J. Plasma Physics (2013), vol. 79, part 1, pp. 69–86. © Cambridge University Press 2012
doi:10.1017/S0022377812000724

Modeling and simulations of high-power microwave devices using the CHIPIC code

JUN ZHOU, D. G. LIU and C. LIAO

School of Physical Electronics, University of Electronic Science and Technology of China, Chengdu 610054, People's Republic of China
(zhoujun@uestc.org)

(Received 15 March 2012; revised 5 July 2012; accepted 9 July 2012; first published online 8 October 2012)

Abstract. The CHIPIC code, a fully electromagnetic particle-in-cell (PIC) code for modeling and simulations of high-power microwave (HPM) devices, is introduced in this paper. It consists of a two-dimensional (2D) code and a three-dimensional (3D) code. The 2D code can model and simulate HPM devices with symmetric structure on 2D Cartesian, cylindrical and polar grids, while the 3D code can model and simulate HPM devices on 3D Cartesian and cylindrical grids. The fields are calculated using the finite-difference time-domain scheme, and the particles are described by the PIC scheme. Various types of boundary conditions have also been implemented for different kinds of applications. In addition, the 3D code is specifically designed for high-performance modeling and computing. It uses the message passing interface and the open specifications for multiprocessing (OpenMP) for parallelization. Its parallel design ensures that it is capable of efficiently executing on a variety of architectures. In order to allow efficient use of parallel architectures, it provides automated partitioning and dynamic load balancing. Even though this code is still in development, it has successfully simulated various real-world HPM experimental devices. Simulation results on some typical HPM devices by using the CHIPIC code are given, which agree well with those obtained from some well-known PIC codes. Direction for future work is also presented.

1. Introduction

High-power microwave (HPM) devices have reached maturity over the years and are considered to be fairly well understood (Gaponov-Grekhov and Granatstein 1994; Barker and Schamiloglu 2001; Benford et al. 2007). These devices include relativistic klystrons (Lau et al. 1990), relativistic backward wave oscillators (BWOs) (Moreland et al. 1996), relativistic magnetrons (Chan et al. 1990), magnetically insulated line oscillators (MILOs; Lemke et al. 1997), magnetically insulated microwave cavities (MIMCs; Geng et al. 2003; Grisham 2009; Stratakis et al. 2011), virtual-cathode oscillators (vircators; Jiang et al. 1999), split-cavity oscillators (SCOs; Lemke et al. 1994), gyrodevices (Chu 2004), and others. Yet, there are important open questions that need to be answered in order to obtain higher levels of efficiency or to improve the performance of the systems. As most of the questions are related to strong nonlinear effects that dominate the behavior of the devices, analytical solutions are almost intractable. Numerical simulations play an important role in this aspect as they allow a thorough investigation of such strong nonlinear effects for different schemes and therefore facilitate the design of HPM devices.

Over the past several decades, there has been significant progress in detailed time-dependent numerical simulations, which are essential for electron-beam systems to describe the nonlinear interaction between the electrons and the electromagnetic field (Antonsen et al. 1999). Researchers have made heavy use of these simulations to predict the behavior of HPM devices. The most accurate one of such models is probably the particle-in-cell (PIC) method (Hockney and Eastwood 1981; Birdsall and Langdon 1985; Birdsall 1991; Verboncoeur 2005). Based on the PIC method, a variety of two-dimensional (2D) and three-dimensional (3D) PIC codes have been written, for instance, MAGIC (Goplen et al. 1995), KARAT (Tarakanov 2001), OOPIC (Verboncoeur et al. 1995), ICEPIC (Blahovec et al. 2000), LSP (Clark and Hughes 2005), VLPL (Pukhov 1999) and VORPAL (Nieter and Cary 2004). These codes can help researchers investigate the complex physical mechanisms involved in various HPM devices, and have been promoting the development of the HPM technology. However, among these powerful codes, there are some disadvantages along with their well-known advantages. For instance, some codes are only user-configurable but not computer-aided, some codes are only 2D or 3D but not both, and

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

some codes only consider Cartesian but not other more suitable coordinate systems.

In this work, we present a computer-aided highly efficient electromagnetic PIC (CHIPIC) code (Zhou et al. 2009) which is designed to be an effective tool for modeling and simulations of HPM devices. It consists of a 2D code and a 3D code. As is well known, most of the HPM devices have an axially symmetric structure, in which case the 2D PIC codes are usually more practical and efficient than the 3D PIC codes because of their much less time consumption. The 2D CHIPIC code is specially designed and developed for this kind of devices. It can model and simulate HPM devices with symmetric structure on 2D Cartesian, cylindrical and polar grids. Also, for the purpose of highly efficient modeling and simulations of the HPM devices without symmetric structure, the 3D parallelized CHIPIC code is in development. It can model and simulate HPM devices on 3D Cartesian and cylindrical grids. 3D CHIPIC uses the message passing interface (MPI) and the open specifications for multiprocessing (OpenMP) for parallelization. Its parallel design ensures that it is capable of efficiently executing on a variety of architectures. In order to allow efficient use of parallel architectures, it provides automated partitioning and dynamic load balancing. Based on the CHIPIC code, various real-world HPM experimental devices have been successfully simulated (Geng et al. 2010; Li and Liu 2010, 2011; Qin et al. 2011; Tang et al. 2012). The designers and researchers who want to largely shorten their design time can benefit a lot from using the CHIPIC code. The model properties such as dimensions (2D or 3D), coordinate systems (Cartesian, cylindrical or polar) and run modes (single or parallel) can be easily selected, and the model information such as structural geometries, grid settings, boundary conditions and material properties can be well specified using the computer-aided design (CAD) system.

This paper is organized as follows. The CHIPIC code is introduced in detail in Sec. 2, including the construction, the fundamental algorithms, the boundary conditions and the parallelization used in this code. The modeling and simulations of various HPM devices by using this code are presented in Sec. 3, and the simulation results are compared with those obtained from some well-known PIC codes. Moreover, conclusions along with the direction for future work are given in Sec. 4.

2. Introduction to the CHIPIC code

2.1. Construction

The CHIPIC code is a fully computer-aided code which combines modeling, calculations and analysis with an integrated environment. It can be run on the WINDOWS, LINUX or UNIX platform. This code is made up of a CAD system and a physical kernel, of which the CAD system is implemented to provide highly efficient modeling and analysis and the physical kernel is developed to provide fast and accurate calculations for many research and design needs. The CAD system is written in object-oriented C++ and uses the open graphics library (OpenGL) for graphical operation, and the physical kernel is written in FORTRAN95 under Intel Visual Fortran Compiler.

The physical kernel is specially designed to be built from a mesh generator, two field solvers, a particle handler, an object manager, and a data manager, of which the mesh generator provides a non-uniform finite-difference mesh with staggered Cartesian, cylindrical or polar Yee grids (Yee 1966), the field solvers advance electromagnetic fields using the standard explicit or time-biased (iterative) algorithm (Godfrey 1980) and apply field boundary conditions, the particle handler advances charged particles with fully relativistic kinematics and applies particle boundary conditions, and the managers store various model information and simulation results, respectively.

Unlike the 2D CHIPIC code, in the 3D CHIPIC code, the mesh with polar Yee grids is not provided, because the cylindrical and polar coordinate systems are identical in the 3D case except for the order of the coordinates. In addition, because the 3D codes are always much more time-consuming than the 2D codes, the MPI and the OpenMP are adopted for parallelization to overcome this main drawback.

2.2. Fundamental algorithms

The numerical simulations of HPM devices consist of solving Maxwell's equations for the electromagnetic fields and the equations of motion for the particles (electrons). For an accurate and self-consistent simulation, one must include the electrons as the source of the electromagnetic fields, and the fields as part of the forces driving the electrons. In the CHIPIC code, the fields are calculated using the finite-difference time-domain (FDTD) scheme (Yee 1966; Taflove and Hagness 2000), and the particles are described by the PIC scheme. The field components are assigned to a spatiotemporal grid, calculated each time step, and used to 'push' the particles. The sources are then prepared and the field is recalculated. Since most of the HPM devices have an axially symmetric structure, we only present some important formulae in the 3D cylindrical coordinate system here.

The electromagnetic fields are governed by Maxwell's equations, written in differential form as

$$\frac{\partial \mathbf{D}}{\partial t} = \nabla \times \mathbf{H} - \mathbf{J}, \tag{2.1}$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}, \tag{2.2}$$

where **E**, **D**, **B**, **H** and **J** are the electric field strength, the electric displacement, the magnetic flux density, the magnetic field strength, and the electric current density,

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

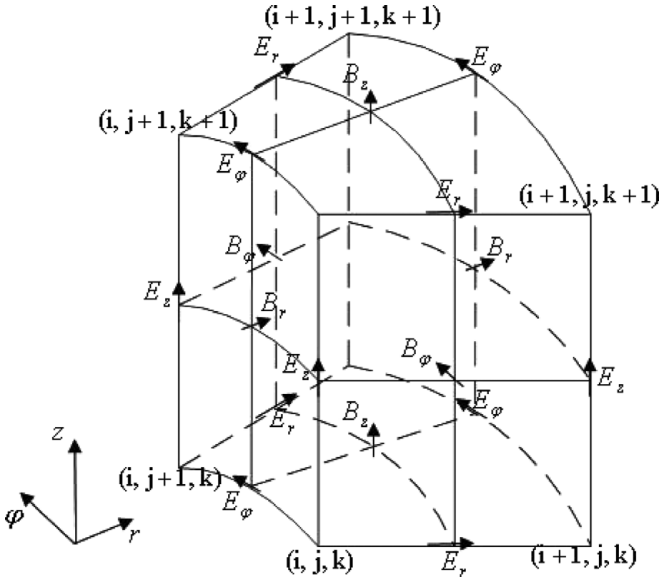


Figure 1. Fields defined on the 3D cylindrical Yee mesh.

respectively. They are all functions of position and time. The components of these field vectors are defined on the 3D Yee mesh, as shown in Fig. 1. That is, **E** and **D** components are defined at the midpoints of the segments connecting mesh nodes, while **B** and **H** components are defined similarly on a mesh displaced by half cell in each dimension. **J** components, which are not shown, are co-located with the corresponding electric field components. The electric field is known at integer time steps, while the magnetic field is known at integer plus a half time steps.

By applying the differential equation

$$\nabla \times \mathbf{A} = \left[\frac{1}{r}\frac{\partial}{\partial r}(rA_\varphi) - \frac{1}{r}\frac{\partial A_r}{\partial \varphi}\right]\mathbf{e}_z + \left[\frac{1}{r}\frac{\partial A_z}{\partial \varphi} - \frac{\partial A_\varphi}{\partial z}\right]\mathbf{e}_r + \left[\frac{\partial A_r}{\partial z} - \frac{\partial A_z}{\partial r}\right]\mathbf{e}_\varphi \tag{2.3}$$

for vector operation, (2.1) and (2.2) are translated to six scalar equations for the electromagnetic field components

$$\frac{\partial D_z}{\partial t} = \frac{1}{r}\frac{\partial}{\partial r}(rH_\varphi) - \frac{1}{r}\frac{\partial H_r}{\partial \varphi} - J_z, \tag{2.4}$$

$$\frac{\partial D_r}{\partial t} = \frac{1}{r}\frac{\partial H_z}{\partial \varphi} - \frac{\partial H_\varphi}{\partial z} - J_r, \tag{2.5}$$

$$\frac{\partial D_\varphi}{\partial t} = \frac{\partial H_r}{\partial z} - \frac{\partial H_z}{\partial r} - J_\varphi, \tag{2.6}$$

$$\frac{\partial B_z}{\partial t} = \frac{1}{r}\frac{\partial E_r}{\partial \varphi} - \frac{1}{r}\frac{\partial}{\partial r}(rE_\varphi), \tag{2.7}$$

$$\frac{\partial B_r}{\partial t} = \frac{\partial E_\varphi}{\partial z} - \frac{1}{r}\frac{\partial E_z}{\partial \varphi}, \tag{2.8}$$

$$\frac{\partial B_\varphi}{\partial t} = \frac{\partial E_z}{\partial r} - \frac{\partial E_r}{\partial z}, \tag{2.9}$$

which can be discretized in center-difference form in the 3D cylindrical (z,r,φ) coordinate system as

$$\frac{(D_z)^{n+1}_{i+1/2,j,k} - (D_z)^{n}_{i+1/2,j,k}}{\Delta t} = \frac{r_{j+1/2}(H_\varphi)^{n+1/2}_{i+1/2,j+1/2,k} - r_{j-1/2}(H_\varphi)^{n+1/2}_{i+1/2,j-1/2,k}}{r_j \Delta r} - \frac{(H_r)^{n+1/2}_{i+1/2,j,k+1/2} - (H_r)^{n+1/2}_{i+1/2,j,k-1/2}}{r_j \Delta \varphi} - (J_z)^{n+1/2}_{i+1/2,j,k}, \tag{2.10}$$

$$\frac{(D_r)^{n+1}_{i,j+1/2,k} - (D_r)^{n}_{i,j+1/2,k}}{\Delta t} = \frac{(H_z)^{n+1/2}_{i,j+1/2,k+1/2} - (H_z)^{n+1/2}_{i,j+1/2,k-1/2}}{r_{j+1/2}\Delta \varphi} - \frac{(H_\varphi)^{n+1/2}_{i+1/2,j+1/2,k} - (H_\varphi)^{n+1/2}_{i-1/2,j+1/2,k}}{\Delta z} - (J_r)^{n+1/2}_{i,j+1/2,k}, \tag{2.11}$$

$$\frac{(D_\varphi)^{n+1}_{i,j,k+1/2} - (D_\varphi)^{n}_{i,j,k+1/2}}{\Delta t} = \frac{(H_r)^{n+1/2}_{i+1/2,j,k+1/2} - (H_r)^{n+1/2}_{i-1/2,j,k+1/2}}{\Delta z} - \frac{(H_z)^{n+1/2}_{i,j+1/2,k+1/2} - (H_z)^{n+1/2}_{i,j-1/2,k+1/2}}{\Delta r} - (J_\varphi)^{n+1/2}_{i,j,k+1/2}, \tag{2.12}$$

$$\frac{(B_z)^{n+3/2}_{i,j+1/2,k+1/2} - (B_z)^{n+1/2}_{i,j+1/2,k+1/2}}{\Delta t} = \frac{(E_r)^{n+1}_{i,j+1/2,k+1} - (E_r)^{n+1}_{i,j+1/2,k}}{r_{j+1/2}\Delta \varphi} - \frac{r_{j+1}(E_\varphi)^{n+1}_{i,j+1,k+1/2} - r_j(E_\varphi)^{n+1}_{i,j,k+1/2}}{r_{j+1/2}\Delta r}, \tag{2.13}$$

$$\frac{(B_r)^{n+3/2}_{i+1/2,j,k+1/2} - (B_r)^{n+1/2}_{i+1/2,j,k+1/2}}{\Delta t} = \frac{(E_\varphi)^{n+1}_{i+1,j,k+1/2} - (E_\varphi)^{n+1}_{i,j,k+1/2}}{\Delta z} - \frac{(E_z)^{n+1}_{i+1/2,j,k+1} - (E_z)^{n+1}_{i+1/2,j,k}}{r_j \Delta \varphi}, \tag{2.14}$$

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

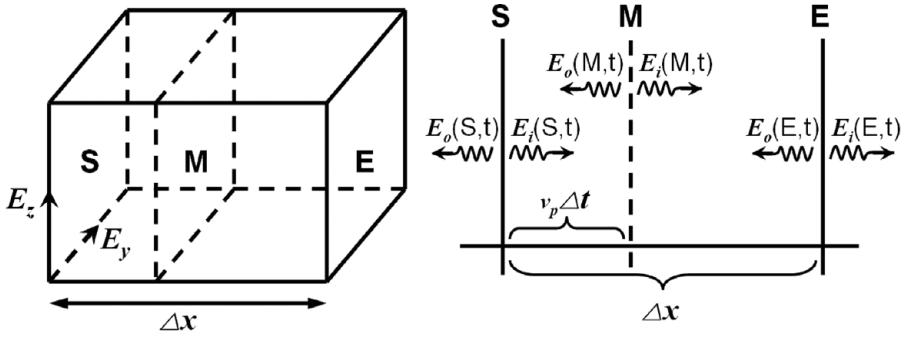


Figure 2. Model for the port-excitation source in waveguide.

$$\frac{(B_\varphi)^{n+3/2}_{i+1/2,j+1/2,k} - (B_\varphi)^{n+1/2}_{i+1/2,j+1/2,k}}{\Delta t} = \frac{(E_z)^{n+1}_{i+1/2,j+1,k} - (E_z)^{n+1}_{i+1/2,j,k}}{\Delta r} - \frac{(E_r)^{n+1}_{i+1,j+1/2,k} - (E_r)^{n+1}_{i,j+1/2,k}}{\Delta z}, \tag{2.15}$$

where Δt is the size of the time step, n is the index of the time step, Δz, Δr and $\Delta\varphi$ are the spatial steps in the z, r and φ directions, i, j and k are the indices of the spatial increments in three directions, respectively. Here, the subscripts denote the coordinates and the superscripts denote the time for the electromagnetic field components. Note that if they are used in the 2D case, the spatial step in the third direction $\Delta\varphi \to \infty$ should be applied in all the equations because of the symmetry. These equations are coupled by the constitutive equations $\mathbf{D} = \varepsilon\mathbf{E}$ and $\mathbf{B} = \mu\mathbf{H}$, where ε and μ are the permittivity and permeability of the media in the HPM devices, respectively.

The particle positions and velocities obey the Newton–Lorentz equations of motion

$$\frac{d}{dt}\gamma m\mathbf{v} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \tag{2.16}$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}, \tag{2.17}$$

where m, q, $\mathbf{x}$ and $\mathbf{v}$ are the mass, the charge, the position and the velocity of the particle, respectively. The relativistic factor γ is given by

$$\gamma = \sqrt{\frac{1}{1 - (|\mathbf{v}|/c)^2}} = \sqrt{1 + \left(\frac{|\mathbf{u}|}{c}\right)^2}, \tag{2.18}$$

$$\mathbf{u} = \gamma\mathbf{v}, \tag{2.19}$$

where c is the speed of light. In finite-difference form, the Newton–Lorentz equation is written as

$$\frac{\mathbf{u}^{n+1/2} - \mathbf{u}^{n-1/2}}{\Delta t} = \frac{q}{m}\left(\mathbf{E}^n + \frac{\mathbf{u}^n}{\gamma^n} \times \mathbf{B}^n\right), \tag{2.20}$$

$$\frac{\mathbf{x}^{n+1} - \mathbf{x}^n}{\Delta t} = \frac{\mathbf{u}^{n+1/2}}{\gamma^{n+1/2}}, \tag{2.21}$$

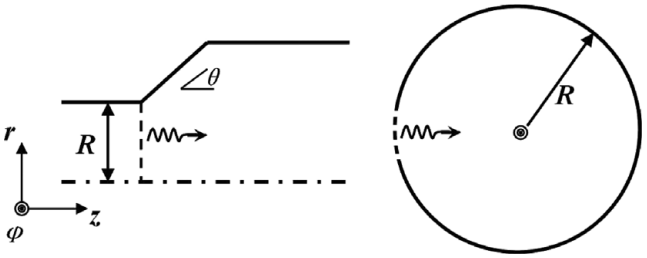


Figure 3. Examples for the port-excitation source in waveguide.

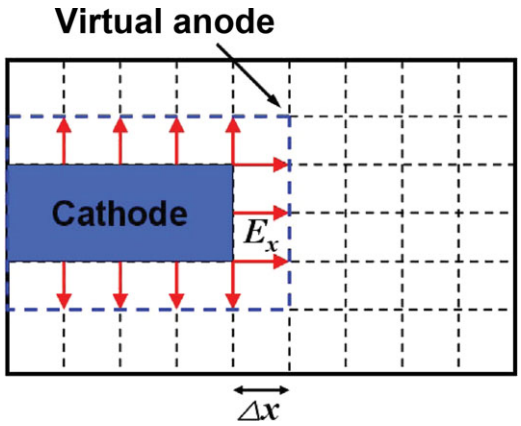


Figure 4. (Colour online) Virtual Child–Langmuir diodes for the space-charge-limited emission.

with

$$\mathbf{u}^n = \frac{\mathbf{u}^{n-1/2} + \mathbf{u}^{n+1/2}}{2}, \tag{2.22}$$

$$\gamma^n = \frac{\gamma^{n-1/2} + \gamma^{n+1/2}}{2}. \tag{2.23}$$

An efficient integration of the scheme, which avoids the full cross-product calculation is due to Boris (1970).

Based on the above formulae for updating the electromagnetic fields and the particle information, it is easy to list the order of a series of calculations. It is also called the leapfrog scheme, which has the additional merits of requiring few operations and minimal storage, since the update can be done in place.

The interaction of continuum particles with the discrete fields occurs through interpolation schemes

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

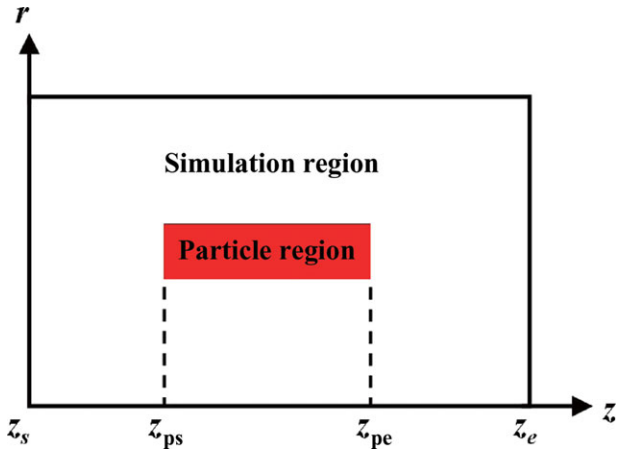


Figure 5. (Colour online) Start and end positions of the simulation region and the particle region in the z direction.

(Verboncoeur 2005). Here, not only the electromagnetic fields which act on the particles, but also the charge and current densities which are caused by the particles and react on the fields, are achieved by the linear weighting interpolation scheme.

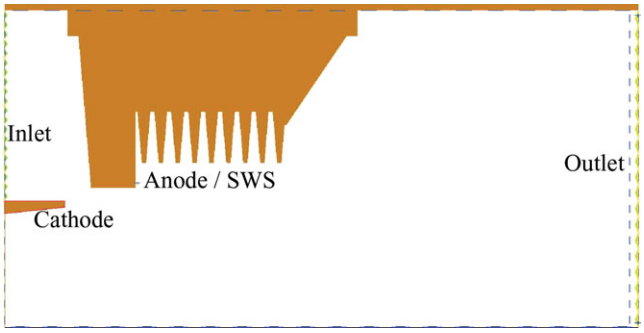


Figure 6. (Colour online) 2D model of the simulated relativistic BWO.

2.3. Boundary conditions

Besides the fundamental algorithms, various types of boundary conditions have also been implemented for different kinds of applications. The common boundary conditions supported are the perfectly conducting metal, the finite-conductivity metal, the dielectric material, the Gedney's perfectly matched layer (PML) wave absorber (Gedney 1996) and the symmetry boundaries. Sinusoidal electromagnetic wave can be driven by the current-density source. In addition, some useful models are

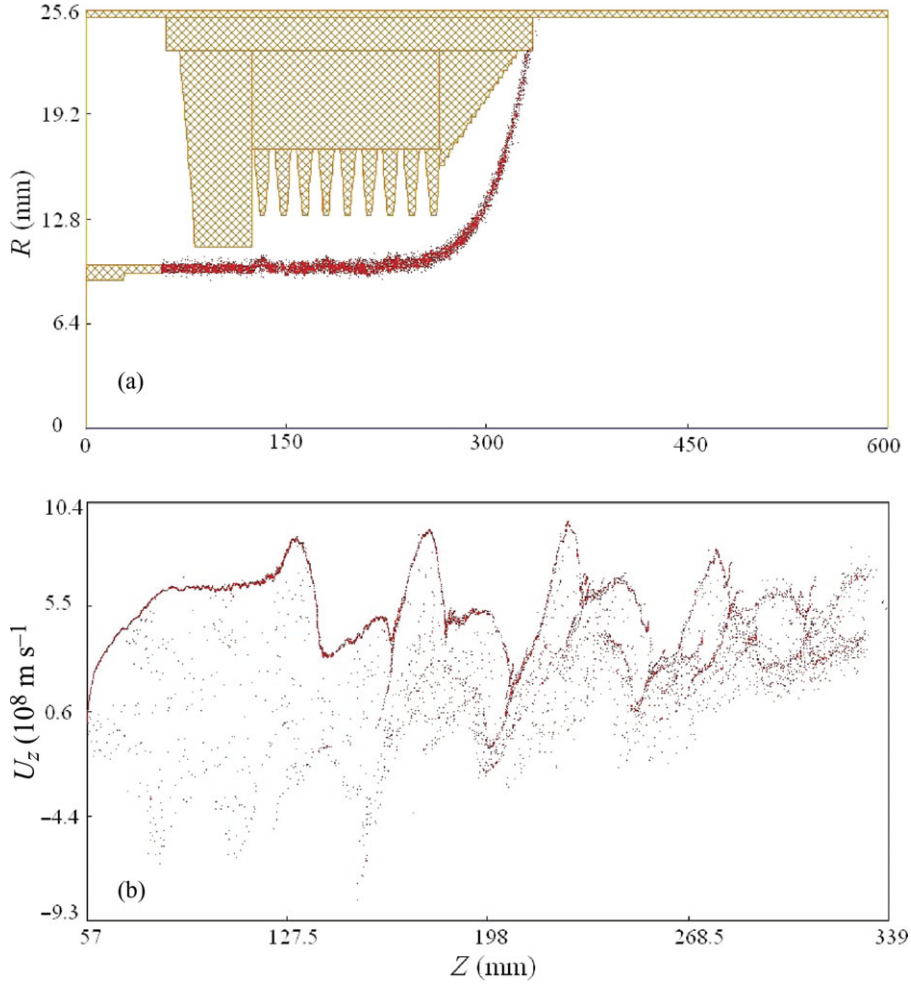


Figure 7. (Colour online) (a) Electron distribution and (b) (z, U_z) phase space of simulated relativistic BWO at 20 ns.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

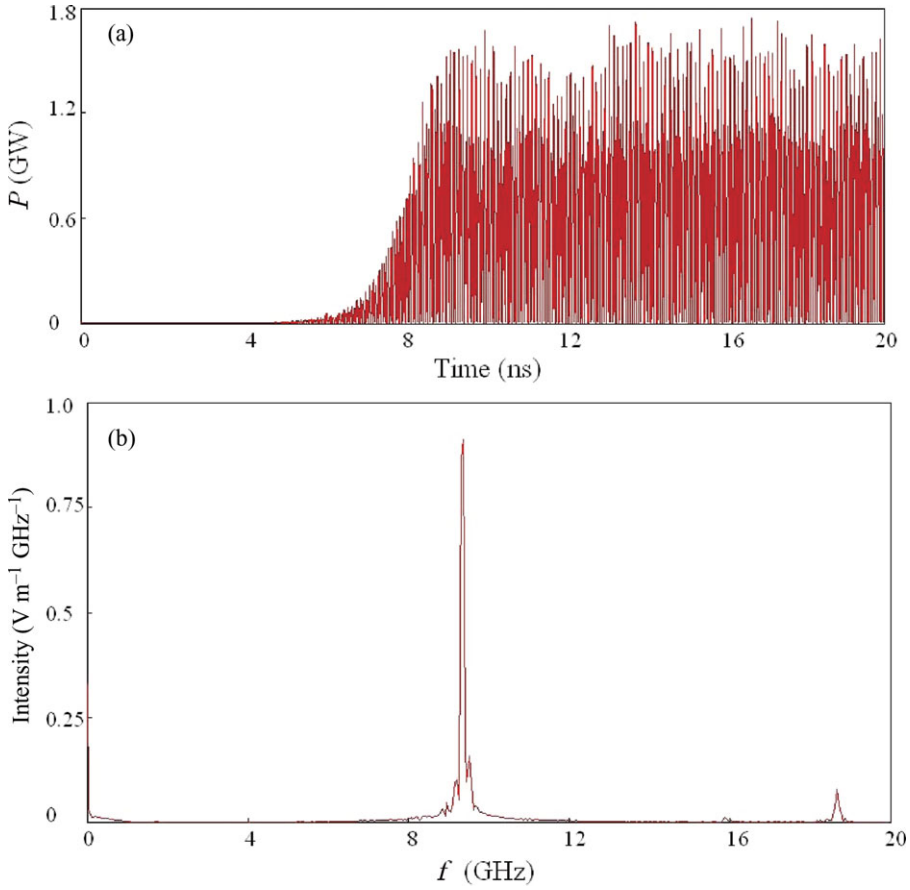


Figure 8. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by simulated relativistic BWO.

supported, too, including the electrostatic Poisson solver, the magnetostatic field coil, the foil and the strut. Here, we only focus on the boundary conditions for the sources of the input pulsed power (high driven voltage) and the relativistic electron beams, which are considered to be most important for the HPM devices.

Pulsed power can usually be launched by the port-excitation source as TEM mode electromagnetic waves in the coaxial waveguides. The algorithm is based on the total field/scattering field system and the propagation property of electromagnetic waves (Zhou et al. 2006). According to the total field/scattering field system, when electromagnetic waves are injected from outside, the total field on the boundary is the sum of the incoming and outgoing (scattering) fields, which can be written as

$$E_t = E_i + E_o = f(t)g(x) + E_o. \tag{2.24}$$

The incoming field E_i is defined by specifying a temporal function $f(t)$ and a spatial profile function $g(x)$. In order to obtain the total field E_t, the key step is the calculation of the outgoing field E_o, in which we should consider the propagation property of electromagnetic waves. Define v_p as the phase velocity of the electromagnetic wave on the port, and it is easy to realize that E_i launched at time t will propagate to a middle plane which has a distance $v_p\Delta t$ from the port plane after one time step Δt. In the same way, E_o on the same middle plane will propagate to the port plane during one time step; it means that E_o on the port plane at time $t+\Delta t$ equals E_o on the middle plane at time t.



Figure 9. (Colour online) 3D model of simulated relativistic BWO. It is in a sectional view for displaying the inner structures.

Since fields are all defined on Yee grids in the FDTD scheme, the calculation of E_o on the middle plane can

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

Table 1. Comparison of simulation results of BWO obtained by using CHIPIC and KARAT codes.

Code	Diode voltage (kV)	Diode current (kA)	Output power (MW)	Frequency (GHz)	Efficiency (%)
CHIPIC2D	700	5.5	750	9.35	19.5
CHIPIC3D	700	5.5	750	9.38	19.5
KARAT	730	5.6	805	9.3	20

be achieved through a linear weighting interpolation scheme, which is shown graphically in Fig. 2. Suppose the port-excitation source is on plane S, which is on the boundary and perpendicular to the x axis. The incoming field E_i propagates along the $+x$ direction and the outgoing field E_o propagates along the $-x$ direction. Plane E is the adjacent parallel plane with a distance Δx from S. Plane M is the middle plane between S and E, and the distance between S and M satisfies $v_p\Delta t \leqslant \Delta x$. The solution of E_o on the port plane at time $t+\Delta t$ can be predicted at time t as

$$E_o(S,t+\Delta t) = E_o(M,t) = \alpha E_o(S,t) + \beta E_o(E,t), \quad (2.25)$$

where $\beta = v_p\Delta t/\Delta x$ and $\alpha = 1-\beta = 1-v_p\Delta t/\Delta x$ are the weighting factors for the contributions of the outgoing fields on E and S, which are inverse to the distances from M. Here, $E_o(S,t)$ is initialized to be zero at time $t=0$ or has been calculated at time $t-\Delta t$, and $E_o(E,t)$ can be calculated by

$$E_o(E,t) = E_t(E,t) - E_i(E,t), \quad (2.26)$$

with

$$E_i(E,t) = \begin{cases} 0, & 0 \leqslant t < \Delta t', \\ E_i(S,t-\Delta t'), & t \geqslant \Delta t', \end{cases} \quad (2.27)$$

where $\Delta t' = \Delta x/v_p$ is the required time for the incoming field to propagate from S to E. Because the incoming field has not reached E until $\Delta t'$, so E_i on plane E is zero before that time.

From (2.24)–(2.27), the solutions of the total and outgoing fields on the port plane in a uniform waveguide can be well obtained from the user-defined incoming field and the total field stored on adjacent Yee grids.

Moreover, in order to make this algorithm more practical, we have considered a special case where the cross-section is not uniform along the normal direction of the port. In this case, an adjustment should be done with consideration of the influence of the cross-section on the fields near the port. When the cross-section is getting larger along the propagation direction, the fields will be 'stretched' to be weaker; in contrast, when the cross-section is getting smaller along the propagation direction, the fields will be 'compressed' to be stronger. So, the field values can approximately be treated as inversely proportional to the size of the cross-section. Define χ as the size ratio of plane S to plane E; we only need to replace $E_t(E,t)$ with $\chi E_t(E,t)$ in (2.26). Figure 3 shows two examples of how to set the ratio: in a conical waveguide excited on the smaller end, the fields near the port are weaker than the fields on the port,

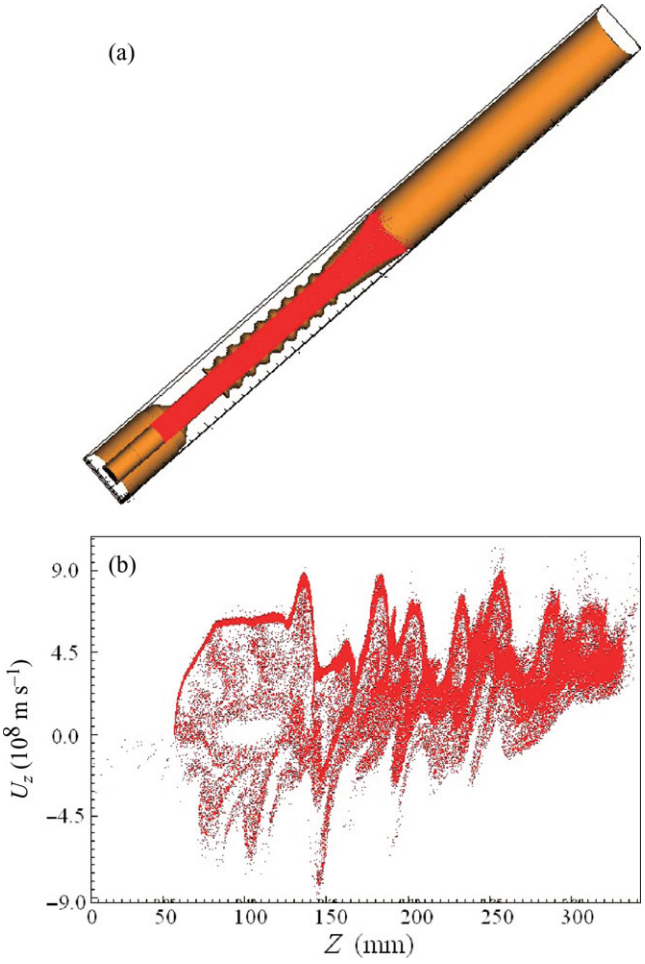


Figure 10. (Colour online) (a) Electron distribution and (b) (z, U_z) phase space of simulated relativistic BWO at 18 ns.

and the ratio is $\chi = R/(R+\Delta z\tan\theta)$; in a cylindrical waveguide excited on the side gap, the fields near the port are stronger than those on the port, and the ratio is $\chi = R/(R-\Delta R)$. Obviously, in a general uniform waveguide, $\chi \equiv 1$.

All things considered, it is easy to obtain the final iterative formulae for the port-excitation source with

$$E_o(S,0) = 0 \quad (2.28)$$

for the initialization at the beginning of the simulation, and

$$E_t(S,t) = E_i(S,t) + E_o(S,t) \quad (2.29)$$

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

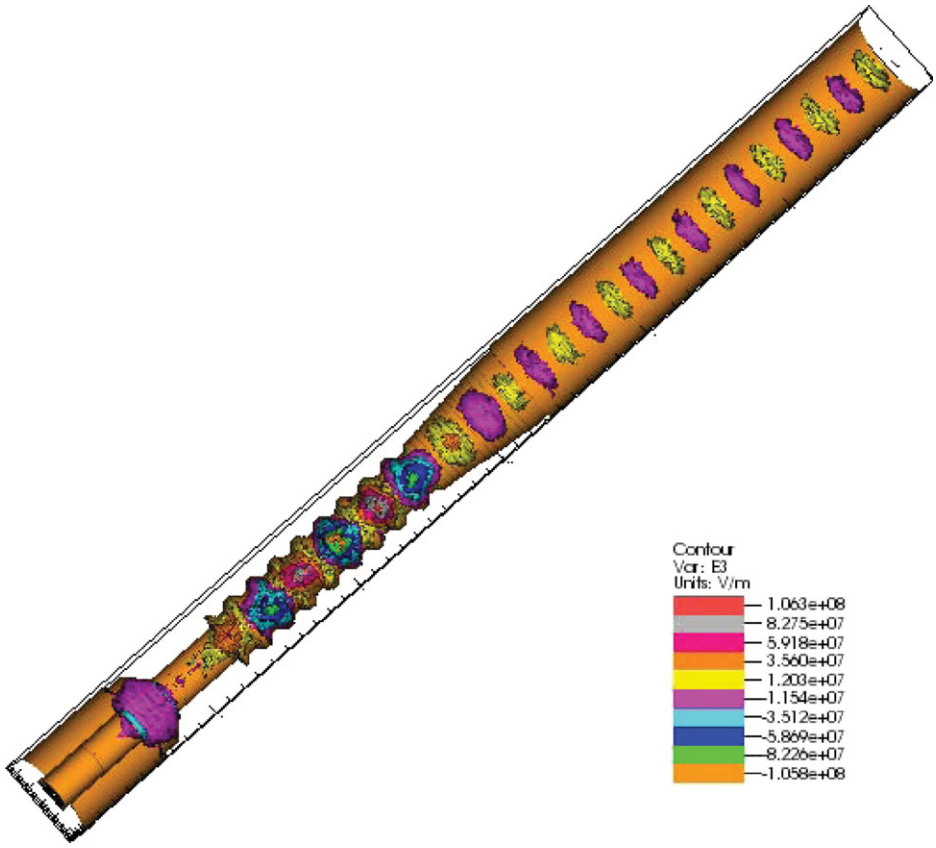


Figure 11. (Colour online) Contour plot of E_z in simulated relativistic BWO at 18 ns.

for the calculation of the total field in each time step, and

$$E_o(S,t+\Delta t) = \begin{cases} \alpha E_o(S,t) + \beta\chi E_t(E,t), & 0 \leqslant t < \Delta t', \\ \alpha E_o(S,t) + \beta\chi E_t(E,t) - \beta E_i(S,t-\Delta t'), & t \geqslant \Delta t', \end{cases} \tag{2.30}$$

for the prediction of the outgoing field which will be used in the next time step.

From the above analysis, we can see that this kind of an excitation source is very compact and suitable for the pulsed power transmission, of which the electromagnetic mode is TEM and the phase velocity is the same as the speed of light, that is $v_p = c$. Moreover, it can also be used for modeling other kinds of waveguide ports, because it is in a general form which only depends on the user-defined functions and the phase velocity. However, it gives perfect transmission only when the phase velocities match. When the user-defined incoming field is zero, it is just an outer boundary for outgoing waves. In this case, it is also suitable for single electromagnetic mode transmission with a single phase velocity. However, if multiple electromagnetic modes are present in a waveguide, inaccuracy is unavoidable, and the PML wave absorber is recommended to be a better choice.

Relativistic electrons can be emitted into the simulation region as part of a beam or through the space-charge limited (SCL) emission. Complex emission geometries are allowed with both the beam and SCL emission models. For the beam emission, the incident current density, the beam voltage and the direction are all user definable. For the SCL emission, the Child–Langmuir law is employed. The Yee grids on the emission surface are regarded as a number of virtual Child–Langmuir diodes with a gap of Δx and an applied voltage of $E_x\Delta x$ between the cathode and the anode, as shown in Fig. 4. Here, E_x is the electric field component in the normal direction x. Based on the Child–Langmuir law, the current density in each diode should obey the three-second power relationship, so the equation of emission current density can be written as

$$J_{\mathrm{CL}} = \frac{4\varepsilon_0}{9\Delta x^2}\left(\frac{2e}{m}\right)^{1/2}(-E_x\Delta x)^{3/2}, \tag{2.31}$$

where e is the electron charge and ε_0 is the vacuum permittivity. According to J_{CL}, a certain amount of macro-particles with finite velocities are emitted from the surface.

Electrons are directly destroyed (absorbed) on the surfaces of metals, dielectric materials, and PML volumes. When electrons pass through other boundaries, such as a port, and leave the simulation region, they are also automatically destroyed.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

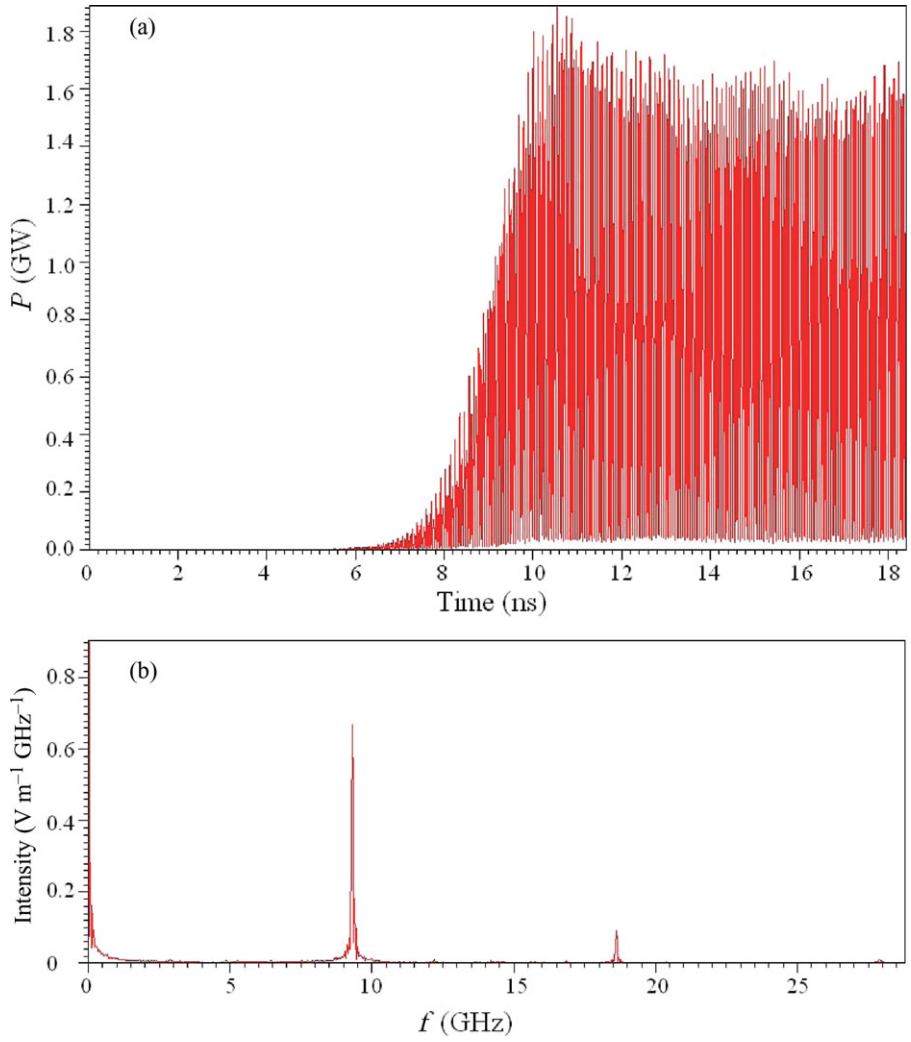


Figure 12. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by simulated relativistic BWO.

2.4. *Parallelization*

From the initial concept, the 3D CHIPIC code is developed with parallel computers in mind, in order to overcome the main time-consuming drawback of 3D PIC codes. It is designed to take full advantage of the speed and power of parallel architectures (Liao et al. 2009). The parallelization of the 3D CHIPIC code is accomplished using MPI and OpenMP. To enable the code to operate efficiently on parallel platforms, it is equipped with automated partitioning and dynamic load balancing (Sasser et al. 1998; Wang et al. 1999; Blahovec et al. 2000).

The MPI parallelization is based on the standard serial leapfrog loop, but has been specially modified to accommodate inter-processor communication requirements. The main modification contains the partitioning procedure and the communication procedure for field and particle information exchange. The partitioning is based on single-domain decomposition along the z coordinate direction. This method of decomposition is largely straightforward for MPI application and seems natural for most of the HPM devices. Once the problem is divided among two or more processors, certain cell and particle information needs to be shared between processors. The cells that need to share information between two processors are called 'border' cells. The initial partitioning is based on the number of cells along the z coordinate direction in the problem and the number of processors. It is usually uniform in default or configured by the user before the simulation. All of the communications are performed on border cells. Because borders are shared between processors, the amount of time spent on these communications is relatively equal across all processors. The simulation region except for the borders contains the majority of the computations, which are the internal field and particle computations. Therefore, if the total workload of the main region is evenly distributed among the allocated processors, the entire problem will be balanced, leading to efficient parallel computation.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

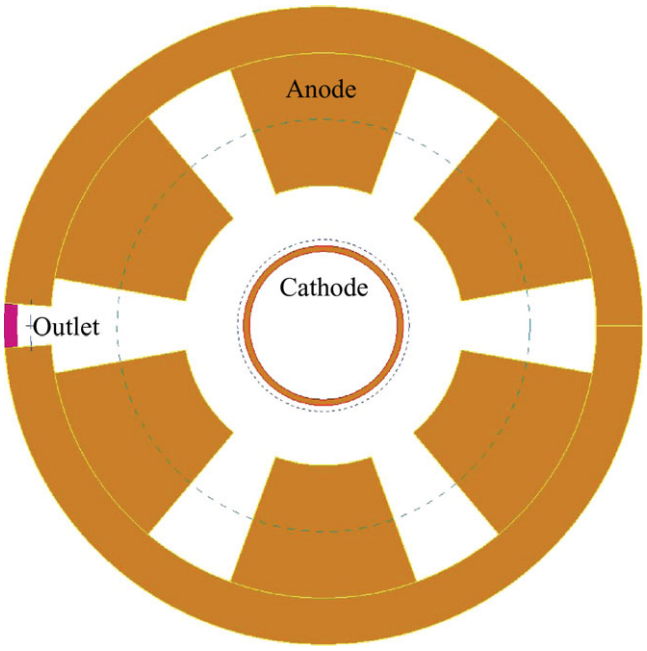


Figure 13. (Colour online) 2D model of the simulated A6 relativistic magnetron.

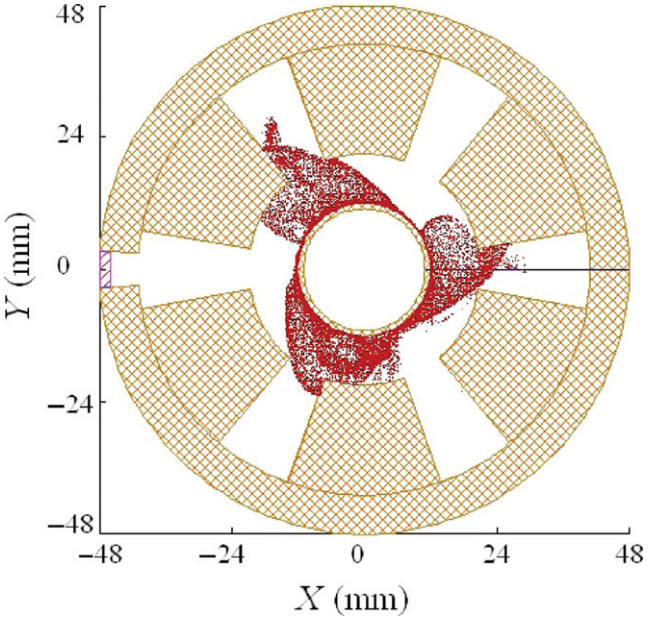


Figure 14. (Colour online) Electron distribution of the simulated relativistic magnetron at 30 ns.

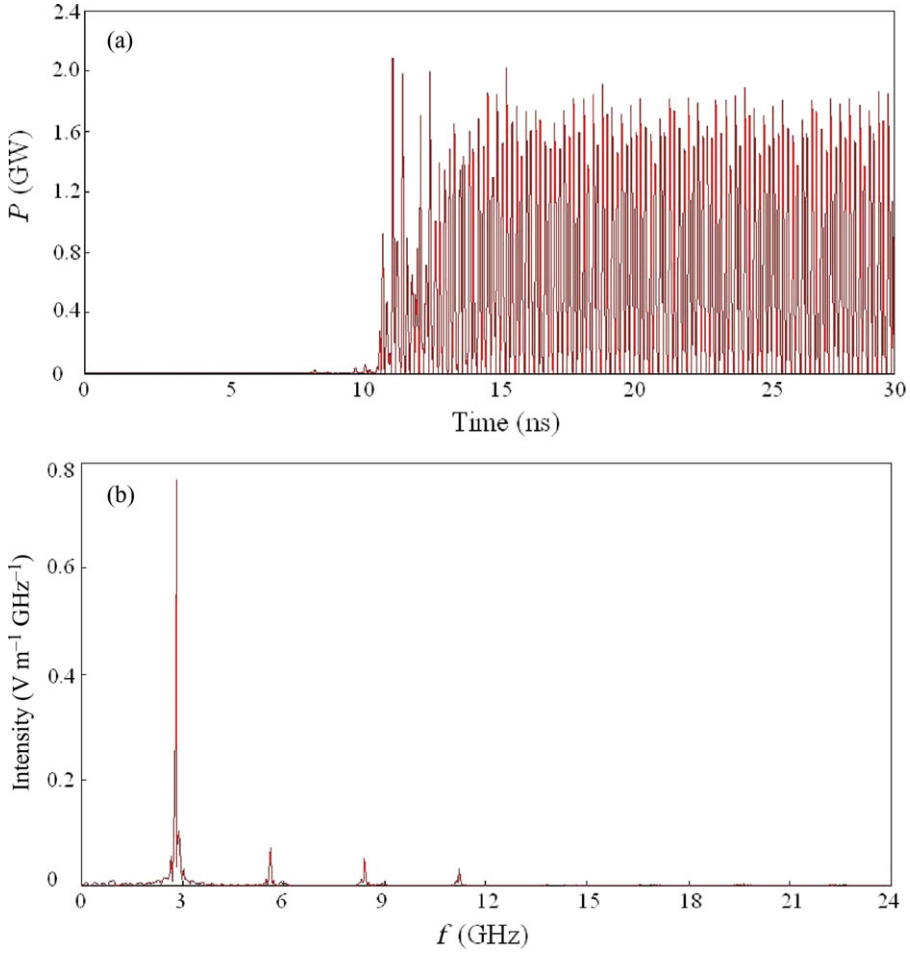


Figure 15. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by the simulated relativistic magnetron.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

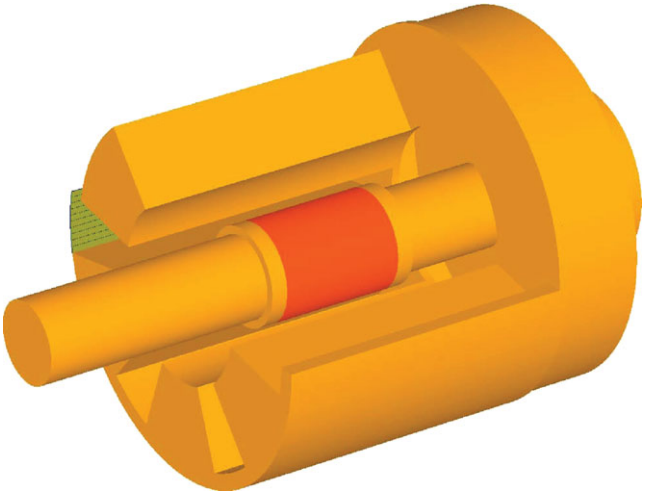


Figure 16. (Colour online) 3D model of the simulated A6 relativistic magnetron. Some outer structures are hidden for displaying the inner structures.

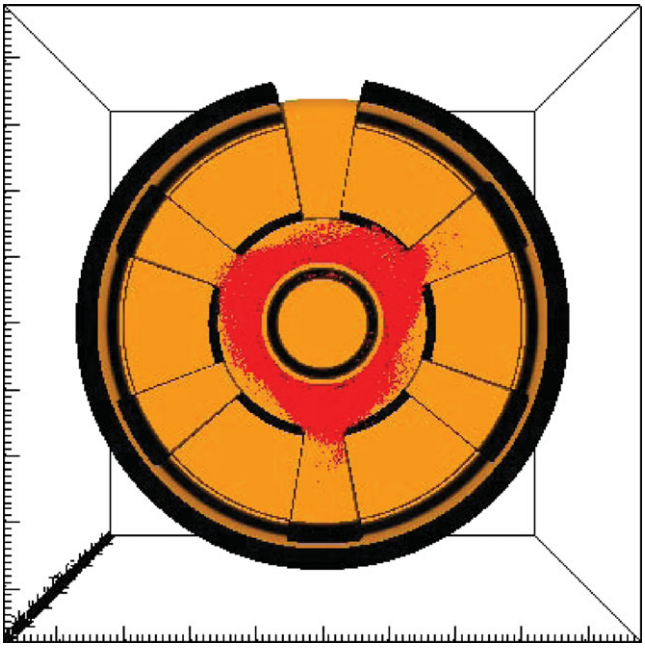


Figure 17. (Colour online) Electron distribution of the simulated relativistic magnetron at 28 ns.

To ensure that the simulation remains balanced, dynamic load balancing and partitioning are implemented to evenly spread the workload across the allocated processors. Because the particles are free to move through the device and may be created or absorbed at boundaries, the workload of each processor may change significantly throughout the duration of the simulation. If the workload becomes unbalanced, it is necessary to dynamically load balance and repartition the simulation during execution. However, one problem is that the repartitioning may produce even more workload and depress the efficiency of parallel computation, because a mass of repartitioning information (including the new partitioning positions and associated information about the grids, fields and particles) must be exchanged between the processors in each repartitioning. If the repartitioning procedure is executed continually, too much time will be consumed by the additional communications. With this consideration, CHIPIC is designed to detect the imbalance periodically using a suitable time interval, which is specified by the user. After each detection, the decision whether to repartition is determined by the degree of imbalance, which is defined as

$$\chi = \frac{W_{\max} - W_{\min}}{W_{\text{ave}}}, \tag{2.32}$$

where $W_{\max}$, $W_{\min}$ and W_{ave} are the maximum, the minimum and the average workload among the processors, respectively. Here, the workload on each processor is determined by monitoring the time required to complete a computational step. If an imbalance develops and χ exceeds the threshold specified by the user, CHIPIC will automatically reload balance and repartition itself.

The repartitioning is carefully designed based on the analysis of workload distribution. As mentioned before, the workload is dominated by the internal field and particle computations. Define T_{fld} and T_{par} as the total time consumed by the field computation and the particle computation in the last computational step, respectively. If $T_{\text{fld}} \geqslant T_{\text{par}}$, it means that the field computation is more dominant. Because field computation is carried out in the entire simulation region, we only need to adopt the default uniform partitioning to make the field computation balanced. In this case, the ith partitioning position can be determined by

$$z_j = z_s + \frac{z_e - z_s}{N} j, \quad (j = 1, 2, \cdots, N-1), \tag{2.33}$$

where N is the number of processors, z_s and z_e are the start and end positions of the simulation region in the z direction, which are indicated in Fig. 5 along with the start (z_{ps}) and end (z_{pe}) positions of the particle region in the z direction.

If $T_{\text{par}} > T_{\text{fld}}$, it means that the particle computation is more dominant, and we should develop another partitioning to make the particle computation as balanced as possible. Suppose the number of macro-particles is N_{par}, and define the required new average particle number for each process as $N_{\text{ave}} = N_{\text{par}}/N$, the first partitioning position can be determined by summing up the particle numbers on the r–φ mesh planes from the first grid along the z direction. Suppose the summed particle number until the mth mesh plane is N_m. Once $N_m > N_{\text{ave}} \geqslant N_{m-1}$, if $N_m - N_{\text{ave}} > N_{\text{ave}} - N_{m-1}$, set the cells on the $(m-1)$th and mth mesh planes as the first border cells, otherwise set the cells on the mth and $(m+1)$th mesh planes. Other partitioning positions can be determined in a similar way.

We have also considered a special case where z_{pe} is not much larger than z_{ps} compared with the size of the

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

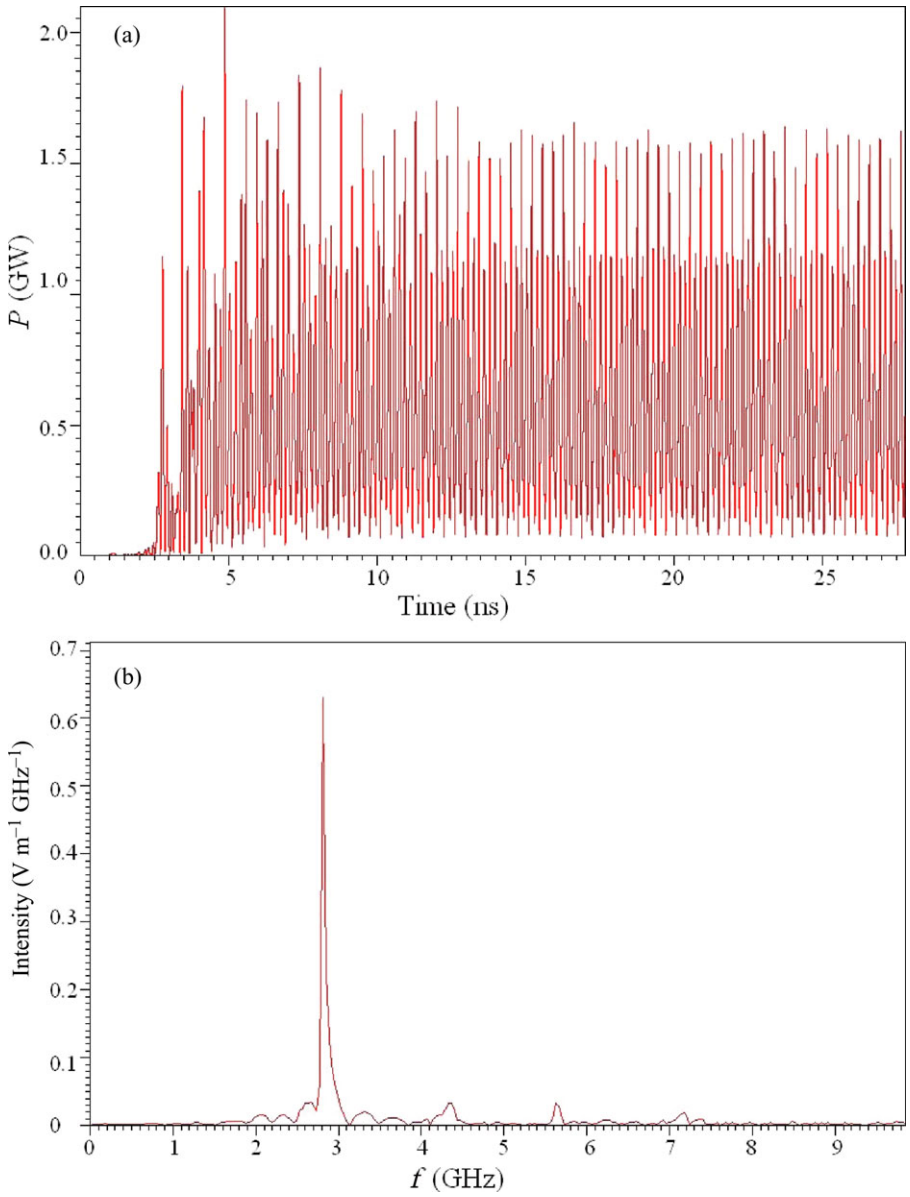


Figure 18. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by the simulated relativistic magnetron.

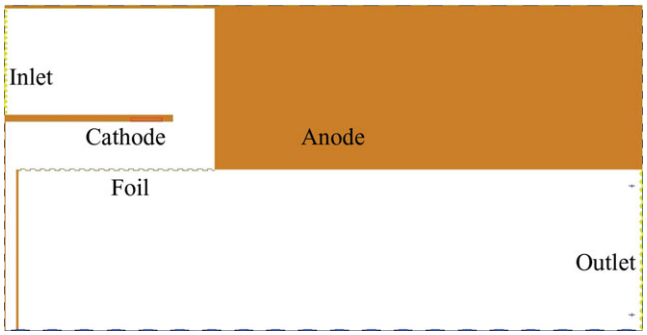


Figure 19. (Colour online) 2D model of the simulated relativistic coaxial vircator.

entire simulation region, which means that most of the macro-particles are restricted in a small particle region. In this case, the above partitioning may cause many small sub-regions which only contain the border cells and several inner cells. When the number of inner cells does not reach a certain scale, the efficiency will be very low because a significant percentage of total time will be consumed by the communications. In order to avoid this problem, only one or two processors are appointed to simulate the particle region in default, and the other processors are appointed to simulate the two outboard regions using the uniform partitioning.

In addition, the OpenMP parallelization is introduced into this code. Because OpenMP supports shared-memory parallel programming particularly for applications with program loops, it is very suitable for PIC codes. It can additionally distribute the calculation of time-consuming loops to multi-threads in each MPI process.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

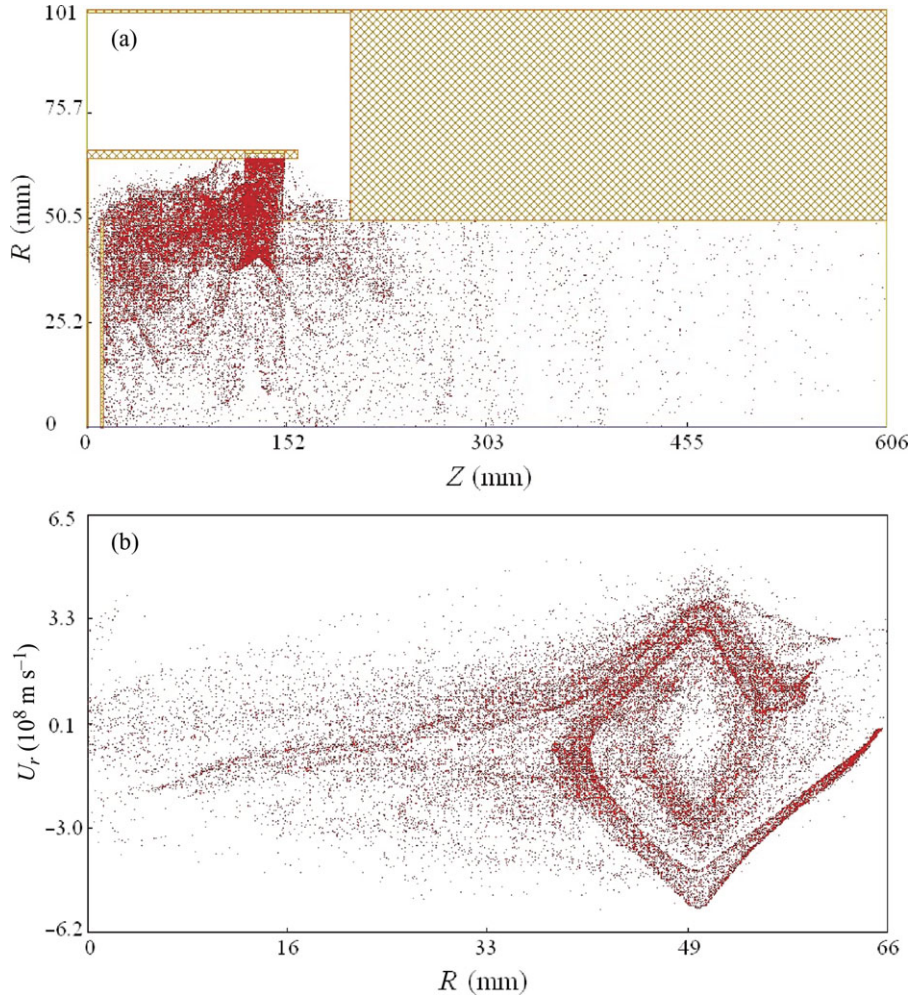


Figure 20. (Colour online) (a) Electron distribution and (b) (r, U_r) phase space of the simulated vircator at 15 ns.

3. Simulations and results

The CHIPIC code has been widely used to develop and design HPM devices, including relativistic klystrons, relativistic BWOs, relativistic magnetrons, MILOs, vircators, SCOs and gyrodevices. In this section, numerical experiments on some typical HPM devices by using this code are given, together with the comparisons with the simulation results obtained from some well-known PIC codes.

3.1. O-type devices

The O-type devices are also called the linear-beam devices, which are characterized by a linear electron beam. The external magnetic field is usually applied parallel to the moving direction of electrons to confine the beam. The relativistic klystrons and BWOs are two typical O-type HPM devices.

The 2D model of a relativistic BWO (Moreland et al. 1996) is shown in Fig. 6, in which a corrugated waveguide slow wave structure (SWS) and an annular electron beam are adopted. The SWS consists of nine periods, with the inner radius, the outer radius and the period space are 1.3, 1.7 and 1.6 cm, respectively. The inner and outer radii of the annular electron beam are 0.95 and 1 cm. This device is driven by a voltage of 500 kV with the rise time of 5 ns. The applied magnetic field is 2.8 T. The whole computational region is divided into 980×60 cells, the size of time step is 0.89 ps, and the steady-state macro-particle number is about 6500. It takes about 20 min on one HP-PC for a 20-ns simulation. Typical results are shown in Figs 7 and 8. The measured diode voltage and diode current are 700 kV and 5.5 kA, respectively. The average output power is about 750 MW, which corresponds to an efficiency of 19.5% according to the input dc power $P_{dc} = UI = 3.85$ GW. The domain frequency of the output signal is 9.35 GHz.

The 3D model of the same relativistic BWO is shown in Fig. 9. The parameters are the same as those of the 2D simulation. The whole computational region is divided into $588 \times 54 \times 10$ cells, the size of time step is 0.20 ps, and the steady-state macro-particle number is about 60 000. For a 18-ns simulation, it takes about 20 h on one HP-PC; in contrast, it takes only 4 h on a small network

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

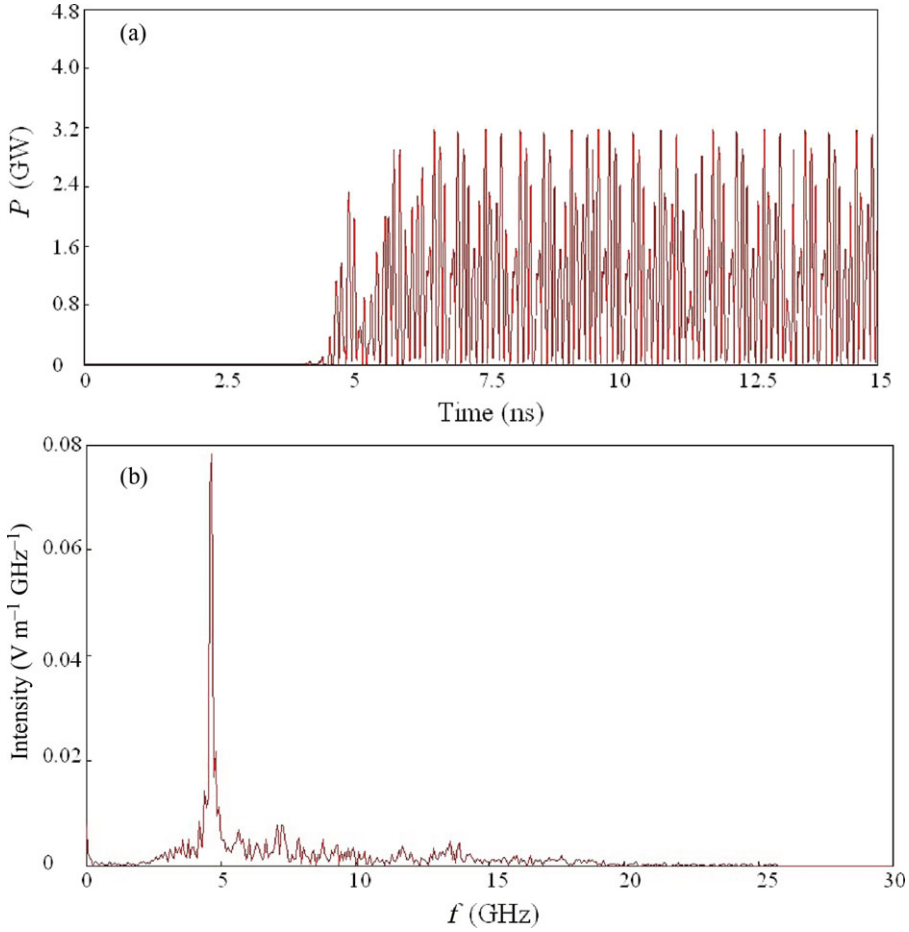


Figure 21. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by the simulated vircator.

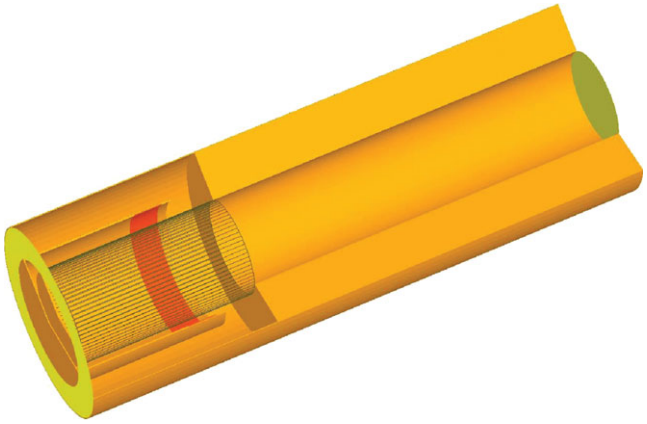


Figure 22. (Colour online) 3D model of the simulated relativistic coaxial vircator. It is in a sectional view for displaying the inner structures.

with six HP-PCs. Typical results are shown in Figs 10–12. The generated microwave is exported through TM_{01} mode in the cylindrical waveguide. The average output power is about 750 MW and the domain frequency of the output signal is 9.38 GHz. These results are very consistent with those obtained from the 2D simulation and agree well with those obtained by using the KARAT code, as summarized in Table 1.

3.2. M-type devices

The M-type devices are also called the crossed-field devices, which means the electric field and magnetic field are crossed. Unlike the O-type devices, the magnetic field is perpendicular to the moving direction of electrons and plays an extremely important role for the fundamental mechanisms of the devices. The relativistic magnetrons and MILOs are two typical kinds of M-type HPM devices.

The 2D model of an A6 relativistic magnetron (Chan et al. 1990) is shown in Fig. 13. The radii of the cathode and the anode shell are 1.2 cm and 4.1 cm, respectively. There are six resonant cavities which are separated by six anode vanes with the inner radius of 2.1 cm. Each cavity has a ringent angle of 20°. One of the cavities is used for power extraction. This device is driven by a voltage of 500 kV with a rise time of 4 ns. The applied external magnetic field is 0.5 T. The 2D polar coordinate system (r–φ) is used to simulate this device. The whole computational region is divided into 40 × 180 cells, the size of time step is 0.98 ps, and the steady-state

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

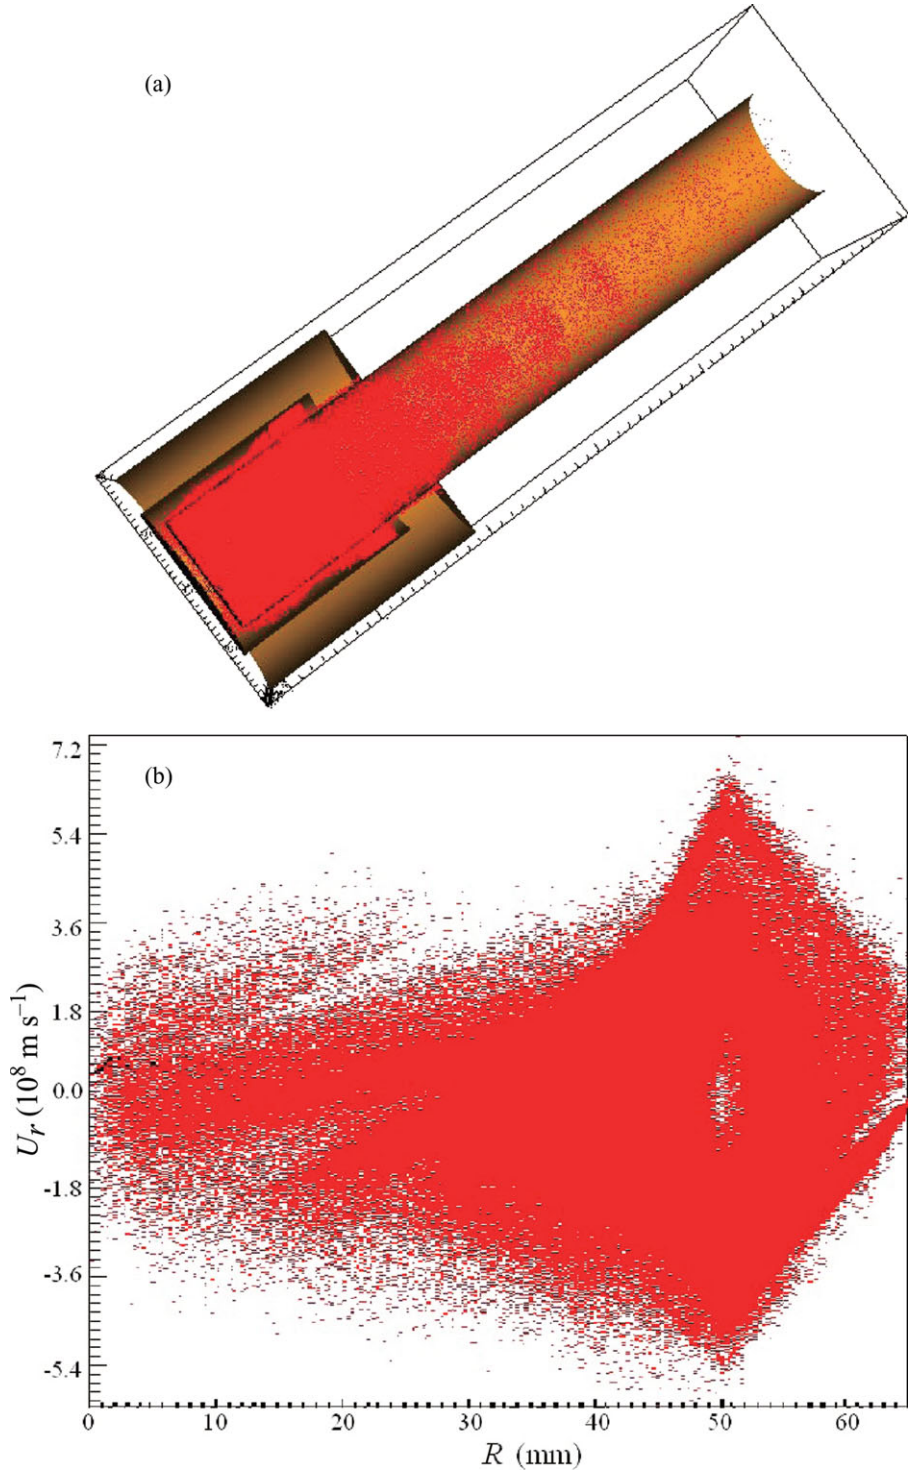


Figure 23. (Colour online) (a) Electron distribution and (b) (r, U_r) phase space of the simulated vircator at 20 ns.

macro-particle number is about 22 000. It takes about 20 min on one HP-PC for a 30-ns simulation. Typical results are shown in Figs 14 and 15. We can see this device is working at π mode, the number of electron spokes is 3, which is half of that of the resonant cavities. The average output power is about 700 MW and the domain frequency of the output signal is 2.80 GHz.

The 3D model of the same A6 relativistic magnetron is shown in Fig. 16. The parameters are the same as those of the 2D simulation, besides the height of the cathode and the anode are 3.0 and 7.2 cm, respectively. The whole computational region is divided into $76 \times 45 \times 180$ cells, the size of time step is 0.23 ps, and the steady-state macro-particle number is about 500 000. For a

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

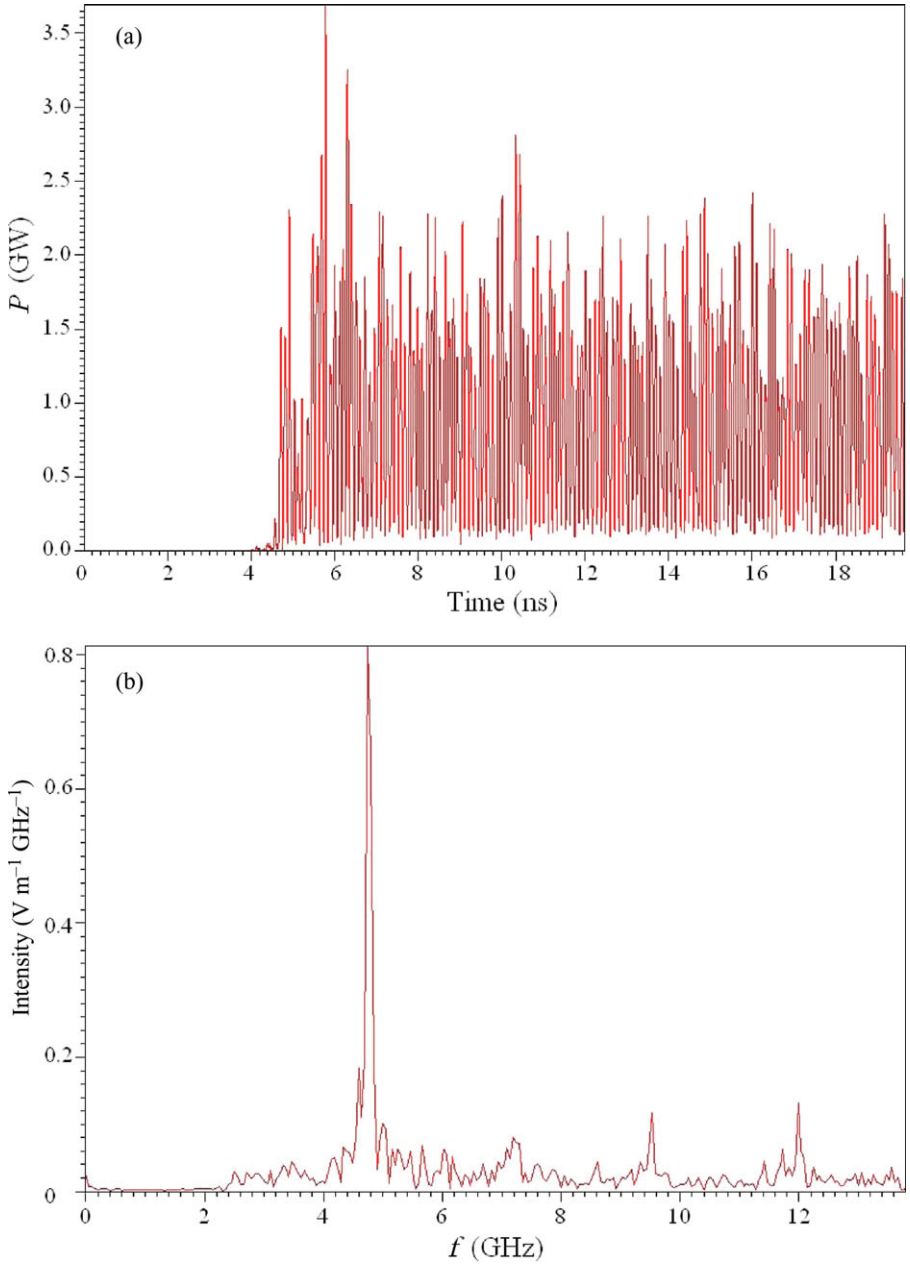


Figure 24. (Colour online) (a) Output microwave power and (b) spectrum of the output signal produced by the simulated vircator.

28-ns simulation, it takes about 26 h on one HP-PC; in contrast, it takes only 7 h on a small network with six HP-PCs. Typical results are shown in Figs 17 and 18. The average output power is about 700 MW and the domain frequency of the output signal is 2.81 GHz. These results are very consistent with those obtained from the 2D simulation and agree well with those obtained by using the MAGIC code, as summarized in Table 2.

3.3. Space-charge-type devices

The space-charge-type devices are another important kind of HPM devices besides the previous two traditional types. Their mechanisms are based on the space-charge effect in conductor waveguides and cavities. The vircators and SCOs are two typical space-charge-type HPM devices.

The 2D model of a relativistic coaxial vircator (Jiang et al. 1999) is shown in Fig. 19. The inner and outer radii of the input coaxial waveguide are 6.7 and 10 cm, respectively. The radii of the anode film and the output cylindrical waveguide are both 5 cm. This device is driven by a voltage of 600 kV with a rise time of 2 ns. The whole computational region is divided into 606×100 cells, the size of time step is 1.89 ps, and the steady-state macro-particle number is about 35 000. It takes about 10 min on one HP-PC for a 15-ns simulation. Typical results are shown in Figs 20 and 21. The measured diode voltage and diode current are 500 kV and 30

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

Table 2. Comparison of simulation results of relativistic magnetron obtained by using CHIPIC and MAGIC codes.

Code	Diode voltage (kV)	Dominant mode	Output power (MW)	Frequency (GHz)
CHIPIC2D	500	π mode	700	2.80
CHIPIC3D	500	π mode	700	2.81
MAGIC	500	π mode	700	2.80

Table 3. Comparison of simulation results of relativistic coaxial vircator obtained by using CHIPIC and KARAT codes.

Code	Diode voltage (kV)	Diode current (kA)	Output power (MW)	Frequency (GHz)	Efficiency (%)
CHIPIC2D	500	30	950	4.65	6.3
CHIPIC3D	500	30	950	4.7	6.3
KARAT	490	30	970	4.6	6.6

kA, respectively. The average output power is about 950 MW, which corresponds to an efficiency of 6.3% according to the input dc power $P_{dc} = UI = 1.5\,\mathrm{GW}$. The domain frequency of the output signal is 4.65 GHz.

The 3D model of the same relativistic coaxial vircator is shown in Fig. 22. The parameters are the same as those in the 2D simulation. The whole computational region is divided into $606 \times 100 \times 24$ cells, the size of time step is 0.41 ps, and the steady-state macro-particle number is about 630 000. For a 20-ns simulation, it takes about 22 h on one HP-PC; in contrast, it takes only 5 h on a small network with six HP-PCs. Typical results are shown in Figs 23 and 24. The average output power is about 950 MW and the domain frequency of the output signal is 4.70 GHz. These results are very consistent with those obtained from the 2D simulation and agree well with those obtained by using the KARAT code, as summarized in Table 3.

4. Summary

This paper introduces the CHIPIC code for modeling and simulating HPM devices with and without symmetric structure. The construction, algorithms and parallelization used in this code are presented. In the CHIPIC code, the electromagnetic fields are calculated using the FDTD scheme, and the particles are described by the PIC scheme. It can model and simulate HPM devices on Cartesian, cylindrical and polar grids. For the purpose of highly efficient modeling and simulations, the MPI and OpenMP are adopted for parallelization. Its parallel design ensures that it is capable of efficiently executing on a variety of architectures. In order to allow efficient use of parallel architectures, it provides automated partitioning and dynamic load balancing. Numerical experiments on some typical HPM devices by using the CHIPIC code are given, which agree well with those obtained from some well-known PIC codes.

Although this code has successfully simulated various vacuum electric devices, it is not capable of dealing with many complicated phenomena related to plasmas and gas. This is mainly because of the complexity of collision processes and the huge amount of calculation. In addition, the parallel algorithm should be significantly improved in order to fulfill the increasing requirements for simulating complex devices.

Therefore, ongoing research includes improvement of parallel processing and addition of models for boundary conditions, plasmas, ionizations and collisions. A version with the Monte Carlo collision model is in development for extending applications of the code. In order to allow recovery if computational nodes fail during lengthy parallel simulations, the checkpoint restart procedure is being considered. Both real-time and post-process graphics developments continue. In addition, we continue to refine and optimize the code.

Acknowledgements

Authors would like to acknowledge the research group of the High Performance Computing Center in the Institute of Applied Physics and Computational Mathematics, Beijing, for the powerful tools for data management and visualization. This work was supported by the Fundamental Research Funds for the Central Universities under contract ZYGX2011J037.

References

Antonsen, Jr, T. M., Mondelli, A. A., Levush, B., Verboncoeur, J. P. and Birdsall, C. K. 1999 Advances in modeling and simulation of vacuum electronic devices. *Proc. IEEE* **87**, 804–839.

Barker, R. J. and Schamiloglu, E. 2001 *High-Power Microwave Sources and Technologies*. New York: IEEE Press.

Benford, J., Swegle, J. and Schamiloglu, E. 2007 *High Power Microwaves*, 2nd edn. New York: Taylor & Francis.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press

Birdsall, C. K. 1991 Particle-in-cell charged-particle simulations, plus Monte Carlo collisions with neutral atoms, PIC-MCC. *IEEE Trans. Plasma Sci.* **19**, 65–85.

Birdsall, C. K. and Langdon, A. B. 1985 *Plasma Physics via Computer Simulation.* New York: McGraw-Hill.

Blahovec, J. D., Bowers, L. A., Luginsland, J. W., Sasser, G. E. and Watrous, J. J. 2000 3D ICEPIC simulations of the relativistic Klystron oscillator. *IEEE Trans. Plasma Sci.* **28**, 821–829.

Boris, J. P. 1970 Relativistic plasma simulation: optimization of a hybrid code. In: *Proc. 4th Conf. on Numerical Simulation of Plasmas*. Washington, DC: Naval Research Laboratory, pp. 3–67.

Chan, H. W., Chen, C. and Davidson, R. C. 1990 Computer simulation of relativistic multiresonator cylindrical magnetrons. *Appl. Phys. Lett.* **57**, 1271–1273.

Chu K. R. 2004 The electron cyclotron maser. *Rev. Mod. Phys.* **76**, 489–540.

Clark, R. E. and Hughes, T. P. 2005 *LSP Users Manual and Reference for LSP Version 8.7.* Newington, VA: ATK Mission Systems Inc.

Gaponov-Grekhov, A. V. and Granatstein, V. L. 1994 *Applications of High-Power Microwaves.* Boston, MA: Artech House.

Gedney, S. D. 1996 An anisotropic perfectly matched layer-absorbing medium for the truncation of FDTD lattices. *IEEE Trans. Antennas Propag.* **44**, 1630–1639.

Geng, Z. H., Liu, P. K., Su, Y. N., Xu, S. X. and Xue, Q. Z. 2010 Design of a Ka band 35 kW CW low-voltage harmonic gyrotron. *J. Infrared Milli. Terahertz Waves* **31**, 41–47.

Geng, R. L., Padamsee, H., Belomestnykh, S., Goudket P., Dykes, D. M. and Carter, R. G. 2003 Suppression of multipacting in rectangular coupler waveguides. *Nucl. Instr. Meth.* A **508**, 227–238.

Godfrey, B. B. 1980 Time-biased field solver for electromagnetic PIC codes. In: *Proc. 9th Conf. on Numerical Simulation of Plasmas*, Evanston, IL.

Goplen, B., Ludeking, L., Smithe, D. and Warren, G. 1995 User-configurable MAGIC for electromagnetic PIC calculations. *Comput. Phys. Commun.* **87**, 54–86.

Grisham, L. R. 2009 Magnetic insulation to improve voltage holding in electrostatic accelerators. *Phys. Plasmas* **16**, 043111.

Hockney, R. W. and Eastwood, J. W. 1981 *Computer Simulation Using Particles.* New York: McGraw-Hill.

Jiang, W., Woolverton, K., Dickens, J. and Kristiansen, M. 1999 High power microwave generation by a coaxial virtual cathode oscillator. *IEEE Trans. Plasma Sci.* **27**, 1538–1542.

Lau, Y. Y., Friedman, M., Krall, J. and Serlin, V. 1990 Relativistic Klystron amplifiers driven by modulated intense relativistic electron beams. *IEEE Trans. Plasma Sci.* **18**, 553–569.

Lemke, R. W., Calico, S. E. and Clark, M. C. 1997 Investigation of a load-limited, magnetically insulated transmission line oscillator (MILO). *IEEE Trans. Plasma Sci.* **25**, 364–374.

Lemke, R. W., Clark, M. C. and Marder, B. M. 1994 Theoretical and experimental investigation of a method for increasing the output power of a microwave tube based on the split-cavity oscillator. *J. Appl. Phys.* **75**, 5423–5432.

Li, W. and Liu, Y. G. 2010 Choosing optimum method for the efficient design of a relativistic magnetron with diffraction output. *J. Appl. Phys.* **108**, 113303.

Li, W. and Liu, Y. G. 2011 Modified magnetic field distribution in relativistic magnetron with diffraction output for compact operation. *Phys. Plasmas* **18**, 023103.

Liao, C., Liu, D. and Liu, S. 2009 Three-dimensional electromagnetic particle-in-cell simulation by parallel computing. *Acta Phys. Sin.* **58**, 6709–6718.

Moreland, L. D., Schamiloglu, E., Lemke, R. W., Roitman, A. M., Korovin, S. D. and Rostov, V. V. 1996 Enhanced frequency agility of high power relativistic backward wave oscillators. *IEEE Trans. Plasma Sci.* **24**, 852–857.

Nieter, C. and Cary, J. R. 2004 VORPAL: a versatile plasma simulation code. *J. Comput. Phys.* **196**, 448–473.

Pukhov, A. 1999 Three-dimensional electromagnetic relativistic particle-in-cell code VLPL (virtual laser plasma lab). *J. Plasma Phys.* **61**, 425–433.

Qin, F., Wang, D., Wen, J., Chen, D. B. and Fan, Z. K. 2011 A novel method to depress higher order mode generation in MILO. *IEEE Trans. Plasma Sci.* **39**, 545–549.

Sasser, G. E., Havranek, J. J., Colella, S. L. and Watrous, J. W. 1998 Modified PIC algorithm for efficient multiprocessor simulations. In: *Proc. 16th Conf. Numerical Simulations of Plasmas*, Santa Barbara, CA.

Stratakis, D., Gallardo, J. C. and Palmer, R. B. 2011 Enhancement of accelerating field of microwave cavities by magnetic insulation. *Nucl. Instr. Meth.* A **643**, 1–6.

Taflove, A. and Hagness, S. C. 2000 *Computational Electrodynamics*, 2nd edn. Norwood: Artech House.

Tang, Y., Meng, L., Li, H., Zheng, L., Yin, Y. and Wang, B. 2012 A dual-frequency coaxial relativistic backward-wave oscillator with a modulating resonant reflector. *Phys Scr* **85**, 055801.

Tarakanov, V. P. 2001 *User's Manual for Code KARAT.* Springfield, VA: Berkeley Research Associates.

Verboncoeur, J. P. 2005 Particle simulation of plasmas: review and advances. *Plasma Phys. Control. Fusion.* **47**, A231–A260.

Verboncoeur, J. P., Langdon, A. B. and Gladd, N. T. 1995 An object-oriented electromagnetic PIC code. *Comput. Phys. Commun.* **87**, 199–211.

Wang, J., Kondrashov, D., Liewer, P. C. and Karmesin, S. R. 1999 Three-dimensional deformable-grid electromagnetic particle-in-cell for parallel computers. *J. Plasma Phys.* **61**, 367–389.

Yee, K. S. 1966 Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Trans. Antennas Propag.* **14**, 302–307.

Zhou, J., Liu, D., Liao, C. and Li, Z. 2009 CHIPIC: an efficient code for electromagnetic PIC modeling and simulation. *IEEE Trans. Plasma Sci.* **37**, 2002–2011.

Zhou, J., Zhu, D., Liu, D. and Liu, S. 2006 Design and realization of waveguide excitation source in particle-in-cell simulation. *High Power Laser Part. Beams* **18**, 2019–2024.

https://doi.org/10.1017/S0022377812000724 Published online by Cambridge University Press