight), x_r y_r z_r are coordinates of the point particles move towards or from (usually a source point), and J is a debug flag ($J = 0/1 = \text{on/off}$).

The optional WWGE card has the form

```
WWGE:n E1E2...Ei...Ej
```

where $n = \text{N/P/E}$ for neutrons/photons/electrons, E_i is the upper energy bound for weight-window group ($E_{i+1} > E_i$), and j is the maximum number of energy groups ($j \leq 15$).

As an example for a point photon source at the origin and in cell 10, the weight-window generator command

```
c Automatically generate weight windows:
c Tally 2, ref cell 10, default lower weight, source coords, no debug
WWG 2 10 0 0 0 0 $ generate weight windows
```

is placed in Block 3 of the input. Near the end of the resulting output file, lines similar to the following will appear.

```

wwp:p 5 3 5
wwe:p 1.0000E+02
wnn 1:p 5.0000E-01 5.2840E-01 3.2808E-01 1.9748E-01 1.1539E-01
        6.7084E-02 3.7644E-02 2.0980E-02 1.1622E-02 6.3944E-03
        3.4216E-03 1.7693E-03 0.0000E+00 -1.0000E+00

```

The ten leading blanks on these lines are edited out (as well as the blank between `wnn` and `1` on the third line), the weights inspected and changed if necessary to ensure there are no spurious fluctuations (caused by incomplete sampling), and then these lines are placed in Block 3 of the input for a second run.

4.4.3 An Example

Consider a point isotropic source emitting 7-MeV photons surrounded by an iron annular spherical shell 30-cm in thickness with an inner radius of 30 cm. The ambient dose 160 cm from the source is sought. Three approaches are used: (1) analog simulation, (2) geometry splitting, and (3) weight windows. The MCNP input file for the analog simulation is shown in Fig. 4.

With this thick iron shield, few source particles penetrate the shield, and hence we need to use some biasing technique to help particles through the shield. This problem is ideally suited for geometry splitting. To implement this, split the 30-cm spherical cell into 10 cells, each 3-cm in thickness. Examination of the output produced when this 10-cell shield problem is run as an analog simulation shows that the photon population in each shield cell decreases by about a factor of two over its neighbor closer to the source. Thus, to use geometry splitting, change the importances of the shield cells to 1 for the innermost cell, 2 for the next, 4 for the next, and so on to the tenth cell with an importance of $2^9 = 256$. This is done with the IMP command

```

c Importances: 1 src cell; 10 shld cells; 2 outer cells; 1 boundary cell
IMP:p 1 1 2m 2m 2m 2m 2m 2m 2m 2m 2R 0 $ cell importances

```

The mean and relative error with such a nonanalog simulation are shown in Figs. 5 and 6. For this nonanalog simulation, the figure-of-merit (FOM) was 4.9 times larger than that for the analog simulation, so that, to achieve the same relative error, the analog simulation would have to be run $4.9^2 = 24$ times longer.

An alternative approach is use weight windows for the 10-cell shield model. First run the problem as an analog problem with the weight window generator command placed in Block 3. Here we use

```

WWG 2 10 0 0 0 0 $ ask WW generator to find weights

```

Then place the generated weight-window cards written near the bottom of the output (with the appropriate blanks removed) into Block 3 and rerun as a weight-window biased simulation. The results are also shown in Figs. 5 and 6. For this weight-window simulation, the figure-of-merit (FOM) was 5.4 times larger than that for the analog simulation.

4.4.4 Exponential Transform

The exponential transform artificially changes the distance to the next collision. In this technique, particles can be moved preferentially towards the tally region and inhibited from moving away from it. The exponential transform stretches the path length between collisions in a preferred direction by adjusting the total cross section as $\Sigma_t(1 - p\mu)$ where p is the stretching parameter, and μ is cosine of the angle between the particle direction and the preferred direction.

p. 2-144,
3-40

The exponential transform biasing is invoked with the EXT and VECT commands. The EXT command has the form

```

EXT:n A1 A2 ... Ai ... Aj

```

where $n = N/P/E$ for neutrons/photons/electrons, and for the i -th cell A_i has the form $Q V m$, and j is the number of cells. Usually $Q = p$, while $Q = 0$ indicates the exponential transform is not to be used (and $V m$ is omitted). The stretching direction is specified by the $V m$ part of A_i and the VECT command (see manual 3-31 for examples).




Figure 5. The mean of the tally for the example problem.

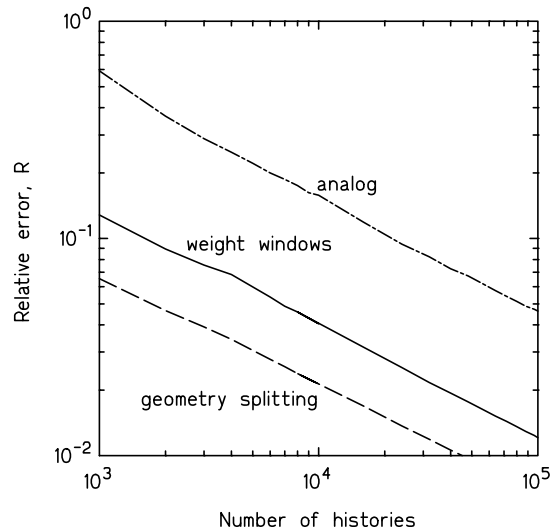


Figure 6. The relative (fractional) error for the test problem.

4.4.5 Energy Splitting/Russian Roulette

In some problems, e.g., finding the high-energy neutron fluence in a pressure vessel, only particles with a certain range of energies are of interest. When such a particle is created, the ESPLT command can be used to split the particle into more daughter particles of the same type. Also, when a particle of energy outside the energy region of interest is created, Russian roulette is used to eliminate some of these particles. An example of the ESPLT command is

p. 2-138,
3-34

```
c   Energy splitting with Russian roulette
c   split to 4 for 1 if parent energy falls below 3 MeV
c   split to 2 for 1 if parent energy falls below 1 MeV
c   split to 1 for 2 if parent energy falls below 0.4 MeV
c   split to 1 for 4 if parent energy falls below 0.1 MeV
ESPLT:n 4 3 2 1 0.5 0.4 0.25 0.1
```

4.4.6 Forced Collisions

The forced collision biasing method increases the sampling of collisions in specified cells, generally those near a DXTRAN sphere or point/ring detector. This method splits particles into collided and uncollided parts. The collided part is forced to interact within the current cell while the uncollided particle exits the cell without collision. The weight windows game is not played at surfaces bounding a cell in which forced collisions are specified. The forced collision option is invoked with the command

p. 2-147,
3-42

```
FLC:n x1 x2 ... xi ... xj
```

where $n = N/P$ for neutrons/photons, j is the number of cells, and x_i controls which particles undergo forced collisions (see manual for details)

4.4.7 Source Biasing

One of the easiest nonanalog techniques to implement is *source biasing*. In MCNP any of the SDEF variables can be biased. For example, source particles can be started with enhanced weights, with preferred energies, and in regions closer to the detector. One of the most useful source biasing techniques is to start particles in preferred directions, generally towards tally regions.

p. 2-148,
3-61

As an example, consider the spherical iron shell problem of Section 4.4.3. Rather than use a spherical F2 detector at 160 cm from the source, place a point detector on the x-axis 160 cm from the point source. (This is a terrible idea compared to using the surface F2 detector, but it illustrates the importance of source biasing.) Then to start particles preferentially towards the detector on the positive x-axis, we might use

```
SDEF  ERG=7.00  PAR=2  VEC=1 0 0  DIR=d1  $ bias source direction
SB1   -31  2.0                                $ exp bias exp[2mu]
```

Here source particles will be emitted with the PDF $p(\mu) = Ce^{K\mu}$ where $\mu = \cos\theta$, the cosine of the angle between the emission direction and the VEC direction (here the x-axis). C is a normalization constant $C = K/(e^K - e^{-K})$, calculated by MCNP. In this example we specify $p(\mu) = Ce^{2\mu}$ so that 50% of all source particles are emitted within 48 degrees of the x-axis. Here the forward-to-backward emission probabilities, $p(1)/p(-1) = e^{-4} \simeq 1/54.5$.

Another approach for source direction biasing is to restrict source emission to a set of nested cones about the bias direction. This discontinuous conical biasing is more time consuming to implement but can produce better results, when optimized, than can the continuous direction biasing. In some problems involving a collimated source, it must be used. Suppose we set up a set of nested cones about the source parameter VEC direction with cosines of the conical half angles $-1 < \mu_1 < \mu_2 < \dots < \mu_n < 1$. We want particles to be emitted in $\mu_{i-1} < \mu < \mu_i$ with probability p_i (here $\mu_0 \equiv -1$ and $\mu_{n+1} \equiv 1$). Then on the SDEF card place the parameter DIR= dn with the following lines placed after the SDEF card:

```
SIn  -1  mu1  mu2  ...  mu_n  1
SPn  0  f1  f2  ...  fn  fn+1
SBn  0  p1  p2  ...  pn  pn+1
```

Here f_i is the fraction of the solid angle of the i-th cone and is calculated as $f_i = [\mu_i - \mu_{i-1}]/2$.

4.5 Final Recommendations

Here are some recommendations for using the various variance reduction techniques.

- Before attempting to use variance reduction techniques for the first time, use the contemplated technique on a simple problem before using it on the practical and more complex problem. You need to get a feel for how the technique works without the confounding complexities of a difficult problem.
- One of the key parameters for assessing the effectiveness of different variance reduction techniques for your problem is the figure-of-merit (FOM). Generally, the better the improvement in the FOM, the better is the variance reduction technique.
- For deep penetration problems, use either cell importances or (preferably) weight windows to keep the particle population high in the cells of interest. Weight windows is more difficult to implement but more effective when done correctly. However, geometry splitting through cell importances is relatively safe and easier to implement.
- Use the CUT, ELPT and PHYS commands when appropriate to avoid time-consuming tracking, physic, or unimportant tally contributions. This can speed up calculational times for some problem by a factor of ten.

5 MCNP Output

The output produced by MCNP provides a wealth of information about the simulation. The skill of the analyst is in using this output to interpret the precision and acceptability of the tally results produced by the Monte Carlo run and to decide what changes need to be made to improve the tally in subsequent runs.

5.1 Output Tables

MCNP provides a wealth of information about the simulation, and a skilled user can elicit much insight from this voluminous output. By default only a small portion of all the possible output is produced. Always output are (1) input file listing, (2) summary of particle loss/creation, (3) summary of KCODE cycles (if KCODE is used), (4) tallies (if used), and (5) tally fluctuations charts. In addition, certain output tables deemed *basic* are always produced—they cannot be avoided. Other *default* tables are also generated unless turned off by the PRINT command. The various MCNP tables are listed in Table 5.

The output is changed from the default with the PRINT command in Block 3 of the input. Examples of the three forms of this command are

```
PRINT          $ produce everything
PRINT  110  20  $ basic & default tables plus Tables 110 and 20
PRINT -110 -20  $ all Tables except Tables 110 and 20
```

Table 5. Output tables available in MCNP. (*d*)=*default*; (*b*)=*basic*

Table No.	Table Description	Table No.	Table Description
10	Source information	120	Importance function analysis
20	Weight windows information	126	Cell particle activity
30	Tally descriptions	128(<i>b</i>)	Universe map
35	Coincident detectors	130	Particle weight balances
40	Material compositions	140	Neutron/photon nuclide activity
50	Cell vols & masses; surface areas	150	DXTRAN diagnostics
60(<i>b</i>)	Cell importances	160(<i>d</i>)	TFC bin tally analysis
62(<i>b</i>)	Forced coll.; expon. transform	161(<i>d</i>)	$p(x)$ tally PDF plot
70	Surface coefficients	162(<i>d</i>)	Cumulative $p(x)$ plot
72(<i>b</i>)	Cell temperatures	170	Source frequency; surface source
85	Electron range & straggling	175	Estimated k_{eff} by cycle
90	KCODE source data	178	Estimated k_{eff} by batch size
98	Physics const.& compile options	180	WWG bookkeeping summary
100(<i>b</i>)	Cross section tables	190(<i>b</i>)	WWG summary
102	S(α, β) nuclide assignment	198	WW from multigroup fluxes
110	First 50 starting histories	200(<i>b</i>)	WW generated windows

5.2 Accuracy versus Precision

With MCNP and its various variance reduction techniques, it is possible (and often the case for novice users) to produce tally results that, while very *precise*, i.e., a small relative error, are not very *accurate*. Technically, *precision* is the uncertainty (as measured by the tally variance) in the tally mean \bar{x} caused by the statistical fluctuations in the individual scores x_i of the simulated histories. By contrast, *accuracy* is a measure of how close the tally mean \bar{x} is to the true physical quantity being estimated. The difference between the true value and the expectation value of the simulation tally is called the *systematic error*, an important quantity but one that is seldom known.

Factors Affecting Accuracy:

- *The MCNP code:* This includes inaccuracies introduced by MCNP in its use of (1) physics models, (2) mathematical models, (3) uncertainties in the nuclear/atomic data, including cross sections, atomic weights, Avogadro's number, etc., and (4) coding errors. MCNP is a very mature code and these sources of error, while always present, are not generally thought to be a major concern for “standard neutron/photon problems.” Many MCNP benchmark validation problems have been analyzed and documented.

- *The MCNP model:* Improper modeling of source energy and angular distributions, poor representation of the actual geometry by the MCNP geometric model, and errors in the material compositions can lead to significant inaccuracies.
- *User errors:* Probably the most important source of inaccuracies (at least for novices) is error introduced by the user in incorrectly using program options or making errors in the input file. Similarly, a novice often misunderstands the difference between a particular tally and the physical quantity being sought.

Factors Affecting Precision:

- *Forward versus adjoint calculations:* For problems with spatially extended sources and a tally in a small region, an adjoint simulation often produces more precise results with few histories compared to a forward simulation.
- *Tally type:* The choice of tally type often greatly affects the precision of the results. For example, point detectors are often less precise than surface detectors in a scattering medium.
- *Variance reduction:* The use of different variance reduction techniques can affect the tally precision tremendously.
- *Number of histories:* The more histories run (and the greater computer effort expended) the better will be the precision of the tallies.

5.3 Statistics Produced by MCNP

MCNP produces a wealth of information about a simulation to allow the user to assess the *precision* (not the accuracy) of the result. While much of the detailed assessment performed by an experienced user depends on careful examination of the many output tables, the initial focus should be on the ten statistical indices calculated by MCNP. In this section we review these ten statistics.

5.3.1 Relative Error

Many beginners examine only the relative error R , and, while this is a very important statistic, it alone cannot decide the acceptability of the tally result. The relative error is the 1-sigma estimated uncertainty in the tally mean, i.e., $R \equiv S_{\bar{x}}/\bar{x}$, the ratio of the standard deviation of the tally mean to the mean. Here is how R is to be used to interpret the tally value:

Table 6. Interpretation of the reactive error R .

Range of R	Quality of Tally
> 0.5	Meaningless
0.2 to 0.5	Factor of a few
< 0.1	Reliable (except for point/ring detectors)
< 0.05	Reliable even for point/ring detectors

The value of R is determined by two quantities: (1) the history scoring efficiency q , which is the fraction of histories producing non-zero x_i 's, and (2) the dispersion in nonzero scores. In almost every tally, the tally PDF $f(x)$ (whose mean the tally is trying to estimate) has a delta-function at $x = 0$ representing the probability a source particle makes no contribution to the tally (e.g., a source particle is absorbed before reaching the tally region).

MCNP breaks R up into two components such that $R^2 = R_{eff}^2 + R_{int}^2$. Here R_{eff} is the spread in R caused by scoring inefficiency and R_{imp} is the intrinsic spread of the non-zero history-scoring events. If every source particle contributes to the tally ($q = 1$) then $R_{eff} = 0$; but as more and more particles produce zero score, R_{eff} increases. By contrast, R_{imp} measures the uncertainty produced

by the spread of nonzero scoring events. If some particles produce zero scores and the remainder produce the *same score*, $R_{imp} = 0$. As the scoring particles have increasingly different scores, R_{imp} increases.

The purpose of variance reduction techniques is to increase the scoring efficiency q and hence to reduce R_{eff} . At the same time we want to decrease the spread in nonzero scores, i.e. to make $f(x)$ more concentrated about its mean so that R_{imp} decreases.

5.3.2 Figure of Merit

Another important statistic generated by MCNP is the *figure of merit* (FOM), defined as

$$\text{FOM} = \frac{1}{R^2 T}, \quad (8)$$

where T is the run time. Since T varies with the computer, the same simulation performed on different machines produces different FOMs. As discussed earlier in Section 4.1.1, the FOM should remain relatively constant (except for fluctuations early in the simulation). For different variance reduction techniques, the one with the largest FOM is preferred.

5.3.3 Variance of the Variance

The estimation of the relative error R is important to indicate the precision of the tally mean. However, how accurate is the estimation of R ? To indicate the accuracy of R , MCNP estimates the relative variance of R , i.e. a variance of a variance (VOV). The VOV is defined as

$$\text{VOV} = \frac{S^2(S_x^2)}{S_x^2} = \frac{\sum_{i=1}^N (x_i - \bar{x})^4}{\left[\sum_{i=1}^N (x_i - \bar{x})^2\right]^2} - \frac{1}{N}. \quad (9)$$

where $S^2(S_x^2)$ is the variance of S_x^2 .

The VOV involves the third and fourth moments of the tally distribution $f(x)$ and is much more sensitive to fluctuations in large history scores than is R , which is based on only the first and second moments of $f(x)$. The proper sampling of infrequent but high scoring events is vital if reliable tally means are to be obtained, and for this reason the VOV is an important indicator of a reliable result.

From Eq. (9), it can be shown that the VOV should decrease as $1/N$. MCNP tests for this $1/N$ behavior in the VOV. Further, the VOV should always be less than 0.1 for all types of tallies.

5.3.4 The Empirical PDF for the Tally

MCNP also constructs the tally PDF $f(x)$ to help assess the quality of the confidence interval estimates for the tally mean. An example is shown in Fig. 7. Examination of the high-end tail of this distribution is very important for problems involving infrequent events with very high score. Three possible outcomes for such problems are possible:

1. Statistically meaningful confidence intervals are produced. This, of course, is always the desired outcome.
2. The sampling of a rare event with a very large score causes the the mean and R to increase and the FOM to decrease significantly. This situation is easily detected by observing the behavior of R and FOM in the tally fluctuation chart (TFC) produced at the end of the MCNP output. See Fig. 8 for a well-behaved example.
3. The third and most troublesome case is one that appears to be converged, based on acceptable statistical behavior of the mean, R , FOM, and the VOV, but in reality the tally mean is substantially underestimated because large scoring histories were inadequately sampled. Detecting this situation of too few large history tallies is difficult. It is for this case that MCNP performs extensive analysis of the high tally tail of the tally PDF.

nps	tally 4				fom	tally 14				fom
	mean	error	vov	slope		mean	error	vov	slope	
16000	2.5565E-19	0.1546	0.0460	0.0	13	1.6147E-20	0.1550	0.0990	0.0	13
32000	2.6267E-19	0.1057	0.0219	0.0	14	1.5614E-20	0.1098	0.0404	0.0	13
48000	2.9321E-19	0.0822	0.0129	10.0	15	1.5964E-20	0.0868	0.0228	0.0	13
64000	2.9096E-19	0.0725	0.0108	10.0	14	1.6062E-20	0.0760	0.0189	0.0	13
80000	2.9088E-19	0.0655	0.0086	10.0	14	1.6037E-20	0.0687	0.0161	4.9	13
96000	2.9487E-19	0.0595	0.0072	10.0	14	1.5578E-20	0.0631	0.0130	2.7	13
112000	2.9758E-19	0.0545	0.0061	10.0	15	1.5749E-20	0.0571	0.0105	3.0	13
128000	3.0167E-19	0.0509	0.0052	10.0	15	1.5970E-20	0.0528	0.0086	2.7	14
144000	3.0142E-19	0.0483	0.0050	10.0	14	1.5824E-20	0.0496	0.0075	2.7	14
160000	3.0284E-19	0.0461	0.0046	10.0	14	1.6205E-20	0.0465	0.0064	2.8	14
176000	3.0391E-19	0.0443	0.0042	10.0	14	1.6276E-20	0.0441	0.0056	3.2	14
192000	3.0143E-19	0.0427	0.0040	10.0	14	1.6351E-20	0.0420	0.0050	3.5	14
200000	3.0080E-19	0.0420	0.0040	10.0	14	1.6317E-20	0.0410	0.0048	3.9	14

Figure 8. Example of a tally fluctuation chart (TFC).

The main difficulty in detecting case 3 above is knowing when you have performed enough histories to make a valid estimate of the confidence interval for the tally mean. The *central limit theorem* (CLT) guarantees the tally mean will appear to be sampled from a normal distribution with a standard deviation σ/N if N is sufficiently large. The confidence intervals estimated by MCNP for the tally are based on this normality assumption. The key question is how large must N be for this assumption to be valid.

For the CLT to hold, the first two moments of the tally PDF $f(x)$, $E(x) = \int_0^\infty xf(x) dx$ and $E(x^2) = \int_0^\infty x^2f(x) dx$, must exist.¹ For the first two moments to exist, $f(x)$ must either have a finite upper tally cutoff, or decrease with x faster than $1/x^3$. It is this behavior of a proper tally PDF that MCNP tests for by analyzing the high-tally tail of the empirical PDF.

MCNP uses the highest scoring histories (the 200 largest) to estimate the slope of the PDF's high-tally tail. This is done by fitting a generalized Pareto function (with parameters a and k), namely

$$f_{\text{Pareto}}(x) = \frac{1}{a(1 + kx/a)^{1+(1/k)}}, \quad (10)$$

to the high tally events. The slope is then estimated from

$$\text{SLOPE} = 1 + \frac{1}{k} \quad (11)$$

On the output plot of the PDF, the Pareto fit is shown by string of s's, and tally mean by a row of m's.

For the high-end tail to be acceptable, i.e., a sufficient number of histories have been run so that the CLT is expected to apply, the SLOPE must be 3 (or greater). If insufficient, rare, high-scoring events have not been tallied, the SLOPE will generally not satisfy this criterion. If too few histories have been run to estimate the slope, the SLOPE is reported as 0; if the PDF falls off faster than $1/x^{10}$, the SLOPE is set to 10 (a "perfect" value).

5.3.5 Confidence Intervals

From the relative error R , MCNP estimates the confidence interval for the tally. Because the estimated mean and estimated uncertainty in the mean are correlated, the mid-point of the confidence

¹For the VOV to be finite, the third and fourth moments must also exist; however, MCNP doesn't enforce this.

interval needs to be shifted slightly from the mean. The amount of this midpoint shift, SHIFT, is proportional to the third central moment, and should decrease as $1/N$. MCNP calculates this refinement for the confidence interval.

5.3.6 A Conservative Tally Estimate

Sometimes a user wishes to make a conservative tally estimate, just in case rare high-tally events may not be completely considered. In the output, MCNP shows what would happen to the mean, R , VOV, confidence interval, etc., if the next history ($N + 1$) were the same as the largest scoring history encountered in the simulation of N histories. If large changes occur, then be very suspicious of the result.

5.3.7 The Ten Statistical Tests

The most valuable tool provided by MCNP for assessing the reliability of results is the suite of 10 statistical tests it performs on the tally. If any of the 10 tests are failed, MCNP automatically produces additional output to aid the user in interpreting the seriousness of the failed test(s). The 10 tests are summarized below.

Tally Mean, \bar{x} :

1. The mean must exhibit, for the last half of the problem, only random fluctuations as N increases. No up or down trends must be exhibited.

Relative Error, R :

2. R must be less than 0.1 (0.05 for point/ring detectors).
3. R must decrease monotonically with N for the last half of the problem.
4. R must decrease as $1/\sqrt{N}$ for the last half of the problem.

Variance of the Variance, VOV:

5. The magnitude of the VOV must be less than 0.1 for all types of tallies.
6. VOV must decrease monotonically for the last half of the problem.
7. VOV must decrease as $1/N$ for the last half of the problem.

Figure of Merit, FOM:

8. FOM must remain statistically constant for the last half of the problem.
9. FOM must exhibit no monotonic up or down trends in the last half of the problem.

Tally PDF, $f(x)$:

10. The SLOPE determined from the 201 largest scoring events must be greater than 3.

If any of these tests fails, a warning is printed in the output and a plot of $f(x)$ is produced. If all ten tests are passed, MCNP then calculates asymmetric and symmetric confidence intervals for the mean at the one-, two-, and three-sigma levels. While these ten statistical tests provided an excellent indication of the reliability of the result, *they are not foolproof*. There is always the possibility that some high-scoring rare event was not sampled in the histories run and that the tally is underestimated. Users must rely on their understanding and insight into the particular problem to avoid such traps.

5.3.8 An Example Problem

Consider a point isotropic source of 7-MeV photons in an infinite medium of iron. The ambient dose equivalent 30 cm from the source is sought. A surface F2 detector and a F5 point detector are both used to estimate this dose. The input file is shown in Fig. 9.

The variation of the tally mean, R , VOV, SLOPE, and FOM with the number of particle histories is shown in Fig. 10. At 10^7 particle histories, the F2 tally passed all 10 statistical tests: the mean and FOM are relatively constant, the relative error R is monotonically decreasing as $1/\sqrt{N}$ (1 y-decades for every 2 x-decades), the VOV is monotonically decreasing as $1/N$ (1 y-decades for every 1 x-decades), and the SLOPE has a “perfect” value of 10. This high slope value is to be expected since there physically is an upper limit to the tally, namely that produced by an uncollided photon reaching the scoring surface. The slope of 10 is a strong indicator of such a tally cutoff.

By contrast, the F5 point detector has not converged. The mean, error, VOV, and FOM all exhibit sudden changes, a result of an occasional photon that scatters very near the point detector and contributes a huge score with the last-flight estimator used by the F5 tally. The SLOPE remains constant at about 2, an indication of a long slowly decreasing high-tally tail. In fact, a point detector in a scattering medium has no tally cutoff. Be very wary of using point/ring detectors in a strongly scattering medium.

If you had performed this simulation for only 5000 histories and, unwisely, looked only at the relative errors, the F5 detector would appear attractive since it has a relative error of about 0.06 (almost near the acceptable value) while the F2 tally has $R > .2$. Recall that every source particle produces a score with a point detector ($q = 1$, $R_{eff} = 0$) and R often starts to decrease properly. The F2 tally, on the other hand, received scores from only 0.8% ($q = 0.0080$) of the source particles leading to $R_{eff} = 0.0024$ and $R_{imp} = 0.0024$ after 10^7 histories. Because of the large fluctuations in the F5 scores, its R_{imp} is much larger (0.127 after 10^7 histories).

The PDFs for these two tallies are shown in Figs. 11 and 12. As expected, the PDF for the F5 tally is spread out over a wide range of scores and has a high-score tail that is poorly defined even after 10^7 histories. The PDF for the F2 tally is much more compact with a well established upper cutoff.

```

Point isotropic 7-MeV photon sources in infinite iron medium
c ***** BLOCK 1: CELL CARDS *****
  1 1  -7.86  -10      imp:p=1  $ iron inside detector shell
  2 1  -7.86  10 -20   imp:p=1  $ iron outside detector shell
  3 0           20     imp:p=0  $ vacuum outside problem boundary

c ***** BLOCK 2: SURFACE CARDS *****
  10 so 30.0          $ detector surface
  20 so 3000.0       $ outer surface of iron

c ***** BLOCK 3: DATA CARDS *****
SDEF  erg=7.00  par=2      $ pt isotropic 7-MeV photon source
mode  p                    $ photon mode only
nps 1000000              $ number of histories to be run
f2:p  10                  $ tally 2:  surface detector at 30 cm
f15:p 30 0 0 -0.3        $ tally 15: pt det 30 cm on x-axis; Ro=.3mfp
c
c --- Photon ambient dose equivalent H*(10mm) Sv cm^2; from ICRP [1987]
de  0.100E-01 0.150E-01 0.200E-01 0.300E-01 0.400E-01 0.500E-01
    0.600E-01 0.800E-01 0.100E+00 0.150E+00 0.200E+00 0.300E+00
    0.400E+00 0.500E+00 0.600E+00 0.800E+00 0.100E+01 0.150E+01
    0.200E+01 0.300E+01 0.400E+01 0.500E+01 0.600E+01 0.800E+01 0.100E+02
df  0.769E-13 0.846E-12 0.101E-11 0.785E-12 0.614E-12 0.526E-12
    0.504E-12 0.532E-12 0.611E-12 0.890E-12 0.118E-11 0.181E-11
    0.238E-11 0.289E-11 0.338E-11 0.429E-11 0.511E-11 0.692E-11
    0.848E-11 0.111E-10 0.133E-10 0.154E-10 0.174E-10 0.212E-10 0.252E-10
c
m1  26000 -1.00000      $ natural iron (density 7.86 g/cm^3)

```

Figure 9. Input file for the example problem




Figure 10. The variation of the various statistics estimated by MCNP for the two tallies of the test problem.




Figure 11. The PDF for the F2 surface tally in the example problem. Heavy line is for 10^6 histories and dotted line for 10^7 .




Figure 12. The PDF for the F5 tally in the example problem. Light line is for 10^6 histories and the heavy line for 10^7 .