Efficient dual space source interpolation method for the numerical solution of self-consistent static beam-wave interactions

A. Gold* and S. Tantawi

SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA

(Received 1 June 2018; published 7 November 2018)

Many numerical solvers model the general nonlinear problem of particle motion in self-generated fields through the computation of particle trajectories and fields in a Lagrangian and Eulerian framework, respectively. The particle-weighting step mapping the source terms between these frameworks is a primary source of noise and computation time. Herein we eliminate this costly step by treating the trajectories as bounds on the particle flow, applying charge conservation to interpolate the source distribution. The result is an order of magnitude improvement in accuracy and an 80 to 120 fold speed-up.

DOI: 10.1103/PhysRevAccelBeams.21.114403

I. INTRODUCTION

Modeling the nonlinear interaction between electromagnetic fields and relativistic charged particle beams is a challenge pervading many areas in modern physics and engineering including accelerator physics, vacuum electronics, and plasma physics. Analytically intractable, the propagation of space-charge dominated beams (where the beam is significantly affected by its own self-fields) is traditionally solved using either large signal codes or the particle-in-cell (PIC) method.

Large signal or parametric interaction codes such as TESLA [1] or MAGY [2] solve reduced systems where approximations have been made to simplify the governing equations. For example, in TESLA the telegrapher's equations are solved for a linearly propagating beam where the fields in the beam tube and the rf cavities are treated separately. The fields are solved using a modal decomposition in a paraxial approximation (the field envelope varies slowly) and interactions between beams and coupling with external cavities are accounted for though special boundary conditions on the wall.

PIC codes such as MAGIC [3], CST Studio [4] and WARP [5] simultaneously solve the equations of motion for charged particles and Maxwell's equations for the resulting electromagnetic fields over discrete time steps. While powerful in their ability to fully capture the physics involved, by solving in the time domain they are inefficient at modeling steady-state problems. Particularly within the

context of radio-frequency sources and accelerator structures, where the relevant problem features tend to span multiple orders of magnitude both spatially and temporally. For example, modeling a klystron requires resolving phenomena at time steps smaller than the radio-frequency (rf) period (on the ps–ns scale) while the simulation must proceed until steady state is reached (μ s–ms scale).

We are developing an alternative solver for steady state static and time harmonic beam-wave interactions that solves for the full fields, including the DC and AC space charge contributions, while solving for the steady state solution directly. It thus combines the generality of the a PIC code with the computational efficiency of a large signal code. Trajectories, external fields and self-fields are computed over the entire problem domain instead of at discrete time steps and iterated until converging to the self-consistent solution, as illustrated in Fig. 1. Not only is this more efficient, taking on the order of ten iterations to converge instead of thousands to millions of time steps, but it allows for a complete accounting of space charge, image currents,

FIG. 1. An iterative steady state beam-wave solver. This paper introduces a new, more efficient source term calculation.

*vriellink@stanford.edu

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

and system boundaries with no assumptions on the relativistic regime of the beam.

A shared challenge amongst all of these approaches, transient and steady state, is the mapping of discrete macroparticle trajectories from the Lagrangian framework to continuous distributions for the charge and current density, ρ and $\mathbf{j}$, in an Eulerian flow framework. This particle-weighting step is necessary to drive the field solver in the next iteration or time-step. Charge is assigned to the mesh nodes of the field solver according to the distribution of particles relative to the mesh. The resulting distributions are sensitive to the mesh and require a large number of particles: fluctuations in the source terms scale as $n^{-1/2}$ where n is the number of particles per mesh cell. To mitigate these issues, shape functions can be used instead of discrete point particles. However, solution accuracy is usually highly sensitive to the shape function properties (shape, width, etc.), which tend to be heuristic [6–8].

To overcome this challenge, we consider the dual space represented by the macroparticle trajectories, treating the macroparticle trajectories not as discrete particles or particle distributions but as bounds between which source terms can be interpolated via charge conservation. This both reduces the number of macroparticles required to accurately model the system and removes the particle-weighting step entirely. Consequently, we demonstrate an order of magnitude improvement in accuracy for the same number of macroparticles simulated relative to a state-of-the-art solver and for the same accuracy, an 80 to 120 fold speed-up due to the removal of the costly particle-weighting process.

II. DUAL SPACE INTERPOLATION

A. Consequences of the dual interpretation

Let us begin by considering the implications of working with the dual space of the particle trajectories. When considering the trajectories as the centroid of a given distribution, each macroparticle has an associated line current density: $\lambda_q(\mathbf{r}) = \lambda \delta(\mathbf{r} - \mathbf{r}_0)$. The beam macro-quantities, ρ and $\mathbf{j}$ are obtained only through accumulating many such distributions—enough to obtain adequate statistics. When treating the trajectories as bounds, they no longer represent any intrinsic charge or current density, yet neither do they suddenly carry any direct information about the beam macro-quantities ρ and $\mathbf{j}$. Rather, the only information we can truly glean from these traces is the particle flow field along the traces ($\mathbf{v}$ of particles in the beam).

Interpolating the beam current density, $\mathbf{j}$ is thus not as straightforward as assuming an approximate distribution function along the beam and solving for its degrees of freedom using the values for $\mathbf{j}$ along the boundaries (represented by the trajectories), the typical interpolation approach. Rather the challenge lies in determining a

method which uses only the beam velocity $\mathbf{v}$ along the trajectories to interpolate $\mathbf{j}$ and ρ .

B. Applying charge conservation

To resolve this issue further information is required, namely an equation governing the evolution of the beam current over the problem domain. For charged particle beams, one possibility is the charge continuity equation, Eq. (1).

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{j} = 0 \quad (1)$$

Consider a region, Ω , bounded by an emission surface, $\partial\Omega_{\text{in}}$, a set of macroparticle trajectories, $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ and an outlet surface, $\partial\Omega_{\text{out}}$. By applying the divergence theorem and noting that for static problems $\frac{\partial \rho}{\partial t} = 0$, we can reformulate Eq. (1) integrated over Ω as a surface integral over $\partial\Omega$, as given by Eq. (2).

$$\int_{\Omega} \frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{j} dV = \int_{\partial\Omega} \mathbf{j} \cdot \hat{n} dS = 0 \quad (2)$$

Consider more closely the various surfaces comprising $\partial\Omega$. In addition to the inlet and outlet surfaces, for an N -dimension region there are also at least N surfaces connecting the trajectories. Assume the current is conserved between trajectories. This assumption is not only valid for laminar beams, but also for nonlaminar beams in the limit where an adequate number of macroparticles are tracked, a statement we shall revisit shortly. In this case, these connecting surfaces can be taken to be parallel to the current flow and the only surfaces which contribute to the integral in Eq. (2) are the inlet and outlet.

$$\int_{\partial\Omega} \mathbf{j} \cdot \hat{n} dS = \int_{\partial\Omega_{\text{in}}} \mathbf{j} \cdot \hat{n} dS + \int_{\partial\Omega_{\text{out}}} \mathbf{j} \cdot \hat{n} dS = 0 \quad (3)$$

This provides a relationship between the known emission current at the inlet, $I_{\text{in}} = -\int_{\partial\Omega_{\text{in}}} \mathbf{j} \cdot \hat{n} dS$ (the negative sign is due to the direction of $\hat{n}$ being opposite to the direction of current flow on $\partial\Omega_{\text{in}}$) and the current density on the outlet surface.

$$\int_{\partial\Omega_{\text{out}}} \mathbf{j} \cdot \hat{n} dS = \int_{\partial\Omega_{\text{out}}} \rho \mathbf{v} \cdot \hat{n} dS = I_{\text{in}}. \quad (4)$$

Finally, we assume a uniform distribution for ρ over the outlet surface, producing an expression for ρ as a function of I_{in} and the beam velocity over the surface. Note that unlike $\mathbf{j}$, $\mathbf{v}$ can be interpolated easily as it is known along the macroparticle trajectories.

$$\rho = \frac{I_{\text{in}}}{\int_{\partial\Omega_{\text{out}}} \mathbf{v} \cdot \hat{n} dS}. \quad (5)$$

C. Computing an interpolation surface

Equation (5) provides a path forward for approximating ρ on the surface, however there remains the significant question of how to define the outlet surface and subsequently, the choice of an appropriate interpolation of $\mathbf{v}$ over this surface. The outlet should of course contain the mesh node at which to interpolate ρ and should span the trajectories. With these satisfied, the choice of outlet surface could depend on the implementation, though one which most closely validates the assumption of uniform ρ will produce the best results.

In this regard, for our 2D azimuthally symmetric implementation for static problems, the surface is defined by an arc which is perpendicular to the velocity of the beam at the bounding trajectories and thus approximately perpendicular to the flow at all points in-between the trajectories. For electrostatic problems with small energy spread, this is also approximately a velocity isosurface so we make the further assumption that $|\mathbf{v}|$ is uniform. Therefore, our computation of ρ and subsequently, $\mathbf{j}$ is straightforward (here A is the surface area), given by Eq. (6).

$$\rho_{\partial\Omega_{\text{out}}} = \frac{I_{\text{in}}}{A_{\partial\Omega_{\text{out}}} |\mathbf{v}|} \quad (6)$$

$$\mathbf{j}_{\partial\Omega_{\text{out}}} = \rho_{\partial\Omega_{\text{out}}} \mathbf{v} \quad (7)$$

D. Illustration in 2D

Figure 2 illustrates how we have implemented this approach for 2D axisymmetric problems. In the example, ρ and $\mathbf{j}$ are to be calculated on the mesh nodes p and q (for

FIG. 2. Calculating source terms from particle trajectories.

future use driving the finite element field solver) using the initial emission current densities, J_1 in between the axis and trajectory 1 and J_2 between trajectories 1 and 2.

Consider first the point p . We begin by computing an arc passing through p and perpendicular to the trajectories at the point where the arc intersects. We then apply Eq. (6) to compute ρ at p (ρ_p) using the interpolated velocity at point p , $\mathbf{v}_p$.

$$\rho_p = \frac{I_2}{A_{\Omega_{\text{out},p,\{1,2\}}} |\mathbf{v}_p|}$$

$$\mathbf{j}_p = \rho_p \mathbf{v}_p = \frac{I_2}{A_{\Omega_{\text{out},p,\{1,2\}}} \hat{n}_{\text{out},p,\{1,2\}}$$

The current at q is slightly more complicated as the node is in a region where macroparticle trajectories have crossed, hence there are contributions from multiple current regions:

$$\rho_q = \frac{I_1}{A_{\Omega_{\text{out},q,\{0,1\}}} |\mathbf{v}_q| + \frac{I_2}{A_{\Omega_{\text{out},q,\{1,2\}}} |\mathbf{v}_q|}$$

$$\mathbf{j}_q = \frac{I_1}{A_{\Omega_{\text{out},q,\{0,1\}}} \hat{n}_{\text{out},q,\{0,1\}} + \frac{I_2}{A_{\Omega_{\text{out},q,\{1,2\}}} \hat{n}_{\text{out},q,\{1,2\}}$$

The case of crossed trajectories underscores an important point regarding the extension of this method to nonlaminar beams. While at first glance, the assumption of current conservation in between trajectories seems highly restrictive, note that this assumption is *local*, between the bounding trajectories, and not a requirement on the beam itself. In the limit of a sufficient number of trajectories, this local laminar flow condition is adequately approximated. The number of macroparticle trajectories required to achieve this limit will depend on how nonlaminar the beam is but in the worst case does not exceed the number of particles required in a PIC code and tends to be at least an order of magnitude less, as demonstrated shortly. Note that the same point holds true for our assumptions of uniform charge density and velocity across the surface: as the number of trajectories simulated increases, this assumption will be increasingly more accurate.

E. Limitations

There are two important situations which this method, as implemented, does not currently handle. If the beam has any uncorrelated initial energy spread, such that there is no longer a surjective mapping between initial position and velocity, the initial current density “between” these trajectories is undefined in our 2D model. To handle such problems, we would need to extend the interpolation to 4D phase-space (6D for 3D problems). Within the context of radio-frequency source design and electron guns, the electrons are generally emitted from rest and the energy spread of the beam is dominated by the effects of the fields rather than the initial energy spread.

The other challenge we encounter is when particles reverse direction midflight. Once again, extending the interpolation to 4D phase-space should alleviate this issue as the reversed beam occupies a different region in phase space than the forward traveling beam. Particles bouncing off the axis of symmetry and secondary emission are not issues and can be handled in the current implementation.

III. CONNECTION TO LIOUVILLIAN INTERPOLATION

There is a nontrivial connection between our method and the related yet distinct phase space interpolation methods recently applied to cosmological N-body simulations and 1D plasma simulations [9,10]. These techniques create a secondary mesh in phase-space where the mesh nodes are the discrete macro-particle positions at a given time. Liouville's theorem on the conservation of phase space density and volume can then be applied to obtain the source term distribution. However, this distribution must still be mapped back to the original mesh through the convolution of the secondary mesh elements with the original mesh elements.

By directly applying the charge conservation law over the appropriate surface, we eliminate the need for the remeshing or particle weighting step entirely. The source term is interpolated directly at the mesh node and varies continuously throughout the mesh element. For electro- and magneto-statics, a connection can be made between Liouville's theorem and our interpolation method using charge continuity. The connection is not general as the conservation laws arise from different symmetries (time invariance in the former, and gauge invariance in the latter). It arises specifically through the properties of the surface over which the source terms are interpolated.

For a Hamiltonian system, Liouville's theorem is given by Eq. (8), where $f(\mathbf{q}, \mathbf{p})$ is the phase space density and $\mathbf{q}, \mathbf{p}$ are the canonical position and momentum. For electro- and magneto- static fields, $\frac{\partial f}{\partial t} = 0$ and $\nabla_{\mathbf{p}} \cdot (f\dot{\mathbf{p}}) = 0$.

$$\frac{\partial f}{\partial t} + \nabla \cdot (f\mathbf{v}) = \nabla_{\mathbf{q}} \cdot (f\dot{\mathbf{q}}) + \nabla_{\mathbf{p}} \cdot (f\dot{\mathbf{p}}) = \nabla_{\mathbf{q}} \cdot (f\dot{\mathbf{q}}) = 0 \quad (8)$$

Integrating over the volume bounded by particle trajectories in 6D phase space and substituting $\dot{\mathbf{q}} = \frac{\mathbf{p}}{\gamma m}$ (here $\mathbf{P} = \mathbf{p} - e\mathbf{A}$ is the kinematic momentum where $\mathbf{A}$ is the magnetic vector potential), we obtain Eq. (9).

$$\begin{aligned} \int_{\Omega_p} \int_{\Omega_q} \nabla_{\mathbf{q}} \cdot (f\dot{\mathbf{q}}) d\mathbf{q} d\mathbf{p} &= 0 \\ \int_{\Omega_p} \int_{\partial\Omega_q} f\dot{\mathbf{q}} \cdot d\mathbf{S}_q d\mathbf{p} &= 0 \\ \int_{\partial\Omega_q} \int_{\Omega_p} f(\mathbf{q}, \mathbf{p}) \frac{\mathbf{p} - e\mathbf{A}(\mathbf{q})}{\gamma m} \cdot \hat{n} d\mathbf{p} d\mathbf{S}_q &= 0 \end{aligned} \quad (9)$$

In our approach we specify the interpolation surface, $\partial\Omega_q$, and make the assumption that over it, $\mathbf{P}$ is a pre-defined, surjective function of $\mathbf{q}$ (laminar flow). As $\mathbf{A}$ is a function only of $\mathbf{q}$, the constraint on $\mathbf{P}$ maps to the canonical momentum, $\mathbf{p}$. Noting that $\rho(\mathbf{q}) = \int_{\Omega_p} f(\mathbf{q}, \mathbf{p}) d\mathbf{p}$, this implies that on $\partial\Omega_q$, $f(\mathbf{q}, \mathbf{p}) = \delta(\mathbf{p} - \mathbf{p}_s(\mathbf{q}))\rho(\mathbf{q})$ and the integral over momentum space can be evaluated, producing Eq. (12).

$$\int_{\partial\Omega_q} \rho(\mathbf{q}) \frac{\mathbf{p}(\mathbf{q}) - e\mathbf{A}(\mathbf{q})}{\gamma m} \cdot \hat{n} d\mathbf{S}_q = 0 \quad (10)$$

$$\int_{\partial\Omega_q} \rho(\mathbf{q}) \dot{\mathbf{q}} \cdot \hat{n} d\mathbf{S}_q = 0 \quad (11)$$

$$\int_{\partial\Omega_q} \mathbf{j} \cdot \hat{n} d\mathbf{S}_q = 0 \quad (12)$$

This derivation underscores a few points on the applicability and extendability of our interpolation method. By using the charge conservation law, we must make assumptions on the momentum distribution of the beam. As mentioned previously, these assumptions preclude the simulation of beams with uncorellated energy spread. For radio-frequency source design, this is not an issue however extending this interpolation method to problems such as the beam-wave interaction in free electron lasers, for example, would likely require the more general Liouvillian interpolation.

IV. BENCHMARKING AND RESULTS

We have implemented the charge conservation interpolation (CC method) for azimuthally symmetric electrostatic problems and benchmarked against COMSOL, a state-of-the-art commercial software [11]. COMSOL computes the steady-state solution using the same global, iterative process illustrated in Fig. 1 such that any significant performance differences exist only in the differing source term computation.

For the steady state solver, COMSOL treats each macro-particle as representing a number of particles per unit time (all taking the same path). Each of these macroparticles represents a discrete spatial distribution corresponding to $\delta(\mathbf{r} - \mathbf{q}_i)$ and the average charge density per unit time in a given element, k , is obtained by integrating these distributions over the element and averaging, as given by Eq. (13) [12].

$$\frac{d\rho_k}{dt} = \frac{1}{V_k} \sum_{i=1}^N f_{\text{rel},i} \int_{\Omega_k} q_i \delta(\mathbf{r} - \mathbf{q}_i) dV. \quad (13)$$

Ignoring the modifications to account for a steady-state solution, spatially this approach is similar to the particle-weighting technique in a particle-in-cell code. It is

FIG. 3. 2D cut-away view of a beam traversing a cylindrical diode with an applied potential, diverging due to space charge. The outer black trajectories obtained using the CC method are in agreement with the white trajectories from COMSOL.

accumulating discrete distributions such that mesh cells traversed by a particle are assigned space charge, and those which particles miss are assigned no charge. Enough particles are required per mesh element to obtain adequate statistics for the calculation of the volume average in Eq. (13).

We begin with a classical azimuthally symmetric problem: the cylindrical diode. Here, an electron beam is emitted from the inner radius and extends to the outer, diverging as it proceeds due to space charge (see Fig. 3). While no analytical solution exists with which to compare, the convergence properties of the numerical solution provide a measure of its quality. Figure 4 plots the Sobolev zero-norm of the normalized solution difference from one iteration to the next. COMSOL, using the conventional particle weighting method to obtain the source terms, does not converge to high accuracy when few particles are simulated as expected. To maintain the convergence rate over all 10 iterations, at least 50 particles are required. The CC method attains the same convergence rate and solution accuracy using only 2 particles.

Figure 5 compares the converged solution using $N < 150$ particles with the solution using 150 macroparticles. As opposed to Fig. 4, a measure of how close the solution is to self-consistency, here we consider how the number of

FIG. 4. Convergence vs iteration: colored curves correspond to the PIC solution with N particles simulated (given by legend) while black corresponds to the CC method and $N = 2$.

FIG. 5. Convergence of the final solution for the cylindrical diode as a function of number of particles simulated.

particles affects the solution. Two promising features of the CC method are demonstrated. First, the normalized solution error is on the order of 10^{-4} with just two particles (a similar result using COMSOL requires 20 particles). Secondly, the solution depends much less strongly on the number of macroparticles compared to the traditional approach, as expected.

The second example is of a hollow beam traversing a parallel disc capacitor. Figure 6, a plot of the convergence as a function of the number of particles simulated, demonstrates that COMSOL requires roughly 40 particles to obtain the same accuracy as the CC method with two particles. Additionally, for similar mesh resolution using both techniques, it takes 116 times longer to obtain this result due to the additional particle weighting step. Even as the solution converges with increasing number of particles, the CC method still outperforms, obtaining an order of magnitude smaller error for the same number of particles.

Finally, we consider a highly nonlaminar beam, where we expect the CC method to require more particle trajectories to approximate laminar flow between trajectories. Figure 7 shows an azimuthally symmetric electron gun with an over-focused beam at zero current (initial iteration).

FIG. 6. Convergence of the final solution for the hollow beam problem as a function of number of particles simulated.

FIG. 7. The converged particle trajectories using the CC method (black) compared with a solution from COMSOL (white) for a highly nonlaminar beam in an electron gun. The inset depicts the initial beam with no space charge.

FIG. 8. The convergence of a test particle trajectory with number of particles simulated for the overfocused electron gun.

Figure 8 plots the integrated error of a test particle trajectory as a function of the number of particles simulated. We plot this instead of the error in the electrostatic potential, which is highly sensitive to the singularities at the sharp edges and dominated by changes at those points. Once again, the CC method outperforms COMSOL by approximately an order of magnitude in terms of overall accuracy, and is 80 times faster for the same accuracy with similar mesh size. The error in the CC method depends more strongly on the number of macroparticles than in the laminar cases due to the strengthened validity of the local laminar flow assumption.

While these examples are electrostatic, necessary to benchmark against a similar approach and hence isolate the difference to only the source term calculation technique, the theory applies to time harmonic problems as well. Future work will expand this initial implementation accordingly and extend this method to 2D problems with an azimuthal symmetry of order $m > 0$ (dipole, quadrupole modes, etc.).

V. CONCLUSION

Concluding, the CC method bypasses the onerous noise inducing particle weighting step common to previous

numerical solvers for the self-consistent beam-wave interaction for static problems. It is both computationally more efficient (no need to compute shape functions or weighting functions) and produces a solution that is highly accurate and relatively insensitive to macroparticle distribution. This is accomplished through the dual interpretation of the macro-particle trajectories, treating them as bounds on the particle flow rather than the centroid of a physical particle distribution.

Not only does this method provide an order of magnitude improvement in accuracy for the same number of macroparticles as compared with traditional particle weighting schemes, it provides accurate solutions for highly nonlinear problems with only 2–3 macro-particle trajectories. Given the similarity between charge conservation and the more general divergence free condition in incompressible fluid flow, this method could apply in modeling other types of fluid flow.

ACKNOWLEDGMENTS

This project was funded by U.S. Department of Energy under Contract No. DE-AC02-76SF00515.

- [1] A. N. Vlasov *et al.*, in *Proceedings of the Fifth IEEE International Vacuum Electronics Conference, Monterey, CA, USA* (IEEE, Monterey, 2004), p. 314.
- [2] M. Botton, T. M. Antonsen, B. Levush, K. T. Nguyen, and A. N. Vlasov, MAGY: A time-dependent code for simulation of slow and fast microwave sources, *IEEE Trans. Plasma Sci.* **26**, 882 (1998).
- [3] MAGIC Tool Suite, Orbital ATK, Dulles, United States, <https://www.cst.com>.
- [4] CST Studio Suite, CST of America, LLC, Framingham, United States, <https://www.cst.com>.
- [5] D. Grote, The WARP code: Modeling high intensity ion beams, *AIP Conf. Proc.* **749**, 55 (2004).
- [6] D. Tskhakaya, K. Matyash, R. Scheneider, and F. Taccogna, The particle-in-cell method, *Contrib. Plasma Phys.* **47**, 563 (2007).
- [7] J. Xiao, J. Liu, H. Qin, and Z. Yu, A variational multi-symplectic particle-in-cell algorithm with smoothing functions for the Vlasov-Maxwell system, *Phys. Plasmas* **20**, 102517 (2013).
- [8] T. D. Arber *et al.*, Contemporary particle-in-cell approach to laser-plasma modelling, *Plasma Phys. Controlled Fusion* **57**, 113001 (2015).
- [9] T. Abel, O. Hahn, and R. Kaehler, Tracing the dark matter sheet in phase space, *Mon. Not. R. Astron. Soc.* **427**, 61 (2012).
- [10] J. Kates-Harbeck, S. Totorica, J. Zrake, and T. Abel, Simplex-in-cell technique for collisionless plasma simulations, *J. Comput. Phys.* **304**, 231 (2016).
- [11] COMSOL Multiphysics v. 5.3, COMSOL Inc, Stockholm, Sweden, www.comsol.com.
- [12] Particle Tracing Module User's Guide v. 5.3., COMSOL Inc, Stockholm, Sweden.