Progress in Nuclear Energy 80 (2015) 110–118

Contents lists available at ScienceDirect

Progress in Nuclear Energy

journal homepage: www.elsevier.com/locate/pnucene

Model and analysis of performance for the method of characteristics direction probabilities with boundary averaging

Zhouyu Liu [a, b, *], Brendan Kochunas [c], Benjamin Collins [c], Thomas Downar [c], Yunlin Xu [c], Hongchun Wu [a]

[a] School of Nuclear Science and Technology, Xi'an Jiaotong University, No. 28 Xianning West Road, Xi'an, Shaanxi 710049, PR China
[b] Shanghai Nuclear Engineering Research & Design Institute, 29 Hongchao Road, Shanghai 200233, PR China
[c] Department of Nuclear Engineering and Radiological Sciences, University of Michigan, 2200 Bonisteel, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:
Received 17 August 2014
Received in revised form
12 November 2014
Accepted 4 December 2014
Available online 5 January 2015

Keywords:
CDP
MOC
Boundary average

ABSTRACT

The work presents a performance model of the method of characteristic direction probabilities (CDP) which integrates the benefits of the collision probability method (CPM) and the method of characteristics (MOC) for solution of the integral form of the Boltzmann Transport Equation and has been implemented in the Michigan PArallel Characteristic based Transport (MPACT) code for 2-D and 3-D transport calculations. The process of boundary averaging reduced the storage and computation but the capability of dealing with complicated geometries is preserved since the same ray tracing information is used as in MOC. The benefits of CDP are demonstrated by the developed performance model which describes the storage, floating point operations and calculation time. The numerical results are given for different cases to show the accuracy, storage, floating point operations and computing time of the CDP compared to the MOC using the performance model. From the cases examined, the boundary average method shows significant improvement on the storage and computational efficiency for three-dimensional cases with sufficient accuracy.

1. Introduction

The method of characteristics (MOC) based on the modular ray tracing technique (Liu et al., 2011) has been implemented in Michigan PArallel Characteristic based Transport code (MPACT) to perform lattice and whole core calculations for LWR applications. The MOC uses a set of discrete ordinates, which is similar to the S_N methods, but MOC is better suited to treat complicated geometries because it only requires an approximation on the spatial variation of the source, and not on the flux itself along the rays. However, the transport sweep needs to be performed along all the characteristics lines for every direction, and this sweeping time will be computationally expensive when calculating three-dimensional problems, in which the number of characteristic rays could be quite large to accurately present very thin regions using burnable absorbers such as IFBA that coat the fuel pin.

In the previous work (Liu et al., 2013), the method of characteristic direction probabilities, which was first proposed by (Hong and Cho, 1999), was implemented in MPACT to minimize the computation efforts. This new transport method couples the desirable features of the MOC and the collision probability (CPM) (Sanchez, 1997). CPM has been widely used in lattice physics codes because it has the capability of treating the complicated geometries and is very efficient when dealing with small size problems. But this method has the drawback that the storage requirements and computing time depend on the square of the number of fine spatial regions in the problem. This is because the collision probability matrix couples all the fine mesh regions. To overcome the drawback of CPM when dealing with big size problems, the interface current method (ICM) (Mohanakrishnan, 1981) was developed which couples the sub-domains with interface current of interface current moments, and within the domains the fine regions are coupled by the CPM. However, compared to the interface current method, CDP doesn't introduce the approximation at the interface of the sub-domain and the anisotropic sources. Another drawback of CPM is that it cannot easily treat anisotropic sources. In the CDP, only fine regions traversed by a characteristic line within a specified sub-domain are coupled which is the most significant difference with

* Corresponding author. School of Nuclear Science and Technology, Xi'an Jiaotong University, No. 28 Xianning West Road, Xi'an, Shaanxi 710049, PR China.
E-mail addresses: liuzhouyu1985@hotmail.com (Z. Liu), bkochuna@umich.edu (B. Kochunas), bscollin@umich.edu (B. Collins), downar@umich.edu (T. Downar), yunlin@umich.edu (Y. Xu), hongchun@mail.xjtu.edu.cn (H. Wu).

http://dx.doi.org/10.1016/j.pnucene.2014.12.004
0149-1970/

CPM. At the same time, the CDP is capable of providing the same accuracy as MOC if the unique boundary sub-domains are the same size as the MOC ray spacing. The only difference in the methods then would be that instead of performing the transport sweep ray by ray of the MOC, the CDP method obtains the outgoing angular flux and fine region flux by direct multiplication of a matrix which contains the collision and transmission probabilities and a vector which includes the coming angular flux and the fine region source. The collision and transmission probabilities in the CDP are derived by integrating the traditional MOC equations along a characteristic line. So in principle, the method of characteristics direction probabilities is mathematically consistent with the conventional MOC. To further improve the efficiency of the CDP, the boundary averaged ray tracing technique was introduced which can reduce the memory required for storing the probabilities and improve computing efficiency.

A performance model is described in this work to explicitly analyze the storage requirement, floating point operations and computing time. Based on the analysis, we can found where and how much the boundary-averaged CDP earns the profit. The examined numerical results proved the consistence of the performance model to the measured results and showed the advantage of the boundary-averaged CDP.

In the following section the basic equations of the MOC are provided along with the derivation of the CDP method. Also described in this section is boundary average scheme. The third section introduces the performance model of the CDP and numerical results are shown in the subsequent section. The final section provides a summary and conclusions.

2. The method of characteristics direction probabilities

2.1. The method of characteristics

The classical method of characteristics for solving partial differential equations has been successfully applied to the Boltzmann Transport equation (BTE) and implemented in several reactor analysis codes. The group-wise form of the BTE for the system R is given by.

$$\Omega\cdot\nabla\varphi_g(\mathbf{r},\mathbf{\Omega})+\Sigma_{t,g}\varphi_g(\mathbf{r},\mathbf{\Omega})=Q_g(\mathbf{r},\Omega), \tag{1}$$

where $Q_g(\mathbf{r})$ is total source including both the fission source term and the scattering source terms.

The MOC equation provides a solution of the Boltzmann Transport equation along a line in a particular direction and it reduces to the total differential Equation (2) which is simplified by removing the energy group subscript g.

$$\frac{d\varphi(\mathbf{r}_0+s\mathbf{\Omega}_m,\mathbf{\Omega}_m)}{ds}+\Sigma(\mathbf{r}_0+s\mathbf{\Omega}_m)\varphi(\mathbf{r}_0+s\mathbf{\Omega}_m,\mathbf{\Omega}_m)=Q(\mathbf{r}_0+s\mathbf{\Omega}_m,\mathbf{\Omega}_m), \tag{2}$$

where $\mathbf{r}_0$ is the starting point of a characteristic line and s is the distance from the initial point to the current point along a specified direction $\mathbf{\Omega}_m$.

When solving the equation, we assume that the source and properties are constant in a small region D_i.

$$Q(\mathbf{r},\mathbf{\Omega}_m)=Q_i(\mathbf{\Omega}_m),\quad \Sigma_t(\mathbf{r}_0+s\mathbf{\Omega}_m)=\Sigma_{t,i},\quad \mathbf{r}\in D_i \tag{3}$$

In this small region if we know the incoming angular flux along the line k which starts at the boundary and which can be written as $\varphi^{in}_{i,k}(\mathbf{\Omega}_m)$, then outgoing angular flux from D_i along the line can be calculated as:

$$\varphi^{out}_{i,k}(\mathbf{\Omega}_m)=\varphi^{in}_{i,k}(\mathbf{\Omega}_m)\exp(-\Sigma_{t,i}s_{i,k})+\frac{Q_{i,k}(\mathbf{\Omega}_m)}{\Sigma_{t,i}}\left[1-\exp(-\Sigma_{t,i}s_{i,k})\right], \tag{4}$$

where $s_{i,k}$ is the length between the outgoing point and the incoming point of the line k in D_i.

The average segment angular flux can then be given as:

$$\overline{\varphi}_{i,k}(\mathbf{\Omega}_m)\cdot s_{i,k}=\frac{Q_{i,k}(\mathbf{\Omega}_m)}{\Sigma_{t,i}}s_{i,k}+\frac{\varphi^{in}_{i,k}(\mathbf{\Omega}_m)-\varphi^{out}_{i,k}(\mathbf{\Omega}_m)}{\Sigma_{t,i}} \tag{5}$$

2.2. Method of characteristics direction probabilities

The three-dimensional characteristics direction probabilities (CDP) include the directional transmission and collision probabilities which are stored for all the unique geometries of the problem. The transmission and collision probabilities are derived by integrating the MOC equation from the incoming boundary to the outgoing boundary. For a given geometry sub-domain (see Fig. 1), the outgoing angular flux of the sub-boundary can be written in terms of the probabilities as:

$$\varphi^{out}_{bo}(\mathbf{\Omega}_m)=\sum_{bi\in N(bo)}T_{bi->bo}(\mathbf{\Omega}_m)\varphi^{in}_{bi}(\mathbf{\Omega}_m)+\sum_{j\in J(bo)}T_{j->out}(\mathbf{\Omega}_m)Q_j(\mathbf{\Omega}_m) \tag{6}$$

where

$$T_{bi->bo}(\mathbf{\Omega}_m)=\sum_{k\in(bo\cap bi)}\frac{A_k\exp\left(-\sum_{j=1}^{i}\left(\Sigma_{t,j}s_{j,k}\right)\right)}{A_{bo}} \tag{7}$$

$$T_{j->bo}(\mathbf{\Omega}_m)=\sum_{k\in(bo\cap j)}\frac{A_k\left[1-\exp\left(-\Sigma_{t,j}s_{j,k}\right)\right]\exp\left(-\sum_{l=j+1}^{i}\Sigma_{t,l}s_{l,k}\right)}{A_{bo}\Sigma_{t,j}} \tag{8}$$

where $\varphi^{in}_{bi}(\mathbf{\Omega}_m)$ and $\varphi^{out}_{bo}(\mathbf{\Omega}_m)$ represent the incoming angular flux and outgoing angular flux of the sub-boundary, respectively, and where k is the characteristic line index and i is the flat source region index along the characteristic line shown in Fig. 1, and where $j=1$ is the first region traversed by the characteristic line k, and where A_k is the cross section area of the characteristic track orthogonal to the characteristic track direction, A_{bo} is the projection area of the sub-boundary, and $k\in(bo\cap bi)$ means the characteristic lines traverse both the bi and bo sub-boundaries.

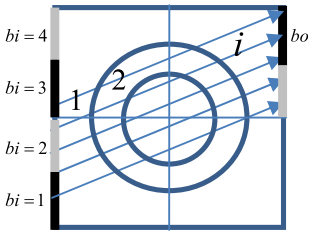

Fig. 1. The modular geometry sub-domain.

In addition to the outgoing angular flux, fine mesh region source of the sub-domain:

$$\overline{\varphi}_i(\boldsymbol{\Omega}_m) = \sum_{bi\in N(i)} (P_{bi->i}(\boldsymbol{\Omega}_m)\varphi_{bi}(\boldsymbol{\Omega}_m)) + \sum_{j\in up(i)} P_{j->i}(\boldsymbol{\Omega}_m)Q_j(\boldsymbol{\Omega}_m), \quad (9)$$

where,

$$P_{bi->i}(\boldsymbol{\Omega}_m) = \sum_{k\in(bi\cap i)} \frac{A_{i,k}}{\Sigma_{t,i}V_i}\exp\left(-\sum_{j=1}^{i-1}\left(\Sigma_{t,j}s_{j,k}\right)\right) \times \left[1-\exp\left(-\Sigma_{t,i}s_{i,k}\right)\right], \quad (10)$$

and,

$$P_{j->i}(\boldsymbol{\Omega}_m) = \begin{cases} \sum_{k\in(i)} \dfrac{(s_{i,k}\Sigma_{t,i} - [1-\exp(-\Sigma_{t,i}s_{i,k})])A_{i,k}}{V\Sigma_{t,i}\Sigma_{t,i}} & j=i \\ \sum_{k\in(i\cap j)} \dfrac{\left[1-\exp\left(-\Sigma_{t,j}s_{j,k}\right)\right]\exp\left(-\sum_{l=j+1}^{i-1}\Sigma_{t,l}s_{l,k}\right)[1-\exp(-\Sigma_{t,i}s_{i,k})]A_{i,k}}{V_i\Sigma_{t,j}\Sigma_{t,i}} & j\neq i \end{cases}, \quad (11)$$

In Eq. (9), $j \in up(i)$ refers to the regions in the up-streaming direction of region i, and in Eq. (11)$k\in(i\cap j)$ refers to the characteristic lines which traverse through both region i and j. Therefore the region average angular flux couples all the incoming angular fluxes which traverse through this region and the fine mesh regions in the up-streaming direction.

2.3. The angular depended boundary average

Similar to the MOC modular ray tracing technique, the ray tracing information is stored for all unique sub-domains as shown in Fig. 2. For purposes of describing the boundary averaging process, it is useful to introduce the coarse ray spacing of the radial direction (see Fig. 3.d_{Cr}) and the coarse ray spacing of z-axial direction (see Fig. 5.d_{Cz}).

The average is performed for all faces, respectively, and it is different for TOP/BOTTOM faces than the other faces. As shown in Fig. 4 for the TOP/BOTTOM faces, the boundary indices coming from the same flat source region between coarse rays are averaged together. So only the coarse ray spacing of the radial direction is useful when the average is determined for the TOP/BOTTOM faces. This averaging method has the benefit that it can determine at the same time all the outgoing fluxes which are outgoing from the same flat source region but in different angles and space, as shown. The average of the other faces is performed on every sub-boundary which is determined by the coarse rays shown in Fig. 5, which also shows that the final sub-boundaries are different for different angles.

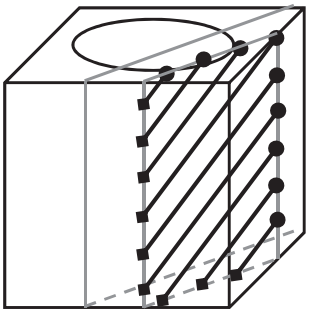

Fig. 2. The characteristic rays of a sub-domain.

3. The performance model

The method of characteristics direction probabilities is essentially an integral form MOC in which the computational effort of sweeping is reduced by storing preprocessed information. The reduction in the computational effort can be quantified by comparing the performance of the CDP to MOC using a performance model which accounts for the storage in memory, the computing operations, and the overall calculation time.

3.1. Memory storage requirements

The non-zeros of the CDP probability matrix are stored for every angle of every unique sub-domain of every energy group. The probabilities include the incoming-outgoing transmission probabilities, region-outgoing transmission probabilities, incoming-region collision probabilities, and region−region collision probabilities. This can be expressed as:

$$N_{nz,id}^{iang} = N_{in->out,id}^{iang} + N_{reg->out,id}^{iang} + N_{in->reg,id}^{iang} + N_{reg->reg,id}^{iang} \quad (12)$$

where $N_{nz,id}^{iang}$ is the number of none-zeros for the $iang$th angle of idth sub-domain.

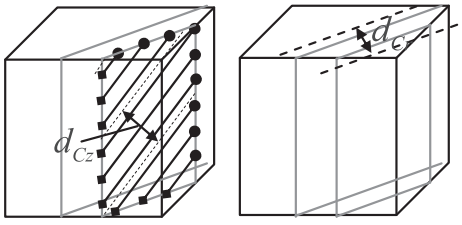

Fig. 3. The coarse ray spacing.

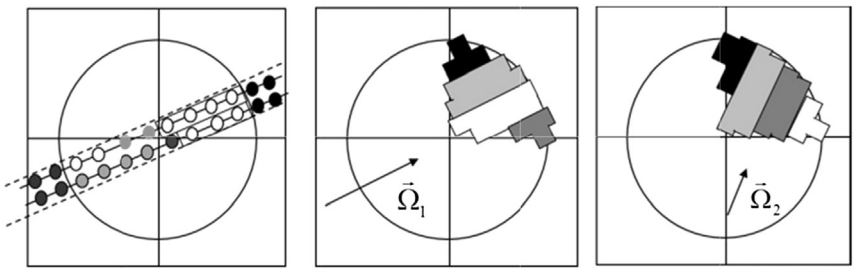

Fig. 4. Boundary index average of the TOP/BOTTOM face.

If no averaging is performed, the expression of each individual term for both with/without average cases can expressed as follows:

$$N^{iang}_{in->out,id} = N^{iang}_{ray,id}, \quad N^{iang}_{reg->out,id} = N^{iang}_{seg,id}, \quad N^{iang}_{in->reg,id} = N^{iang}_{seg,id}, \quad N^{iang}_{reg->reg,id} = N_{c,id}N_{reg,id} \tag{13}$$

where, $N^{iang}_{ray,id}$ is the number of characteristic rays for the *iang*th angle of *id*th sub-domain, $N^{iang}_{seg,id}$ is the number of segments for the *iang*th angle of *id*th sub-domain, and $N_{reg,id}$ is the number of fine regions.

Every outgoing angular flux has a corresponding incoming angular flux along its characteristics ray, so the number of incoming-outgoing transmission probabilities is equal to the number of characteristics rays. Along a single ray, every flat source region traversed will contribute to the outgoing angular flux, and therefore the total number of region-outgoing transmission probabilities will be the same as the number of segments. According to the reciprocity relations, the number of incoming-region collision probabilities will be equal to the number of segments as well. It is more complicated to determine the number of region–region probabilities since in the CDP every region just couples the up-streaming regions which vary with the number of angles and spacing of the rays. However, for a given sub-domain the region–region probabilities of different azimuthal angles will not change much, and therefore can be assumed as a constant times the number of flat source region, where the constant is set as N_c and can be determined by measurement for a particular angle with a relatively coarse ray spacing. The number of none-zero probabilities can then be expressed as:

$$N^{iang}_{nz,id} = N^{iang}_{ray,id} + 2N^{iang}_{seg,id} + N_{c,id}N_{reg,id}. \tag{14}$$

For the average case, a multi-to-one relation of the incoming-outgoing probabilities can be used where one outgoing angular flux no longer has one corresponding incoming angular flux but is coupled several incoming angular fluxes. A similar relation can be used for the region-out transmission probabilities and incoming-region collision probabilities. It is interesting to note that the number of region–region does not change since the angles and ray spacing doesn't change. Therefore, the number of none-zeros for the average case can be written as:

$$N^{iang}_{in-out,id} = \sum_{ibo=1}^{BO^{iang}_{id}} Nbi^{iang}_{ibo,id}, \quad N^{iang}_{reg->out,id} = \sum_{ibo=1}^{BO^{iang}_{id}} Nr^{iang}_{ibo,id},$$
$$N^{iang}_{in->reg,id} \approx N^{iang}_{reg->out,id}, \quad N^{iang}_{reg->reg,id} = N_{c,id}N_{reg,id}, \tag{15}$$
$$N^{iang}_{nz,id} = \sum_{ibo=1}^{BO^{iang}_{id}} Nbi^{iang}_{ibo,id} + 2\sum_{ibo=1}^{BO^{iang}_{id}} Nr^{iang}_{ibo,id} + N_{c,id}N_{reg,id}$$

where BO^{iang}_{id} is the number of outgoing sub-boundaries of *id*th sub-domain and $Nbi^{iang}_{ibo,id}$ is the number of the incoming sub-boundaries of the *ibo*th outgoing sub-boundary of the *id*th sub-domain.

The number of none-zeros is reduced when averaging is performed on the boundary, especially for 3-D cases. Eq. (14) and Eq. (15) provide the number of none-zeros for a particular angle of a sub-domain for one energy group. The total number of none-zeros can then be written:

$$N^{total}_{nz} = N_g \sum_{id=1}^{N_d} \sum_{iang=1}^{N_{ang}} N^{iang}_{nz,id} \tag{16}$$

3.2. The operation comparison of the MOC and CDP

The operation count for MOC and CDP requires explicit comparison the floating point operations (FLOPs) required for ray tracing. However, in addition to the sweep kernel, additional operations are required in CDP for generating the probabilities. Because the FLOPs are primarily dependent on the details of the coding, the pseudo code of the sweep kernel is given in Fig. 6.

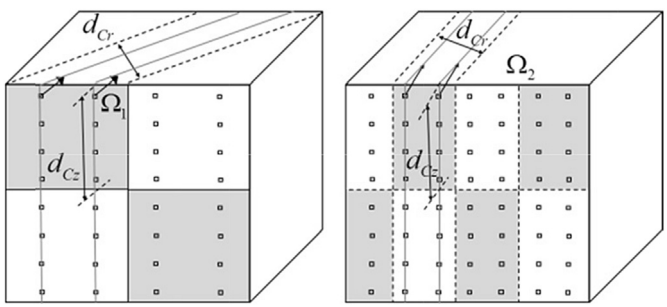

Fig. 5. Boundary index average of the other faces.

```
LOOP angles
  LOOP over all domains
    -Fetch the incoming BC of this domain
    !For the CDP
    - Matrix-vector multpl is performed   !CDP line
    !For the MOC
    - LOOP over rays in this domain       !MOC line

    -Pass the outgoing BC to the next domain
  END LOOP over all domains
  -Update the boundaries
END LOOP over angles
```

Fig. 6. The sequence of the sweep kernel.

```
DO iray=1,nrays(imod,iang)
  !Get the segment information in the ray
  DO imseg=1,nseglray(iray,imod,iang)
    ireg_seg(imseg)=ifirstreg(imod)
    ireg=ifirstreg(imod)+ireg(imseg,iray,imod,iang)
    tau_seg(imseg)=-xstr(ireg)*hseg(imseg,iray,imod,iang)
  ENDDO
  !Evaluate Exp()
  DO imseg=1,nseglray(iray,imod,iang)
    ix=FLOOR(table_rx*tau_seg(iseg))
    exparg(imseg)=expTable(1,ix)+exptable(2,ix)*tau_seg(iseg)
  ENDDO
  !Ray tracing
  DO imseg=1,nseglray(iray,imod,iang)
    ireg=ireg_seg(imseg)
    phid=(phic-qbar(ireg))*exparg(imseg)
    phio=phic-phid
    phiLocal(ireg)=phiLocal(ireg)+phid*dlr(im,iang)
  ENDDO
ENDDO
```

Fig. 7. The sweep kernel of MOC.

Fig. 6 shows the sweep sequence for both CDP and MOC, which also indicates that the only difference between them is the process of ray tracing; the CDP is a matrix-vector multiplication while the MOC traces ray by ray. Therefore only the operations of that part will be counted and compared in the sweep.

The operation count of the MOC kernel is determined from the Fig. 7 which shows the detailed ray tracing kernel of the MOC with rectangular boxes marking all operations in each individual part of the kernel. It should be noted that the plus operation marked by the circle will not be counted because it is just taken as an index of array xstr(:). Therefore the total operation count of the MOC sweep kernel can be expressed as:

$$F_{moc} = 9N_{swp}N_gN_{seg}$$

The CDP sweep kernel is shown in Fig. 8. Since the dimension of the matrix "n" is equal to $N_{BO} + N_{reg}$, the total number of operations can be expressed as:

$$F = N_{swp}N_g\left(2N_{nz}^{ig} + N_{BO} + N_{reg}\right) \tag{17}$$

```
!For the CDP
!Matrix-vector multply
INTEGER :: n,nnz,i,j,k
INTEGER :: ia(n+1),ja(nnz)
REAL :: aa(nnz),x(n),y(n),temp
DO i=1,n
  tmp=0.0 SDK
  DO k=ia(i),ia(i+1)- 1
    j=ja(k)
    tmp=tmp+aa(k)*x(j)
  ENDDO
  y(i)=y(i)+tmp
ENDDO
```

Fig. 8. The kernel of CDP.

Instead of the verbose code for generating the probabilities, a brief description of the processing of the probabilities is shown in Fig. 9. The total operation count can be derived by the operations required for every step.

$$\begin{aligned}
F_{GP,id}^{iang,ig} &= \left[6N_{seg,id}^{iang}\right] + \left[\sum_{iray=1}^{N_{ray,id}^{iang}} Ns_{iray,id}^{iang}\left(Ns_{iray,id}^{iang}-1\right)/2\right] \\
&\quad + \left[N_{ray,id}^{iang}\right] + \left[4N_{seg,id}^{iang}\right] \\
&\quad + \left[\begin{array}{l}\sum_{iray=1}^{N_{ray,id}^{iang}} Ns_{iray,id}^{iang}\left(Ns_{iray,id}^{iang}+1\right)/2 + 2N_{seg,id}^{iang} \\ +3\left(\sum_{iray=1}^{N_{ray,id}^{iang}} Ns_{iray,id}^{iang}\left(Ns_{iray,id}^{iang}+1\right)/2 - N_{seg,id}^{iang}\right)\end{array}\right] \\
&\quad + \left[N_{in->out,id}^{iang} + 2N_{reg->out,id}^{iang} + N_{in->reg,id}^{iang}\right. \\
&\quad \left. + 2N_{reg->reg,id}^{iang}\right] + \left[N_{in->out,id}^{inv,iang} + 2N_{reg->out,id}^{inv,iang}\right. \\
&\quad \left. + N_{in->reg,id}^{inv,iang} + 2N_{reg->reg,id}^{inv,iang}\right] \\
&= 2.5\sum_{iray=1}^{N_{ray}}\left(Ns_{iray,id}^{iang}\right)^2 + 10.5N_{seg,id}^{iang} + N_{ray,id}^{iang} \\
&\quad + 2N_{in->out,id}^{iang} + 6N_{reg->out,id}^{iang} + 4N_{reg->reg,id}^{iang}
\end{aligned} \tag{18}$$

The total number of operations is summarized in Eq. (18) where each term in the square brackets is the number of FLOPs for every step. It should be noted that for a particular angle, the probability matrix of both forward and backward direction are set together which minimizes the FLOPs required. By summing over all angles and multiplying by the number of energy groups, the total FLOPs are:

$$F_{id} = N_g\sum_{iang} F_{GP,id}^{iang,ig} \tag{19}$$

```
LOOP over angles
  LOOP over energy groups
    LOOP over rays
      -Evaluate Exp() of every segment
      -Evaluate Exp() between segments
      -Set T_bi->bo in density matrix
      -Set T_req->bo in density matrix
      -Set P_bi->req in density matrix
      -Set P_req->req in density matrix
    END LOOP over rays
    -Set and store all None-zeros into the
      sparse matrix for the forward direction
    -Set and store all None-zeros into the
      sparse matrix for the backward direction
  END LOOP over energy groups
END LOOP over angles
```

Fig. 9. The process of generating the probabilities.

3.3. The computing time and performance of the algorithm

The standard metric for evaluating the performance of an algorithm can be expressed as:

$$P = \frac{F}{T}, \tag{20}$$

where P means the performance which stands for the FLOPs performed per second of the processor by an algorithm; F is the total number of floating point operations; T is the total computing time. The floating point operations can be determined by measurement or using the methods introduced in the previous section, and the total computing can be evaluated by considering both the memory access time and the time to perform operations. An expression accounting for both the operation time and the memory access time can be written as:

$$T = Ft_F + Mt_{M,avg} \tag{21}$$

$$t_{M,avg} = (t_{hit} + r_{miss}t_{miss}) \tag{22}$$

where F is the number of floating point operations; M is the number of accesses to the memory; t_F is the time to perform a FLOP; $t_{M,avg}$ is the average time to perform a memory access; t_{hit} is the access time to the high-level memory; r_{miss} is the ratio of missing to access the data in the high level memory; t_{miss} is the penalty of missing to access the data in the high level memory.

Substitute the time model into Eq. (20):

$$P = \frac{1}{t_F + \frac{M}{F}(t_{hit} + r_{miss}t_{miss})},$$

The computational intensity can be defined as the ratio $q = F/M$ which is the number of FLOPs per memory access. The performance can then be written as:

$$P = \frac{1}{t_F + \frac{1}{q}(t_{hit} + r_{miss}t_{miss})} \tag{23}$$

Eq. (23) provides a simple expression to relate the performance of an algorithm to its computational intensity and memory access characteristics.

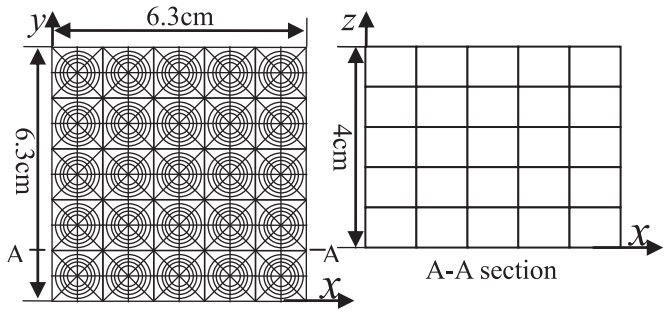

Fig. 10. Test problem configuration.

Table 1
The Eigenvalue of the problem.

	MOC	CDP				
		No_avg	Case 1	Case 2	Case 3	Case 4
k-eff	1.32624	1.32624	1.32623	1.32620	1.32616	1.32576

4. Numerical results

The accuracy and performance of the boundary-averaged CDP method proposed here were evaluated using two test problems. The first is a 5 × 5 − pin problem and the second is a standard reactor benchmark published by Takeda and Ikeda (1991).

4.1. The 5 × 5-pin test problem

This test problem is a mult-pin case with 5 by 5 pins in the X–Y plane and 4 levels in the Z-axial direction as shown in Fig. 10. A standard set of UO_2/Water cross sections were used from the C5G7 benchmark. The analysis was performed based on the following discretization: 40 flat source regions in the pin, 1 angle in one octant and 0.03 cm ray spacing which introduces 31,960 rays and 176,832 segments in one pin for all angles. In order to the explicitly show the performance of the boundary averaging, 4 different average cases were performed:

1. Case 1: coarse ray spacing for TOP/BOTTOM face is 0.03 cm; coarse ray spacing in radial direction is 0.03 cm and coarse ray spacing in Z-axial direction is 1.0 cm for other faces.
2. Case 2: coarse ray spacing for TOP/BOTTOM face is 0.06 cm; coarse ray spacing in radial direction is 0.06 cm and coarse ray spacing in Z-axial direction is 1.0 cm for other faces.
3. Case 3: coarse ray spacing for TOP/BOTTOM face is 0.1 cm; coarse ray spacing in radial direction is 0.1 cm and coarse ray spacing in Z-axial direction is 1.0 cm for other faces.
4. Case 4: coarse ray spacing for TOP/BOTTOM face is 2.0 cm; coarse ray spacing in radial direction is 0.1 cm and coarse ray spacing in Z-axial direction is 1.0 cm for other faces.

The results for all cases are shown in Table 1 which also provides the MOC result. As expected, if no averaging is used in CDP, then the MOC and CDP results are identical. Also as expected, the results change as the boundary averaging changes, but even in the most aggressive averaging case the difference is only about 50 pcm. The results suggest that the coarse ray spacing in case 3 with 40 sub-boundaries on the TOP/BOTTOM face are sufficient to provide an accurate result.

The performance of the algorithm was then evaluated and Table 2 shows the number of non-zero elements for each case. Several conclusions are apparent from this table. First, the boundary averaging can significantly reduce the required storage. As shown in the Table case 4 requires only 3.23% of the storage of that without averaging. The second conclusion is that there is

Table 2
Non-zero elements of the probability matrix.

	NBCs	Nin-out	Nreg-out	Nin-reg	Nreg-reg	Nnz_1G	
No_average	31,960	31,960	176,832	176,832	2400	388,024	100.00%
Case 1	3760	7040	21,240	21,240	2400	51,920	13.38%
Case 2	2016	3792	11,872	11,872	2400	29,936	7.71%
Case 3	1424	2688	8648	8648	2400	22,384	5.77%
Case 4	488	2672	3744	3704	2400	12,520	3.23%

Table 3
Floating point operations summary.

	F_GP (estimated)	F_GP (measured)	Ratio/%	F_Swp (estimated)	F_Swp (measured)	Ratio/%	F_Swp (/Domisns/Nswp)
MOC	–	–		1.1140E+11	1.1140E+11	100.00	1.1140E+07
No_Avg	2.2983E+07	2.2983E+07	100.00	5.6605E+10	5.6610E+10	99.99	5.6605E+06
Case 1	1.9541E+07	1.9540E+07	100.01	7.5768E+09	7.5770E+09	100.00	7.5768E+05
Case 2	1.9323E+07	1.9320E+07	100.01	4.3770E+09	4.3770E+09	100.00	4.3770E+05
Case 3	1.9247E+07	1.9250E+07	99.99	3.2782E+09	3.2780E+09	100.01	3.2782E+05
Case 4	1.9143E+07	1.9140E+07	100.02	1.8318E+09	1.8320E+09	99.99	1.8318E+05

Table 4
The sweep performance of the MOC and CDP.

	FLOPs	t/s	FLOPs speedup	Time speedup	No. of DCA	q	r/%	P/MFLOPs
MOC	1.1140E+11	368.635	–	–	5.1999E+11	0.214	1.21	302.21
No_average	5.6610E+10	78.032	2.0	4.7	1.1156E+11	0.507	1.91	725.41
Case 1	7.5770E+09	11.276	14.7	32.7	1.5054E+10	0.503	1.69	671.94
Case 2	4.3770E+09	6.861	25.5	53.7	8.8781E+09	0.493	1.56	637.95
Case 3	3.2780E+09	5.29	34.0	69.7	6.7116E+09	0.488	1.76	619.71
Case 4	1.8320E+09	2.914	60.8	126.5	3.6710E+09	0.499	2.18	628.61

significant reduction in the incoming-outgoing and region-−region collision probabilities as the boundary averaging is coarsened, however, the number of region−region probabilities is the same.

The number of floating point operation is summarized in Table 3, where F_GP is the FLOPs of the probability generation and F_Swp is the FLOPs performed by the ray tracing process. The last column is the FLOPs needed by one sub-domain per sweep. Both the estimated and measured FLOPs are listed here where the estimated FLOPs are calculated by the equations in the previous section and the measured FLOPs are directly obtained by measurement using the TAU (Tuning and Analysis Utilities) and PAPI (Performance Application Programming Interface). The results in this Table provide several insights about the performance of the algorithm.

1) There is very good agreement in the measured and predicted performance which provides confidence in the models introduced in Section 3.2;
2) The FLOPs required by the CDP without averaging is less than 1/2 of that of the MOC;
3) The coarsening of the averaging can significantly reduce the FLOPs performed by the CDP, primarily because the number of none-zeros and therefore the storage is significantly reduced;
4) The FLOPs performed by the generation of the probabilities are almost twice that of the MOC sweep and it does not change much for different averaging cases, primarily because the FLOPs is dominated by the number of segments which is the same for different averaging cases.

The overall performance of the MOC and CDP algorithm is summarized in Table 4, where DCA indicated the number of data cache accesses; q is the computational intensity which is determined by dividing FLOPs with DCA; and r is the L1 cache miss rate. As indicated in the Table, it is clear that the overall speedup is larger than the FLOPs speedup, which is about 2 times for every case. This is primarily because the performance of the processor for the CDP is almost twice better than that for the MOC as shown in the last column. As indicated in Eq. (23), the performance is dependent on the computational intensity and the high-level cache miss ratio. In this test problem, the high-level cache miss ratio is almost the same which is about 1%−2%, however the computational intensity of the MOC is about half that of the CDP. Therefore we can conclude that the CDP not only reduces the FLOPs performed but also improves the overall performance of the processors.

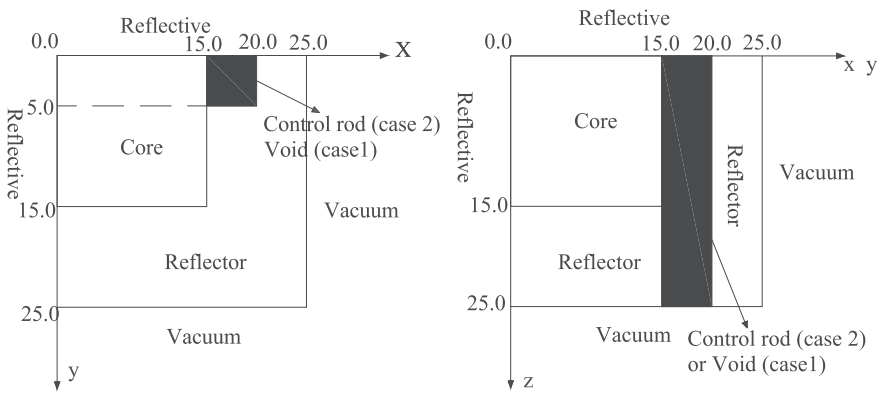

Fig. 11. KUCA core benchmark core configuration.

Table 5
Comparison of *k*-eff for KUCA Benchmark.

Method	UnRodded	Rodded	CR-worth
Ref Monte-Carlo	0.9780 ±0.0006	0.9624 ±0.0006	1.66E-02 ±0.09E-2
MOC	0.9776	0.9627	1.59E-02
CDP_case1	0.9776	0.9627	1.59E-02
CDP_case2	0.9773	0.9624	1.59E-02
CDP_case3	0.9769	0.9619	1.59E-02

4.2. KUCA Takeda core benchmark

The Kyoto University Critical Assembly (KUCA) benchmark was published by (Takeda and Ikeda, 1991). This problem models a small Light Water reactor core of dimensions 50 cm × 50 cm × 50 cm, with three material regions including fuel, control rod, and reflector. Fig. 11 shows the core configuration and the 2 group cross section are provided in Appendix A. The control rod worth was evaluated for this problem by inserting and removing the control rod.

The two cases were performed with identical discretizations in which the sub-domain size was 1 cm × 1 cm × 1 cm with 32 fine-regions in each sub-domain. The ray spacing was 0.03 cm with the coarse ray spacing of 0.1 cm in the radial plane and 0.3 cm in the axial direction. An S4 angular quadrature sets was used in both cases.

Four calculating results are given including the MOC result and CDP results with three different average cases:

1) CDP_Case1: CDP calculation without average,
2) CDP_Case2: CDP averaged with the coarse ray spacing of 0.03 cm in the radial plane and 0.15 cm in the axial direction,
3) CDP_Case3: CDP averaged with the coarse ray spacing of 0.1 cm in the radial plane and 0.3 cm in the axial direction.

The reference results of Monte-Carlo method were provided in (Takeda and Ikeda, 1991) and the *k*-eff and control rod worth $(k - k')/(kk')$ are compared to these results in Table 5. As expected, the CDP without average is consistent to the MOC, and the average introduces some discrepancy in *k*-eff. Because this problem has a significant axial flux distribution compared to the previous test problems, the use of a smaller number of flat source regions on the TOP/BOTTOM interfaces introduces a larger discrepancy in *k*-eff. The overall conclusion from these results is that CDP with boundary averaging can provide satisfactory results for this problem with the above discretization.

The computational performance of the MOC and CDP for the KUCA problem is compared in Table 6. As indicated, CDP shows a reduction in the execution time of from 3 to 30. The execution time is reduced because of the smaller probability matrix and higher processor performance. It also gives the average memory allocation of every processor when calculating this problem, in which 25 processors are used in special decomposition for every case. Because the boundary angular fluxes needs to be stored for communication, and the amount of boundary conditions dominates the memory storage for there-dimensional MOC calculations, it sometimes makes the calculation impossible because of the insufficient memory for the boundary angular fluxes. The CDP needs to store the probability matrix, so it increases the memory allocation when the average is not performed. From the table, it can be concluded that the average significantly speedups the computing time and remarkably reduces the memory requirement of the boundary angular fluxes as well.

5. Summary and conclusions

The previous research investigated the method of characteristic direction probability (CDP), which combines the geometry flexibility of the MOC and computing efficiency of the CPM, as a means to reduce the computational efforts for 3D problems. In the CDP method, the probabilities only couple the fine mesh regions traversed by the characteristic lines in a particular direction within a sub-domain. However the probability matrix can still be large in CDP when the ray spacing is small, so the additional feature was introduced of boundary averaging which considerably reduced the number of outgoing/incoming angular fluxes transferred between domains.

From the numerical results, it can be concluded that the CDP without boundary averaging is exactly identical to the MOC and the accuracy impact of the boundary averaging is depended on the coarse ray spacing. The accuracy gets worse while the coarse ray spacing increases. So there is a balance between the accuracy losing and the computing time benefit.

A performance model was introduced in order to explicitly analyze and compare the performance of MOC with the CDP on the storage, floating point operations and computing time for various boundary averaging cases. It shows the benefits from the CDP in detail and the numerical results of two test problems indicate that the boundary averaging method can significantly reduce the storage and computing time, and yet maintain sufficient accuracy.

Acknowledgements

This research was carried out under the financial support of the State Scholarship Fund of China (No. 2011628095). Thank the NERS of the University of Michigan for providing the active research environments.

Appendix A. Cross Sections

Table 6
Computing time comparison.

Case	Method	Iterations	Memory per-processor	Total time/s	Speedups
Void	MOC	918	188.87 MB	3357.8	
	CDP_case1	918	210.48 MB	1041.04	3.2
	CDP_case2	918	51.09 MB	273.69	12.3
	CDP_case3	918	18.33 MB	119.28	28.2
Rodded	MOC	927	188.87 MB	3494.5	
	CDP_case1	927	210.48 MB	916.01	3.8
	CDP_case2	918	51.09 MB	272.08	12.8
	CDP_case3	918	18.33 MB	116.7	29.9

Table A.1
KUCA core benchmark two-group cross sections and energy group structure for KUCA

Region	Group	Σ_{abs}/cm - 1	$\nu\Sigma_f$/cm - 1	Σ_{1-1}/cm - 1	Σ_{1-2}/cm - 1	χ
Core fuel	1	8.52709E-03	9.09319e-3	1.92423e-1	2.28253e-2	1.0
	2	1.58196E-01	2.90183e-1	0.00000E+0	8.80439e-1	
Control rod	1	1.74439E-02	0.00000E+0	6.77241e-2	6.45461e-5	–
	2	1.82224E-01	0.00000E+0	0.00000E+0	3.52358e-2	
Reflector	1	4.16392E-04	0.00000E+0	1.93446e-1	5.65042e-2	–
	2	2.02999E-02	0.00000E+0	0.00000E+0	1.62452E+0	
Empty (void)	1	4.65132E-05	0.00000E+0	1.27700e-2	2.40997e-5	–
	2	1.32890E-03	0.00000E+0	0.00000E+0	1.07387e-2	

Table A.2
Fuel-clad macroscopic cross-sections for C5G7 benchmark

	Transport cross-section/cm^{-1}	Absorption cross-section/cm^{-1}	Capture cross-section/cm^{-1}	Fission cross-section/cm^{-1}	Nu	Chi
Group 1	1.77949E-01	8.02480E-03	8.12740E-04	7.21206E-03	2.78145E+00	5.87910E-01
Group 2	3.29805E-01	3.71740E-03	2.89810E-03	8.19301E-04	2.47443E+00	4.11760E-01
Group 3	4.80388E-01	2.67690E-02	2.03158E-02	6.45320E-03	2.43383E+00	3.39060E-04
Group 4	5.54367E-01	9.62360E-02	7.76712E-02	1.85648E-02	2.43380E+00	1.17610E-07
Group 5	3.11801E-01	3.00200E-02	1.22116E-02	1.78084E-02	2.43380E+00	0.00000E+00
Group 6	3.95168E-01	1.11260E-01	2.82252E-02	8.30348E-02	2.43380E+00	0.00000E+00
Group 7	5.64406E-01	2.82780E-01	6.67760E-02	2.16004E-01	2.43380E+00	0.00000E+00

	To group 1/cm^{-1}	To group 2/cm^{-1}	To group 3/cm^{-1}	To group 4/cm^{-1}	To group 5/cm^{-1}	To group 6/cm^{-1}	To group 7/cm^{-1}
Group 1	1.27537E-01	4.23780E-02	9.43740E-06	5.51630E-09	0.00000E+00	0.00000E+00	0.00000E+00
Group 2	0.00000E+00	3.24456E-01	1.63140E-03	3.14270E-09	0.00000E+00	0.00000E+00	0.00000E+00
Group 3	0.00000E+00	0.00000E+00	4.50940E-01	2.67920E-03	0.00000E+00	0.00000E+00	0.00000E+00
Group 4	0.00000E+00	0.00000E+00	0.00000E+00	4.52565E-01	5.56640E-03	0.00000E+00	0.00000E+00
Group 5	0.00000E+00	0.00000E+00	0.00000E+00	1.25250E-04	2.71401E-01	1.02550E-02	1.00210E-08
Group 6	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	1.29680E-03	2.65802E-01	1.68090E-02
Group 7	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	8.54580E-03	2.73080E-01

Table A.3
Moderator macroscopic cross-sections for C5G7 benchmark

	Transport cross-section/cm^{-1}	Absorption cross-section/cm^{-1}	Capture cross-section/cm^{-1}
Group 1	1.59206E-01	6.01050E-04	6.01050E-04
Group 2	4.12970E-01	1.57930E-05	1.57930E-05
Group 3	5.90310E-01	3.37160E-04	3.37160E-04
Group 4	5.84350E-01	1.94060E-03	1.94060E-03
Group 5	7.18000E-01	5.74160E-03	5.74160E-03
Group 6	1.25445E+00	1.50010E-02	1.50010E-02
Group 7	2.65038E+00	3.72390E-02	3.72390E-02

	To group 1/cm^{-1}	To group 2/cm^{-1}	To group 3/cm^{-1}	To group 4/cm^{-1}	To group 5/cm^{-1}	To group 6/cm^{-1}	To group 7/cm^{-1}
Group 1	4.44777E-02	1.13400E-01	7.23470E-04	3.74990E-06	5.31840E-08	0.00000E+00	0.00000E+00
Group 2	0.00000E+00	2.82334E-01	1.29940E-01	6.23400E-04	4.80020E-05	7.44860E-06	1.04550E-06
Group 3	0.00000E+00	0.00000E+00	3.45256E-01	2.24570E-01	1.69990E-02	2.64430E-03	5.03440E-04
Group 4	0.00000E+00	0.00000E+00	0.00000E+00	9.10284E-02	4.15510E-01	6.37320E-02	1.21390E-02
Group 5	0.00000E+00	0.00000E+00	0.00000E+00	7.14370E-05	1.39138E-01	5.11820E-01	6.12290E-02
Group 6	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	2.21570E-03	6.99913E-01	5.37320E-01
Group 7	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	0.00000E+00	1.32440E-01	2.48070E+00

References

Hong, S.G., Cho, N.Z., 1999. Method of characteristic direction probabilities for heterogeneous lattice calculation. Nucl. Sci. Eng. 132, 65–77.

Liu, Z., Wu, H., Cao, L., Chen, Q., Li, Y., 2011. A new three-dimensional method of characteristics for the neutron transport calculation. Annu. Nucl. Energy 38, 447–454.

Liu, Z., Kochunas, B., Collins, B., Downar, T., Wu, H., 2013. The method of modular characteristic direction probabilities in Mpact. In: International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013), Sun Valley, Idaho, USA, May 5-9, 2013, on CD-ROM. American Nuclear Society, LaGrange Park, IL.

Mohanakrishnan, P., 1981. Angular current approximations in neutron transport calculations using interface current—a review. Prog. Nucl. Energy 7, 1–10.

Sanchez, R., 1997. Approximate solutions of the two dimensional integral transport equation by collision probability methods. Nucl. Sci. Eng. 64, 384–404.

Takeda, T., Ikeda, H., 1991. 3-D neutron transport benchmarks. J. Nucl. Sci. Technol. 28 (7), 656–669.