LA-UR 93-1392

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE HYBRID SN/MONTE CARLO RESEARCH AND RESULTS

AUTHOR(S) Randal S. Baker

SUBMITTED TO Advanced Monte Carlo Computer Programs for Radiation Transport, 27-29, April 1993, Saclay, France

DISCLAIMER

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

HYBRID S_N/MONTE CARLO RESEARCH AND RESULTS

R. S. Baker
Radiation Transport Group, MS B226
Los Alamos National Laboratory
Los Alamos, New Mexico 87545 USA

ABSTRACT

The neutral particle transport equation is solved by a hybrid method that iteratively couples regions where deterministic (S_N) and stochastic (Monte Carlo) methods are applied. The Monte Carlo and S_N regions are fully coupled in the sense that no assumption is made about geometrical separation or decoupling. The hybrid Monte Carlo/S_N method provides a new means of solving problems involving both optically thick and optically thin regions that neither Monte Carlo nor S_N is well suited for by themselves. The hybrid method has been successfully applied to realistic shielding problems.

The vectorized Monte Carlo algorithm in the hybrid method has been ported to the massively parallel architecture of the Connection Machine. Comparisons of performance on a vector machine (Cray Y-MP) and the Connection Machine (CM-2) show that significant speedups are obtainable for vectorized Monte Carlo algorithms on massively parallel machines, even when realistic problems requiring variance reduction are considered. However, the architecture of the Connection Machine does place some limitations on the regime in which the Monte Carlo algorithm may be expected to perform well.

INTRODUCTION

TWODANT/MC[1-3] is an extension to the existing S_N code TWODANT. TWODANT/MC solves the neutral particle transport equation by a hybrid method that iteratively couples regions where deterministic (S_N) and stochastic (Monte Carlo) methods are applied. Unlike previous hybrid methods,[4] the Monte Carlo and S_N regions are fully coupled in the sense that no assumption is made about geometrical separation or decoupling. TWODANT/MC provides a new means of solving problems involving both optically thick and optically thin regions that neither Monte Carlo nor S_N is well suited for by themselves. TWODANT/MC is capable of solving forward, inhomogeneous source problems in X - Y and R - Z geometries. This capability includes multigroup problems involving upscatter and fission.

The fully coupled Monte Carlo/S_N technique used in TWODANT/MC consists of defining spatial and/or energy regions of a problem in which either a Monte Carlo calculation or an S_N calculation

is to be performed. The Monte Carlo region may comprise the entire spatial region for selected energy groups, or may consist of a rectangular area that is either completely or partially embedded in an arbitrary S_N region. The Monte Carlo and S_N regions are then connected through the common angular boundary fluxes and scattering/fission sources, which are determined iteratively using the response matrix technique.

The hybrid method has been implemented in TWODANT/MC by adding special-purpose vectorized Monte Carlo subroutines, and linkage subroutines to carry out the interface flux iterations. The common angular boundary fluxes are included in TWODANT as interior boundary sources, leaving the logic for the S_N solution of the transport flux unchanged, while the diffusion synthetic accelerator remains effective in accelerating the S_N calculations. The physical description of the Monte Carlo region is derived from the standard TWODANT input file, reducing the required user input at some expense in generality, since this limits the Monte Carlo region to two-dimensional $X - Y$ and $R - Z$ grids. The S_N cross sections are used to form multigroup cross sections for the Monte Carlo region.

THEORY

Let ${}^{*}\psi^{in}$ represent the flux entering the Monte Carlo region from the S_N region, and ${}^{+}\psi^{out}$ the flux entering the S_N region from the Monte Carlo region. Then, the outgoing flux from the Monte Carlo region is related to the incoming flux from the S_N region by

$$ {}^{+}\vec{\psi}^{out} = {}^{+}\mathbf{R}^{*}_{mc} \, {}^{*}\vec{\psi}^{in} + {}^{+}\vec{S}^{out}, \tag{1} $$

where ${}^{+}S^{out}$ is the exiting flux from the Monte Carlo region under vacuum boundary conditions (${}^{*}\psi^{in} = 0$) and ${}^{+}\mathbf{R}^{*}_{mc}$ is the Monte Carlo response matrix.

The fluxes entering (${}^{*}\psi^{in}$) and leaving (${}^{+}\psi^{out}$) the Monte Carlo region are represented numerically as N_{in} and N_{out} length vectors, respectively. Since the Monte Carlo/S_N boundary may consist of both spatial and energy interfaces, N_{in} does not necessarily equal N_{out}. That is, although the number of discretized states along the spatial boundary (energy x spatial x discrete direction) are the same for ${}^{*}\psi^{in}$ and ${}^{+}\psi^{out}$, the number of discretized states along the energy boundary (energy x spatial x angular moment) are not, since a different number of Monte Carlo energy groups may be reachable from collisions in the S_N groups than vice versa (see Figure 1 for an example). Here, we use an asterisk (*) to indicate a dimension of length N_{in} and a plus (+) for a dimension of length N_{out}.

Equation (1) is solved iteratively by

$$ {}^{+}\psi^{out(p+1)} = {}^{+}\mathbf{R}^{*}_{mc} \, {}^{*}\psi^{in(p)} + {}^{+}S^{out}, \tag{2} $$

where we set ${}^{+}\psi^{out(1)} = {}^{+}S^{out}$, and ${}^{*}\psi^{in(p)}$ is obtained by using an S_N solver (TWODANT) with the prescribed boundary flux ${}^{+}\psi^{out(p)}$. Note that for problems involving upscatter/fission, the S_N calculation includes multiple outer iterations, as usual.

Since ${}^{+}\mathbf{R}^{*}_{mc}$ is an $N_{out} \times N_{in}$ sized matrix and, for multigroup problems in two dimensions, N_{out} and N_{in} are typically on the order of 10^5 to 10^6, the direct calculation/storage of the Monte Carlo response matrix is not feasible. Instead, the response matrix is represented by P orthogonal

interface basis vectors,[5] where P typically ranges from 4 to 8. This reduces the size of the Monte Carlo response matrix to $N_{out} \times P$. Furthermore, after $P+1$ iterations, a $P \times P$ representation of the overall S_N/Monte Carlo response matrix can be obtained. Then, since P is small, the converged interface boundary fluxes can be predicted using direct inversion.

IMPLEMENTATION

Each interface state on the Monte Carlo/S_N boundary corresponds to a unique spatial cell, energy group, and angular direction, for spatial interfaces, or angular moment, for energy interfaces. An angular bin $\Delta\vec{\Omega}^m$ is used to assign a range of continuous Monte Carlo directions to the discrete S_N direction $\vec{\Omega}^m$. Since S_N quadrature sets specify the angular weight (the area subtended on the unit sphere by $\Delta\vec{\Omega}^m$) for each S_N direction, but not its shape, there is considerable freedom in defining the $\Delta\vec{\Omega}^m$. The system we use is depicted in Fig. 2. Note that as a particle streams in $R-Z$ geometry, $\vec{\Omega}_z$ remains constant. The bin configuration we have chosen preserves this property, since streaming particles can only flow into bins whose quadrature directions have the same $\vec{\Omega}_z$.

As explained in Ref. 1, the Monte Carlo region ideally comprises a zone of low scattering material, surrounded by a boundary layer (~ 1 mfp thick) of highly scattering material. Although eliminating the boundary layer from the Monte Carlo region would speed up the Monte Carlo calculations, it would also require using a higher quadrature order in the S_N calculation for accurate results, since the angular flux near a low scattering region can be a rapidly varying function of angle and position. By including a boundary layer with the Monte Carlo region, a lower S_N quadrature order can be used. Note that for multigroup problems, the mean free path may vary from group to group, resulting in a Monte Carlo region of varying size (see Fig. 1).

The Monte Carlo routines in TWODANT/MC use the macroscopic multigroup S_N cross sections, so that no additional cross section input is required. The treatment of angular scattering is somewhat complicated, however, because the direct sampling of the group to group transfer moments is usually not possible since they may be negative. Instead, TWODANT/MC generates I equiprobable bins, with L bin boundaries, such that the L Legendre moments of the angular scattering cross sections are (approximately) conserved. That is, in a P_1 approximation, the code will generate three equiprobable bins, with boundaries between $\mu = -1$ and $\mu = +1$, for sampling angular scattering directions. This approach is the same one currently used in the multigroup version of MCNP, and is further described in Ref. 6.

The special purpose Monte Carlo routines implemented with TWODANT are basically analog routines since, ideally, the Monte Carlo region is low scattering, and should not require extensive variance reduction methods. However, several basic variance reduction techniques have been incorporated. These include implicit capture, weight cutoffs, geometry splitting/Russian roulette, energy splitting/Russian roulette, and source direction biasing.

To increase the efficiency of the Monte Carlo calculations, the Monte Carlo routines have been vectorized through the formation of event based stacks,[7] where each stack ideally consists of a group of 64 particles undergoing an identical event, such as a collision. On the Cray, assembly language routines are used to encode particle destination tag words, and to move particles between stacks. Vectorization resulted in a factor of four to five increase in speed of the vectorized Monte Carlo routines over the non-vectorized Monte Carlo routines on a Cray Y-MP.

RESULTS

The hybrid Monte Carlo/S_N method has been successfully applied to a realistic cylindrical duct shielding problem. The problem consists of a duct (void) with an isotropic ($\vec{\Omega}_z > 0$) neutron source incident at $z = 0$. The duct is surrounded by an iron liner, and an iron and water shield (see Fig. 3). Thirty energy groups are used in the problem, where the source neutrons are located in group 2 (13.5–15.0 MeV).

For the TWODANT/MC calculations, the Monte Carlo option was used only for the source energy group (Group 2), while pure S_N was used in all remaining energy groups. The void region along the centerline was designated as the Monte Carlo region, along with an additional 1 mfp boundary layer, which, for Group 2, was equivalent to 6 cm. Thus, the actual Monte Carlo region consisted of the void region, the iron liner (2 cm), and 4 cm of the shield (see Fig. 3). S_N was used in the remainder of the shield region for Group 2. The S_N calculation used an S_{12} quadrature order, with a P_4 Legendre expansion.

The neutron fluxes from the TWODANT/MC calculations, as well as the results from a multigroup MCNP calculation, are shown at Table 1, for selected regions of interest. The relative error in the MCNP results is shown in parentheses, while the difference (in terms of standard deviations) from the MCNP results is shown for the TWODANT/MC results. Note that the S_N/Monte Carlo results are all within three standard deviations of the multigroup MCNP results, and that most fall within two standard deviations.

MASSIVELY PARALLEL SIMD IMPLEMENTATION

The Monte Carlo algorithm in TWODANT/MC has been ported to the massively parallel SIMD architecture of Thinking Machines Corporation's CM-2/CM-200.[8] Although the coding was rewritten in CM FORTRAN, the Monte Carlo algorithm itself, which had already been vectorized as described above, was essentially unchanged. The most significant problem encountered in porting the algorithm to the CM was the expense of communications. With the event stack method of vectorization, we have found that the time spent in moving particles between different event stacks typically averages around 10% for most problems on a Cray Y-MP with shared memory. Due to the distributed memory architecture of the CM, however, particle movement between stacks usually requires general ("router") communication, since a particle located in a particular position in Stack A may need to be moved to a different position in Stack B, and the particle movement pattern is essentially random. Timing tests revealed that the problem time on the CM was entirely dominated by the time spent in moving particles between stacks, with an adverse affect on performance.

To remedy this, a two-step process was used to move particles between stacks. If in moving a particle from Stack A to Stack B, the corresponding location in Stack B (i.e., the element in Stack B located on the same VP (Virtual Processor) as Stack A's particle) is empty, the particle is moved directly into Stack B without any required communications. This is possible since all event stack arrays are allocated as conformable arrays. If the corresponding location in Stack B is already filled, only then are general communications used to move the particle to an empty location in Stack B. Typically, 80% of the particle movement occurs "in place" for most problems, and only 20% requires general communications. However, even with these changes, communications still consumes over 50% of the total time for most problems.

To compare the performance of the vectorized Y-MP version of the algorithm against the parallelized CM version, we selected a sample two-dimensional well logging problem in $R-Z$ geometry (see Fig. 1), which we solved entirely by the Monte Carlo method. Given an (α,n) neutron source (AmBe) located in a tool positioned in the center of a borehole, the problem is to transport particles from the source out into the surrounding formation, then back into one of two ^{3}He detectors located in the tool. The detectors are protected by shields to prevent any direct contribution from the source. From the relative response of the two detectors, an analyst can then determine the porosity of the surrounding formation. The problem was run with geometry splitting/Russian roulette, implicit capture, source direction biasing, a .4 eV energy cutoff (for epithermal neutron porosity tools, the detectors are surrounded by cadmium foil), and thirty energy groups. At Figure 5, we present the speed (Source Particles/Minute (SPM)) for various VP ratios and number of source histories/batch on the CM-2, normalized to the speed of the Y-MP version (single processor). The abscissa gives the VP ratio, where a VP ratio of 2 corresponds to a stack size of 4096 particle events, since all CM-2 calculations were performed on a 2048 Processing Element (PE) machine. The ordinate is the ratio of the speed on the CM-2 to that of the Y-MP. The lines represent varying numbers of source histories/batch, from 80,000 to 2,560,000. We see that at low VP ratios (< 16), the performance of the CM-2 suffers. As the VP ratio increases for a fixed number of source histories/batch, the performance of the CM-2 increases, passes through a maximum, and then decreases. Also, although the performance at low VP ratios is relatively unaffected by the number of source histories/batch, that at higher VP ratios varies dramatically. In general, however, the performance of the massively parallel algorithm on the CM-2 is comparable to that of the vectorized version on a single processor Y-MP.

The performance of the CM-2 on the left hand side of the graph is due to the nature of the CM-2 hardware, since each Floating Point Accelerator associated with a PE has a vector length of 4. Thus, efficient use of the CM-2 requires a minimum VP ratio of 4. To amortize the overhead associated with code execution, actual peak efficiency for the CM-2 is not usually achieved until VP ratios of 32 or 64. This is evident from the top (solid) line in Figure 5.

However, as the VP ratio goes up, the stack length also increases, and thus more particles are necessary to fill a stack. At some point, one runs out of source particles, and is forced to execute a less than full stack. This results in a degradation in stack efficiency. Obviously, the fewer the amount of source particles, and the more the last remaining particles "straggle", the lower the overall stack efficiency. As an example, the average collision stack efficiency for 80,000 source particles/batch ranged from 79%, at a VP ratio of 2, to just 6% at a VP ratio of 128. In contrast, the collision stack efficiency for 2,560,000 source particles/batch only varied from ~100%, at a VP ratio of 2, to 72%, at a VP ratio of 128. Stack efficiency accounts for the drop off in performance on the right hand side of the graph.

Currently, for actual hybrid Monte Carlo/S_N calculations, the Monte Carlo calculation runs on the Connection Machine, while the S_N runs on the SUN front end. Thus, the total calculation time on the CM for hybrid problems is greater than that on a Cray, due to the much slower processing of the S_N calculation by the SUN. However, we have recently developed a version of the code where the S_N calculation is performed on a Cray, the Monte Carlo calculation on a CM-2, and PVM (Parallel Virtual Machine) is used to link the calculations. This version is operational and undergoing testing, and we hope to report results in the near future.

CONCLUSIONS

TWODANT/MC provides a new means of solving inhomogeneous source problems in $X - Y$ or $R - Z$ geometry that neither Monte Carlo nor S_N is well suited for by themselves. These problems typically consist of those with a singular type source, such as a point or beam source, and/or large amounts of streaming through voids, which pose difficulties for S_N calculations, and also contain regions with large amounts of scattering, which increase the expense of Monte Carlo calculations. For these types of problems, the hybrid Monte Carlo/S_N method can perform more efficiently than either the S_N or Monte Carlo methods can alone. Since the Monte Carlo and S_N regions are fully coupled through the response matrix technique, the hybrid Monte Carlo/S_N method can be used even in situations where the Monte Carlo and S_N regions are closely coupled.

ACKNOWLEDGEMENTS

The hybrid Monte Carlo/S_N method was developed in conjunction with Raymond E. Alcouffe, Los Alamos National Laboratory, and William L. Filippone, University of Arizona. The cylindrical duct test problem and MCNP results were provided by William T. Urban, Los Alamos National Laboratory.

REFERENCES

1. "The S_N/Monte Carlo Response Matrix Hybrid Method", W. L. Filippone, and R. E. Alcouffe, *Nucl. Sci. & Eng.*, **100**, p. 209 (1988).

2. "The Multigroup and Radial Geometry Formulation of the Monte Carlo/S_N Response Matrix Method", R. S. Baker, W. L. Filippone, and R. E. Alcouffe, *Nucl. Sci. & Eng.*, **105**, p. 181 (1990).

3. "Extension of the Fully Coupled Monte Carlo/S_N Response Matrix Method to Problems Including Upscatter and Fission", R. S. Baker, W. L. Filippone, and R. E. Alcouffe, Int. Topical Meeting on Advances in Math., Comp., and Reactor Phy., **5**, p. 21.2 3 1, Pittsburgh, PA (1991).

4. "Domino, A General Purpose Code for Coupling Discrete Ordinates and Monte Carlo Radiation Transport Calculations", M. B. Emmett, C. W. Burgart, and T. J. Hoffman, ORNL 4853 (1973).

5. "The Monte Carlo/S_N Hybrid Method With Orthogonal Interface Basis Vectors", W. L. Filippone, R. S. Baker, and R. E. Alcouffe, to be published, Joint Intl. Conf. on Math. Methods and Supercomputing in Nuclear Applications, Karlsruhe, Germany (1993).

6. "Transfer Matrix Treatments for Multigroup Monte Carlo Calculations — The Elimination of Ray Effects", T. J. Carter and C. A. Forest, *Nucl. Sci. & Eng.*, **59**, p. 27 (1976).

7. "Vectorized Monte Carlo Photon Transport", F. W. Bobrowicz, J. E. Lynch, K. J. Fichet, and J. L. Labor, LANL Report LA 9752 MS (1983).

8. "Implementation of a Monte Carlo Algorithm for Neutron Transport on a Massively Parallel SIMD Machine", R. S. Baker, to be published, Joint Intl. Conf. on Math. Methods and Supercomputing in Nuclear Applications, Karlsruhe, Germany (1993).

Table 1

Neutron Fluxes for Cylindrical Duct Problem

Region	Flux	MCNP/MG	TWODANT/MC
I	Group 2	2.06-3 (.003)	2.05-3 (-1.6σ)
	Total	5.14-3 (.006)	5.15-3 (0.3σ)
II	Group 2	3.31-7 (.022)	3.35-7 (0.5σ)
	Total	9.16e-6 (.015)	9.04-6 (-3.0σ)
III	Group 2	4.42-5 (.019)	4.28-5 (-1.7σ)
	Total	2.28-4 (.021)	2.20-4 (-1.7σ)
IV	Group 2	1.11-8 (.042)	1.03-8 (-1.7σ)
	Total	4.07-7 (.038)	3.92-7 (-1.0σ)
V	Group 2	7.18-6 (.041)	6.92-6 (-0.9σ)
	Total	2.20-5 (.059)	1.91-5 (-2.2σ)
VI	Group 2	3.98-5 (.028)	3.82-5 (-1.4σ)
	Total	5.01-5 (.052)	4.63-5 (-1.5σ)

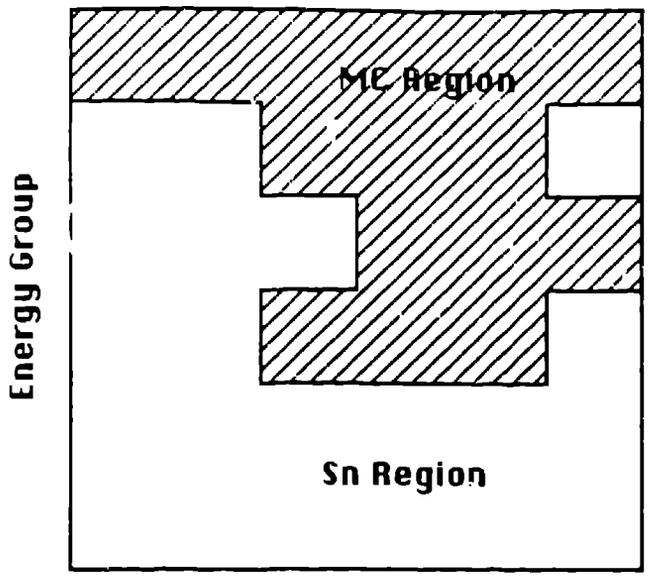
MC Region
Energy Group
Sn Region

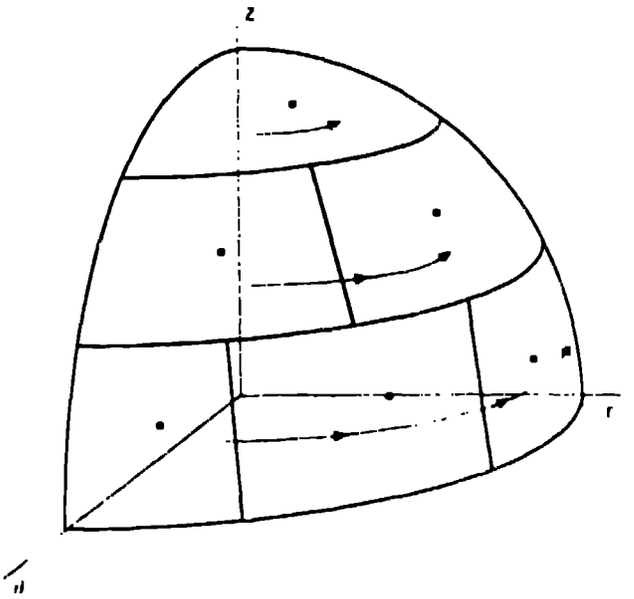

Figure 5–3 The S_N Angular Bin Arrangement for $N = 6$

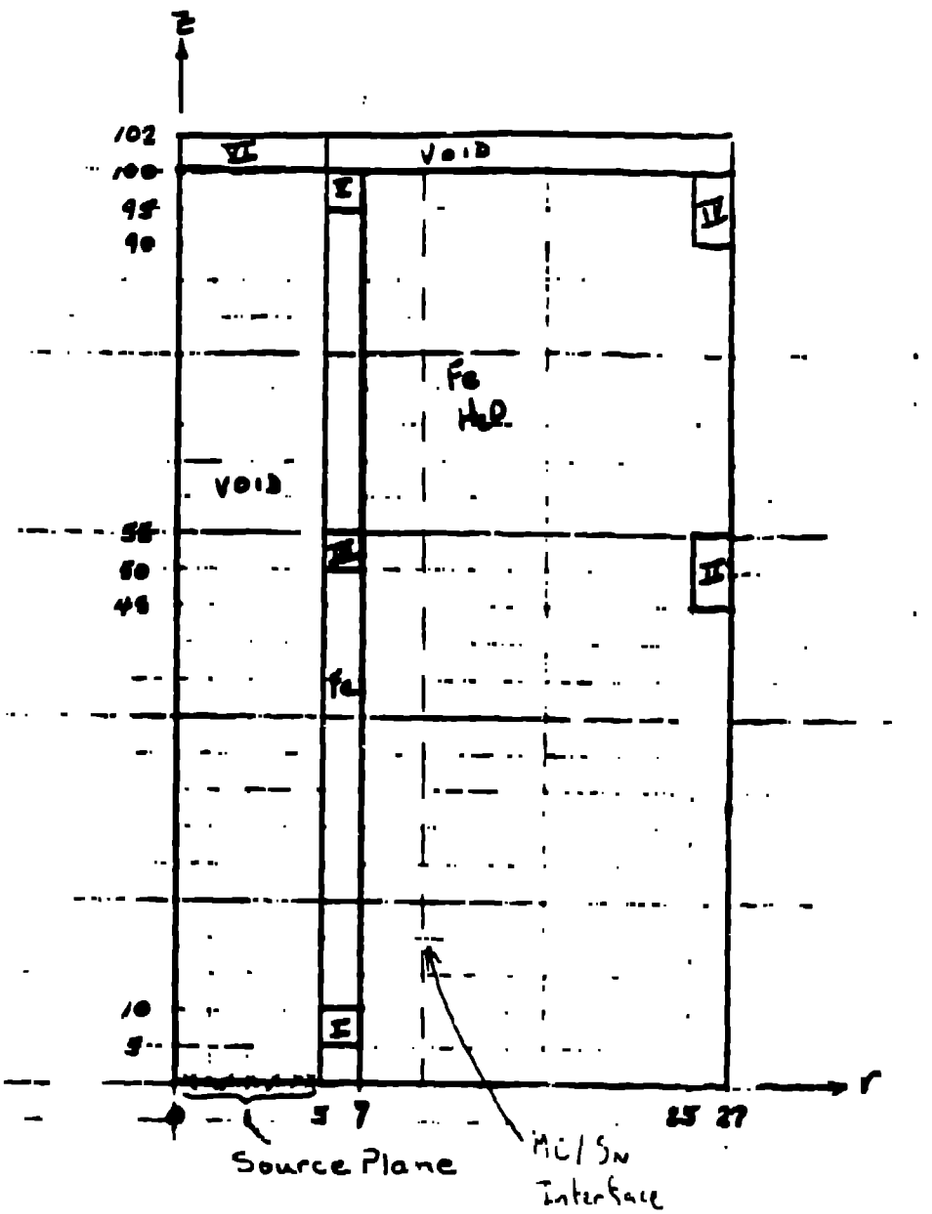

Fig. 1. Schematic of Test Problem 1 with dimensions in centimeters. (not to scale.)

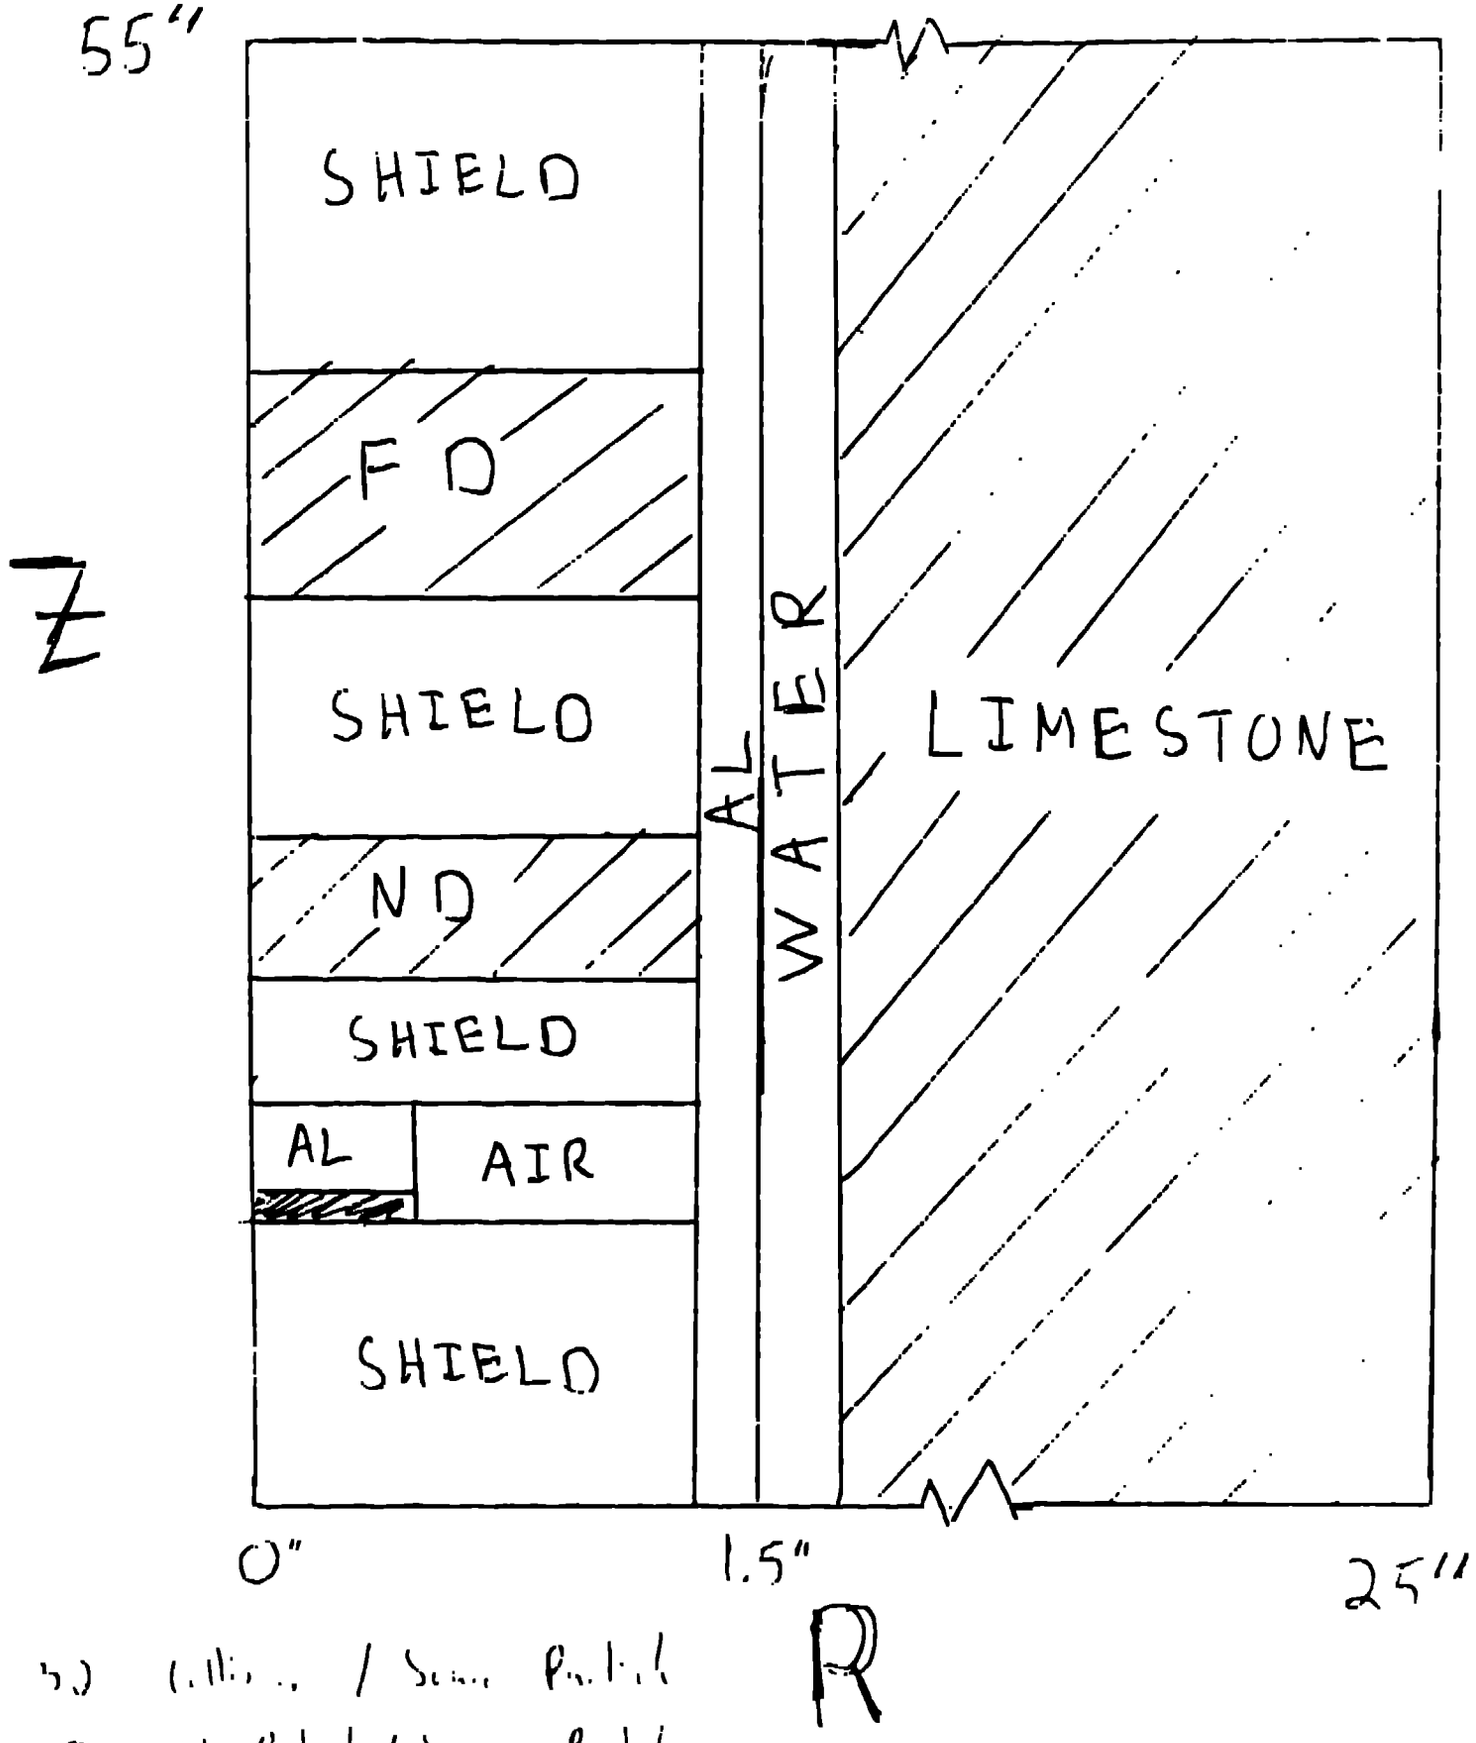
55"
SHIELD
F D
SHIELO
ND
SHIELD
AL
AIR
SHIELO
AL
WATER
LIMESTONE
Z
0"
1.5"
25"
R

2048 FPU CM2

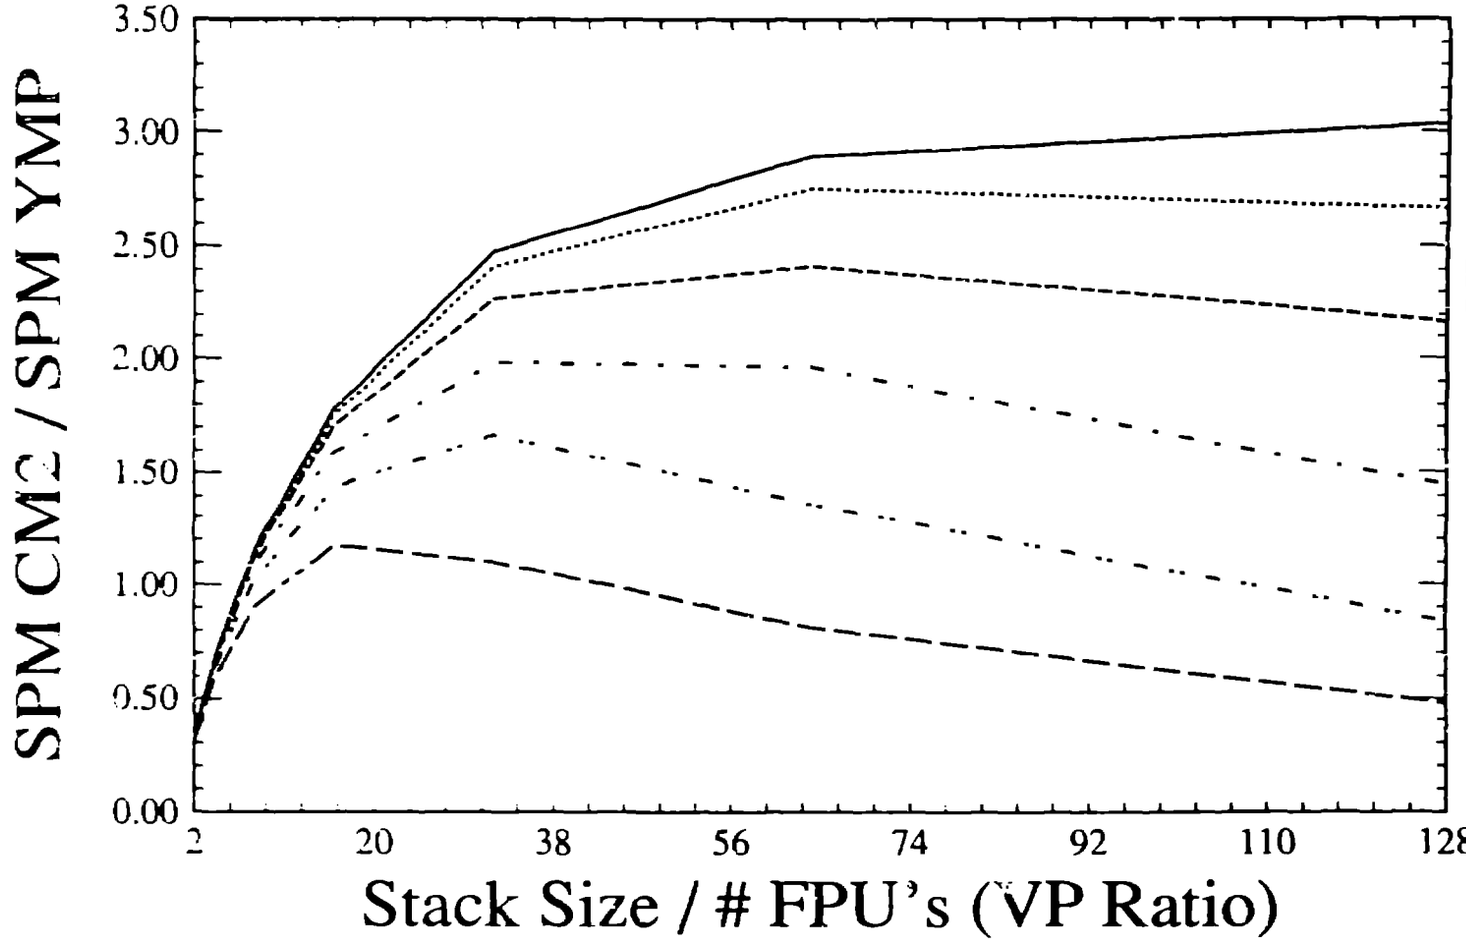
SPM CM2 / SPM YMP
3.50
3.00
2.50
2.00
1.50
1.00
0.50
0.00
2
20
38
56
74
92
110
128
Stack Size / # FPU's (VP Ratio)
2560K p/batch
1280K p/batch
640K p/batch
320K p/batch
160K p/batch
80K p/batch