SLAC-402
UC-414
(A)

A HIGH-BRIGHTNESS THERMIONIC MICROWAVE ELECTRON GUN*

Michael Borland

Stanford Linear Accelerator Center
Stanford University, Stanford, CA 94309

February 1991

Prepared for the Department of Energy
under contract number DE-AC03-76SF00515

*Ph.D. Thesis

Abstract

In a collaborative effort by SSRL, AET Associates, and Varian Associates, a high-brightness microwave electron gun using a thermionic cathode has been designed, built, tested, and installed for use with the SSRL 150 MeV linear accelerator. This thesis discusses the physics behind the design and operation of the gun and associated systems, presenting predictions and experimental tests of the gun's performance.

The microwave gun concept is of increasing interest due to its promise of providing higher-current, lower-emittance electron beams than possible from conventional, DC gun technology. In a DC guns, accelerating gradients are less than 8 MV/m, while those in a microwave gun can exceed 100 MV/m, providing much more rapid initial acceleration, thereby reducing the deleterious effects of space-charge. Microwave guns produce higher momentum beams than DC guns, thus lessening space-charge effects during subsequent beam transport. Typical DC guns produce kinetic energies of 80-400 KeV, compared to 2-3 MeV for the SSRL microwave gun.

"State-of-the-art" microwave gun designs employ laser-driven photocathodes, providing excellent performance but with greater complexity and monetary costs. A thermionic microwave gun with a magnetic bunching system is comparable in cost and complexity to a conventional system, but provides performance that is orders of magnitude better.

Simulations of the SSRL microwave gun predict a normalized RMS emittance at the gun exit of $< 10\ \pi \cdot m_e c \cdot \mu m$ for a beam consisting of approximately 50% of the particles emitted from the gun, and having a momentum spread of ± 10 %. These emittances are for up to $5 \times 10^9 e^-$ per bunch. Chromatic aberrations in the transport line between the gun and linear accelerator (GTL) increase this to typically $< 30\ \pi \cdot m_e \cdot \mu m$.

The SSRL microwave gun was designed to have a longitudinal phase-space suited to magnetic bunch compression. Simulations predict that peak currents of several hundred amperes are achievable.

Acknowledgements

As with any project of the complexity of the SSRL microwave gun and its associated systems, the task of going from the conceptual stage to the point where the system is turned over for operational use involves a great many people. It is my pleasure to acknowledge the contributions of those who participated in the project that I am reporting on in this thesis. With so many people deserving of thanks, there is the danger that I have overlooked someone or failed to fully recognize his contributions. I have done my best to avoid this, and apologize in advance if I fail in this regard.

First and foremost, I wish to thank my thesis advisor Helmut Wiedemann for introducing me to the exciting field of accelerator physics, for providing me with guidance and instruction as I learned about the field, and for giving me the rare opportunity to take on the microwave gun project. I am particularly grateful for Helmut's patience, support, and confidence in me when others thought that more conventional technology and a more experienced person were needed. Virtually all of the work that I report on here has benefited from Helmut's suggestions, ideas, scrutiny, and criticism.

Jim Weaver of SSRL participated in the design of the gun and many components of the gun-to-linac (GTL) system, in addition to designing the RF system for the gun and preinjector. I am convinced that without Jim's tireless attention to detail and his sound design ideas, the gun, the alpha-magnet chamber, and the chopper would not have been as successful as they were. I am also grateful to Jim for converting inches to millimeters when speaking to me and for laughing at most of my jokes.

Michael Baltay of SSRL did most of the mechanical engineering for the GTL and directed the assembly and alignment as well. All of the GTL magnets were designed in collaboration between Mike and myself (with many helpful ideas from Helmut, John

Voss, and others). Mike also helped me to assemble the necessary hardware for the magnetic measurements and helped with alignment of the measurement equipment.

Other SSRL and Injector people made significant contributions to the project: John Voss was the project manager for the Injector, and as such helped to organize the effort and push us toward completion. Clarence Chavis designed and directed the building of the modulators, and he and Jim Haydon frequently came to the rescue when something broke. Harry Morales oversaw the vacuum system work in the GTL and contributed to the design of many components. Paul Golceff's meticulousness when it came to the gun vacuum system and his prominent participation in the assembly of the GTL line, alpha-magnet chamber, and chopper are very much appreciated. Defa Wang assembled much of the equipment used for running the RF system, did cold test measurements on RF system components, and calibrated the gun RF diodes; he also proved tireless in the RF processing of the gun and linac, and in the commissioning of the chopper and linac. Jim Sebek did the TDR of the chopper, tested and refined the toroid design, stayed all night several times during the commissioning of the chopper, and, with Bob Hettel, helped to track down noise problems in many of the diagnostic signals. Ray Ortiz found and packaged the coaxial switches that proved so useful during the commissioning and during experiments, and also assembled the necessary circuits for running the gun filament. Bill Lavender's "test" control program for the GTL magnets ended up being the principle means of controlling the GTL during the commissioning and experiments. Brad Youngman, Tom Sanchez, and Ken Ruble participated in the design of the alpha-magnet chamber, chopper, and other GTL components.

Several of Helmut's other graduate students participated in the gun and GTL commissioning as well. Louis Emery in particular was indispensible during the crucial early gun experiments at SSRL, when many remained to be convinced that the gun could function reliably. In addition, it was very valuable to be able to refer to the LaTeX input for his thesis in preparing mine. James Safranek was responsible for the phosphor screens, TV cameras, and related hardware, and assumed his share of those seemingly endless RF processing shifts.

Substantial credit for the success of the microwave gun goes to Eiji Tanabe of AET Associates and Luther Nelson of Varian Associates. While I performed all of the simulations, evaluated the predicted performance of the gun, and modified gun parameters in order to achieve what SSRL needed, Eiji provided the practical knowledge of what could be done and how it could be done, along with suggestions of what to try when a given idea didn't work out or wasn't easily realizable. Luther came up with the idea of using a toroidal tungsten spring and a metal-coated ceramic annulus around the cathode, eliminating the need for a RF choke and greatly simplifying the design; he also provided technical support in the assembly and testing of the gun.

Other Varian people contributed to the gun project as well. Mike Green supplied the cathodes for the prototype and consulted on many occasions on issues relevant to the cathode. Steve Skellenger did the initial RF processing of the gun and played a major part in many of the early experiments at Varian. Ralph Tol machined the gun cavity and many related components. Bill Leong tuned the gun and performed the bead-drop measurements. Ray McIntrye backed the Varian effort and generously provided time for initial testing in one of Varian's test cells.

Roger Miller of SLAC served as a reviewer and advisor on many aspects of the gun and GTL project, including the gun itself, the instrumentation, the magnets, and the chopper. I am grateful for his advice and for many stimulating conversations.

Radiation safety concerns for the project were addressed by Ed Benson, Nisy Ipe, and Ralph Nelson of SLAC.

On a more personal note, I wish to thank my Mother and Father for their support and understanding throughout my long education, and for being so proud of me.

Finally, thanks to Kim Adams, for everything we share.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The history of accelerator physics is inseparable from the history of the application of particle beams. The production of high-brightness electron beams is one of increasing interest among accelerator physicists precisely because of increasing interest in the applications of such beams. These applications provide the motivation for the research reported on in this thesis.

In this introductory chapter, I will indicate the nature of the these applications to give the reader some appreciation of the motivation for research into microwave electron guns. Prior to this, I review fundamental concepts—such as phase-space, Liouville's theorem, and emittance—that are necessary to the understanding of high brightness. I also review issues relevant to the production of high-brightness photon beams from synchrotron radiation emitted by high-energy electrons, including a discussion of coherent radiation and free electron lasers.

Microwave, or "RF", electron guns[1, 2] are a relatively recent development in the production of high-brightness beams. Their general characteristics and brief history are reviewed in the third section of this chapter.

I end this chapter with an overview of the main body of the thesis.

1.1 Review of Fundamental Concepts

In order to understand the meaning of "high-brightness", it is necessary to first understand several prior concepts. The first of these is the concept of emittance. The emittance of a beam is related to the phase-space volume occupied by the beam, or some fraction of it. To properly define the emittance requires a discussion of phase-space and Liouville's theorem.

1.1.1 Phase-Space and Liouville's Theorem

Phase-space refers to the six-dimensional space in which the spatial position and dynamical properties of any particle are defined. For example, the coordinates of phase-space may be taken to be the Cartesian coordinates (q_1, q_2, q_3) and the corresponding momenta (p_1, p_2, p_3). In classical mechanics, these six coordinates uniquely and completely describe the state of any particle. (More will be said about the choice of coordinates below.)

Suppose that at some time, t=0 say, an arbitrary closed surface S in phase-space is chosen, and that it bounds a volume V. Allow S to evolve in time as if it were anchored to imaginary particles that lie on S at t=0. That is, consider S(t) to be defined as the surface occupied by those particles. Liouville's theorem[3] states that the phase-space volume, V(t), bounded by S(t) is constant, provided that only conservative forces act on the particles.

Proof of Liouville's Theorem

Perhaps the simplest and most intuitively appealing proof of Liouville's theorem is that given by Weiss[4]. For simplicity in notation, consider a two-dimensional phase-space, with coordinates (q, p), so that V(t) is the area bounded by a closed curve S(t).

In order to calculate the total time derivative of V(t), one need only look at the motion of the boundary, which again may be thought of as determined by the motion of hypothetical boundary particles that are defined by the choice of S(t=0). Let df represent the outward vector for an infinitesimal segment of the boundary. A motion of this boundary element by $(\Delta q, \Delta p)$ will increase the area bounded by S by

$$df_q \Delta q + df_p \Delta p. \tag{1.1}$$

Dividing by Δt and integrating over S, one sees that

$$\frac{dV}{dt} = \int_S (\dot{q}, \dot{p}) \cdot \mathbf{df}. \tag{1.2}$$

Use of the Divergence Theorem allows one to convert the surface integral into an integral over the region of phase-space bounded by S:

$$\frac{dV}{dt} = \int_V \nabla \cdot (\dot{q}, \dot{p})\, dq\, dp = \int_V (\partial_q \dot{q} + \partial_p \dot{p})\, dq\, dp. \tag{1.3}$$

For a conservative system[3], Hamilton's equations,

$$\dot{q} = \partial_p H \tag{1.4}$$

$$\dot{p} = -\partial_q H, \tag{1.5}$$

are applicable. Since the boundary S moves as if anchored to particles, these equations for particle motion specify the motion of the boundary, and may meaningfully be used in equation (1.3), yielding

$$\frac{dV}{dt} = \int_V (\partial_q \partial_p H - \partial_p \partial_q H)\, dq\, dp = 0. \tag{1.6}$$

Hence, the phase-space area bounded by S(t) is constant as the system evolves.

Implications of Liouville's Theorem

Most statements of Liouville's theorem[5, 6] make use of the particle distribution function, $\Psi(\mathbf{q}, \mathbf{p}, t)$, which gives the density of particles in phase space. Since the paths of particles cannot cross in phase-space, any particle inside S at t=0 will remain inside S. Thus, if one accepts that the volume bounded by any surface S(t) is constant, then it follows that the integral of Ψ over that volume is also constant, since this integral gives the number of particles in the volume and since particles cannot cross S. Since the volume bounded by S(t) and the number of particles inside S(t) are both constant, the average density of particles inside S(t) must also be constant.

To deduce a further result, assume that S(t) bounds an *infinitesimal* volume $dV = d\mathbf{q}d\mathbf{p}$, so that one can ignore the variation of Ψ over the volume. Then the statement that the average of Ψ over V(t) is constant implies that the value of Ψ at the center of the volume (or for any other point inside it) is constant. That is, if one chooses any point $(\mathbf{q}, \mathbf{p})$ in phase space at t=0, then as one travels with a particle starting at that point and moving under the influence of conservative forces, the value of Ψ evaluated at the position of the particle is constant. Writing this in mathematical form, one obtains

$$\frac{d\Psi}{dt} = \frac{\partial\Psi}{\partial t} + \frac{d\mathbf{q}}{dt} \cdot \frac{\partial\Psi}{\partial\mathbf{q}} + \frac{d\mathbf{p}}{dt} \cdot \frac{\partial\Psi}{\partial\mathbf{p}} = 0, \tag{1.7}$$

which is the mathematical result commonly referred to as Liouville's theorem[5, 6].

Liouville's Theorem and Real Beams

In accelerator physics, one deals with the properties of ensembles of large numbers of particles, referred to as "beams" or "bunches". The real phase-space distribution of a bunch is the sum of many delta-functions:

$$\Psi_o(\mathbf{q}, \mathbf{p}, t) = \sum_{i=1}^{N} \delta(\mathbf{q} - \mathbf{q}^{(i)}(t))\delta(\mathbf{p} - \mathbf{p}^{(i)}(t)), \tag{1.8}$$

where N is the number of particles, and $(\mathbf{q}^{(i)}, \mathbf{p}^{(i)})$ are the phase-space coordinates of the i^{th} particle.

What does Liouville's theorem tell us about the evolution of such a bunch? Contrary to the impression given by some discussions[5], there is nothing in the derivation[6] of equation (1.7) that invalidates it for a distribution of this type. In my discussion, I have taken care to refer to "imaginary" particles in defining boundary surfaces, in order to emphasize that Liouville's theorem is not dependent for its validity on having an infinite number of particles or a smooth continuous distribution of particles. Liouville's theorem stated in terms of the constancy of the volume inside a closed, evolving surface in phase-space is clearly a result that is valid regardless of what sort of actual beam distribution one has. This applies just as well to Liouville's theorem as stated in equation (1.7).

One caveat needs to be issue in this regard[7]. The derivation of Liouville's theorem implicitly assumed that the Hamiltonian was macroscopic in nature, and in particular that particle-particle forces could be included in the Hamiltonian in a way that did not refer to the individual coordinates of particles. If particle-particle collisions are considered, then the real Hamiltonian is a function of 6N variables, where N is the number of particles. In this case, Liouville's theorem is strictly valid only in 6N dimensional phase-space. In 6 dimensional phase-space, particle-particle collisions will increase the phase-volume occupied by the beam and *appear* to violate Liouville's theorem. For practical purposes, the phase-volume in 6 dimensional phase-space is what matters, so this can be an important issue. For sparsely populated beams, the issue becomes even more relevant, since then the particle fields cannot be smoothed into a macroscopic field.

Even given this conclusion, a distribution as defined by equation (1.8) is unwieldy and contains more information than is needed or useful. In the limit of a very large number of particles, Ψ_o can generally (though not always) be approximated by a smooth, continuous function of $\mathbf{q}$ and $\mathbf{p}$. The applicability of a smooth distribution depends on practical considerations of how well one wants the smooth distribution to match the actual distribution. Since real bunches are always confined to a limited volume in phase-space, a practical way to gauge whether a bunch is well-approximated by a smoothed distribution is to ask whether an arbitrary phase-space volume inside the bunch that is small compared to the total phase-space volume contains a number of particles that is much larger than 1.

1.1.2 Beam Emittance

I mentioned above that the emittance is related to the volume occupied by a bunch, or some part of it, in phase-space. The intervening discussion of Liouville's theorem indicates why the emittance is an important concept for accelerator physics. Ignoring dissipative effects such as synchrotron radiation[8], all of the forces in an accelerator are conservative. Hence, Liouville's theorem is applicable, and potentially provides a means of describing a bunch in terms of a conserved property that applies to the whole bunch, rather than in terms of the coordinates of the individual particles. Indeed, the standard analysis of beam evolution in terms of the Twiss parameters[8, 9] makes use of the emittance in order to simplify the computation of bunch properties along an accelerator.

Commonly-Employed Simplifications

In this analysis, certain simplifications are commonly employed. First, instead of dealing with six-dimensional phase-space, it is common to deal with the 2-dimensional projections of the phase-space coordinates into the (q_1, p_1), (q_2, p_2), and (q_3, p_3) planes. In the absence of coupling between these planes, Liouville's theorem is valid for each of the 2-dimensional phase-spaces.

In accelerator physics, it is common to have a bunch that travels largely in one direction, i.e., as a well-collimated beam. Hence, instead of a Cartesian coordinate system, one employs a curvilinear coordinate system which follows the path of an ideal, central particle[10]. At each point along this path, one defines a locally Cartesian system in such a way that one axis (z) lies along the direction of motion of the central trajectory, while the other two axes (x and y), are perpendicular to the central trajectory, so that (x, y, z) is a right-handed system. Corresponding to each of these spatial coordinates is a momentum, thus completing the six-dimensional phase-space.

Discussion of the emittance per se is usually confined to the transverse planes, (x, p_x) and (y, p_y). For the longitudinal plane, (z, p_z), it is more common to speak of the bunch length and momentum spread without defining an emittance. Hence, I will confine myself to the transverse planes in what follows. I will also write the equations only for the (x, p_x) plane, though they are equally valid for the (y, p_y) plane.

Another simplification commonly used pertains to the method used for computing the area occupied by the bunch in each phase-plane. A seemingly straight-forward definition of the emittance would be: the smallest contiguous phase-space area containing a specified fraction, say 90%, of the particles. While this would give an accurate measure of the phase-space area occupied by the bunch, it is difficult to use in practice, and does not lend itself to analytical treatments. (In addition, it may not accurately characterize the effective phase-space area occupied by a bunch, due to filamentation and non-linear correlations in phase-space.)

The RMS Emittances

For this reason, and for others to be seen presently, the most commonly used means of measuring the phase-space area occupied by a bunch is the "normalized RMS emittance", defined as

$$\varepsilon_n = \pi m_e c \sqrt{\langle x^2\rangle\langle p_x^2\rangle - \langle p_x x\rangle^2}, \tag{1.9}$$

where angle brackets indicate averages of the bracketed quantities over the entire bunch. The factor of $m_e c$ indicates that the momentum is dimensionless, i.e., $p_x = \beta_x\gamma$. The reason for the factor of π will be seen presently. Both of these factors are absorbed in the units of ε_n, which in the present work are $\pi \cdot m_e c \cdot \mu m$.

ε_n is referred to as the normalized emittance to distinguish it from the "geometric RMS emittance", defined as

$$\varepsilon = \pi \sqrt{\langle x^2\rangle\langle x'^2\rangle - \langle x' x\rangle^2}, \tag{1.10}$$

where $x' = p_x/p_z = \beta_x/\beta_z$ is the slope of a particle's trajectory. The units of ε are $\pi \cdot \mu m$ or $\pi \cdot mm \cdot mrad$. If p_z is nearly the same for all particles, then

$$\varepsilon_n \approx \langle p_z\rangle m_e c \varepsilon. \tag{1.11}$$

In most applications, the geometric emittance is a more important quantity, since the divergence is what is relevant to the optical properties of the bunch or any radiation produced by it. (I will return to this issue later in this chapter.) Note, however, that if the bunch is accelerated, so that $\langle p_z\rangle$ is increased, the normalized emittance will be unchanged while the geometric emittance will decrease. Hence, the "phase-space" formed by (x, x') is really not a phase-space in the strict sense (it is sometimes referred to as "trace-space", instead[11]). p_x is the momentum conjugate to x, while x' is a ratio of two momenta. Because of these considerations, the normalized emittance is most relevant to the comparison of different accelerators that produce bunches of different longitudinal momenta.

Some authors[12, 13, 11, 14] prefer to define the RMS emittance with an additional factor of 4, in order to obtain a measure of the phase-space area occupied by a larger (though not necessarily well-defined) fraction of the beam. This definition is also put forward because the emittance so defined is equal to the total phase-space area for a uniformly populated ellipse in (x, p_x) space (a projection of the "K-V" distribution[7, 14]). I have used the definition given in equation (1.9) principly because it is the definition used in the electron storage ring community.

For electron storage rings[8], a gaussian phase-space distribution is found to approximate the actual phase-space to a high degree. Such a phase-space distribution may be specified as

$$\Psi(x, p_x) = \frac{m_e c}{2\varepsilon_n} \exp\left(-\frac{\pi^2 m_e^2 c^2 \langle x^2\rangle\langle p_x^2\rangle}{2\varepsilon_n^2} \left\{ \frac{x^2}{\langle x^2\rangle} - 2\frac{\langle x p_x\rangle}{\langle x^2\rangle\langle p_x^2\rangle} x p_x + \frac{p_x^2}{\langle p_x^2\rangle} \right\} \right) \tag{1.12}$$

where the normalization is such that

$$\int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dp_x \Psi(x, p_x) = 1. \tag{1.13}$$

The interested reader may verify that if this distribution used to compute the right-hand side of equation (1.9), then the parameter ε_n in equation (1.12) is indeed the normalized RMS emittance.

It is now possible to explore the connection between the RMS emittance and the area occupied by the bunch in phase space. To simplify the analysis, let $\langle x p_x\rangle = 0$, so that $\varepsilon_n = \pi m_e c\sqrt{\langle x^2\rangle\langle p_x^2\rangle}$. (This simple equation for the normalized RMS emittance for an uncoupled phase-space distribution is one of the appeals of this definition of the emittance.) The distribution in equation (1.12) becomes

$$\Psi(x, p_x) = \frac{1}{2\pi\sqrt{\langle x^2\rangle\langle p_x^2\rangle}} \exp\left(-\frac{1}{2} \left\{ \frac{x^2}{\langle x^2\rangle} + \frac{p_x^2}{\langle p_x^2\rangle} \right\} \right). \tag{1.14}$$

Consider an ellipse in (x, p_x) space defined by

$$\frac{x^2}{\langle x^2\rangle} + \frac{p_x^2}{\langle p_x^2\rangle} = K^2, \tag{1.15}$$

where K is a dimensionless constant. The area of this ellipse is

$$A(K) = \pi m_e c K^2 \sqrt{\langle x^2\rangle\langle p_x^2\rangle}, \tag{1.16}$$

where I have included the "units" ($m_e c$) of p_x. For K=1, the area is equal to the normalized RMS emittance, which is the motivation for including the factor of π in the definition of the emittance (equations (1.9) and (1.10)). The fraction of the bunch within this ellipse is readily computed:

$$\mathrm{F(K)} = \int_0^{\mathrm{K}} \mathrm{k\,dk} \exp\left(-\frac{\mathrm{k}^2}{2}\right) = 1 - \exp\left(-\frac{\mathrm{K}^2}{2}\right). \tag{1.17}$$

For K=1, the ellipse has an area equal to the normalized RMS emittance, and contains 39.35% of the particles. The maximum x coordinate of the ellipse is the RMS value of x, while the maximum p_x coordinate of the ellipse is the RMS value of p_x. For K=2, the ellipse has an area equal to four times the normalized RMS emittance, and contains 86.47% of the particles.

Virtues and Pitfalls of RMS Emittances

One virtue of the geometric and normalized RMS emittances is that both are constant for propagation through a system with no acceleration and where all transverse forces are linear in x (for the normalized RMS emittance, this is true only if the longitudinal momentum spread is zero[11]). For a bunch with no longitudinal momentum spread, the normalized RMS emittance is constant for a linear system, even with acceleration. Because of this, the degradation of the RMS emittances in a beamline is an indication of the severity of non-linear effects and cross-plane coupling in the beamline. This is discussed in more detail in Chapter 5.

In Chapter 5, I also discuss how the geometric RMS emittance can be measured experimentally. In the absence of noise and longitudinal momentum spread, measurements of RMS beam-sizes can be used to *exactly* measure the RMS geometric emittance, which is yet another of the appeals of the using the RMS definition of the emittance.

One problem with the RMS emittances is that they do not measure the actual phase-space area occupied by a fixed fraction of the bunch for an arbitrary phase-space distribution. In fact, it is easy to construct phase-space distributions occupying zero area while having non-zero RMS emittances. One solution to this problem is to use higher-moment descriptions of the beam phase-space, in order to correct the area estimate for higher-order correlations, thus producing an emittance that corresponds more closely to the actual phase-space area occupied by the beam[15]. One can also use such analysis to evaluate and potentially eliminate higher-order phase-space correlations (which are caused by non-linear forces), thus decreasing the RMS emittance.

In the final analysis, one must realize that for non-gaussian beams the RMS emittances are simply convenient but potentially crude estimates of the actual phase-space area occupied by a bunch. In fact, there is simply no way to accurately describe a complicated phase-space distribution with a single number (like the RMS emittance) that will be appropriate to all cases and relevant in all applications. Whenever a precise description of the bunch phase-space is required, it is necessary to specify the distribution itself, either in some functional form or in terms of a representative sample of particles from the distribution. From the standpoint of gun and accelerator simulations, the later method is indeed the method used to specify and compute the evolution of the bunch. That is, in gun and accelerator simulations, one frequently simulates the bunch by a number of "macro-particles" which are representative of the actual distribution. This technique is used extensively in the present work.

If an accurate description of a complicated beam phase-space is to be provided by the accelerator physicist to those interested in applications of the beam, then that description may need to go beyond the RMS emittances and deal directly with the distribution. In many cases, particularly for guns and linear accelerators, RMS and other averaged beam properties must be recognized as approximate characterizations of the phase-space, suitable for approximate calculations only.

1.1.3 Beam Brightness

Bunch Length and Peak Current

The transverse emittances of a bunch, or its transverse phase-space distribution, do not fully characterize the bunch and its usefulness for applications. Missing is information about the number of particles in the bunch and their distribution in longitudinal phase-space. As noted above, the longitudinal phase-space of a bunch is most often characterized in terms of the momentum spread and the bunch-length. Bunch length can be specified either as the literal, spatial length of the bunch in the longitudinal dimension, or as the time-length of bunch, i.e., the time it takes for the particles in the bunch to pass by a fixed location along the beamline.

Some authors[14] define the bunch length as an RMS value about the centroid. For present purposes, I prefer to specify the bunch length as the time-interval occupied by the central 90% of the particles in the bunch. The reason is that the RMS bunch length and the 100% bunch length are sensitive to straggling particles that comprise a small fraction of the bunch and hence are unimportant. As a result, these are unreliable and potentially misleading measures of the bunch length.

Having defined the bunch length, one can go on to define the average current during the bunch, referred to as the "peak current"

$$I_{peak} = \frac{Q}{\delta t}, \tag{1.18}$$

where δt is the 90% bunch length and Q is the amount of charge in the interval δt. More specifically, $\delta t = t_2 - t_1$, where

$$\int_{t_m - t_1}^{t_m} \Psi(t)dt = \int_{t_m}^{t_m + t_2} \Psi(t)dt = 0.45. \tag{1.19}$$

t_m is the "median" time, that is, the time at which half the particles in the bunch, specified by the temporal distribution $\Psi(t)$, have passed by some specified point in the beamline. (Note that $\Psi(t)$ is normalized to unit area.) While this is not the only way one might define the 90% bunch length, it has the advantage of being readily computed and robust (in the sense of being insensitive to outlying particles or noisy distributions).

In contrast to the peak current, the "average current" is

$$I_{ave} = \frac{Q_b}{T_b} \tag{1.20}$$

where Q_b is the total charge per bunch and T_b is the time between bunches (equal to the RF frequency for a RF gun). In general, the average current is much less than the peak current, since the distance between bunches is much greater than the bunch length.

Brightness

The normalized brightness is the proportional to average current density in phase-space during the bunch, defined[16] in terms of the peak current and emittance as

$$B_n = \frac{2I_{peak}}{\varepsilon_{x,n}\varepsilon_{y,n}}(\pi m_e c)^2, \tag{1.21}$$

where $\varepsilon_{q,n}$ is the normalized RMS emittance for the q-plane. The units of the normalized brightness that I will use are $A/mm^2/mrad^2$.

The motivation for this definition[7] is that the volume of an uncoupled ellipse in four-dimensional phase space is $V_4 = \pi^2/2abcd$, where a, b, c, and d are the semi-major axes of the ellipse in each of the dimensions. Hence, for a uniformly filled ellipse (i.e., a K-V distribution), $V_4 = \pi^2/2x_{max}p_{x,max}y_{max}p_{y,max}$, so that the current density in phase-space is

$$\frac{I}{V_4} = \frac{2I}{\pi^2 x_{max}p_{x,max}y_{max}p_{y,max}}, \tag{1.22}$$

which corresponds to equation (1.21 except for the factor of π. (In retrospect, it would have been preferable to leave the factor of π^2 out of equation (1.21).)

To increase the brightness, one needs to increase the peak current, decrease the emittance, or both. To increase the peak current, a process known as "bunching" or "bunch compression" is often employed, which involves increasing the peak current by compressing the bunch into a shorter time-length. (This is discussed in Chapter 4.) Increasing the emittance is often a matter of mitigating emittance-diluting effects, such as non-linear fields. More will be said about about these issues in Section 1.3.

Having reviewed these concepts, many of which will be used throughout this work, I now proceed with a discussion of applications of high-brightness electron beams.

1.2 Applications of High-Brightness Beams

While the problem of producing high-brightness electron beams is a challenging one for the accelerator physicist, these beams are not pursued as ends in themselves. Among the applications that require such beams are linear colliders[17], wake-field and other two-beam acceleration schemes[18, 19], and radiation-producing devices, such as Free Electron Lasers (FELs)[20, 21]. Of these applications, I will discuss only those related to the production of radiation.

1.2.1 Synchrotron Radiation

Synchrotron radiation is the radiation emitted when a charged particle (usually an electron or positron) is subjected to accelerating forces imposed by external magnetic or electric fields. In this section, I will review the properties of synchrotron radiation, taking results from some of the many excellent discussions that appear in the literature[22, 23, 24, 25]

Bending-Magnet Radiation

The simplest way to produce synchrotron radiation is to send a high-energy electron beam through a uniform magnetic field transverse to the direction of travel of the beam. In such a circumstance, the beam undergoes acceleration perpendicular to its direction of motion, which bends its path into an arc of a circle of some fixed radius, ρ. This is commonly referred to as "bending-magnet radiation", since it is produced by the magnets used to "bend" the central trajectory of an electron beam (often into a closed path, as in an electron storage ring[8]). This radiation is characterized by the "critical frequency",

$$\omega_c = 2\pi f_c = \frac{3c\gamma^3}{2\rho}, \tag{1.23}$$

where γ is the electron energy in units of the rest mass. The spectral distribution of bending magnet radiation is such that half the radiated power is below the critical frequency and half above. In addition, the power spectrum of the radiation observed in the bend plane is peaked at $\omega \approx 0.83\omega_c$. The wavelength corresponding to the critical frequency is

$$\lambda_c = \frac{4\pi\rho}{\gamma^3}. \tag{1.24}$$

It is well known that the radiation from a relativistic particle undergoing instantaneously circular motion is emitted primarily in the forward direction. If one defines the horizontal plane to be the plane in which the magnet bends the beam trajectory, then the radiation emitted from a bending magnet is spread over a horizontal angle essentially equal to the bend angle, assuming that the bend angle is large compared to $1/\gamma$. In the vertical plane, however, the "opening angle" of the radiation is much narrower. For $\omega = \omega_c$, the RMS vertical opening angle of the radiation power distribution is $\psi \approx 0.57/\gamma$.

The instantaneous total synchrotron radiation power for a single electron is

$$\mathrm{P_{se}} = 4.611 \times 10^{-20}\mathrm{Watt} \cdot \mathrm{meter}^2 \times \frac{\beta^4\gamma^4}{\rho^2}, \tag{1.25}$$

Hence, if a beam consisting of bunches of $\mathrm{N_e}$ electrons with a repetition rate of f, passes through a bending magnet with bending radius ρ that bends through an angle θ, the average radiation power is

$$\begin{aligned} \langle \mathrm{P} \rangle &= \mathrm{P_{se}}\frac{\theta\rho}{\beta c}\mathrm{N_e f} && (1.26) \\ &= 1.538 \times 10^{-28}\mathrm{Joule} \cdot \mathrm{meter} \times \frac{\beta^3\gamma^4}{\rho}\theta\mathrm{N_e f} && (1.27) \end{aligned}$$

As an example, the SSRL pre-injector linac[26] delivers a 120 MeV beam ($\gamma \approx 235$), which is subsequently deflected by 0.72 rad by a bending magnet with $\rho \approx 0.6$m, so that $\lambda_c \approx 0.6\mu$m, which is in the visible part of the spectrum. Typically 2×10^9 electrons are accelerated per pulse, with 10 pulses per second, so that the average radiation power is a mere 11 nW. However, as will be seen in Chapter 4, each electron bunch has a length of order 1 ps, so that the peak radiation power is of order 1 kW. The "natural" RMS opening angle of the radiation (i.e., ignoring beam emittance effects, which are discussed below) is 2.4 milli-radians.

For a 120 MeV beam with 2×10^9 electrons per bunch in approximately 1 ps bunches, the peak instantaneous beam power is of order 40 GW. This is seven orders of magnitude greater than the peak instantaneous radiation power produced in the example. Clearly, a more efficient way of converting electron beam power into radiation power is desirable. From equation (1.27), it would appear that the most straight-forward ways to do this are to increase the field strength (i.e., decreasing ρ) and increase the path length, $l = \rho\theta$, in the bending magnet (these are not, of course, independent quantities).

Undulators and Wigglers

The power from a bending magnet is fanned out in the bend plane over an angle equal to the bending angle. Hence, the radiation power density per unit solid angle does not increase if one has a longer bending magnet with the same bending radius (which has the same fields but a greater bending angle). However, by using a series of bending magnets of equal but opposite bending angles, one *can* increase the radiation power per unit solid angle. Such a device is called a "wiggler" or "undulator" magnet[25]. By reducing the fanning-out of the radiation caused by the deflection of the beam path, these devices not only allow one to extract more power from a beam, but also to concentrate that power into a narrower solid angle, thus increasing the brightness of the radiation.

The dominant field component in such magnets is a transverse field described by

$$B_y = B_o \cos(2\pi z/\lambda_w), \tag{1.28}$$

where λ is the periodic length of the magnet. (As in the last section, z is the longitudinal coordinate, while x and y are the horizontal and vertical coordinates, respectively.) If B_o is not too large, the particle trajectory in the y=0 plane is sinusoidal, i.e.,

$$x = a \cdot \cos(2\pi z/\lambda_w), \tag{1.29}$$

where

$$a = \frac{eB_o\lambda_w^2}{4\pi^2 m_e c\gamma}. \tag{1.30}$$

The maximum slope of x(z) is thus

$$x'_{max} = a\frac{2\pi}{\lambda_w} \equiv \frac{K}{\gamma}, \tag{1.31}$$

where K is the usual undulator strength parameter. The distinction between wigglers and undulators is that for wigglers, $K \gg 1$, whereas for undulators K is less than or of order 1. That is, an undulator is a wiggler that causes only very slight motion of the beam, so slight that the natural opening angle of the radiation is not greatly degraded in either plane due to the fanning-out effect that occurs in ordinary bending magnets.

From equation (1.25), one sees that the instantaneous radiation power is proportional to $1/\rho^2$. Since $1/\rho \sim B_y$, the highest instantaneous power is radiated when the electron is at the crests of its sinusoidal trajectory, where the acceleration is greatest. For $K \gg 1$, x'_{max} is much greater than the natural opening angle of the radiation at the crest, which means that a distant, on-axis observer principly sees radiation emitted from near the crest of the oscillations (where the electron is traveling with $x' < \gamma$), which enhances the dominance of the radiation emitted at the crest. Hence, for $K \gg 1$, it is plausible that the radiation seen by a distant, on-axis observer will take the form of a series of pulses, one emitted from each crest in the electron's oscillations. The frequency spectrum of this radiation is dominated by the instantaneous spectrum at the crests, and one can show that the spectrum from a wiggler is indeed very much like that from a bending magnet with field B_o.

If $K \approx 1$, however, the same observer will receive radiation from a significant portion of the electron's oscillation, and one finds the the spectrum of this radiation is related to the frequency of the electron's oscillatory motion. In the average rest-frame of the electron, the electron executes transverse oscillations with frequency

$$f_e \approx \frac{\gamma\beta c}{\lambda_u}, \tag{1.32}$$

where the factor of γ is due to the Lorentz-contraction of the undulator period as seen in the moving frame. In its rest frame, then, the electron emits dipole radiation at this frequency. In the laboratory frame, this frequency is Doppler shifted, so that a spectrum of frequencies is produced:

$$f(\theta) = \frac{c}{\lambda_u}\frac{2\gamma^2}{1 + K^2/2 + \gamma^2\theta^2}, \tag{1.33}$$

where θ is the angle in the x-z plane relative to the axis of the undulator. For

$\theta < 1/\gamma$, i.e., near the axis, the wavelength of the radiation is simply

$$\lambda_1 = \lambda_u \frac{1 + K^2/2}{2\gamma^2}, \tag{1.34}$$

so that for a high-energy electron (i.e., $\gamma^2 \gg 1$), the wavelength of the radiation will be much shorter than the undulator period. For very weak undulators, the radiation spectrum is dominated by this "first-harmonic" radiation. For $K \geq 1$, the electron motion in the rest frame contains significant frequency components other than the fundamental oscillation, which generate additional frequencies in the rest frame at even and odd harmonics of the fundamental. The odd harmonics are a result of transverse oscillations, while the even harmonics are a result of longitudinal oscillations relative to the average rest frame. The odd harmonics of the motion produce spectral peaks at frequencies that are, of course, odd multiples of f_1. The odd harmonics are of the greatest interest, since the radiation in this case is confined to a narrow forward cone with an RMS divergence of

$$\sigma'_{r,k} = \frac{1}{\gamma}\sqrt{\frac{1 + K^2/2}{2nN}} = \sqrt{\frac{\lambda_1}{nL}} = \sqrt{\frac{\lambda_n}{L}} \tag{1.35}$$

where N is the number of undulator periods, L is the total undulator length, n is the (odd) harmonic number, and $\lambda_n = \lambda_1/n$ is the wavelength of the n^{th} harmonic. The subscript "r" is used to emphasize that this angular divergence is an intrinsic property of the radiation, separate from the electron beam divergence.

This result assumes that the radiation from the undulator may be approximated as coming from a source at the longitudinal center of the undulator and that the light is observed from a distance that is large compared to the length of the undulator[27]. If this assumption is made, then for a zero-emittance electron beam (e.g., a single electron), the radiation produced in an undulator has an angular divergence given by equation (1.35), as well as an apparent source size (due to the length of the undulator), given by

$$\sigma_{r,k} = \frac{1}{4\pi}\sqrt{\lambda_n L} \tag{1.36}$$

Hence, the geometric RMS emittance of the photon beam emitted by a single electron passing through an undulator is

$$\varepsilon_{r,k} = \pi \sigma_{r,k} \sigma'_{r,k} = \frac{\lambda_n}{4}. \tag{1.37}$$

For an electron beam with a gaussian transverse distribution, characterized by a spatial parameter σ_e and an angular parameter σ'_e, the effective photon beam parameters are obtained by adding in quadrature with the single-electron parameters[23, 27, 28] :

$$\begin{aligned} \sigma_s^2 &= \sigma_e^2 + \sigma_r^2 \\ \sigma_s'^2 &= \sigma_e'^2 + \sigma_r'^2 \end{aligned}$$

Thus, the emittance of weak undulator radiation for a non-zero emittance electron beam is

$$\varepsilon_s = \pi\sqrt{\sigma_e^2 + \sigma_r^2}\sqrt{\sigma_e'^2 + \sigma_r'^2}. \tag{1.38}$$

From this, one concludes that the emittance of the photon beam is maintained at its minimum value if

$$\sigma_e^2 \ll \sigma_r^2 \tag{1.39}$$

and

$$\sigma_e'^2 \ll \sigma_r'^2. \tag{1.40}$$

This implies that

$$\epsilon_e \ll \epsilon_r. \tag{1.41}$$

(Note that the latter condition is necessary but not sufficient to fulfill equations (1.39) and (1.40).)

This result two important implications. First, if one desires a low-emittance photon beam, then the best one can do is to supply an electron beam with emittance significantly less than the wavelength being produced, and with $\sigma_e/\sigma'_e = \sigma_r/\sigma'_r = 4\pi L$ (where L is, again, the length of the undulator). Second, as will be seen presently, if one maintains the intrinsic photon beam emittance, then (if other conditions are also satisfied) the radiation will be spatially coherent. This is path to Free Electron Lasers.

1.2.2 Coherent Radiation

In order to discuss the possibility that undulator radiation might be coherent, it is helpful to review the meaning of coherence[29, 30]. There are two varieties of coherence, spatial coherence and temporal (or longitudinal) coherence.

Spatial Coherence of Transversely Extended Sources

Spatial coherence refers to the constancy of the phase across a wave-front of light. In particular, spatially coherent light will form interference fringes when used in a Young's two-slit experiment. Figure 1.1 shows an experiment in which an on-axis point source is used to illuminate two slits, spaced by 2A, a distance D from the source. Suppose further that the light from these slits falls on a screen at a distance S from the slits. If the width of the individual slits is small compared to D, L, and A, and if A is small compared to D and L, then interference fringes are formed on the screen, spaced by

$$\Delta x_f = \frac{\lambda S}{2A}. \tag{1.42}$$

If the point source is moved transversely (i.e., to an off-axis position), then the center of the fringe pattern moves as well, to

$$x_c = -\frac{x_s S}{D}, \tag{1.43}$$

where x_s is the distance of the source from the axis.

If one had two equal-strength point sources that individually produced fringe patterns offset by $\Delta x_f/2$ relative to one another, then the combined irradiance would be flat—i.e., no fringes would be seen. (One must, of course, add the electric fields and not the irradiances to see this.) This occurs when the two point sources are off axis at

$$x = \pm\frac{\lambda D}{8A}. \tag{1.44}$$

Hence, even though these two sources might be radiating in phase, the combined source does not produce interference fringes and is said to be spatially incoherent. As the point sources are moved toward one another, the fringes gradually reappear. To obtain a high degree of spatial coherence for two point sources offset by $\pm x_s$, one wants the centers of the two fringe patterns to be spaced by, say, less than $\Delta x_f/4$:

$$\frac{2x_s S}{D} < \frac{1}{4}\Delta x_f = \frac{\lambda S}{8A}. \tag{1.45}$$

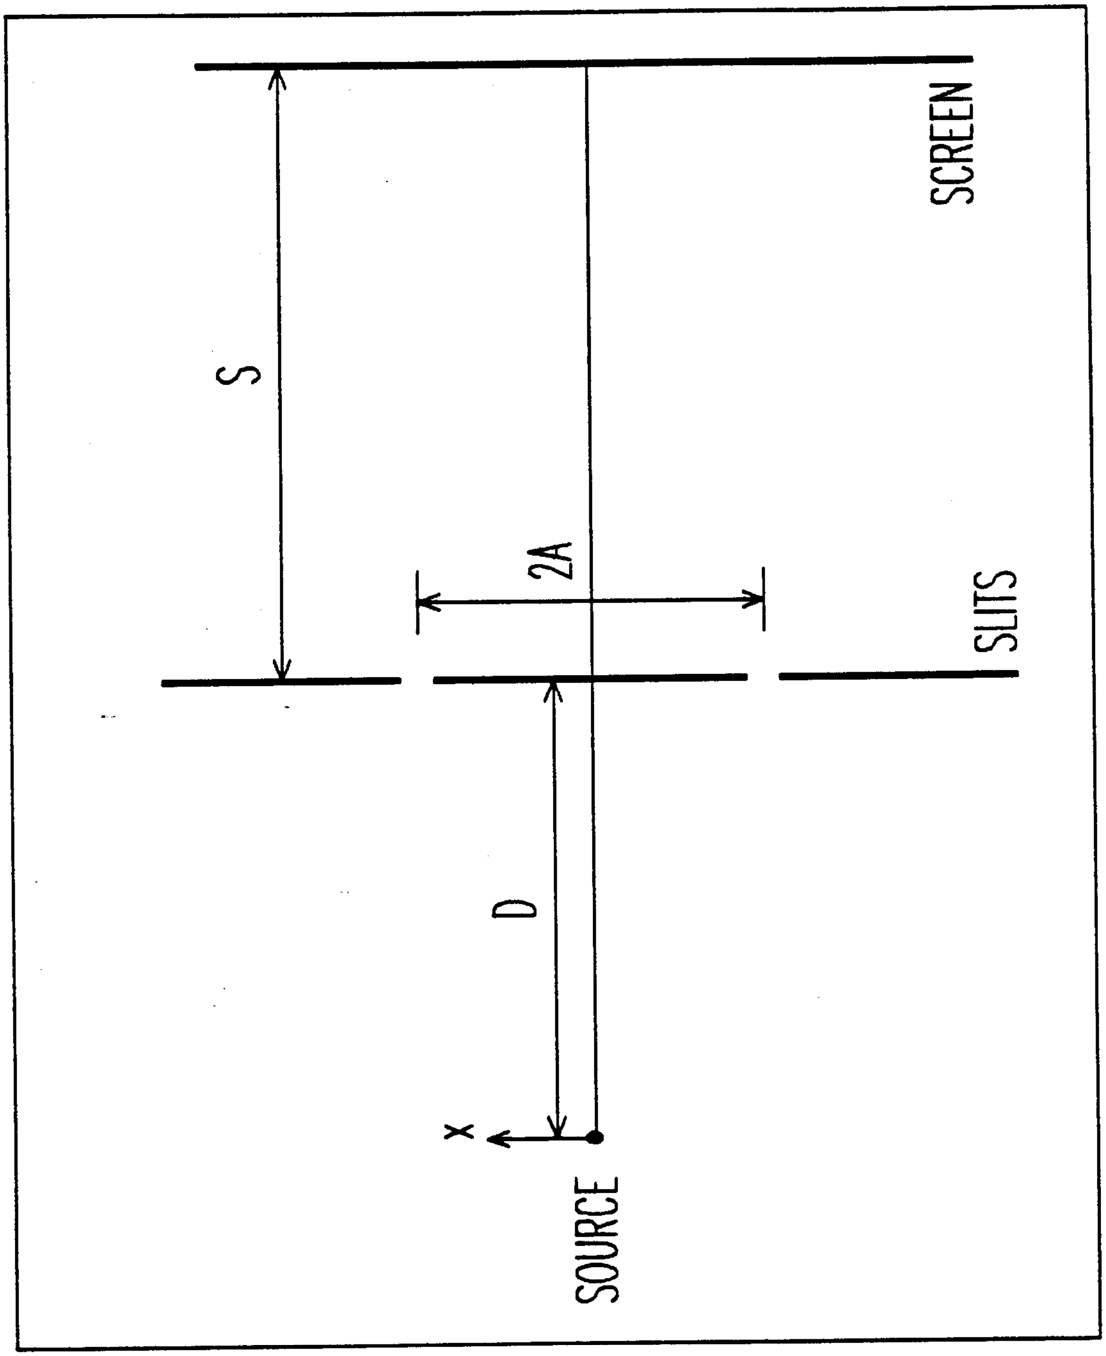

Figure 1.1: Young's Two-Slit Experiment

This can be rewritten as

$$x_s \frac{A}{D} = x_s \theta_s < \frac{\lambda}{16}, \tag{1.46}$$

where I have used θ_s to denote the angle of a ray that starts on-axis and passes through one of the slits.

If one now lets x_s represent the spatial size of an extended source and θ_s the characteristic angular divergence of the source, then one concludes that radiation with emittance $\varepsilon_r = \pi x_s \theta_s$ exhibits spatial coherence at certain wavelengths, namely if

$$\lambda > 16\varepsilon_r/\pi. \tag{1.47}$$

A more rigorous analysis shows[24] that for a gaussian beam, complete spatially coherence (i.e., no washing-out of fringes) is obtained if

$$\lambda \geq 4\varepsilon_r = 4\pi\sigma_r\sigma_r'. \tag{1.48}$$

While I have discussed coherence in the context of Young's experiment, and will continue to do so below, this is only for clarity. What the appearance of fringes in Young's experiment attests to is the coherence of the radiation from the source at a certain distance from the source (namely, the position of the slits) and within a certain region (namely, between the slits).

If one refers back to equation (1.37) for the emittance of single-electron undulator radiation, one sees that this condition for spatial coherence is in fact satisfied, so that single-electron radiation from an undulator is spatially coherent. In addition, one sees that if the electron-beam emittance is significantly less than $\lambda/4$, then the radiation is spatially coherent even for a beam of many electrons. One conclusion that can be drawn from this is that for a *zero-length* electron bunch that satisfies equation (1.48), the on-axis intensity is the coherent superposition of the radiation from each electron, so that the flux on-axis will be N^2 times the flux for a single electron.

Spatial Coherence of Longitudinally Extended Sources

Until now, I have considered the sources to lie in a plane of constant z (i.e., in a plane parallel to the planes of the slits and the screen). However, when an electron bunch acts as a source of synchrotron radiation, it has a longitudinal extent as well as a transverse extent. Imagine then that one has two equal-strength, on-axis point sources, one at z=0 and another at $z = \Delta z$, with $\Delta z \ll D$. Since both sources are on-axis, each individually illuminates both slits with the same phase. Hence, spatial coherence is maintained. However, depending on Δz, the coherent flux passing through the slits may be diminished or increased relative to that for a single source, due to interference between the two sources.

The optical path length difference for light from the two sources to the slits is simply Δz, to first order. Hence, if $\Delta z \ll \lambda$ and the sources emit in phase, then constructive interference will occur at the slits. Clearly, for N equal-strength, in-phase sources, confined to a longitudinal interval $\Delta z \ll \lambda$, the peak on-axis intensity is increased by N^2. Thus, if an electron-bunch is short compared to the wavelength and if the transverse coherence condition (equation (1.48)) is also satisfied, then the on-axis flux will be N^2 times the single-electron flux, preserving the result obtained for a zero-length electron bunch. This provides a dramatic increase in the radiation power (in regions of constructive interference), and is therefore a highly desirable result.

Consider also that if the two sources are out of phase by 180° and spaced by $\lambda/2$, then they will constructively interfere. Hence, if two electron bunches, each short compared to λ, are spaced by $(n + 1/2)\lambda$ (where n is an integer) and experiencing opposite acceleration, then they will radiate in phase at the slits (or some distant on-axis point) and hence their intensities will interfere constructively. In this way, a train of mutually-coherent sources can be created.

A typical undulator period, λ_u, is 2.5 cm. For a 120 MeV beam such as that produced by the SSRL pre-injector, the first-harmonic wavelength (assuming $K \ll 1$) is $\lambda_1 = 0.2\mu$m. A relativistic electron bunch with a bunch-length of 0.2μm would have $\delta t = 0.6 \times 10^{-3}$ps. While bunch-lengths of order 1 ps are possible (see Chapter 4), a sub-femto-second bunch length is not within the realm of current accelerator technology. Hence, it would seem that the promise of coherent undulator radiation is out of reach. The solution to this problem is called a Free Electron Laser (FEL), and I will discuss it in the next section.

Temporal Coherence

First, I wish to complete the discussion of coherence by explaining temporal coherence. Temporal coherence refers essentially to the monochromaticity of light, and can be understood by returning to Young's two-slit experiment with a single on-axis source. If one modifies the experiment by placing refractive material over one of the slits, then the optical path length for light going through that slit is increased, resulting in a phase-shift at the screen relative to the light going through the other slit. Suppose that the light going through the refractive material is delayed by time Δt. For a perfectly monochromatic source, this phase shift is irrelevant, and merely results in a shifting of the fringes. A perfectly monochromatic source is said to be completely temporally coherent, or to have $\Delta t_{coh} = \infty$.

However, for a realistic source, the spread in frequencies means that for too long of a delay, $\Delta t \geq \Delta t_{coh}$, the interference fringes will no longer appear. Instead, one will see the uniform illumination characteristic of incoherent light. To see this, imagine that a source has a spread $\Delta\omega$ of frequencies about ω_o. If these frequencies are in phase at the source at time t=0, then at $t = \pi/\Delta\omega$, the phase spread for light emerging from either of the slits will be 2π. If light from such a source is used in a modified Young's two-slit experiment with a refractive plate giving $\Delta t \geq \pi/\Delta\omega$, then the interference pattern will be washed out because the phase-shift will shift the fringes for the outer frequencies by one-half of the fringe spacing. Hence, $\Delta t_{coh} = \pi/\Delta\omega$ is the time over which the source is said to maintain coherence with itself. The coherence time can be improved to the desired degree by employing a mono-chromator with a sufficiently narrow band-pass.

The spectral broadness of undulator radiation at harmonic n can be estimated by taking the Fourier transform of a sinusoidal field oscillation of Nn periods, where N is the number of undulator periods. One finds that the FWHM of the spectral power distribution around each harmonic is

$$\Delta\omega_n \approx \frac{\omega_n}{nN}, \tag{1.49}$$

implying coherence times of

$$\Delta t_{coh} = \frac{N\pi}{\omega_1} = \frac{N\lambda_1}{2c} \tag{1.50}$$

and coherence lengths of

$$\Delta l_{coh} = c\Delta t_{coh} = \frac{N\lambda_1}{2}. \tag{1.51}$$

Since typically $N \gg 1$, one sees that undulator radiation is temporally coherent over many wavelengths.

1.2.3 Free Electron Lasers

Energy Transfer Condition

Imagine that an electron beam is sent through an undulator, and that the fundamental radiation wavelength is λ_1. Further imagine that a laser beam of the same wavelength is also made to pass through the undulator, so that it overlaps the electron beam. Electrons oscillate at the same frequency as the laser field (otherwise, the electrons wouldn't radiate at the same frequency).

The transverse motion of the electrons is determined primarily by the undulator fields[21], so that for any electron

$$\begin{aligned} \frac{d\gamma}{dt} &= -\frac{e}{m_e c^2} v E \cos(2\pi z(t)/\lambda_u) \cos(\omega t - kz(t)) \\ z(t) &= \beta^*(t - t_o), \end{aligned}$$

where $\beta^* c$ is the average longitudinal velocity of an electron which enters the undulator and laser fields at $t = t_o$. v and E are positive quantities, v being the peak transverse electron velocity and E the peak transverse field strength of the laser beam. (This ignores the transverse motion that individual electrons have in addition to the motion imposed by the undulator, consistent with the assumption of a small beam-emittance.)

When the phase of the undulations and of the laser field are such that the cosines have the same sign, then energy is transferred from the electron to the laser beam. The energy transfer is greater when the transverse electron velocity is greater, i.e., when the electron is near the zero crossing of its undulating motion. If the cosines are of opposite sign, then the field does work on the electron and thus energy is extracted from the laser beam. Clearly, to amplify the laser beam, one wants the former condition to hold. However, since the electrons do not travel exactly at the speed of light, it would seem that the condition cannot be maintained, and that the electron will alternately gain and loose energy as it falls behind the wave.

To see how net energy transfer to the wave is possible, note that there will be such transfer if

$$A = \int_0^L \cos(2\pi z/\lambda_u)\cos(kz - \omega t(z))\,dz > 0, \tag{1.52}$$

where I am now considering t a function of z (the electron position), and where L is the length of the undulator. Rewriting this using equation (1.52) one obtains

$$A = \int_0^L \cos(2\pi z/\lambda_u)\cos(2\pi z/\lambda_1(1/\beta^* - 1))\,dz \tag{1.53}$$

To maximize A, one equates the factors multiplying z in the cosines, obtaining

$$\frac{\lambda_1}{\lambda_u} = \frac{1}{\beta^*} - 1, \tag{1.54}$$

which is the "phase-slip" condition[20]. If the phase-slip condition is satisfied, then as the electron falls back, its transverse velocity is falling as well, so that when it has fallen back to where the laser field has changed sign, its velocity has also changed sign. The crests of the electron motion, where the velocity is zero, coincide with zero-crossings of the laser fields. The phase-slip condition, together with the electron beam energy and equation (1.34) for the radiation wavelength, must be self-consistently solved in order to find the conditions on K and γ necessary to achieve energy transfer to the laser beam for a given λ_u. Alternatively, one can find the "synchronous velocity", $v_s = \beta^* c$ that electrons must have in order to give (or receive) net energy from the laser fields.

Micro-Bunching

Most of the elements necessary for an intuitive understanding of FELs have been reviewed in the previous sections. I showed that under certain conditions, undulator radiation could be spatially coherent and that the on-axis flux would scale like N^2. It seemed, however, that to realize this would require an unrealistically short electron bunch. FELs succeed in spite of this because the radiation field interacts with the bunch to produce "micro-bunching", i.e., longitudinal density modulations on the scale of the light wavelength[21, 20].

For an initially longitudinally uniform beam traveling at the synchronous velocity, half the particles will loose energy and half will gain energy. Unless the density of the beam can be modulated on the scale of λ_1, there will be no amplification of the laser beam. However, since the electron beam energy is being modulated at the wavelength of the laser beam, so is the velocity. This velocity modulation has the same spatial frequency as the laser fields, and is thus just what is needed allow coherent power to be generated. (see Section 1.2.2, page 21).

Note that the derivation of the phase-slip condition ignored the fact that as electrons gain or loose energy, their velocities change. In order to get gain, the beam must initially travel somewhat faster than the synchronous velocity, so that bunching is not symmetric about the null of the laser field[20].)

Electron Beam Requirements for FELs

Having given a brief account of the physics at work behind the generation of coherent radiation in FELs, I now list the general beam-quality requirements for an FEL:

- The electron beam emittance should be less than the natural emittance of the undulator radiation. More precisely, equations (1.39) and (1.40) should be satisfied.

- The undulator parameter K, the undulator wavelength, and the initial beam momentum should be chosen to give an initial beam velocity somewhat greater than the synchronous velocity, so that micro-bunching occurs in the region where the electron beam looses energy to the laser beam. (See [20] for a more precise statement.)

- The bunch length, δt, should be as long as practical in order to provide more micro-bunches, which results in greater gain. Similarly, the peak current should be high, in order to give as much charge per micro-bunch as possible.

- The initial momentum spread of the bunch should be small enough to be within the "buckets" created by the laser field. If the initial momentum spread is too large, the gain is decreased because not all particles have the proper velocity relative to the synchronous velocity. More specifically[20], one wants

$$\frac{\Delta\gamma}{\gamma} \leq \frac{1}{2\mathrm{N}}, \tag{1.55}$$

where N is the number of undulator periods.

1.3 RF Guns

A RF or microwave electron gun [1, 2] is an electron source consisting of an electron emitter (the "cathode") immersed in the radio-frequency fields inside a metal-walled cavity. The SSRL RF gun is discussed in detail in Chapter 2. For the present, I will make some remarks about RF guns in general and about their history.

1.3.1 Varieties of RF Guns

There are two principle varieties of RF guns in use today[14]: "thermionic" and "photocathode" RF guns. They are distinguished by the type of cathode used or, more precisely, by the way the cathode is made to emit electrons.

Thermionic RF guns employ cathodes that must be heated in order to obtain emission of electrons. In the simplest design, a pillbox cavity[31] might be used with a thermionic cathode being part of one end-wall, and with a beam exit-hole in the opposite end-wall. When the RF fields in the cavity are in the accelerating phase, electrons are accelerated off of the cathode. With proper design, a large fraction of these electrons exits the cavity before the fields go into the decelerating phase. Those that do not exit the cavity are decelerated and turned around, and may return to impact the cathode (a phenomenon referred to as "back-bombardment"). As long as RF power is supplied to the gun, this cycle is repeated every RF period, resulting in a train of bunches spaced at the RF period.

Commonly-used cathode materials for thermionic guns are LaB_6 and dispenser cathodes (a tungsten matrix with work-function-lowering compounds added). Typical operating temperatures for LaB_6 are 1600°C[32], while the dispenser cathode for the SSRL RF gun is run at 950°C. Typical current densities for both these types of cathodes are in the 10-30 A/cm^2 range, though LaB_6. is capable of up to 200 A/cm^2[2] and advanced dispenser cathodes of up to 140 A/cm^2[33].

Problems with thermionic guns stem from the fact that emission occurs throughout the accelerating phase of the RF, and during every RF period. This results in a beam with a large momentum spread and a large time-spread as well. These issues are discussed in more detail below and in Chapter 2

Photocathode (or "laser-driven") RF guns[34] employ a photoemitting cathode material that is pulsed by a laser. Commonly-used cathode materials[14] include LaB_6, Cs_3Sb, and CsK_2Sb. Typically, the laser pulse is much shorter than the RF period, so that emission occurs only over a short phase interval. Thus, a photocathode RF gun can produce a much shorter bunch than a thermionic RF gun working at the same RF frequency, since in thermionic guns electron emission occurs continuously during the accelerating phase of the RF.

While it is not my purpose here to discuss the physics of photocathode RF guns in detail (see [13, 35, 36, 37]), I wish to indicate the reasons that such guns out-perform thermionic systems.

The high current densities possible from a photoemitter (as much as 600 A/cm^2[38]) result in significantly higher charge per bunch for photocathode RF guns than for thermionic RF guns. Because of the shorter phase-interval during which electrons are emitted in a photocathode RF gun, RF focusing effects (see Chapter 2) are greatly reduced relative to a thermionic RF gun, resulting in smaller emittances. Since the current density is very high, the pulse from a photocathode RF gun need not be greatly compressed in order have high peak currents; hence, one sees that photocathode based systems generate longer bunches than thermionic based systems, since the later systems *must* compress to very short bunches in order to achieve high peak currents. As noted in Section 1.2.3, a long, high peak-current bunch is preferred for FEL work.

Photocathode RF guns also have the advantage of being free of the cathode back-bombardment problem that can plague thermionic RF guns when long RF pulses or high repetition rates are used[39]. In addition, the use of the laser to trigger emission permits a more flexible bunch pattern, at least in principle. For example, if only N bunches are desired per RF pulse, then one need only fire the laser N times per pulse. The firing can be delayed until the cavity has fully charged, so that the fields are not changing between bunches. In a thermionic RF gun, emission occurs continuously, even while the cavity is filling, giving a train of bunches that vary in momentum until the beam-cavity system has come to equilibrium. This has implications for FEL use, where a small momentum spread is required.

The superior performance of a photocathode RF gun comes at the expense of the greater complexity of a photocathode-based system, which requires a complicated and expensive high-power, RF-synchronized laser and, for the highest-performing systems using Cesiated cathodes, a mechanism to withdraw and re-Cesiate the cathode at intervals [40]. While a thermionic RF gun with a magnetic bunching system is comparable in complexity and cost to a DC-gun-based system with a gap-and-drift buncher[41], a photocathode RF gun is substantially more complicated and much more expensive. Given that the thermionic RF gun system can out-perform a conventional DC gun system (excluding sophisticated state-of-the-art DC gun systems, employing several sub-harmonic bunchers, such as the SLC injector[41]) by orders of magnitude in brightness and peak-current, there is clearly a role for thermionic RF guns. Data to substantiate this claim is given in Chapter 4.

1.3.2 A Brief History of RF Guns

Kapitza and Melekhin[42] report that in 1948 the first microtron was constructed, and that it used field-induced emission from the gap of the accelerating cavity to generate a beam. In the broadest sense of the term, then, this was the first RF gun. In 1959, Melekhin proposed the use of a hot cathode placed off-axis in the microtron cavity. The primary difference between microtron guns and modern RF guns is that in microtrons, the cavity must not only accelerate the beam off of the cathode, but it must also provide acceleration each time the beam returns to the cavity. Hence, the microtron cavity must have entrance and exit beam holes, which greatly complicates the placement of the cathode and the initial trajectories of emitted electrons. Without reviewing microtron performance in detail, I will simply state that it does not equal that of a modern RF gun.

R.B.Neal would appear to be the first to use a RF gun other than a microtron gun. In a 1953 report[43] on work done at Stanford's Microwave Laboratory, Neal discusses experiments done with a hot cathode inserted in the first cell of a 2856-MHz linear accelerator section. These experiments were done in order to explore certain aspects of electron capture by RF fields, and not for the purpose of developing a new type of gun.

In 1975, Y. Minowa of Japan's Mitsubishi Electric filed for a Japanese patent on a RF gun[1], though it is unclear when an operational gun was actually built. Successful operation and experimental characterization of a multi-cell Mitsubishi RF gun with a LaB_6 cathode was reported at the 1989 Linear Accelerator conference in Nara, Japan[44].

The RF gun is usually considered to have been invented by G.Westenskow and J.M.J.Madey at Stanford's High-Energy Physics Laboratory (HEPL). Their design work[2] on a gun employing a pillbox-like cavity and a LaB_6 cathode was first reported in the literature in 1984, and the gun was installed and operated in 1985[39]. The HEPL system was the first to combine a RF gun with an alpha-magnet-based magnetic buncher[32, 45], resulting in a simple, compact system.

In 1985, J.S. Fraser and R.L. Sheffield of Los Alamos proposed use of a laser-pulsed cesium antimonide cathode in a multicell RF gun[34]. Experimental characterization of a prototype one-cell gun with such a cathode was reported in 1987[40].

The number of proposed RF gun projects increased rapidly after 1985. A 1990 review article by C. Travier[14] gives the number of RF gun projects world-wide as 22. Those readers interested in a listing of the parameters of existing and proposed RF gun projects are referred to this article and to Chapter 4.

1.3.3 Factors That Degrade Electron Beam Brightness

The degradation of electron-beam emittance, and hence brightness, in electron acceleration and bunching is discussed in many places in the literature[46, 47, 16]. Knowledge of these effects is necessary if one is to appreciate the advantages and disadvantages of RF guns. For the present, I will list and briefly discuss some of the effects involved, leaving the detailed discussion for the chapters that follow. Corresponding to each of these effects is the brightness increase to be gained by eliminating its influence.

Effects that Degrade Emittance

- *Thermal velocity spread at the cathode* is almost always a negligible effect, being overwhelmed by other effects. This is the case for the SSRL RF gun.

- *Non-linear space-charge forces* are a particular problem for high-brightness beams, which by definition have high charge density. These effects are mitigated by accelerating as rapidly as possible to relativistic energy, where the beam is "stiffer". For a beam of constant radial cross-section, the radial acceleration scales like $1/\gamma^3$ [7]. This is part of the motivation for RF guns, as I will discuss presently.

- *Time-dependent focusing forces due to RF fields* are a particular problem for thermionic RF guns, as mentioned above. These also have a non-negligible effect in photocathode RF guns[15] and in DC-based injectors, since time-dependent RF fields and long phase-length bunches are present during initial bunching. This is the dominant emittance-increasing effect in the SSRL RF gun.

- *Non-linear tranverse forces in RF cavities and from DC magnets* are avoidable, in general, though not always in practise, by increasing the apertures of the magnets or RF cavities, or by properly shaping the metal surfaces in the cavity. In the SSRL RF gun, non-linear tranverse fields in the gun are responsible for significant emittance increase and for broadening of the momentum spectrum. Solution of this problem by a simple cathode modification is discussed briefly in Chapter 4.

- *Chromatic aberrations from DC magnets* are a problem for thermionic RF guns, which emit a very broad momentum spectrum, since particles starting from the cathode at different phases are given different energies. In order to increase intensity, it is often desirable to use particles emitted over a substantial phase-interval, which entails dealing with a larger momentum spread. For the SSRL RF gun, using electrons emitted over $40°$ of phase entails accepting a momentum spread of about $\pm 10\%$.

1.3.4 Advantages and Disadvantages of RF Guns

Advantages

Until now, I have not discussed why RF guns are being adopted in preference to well-established DC gun technology. However, from the foregoing discussion it will have been anticipated that RF guns are capable of delivering very high-brightness beams.

The principle advantage of RF guns stems from the rapid acceleration that can be achieved with the strong electric fields possible in an RF cavity. The maximum fields possible in a DC gun are about 8 MV/m[2], while the peak surface fields in a 2856 MHz RF can be as high as 240 MV/m[44]. Typical output energies for DC guns are from 0.08 MeV to 0.4 MeV (kinetic energy), with the upper ranges being reached only by the most state-of-the-art designs. In contrast, the SSRL RF gun delivers a 2–3 MeV beam. This higher beam momentum greatly decreases the effect of space-charge forces in subsequent bunching and transport.

Another advantage of thermionic and photocathode RF guns is that bunching (in order to decrease the bunch length and hence increase peak current) can make use of relatively a simple magnetic bunching system (see Chapter 4), rather than the more complicated gap-and-drift prebuncher[41] and sub-velocity-of-light RF "buncher" acceleration section required in conventional injectors.

Disadvantages

The principle disadvantage (or difficulty) for a thermionic RF gun is the time dependence of focusing and energy gain in the gun, due to the time-varying nature of the RF fields. DC guns have no RF focusing in the gun *per se*, but it is inevitably encountered in subsequent bunching and acceleration. Similarly, while a DC gun emits a highly mono-energetic beam, that energy must be modulated (in order to modulate the velocity) for bunching. Hence, the time-dependence of the gun fields is not necessarily a *net* disadvantage relative to DC-based systems, but it does decrease the achievable brightness, and is one of the principle motivations for use of a photocathode.

A more significant problem for thermionic RF guns is that of cathode lifetime and survivability in the presence of back-bombardment. For sufficiently high repetition rates and long RF pulse lengths, back-bombardment can damage the cathode and degrade the gun's performance. In addition, back-bombardment can lead to current variation during the RF pulse, which causes problems for FEL applications[39]. The SSRL system is run with a 2 μs RF pulse and a repetition rate of 10 Hz, so that back-bomdardment is not a serious issue.

Another issue in thermionic RF guns is that of cathode thermal isolation. Since thermionic cathodes have operating temperatures typically in excess of 1000° C, the cathode support stem cannot be in direct contact with the metal surfaces of the cavity.

At the same time, it is necessary that the cathode be in electrical contact with the cavity. This is usually done using an RF choke[32], which can be problematical, leading to distortions of the fields in the cavity and unanticipated power losses. For the SSRL RF gun, a simpler concept was employed, as discussed in Chapter 2.

Conclusion

One cannot conclude from this discussion whether RF gun based preinjectors will in fact provide brighter beams than conventional preinjectors. Such a conclusion requires detailed simulations of the effects discussed in this section and the last

In Chapter 4, I present comparative data for the SSRL preinjector, other RF-gun-based preinjectors, and DC-gun-based preinjectors. It will be seen that, in terms of beam brightness, a thermionic RF gun system with magnetic bunching can outperform all but the most state-of-the-art DC-gun-based injectors, which employ high-performance DC guns with multiple frequency sub-harmonic bunching. Thermionic RF gun systems can outperform conventional "low-technology" injectors by orders of magnitude in brightness and peak current. That this can be done using low-cost magnetic bunching technology, with a system using a single RF frequency (the gun frequency is the same as the linear accelerator frequency), and with only velocity-of-light accelerator section, is a significant improvement in the simplicity and affordability of high-brightness injectors.

1.4 Overview of Thesis

Figure 1.2 shows a schematic layout of the SSRL RF gun project, including the gun, the gun-to-linac transport line (GTL), and the beginning of the first of three SLAC-type[48] 2856-MHz accelerating sections. Electrons start at the gun cathode and are accelerated by the fields in the gun cavity, forming a beam that is transported through the first part of the GTL to the bunch-compression alpha magnet. After emerging from the alpha-magnet, the train of bunches is "chopped" by the traveling-wave beam-chopper, which admits 3-5 bunches into the linear accelerator. (A more detailed discussion of the GTL will be found in Chapter 5.) The plan of this thesis in large part follows the path that electrons take—that is, it proceeds from the gun, through the alpha-magnet, to the linear accelerator.

In this section, I present a brief overview of the thesis. To do so, I must anticipate much that will not become completely clear until latter in the thesis. It is also appropriate at this time to be explicit about my individual contributions to the project, having noted the contributions of others in the acknowledgements.

Chapter 2 is in part a computer-aided explanation of the detailed workings of the RF gun. It starts with an overview of the concepts behind the gun and the goals of the design. The capabilities of various relevent simulation programs are discussed, as well as my methodology of applying the codes. The codes are used to explore design alternatives, to understand the physics at work in the gun, and, finally, to predict expected gun performance.

The simulation results presented in Chapter 2 are my own work (though the codes used were in some cases created by others, who are acknowledged at appropriate points). My contribution included evaluation and generation of gun design alternatives and modification of the design to satisfy the project goals. In particular, I determined the size and shape of the focusing noses (necessary to obtain good control of the tranverse beam size in the gun over a wide range of currents) as well as the necessary on-axis field ratio between the cells (necessary to obtain a longitudinal phase-space suited to magnetic compression). I created the code `rfgun` to provide a fast, accurate design tool, and took the leading role in the commissioning of the gun and GTL at SSRL.

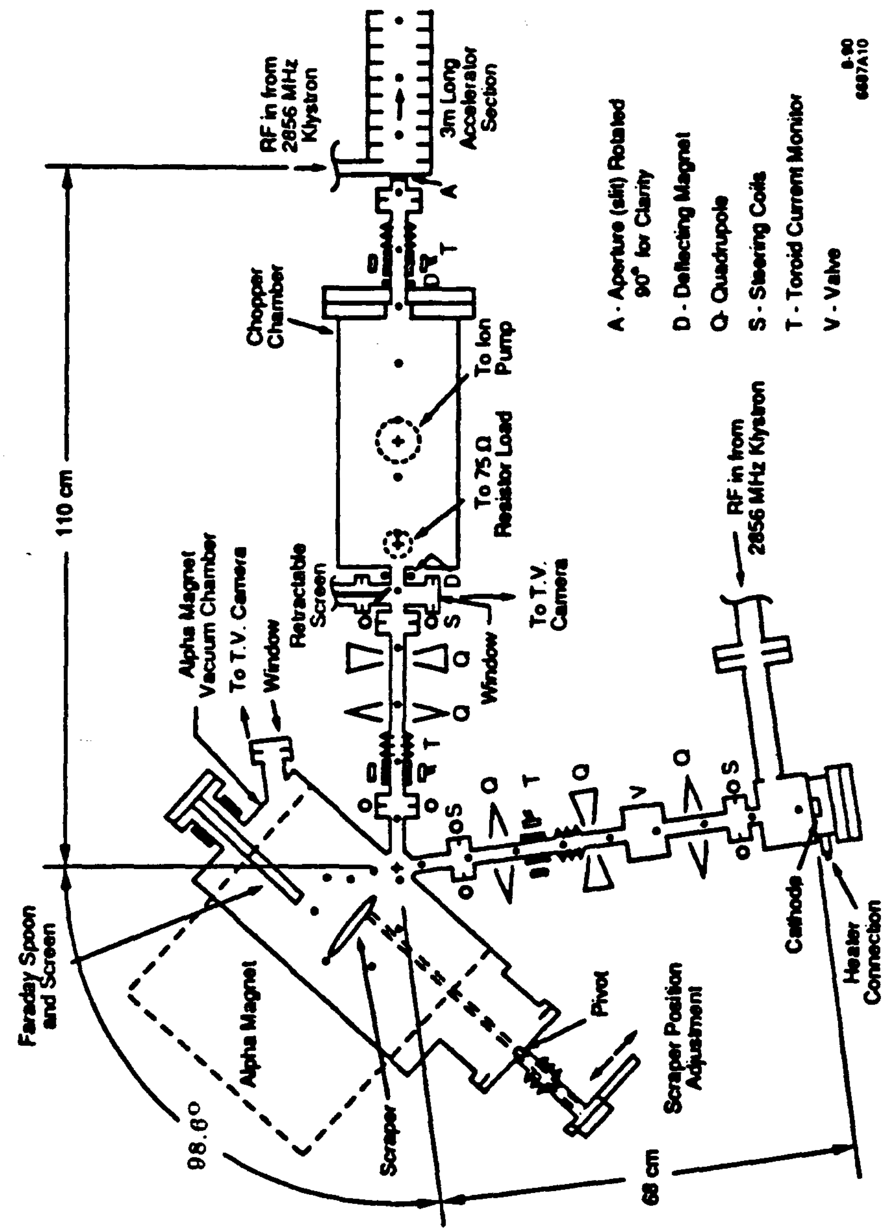

Figure 1.2: Schematic Layout of the Gun-to-Linac Region of the SSRL Preinjector

Design of the gun cannot be separated from the design of the GTL itself, which was also my responsibility. This included determination of the GTL optics, the length of the transport line, and the requirements for the alpha-magnet. While I will not discuss these details here, I did the physics design of all the GTL magnets (quadrupoles, steering magnets, and the alpha-magnet) and the beam-chopper (including the permanent magnet deflectors). My work included magnetic characterization of the magnets using a computer-aided measuring setup that I assembled and wrote the code for.

Chapter 3 goes into the properties of the alpha-magnet[45] in great detail. While the alpha-magnet was invented in 1963 and has been widely used in bunching and other applications, to my knowledge this is the first work in which full second and third order transport matrices are given. These matrices, along with highly accurate zero and first order results, are calculated by a program that I wrote, using a method that I developed and which I explain in the chapter. The method is applicable to finding matrices up to third order for any beam transport element (as is the code, if provided with the appropriate equation of motion). I also report on beam-optics experiments that I performed to test some of these calculations.

Chapter 4 discusses the subject of longitudinal dynamics, in the alpha-magnet-and-drift buncher and in linear accelerators. I present a new way of looking at the problem of matching the injected bunch longitudinal phase space, using contour maps of final phase and momentum as a function of initial phase and momentum. I discuss the well-known general principles of magnetic bunch compression, and employ results from Chapter 3 in order to show how an alpha magnet can be used for magnetic bunching, and under what circumstances. Using these ideas, I employ computer methods to find the optimal bunch-compression parameters for the SSRL RF gun. This is followed by detailed predictions of the performance of the SSRL pre-injector in terms of peak current, emittance, and brightness, as well as comparison with other projects. For this, I used the tracking code **elegant**[49], which I wrote specifically for the SSRL project.

Chapter 5 concentrates on experimental characterization of the gun and the pre-injector. It starts with a detailed walk through the GTL and continues with a discussion of the GTL optics. The remainder of the chapter gives experimental results and comparable results of *simulated* experiments. These experiments include measurement of the momentum spectrum, the emittance, and the bunch length. I performed all of the experiments and data analysis, as well as the simulations.

Computer simulations are employed extensively in the present work, because the nature of the problems dictates that numerical methods be used for precise solutions. In order to more clearly separate my own contributions from those of others, I have adopted the convention of putting the names of computer programs that I have written in type-writer type face (e.g., "`program`") while putting the names of computer programs written by others in upper-case letters (e.g., "PROGRAM"). (This is not always adhered to in figure labels.) In addition, I have mentioned, where appropriate, what modifications I have made to programs written by others.

Some of my programs (albeit those with the least literal physics content) are described in Appendix A, and others are either described in the text or will be described in forth-coming publications[49, 50]. While these programs represent a substantial amount of original work, I have decided to concentrate on the physics that the codes predict rather than the details of the codes themselves. Appendix A is included in part to advance a new philosophy of code integration, in part as an example of the implementation of that philosophy, and in part as documentation for that implementation.

Chapter 2

Gun Design and Simulations

As discussed above, one of the RF gun's principle advantages, and indeed the essential reason for using an RF gun rather than a DC gun, is the rapid acceleration of electrons, which greatly lessens space-charge induced emittance degradation. Unfortunately, the very fact that electrons emitted from the RF gun cathode are accelerated from thermal velocities to velocities approaching the speed of light means that analytical approaches to the electron dynamics in the gun are unlikely to be fruitful. Additional complications arise from the time-varying nature of the accelerating fields, and from the continuous emission of electrons from the cathode.

Because of these complications, it is necessary to employ computer programs to simulate the detailed operation of the gun and to evaluate alternative designs. In this chapter, I will discuss many aspects of the physics of the gun as explored with various computer codes. In particular, I will concentrate on steady-state simulations of single-bunch evolution in the gun. (Multiple-bunch simulations will be addressed in future publications.) In addition to discussing the physics behind the codes and the methodology of the simulations, I will discuss the design criteria for the gun and how the codes were used to explore alternative designs. Finally, I will present and discuss collated results of simulations of the gun as it was built, giving a picture of many aspects of gun performance. Experimental results are reported in a later chapter.

The primary computer codes that I employed in this context were MASK[51], SUPERFISH[52], and `rfgun`. SUPERFISH is a well-established code that calculates the frequencies and the field distributions for TM modes of resonant cavities. In the present work, it found application in finding higher-order mode frequencies, in "tuning" simulated cavity shapes to the desired frequency, and in computing field distributions for use by `rfgun`.

`rfgun` is a code that I wrote specifically for the SSRL RF gun project. It uses the longitudinal field profile generated by SUPERFISH and an approximation to the radial electric field and azimuthal magnetic field to calculate beam evolution in the absence of space-charge effects. Like other such codes, `rfgun` integrates the equations of motion for discrete "macro-particles", each of which represents many electrons. Its primary advantages are speed and simplicity of use, which make it a valuable design tool when coupled with the slower, more cumbersome, but also more accurate code MASK. `rfgun` also allows the investigation of the importance of various non-linear field terms by allowing the user to turn such terms on or off at will.

MASK is a "particle-in-cell" code that self-consistently integrates Maxwell's equations for the electromagnetic field and the Lorentz equation for simulated macro-particles, including the effects of space-charge. MASK's advantage over `rfgun` is that it can simulate space-charge and higher-order cavity modes, at the expense of greater complexity and greatly reduced speed. MASK is also more accurate in predicting the effects of non-linear fields near the cathode, which are poorly handled by `rfgun`'s off-axis expansion.

2.1 Gun Design Overview

The SSRL RF gun was designed as part of a larger project, the SSRL 3 GeV Injector for the storage ring SPEAR[26], and hence was required first of all to meet the needs of that project. The primary need of the Injector project was for a reliable high-current electron source that could be matched to the subsequent linear accelerator sections in such a way as to produce a beam with less than 0.5% momentum spread at 120 MeV/c. (This is discussed in more detail in Chapter 4.) The basic goal for the gun was to be able to provide 10^9 usable electrons per gun bunch, which, assuming operation at 10 pulses per second with the equivalent of two bunches accelerated per pulse and a very conservative filling efficiency of 10%, would allow filling of SPEAR to a (quite high) current of 100 mA (5×10^{11} electrons) in under ten minutes. I found in the course of my design studies that this goal was relatively easy to meet, requiring reasonable RF power, cathode current, and bunching.

2.1.1 Design Characteristics

Before examining the design criteria in detail, it is helpful to review the general characteristics of the gun[53, 54]. In doing so, I will necessarily mention many points that I will not discuss in detail until later.

The SSRL RF gun consists of a thermionic cathode mounted in the first cell of a 1-$\frac{1}{2}$ cell side-coupled 2856 MHz velocity-of-light standing wave structure, as illustrated in Figure 2.1. The gun was designed in collaboration with Varian Associates, and the basic cavity design is one used in Varian Medical accelerators. The modifications to the cavity were purposely kept to a minimum in order to reduce the magnitude of the research and development effort. This is not without its costs in terms of beam quality, since the Varian cavity is not optimized for elimination of non-linear RF fields.

Consisting as it does of three coupled resonant cavities, the gun has three possible "structure" modes[55] with frequencies near the fundamental frequency (as distinguished from higher-order modes of the individual cells, which are infinite in number). This subject is discussed in the references, and to be brief I will simply state that the gun is operated in the $\pi/2$ mode, which means that there is a phase-shift of 90°

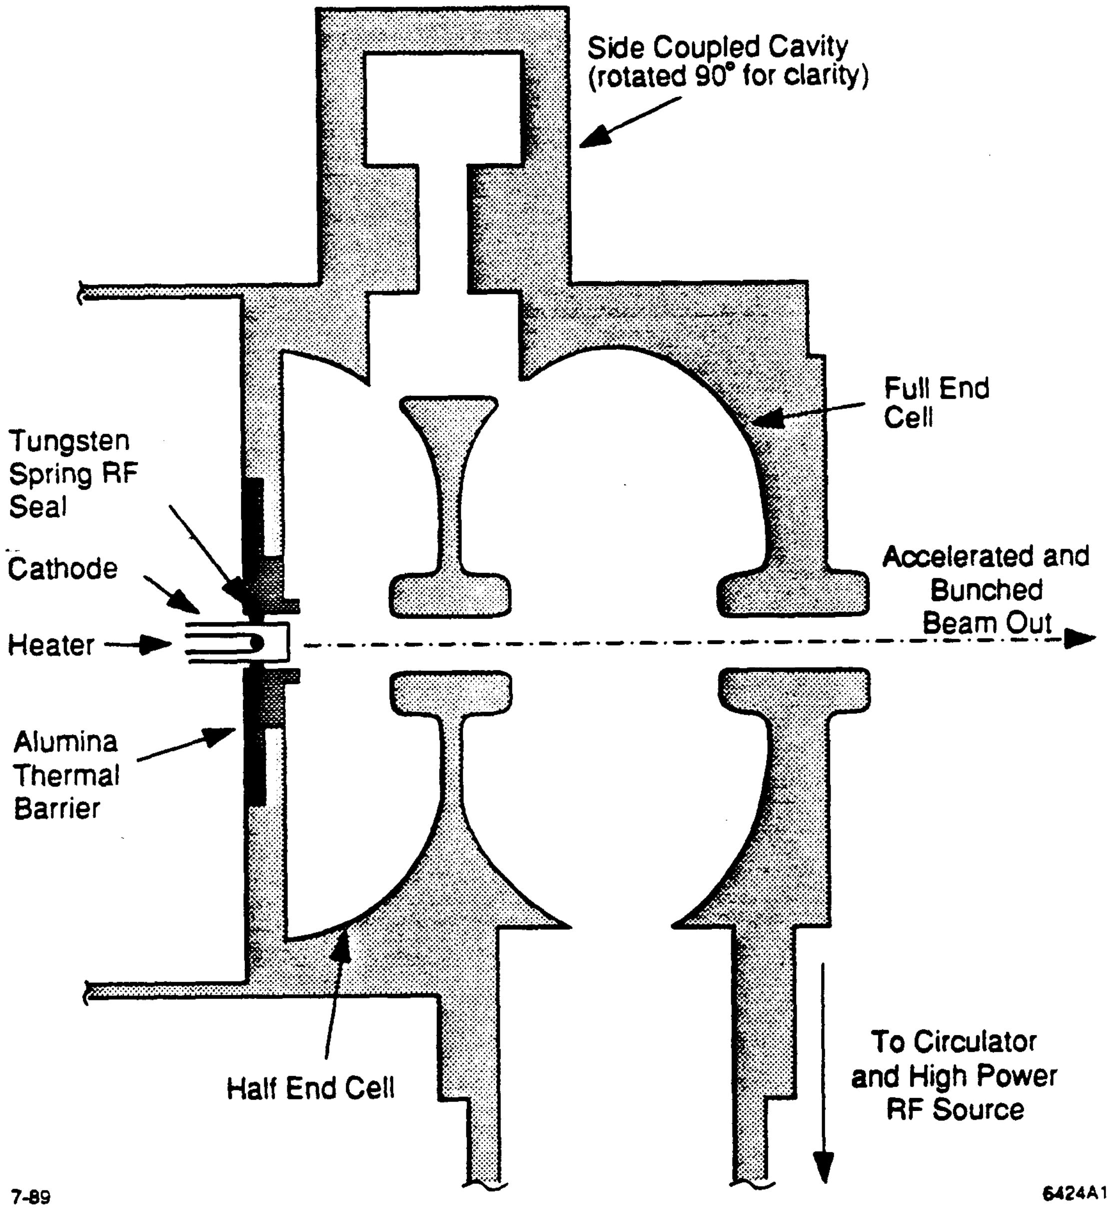

Figure 2.1: Cross Sectional View of the SSRL RF Gun

between the first cell and side-coupling cell, and between the side-coupling cell and second cell. Hence, the fields in the first and second cells are 180° out of phase. RF is fed into the gun through a port in the second cell, via rectangular wave guide, as shown in Figure 2.1, and fills the first cell through the coupling cell.

The design of the Varian side-coupled structure is such that the length of a full cell is one-half the free-space wavelength of 2856 MHz radio waves:

$$\mathrm{L} = \frac{\lambda}{2} = \frac{\mathrm{c}}{2\mathrm{f}}, \tag{2.1}$$

where L represents the periodic length of the structure and $\lambda = 10.497\mathrm{cm}$. For an accelerator consisting of a long chain of such cells, a relativistic particle that arrives at the center of one cell at the RF crest is guaranteed to arrive at the center of all subsequent cells at the RF crest in each cell. Hence, it will achieve maximum acceleration from each cell[56]. This is why this structure is referred to as a "velocity-of-light" structure. The RF gun is a very short version of such a structure, with the electrons being emitted from a cathode mounted in the end-wall. (The end wall in the first cell does not change the fields in the remaining part of the gun, since the end wall is placed at a location of symmetry.)

The cathode for the RF gun is a Varian dispenser cathode[33], with a flat, circular emitting surface of 6mm diameter, capable of current densities of up to $140\mathrm{A/cm}^2$. It is mounted in a modified version of Varian's high-voltage isolation mounting, which is used in Varian's DC guns. In order to provide focusing of the electron beam in the first cell, we have put a metalized ceramic annulus around the cathode. This annulus and the cathode must be in RF electrical contact with the metal walls of the RF cavity in order to avoid distortion of the RF fields. One way to achieve this is using an RF choke[32]. For the SSRL gun, a simpler concept was developed.

RF electrical contact between the cathode stem and the annulus is achieved by a toroidal tungsten spring around the cathode stem. This spring fits snugly into a toroidal cavity in the annulus. Electrical contact between the outer diameter of the annulus and the cavity is achieved through a knife-edge on the annulus that bites into a soft metal O-ring that rests in a toroidal channel in the back side the cavity wall. The reason for making the annulus out of ceramic is to provide a heat-barrier to prevent heat from being conducted too readily into the metal walls of the cavity, which would cool the cathode. If this "heat leak" were too great, the cathode might not achieve the temperature ($> 900°C$) necessary for emission. The tungsten spring serves a similar purpose, in that it connects the cathode and annulus electrically without providing an easy path for heat flow.

Pyrometric measurements, shown in Figure 2.2, demonstrate that it is an easy matter to achieve cathode temperatures in excess of 1050°C, which is the approximate temperature required to "convert" the cathode (which refers to a chemical change that must occur before good emission is obtained). In operation, the cathode temperature is closer to 950°C. The Figure shows measurements that I took using two different pyrometers, corrected for the emissivity of tungsten. I found that the temperature variation across the cathode was less than 5°C over the entire cathode surface, and less than the measurement resolution of about 2°C out to about 90% of the cathode radius.

2.1.2 Gun Operating Cycle

In the steady-state, as long as RF is supplied to the gun, the gun operating cycle repeats itself every 350 ps (i.e., at the RF frequency). Electrons that start from the cathode during the accelerating phase of the RF in the first cell are initially moving at thermal velocities. The RF fields accelerate these electrons rapidly, as a result of which a beam is injected into the second cell (provided the fields are high enough). By virtue of the time it takes for the beam to get from the cathode to the second cell, the particles arrive in the second cell during the accelerating phase of the fields in that cell. Hence, the second cell continues the acceleration of the beam that began in the first cell. In a matter of about 330–360 ps after being emitted from the cathode (how long depends on the field), the first electrons of the bunch are ejected from the gun. During that same time, electrons that do not make it out of the first cavity are back-accelerated into the cathode, causing additional heating of the cathode surface. In addition, some particles that do not make it out of the second cell (because they entered too late in the accelerating phase of that cell, with too little momentum), will be back-accelerated into the first cell, and contribute to the back-bombardment of the cathode. I will return to these points in the next section, where I will show a series of

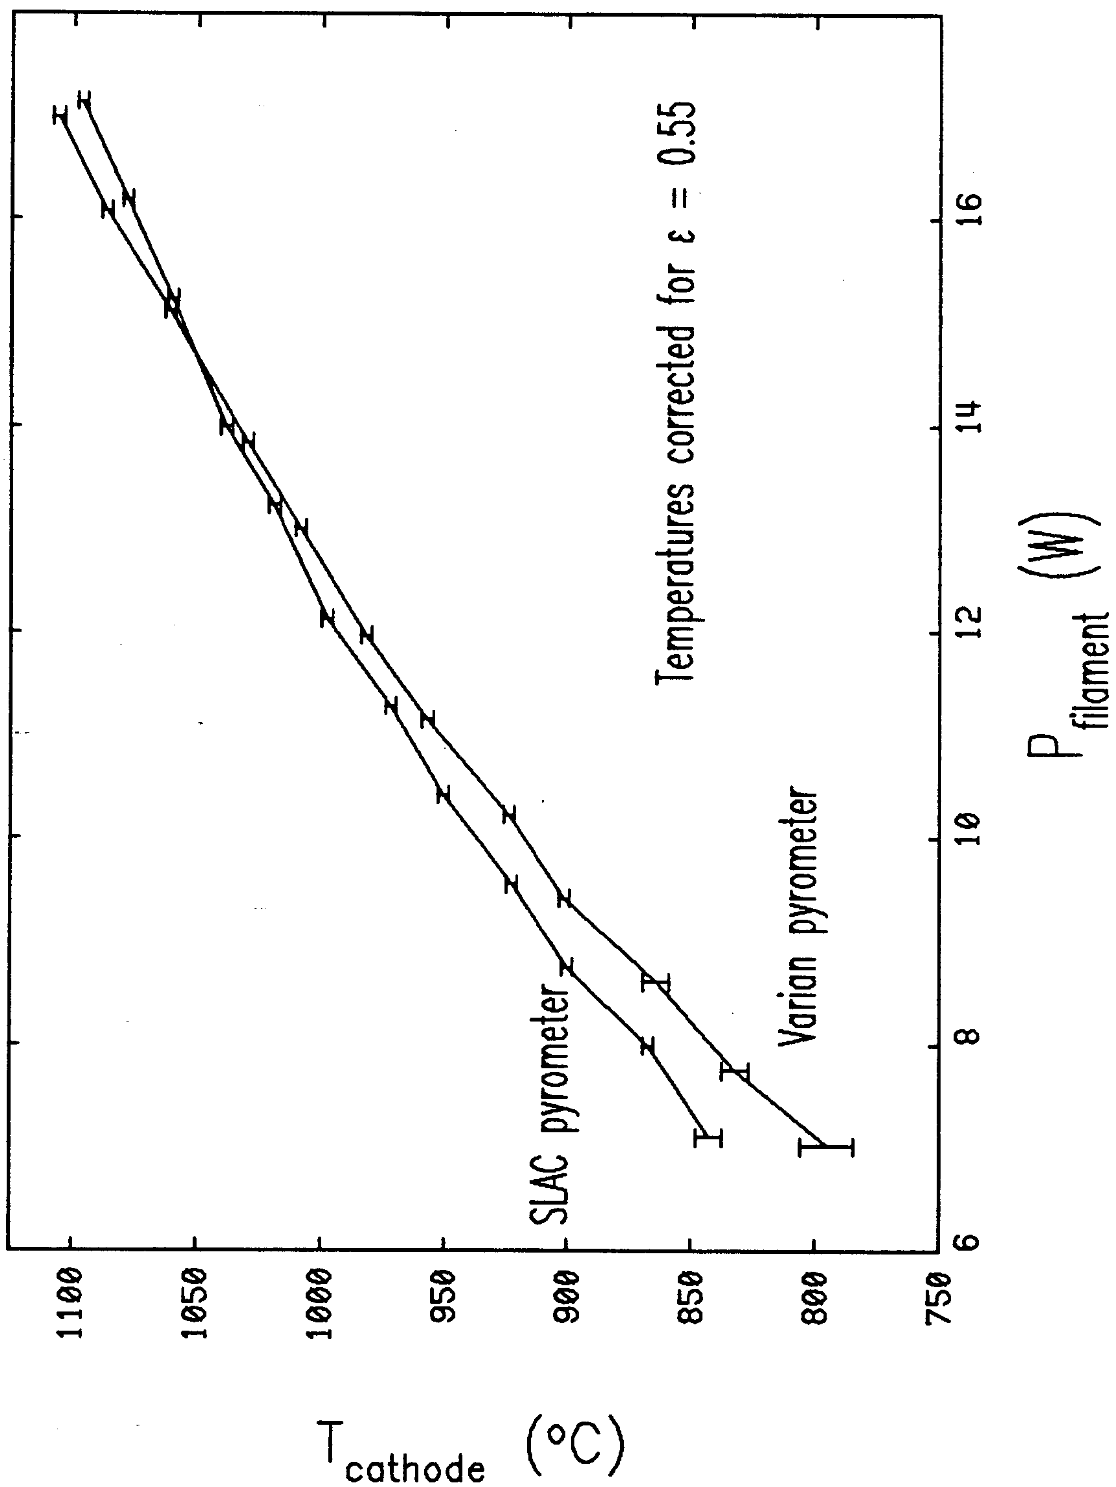

Figure 2.2: Pyrometric Measurements of the Gun Cathode Temperature

"beam snapshots" (from MASK simulations) at various points in the operating cycle.

Those particles that exit the gun comprise the "gun current" or "the beam". Because of the RF nature of the accelerating fields, the gun current for each cycle has a large spread in momentum, from some maximum down to essentially zero. However, we shall see that about half the particles in the beam have at least 80% of the peak momentum. Similarly, there is a wide spread in exit times, but about half the particles are within 25 ps or so of the particle with the peak momentum.

2.1.3 Matching to the RF Source

There are two principle limitations on the gun current. One is the current density that is available from the cathode. The other is the amount of RF power that is available to accelerate that current. Extraction of current from the gun requires a certain minimum electric field level in the cavity, otherwise current emitted from the cathode will not be accelerated sufficiently rapidly to make it out of the gun before the RF fields reverse sign. In order to maintain the electric fields in the cavity and accelerate electrons, one must supply sufficient RF power to compensate the power that goes into the beam, as well as the power that must be dissipated in the cavity walls to maintain the electric fields.

Early design studies on the gun indicated that 1–1.5 A of current at an average kinetic energy of 2 MeV was feasible. In addition, it was anticipated that 5 MW of RF power could be supplied to the gun. The beam power is simply the product of beam energy and current, from which one concludes that 2-3 MW of RF power must be supplied "for the beam" in this case. An additional power loss of about 1 MW occurs in the walls of the RF cavity, as a result of creating the electric fields that provide acceleration to 2 MeV. This makes a total of 4 MW, which is conservatively below what we anticipated would be available. The cavity must be "matched" to the RF source in order to make best use of the available RF power in the presence of beam. The cavity was thus made to be "over-coupled", with a normalized load impedance of $\beta = \mathrm{P_{beam}}/\mathrm{P_{wall}} + 1 \approx 4.15$. (This is topic is discussed in the references[57, 58].)

2.1.4 On-Axis Field Profiles

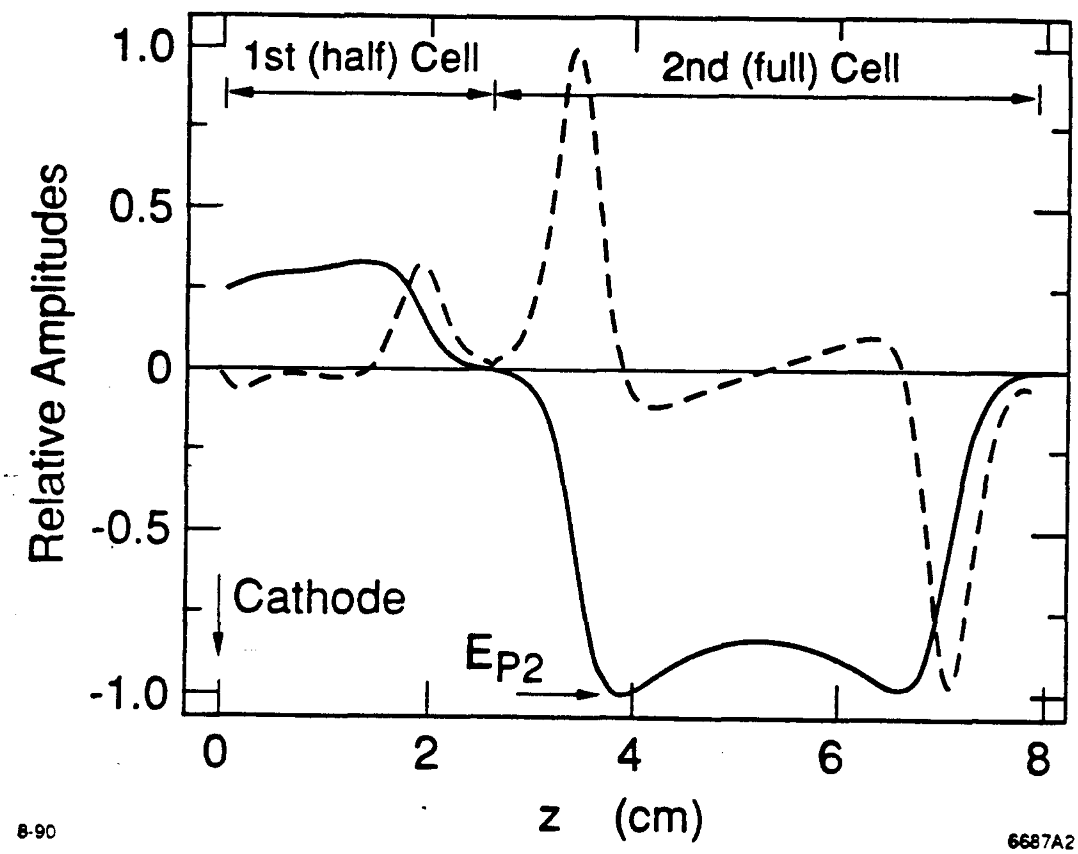

Figure 2.3: Gun On-Axis Longitudinal and Radial Electric Field Profiles

Computer studies also showed that in order to optimize the bunch longitudinal phase-space for magnetic compression, it would be necessary to delay the arrival of electrons in the second cell. The most straight-forward way to do this was by accelerating the electrons less rapidly in the first cell. Hence, the field amplitude in the first cell is approximately one-third of that in the second cell, as shown in Figure 2.3 (calculated with SUPERFISH). This was accomplished by modifying the coupling slots that connect the on-axis cells and the coupling cell[44].

The on-axis fields in the gun are related to those shown in Figure 2.3 (which is for the first RF gun, as built) by the multiplicative factor $E_{p2}\cos(\omega t)$, where E_{p2} is the peak, on-axis field in the second cell. E_{p2} provides a convenient measure of the excitation level of the RF fields, and I shall use it for this purpose throughout my discussion. Another important quantity related to the excitation level of the cavity is the peak surface field, E_{ps}. This field is of course proportional to E_{p2}, and since there are breakdown limitations on how high E_{ps} may be, there are limits on how high E_{p2} may be. For the RF gun as it was built, SUPERFISH gives $E_{ps} \approx 2 \cdot E_{p2}$. For 2856 MHz RF in a copper cavity with a surface finish of the quality that is achieved in the RF gun, $E_{ps} \leq 240$MV/m is feasible[44]. Typically, Varian operates its accelerators with $E_{ps} \leq 165$MV/m[44]. We will see below that the optimum operation of the RF gun is comfortably below the breakdown limit.

2.1.5 Design Goals

Having given an overview of the gun and the concepts involved, I am now in a position to list and briefly discuss the design goals [53, 54]:

1. There should be at least 10^9 "usable" electrons per S-Band bunch, for cathode current densities of less than 100A/cm^2, i.e., at least 10^9 electrons per bunch with momentum greater than 80% of the maximum momentum. This momentum window was established because it seemed feasible to transport a $\pm10\%$ momentum spread from the gun to the linear accelerator without excessive losses, and because this momentum range turns out to contain about half the current that exits the gun.

2. The average momentum in the useful beam should be 2 to 3 MeV/c, for surface fields below the breakdown limit. The primary reason for choosing this momentum range was to reduce the influence of space charge after the gun. Operating in this momentum range also results in more efficient extraction of current from the gun than does operating at lower momenta.

3. The longitudinal phase-space should be suitable for magnetic compression. As I will discuss in detail in Chapter 4, this implies that the momentum-vs-time characteristic of the gun longitudinal phase-space should be near-linear and monotonically decreasing with time.

4. The focusing structure around the cathode should provide a gently-converging beam for a wide range of current densities. Since it was initially not known what current densities would be achievable in normal operation, it was desirable that low current beams should not be over-focused and that high current beams should show some convergence. In the absence of focusing noses, the beam would fill the exit tube in the first cell or even hit the cell noses (for high current density).

5. The normalized RMS emittance for the useful beam, defined by equation (1.9), should be less than $15\pi \cdot \mathrm{m_e c} \cdot \mu\mathrm{m}$ over the entire range of current densities, where the averages are taken over the useful electrons in the beam. This corresponds to a geometric emittance of less than about $3\pi \cdot \mathrm{mm} \cdot \mathrm{mrad}$, and was selected based on what seemed feasible from initial studies.

6. The average beam power returning to hit the cathode should be manageable, i.e., not greater than the filament power (about 11W) used to heat the cathode, and preferably below 5 W, in order to ensure stable operation and long cathode lifetime.

2.2 Simulation Codes and Methodology

As I indicated in the introduction, the nature of particle motion in the RF gun is such that analytical methods are of little use if one wants detailed, accurate predictions. (This is less so for photocathode RF guns, where the beam is emitted in a short pulse, triggered by a short laser pulse[37, 12, 13].) Electrons go from thermal velocities at the cathode to relativistic velocities in the space of a few centimeters, meaning that neither non-relativistic nor highly-relativistic approximations are adequate. What is more, the rate at which an electron is accelerated depends on the phase at which it is emitted from the cathode. Depending on when it is emitted and what the fields are in the gun, an electron may exit the gun with $\beta = 0.98$ or $\beta = 0.01$, or it may not exit the gun at all, returning rather to hit the cathode (again, with a wide range of possible velocities). Some electrons even oscillate between the first and second cells one or more times before finally exiting the gun or hitting the cathode. When one adds to this complexity the additional complexity of space-charge effects, the problem is even more clearly out of the realm of analytical solution. Rather than attempt to find approximate analytical tools, then, I have employed numerical methods exclusively.

I will not attempt to explain the detailed workings of the codes that I have used. This is treated in the references[51, 52, 50]. While the version of MASK that I use is non-standard, the modifications I have made are primarily to the user interface. In addition, I have added a number of capabilities that were necessary for simulation of the RF gun. My version of MASK has the capability to simulate a cathode with emission limited at some specified uniform current density. It also allows one to easily inject simulation particles from one MASK run into another MASK run, which proved necessary because the two cells of the gun had to be simulated separately.

2.2.1 Tuning and Boundary Conditions

Before simulating the gun with MASK or `rfgun`, one must first check the cavity profile using SUPERFISH, in order to determine that the computed resonant frequency for the fundamental mode is as expected. Because of the off-axis coupling cell, it is not possible to run SUPERFISH for the entire gun, since SUPERFISH is for cylindrically symmetric cavities only. Hence, I have run SUPERFISH for the first and second cells separately. Simulation of the entire gun cavity could be done with a three-dimensional code like MAFIA[59], but I have not done this.

If the cavity had a uniform $\pi/2$-mode, then the longitudinal field at the junction between the cells (i.e., at $z = \lambda/4$, measuring from the end wall of the first cell) would necessarily be zero for the $\pi/2$ mode, since the fields in the two cells are 180° out of phase. In this case, one would use Dirichlet boundary conditions at $z = \lambda/4$ for both cells. The boundary condition at $z = 3\lambda/4$ is less clear, since there is no following cell to provide symmetry—the boundary condition is a combination of Neumann and Dirichlet. Dirichlet boundary conditions would be appropriate at $z = 3\lambda/4$ only if there were following cells to provide the necessary symmetry. What is more, since the RF gun has a non-uniform $\pi/2$-mode, one cannot conclude that Dirichlet boundary conditions are appropriate at the $z = \lambda/4$ boundary either. Because one cannot decide exactly what the boundary conditions should be without first simulating both cells (which SUPERFISH cannot do) and the structure following the second cell, the problem is in fact only solved by implicitly giving up the approximation of independent cells. Rather than do this, I elected to use Dirichlet boundary conditions at $z = \lambda/4$ and $z = 3\lambda/4$, since these are most likely to be closest to the actual boundary conditions.

To see that this is a justifiable approximation, consider that the cutoff frequency of the beam tube (which has a radius of $R_t = 3.8$mm) is[31]

$$f_{\mathrm{cutoff}} = \frac{2.405c}{2\pi R_t}, \tag{2.2}$$

or approximately 30 GHz, compared to $f = 2.856$GHz for the fundamental mode of the gun. Hence, the fields in the beam tube should fall off rapidly in moving into the beam tube from either cell. The $1/e$ distance is given by

$$d = \frac{c}{2\pi\sqrt{f_{\mathrm{cutoff}}^2 - f^2}}, \tag{2.3}$$

which comes out to 1.6mm for the present case, compared to about 7mm for the distance from the beginning of any cell nose to the nearest boundary plane. Hence, one expects that the fields in the cells and the resonant frequency of the cells will be insensitive to the boundary conditions, and this is indeed what I have found to be the case. I find that the resonant frequencies of both cells change by less than 400kHz in changing from Dirichlet to Neumann boundary conditions in SUPERFISH. The on-axis longitudinal fields differ visibly only near the boundaries, and then only by a small fraction of the peak field.

The actual dimensions of the gun cavity are different from the original design. In particular, the re-entrant noses (not the focusing noses, but those belonging to the Varian cell design) were of different lengths than specified in the design. This was a result of cavity tuning during machining. All of the data I present will be for the first RF gun as it was actually built (which I refer to as "the gun as built"), unless otherwise noted.

I found that when the cavity shapes for the gun as built were put into SUPERFISH, the frequencies of the cells were different from each other and from measurements, with the second cell calculated frequency being about 2838 MHz. This discrepancy is most probably a result of imprecise knowledge of the exact actual cavity dimensions and the effect of the coupling slots and the coupling cell. Upon modeling the second cell in MASK, I found that I obtained very nearly the same frequency as in SUPERFISH. I decided to tune the first cell to 2838 MHz also, as this frequency is within 0.5% of the goal of 2856 MHz (I will show below that the error introduced by this is small). Even after tuning the first cell to within 1 MHz of 2838 in SUPERFISH (which was done by slight alterations of the upper radius of the cell nose), I found that the frequency given by MASK was 30 MHz low. This is probably attributable to the coarseness of the mesh in MASK, implying that the agreement obtained for the second cell was fortuitous. I found it necessary to insert an artificial tuning plug into the MASK simulation of the first cell (see below). This plug is far from the beam and makes no significant change in the on-axis fields.

Figure 2.4 shows the SUPERFISH-generated field lines for the first and second cell. Table 2.1 lists SUPERFISH-generated parameters for both cells, some of which are self-explanatory, others of which are explained below.

While SUPERFISH calculates cavity resonant frequencies explicitly, MASK calculates only the time-dependent field evolution. In order to find the resonant frequency of a cavity simulated in MASK, I "hit" the cavity with a relatively broad-band signal and looked at the frequency of the ringing. How this was done is discussed below. For

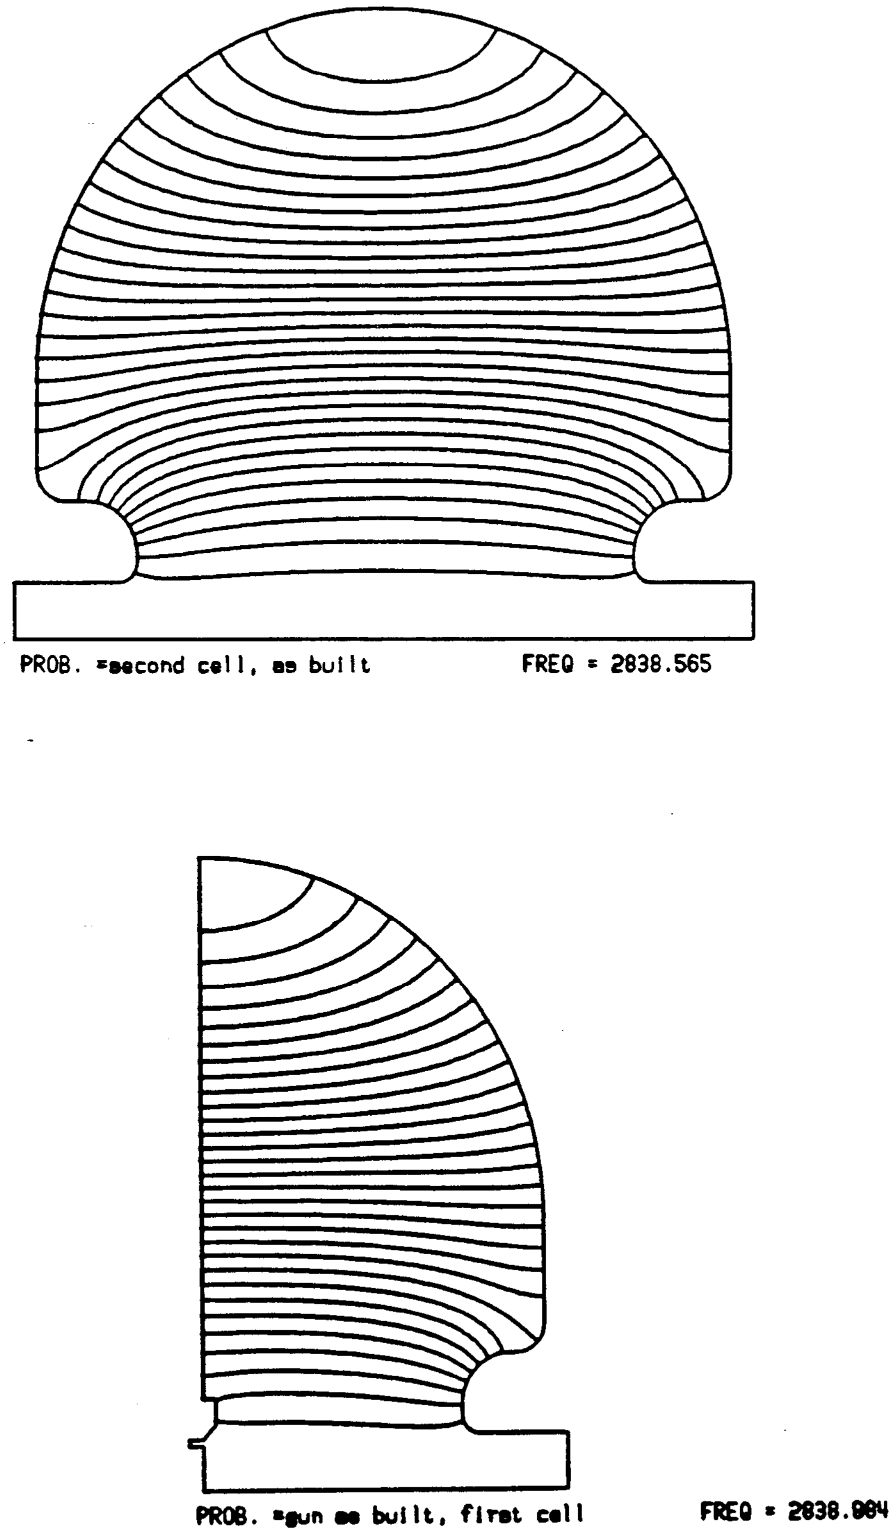

Figure 2.4: SUPERFISH Field Line Plots for RF Gun Cells

Table 2.1: **SUPERFISH-generated Cell Parameters**

quantity	Cell 1	Cell 2	unit
length	2.624	5.248	cm
Q	11777	18671	
shunt impedance	2.44 100	8.45 161	$M\Omega$ $M\Omega/m$
K_i	0.0831	0.1473	$mJ/(MV/m)^2$
E_{ps}/E_{pi}	1.79	1.92	
transit time factor	0.766	0.787	
effective length V/E_{peak}	1.82	3.45	cm

now, suffice it to say that the simulated frequencies of the first and second cells were 2834.5 MHz and 2837.8 MHz, respectively, where I determine the frequency from the time between subsequent zero-crossings of the electric field at some fixed point in each cell. This is a valid procedure provided there is only one mode that is appreciably excited, which I took pains to ensure was the case, as I discuss below. Fourier analysis would have required simulating the cells for 1 μsec in order to determine the frequencies to within 1 MHz. This would have taken about 20 days of dedicated computation by SSRL's VAX 8810. While these frequencies are the same to within 0.1%, the difference is not negligible and must be compensated for.

The final cavity profiles used in MASK are shown in Figures 2.5 and 2.6. The solid lines show the desired profile (which is the actual profile, except for the tuning artifice in the first cell), while dots show the grid points that were filled with "metal" in MASK in order to achieve that profile. The choice of boundary conditions in MASK is even more complicated than for SUPERFISH, since the beam-induced fields have no definite symmetry and are not constrained to frequencies below cutoff. I chose to use Neumann boundary conditions in all of my MASK simulations, because the copy of the code that I have does not implement Dirichlet boundary conditions. I have verified that the influence of the boundaries on the beam is negligible by simulating the first cell with a long exit-tube, and comparing the results to a simulation which ends 0.43 mm (one longitudinal grid spacing) after $z = \lambda/4$.

Figure 2.7 shows on-axis longitudinal field amplitudes from SUPERFISH and

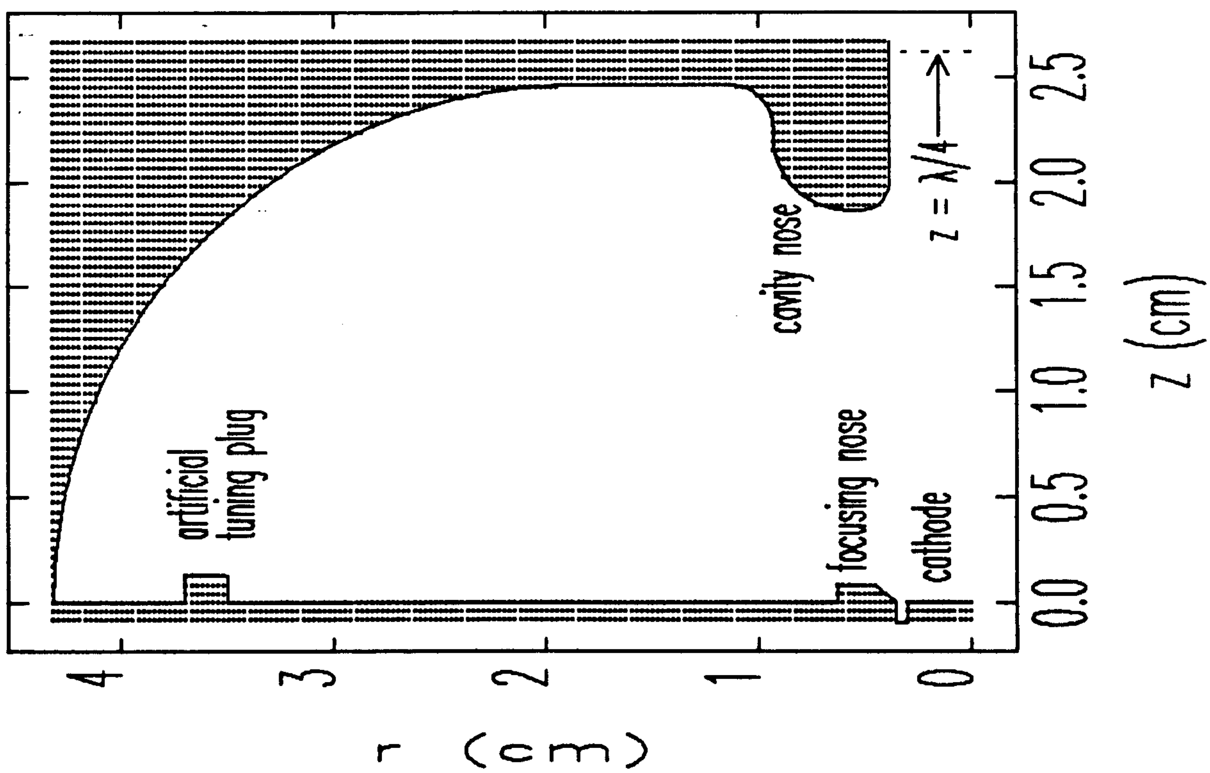

Figure 2.5: Profile Used in MASK for the First Cell

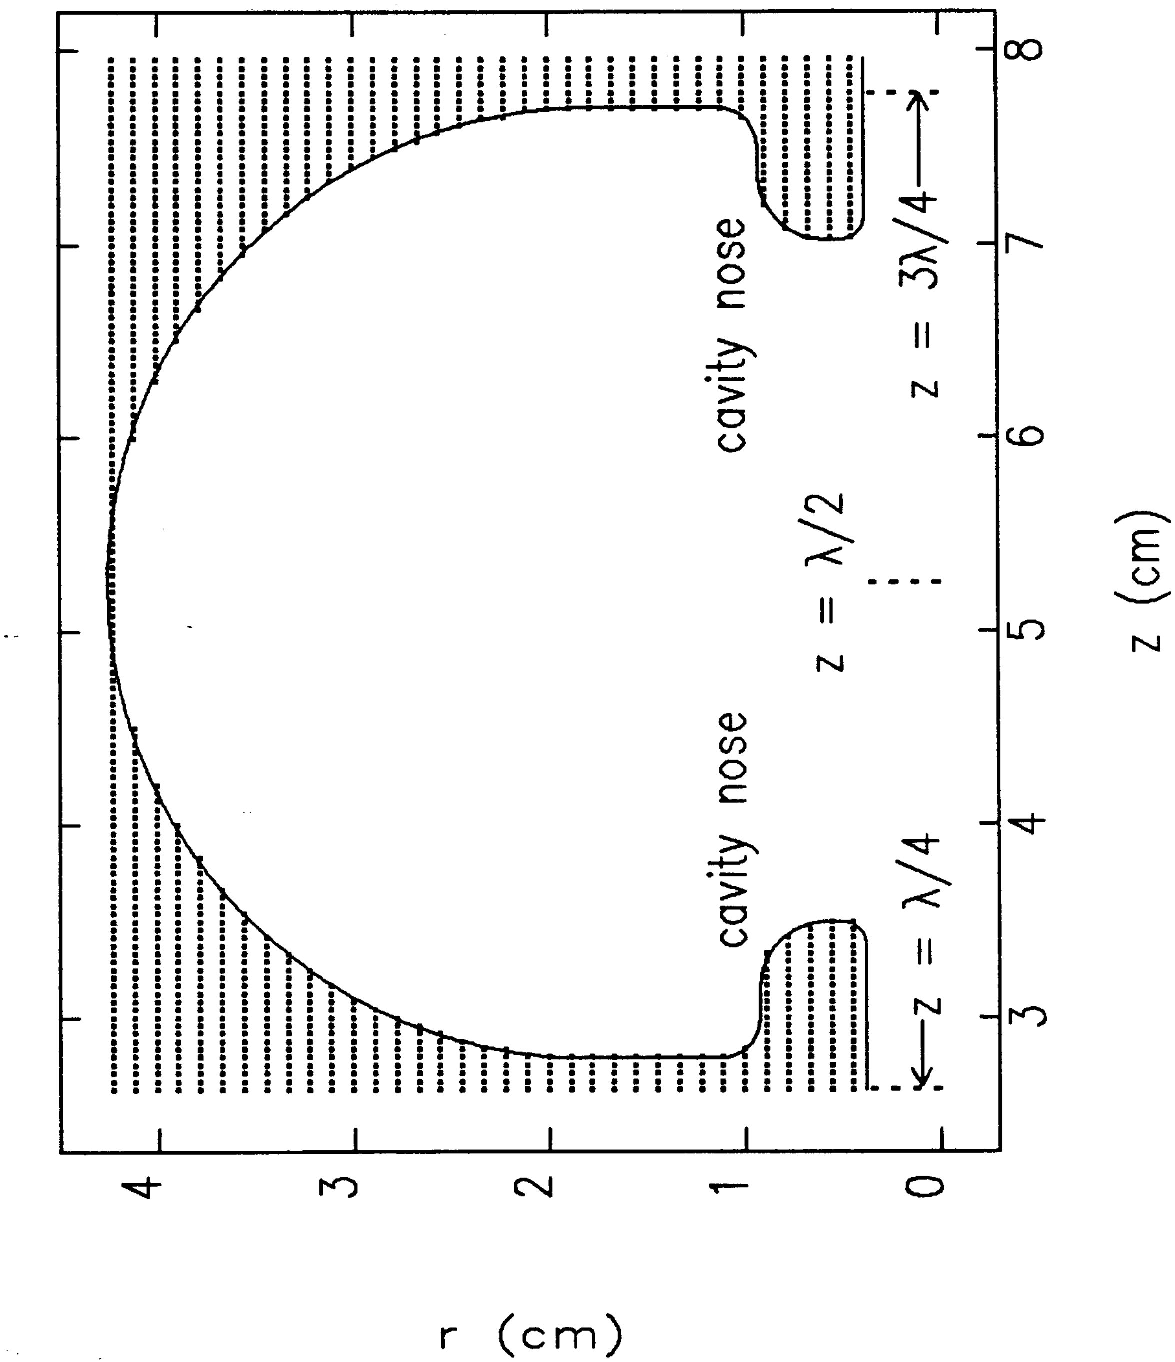

Figure 2.6: Profile Used in MASK for the Second Cell

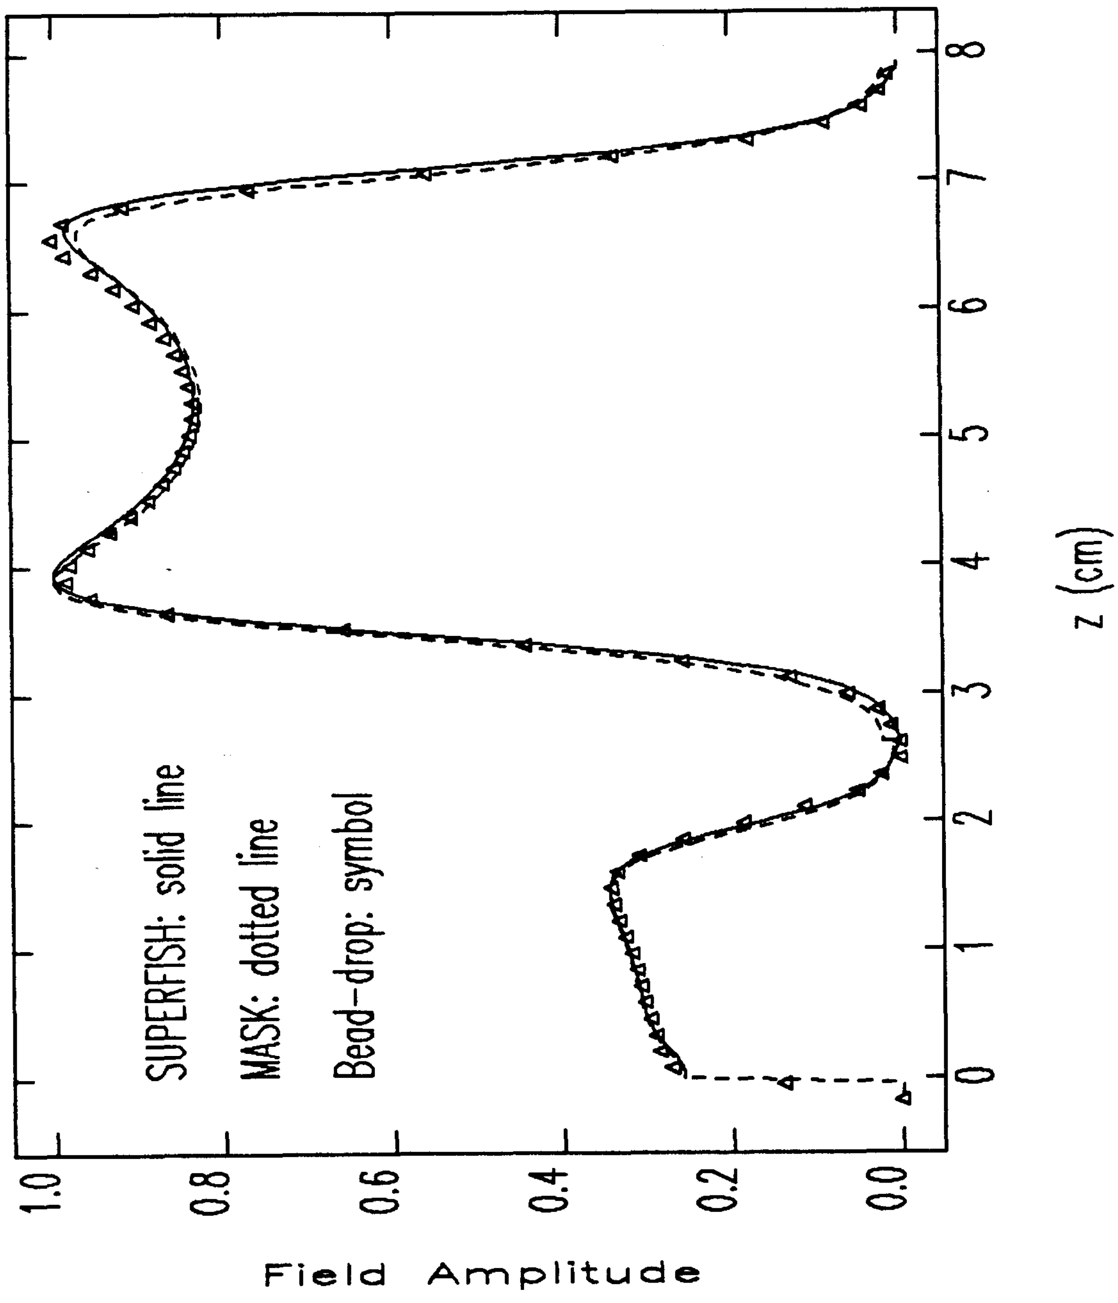

Figure 2.7: Longitudinal Field Profiles from SUPERFISH, MASK, and Bead-Drop Measurement

MASK, along with the results of a bead-drop (or "bead-pull") measurement[60] taken at Varian. Each profile is normalized so that the maximum value is 1. Note the small differences between SUPERFISH and MASK at the cell boundary. These are a result of the boundary conditions used in MASK, as just mentioned. There are also some differences near the electric-field peaks, due to inaccuracy in simulating the cell noses and beam-pipe radius in MASK (a result of the coarseness of the grid).

2.2.2 Gun Cavity Parameters

While SUPERFISH gives parameters for the individual cells directly, it does not give results for the gun as a whole. I will digress briefly to show to obtain such results. These will prove useful in analysis of experimental results in Chapter 5. The definition of the Q for either cell is[56]

$$\mathrm{Q_i} = \frac{\omega \mathrm{U_i}}{\mathrm{P_i}}, \tag{2.4}$$

where $\mathrm{U_i}$ is the stored energy and $\mathrm{P_i}$ is the power dissipated in the cavity walls, for the $\mathrm{i^{th}}$ cell, where i is 1 or 2 for the first or second on-axis cell (I ignore the coupling cell, since there is no field stored in it in the $\pi/2$ mode, as discussed earlier in this chapter). The stored energy may be expressed in terms of the peak, on-axis electric field as

$$\mathrm{U_i} = \mathrm{K_i} \cdot \mathrm{E_{pi}^2}, \tag{2.5}$$

where the $\mathrm{K_i}$ are constants that can be deduced from SUPERFISH output (see Table 2.1). Since $\mathrm{E_{p1}} = \mathrm{E_{p2}}/\alpha$, where $\alpha \approx 3$ is a constant, the total stored energy is

$$\mathrm{U} = \mathrm{U_1} + \mathrm{U_2} = \mathrm{U_2}(1 + \frac{\mathrm{K_1}}{\mathrm{K_2}\alpha^2}). \tag{2.6}$$

This form is convenient, since if $\mathrm{K_1}$ and $\mathrm{K_2}$ are of the same order, as one expects, then U is dominated by $\mathrm{U_2}$. Similarly, the total power lost in the cavity walls is

$$\mathrm{P} = \mathrm{P_1} + \mathrm{P_2} = \mathrm{P_2}(1 + \frac{1}{\alpha^2}\frac{\mathrm{K_1 Q_2}}{\mathrm{K_2 Q_1}}) \tag{2.7}$$

Hence, the predicted Q of the cavity in the $\pi/2$ mode is

$$\mathrm{Q} \;=\; \omega\frac{\mathrm{U_2}}{\mathrm{P_2}}\frac{1 + \frac{\mathrm{K_1}}{\mathrm{K_2}\alpha^2}}{1 + \frac{1}{\alpha^2}\frac{\mathrm{K_1 Q_2}}{\mathrm{K_2 Q_1}}} \tag{2.8}$$

$$= \quad Q_2 \frac{1 + \frac{K_1}{K_2 \alpha^2}}{1 + \frac{1}{\alpha^2} \frac{K_1 Q_2}{K_2 Q_1}},$$

from which one sees that for K_1 and K_2 of the same order and $\alpha^2 \gg 1$, then $Q \approx Q_2$. Table 2.2 lists the results of cold-test measurements performed at Varian and along with inferred properties for the gun, based on this analysis. For some of the inferred properties, I have used the values of K_1, K_2, and Q_1/Q_2 calculated with SUPERFISH. I have also listed certain "desired" values, along with predictions based on SUPERFISH results, with $\alpha = 2.9$ (the measured value) used where necessary.

I have calculated the peak electric field in the second cell as a function of total wall power using

$$E_{p2} = \sqrt{\frac{U_2}{K_2}} = \sqrt{\frac{Q_2}{\omega K_2}} \sqrt{P_2}, \tag{2.9}$$

or (using 2.7)

$$E_{p2} = \sqrt{\frac{Q_2}{\omega K_2 (1 + \frac{1}{\alpha^2} \frac{K_1 Q_2}{K_2 Q_1})}} \sqrt{P}. \tag{2.10}$$

Table 2.2: **Measured and Desired Cavity Parameters for the RF Gun**

quantity	desired (predicted)	measured (inferred)	unit
Q for $\pi/2$-mode	(18008)	14000	
frequency at 20° C, air			
0-mode	—	2922.975	MHz
$\pi/2$-mode	2855.8	2855.835	MHz
π-mode	—	2802.960	MHz
$\alpha = E_{p2}/E_{p1}$	2–3	2.9	
β	4	4.15	
$E_{peak,2}/\sqrt{P_1 + P_2}$	(79.9)	(70.5)	$MV/m/MW^{\frac{1}{2}}$

2.2.3 Methodology of MASK Simulations

I return now to the discussion of simulations, and in particular MASK simulations. My methodology in using MASK was heavily influenced by the need to economize computer time. A single cycle of the RF gun with beam takes approximately 30 hours of CPU on SSRL's VAX 8810. Part of the reason for this is the need to use a large number of simulation macro-electrons. The emission algorithm in MASK is such that charge is emitted at each step in order to obtain the desired current density. Hence, if the macro-electron weight (the number of electrons that each macro-electron represents) is made too large in an attempt to decrease the number of macro-electrons, the simulation may end up emitting no macro-electrons at all, because one macro-electron per time-step may exceed the allowed current density. Thus, having many simulation particles is a result in part of having a small time-step. The time-step, Δt, is chosen under the constraint of the Courant stability condition [61] (for integration of Maxwell's equations), which requires that

$$\Delta t \leq \frac{\mathrm{minimum}(\Delta r, \Delta z)}{c\sqrt{2}}, \tag{2.11}$$

where Δr and Δz are the grid spacings in r and z, respectively. In order to accurately simulate the fields in the vicinity of the cathode, I chose $\Delta r = 0.25$mm, which gives 12 grid points across the cathode and 3 grid points spanning the recess between the cathode and focusing annulus. $\Delta z = 0.43$mm was also chosen, based on the need to have 64 grid points between the recess around the cathode and the end of the first cell (the power of 2 is required by MASK.) Hence, the time step would need to be less than about 0.6 ps. I found that a smaller time-step was needed in order to get stability (perhaps a result of the fact that MASK uses single-precision), and chose a conservative value of 0.171 ps, which is convenient in that it gives 2048 time-steps per 2856 MHz RF period.

As I will discuss below, the longitudinal mesh spacing was belatedly discovered to be somewhat larger than needed to accurately fulfill the boundary conditions on the slope of E_z at the cathode. However, this has no significant effect on the results. Test runs with smaller radial and longitudinal mesh sizes were found to give virtually identical results to those with the mesh sizes listed in the last paragraph.

I then chose the ratio of the macro-electron weight to the current density in order to ensure that several macro-electrons were emitted per time step. The number of macro-electrons emitted per time step is

$$N_s = \frac{J\pi R_c^2}{eN_m}\Delta t, \tag{2.12}$$

where N_m is the number of electrons represented by each macro-electron. In order to obtain good statistics, it is desirable that there be several thousand macro-electrons in the simulated useful beam. For examining the results of filtering small momentum spreads from the beam, it is necessary to have even more macro-electrons. Roughly speaking, N_u useful-beam macro-electrons requires $2N_u$ macro-electrons total exiting the second cell, which requires roughly $4N_u$ macro-electrons emitted from the cathode during the accelerating phase in the first cell (the reasons for these factors will be discussed below). This analysis would lead one to conclude that, for 2048 time-steps per cycle, $N_s = N_u/256$. I chose $N_s = 15$, which gives about 4000 useful-beam macro particles and gives good statistics even for analysis of small momentum intervals. In retrospect, this is probably higher that it needs to be, but I have used in throughout in order to avoid changing simulation parameters which would confuse comparison of different MASK runs.

Because of the long running times, I decided to simulate only the steady-state behavior of the gun, that is, the behavior after the RF fields and beam current have come to stability. This happens in the last half of the 2μs RF pulse used at SSRL. Even taking $N_s = 1$ would not help to decrease the running time sufficiently to allow simulation of the entire operating cycle of the gun in a reasonable time. There are a number of assumptions upon which the validity of this procedure rests. It assumes that it is sufficient to simulate only the $\pi/2$ structure mode, which allowed simulation of the cells separately. If, for example, the beam drives significant power into the zero or π modes, this approximation would be invalid. Since MASK cannot simulate the coupling cell, there was little choice about this.

This methodology also assumes that each bunch sees the cavity with only the fundamental mode excited—i.e., that the higher-order-mode fields excited by previous bunches have no significant influence on any particular bunch. This is equivalent to saying that the only significant effect of previous bunches is to remove power from the fundamental cell mode. In the steady-state, this power is assumed to be replaced by the RF power source, so that each bunch sees that same fields; hence, if the assumptions are valid, only one cycle needs to be simulated with particles. Similarly, this methodology assumes that electrons left in the gun from the previous operating cycle have no significant influence on newly emitted particles. These assumptions can be checked by simulating the gun for several operating cycles; results of such a simulation are present in the next section. Normally, however, I simulate only a single operating cycle with beam, in addition to simulating several RF cycles for the build-up of the cell fields.

Since MASK is a time-dependent code, the cavity fields must be built up by means of some suitable simulated RF power source. The program allows the simulation of both RF ports and antennae. I chose to use an antenna because of the greater simplicity. Since I wished to simulate only the steady-state behavior of the gun, it made no difference how the cavity was driven. (I was not attempting to simulate the evolution of the beam during the charging of the cavity.) In driving an RF current in the antenna, MASK creates fields in the vicinity of the antenna that propagate throughout the cavity via Maxwell's equations. The sinusoidal RF current was modulated by the envelope shown in Figure 2.8, in order to avoid excitation of higher-order modes (the particular shape is composed of two cubic splines, using a standard feature of MASK). In this way, the Fourier amplitudes excited in the first three higher-order modes were kept below 10^{-3} of the fundamental. In addition, the antenna was placed at the intersection of nodes of the first two higher-order modes, to reduce the excitation of these modes even further (shaping the current envelope is by far the most important consideration).

In general, a simulation of the gun with MASK consisted of the following steps:

1. The fields in the first cell were excited to the desired level, or else the fields saved on disk from a previous run were read in (and optionally scaled) . Once the driving current envelope has fallen to zero, the fields can be saved for use in subsequent runs, or used immediately in the next stage.

2. Particle emission from the cathode, limited to some fixed current density and by space-charge forces, was then allowed for one RF period, beginning during the accelerating phase of the fifth RF period since the excitation began. During this period, MASK "pushes" the macro-electrons through the first cell, in addition

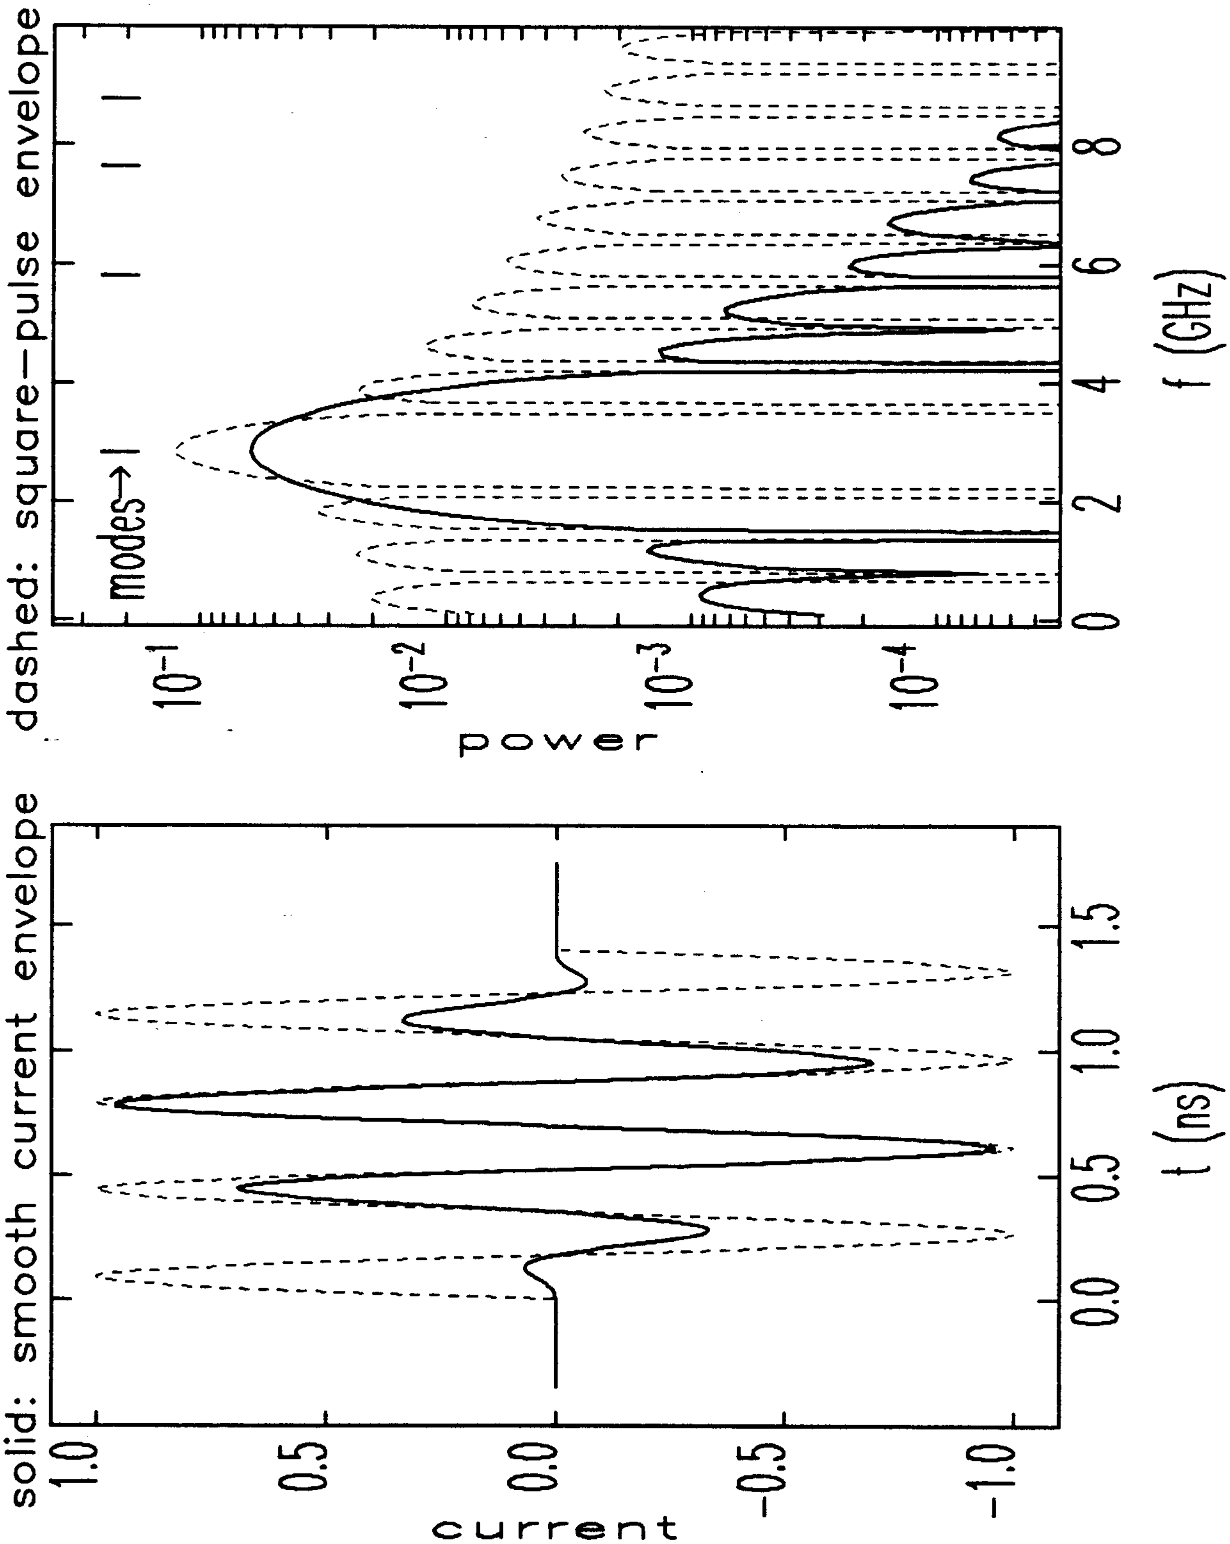

Figure 2.8: RF Current Waveform for Exciting Cells in MASK

to simultaneously integrating Maxwell's equations. As macro-electrons pass $z = \lambda/4$, the end of the first cell, their phase-space coordinates are saved to disk.

3. The fields in the second cell were excited to the desired level or read in from disk, as for the first cell. Of course, these fields should be 180° out of phase with the first-cell fields. (In fact, they were *driven* 178° out of phase, to compensate for the difference in frequency, as discussed below.)

4. Macro-electrons from the simulation of the first cell were injected at $z = \lambda/4$ in the second cell. Since both simulations start at time $t = 0$, each macro-electron is injected into the second simulation when the time counter is the same as it was when the macro-electron left the first simulation.

The measured field ratio in the gun is $\alpha = 2.9$. The MASK simulations reported on here used $\alpha = 3.0$, as a result of a mis-reading of the data that was not noticed until the simulations had already been run. Because the simulations take so long, time does not permit me to repeat them. I will show below that this error has only a small effect on the results. Given the many other approximations being made (e.g., independent bunches, ignoring the three-dimensional nature of the cavity with coupling holes, having the simulation cell frequency differ from the actual frequency), this is not a serious error.

2.2.4 Cathode Simulation

Cathode emission was simulated by injecting a total charge of $\pi R_c^2 J_{emitted} \Delta t$ during each time interval, where R_c is the cathode radius, $J_{emitted}$ the emitted current density, and Δt the time interval (i.e., the time-step for the integration). The space-charge limitation on the current density was assumed to follow the Gaussian emission law (a standard feature of MASK), namely,

$$-\left(J_{emitted}\Delta t + Q_{surf}\right) \leq -\epsilon_o E_{z,cathode} \tag{2.13}$$

where $E_{z,cathode}$ must be negative for the emission of electrons (negative J), and where $Q_{surf} \leq 0$ is the existing "surface charge" (i.e., charge just above the cathode surface) due to previously emitted macro-electrons. In other words, the emitted charge density cannot be so high that, in combination with previously emitted charge density, the newly emitted charge depresses the field at the cathode so much that the sign of the field is such as to push electrons back into the cathode. In addition I required that

$$-J_{emitted} \leq -J_{limit}, \tag{2.14}$$

where $|J_{limit}|$ is the user-specified maximum current density available from the cathode, which in actual operation is controlled by controlling cathode temperature. In addition to adding this feature, I modified MASK to emit a bona-fide uniform random distribution from the circular cathode surface (the standard code uses an approximation to such a distribution). Each macro-electron emitted in the simulations has the same charge, in order to make the interpretation of graphs and other data more straight-forward. This implies that in order to have a uniform emission over the surface, fewer macro-electrons should be emitted near $r = 0$ than near $r = R_c$. The reader will see this effect in some of the graphs that follow.

It is interesting to look at the details of this emission process. SUPERFISH calculations give the peak field at the cathode as $E_{pc} = 0.258E_{p2}$ (assuming $\alpha = 2.9$). Hence, for a typical operating value of $E_{p2} = 75$ MV/m, E_{pc} is about 20 MV/m. As the field at the cathode passes into the accelerating phase for electrons, emission begins. Around this time (call it t=0), the field at the cathode may be approximated as $E_c \approx -E_{pc}\omega t$. Hence, if the integration time step is Δt, and if I momentarily assume $J_{limit} \rightarrow -\infty$, then

$$J_{emitted}(t = n\Delta t)\Delta t = -Q_{surf}(t = n\Delta t) - \epsilon_o E_{pc}\omega n\Delta t \tag{2.15}$$

where n is the number of time steps since t=0. Emission can occur as long as $J_{emitted} < 0$.

One expects that, initially, the surface charge will simply be the sum of all charge emitted since t=0, since the fields have not had time to accelerate that charge away from the cathode. It is possible to approximate the distance a macro-electron will move and confirm this expectation. Ignoring the variation of the field with position, the Lorentz equation gives

$$\frac{dp}{dt} \approx \frac{eE_{pc}\omega t}{m_e c}, \tag{2.16}$$

where $p \equiv \beta\gamma$. For a macro-electron that starts from the cathode at t=0,

$$p(t) \approx \frac{eE_{pc}\omega t^2}{2m_e c}. \tag{2.17}$$

For a non-relativistic particle, the velocity is approximately pc, so that the position is

$$z(t) \approx \frac{eE_{pc}\omega t^3}{6m_e}. \tag{2.18}$$

For $E_{pc} = 20\mathrm{MV/m}$,

$$z(t) \approx 10^{-8}t^3, \tag{2.19}$$

where z is in meters and t in pico-seconds. The cathode radius, $R_c = 3\mathrm{mm}$, provides a natural length scale, and one sees that it will take about 70 ps for the first macro-electron to move by R_c. In other words, the effects of space-charge on emission might be expected to be quite large, since it takes so long for emitted charge to be moved away from the cathode.

Since the charge is moved away from the cathode so slowly, one can assume that the surface charge is just the sum of all previously emitted charge:

$$Q_{surf}(t = n\Delta t) = \sum_{i=1}^{n-1} J_{emitted}(t = i\Delta t)\Delta t, \tag{2.20}$$

so that

$$J_{emitted}(t = n\Delta t) = -\sum_{i=1}^{n-1} J_{emitted}(t = i\Delta t) - \epsilon_o E_{pc}\omega n. \tag{2.21}$$

Subtracting this expression from the same expression with $n \to n-1$, one obtains

$$J_{emitted}(t = n\Delta t) = -\epsilon_o E_{pc}\omega. \tag{2.22}$$

For $E_{pc} = 20\mathrm{MV/m}$, one obtains $J_{emitted} = -318\mathrm{A/cm}^2$. Since the cathode is capable of no more than 140 $\mathrm{A/cm}^2$, we see that the space charge limitation will never come into play. I will hereafter use $J = |J_{limit}|$ to refer to the cathode current density assumed in any simulation.

MASK supports the inclusion of a initial thermal velocity distribution for the macro-electrons emitted from the cathode. I found that inclusion of this effect made no significant difference in the results. To see why, consider that the RMS thermal momentum ($\beta\gamma$) is given by the Maxwell velocity distribution[62] as

$$\mathrm{p_{rms}} = \sqrt{\frac{3\mathrm{kT}}{\mathrm{m_e c^2}}}. \tag{2.23}$$

A cathode temperature of 1000°C or 1273°K gives $\mathrm{p_{rms}} \approx 10^{-3}$. The RMS radial momentum will be of the same order. During acceleration, this radial momentum be adiabatically damped. Assuming acceleration to a typical value of $\mathrm{p_z} = 5$, the slope due to the thermal velocity would be of order $\mathrm{p_{rms}/p_z} \approx 0.2$mrad. A typical angular half-width for an RF gun beam is 20 mrad, and hence one sees that thermal effects may safely be ignored. This is confirmed by MASK simulations done with and without thermal effects. While adding thermal effects to MASK entails little cost in computer resources, I have ignored thermal effects in all the MASK simulations reported on here, since this permits cleaner comparison with **rfgun**, which did not originally support an initial thermal velocity distribution.

A number of different types of output are available from MASK. I have upgraded MASK extensively to improve the ease with which one can make use of these outputs; in particular, I have used my self-describing data format (**awe** format) extensively. This is discussed in more detail in Appendix A.

MASK allows one to sample the cavity fields along any line in r or z at intervals in time, as well as to sample the fields at a fixed location as a function of time. The initial coordinates of macro-electrons emitted from the cathode may be saved, as well as the phase-space coordinates (time, positions, and velocities) of every macro-electron as it crosses a number of user-defined planes of constant z coordinate. This later capability is the principle means of getting information about beam energy and emittance, and for keeping track of particles that return to hit the cathode. It is also used for the re-injection of macro-electrons from first-cell simulations into second-cell simulations. MASK keeps a record of particles lost by hitting metal boundaries or exiting the simulation region, so that the total power going into particles can be calculated. This is important for assessing how much power is needed to produce a given beam, since not all the power that goes into electrons goes into the electrons *in the beam.* MASK can also produce output containing "beam-snapshots", which record the positions of all particles at a given time during the simulation. Examples of all of these output facilities will be seen in the next section.

2.2.5 Compensation of Cell Frequency Mismatch

I mentioned above that the phase-shift between the driving RF currents for the two cells must be different from 180° to compensate for the difference in resonant frequencies between the first and second cells[63]. The first simulated cell resonates at 2834.5 MHz, and the second at 2837.8 MHz. Particles from the simulation of cell 1 are injected into the simulation at cell 2 during the fifth RF cycle since the beginning of the RF excitation. Hence, one wants the fields in the cells to be 180° out of phase during the fifth RF cycle. The phase in each cell is

$$\phi_i = \omega_i t + \Delta\phi_i, \tag{2.24}$$

where $\Delta\phi_1 = 0$ and $\Delta\phi_2 \approx 180°$. Setting $t = 5T_o$, with T_o the average RF period of the cells, one obtains the phase difference at the beginning of the fifth RF period

$$\phi_2 - \phi_1 = 5T_o(\omega_2 - \omega_1) + \Delta\phi_2, \tag{2.25}$$

requiring $\phi_2 - \phi_1 = \pi$, one obtains the necessary value of $\Delta\phi_2$

$$\Delta\phi_2 = \pi - 5T_o(\omega_2 - \omega_1). \tag{2.26}$$

Using the values of the simulated cell frequencies that I gave just above gives $\Delta\phi_2 = 3.105$, or 177.9°. I used this value in of my all MASK simulations of the gun as built. I shall use 2836 (the average frequency) as the nominal frequency of the MASK simulations. In the next section, I shall show what the effect of this frequency error is.

2.2.6 The `rfgun` Program

Having dealt with MASK at some length, I now turn to **`rfgun`**. As I have mentioned above, **`rfgun`** uses the longitudinal field profiles generated by SUPERFISH (or any equivalent program) to calculate particle motion in the gun. **`rfgun`** works in cartesian coordinates (x, y, z), allowing the imposition of both transverse and solenoidal external magnetic fields. The cavity fields are assumed to have a perfect sinusoidal variation in time, at some user-specified frequency. (Since the fields for the whole cell are specified at once, there is no need to worry about the individual cell frequencies.) Macro-electrons are emitted from the cathode at discrete phase-intervals and from discrete radial positions on the cathode, with macro-electron charges adjusted to achieve approximately uniform current density. The solution for macro-electron motion proceeds from the Lorentz equation, which is integrated numerically using fourth-order Runge-Kutta[61]. Like MASK, **rfgun** supplies the user with phase-space coordinates of exiting particles as well as those of particles that return to hit the cathode. Unlike MASK, **rfgun** does not include space-charge, nor does consider the positions of metal structures that macro-electrons might hit (a maximum radius may be specified, however, to simulate the collimating effect of the beam tube).

2.2.7 Off-Axis Field Expansion

It is well known[31] that the modes for a cavity like one of the RF gun cells can be separated into independent transverse-electric (TE) and transverse-magnetic (TM) modes. TM modes have $B_z = 0$ and involve only E_z, E_r, and B_ϕ, whereas TE modes have $E_z = 0$ and involve only B_z, B_r, and E_ϕ. While the gun cavity can support TE modes, the fundamental mode is TM, which has a non-zero accelerating field E_z. (Note that a cylindrically symmetric beam will induce only TM modes, since it always produces $E_r \neq 0$ and $B_\phi \neq 0$, but never $E_\phi \neq 0$ or $B_r \neq 0$.)

Knowing E_z along the axis allows one to compute series expansions in r for E_z, E_r, and B_ϕ. One starts with Maxwell's equations (in MKSA units):

$$\nabla \cdot \mathbf{E}(z,r,t) = 0 \tag{2.27}$$

$$\nabla \times \mathbf{E}(z,r,t) = -\partial_t \mathbf{B}(z,r,t) \tag{2.28}$$

$$\nabla \times \mathbf{B}(z,r,t) = \mu_o \epsilon_o \partial_t \mathbf{E}(z,r,t). \tag{2.29}$$

Expressing these in cylindrical coordinates (r, ϕ, z) and assuming that the electric fields vary like $\sin(\omega t)$ (and hence that $B_\phi \sim \cos(\omega t)$), one obtains

$$\partial_r (r E_r(z,r)) + r \partial_z E_z(z,r) = 0, \tag{2.30}$$

$$\partial_z E_r(z,r) - \partial_r E_z(z,r) = \omega B_\phi(z,r), \tag{2.31}$$
$$\partial_r(rB_\phi(z,r)) = r\mu_o\epsilon_o\omega E_z(z,r). \tag{2.32}$$

and

$$-\partial_z B_\phi(z,r) = \mu_o\epsilon_o\omega E_r(z,r), \tag{2.33}$$
$$\tag{2.34}$$

Because of cylindrical symmetry, $E_r(z,r)$ must be an odd function of r, and hence from equation (2.30) one sees that $E_z(z,r)$ must be an even function of r. Equation (2.33) implies that $B_\phi(z,r)$ has the same symmetry as $E_r(z,r)$. Thus, I can expand these functions as

$$E_z(z,r) = \sum_{i=0}^{\infty} \tilde{E}_{z,2i}(z) r^{2i} \tag{2.35}$$
$$E_r(z,r) = \sum_{i=0}^{\infty} \tilde{E}_{r,2i+1}(z) r^{2i+1} \tag{2.36}$$
$$B_\phi(z,r) = \sum_{i=0}^{\infty} \tilde{B}_{\phi,2i+1}(z) r^{2i+1}, \tag{2.37}$$

where the tildes are used to emphasize that the coefficients do not necessarily have the same units as the fields.

Inserting expressions (2.35) and (2.36) into the divergence equation, (2.30) and equating terms with the same power of r, one obtains

$$\tilde{E}_{r,2i+1}(z) = -\frac{1}{2i+2}\partial_z \tilde{E}_{z,2i}(z) \quad i = 0,1,\ldots\infty. \tag{2.38}$$

Inserting expressions (2.35) and (2.37) into (2.32), one obtains

$$\tilde{B}_{\phi,2i+1}(z) = \frac{\mu_o\epsilon_o\omega}{2i+2}\tilde{E}_{z,2i}(z) \tag{2.39}$$

Using equations (2.38) and (2.39) in (2.31) one obtains

$$\sum_{i=0}^{\infty}\left[-\frac{r^{2i+1}}{2i+2}\left(\partial_z^2 + k^2\right)\tilde{E}_{z,2i}(z) - 2ir^{2i-1}\tilde{E}_{z,2i}(z)\right] = 0, \tag{2.40}$$

where $\omega = kc$. Equating terms of the same power in r yields

$$\tilde{E}_{z,2i+2}(z) = -\frac{1}{(2i+2)^2}\left(\partial_z^2 + k^2\right)\tilde{E}_{z,2i}(z), \tag{2.41}$$

which completes the solution for the radial expansion coefficients by providing a recursion relation that gives all the $\tilde{E}_{z,2i}$'s, from which all the $\tilde{E}_{r,2i+1}$'s and $\tilde{B}_{r,2i}$'s can be obtained via equations (2.38) and (2.39).

It is useful to work out these expressions for the first few terms:

$$\tilde{E}_{r,1}(z) = -\frac{1}{2}\partial_z \tilde{E}_{z,0}(z), \tag{2.42}$$

$$\tilde{E}_{r,3}(z) = \frac{1}{16}\left[\partial_z^3 + k^2\partial_z\right]\tilde{E}_{z,0}(z), \tag{2.43}$$

$$\tilde{E}_{z,2}(z) = -\frac{1}{4}\left[\partial_z^2 + k^2\right]\tilde{E}_{z,0}(z), \tag{2.44}$$

$$\tilde{B}_{\phi,1}(z) = \frac{\omega}{2c^2}\tilde{E}_{z,0}(z), \tag{2.45}$$

and

$$\tilde{B}_{\phi,3}(z) = -\frac{\omega}{16c^2}\left[\partial_z^2 + k^2\right]\tilde{E}_{z,0}(z). \tag{2.46}$$

`rfgun` includes terms up to third order in r, with the option to use only the linear terms in r. This allows investigation of the importance of higher-order terms. Note that $\tilde{E}_{z,0}(z) = E_z(z, r = 0)$, where values of the latter function are given by SUPERFISH. Since it has the units of an electric field, I'll drop the tilde from $\tilde{E}_{z,0}(z)$ in what follows.

2.2.8 Non-Linear Field Terms

To see how influential the nonlinear terms might be, note that the importance of the nonlinear magnetic fields is related to their magnitude relative to the linear terms, which is characterized by the function

$$T_1(z) = \frac{R_c^2}{8E_{p2}}\left[\partial_z^2 + k^2\right]E_{z,0}(z), \tag{2.47}$$

where R_c is the cathode radius, which gives an upper limit on the radius of any particle that makes it out of the gun (as will be seen in the next section). Note that the average radius of particles emitted from a uniformly-emitting circular cathode is $2R_c/3$. If $T_1(z)$ is small compared to $E_{z,0}(z)/E_{p2}$, then nonlinear magnetic fields are

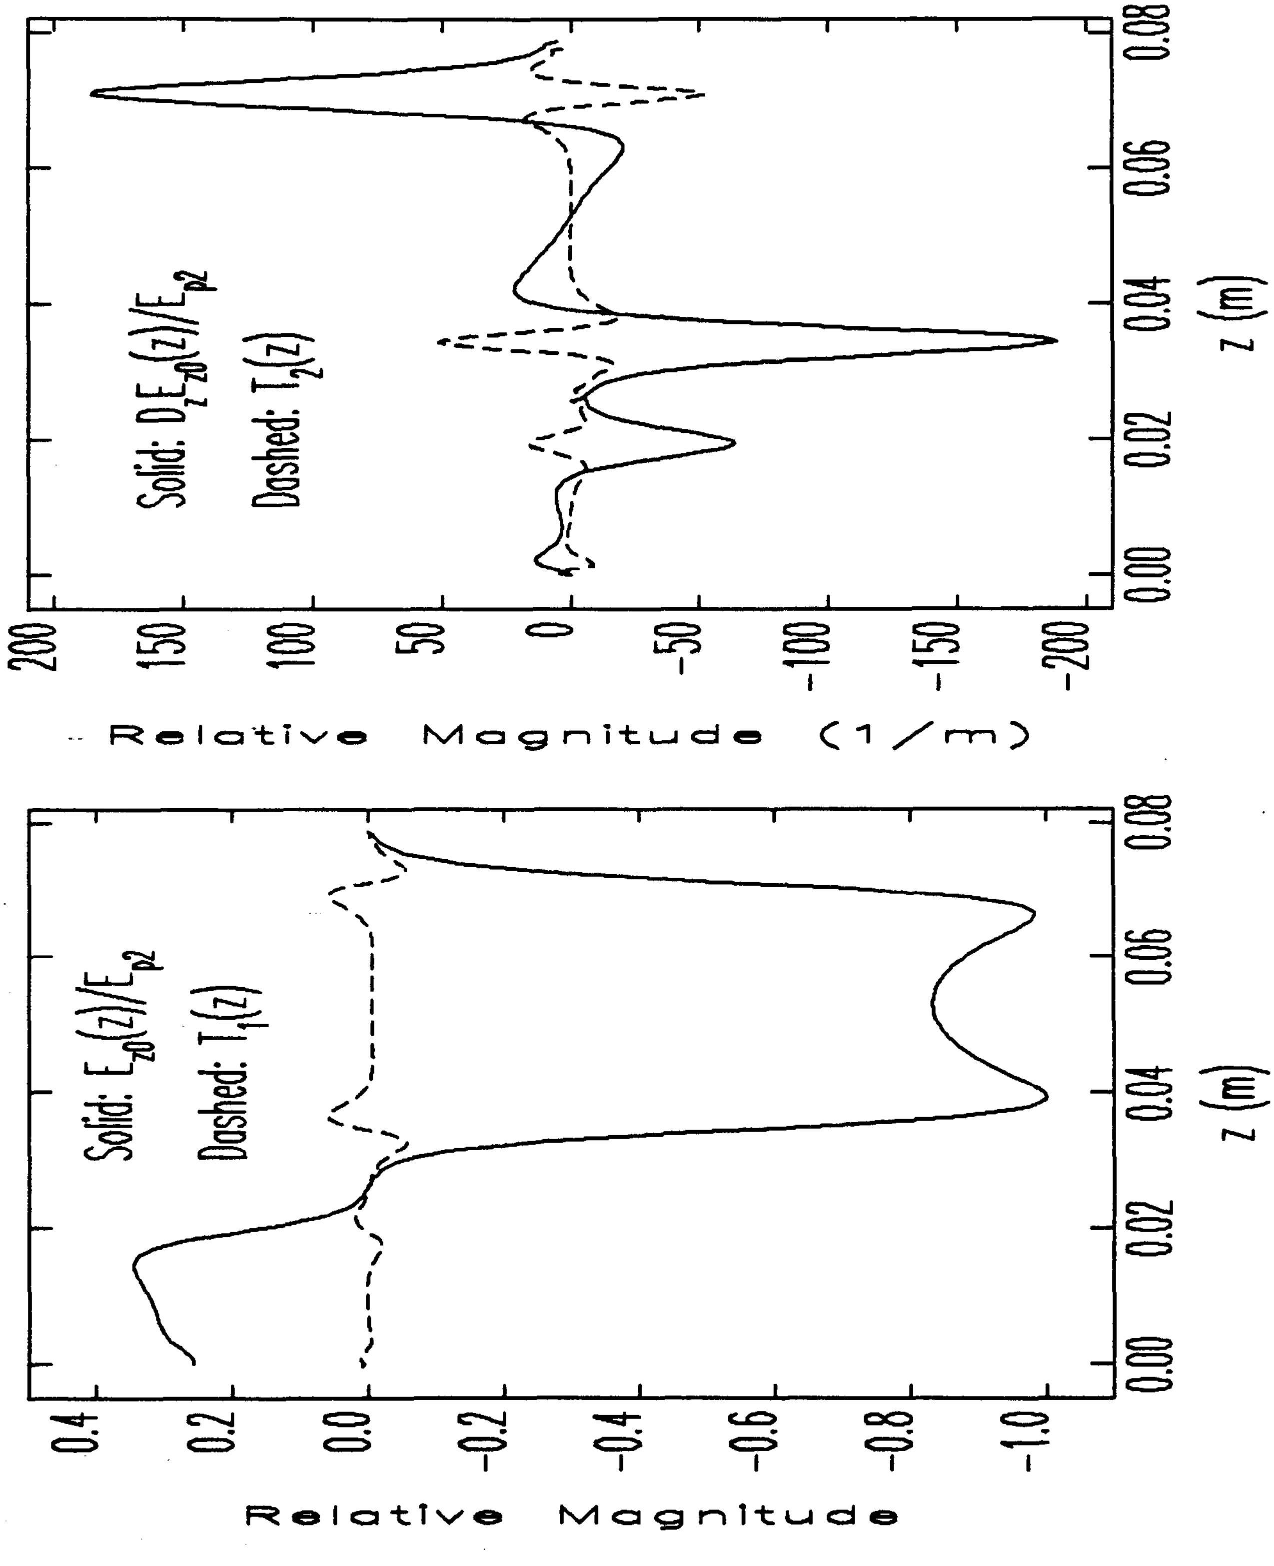

Figure 2.9: Nonlinear Field Terms in the RF Gun

unimportant. $T_1(z)$ is also related to the importance of the nonlinear longitudinal electric fields, though $4T_1(z)$ is a more accurate measure.

Similarly, importance of the nonlinear radial electric fields is characterized by the function

$$T_2(z) = \frac{R_c^2}{8E_{p2}} \left[\partial_z^3 + k^2\partial_z\right] E_{z,0}(z). \tag{2.48}$$

If $T_2(z)$ is small compared to $\frac{\partial}{\partial z}E_{z,0}(z)$, then nonlinear radial electric fields are unimportant.

Figure 2.9 compares $T_1(z)$ and $T_2(z)$ to the relevant lower-order terms. One sees that the non-linear magnetic fields are quite small, and one infers that the non-linear longitudinal electric fields are small as well. The non-linear radial electric fields are more significant, particularly in the vicinity of the cathode and near the cell noses. Of these, the fields near the cell noses are less important, since most particles of interest pass the cell noses well off the RF crest.

Finally, I note that $T_1(z)$ is completely dominated by the $\partial_z^2 E_{z,0}(z)$ term, while $T_2(z)$ is completely dominated by the $\partial_z^3 E_{z,0}(z)$ term. This is as expected, since the fields (and their derivatives) are changing rapidly on the scale of $1/k \approx 1.7$cm. I shall return to the issue of non-linear fields, specifically their effect on the beam, in the next section. For now, I simply point out that if one wants to reduce the effect of non-linearities, one must reduce $\partial_z^2 E_{z,0}(z)$ and $\partial_z^3 E_{z,0}(z)$. In particular, it is important to avoid non-linear terms in z near the cathode, where the beam is particularly sensitive (because of its relatively large radial size and its low momentum). This was not attempted for the SSRL RF gun design.

2.2.9 Boundary Conditions for `rfgun`

The data shown in Figure 2.9 was assembled from SUPERFISH output for the individual cells. SUPERFISH gives values of E_z along the axis at equi-spaced intervals in z. However, these values are calculated with boundary conditions for the individual cells that do not properly reflect the boundary conditions appropriate to the full cavity, as discussed above. Hence, the two solutions do not join smoothly at $z = \lambda/4$. One is tempted to smooth the solutions in order to eliminate the discontinuities in the fields and its derivatives at this point, but I have found that this can be done only at the expense of smoothing away real, significant derivative information in the rest of the cavity. Hence, the data in Figure 2.9 as well as that used in RF gun simulations is simply composed from concatenating the solutions for the individual cells, with appropriate scaling to get $\alpha = 2.9$ and a $\pi/2$-mode configuration.

In order to calculate the non-linear field terms, **rfgun** takes z derivatives of $E_{z,0}(z)$, up to the fourth derivative (so that the third derivative can be interpolated in between data points). Each derivative is calculated using the second-order formula[64]

$$\partial_z F(z) = \frac{F(z+\Delta z) - F(z-\Delta z)}{2\Delta z} + \mathcal{O}(\Delta z^2). \tag{2.49}$$

There is no problem in the interior of the cells, where each point has two neighboring points with well-defined field values, but problems do arise at the boundaries. In order to be able to calculate the required derivatives at the beginning and end of the gun, **rfgun** needs data outside these intervals.

I have provided this data for the beginning of the gun by applying an idea gleaned from the above analysis of non-linear fields. Since z=0 is a metallic surface for $r < R_c$, $E_r(z=0) = 0$ for $r < R_c$, and equation (2.38) implies that

$$\left(\partial_z \tilde{E}_{z,2i}(z)\right)_{z=0,r<R_c} = 0. \tag{2.50}$$

Taking the partial of equation (2.41) with respect to z, one then sees that

$$\left(\partial_z^3 \tilde{E}_{z,2i}(z)\right)_{z=0,r<R_c} = 0. \tag{2.51}$$

Taking ∂_z^3 of (2.41), one then sees that

$$\left(\partial_z^3 \tilde{E}_{z,2i+2}(z)\right)_{z=0,r<R_c} = 0 = \left[-\frac{1}{(2i+2)^2}\left(\partial_z^5 + k^2\partial^3\right)\tilde{E}_{z,2i}(z)\right]_{z=0,r<R_c}, \tag{2.52}$$

and hence that

$$\left(\partial_z^5 \tilde{E}_{z,2i}(z)\right)_{z=0,r<R_c} = 0. \tag{2.53}$$

Proceeding in this fashion, it is clear that

$$\left(\partial_z^{2n+1} \tilde{E}_{z,0}(z)\right)_{z=0,r<R_c} = 0 \qquad n = 0, 1, \ldots \infty, \tag{2.54}$$

and hence that $\tilde{E}_{z,2i}(z)$ is an even function of z. Therefore, one can easily supply **rfgun** with the necessary information by taking $E_{z,0}(-z) = E_{z,0}(z)$.

For the end of the gun ($z = 3\lambda/4$), the boundary conditions are ambiguous, as I discussed above. Since the field amplitude is small here and the beam is at full energy by the time it is in this region, what one does in at this boundary is of little importance. This being so, it is reasonable to simply extend the beam tube beyond the nominal end of the gun in order to obtain information necessary for taking derivatives up to the end of the gun. The primary benefit of doing this is aesthetic, and I have not done the equivalent in MASK, since this would force me to use a coarser mesh in the part of the gun that is really important.

2.3 Simulation Results and Predictions

Having outlined the capabilities of the codes and the methodology of their use, I now present results of my simulations of the SSRL RF gun. In addition to giving predictions of how the actual gun is expected to perform, I also report on simulations done to evaluate design alternatives, to check the assumptions of my methods, to compare the predictions of **rfgun** and MASK, and to evaluate the importance of different effects (e.g., space-charge, non-linear fields) on the beam. Experimental results and comparison of these with simulations will be presented in a Chapter 5.

2.3.1 Effects of the Cell Field Ratio

I will discuss **rfgun** results first, and in particular the effect of varying $\alpha = E_{p2}/E_{p1}$. I have done a series of **rfgun** simulations for a range of α values from 1 to 4. For each series, I varied E_{p2} in 20 MV/m steps from 40 to 120 MV/m (the latter being close to the break-down limit in the second cell).

Longitudinal Phase-Space

Figure 2.10 shows longitudinal phase space results for $\alpha = 1, 2, 3$, and 4. For each value of α, a series of curves appears, one for each value of E_{p2}. Each curve represents the momentum ($p = \beta\gamma$) for macro-electrons as a function of the time of exit from the gun (i.e., the time at which the plane $z = 3\lambda/4$ is crossed). $t = 0$ corresponds to the emission of the first particle from the cathode, which occurs just as the field enters the accelerating phase in the first cell. For these simulations, I directed **rfgun** to emit macro-electrons at 2° phase intervals, the individual macro-electrons being represented by the points on the graphs. The macro-electrons where emitted from r=0, and hence no radial motion was involved.

Note that for $\alpha = 1$, p(t) is non-monotonic over most of the range of E_{p2}, exhibiting a distinctly sinusoidal shape for higher values of E_{p2}. This sinusoidal shape results from the arrival of particles in the second cell ahead of the accelerating crest in that cell. Obviously, if the first particles arrive ahead of the crest, then those that follow will arrive nearer to the crest, and thus gain more momentum. (This issue is

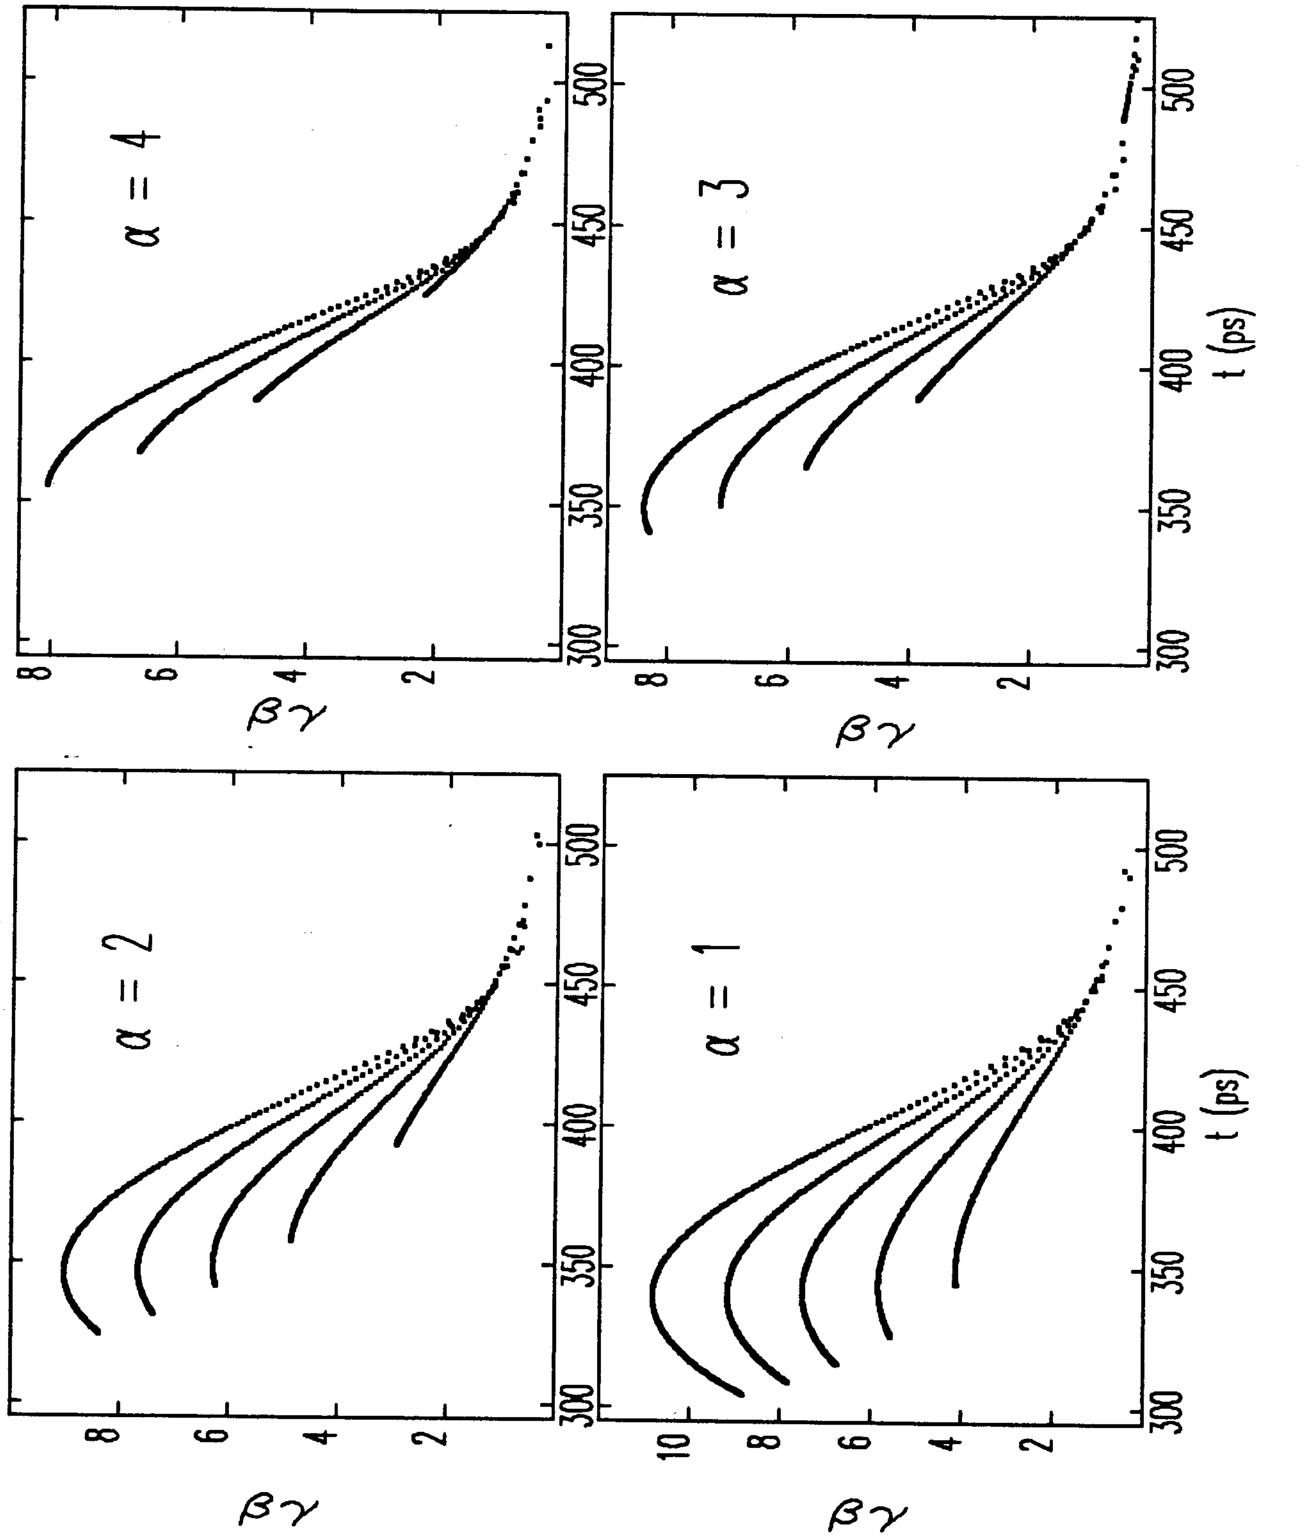

Figure 2.10: Gun Longitudinal Phase Space for Various Values of α

discussed in Chapter 4, as is the need for a monotonic p(t) curve in order to allow the use magnetic of bunch compression.) As α is increased, p(t) becomes monotonic for increasingly higher values of E_{p2}, and also for increasingly high momentum levels. For $\alpha = 4$, p(t) is monotonic for the full range of $E_{p2} > 60\text{MV/m}$, but for $E_{p2} = 40$ the fields in the first cell are too weak to deliver any significant beam to the second cell; hence, few particles exit the gun. For $\alpha = 3$, which is approximately what was achieved in the gun as built, p(t) is reasonably monotonic over $60 < E_{p2} < 100$.

Effective Cathode Area

Figure 2.11 shows the dependence on α and E_{p2} of several beam properties. Each set of connected points corresponds to the value of α indicated by the plotting symbol. The effective cathode area A_{eff} is a measure of how efficiently charge is extracted from the gun. It is defined by

$$A_{\text{eff}} = \frac{I}{J} = \frac{Q}{T_o J}, \tag{2.55}$$

where Q is the total charge exiting the gun during one RF period (i.e., the total charge "in the beam"), T_o is the RF period, and J is the current density. As one might expect, the effective cathode area increases as E_{p2} is increased and as α is decreased. This is due to the increase in the field in the first cell that accompanies both of these changes, which results in more rapid acceleration of charge from the cathode and hence more efficient extraction of beam from the gun. Recall that the physical area of the cathode is $A_c = 0.28\text{cm}^2$. One expects $A_{\text{eff}} < \frac{1}{2}A_c$, since the RF field in the first cell is only in the accelerating phase half the time. A further decrease in A_{eff} results from the fact that not all particles that leave the cathode make it out of the first cell. Many that do not exit the first cell before the RF goes into the decelerating phase will not exit the first cell at all, their momentum being insufficient to overcome the decelerating fields. These particles return to hit the cathode.

Figure 2.11 also shows the normalized beam power and the average momentum in the beam. As expected, both of these increase with increasing E_{p2} and decreasing α (i.e., increasing E_{p1}).

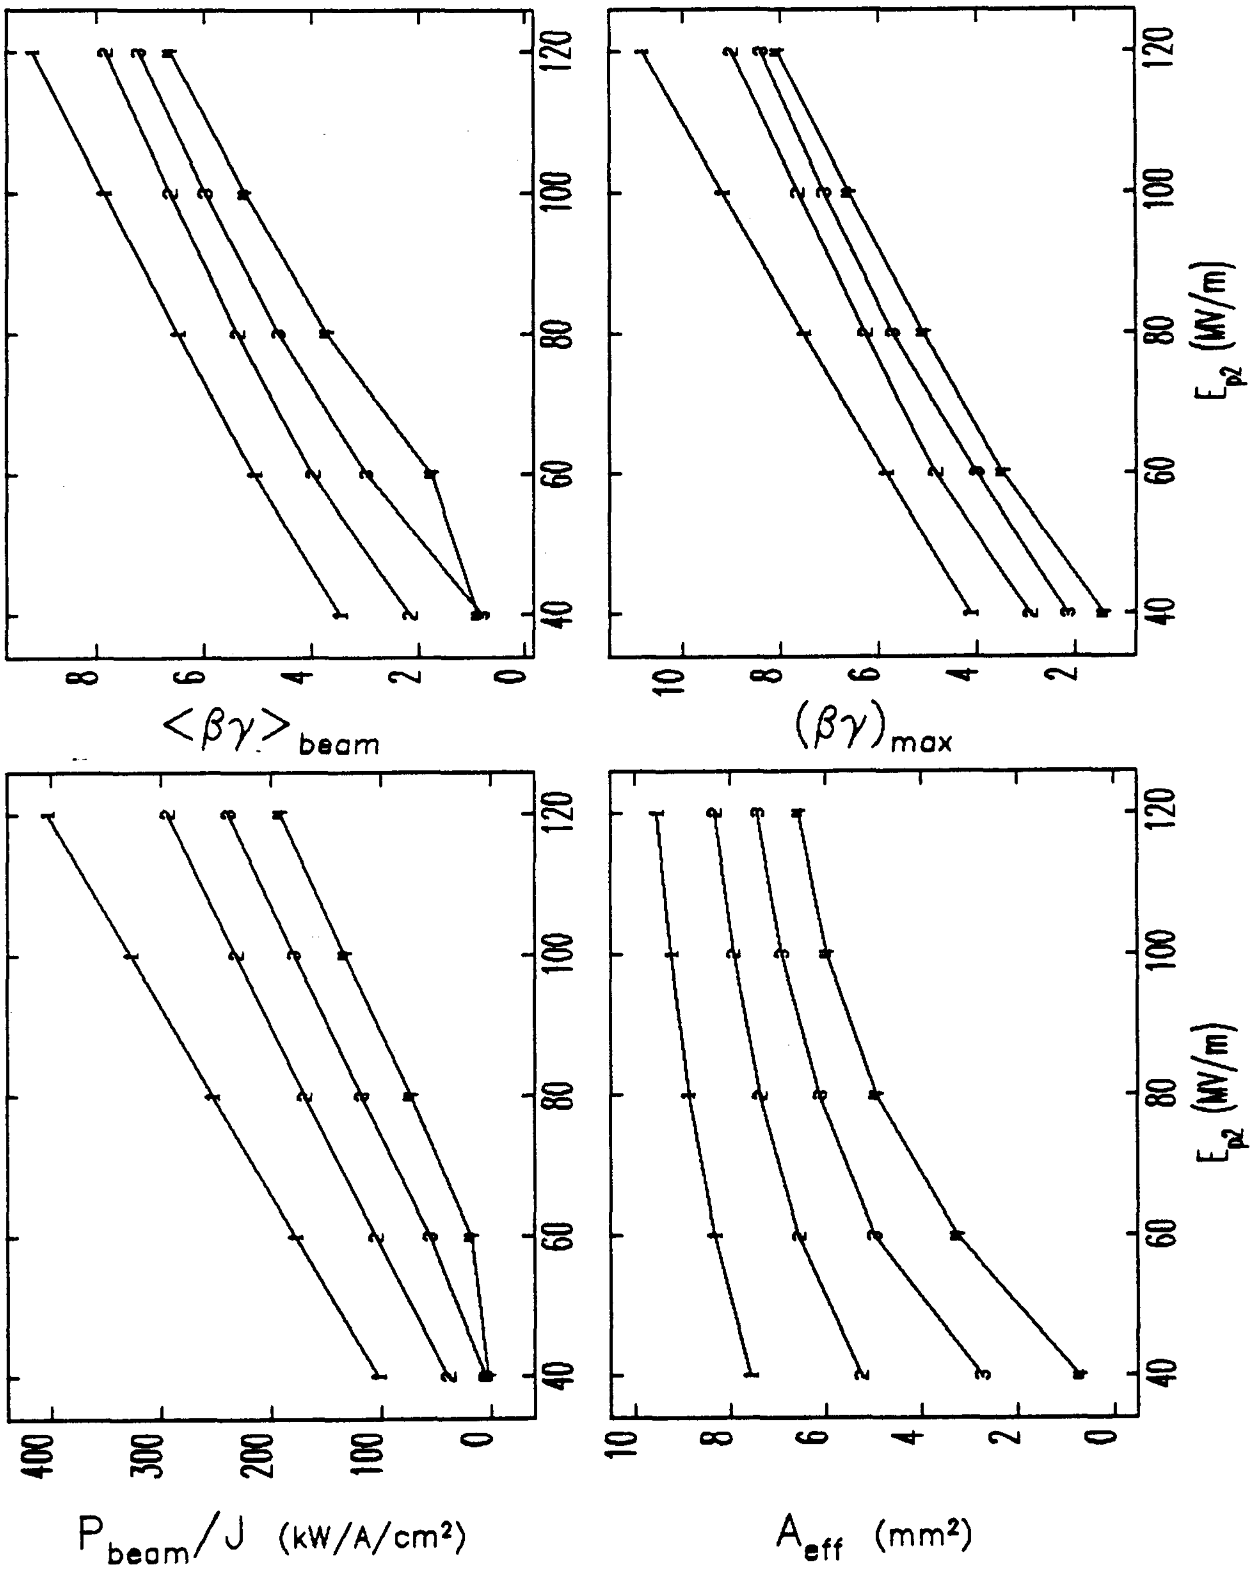

Figure 2.11: Dependence of Beam Properties on α and E_{p2}

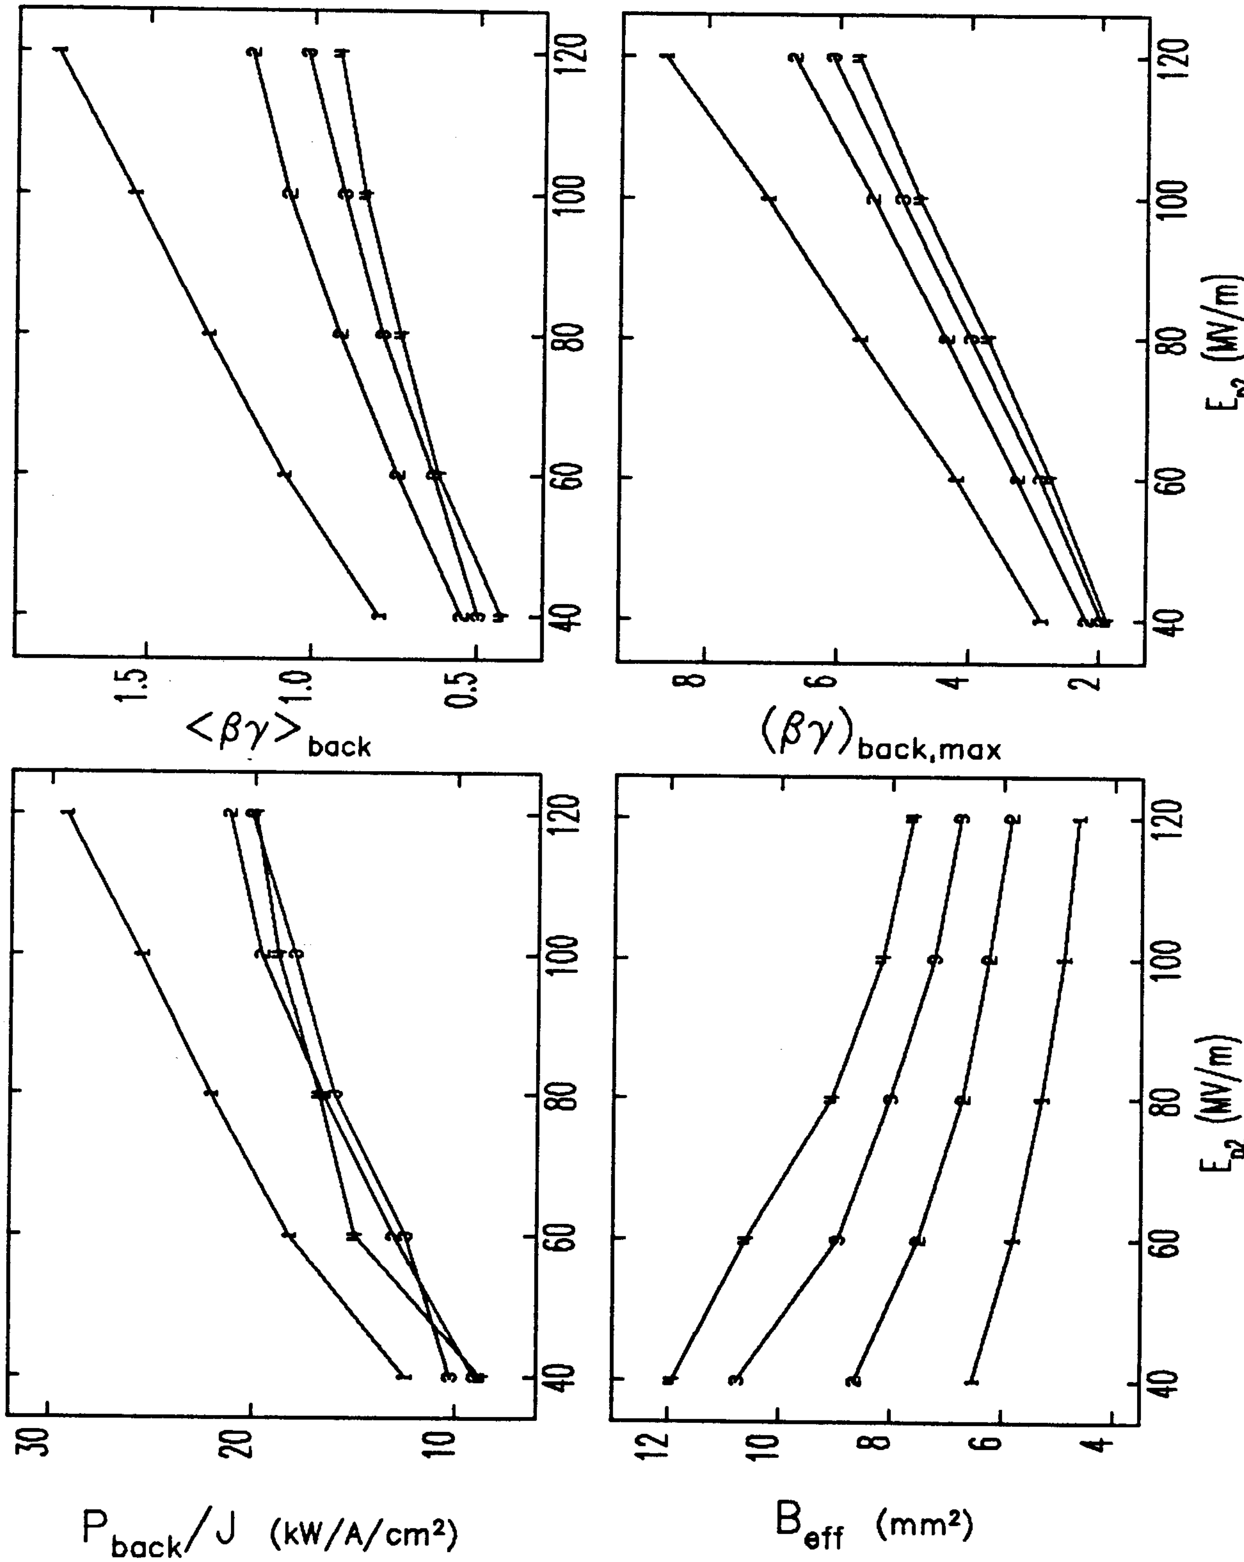

Figure 2.12: Dependence of Back-Bombardment Properties on α and E_{p2}

Back-Bombardment Power

Figure 2.12 shows corresponding quantities for the particles that return to hit the cathode. B_{eff} is the effective back-bombarder area, defined by

$$B_{\mathrm{eff}} = \frac{Q_{\mathrm{back}}}{T_o J}, \tag{2.56}$$

where Q_{back} is the charge that returns to hit the cathode during one RF period. One expects that $A_{\mathrm{eff}} + B_{\mathrm{eff}}$ will be constant, and this is found to be the case by comparison with the last figure (slight discrepancies are due to particles that do not exit the gun or hit the cathode within the total time interval of the simulation).

The Figure also shows the back-bombardment power normalized to the cathode current density. Paradoxically, there is a strong dependence on E_{p2}, but no clear trend with α.

It is actually unclear what one should expect to see here, since increasing E_{p2} or decreasing α is expected to increase the efficiency of charge extraction (as the graphs of A_{eff} confirm), thus decreasing the number of back-bombarding electrons (as the graphs of B_{eff} confirm). However, increasing E_{p2} or decreasing α is also expected to increase the momentum of any electrons that do return to hit the cathode. This expectation is confirmed by Figure 2.12. What one sees is that for constant E_{p2} and $\alpha \geq 2$, the increase in the amount of back-bombarding charge is compensated by the decrease in the average kinetic energy carried by each particle. As E_{p2} is changed, a different effect comes into play, namely the change in the back-bombardment power due to highly energetic electrons returning from the second cell. This accounts for the difference between increasing E_{p2} and decreasing α. The Figure shows the maximum momentum of any back-bombarding electron as a function of α and E_{p2}, confirming this analysis.

2.3.2 Effects of the RF Frequency

Next, I investigate the effect of changing the RF frequency while keeping the field profile the same, choosing $\alpha = 2.9$ for this and all subsequent `rfgun` studies. Note that it is sensible to imagine changing the RF frequency while keeping the on-axis fields the same, since this can be done by means of a tuning device far from the axis,

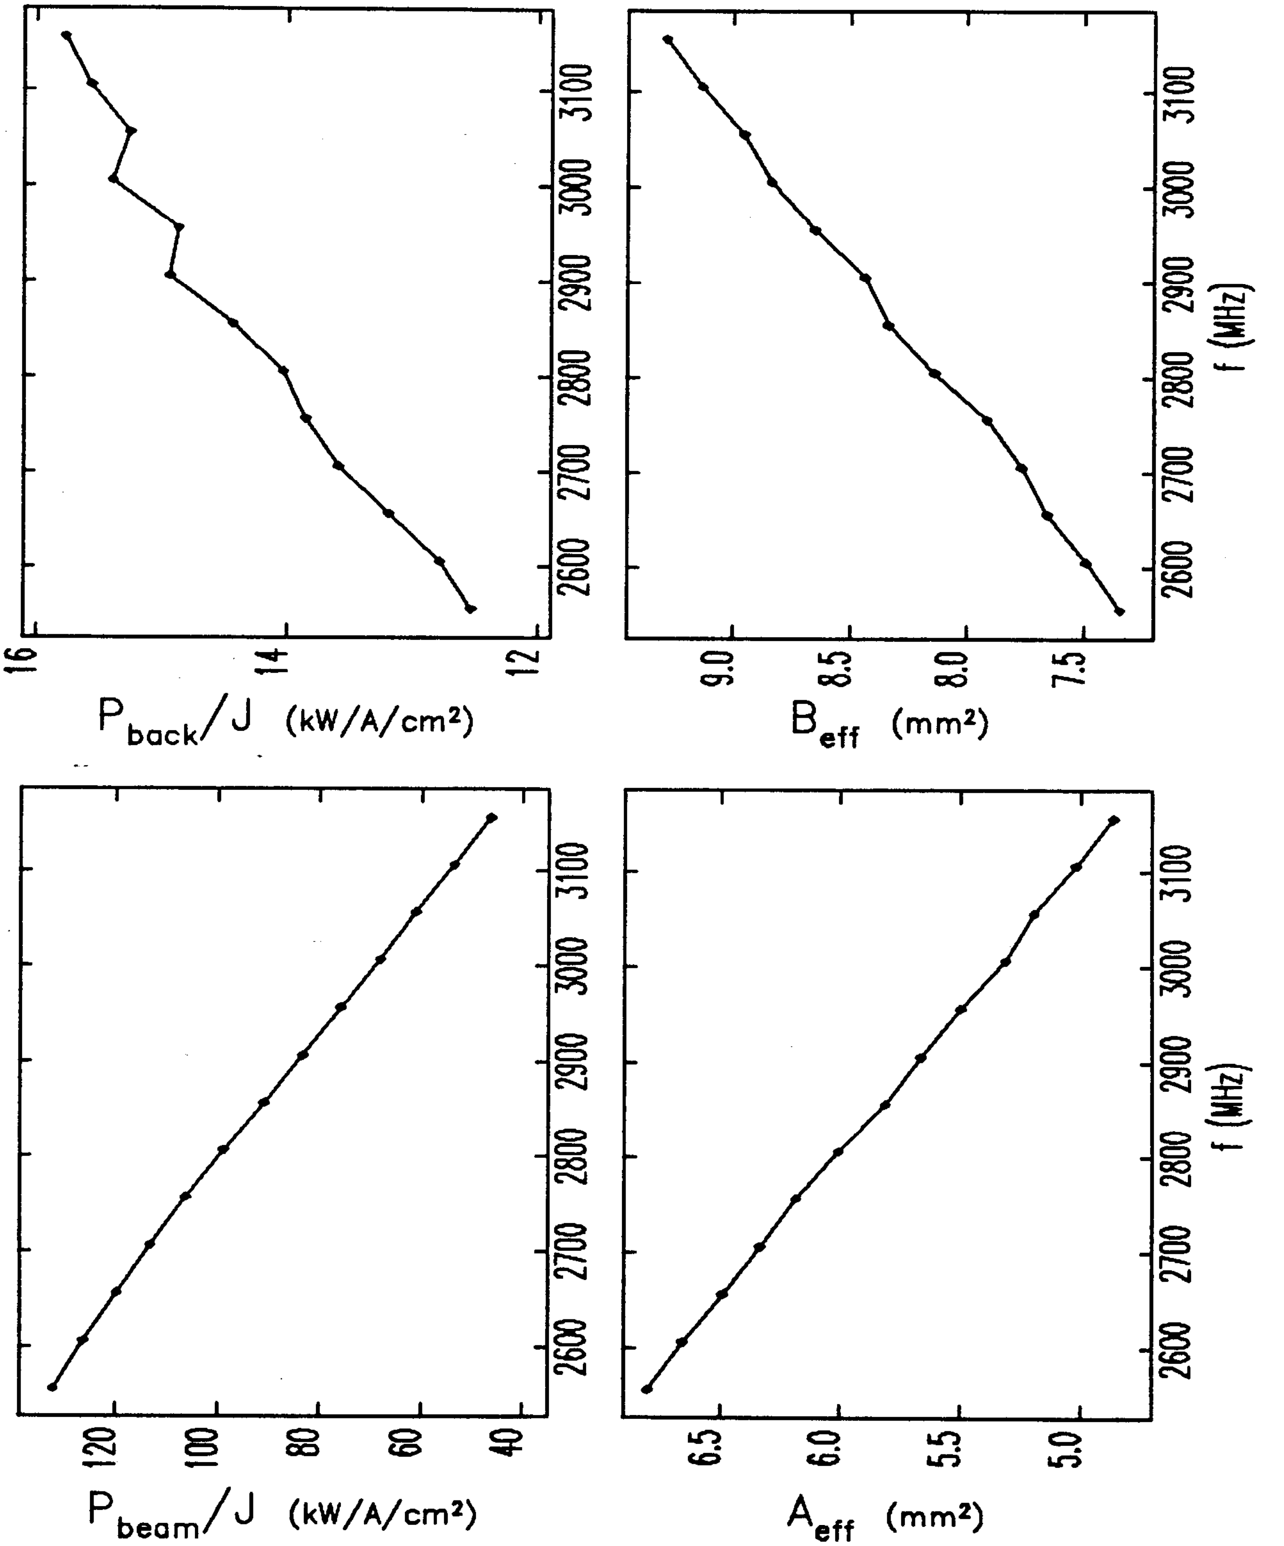

Figure 2.13: Dependence of Beam Properties on RF Frequency

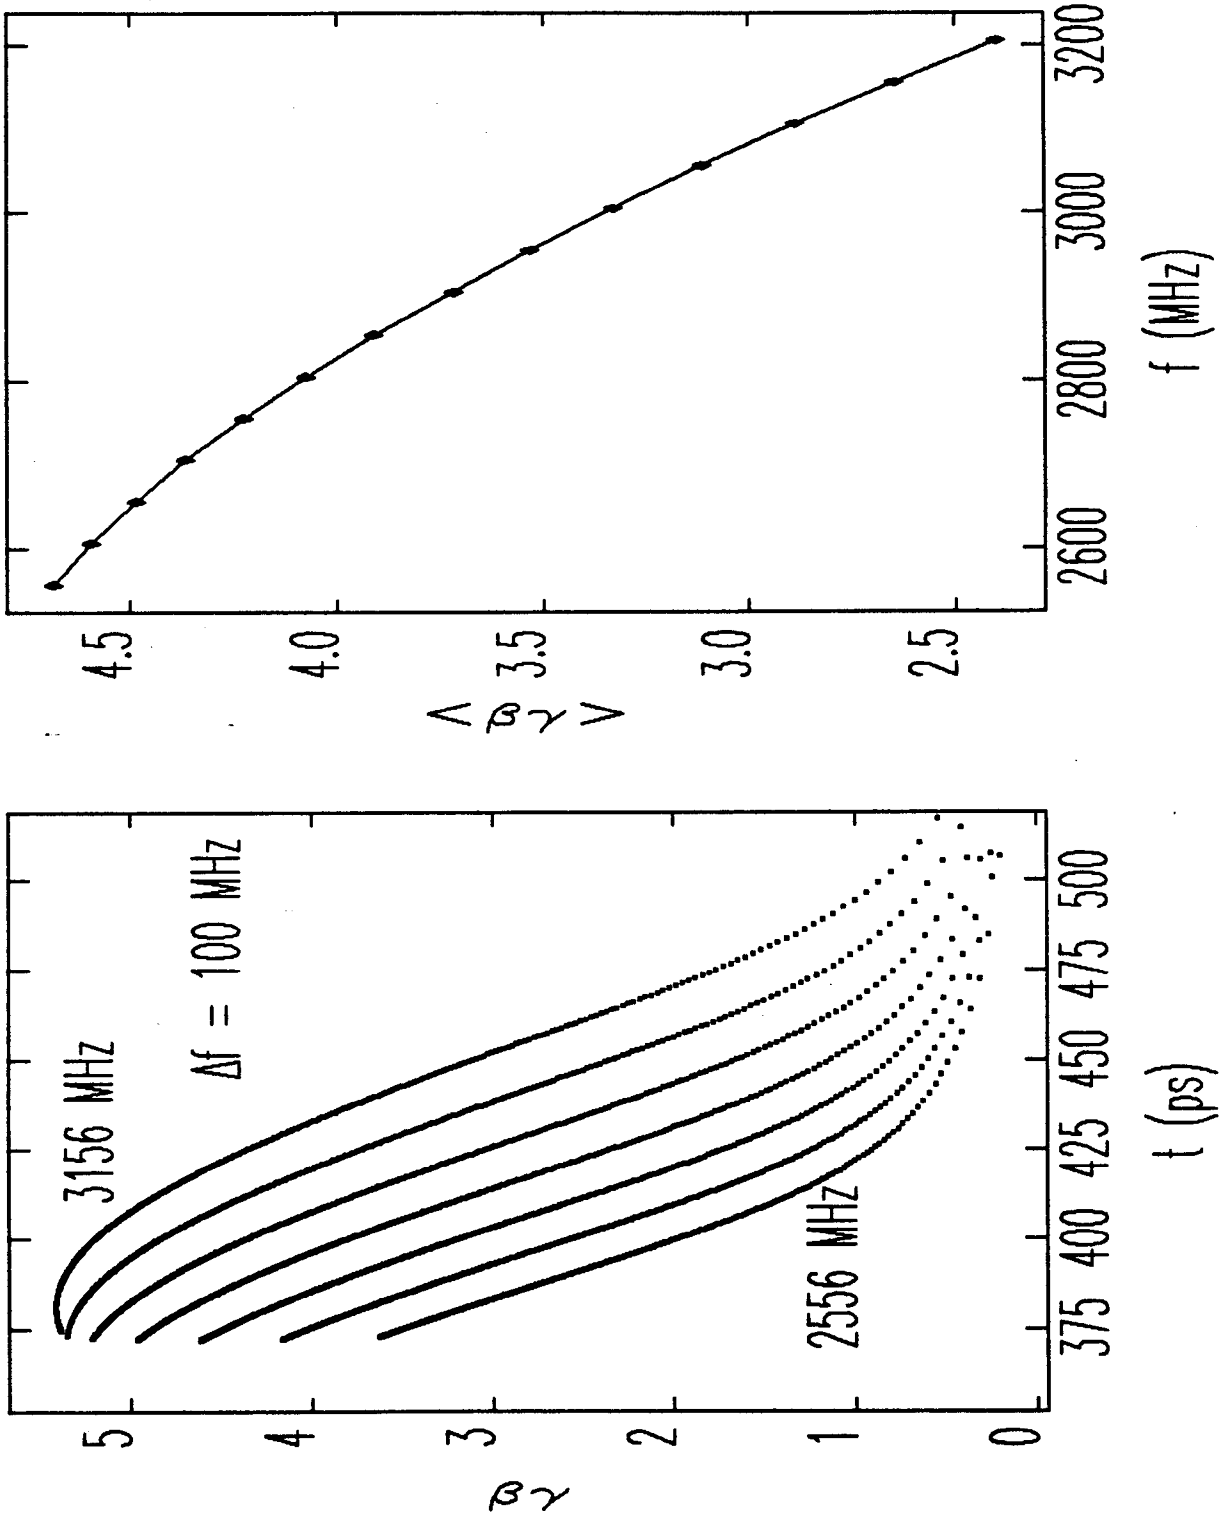

Figure 2.14: Dependence of Longitudinal Phase-Space on RF Frequency

as I did for the MASK simulations of the first cell. Figure 2.14 shows the longitudinal phase-space as a function of the RF frequency, while Figure 2.13 shows other beam properties as a function of frequency. Higher frequencies are equivalent in some ways to having a phase-shift of less than $\pi/2$ between the first and second cells, and hence the macro-electrons are seen to arrive nearer to the RF crest in the second cell. This is why the longitudinal phase-space exhibits an increasingly sinusoidal shape as the frequency is increased. For lower frequencies, the macro-electrons arrive further behind the crest in the second cell; hence, the total momentum gain is less, and the monotonic p(t) curve is retained. Lower frequencies are equivalent to having a lower-than-velocity-of-light structure.

For assessing the effect of frequency errors in MASK, it is only the region around 2856 MHz that is of interest. One sees from the Figure that the phase-space curves for $\pm$100MHz around 2856 MHz are not greatly different from those for 2856 MHz. That is, there is no dramatic effect on the slope or curvature. I find that the change in the average momentum between 2856 and 2836 MHz is 1.7%, while the change in the maximum momentum is 1.0%. The change in the total charge is similarly small, being 2.0%. Hence, the errors made in using 2836 as the frequency in MASK are negligible. Uncertainty in the value of α and in the knowledge of the exact field distribution, plus the effects of higher-order modes induced by the beam, will contribute errors as large as those introduced by the frequency errors.

2.3.3 Effects of Non-Linear Field Terms

Next, I look into the effects of field non-linearities. In particular, Figure 2.15 shows the normalized RMS emittance and normalized transverse brightness for various initial phase intervals (explained presently), where the normalized *transverse* brightness is defined as

$$\mathrm{B}_{\mathrm{n}}^{*} = \frac{\mathrm{Q}}{\mathrm{J}\epsilon^{2}}, \tag{2.57}$$

and the normalized RMS emittance is as given in equation (1.9). Note that $\mathrm{B}_{\mathrm{n}}^{*}$ is *not* the same as the normalized brightness, B_{n} defined in equation (1.21). $\mathrm{B}_{\mathrm{n}}^{*}$ differs in begin normalized to the current density and in having no reference to the bunch length.

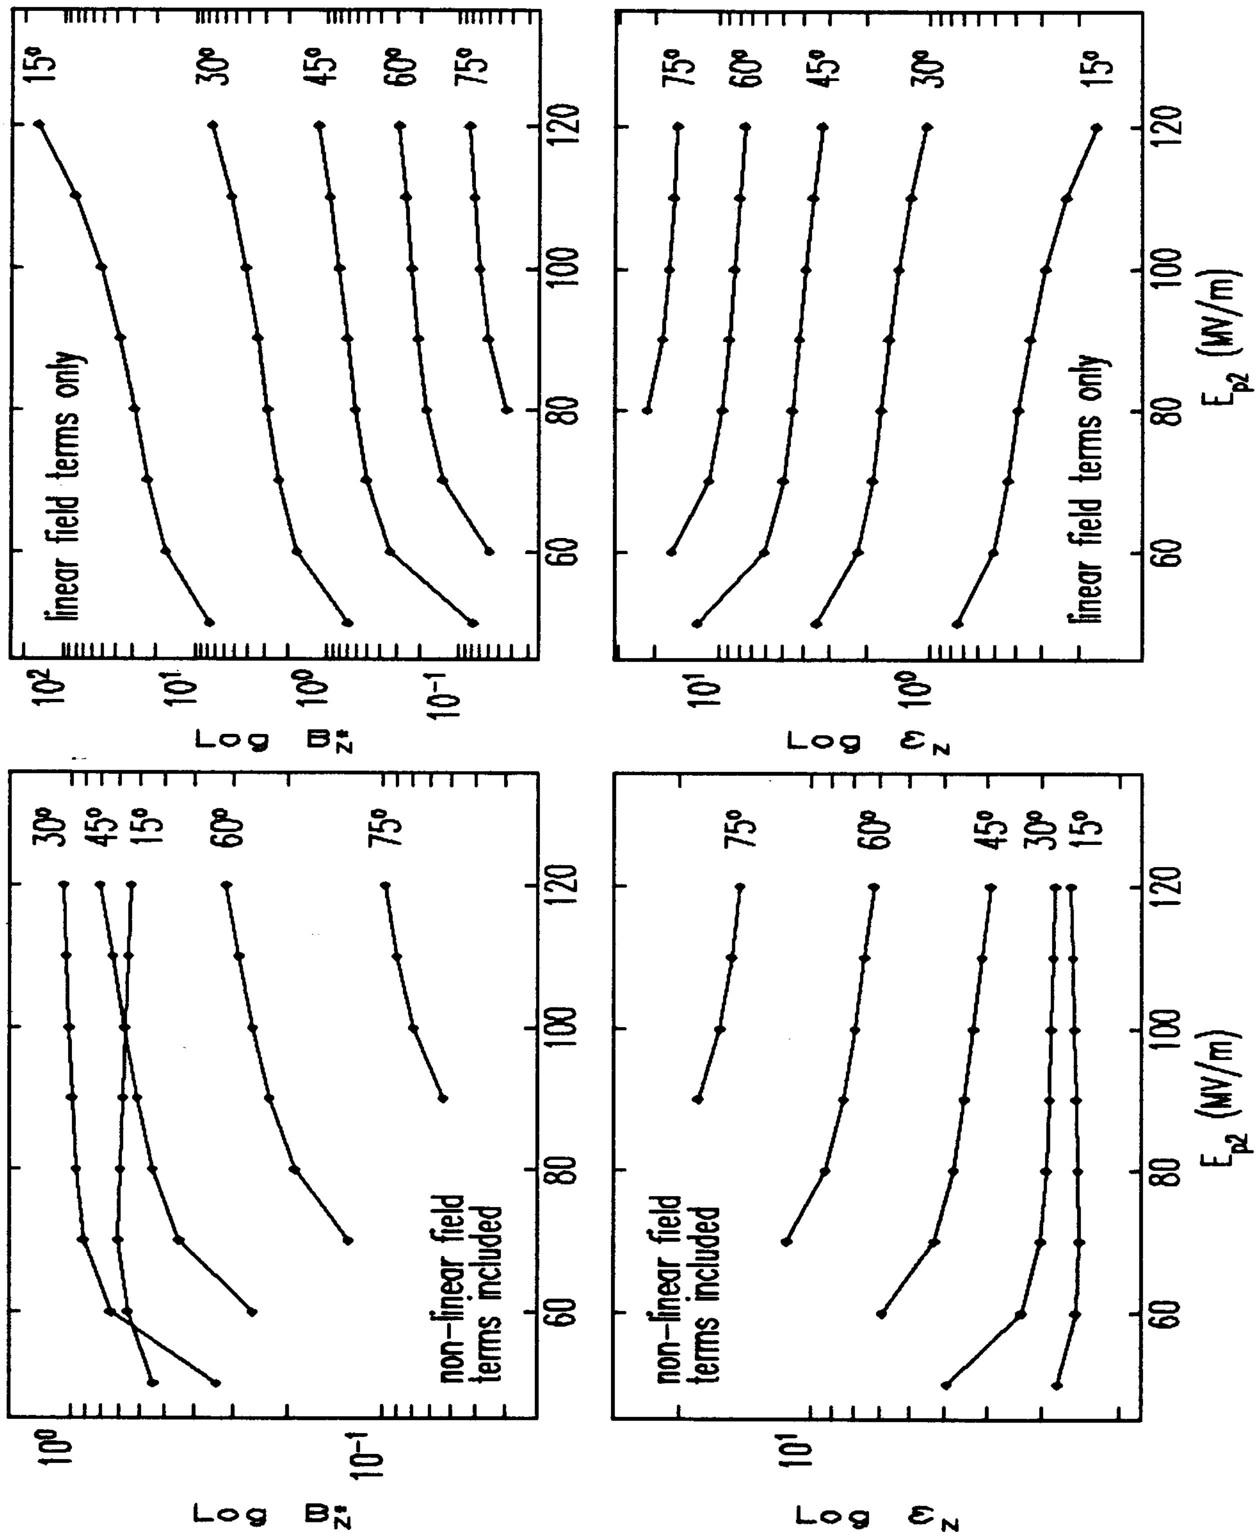

Figure 2.15: `rfgun` Predictions for $\alpha = 2.9$ and Various Initial Phase Intervals

The units for Q, J, and ϵ used for the graphs are pC, A/cm^2, and $\pi \cdot m_e c \cdot \mu m$, respectively. Results are shown for calculations that include non-linear RF field terms (see equations (2.35) through (2.37)), and for calculations that include only linear terms in E_r, B_ϕ, and E_z. In order to make a valid comparison between the linear and non-linear cases, it is necessary to ensure that one is looking at the emittance of the same particles in each case. That is, one wants to compare the emittance of particles emitted over the same range of initial phase. This has been done in the Figures, where I plot the results for the linear and non-linear cases for particles emitted during the first $\Delta\phi$ degrees of phase (measured from the beginning of the accelerating phase). As one would expect, the non-linear fields increase the emittance. Just how this occurs, and what the significance of the brightness is, will be discussed shortly.

RF Focusing

To better understand the effects of non-linear fields and time-dependent focussing forces, it is useful to look at a pair of **rfgun** runs in more detail. In particular, I ran **rfgun** at $E_{p2} = 75\text{MV/m}$ with the initial particle distribution given by

$$x_{ij} = i \cdot \delta r, \quad i = 0 \ldots N_r \tag{2.58}$$

$$y_{ij} = 0 \tag{2.59}$$

$$\phi_{ij} = j \cdot \delta\phi, \quad j = 0 \ldots N_\phi, \tag{2.60}$$

where $\delta r = R_c/(N_r - 1)$ and where ϕ_{ij} is the phase of emission from the cathode, $\phi = 0$ being the beginning of the accelerating phase in the first cell. Taking such an initial distribution allows one to see with particular clarity what the effects of the initial coordinates are on the final coordinates. Figure 2.16 shows the resultant transverse phase-space for linear and non-linear fields, with $N_r = 6$, $N_\phi = 1441$, and $\delta\phi = 0.125°$. While at first glance the non-linear fields would appear to decrease the emittance, this is not so. For the linear case, macro-electrons leaving the cathode at the same phase lie along a line in x–x′ space. Hence, the emittance for such a group of macro-electrons is zero, since $\langle x'x \rangle^2 = \langle x'^2 \rangle \langle x^2 \rangle$. For the non-linear case, macro-electrons leaving the cathode at the same phase do not lie along a line in x–x′ space. This increases the emittance by decreasing $\langle x'x \rangle$. This will be seen more clearly in other data, below.

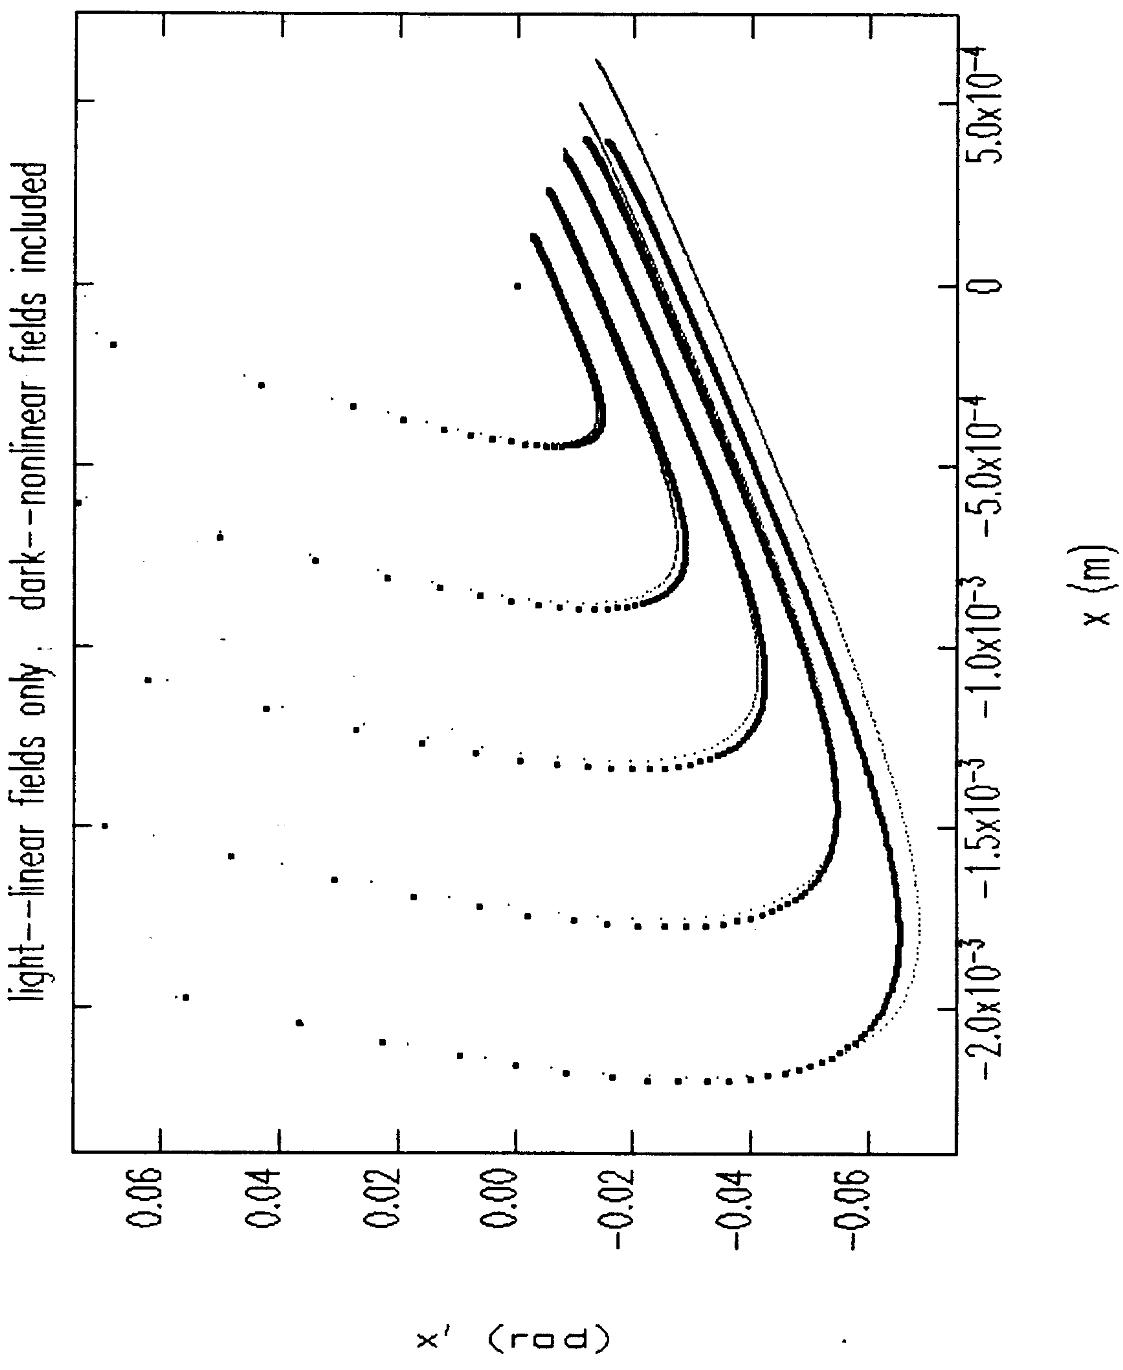

Figure 2.16: **rfgun** Transverse Phase-Space Results for E_{p2} = 75 MV/m and α = 2.9

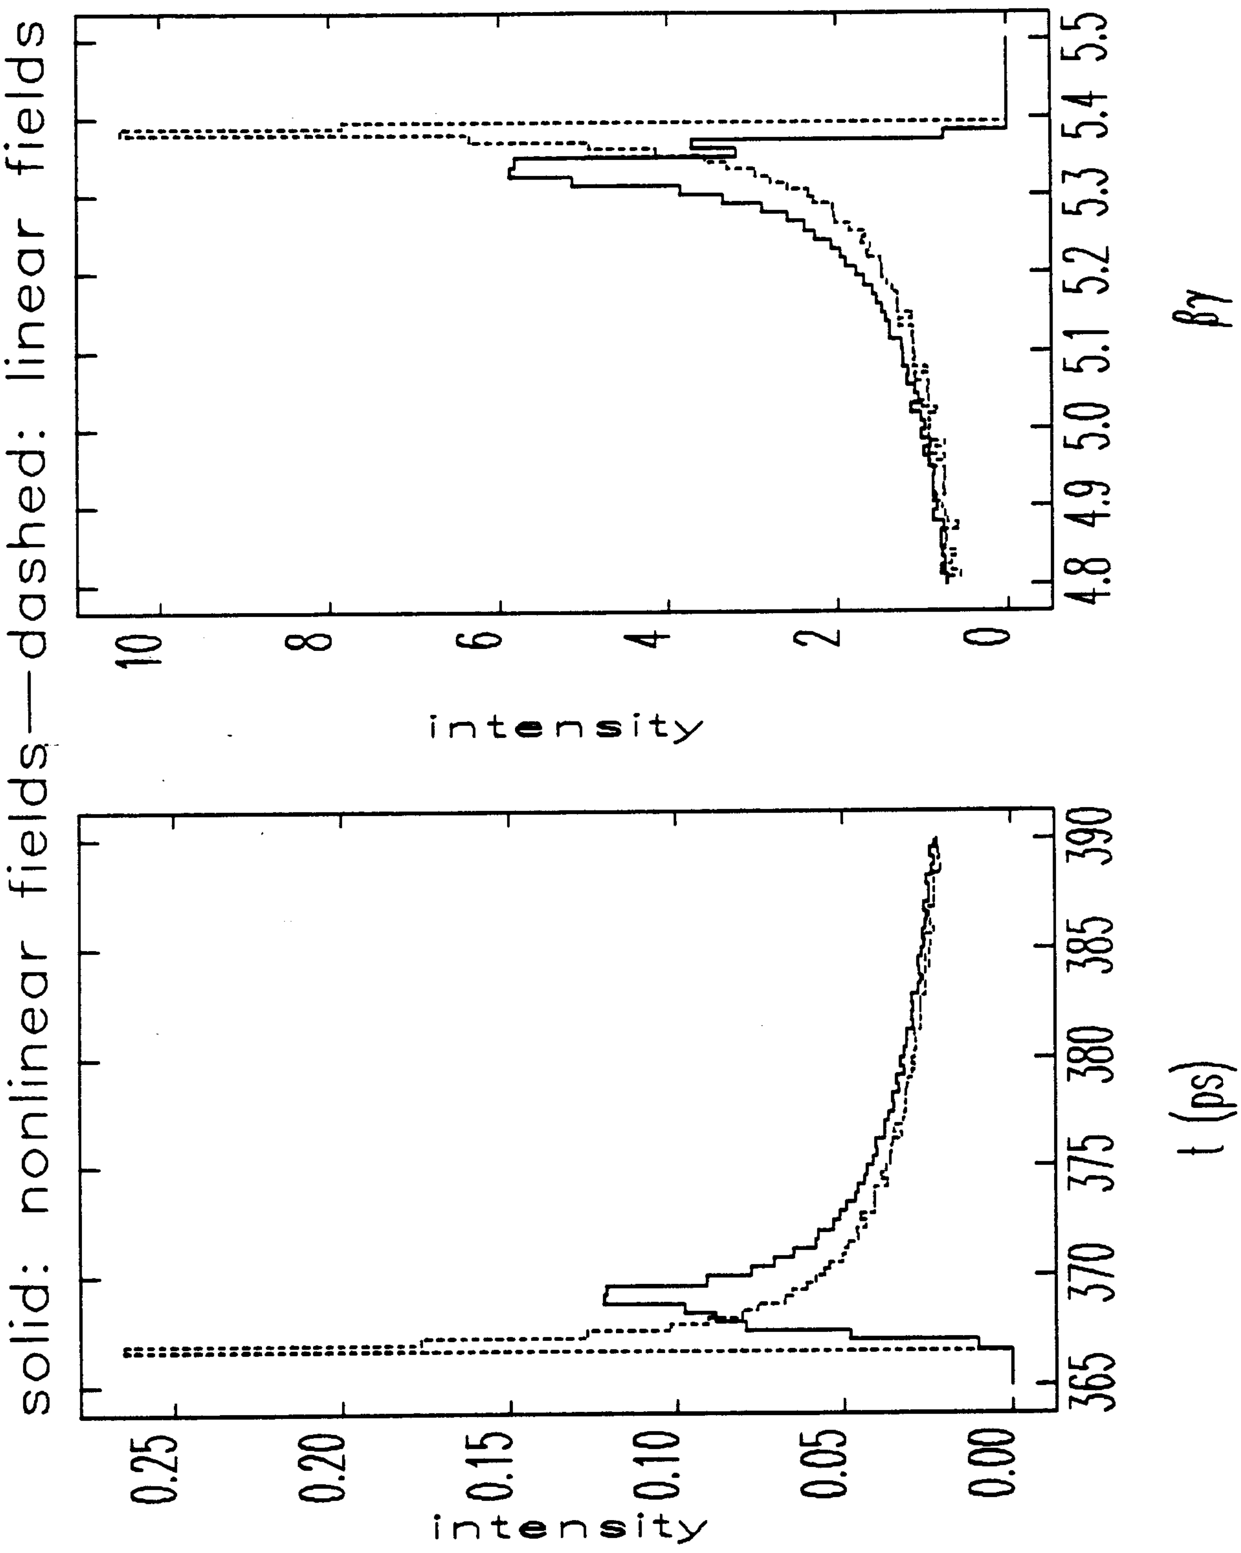

Figure 2.17: **rfgun** Longitudinal Phase-Space Results for $E_{p2} = 75$ MV/m and $\alpha = 2.9$

This Figure also illustrates the effects of "RF focusing" on the emittance. The first macro-electrons emitted from the cathode are those on the right side of the figure ($x > 0, x' < 0$). Subsequent particles are focused differently by the time-varying (and hence initial-phase-dependent) E_r and B_ϕ fields. As a result, the RF sweeps the particles in a clockwise sense in the Figure. This is all consistent with Figure 2.15 where one sees that the emittance increases as a larger initial phase interval is considered, and that, for the linear case, the emittance falls increasingly rapidly as $\Delta\phi$ is decreased, whereas for the non-linear case, the emittance seems to reach a lower limit as $\Delta\phi$ is decreased. The brightness clearly saturates in the non-linear case, whereas it does not do so in the linear case. From Figure 2.16 it is also clear that curvature of the x–x' path traced by particles starting at the same radius is a result of the sinusoidal nature of the RF fields, rather than non-linearities in r.

Longitudinal Phase-Space

Figure 2.17 shows the longitudinal phase-space for the non-linear and linear cases. The non-linear fields produce a broadening of the momentum peak and of the time distribution as well, since particles starting at different radii experience different accelerating fields. In the linear case, all particles starting at the same phase receive, to first order, the same momentum and take the same time to exit the cavity. I infer from the sharpness of the momentum and time peaks for the linear case that any deviations in momentum gain or time-of-flight that result from longer path-lengths due to transverse motion are small. The momentum and exit-time distributions predicted by MASK are considerably broader than those predicted by `rfgun`, even with non-linearities included in `rfgun`. The reasons for this will be seen later in this chapter.

Transverse Phase-Space Evolution

It is instructive to look at the evolution of the transverse phase-space as the beam travels through the gun. To do this, I have run `rfgun` with an initial particle distribution defined by equations (2.59) through (2.60), with $N_r = 41$, $N_\phi = 19$, and $\delta\phi = 10°$. Figure 2.18 shows the resultant transverse phase-space at a series of z positions in first cell. The lines connect particles emitted at the same phase but different cathode

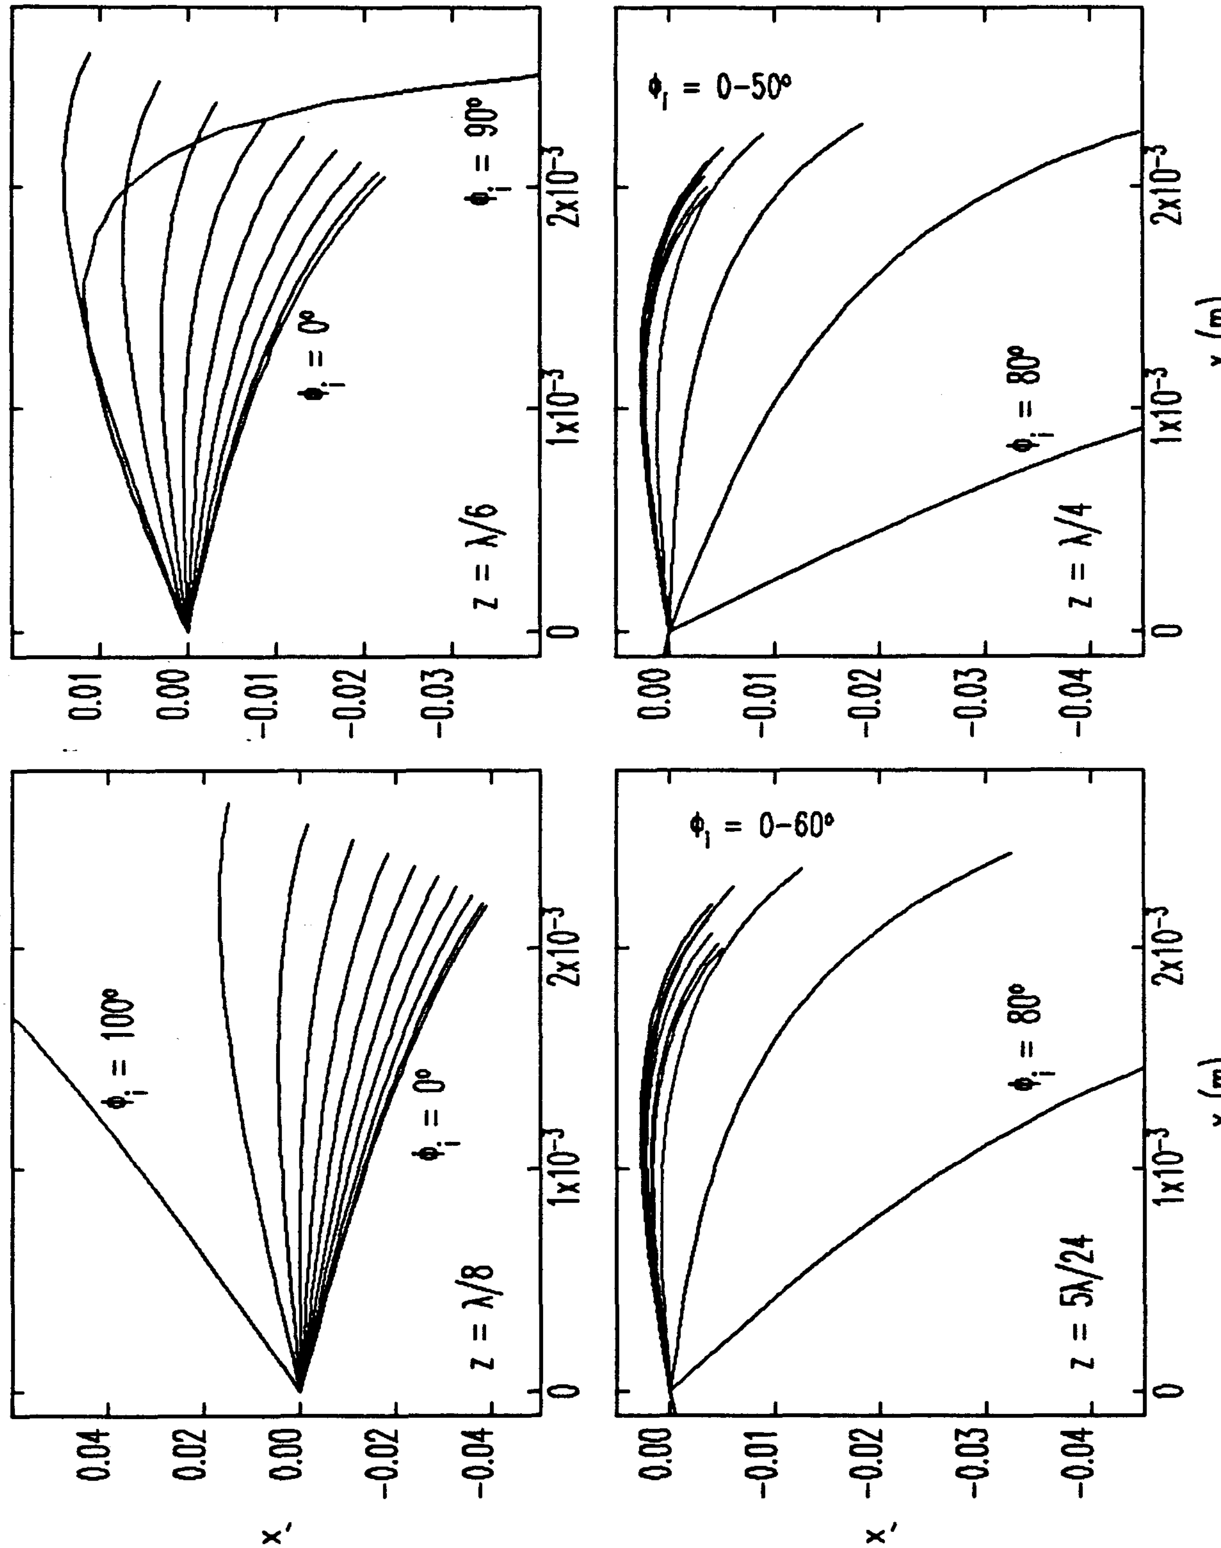

Figure 2.18: **rfgun**: Transverse Phase-Space Evolution in the First Cell, for $E_{p2} = 75$ MV/m and $\alpha = 2.9$

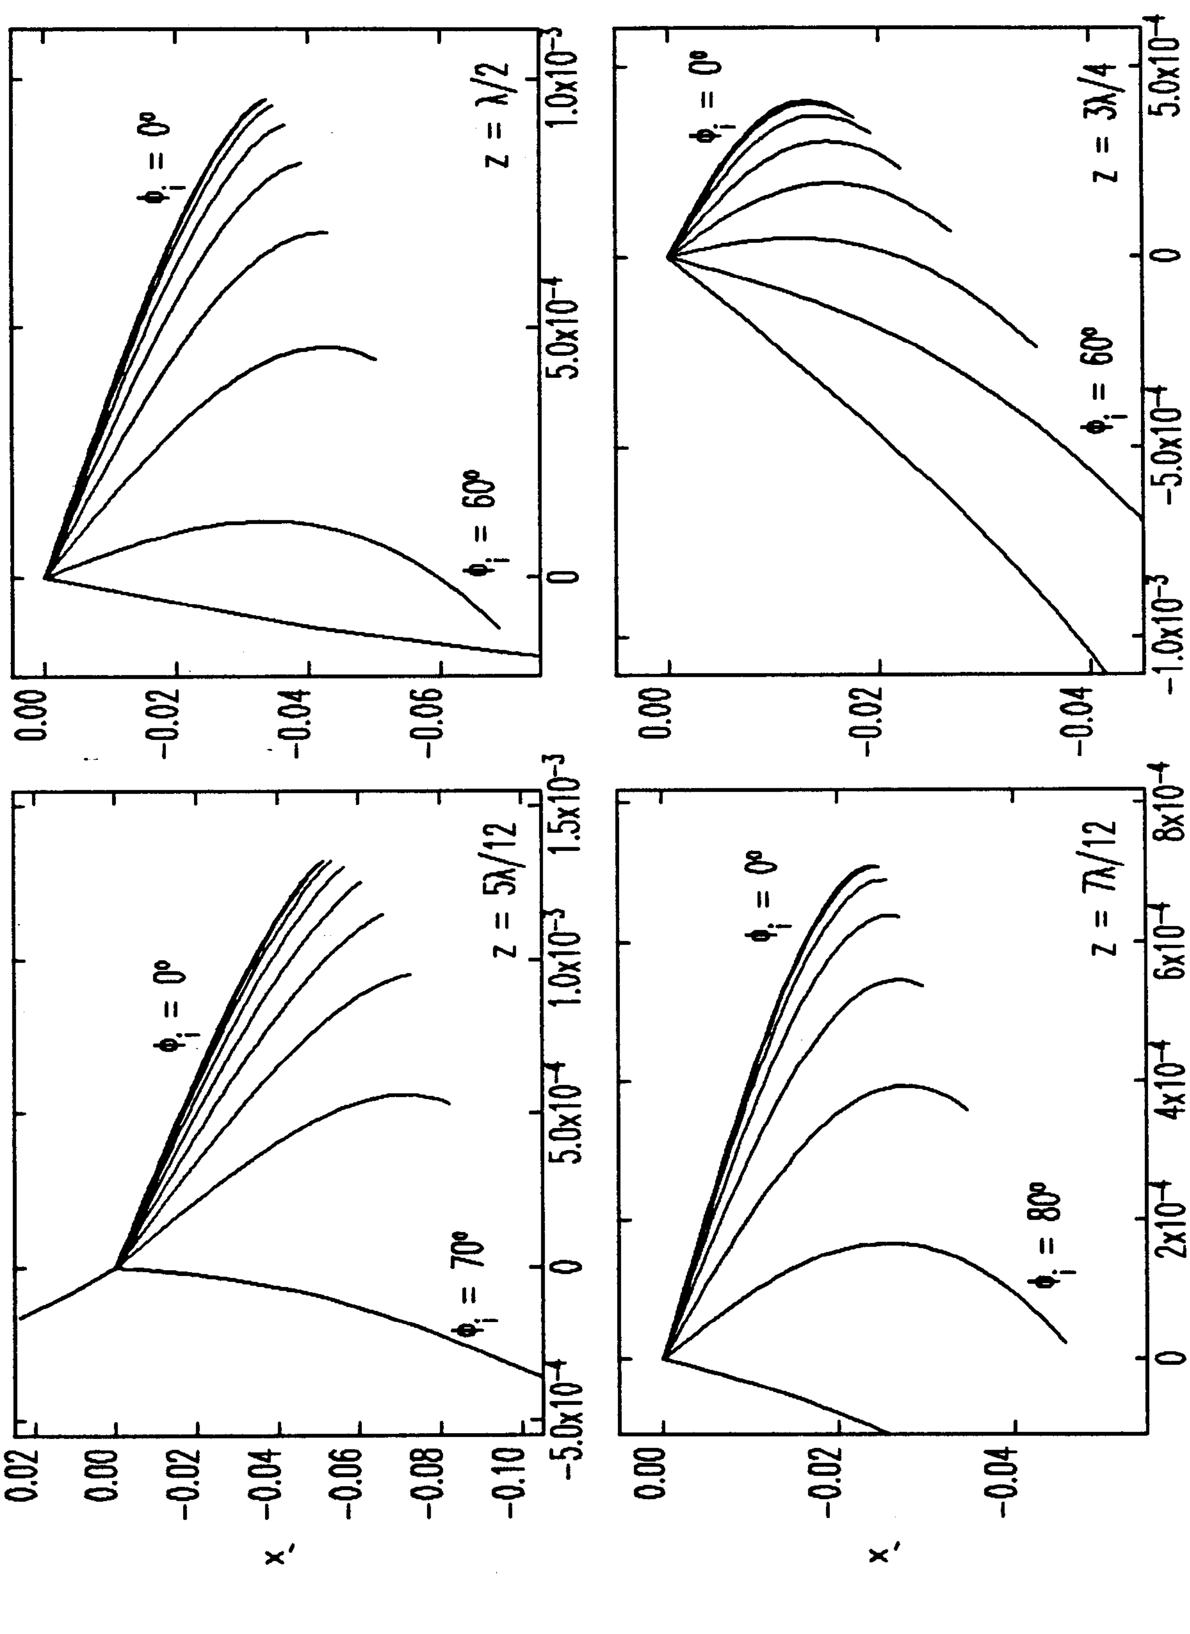

Figure 2.19: **rfgun**: Transverse Phase-Space Evolution in the Second Cell, for $E_{p2} = 75$ MV/m and $\alpha = 2.9$

radii. (All of the graphs have fewer than N_ϕ lines, because particles emitted with ϕ increasingly near to 180° traveling increasing short distances in the first cell, and hence do not reach the z-planes of these phase-space plots.)

Even from the first of these plots one can see both non-linear and RF focusing effects, the later causing the fanning out of the lines, while both contribute to the curvature (I shall show below why this last point is true). As the bunch proceeds, one sees that curvature is more severe for the particles emitted at later phases, a result primarily of the large phase-spread these particles end up with. One also sees that the RF focusing is not simply fanning out the beam according to initial phase, but is also "mixing" the beam in transverse phase-space. This is simply a result of the sinusoidal variation of the focusing forces (i.e., if the variation in time were monotonic, there would be no mixing). The number of lines is seen to decrease with successive graphs, due to the slowing down and back-accelerating of the later parts of the beam.

The evolution of the beam in the second cell is shown in Figure 2.19. One sees that here is a dramatic increase in the curvature of the lines, i.e., the effect of non-linear fields, that occurs in this cell. While would appear that this change is a result of fields in the center of the second cell, this is an incorrect conclusion. What happens is that as the beam continues to converge, the effect of non-linearities from the first cell and the beginning of the second cell becomes much more evident. As one sees from Figure 2.9, the non-linear fields in the center of the second cell are very small.

Relative Importance of Different Non-Linear Field Components

I stated earlier that the curvature of the lines is due only partly to non-linearities, and partly to the time-variation of the fields. The reason is that the non-linear E_z terms cause particles starting at the same phase but at different radii to be accelerated at different rates, and thus to go through the gun at different phases relative to the RF. `rfgun` allows one to selectively "turn off" the non-linear E_z, E_r, and/or B_ϕ fields. By turning off the non-linear E_z terms in RF gun, I have verified that this is significant effect.

Figure 2.20 shows the effect of turning off each of E_r, E_z, and B_ϕ in turn. Each graph has both the results for all non-linear fields (in the lighter pen) along with the results with one non-linear field eliminated. Non-linear E_z terms have a dominant

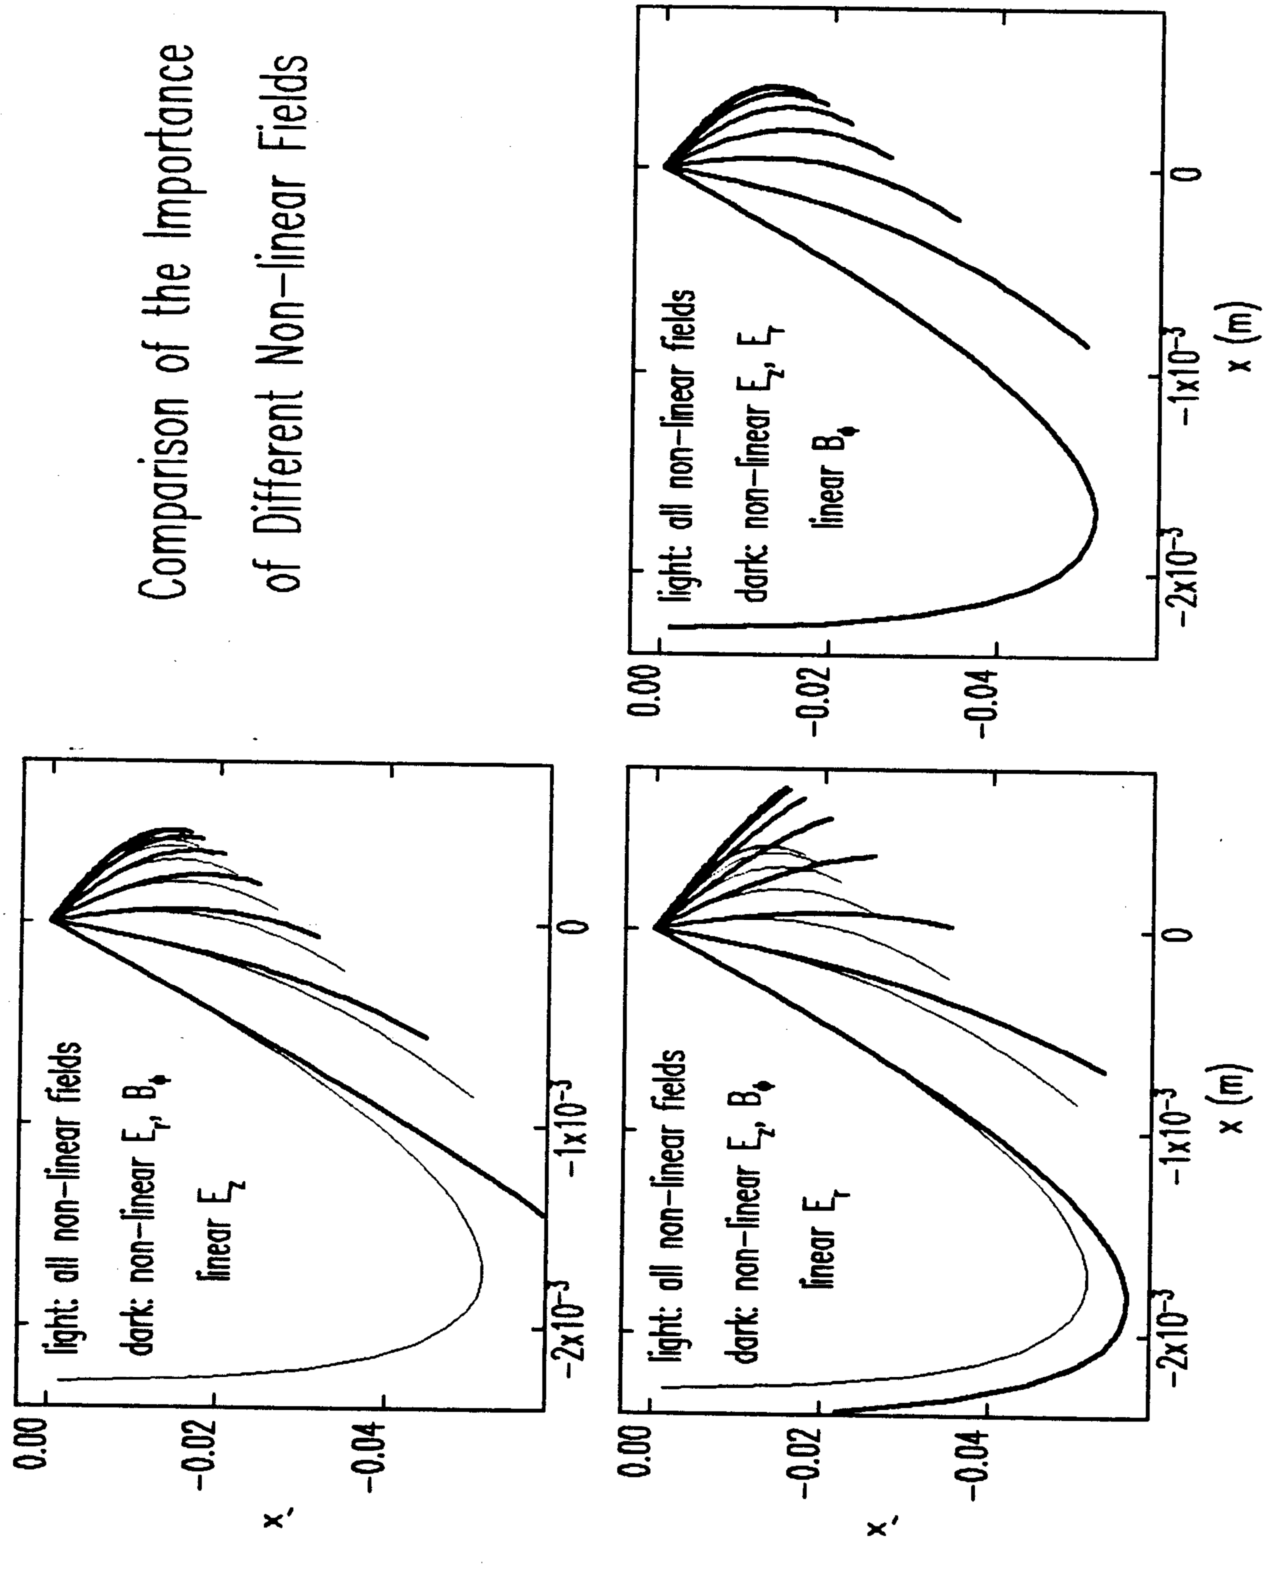

Figure 2.20: rfgun: Effects of Different Non-Linear Fields

effect on later particles, while E_r is the dominant source of curvature for the particles emitted closer to $\phi = 0$. This is plausible, since the velocity spread for the later particles (which have lower momenta) will be larger, since they are less relativistic. The effect of nonlinear B_ϕ terms is seen to be insignificant (the differences caused by removing non-linear B_ϕ terms cannot be seen on the graph, though there are differences). This is also plausible, since B_ϕ is 90° out of phase with the accelerating field (i.e., E_z), while the particles themselves are largely in phase with the accelerating field; in addition, one sees from Figure 2.9 that the non-linear B_ϕ terms are smaller compared to the linear B_ϕ terms than are the non-linear E_r terms compared to the linear E_r terms.

2.3.4 Effects of Parameter Errors in MASK Runs

Next, I look at the effect of using $\alpha = 3$ and $f = 2836$MHz in MASK. I use `rfgun` to do this evaluation, since it is faster (and since I don't have simulation cells tuned for 2836MHz for use in MASK). In particular, Figure 2.21 shows phase-space plots for $z = \lambda/4$ and $z = 3\lambda/4$ for $\alpha = 3.0$ and $f = 2856$MHz, and for $\alpha = 3.0$ and $f = 2836$MHz, compared to the results for $\alpha = 2.9$ and $f = 2856$MHz. One sees that while there are effects, they are confined to the particles that come later in the beam—i.e., the highest momentum particles seem least effected. The principle effect is a rotation in phase-space. The curvature of the lines is not noticeably changed.

Figure 2.22 shows the normalized emittance and normalized brightness for $\alpha = 3$ and $f = 2836$MHz, along with those for $\alpha = 2.9$ and $f = 2856$MHz. The emittance is somewhat smaller, and the brightness correspondingly larger, for the former than for the later. Thus, one expects that the MASK simulations will under-estimate the emittance by perhaps as much as 15% (though generally less) in the range $E_{p2} \geq 60$MV/m, with the error decreasing as E_{p2} increases, and decreasing as a smaller initial phase interval is taken. This difference is overwhelmingly a result of the difference in α, rather than the difference in frequency. This indicates that the difference is due to the more-rapid acceleration for $\alpha = 3$ than for $\alpha = 2.9$, which reduces the effect of non-linear E_r fields by increasing the momentum of the particles.

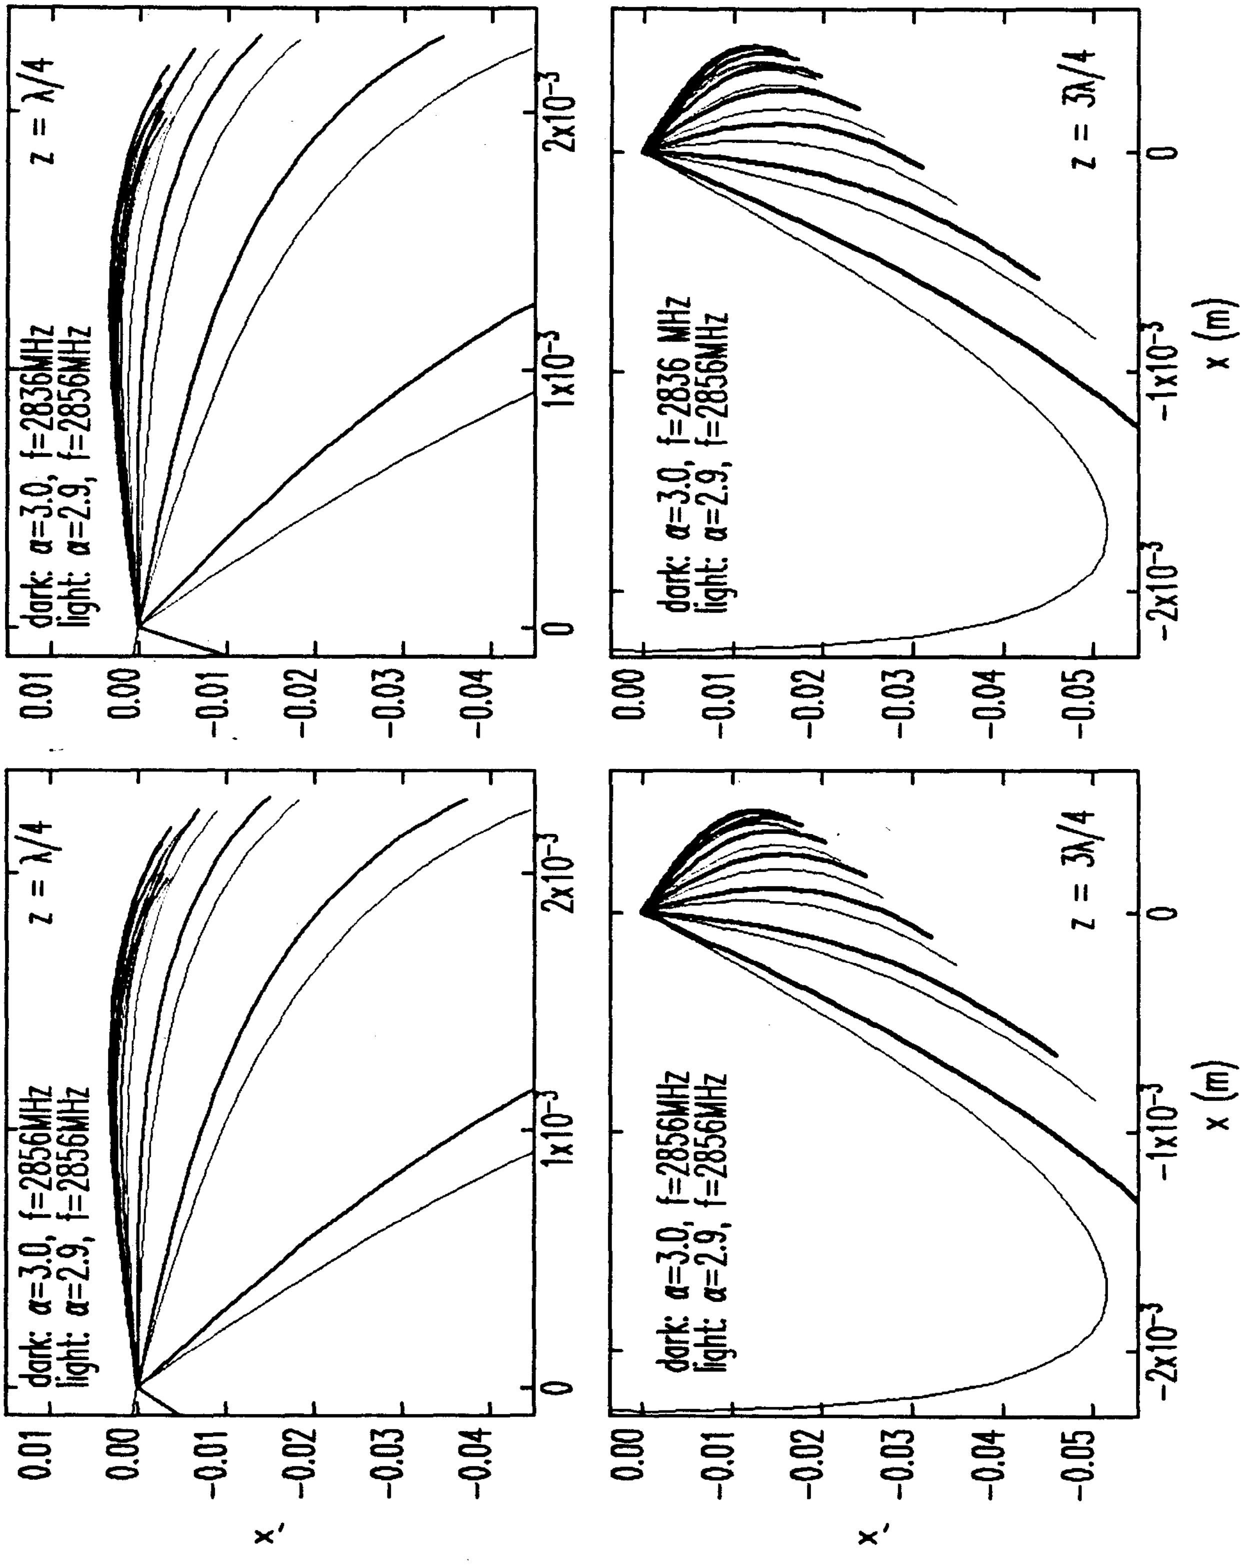

Figure 2.21: **rfgun**: Effects of α and Frequency Errors

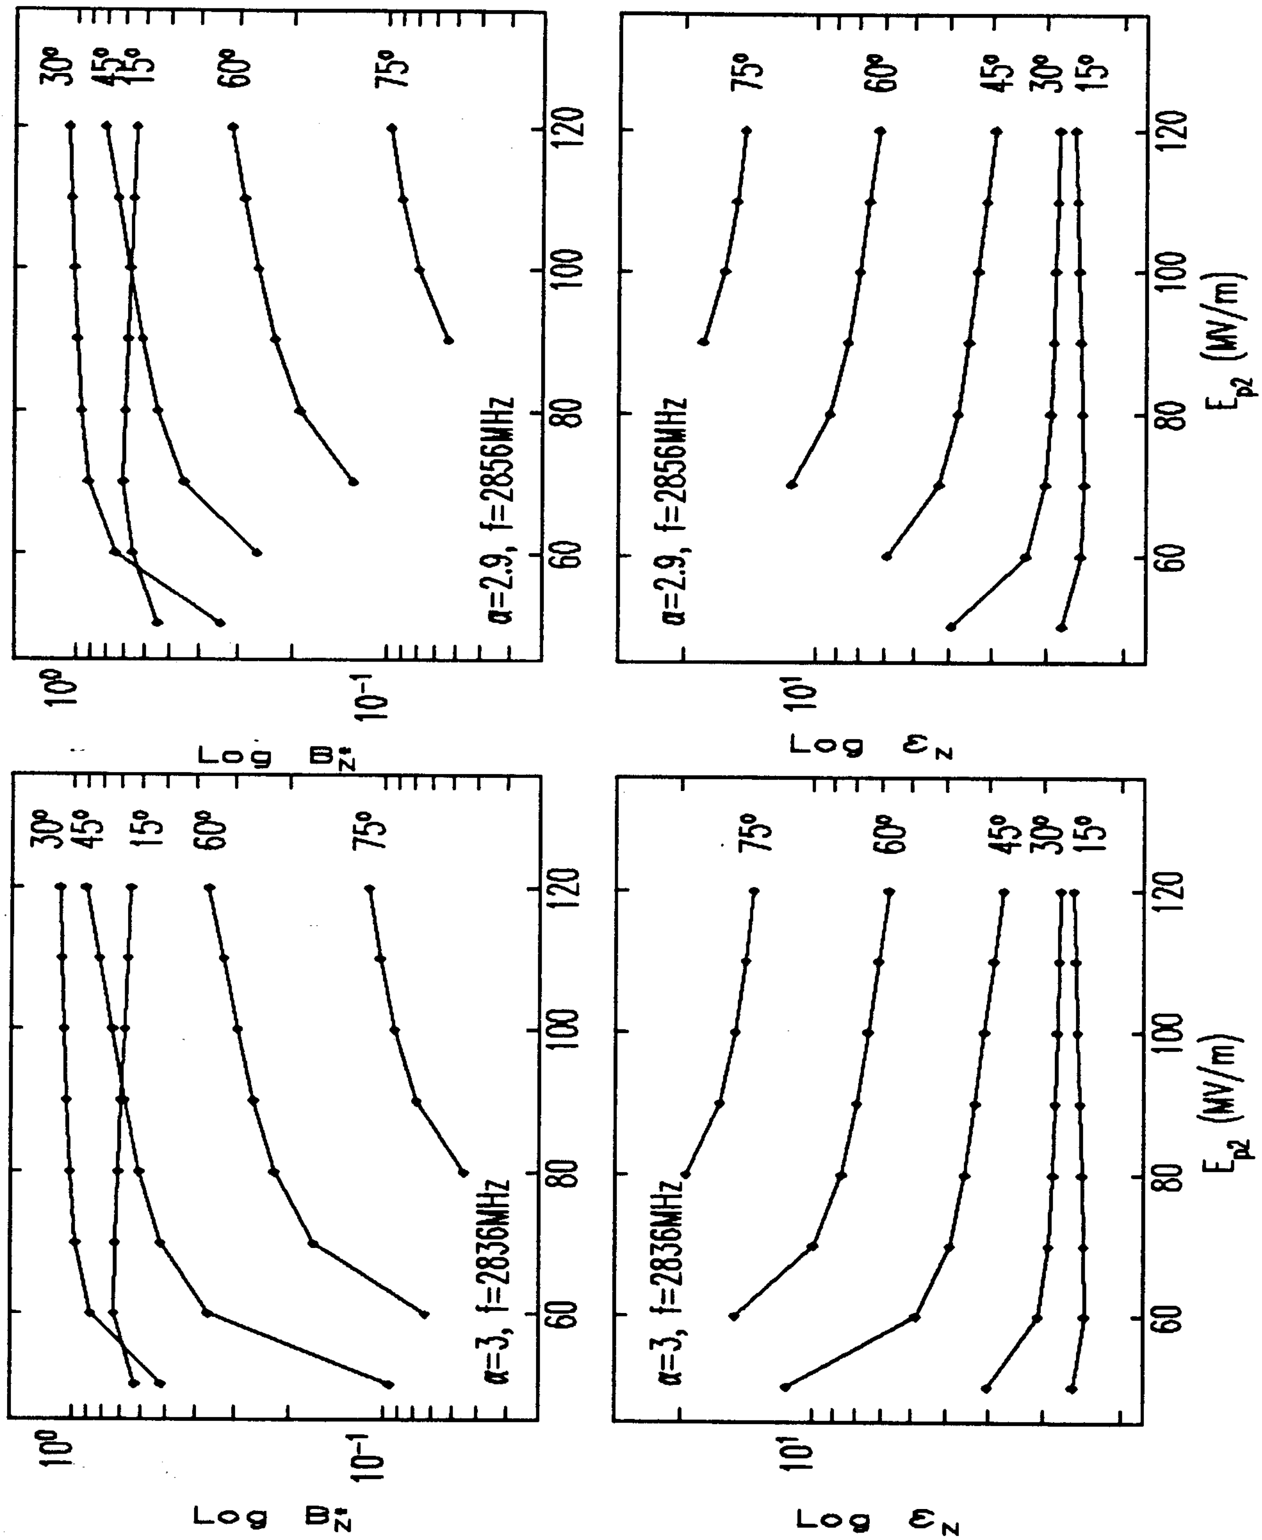

Figure 2.22: Emittance and Brightness for $\alpha = 2.9, \mathrm{f} = 2856$ MHz, and $\alpha = 3$, f = 2836 MHz.

2.3.5 rfgun Predictions versus Momentum Spread

In the foregoing, I have looked at properties of subsets of the beam based on initial phase. This is useful, but in reality one cannot choose such subsets directly. However, since there is a high degree of correlation between initial phase, ϕ, and final momentum, one can to a large extent filter for initial phase by momentum filtration. In the above, I have chosen to work in terms of ϕ directly, because it makes the analysis more straight-forward. Because the non-linearities affect the momentum distribution as well as the transverse coordinates, the final momentum interval corresponding to a given $\Delta\phi$ for the linear case is different from the final momentum interval corresponding to the same $\Delta\phi$ for the non-linear case. Hence, if the above analysis were done with momentum filtration, the difference between the linear and non-linear cases would have been blurred.

Because MASK does not provide the initial phase of particles, it is not impossible to do the analysis based on initial phase for MASK results. Hence, for the MASK results given below, I employ momentum filtration. For comparison, I do the same for the non-linear case in `rfgun`. Figure 2.23 shows the normalized RMS emittance, normalized transverse brightness, and normalized charge per bunch, for fractional momentum intervals defined by

$$\frac{1-f}{1+f}p_{max} \leq p \leq p_{max}, \tag{2.61}$$

where p_{max} is the maximum momentum in the beam and $\pm f$ is the fractional momentum range about $p_{max}/(1+f)$. The central momentum for the interval is $p_{max}/(1+f)$, and is not the same as the average momentum of the particles in the momentum interval. For typical gun operating parameters of $E_{p2} = 75\text{MV/m}$ and $J = 10\text{A/cm}^2$, `rfgun` predicts a normalized RMS emittance of less than about 4 $\pi \cdot m_e c \cdot \mu m$ and useful charge of as much as 100 pC, for momentum spread of less than $\pm 10\%$. Other relevant data for comparison of `rfgun` results with MASK can be gleaned from the data for $\alpha = 3$ presented in Figures 2.11 and 2.12.

The brightness is useful in comparing `rfgun` and MASK results. The merit of this quantity is that it is related to the density of particles in phase-space, rather than simply the area. It should thus be less insensitive than the emittance to momentum filtration "errors" (i.e., the inclusion of different subsets of the beam in the same

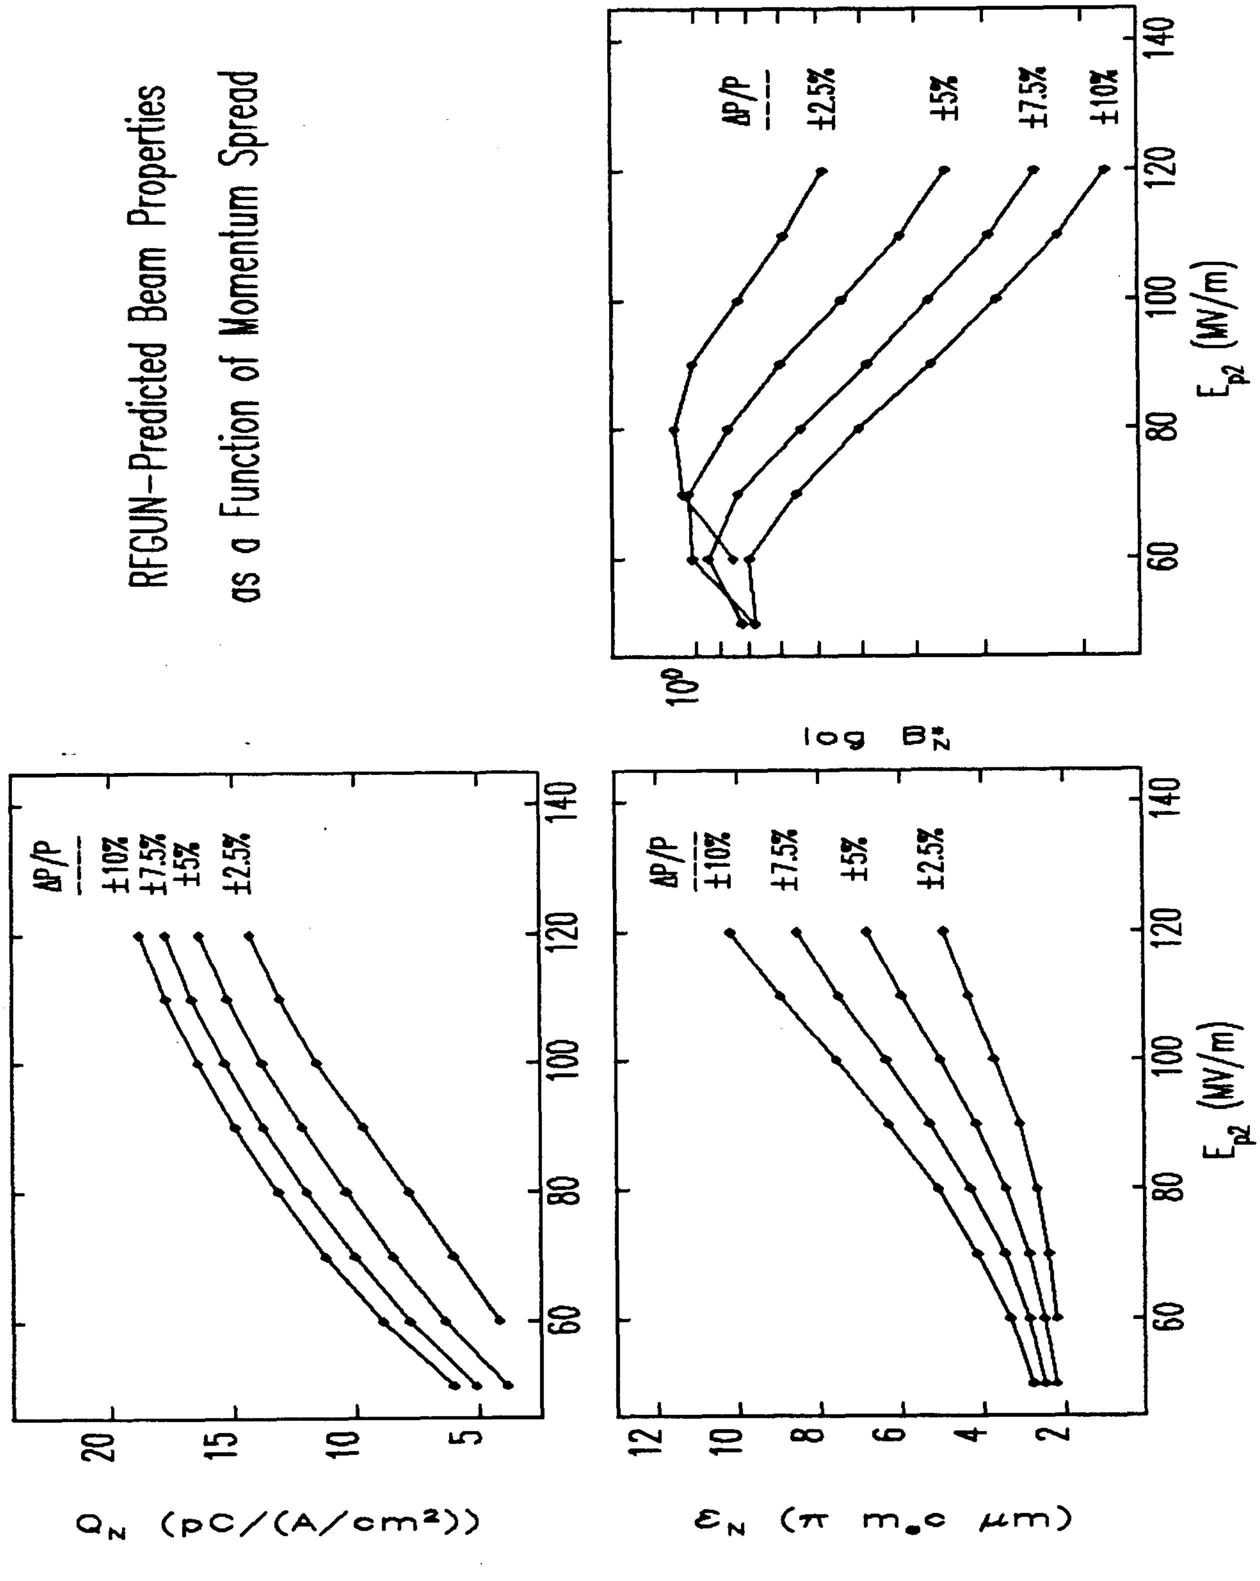

Figure 2.23: `rfgun` Predictions for $\alpha = 2.9$ and Various Final Momentum Intervals

momentum fraction f). Of course, the emittance numbers are important as well, and having them as a function of the momentum interval is also important, since in reality one is constrained by the momentum acceptance the beamline after the gun, and since chromatic effects in this beamline will worsen with larger momentum intervals. (This is discussed in Chapter 5.)

2.3.6 MASK Beam Snapshots

`rfgun` is able to simulate the gun with non-linear fields and with accurate longitudinal and transverse dynamics. However, it does not include any of the effects of space-charge. To look into these effects, I next discuss the results of MASK simulations. I have discussed my methodology in using MASK in the previous section. Here, I concentrate on what MASK predicts and on tests of the validity of my methodology.

Figures 2.24 and 2.25 show a series of beam snap-shots taken at various RF phases during the RF gun cycle, where a phase of 0° marks the beginning of the accelerating phase in the first cell. These were made for the nominal operating parameters of $E_{p2} = 75$MV/m and $J = 10$A/cm^2. The graphs in Figure 2.24 have a vertical to horizontal aspect ratio of 1, while those in Figure 2.25 have an aspect ratio of 3. The dots represent individual macro-electrons, while the solid line is the actual cavity shape (which differs slightly from the mesh approximation used in the simulation, as seen in Figures 2.5 and 2.6). (In some cases, dots appear inside the "metal" of the cavity walls or outside the simulation boundary; this is because MASK dumps the particle coordinates before checking for particles that have been lost.)

The beam in these graphs appears to be hollow in part because each macro-electron represents a ring of charge and because each ring represents the same amount of charge, so that more macro-electrons are needed at larger radii in order to achieve the same current density. The desired current density is a constant, J. Hence, the number of particles inside a radius R is $\pi R^2 J$, and the number within an annulus of width ΔR about radius R is $2\pi R \Delta R J$. Hence, the number of macro-electrons within an interval ΔR about R increases linearly with R. Another reason for the hollow appearance is that non-linear focusing terms that cause an increase in the radial field magnitude with increasing r, tend to produce an increasing particle density with

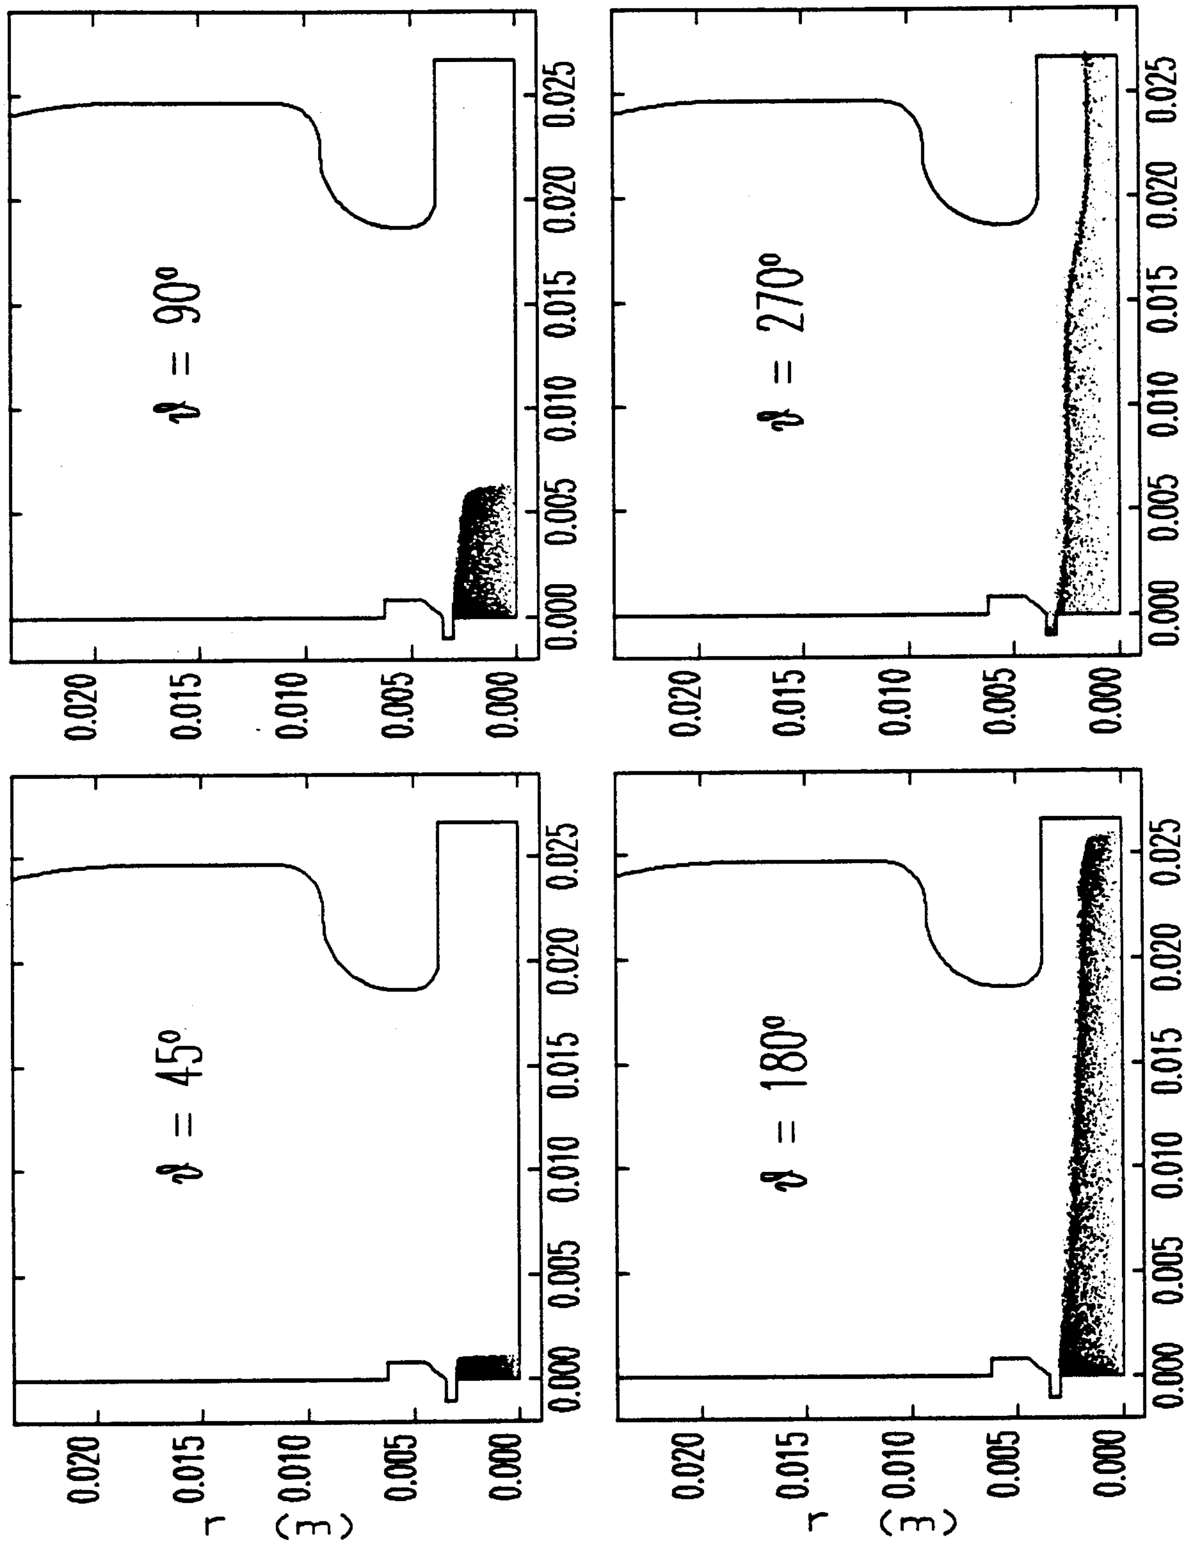

Figure 2.24: MASK Beam Snapshots at Various Phases—First Cell

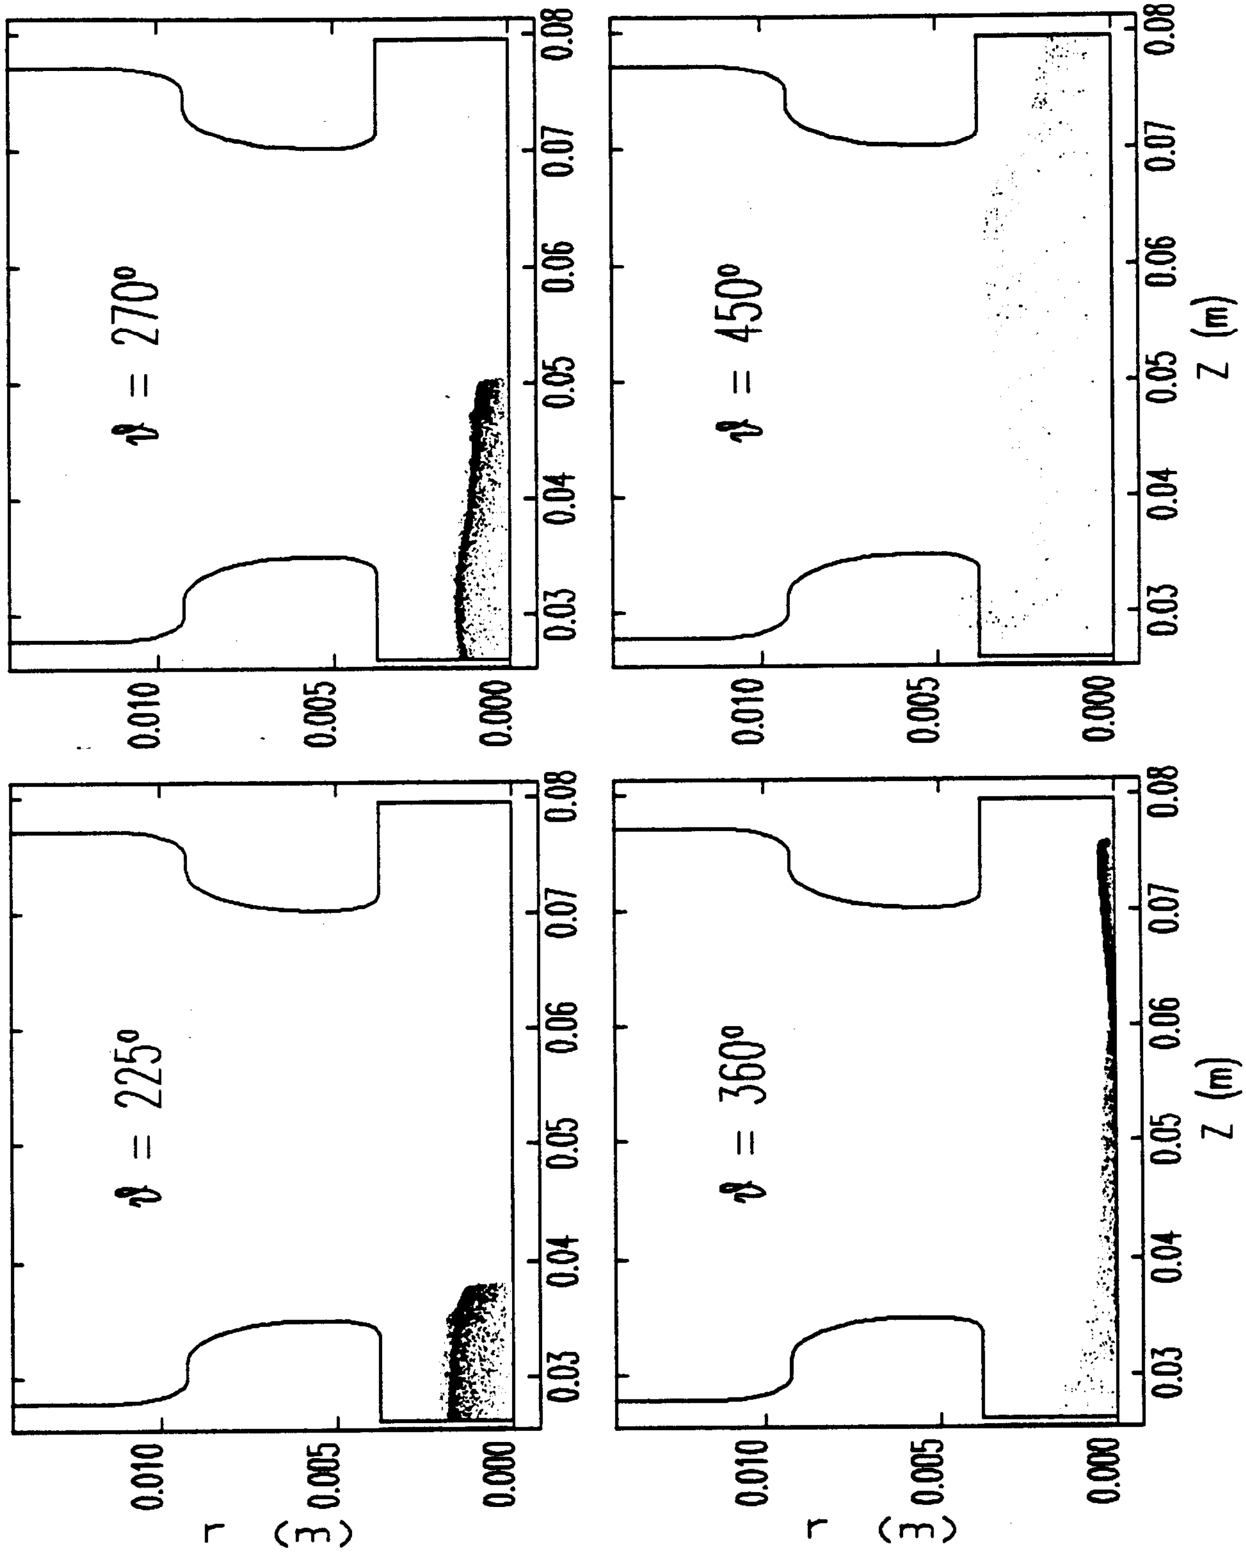

Figure 2.25: MASK Beam Snapshots at Various Phases—Second Cell

increasing r (see Sections 2.3.11 and 2.3.12).

Notice that the lead particles in the beam have not moved more than about 6mm after the first 90° of the cycle, while after another 90° of the cycle, the lead particles are about to exit the first cell, having traveled more than 25mm. This testifies to the rapid increase in velocity that accompanies the particles' trip across the first cell. Indeed, for $E_{p2} = 75$MV/m, the lead particles have $\beta \approx 0.8$ upon exiting the first cell. Note that the beam travels most of the length of the second cell in 180°, even though the second cell is twice the length of the first. After 270° of the cycle have passed, the RF in the first cell is at the decelerating crest. One sees that there are still many particles in the first cell. These particles are in fact being accelerated back into the cathode.

The relatively slow initial motion of the particles also underlines the importance of the cathode region in determining beam properties, since particles spend a disproportionate amount of time in the region of the cathode. It is this initially slow motion in a region with large non-linear fields that, for example, leads to the large effect of the non-linear terms in E_z on the momentum spread. As the beam travels through the first 5mm of the first cell, the front edge of the beam takes on a cupped shape, due to the non-uniformity of the longitudinal field across the cathode.

The Figures clearly show the effect of the focusing noses in producing a converging beam: the beam radius has decreased by a factor of about 2 by the time the beam exits the first cell. If the focusing noses were not in place around the cathode, the beam would fill the aperture of the beam tube. With the focusing noses, the beam converges and passes easily through the beam tube, even for very high current densities. The transverse beam size continues to decrease as the beam travels through the second cell, partially due to additional focusing forces encountered in passing the first cavity nose in the second cell. These are visible in Figure 2.25 by virtue of the "kink" they produce in the radial beam envelope, as seen in the graph for 270°.

2.3.7 Calculating Emittance from Cylindrical Coordinates

Since MASK (unlike `rfgun`) works in cylindrical coordinates, the dynamics of macroelectrons is calculated in terms of radial, azimuthal, and longitudinal momenta. Since

I do not impose any external magnetic fields (e.g., a solenoidal field along the axis), the azimuthal momenta are identically zero. Hence, the relevant phase-space coordinates of any macro-electron are $(\mathrm{z}, \mathrm{r}, \mathrm{p_z}, \mathrm{p_r})$. Macro-electrons that pass through the origin still have positive r coordinates, but have $\mathrm{p_r}$ reversed in sign. To see how to obtain the emittance, note that the x and y coordinates are related to r by

$$\mathrm{x} = \mathrm{r}\cos\theta_s \quad \text{and} \quad \mathrm{y} = \mathrm{r}\sin\theta_s, \tag{2.62}$$

where the subscript *s* stands for the *spatial* coordinates. Similarly, the x and y momenta are related to $\mathrm{p_r}$ by

$$\mathrm{p_x} = \mathrm{p_r}\cos\theta_p \quad \text{and} \quad \mathrm{p_y} = \mathrm{p_r}\sin\theta_p, \tag{2.63}$$

where the subscript *p* stands for the momentum. Because there is no beam rotation, however, one must take $\theta_s = \theta_p$. The normalized emittance in the x plane is given by equation (1.9), which implies

$$\epsilon_{nx} = \pi \mathrm{m_e c}\sqrt{\langle \mathrm{r}^2\cos^2\theta\rangle\langle \mathrm{p_r^2}\cos^2\theta\rangle - \langle \mathrm{p_r r}\cos^2\theta\rangle^2}. \tag{2.64}$$

Averaging over θ, one sees that

$$\epsilon_{nx} = \frac{\pi}{2}\sqrt{\langle \mathrm{r}^2\rangle\langle \mathrm{p_r^2}\rangle - \langle \mathrm{p_r r}\rangle^2}. \tag{2.65}$$

Clearly, $\epsilon_{nx} = \epsilon_{ny}$.

2.3.8 Tests of the Independent Bunch Assumption

I mentioned above that the MASK simulations are done with the implicit assumption that each bunch is independent. To test this assumption, I simulated the first cell for five cycles with beam (in addition to the cycles necessary to excite the cell). It was not possible to simulate both cells, since the frequency mismatch between the two cells in the simulation would have made the phase between the cells drift, thus obscuring the effect. In any case, one expects that the predominant effect will come in the first cell, where the energy of the beam is lowest and where more charge is present during more of the RF cycle. In order to bring out any effect, I deliberately chose a very high current density of $80\mathrm{A/cm^2}$, though 10–20 $\mathrm{A/cm^2}$ is the range used

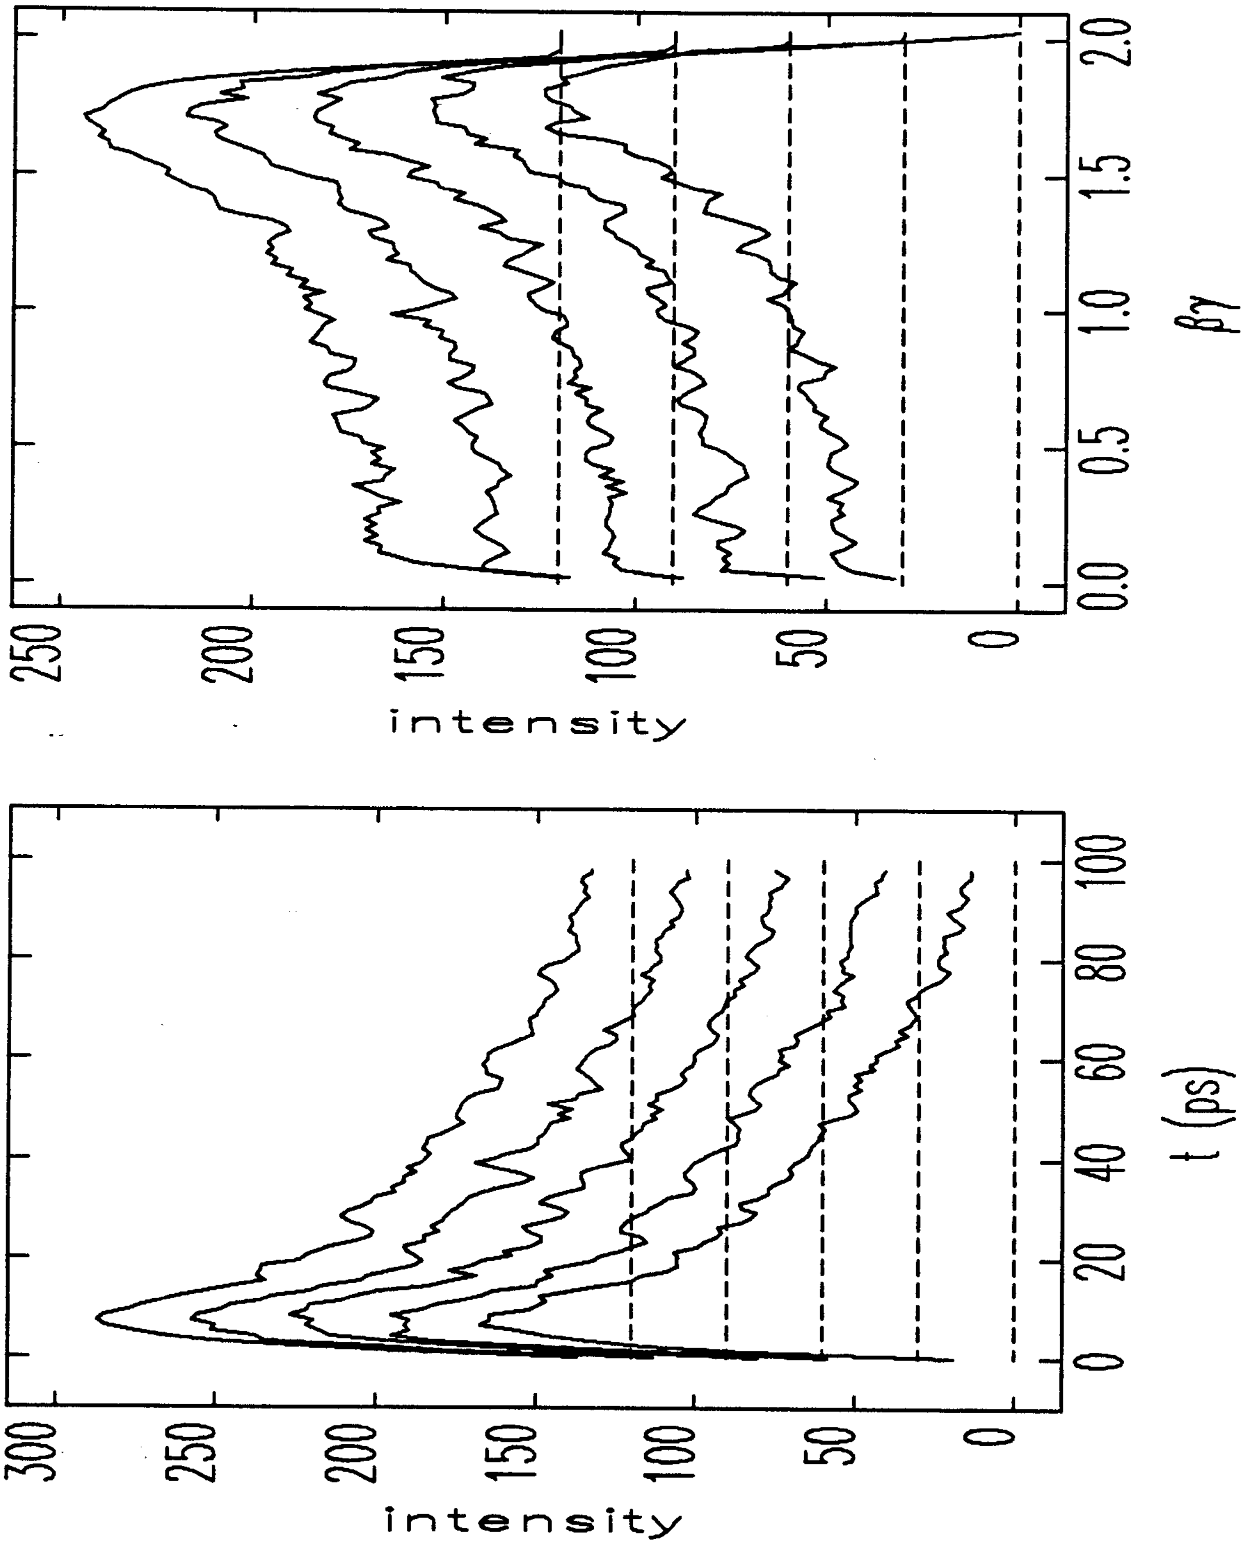

Figure 2.26: Exit-Time and Momentum Histograms as Functions of Bunch Number

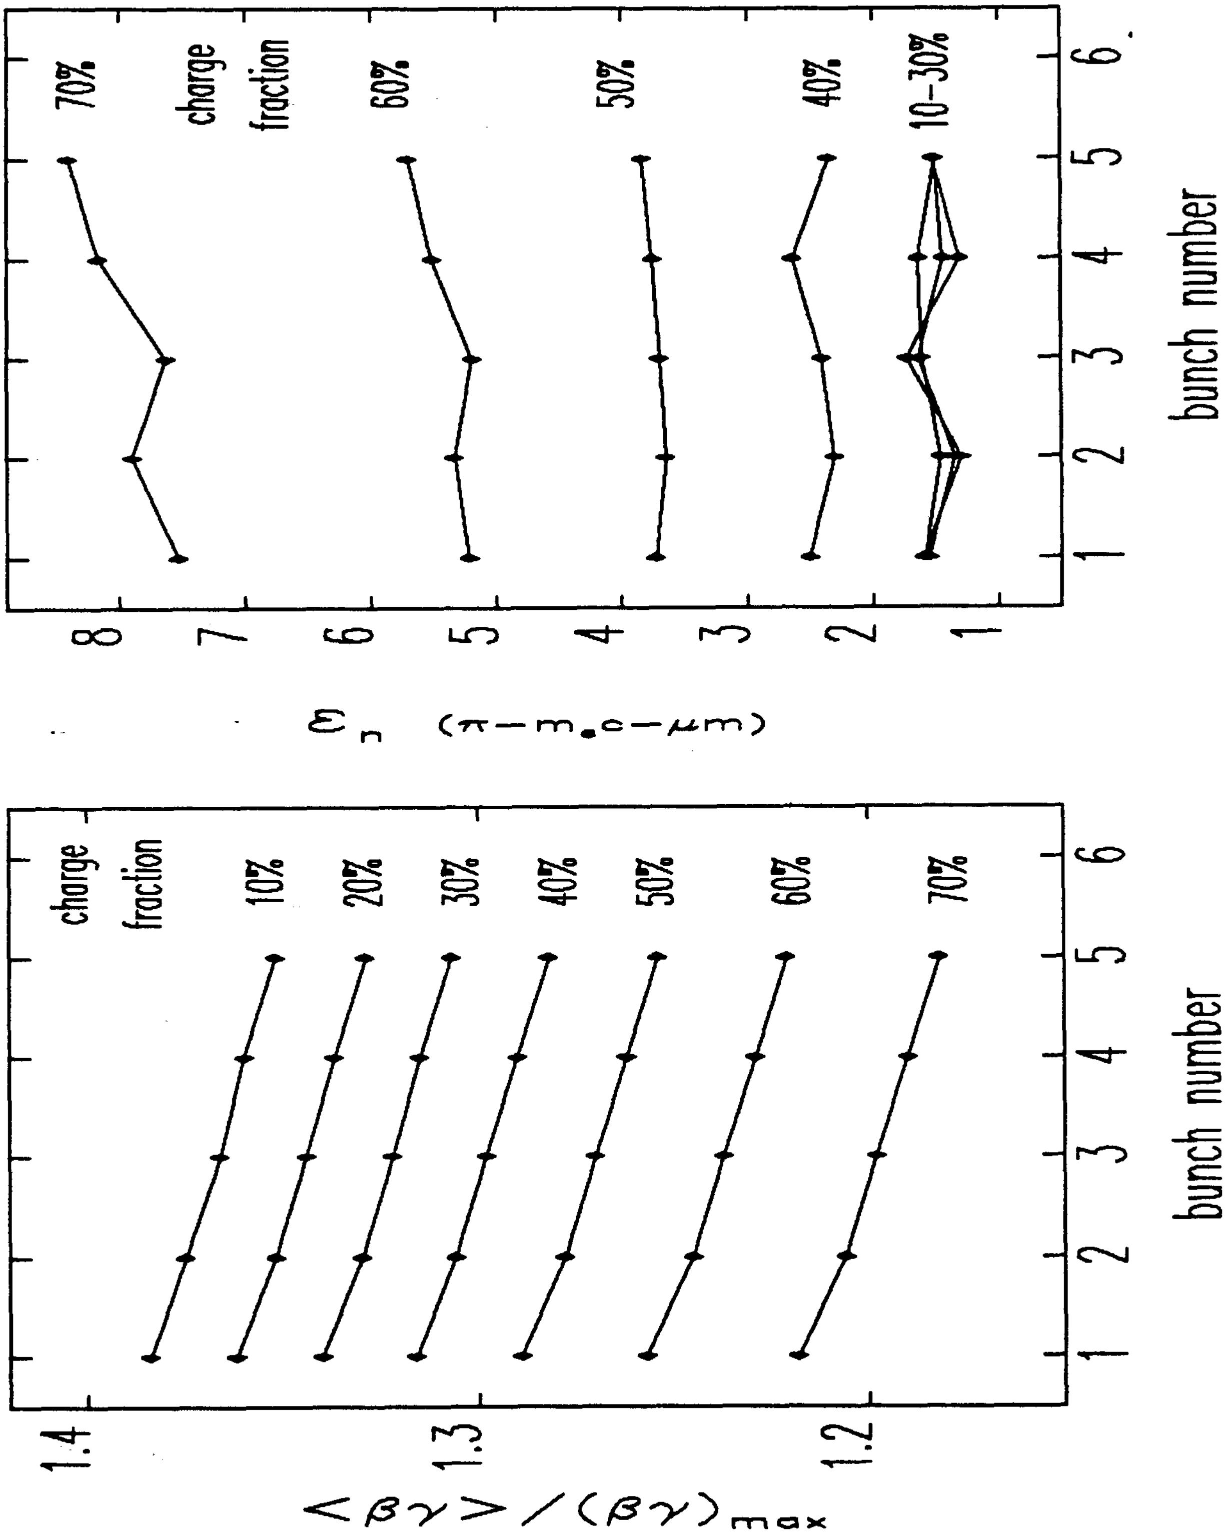

Figure 2.27: Normalized Average Momentum and Normalized Emittance as Functions of Bunch Number

in actual operation of the gun at SSRL. In Figure 2.26, I show histograms of exit-time (relative to the first particle in the bunch) and momentum for the five bunches, offset for clarity, with later bunches being offset by a larger, positive amount. No dramatic changes are evident, though the momentum peak is clearly occurring at smaller momenta for later bunches. This is a result of the extraction of energy from the cavity by previous bunches.

Figure 2.27 shows additional data. Here, I am compare the normalized average momentum, $\langle\beta\gamma\rangle/(\beta\gamma)_{\mathrm{max}}$, and the normalized emittance for successive bunches, for various fractions of the total charge in the first bunch, starting with the most energetic particles in each bunch. This is less ambiguous than using momentum-spread intervals, since in the present case the average momentum and the momentum distribution are changing. This is roughly equivalent to choosing the same initial phase interval (ignoring longitudinal mixing caused by non-linear E_z terms.)

One sees that the effect on the normalized emittance is not dramatic, producing a spread of about 0.5 $\pi \cdot m_e c \cdot \mu$m and no clear trend toward an increase for less than 60% of the charge in the beam. There is an trend in the normalized average momentum, but the regularity of the trend suggests that it is simply a result of the natural change in the momentum distribution as a function of field level in the cell. If the effect were due to high-order modes, one would expect it to display less regularity. Without proving this, I believe the data presented confirm the reasonableness of using MASK in the single-bunch, assumed-steady-state mode, especially for current densities significantly less than 80 A/cm^2.

2.3.9 Transverse Beam Evolution

In order to better understand the bunch evolution within the gun, I have done MASK runs with "emittance windows" at various locations inside the gun. These windows are user-defined planes of constant z, such that whenever a macro-electron passes one of the planes while traveling in a specified direction (i.e., toward positive z or negative z), MASK dumps the macro-electron's phase-space coordinates (i.e., radius, time, and radial and longitudinal momentum). Because MASK checks the z coordinate of each macro-electron against the z coordinate of each window at every time step, using too many windows is expensive in terms of CPU time. Hence, I placed only four windows in the first cell and four in the second. In the first cell, one window was in front of the cathode, to keep track of back-bombarding macro-electrons. The other three were equispaced by $\Delta z = \lambda/12$, being placed at $z = \lambda/12, \lambda/6$, and $\lambda/4$. The last of these was used to record macro-electron coordinates for re-injection into second-cell simulations. In the second cell, the windows were similarly placed. One window was at $z = \lambda/4$, in order to keep track of macro-electrons that get back-accelerated into the first cell. The others equispaced by $\Delta z = \lambda/6$, being placed at $z = \lambda/4 + \lambda/6, \lambda/4 + \lambda/3$, and $3\lambda/4$. The last window was used to record the macro-electron coordinates at the exit of the gun.

Figure 2.28 shows the MASK-calculated transverse phase-space distributions in the first cell from a simulation with $J = 10^{-5}$A/cm^2 (i.e., essentially turning off the space-charge) and $E_{p2} = 75$MV/m. Comparison with Figure 2.18 shows that the MASK results are qualitatively similar to the `rfgun` results, but quantitatively different.

2.3.10 Accuracy of MASK Field Calculations

These differences are a result of differences in the fields calculated by `rfgun` and MASK. Recall that `rfgun` uses an off-axis expansion, starting from SUPERFISH-calculated values of $E_z(z, r = 0$.

Figure 2.29 shows derivatives of MASK- and SUPERFISH-calculated fields for the mesh spacings given earlier in this chapter (i.e., $\Delta z = 0.43$mm and $\Delta r = 0.25$mm). SUPERFISH predicts somewhat smaller $\partial_z^2 E_z(z, r = 0)$ and $\partial_z^3 E_z(z, r = 0)$ near the cathode and in the vicinity of the cell noses.

Figure 2.30 compares the longitudinal fields calculated by MASK to those calculated by URMEL[65], showing that $E_z(z, r)$ has a increasingly large apparently linear term in z as r increases toward R_c. Thus, the fields calculated by MASK do not exactly satisfy (2.54). (This comparison could not be done between MASK and SUPERFISH, because SUPERFISH uses an adaptive, and hence irregular, triangular mesh, which makes it difficult to obtain the off-axis fields. The URMEL fields could not be used in `rfgun` because they are too noisy to permit accurate higher-order

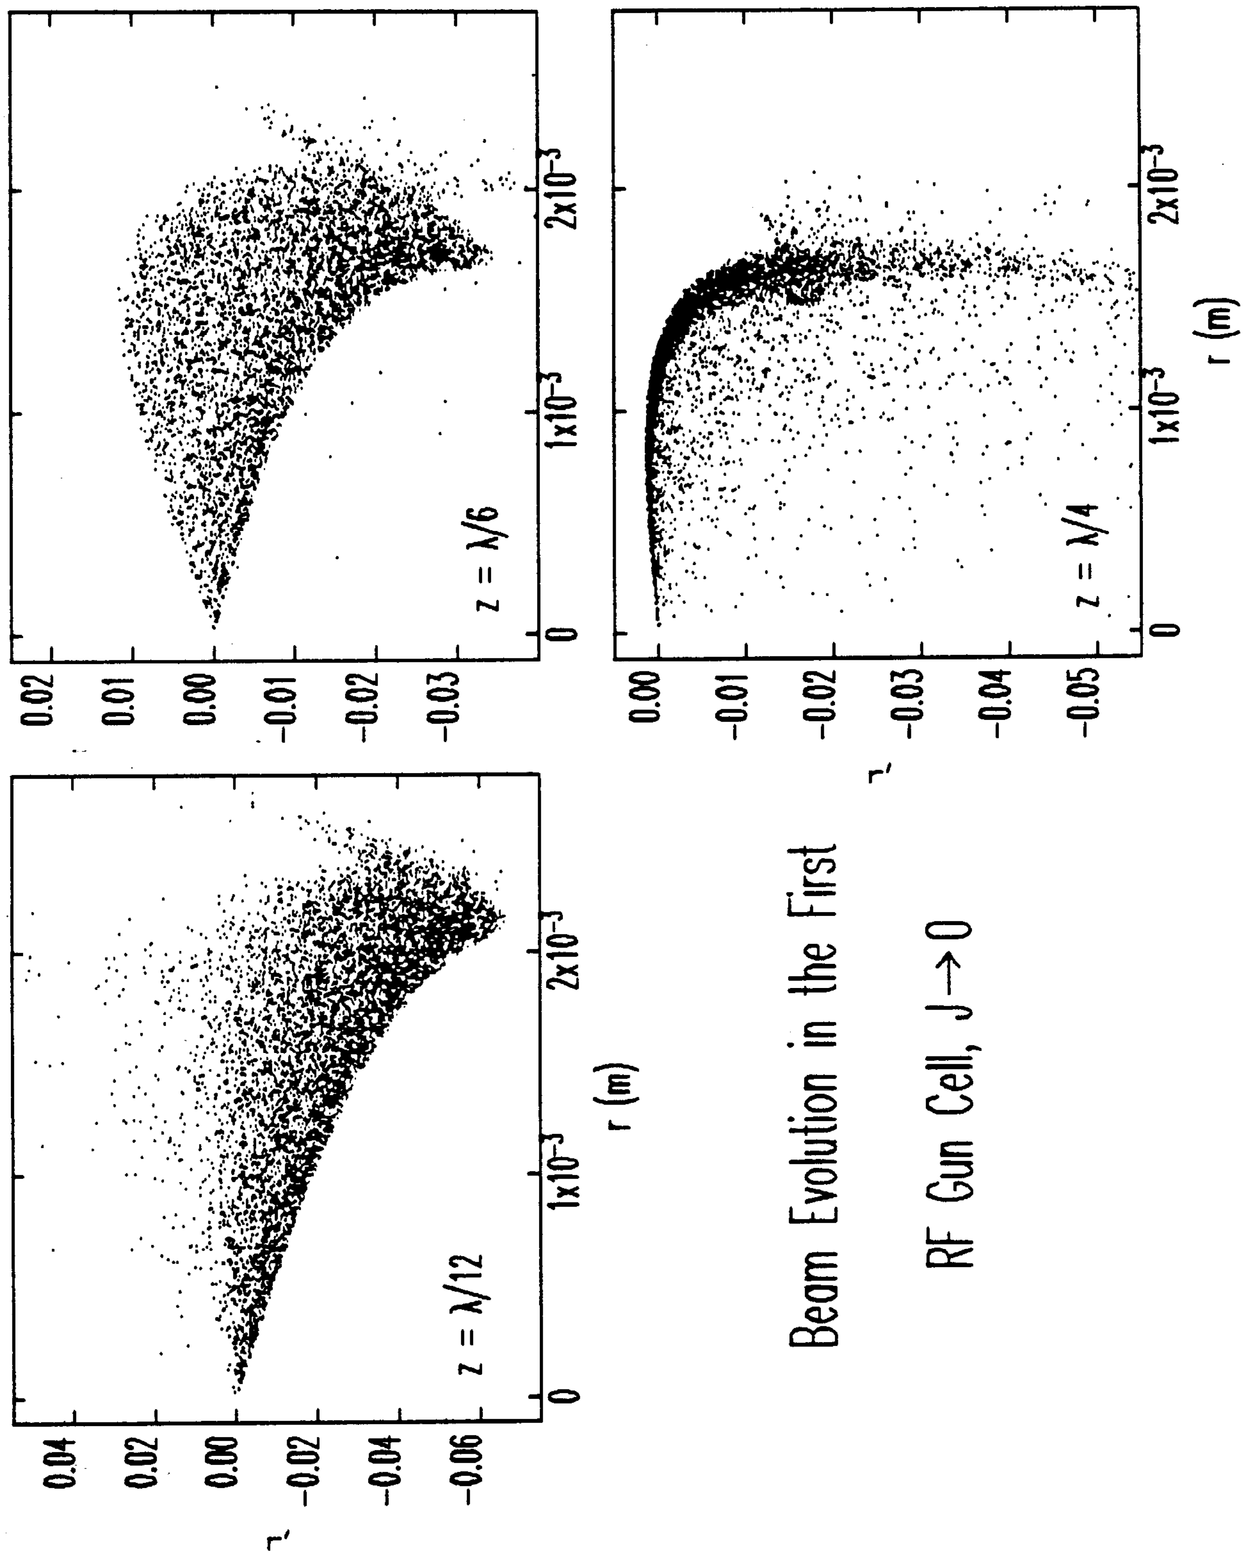

Figure 2.28: MASK Transverse Phase-Space Evolution in the First Cell, for J → 0

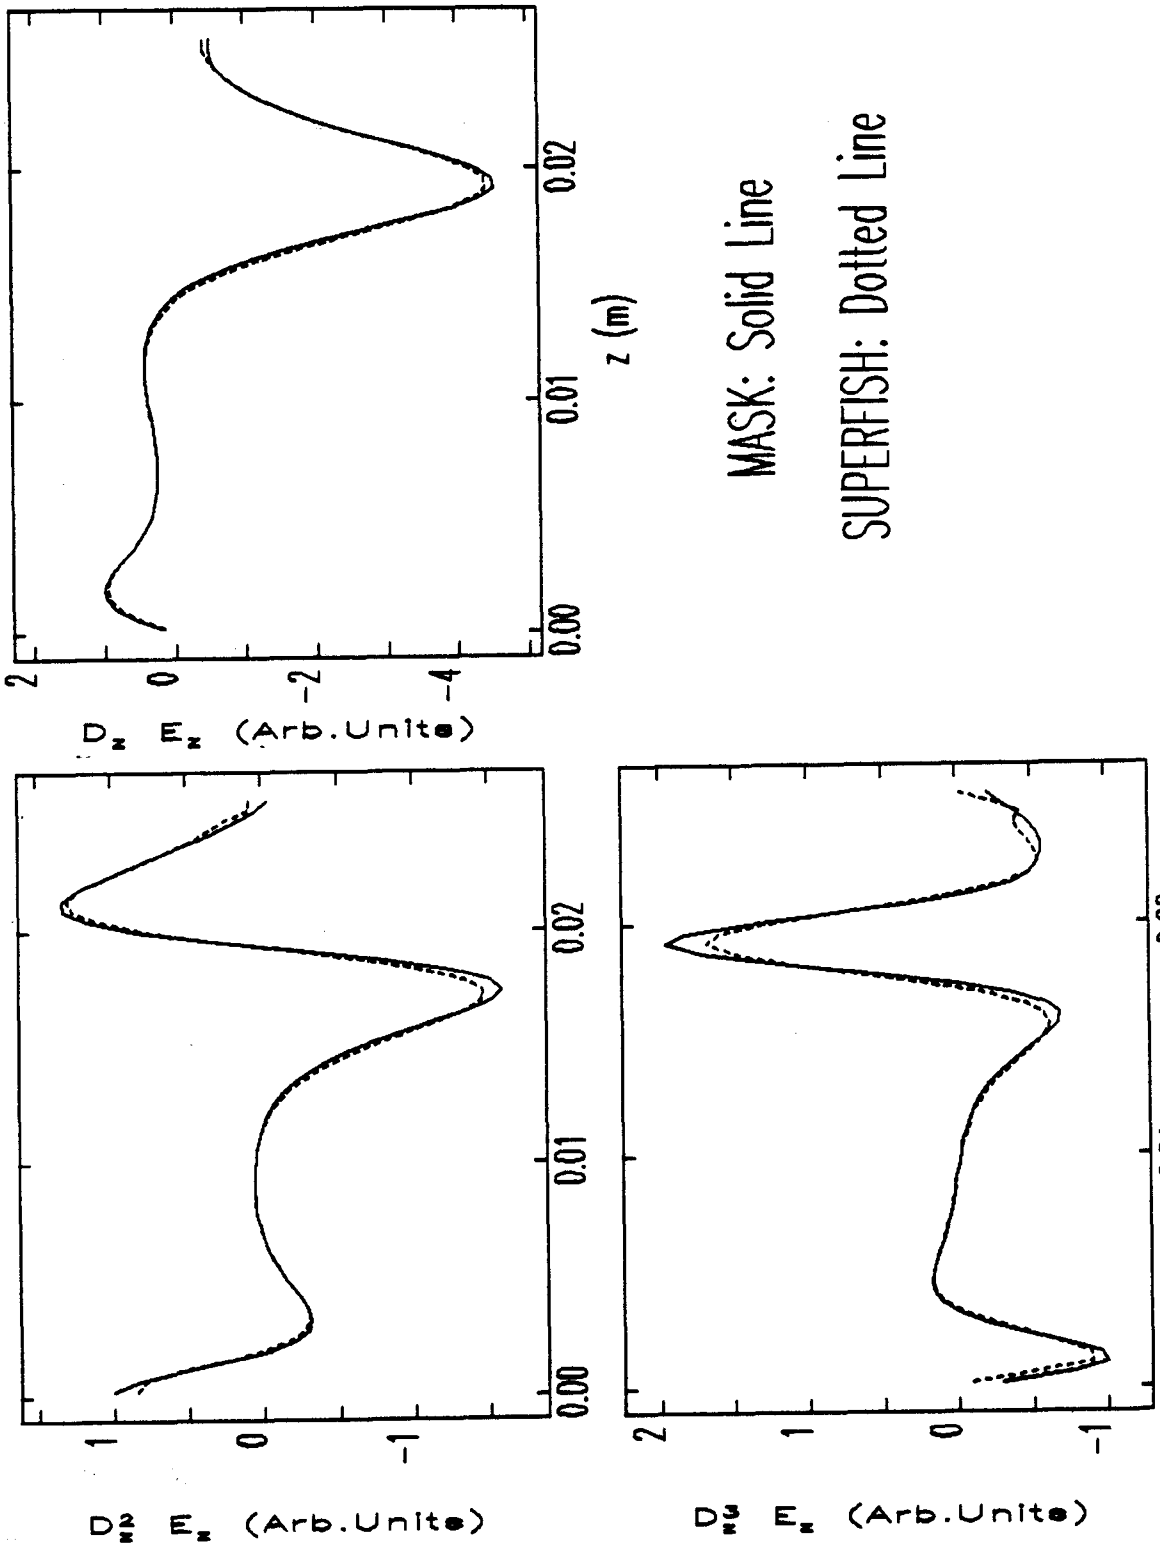

Figure 2.29: Comparision of Derivatives of On-Axis Longitudinal Fields in the First Cell as Calculated by MASK and SUPERFISH

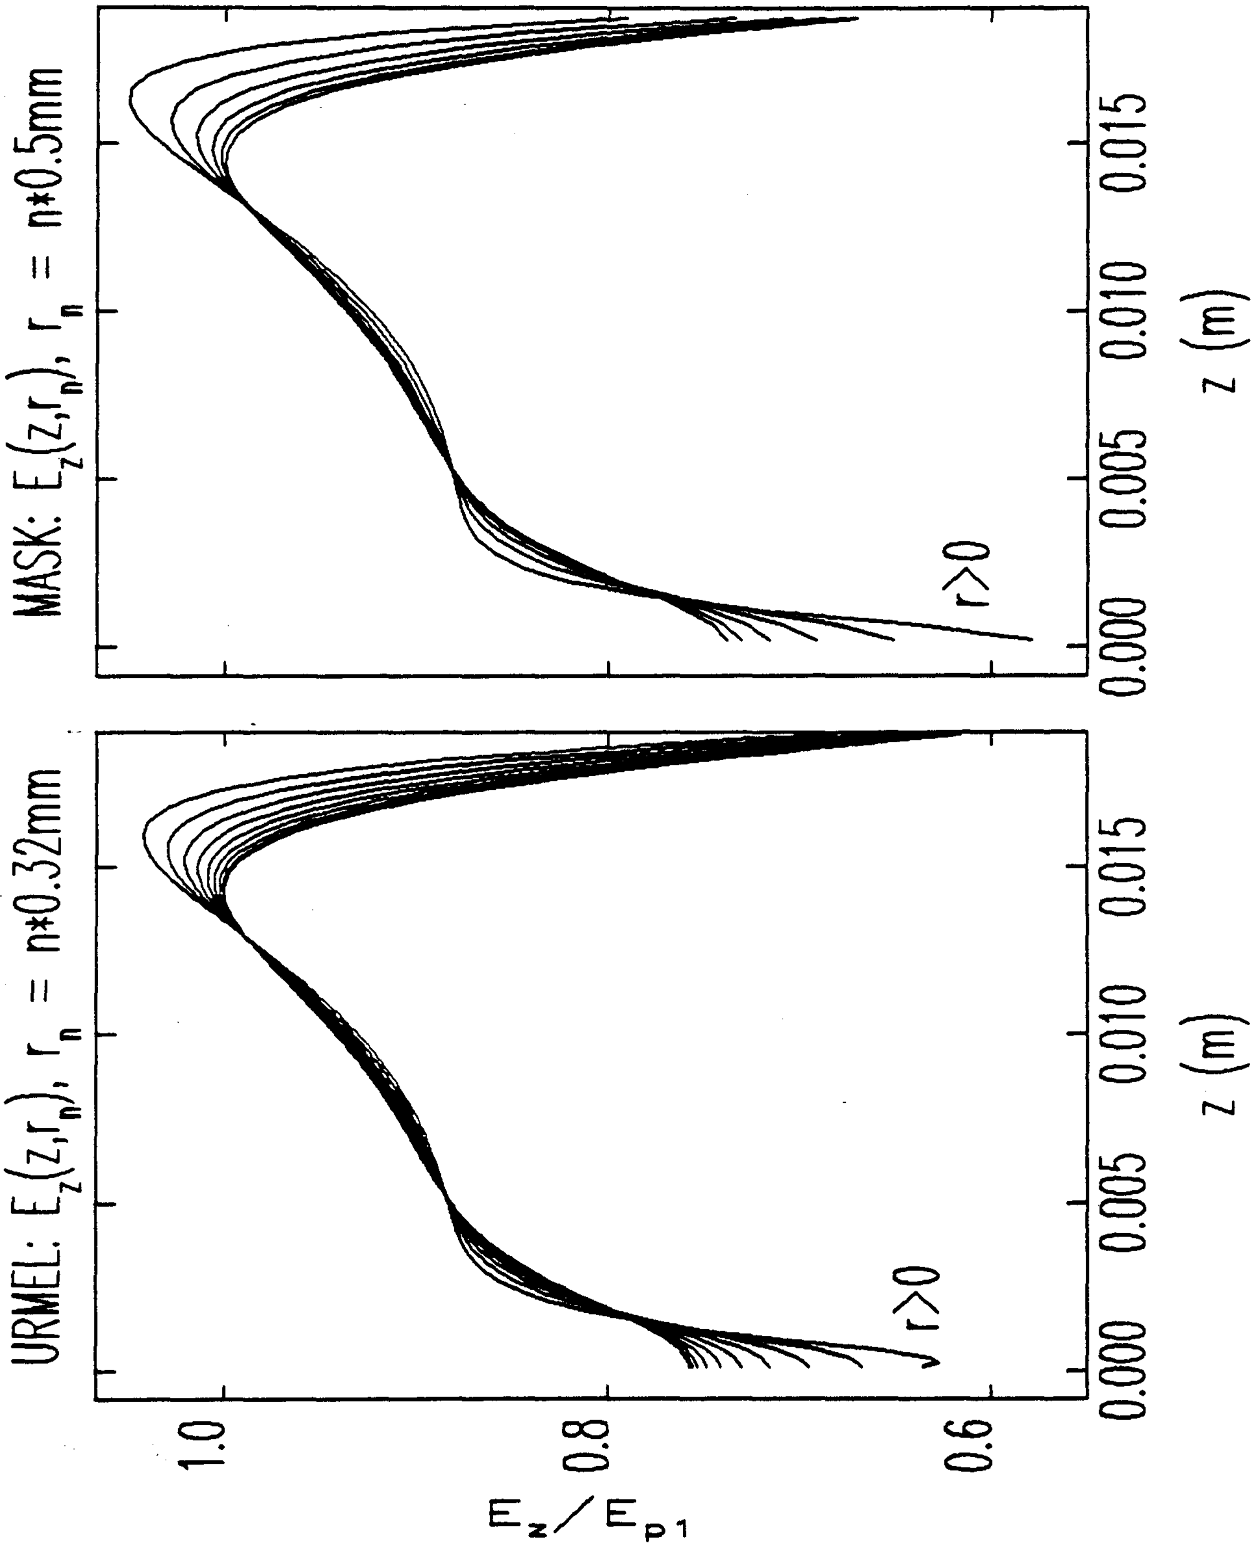

Figure 2.30: Longitudinal Fields in the First Cell as Calculated by MASK and URMEL

numerical differentiation.) If the longitudinal mesh-spacing is decreased in MASK, the MASK and URMEL results come into much better agreement, but the predicted particle motion is not greatly changed. Hence, another explanation must be sought for the differences between MASK and **rfgun** predictions.

2.3.11 Adequacy of Off-Axis Expansion

The real source of the discrepancies between MASK and **rfgun** is the inadequacy of the off-axis expansion used in **rfgun**, which for E_r is only third order in r. Figure 2.31 shows four different calculations of $E_r(z)$ at $r = 2.87$mm, which is close to the cathode radius (3mm). An explicit MASK result was obtained by running MASK with a finer mesh ($\Delta z = 0.215$mm and $\Delta r = 0.164$mm) and sampling $E_r(z)$ at r=2.87mm directly; MASK predicts larger a E_r than any of the other calculations. The other three curves are various calculations of $E_r(z)$ at the same radius using off-axis expansions. Two third-order calculations are shown, one starting with SUPERFISH-calculated on-axis longitudinal fields, $E_{z0}(z)$, the other starting with the same result from MASK. The other expansion is a fifth-order expansion starting with the SUPERFISH-calculated $E_{z0}(z)$ (the MASK data was too noisy to allow a fifth-order expansion, as the noisiness of the third-order expansion shows). Clearly, the fifth-order expansion is the closest to the explicit MASK result. Hence, the conclusion is that MASK is making more accurate predictions of the effect of non-linear fields than **rfgun** is making.

In order to test this diagnosis, I used the $E_z(z, r = 0)$ profile from MASK in **rfgun**, and repeated some of the analysis done above. Figure 2.32 shows the normalized RMS emittance and the normalized brightness for the two cases. A significant, though hardly dramatic, change in the predicted emittance is obtained when using the MASK-calculated on-axis field profile. For larger initial-phase intervals, the predicted emittance is smaller, while for smaller initial-phase intervals, it is larger. The brightness follows the opposite pattern, as expected.

Figure 2.33 shows a comparison of MASK results for $J \to 0$ with **rfgun** results obtained using the MASK-calculated $E_z(z, r = 0)$. One sees that MASK predicts larger emittances than **rfgun**, though not dramatically larger. The reason for this

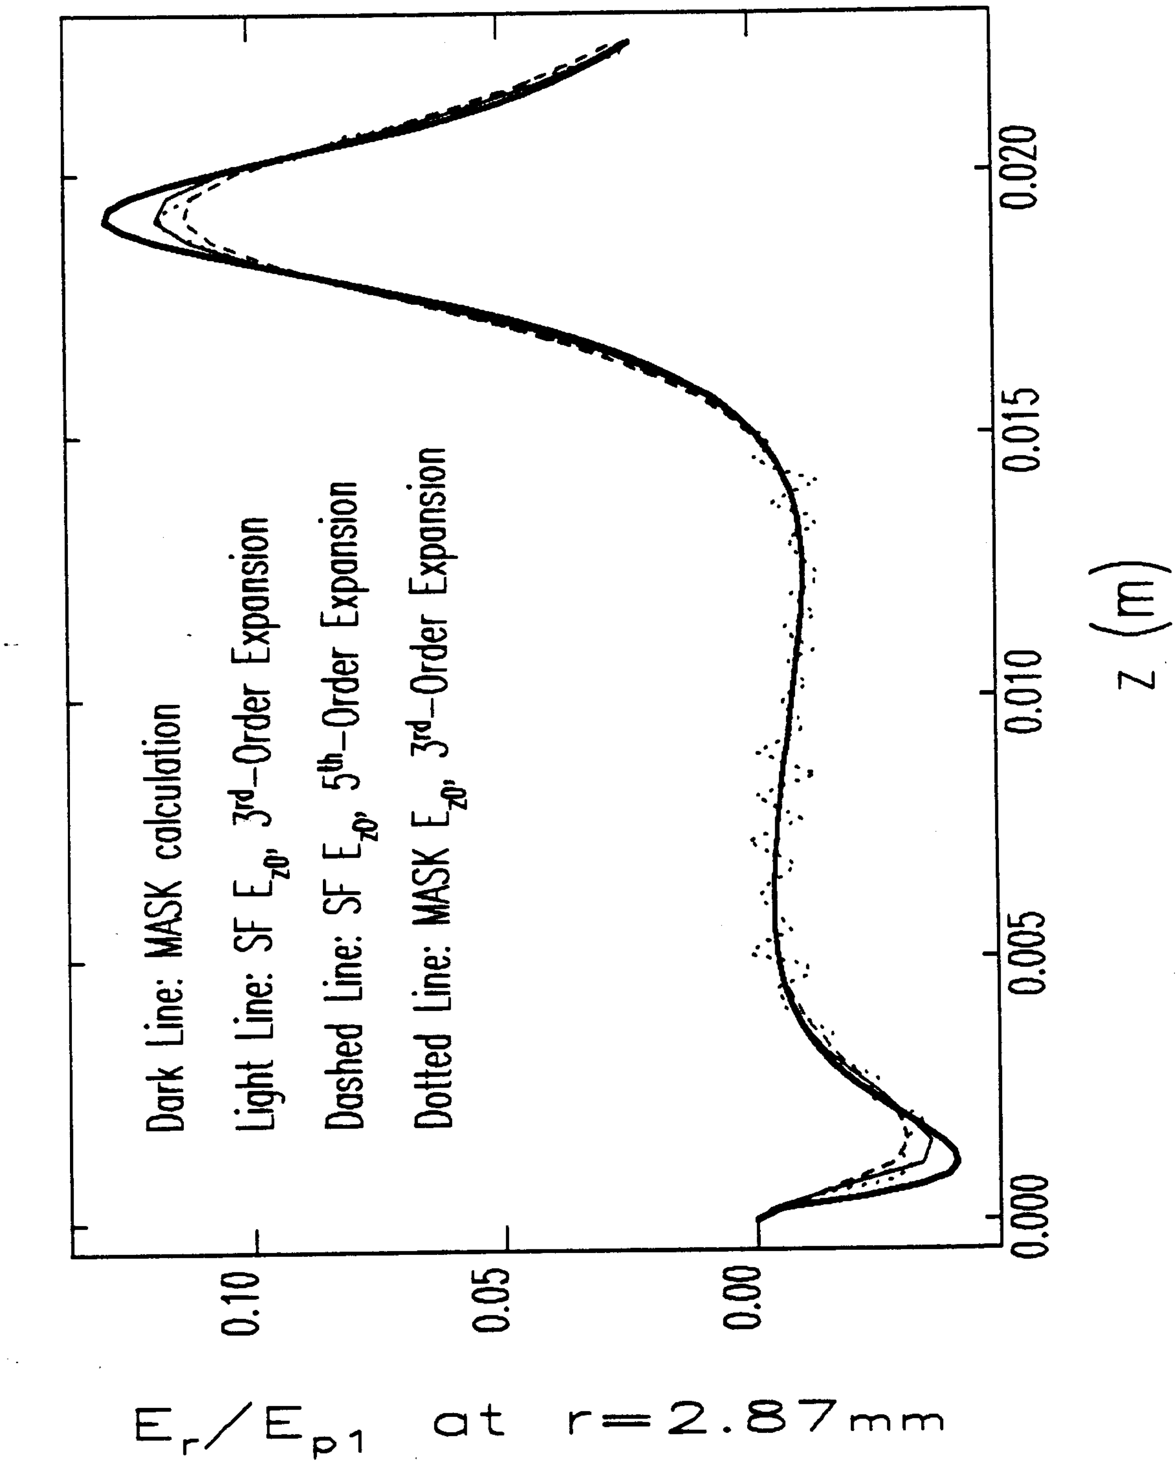

Figure 2.31: Comparison of $E_r(z)$ for r=2.87mm, as Calculated by MASK and Using Off-Axis Expansions of Various Orders

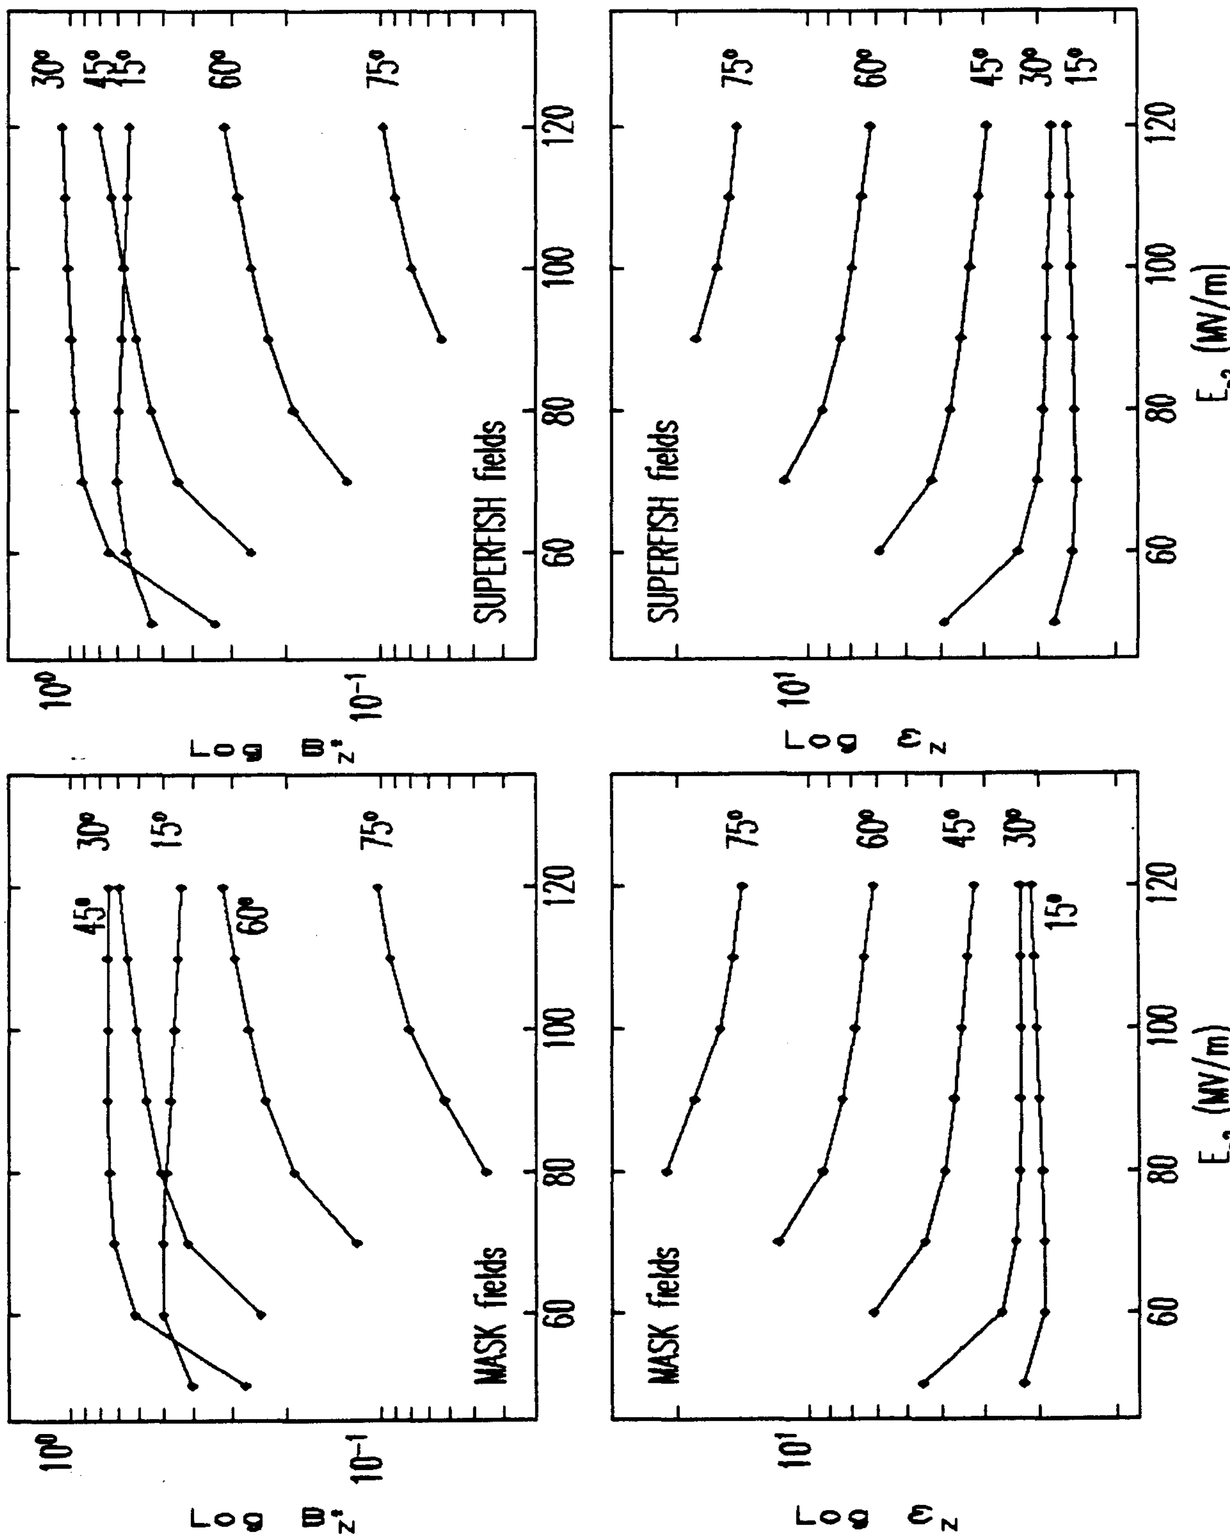

Figure 2.32: Comparison of **rfgun** Emittance Predictions for Various Initial Phase-Intervals, for MASK- and SUPERFISH-Calculated Fields

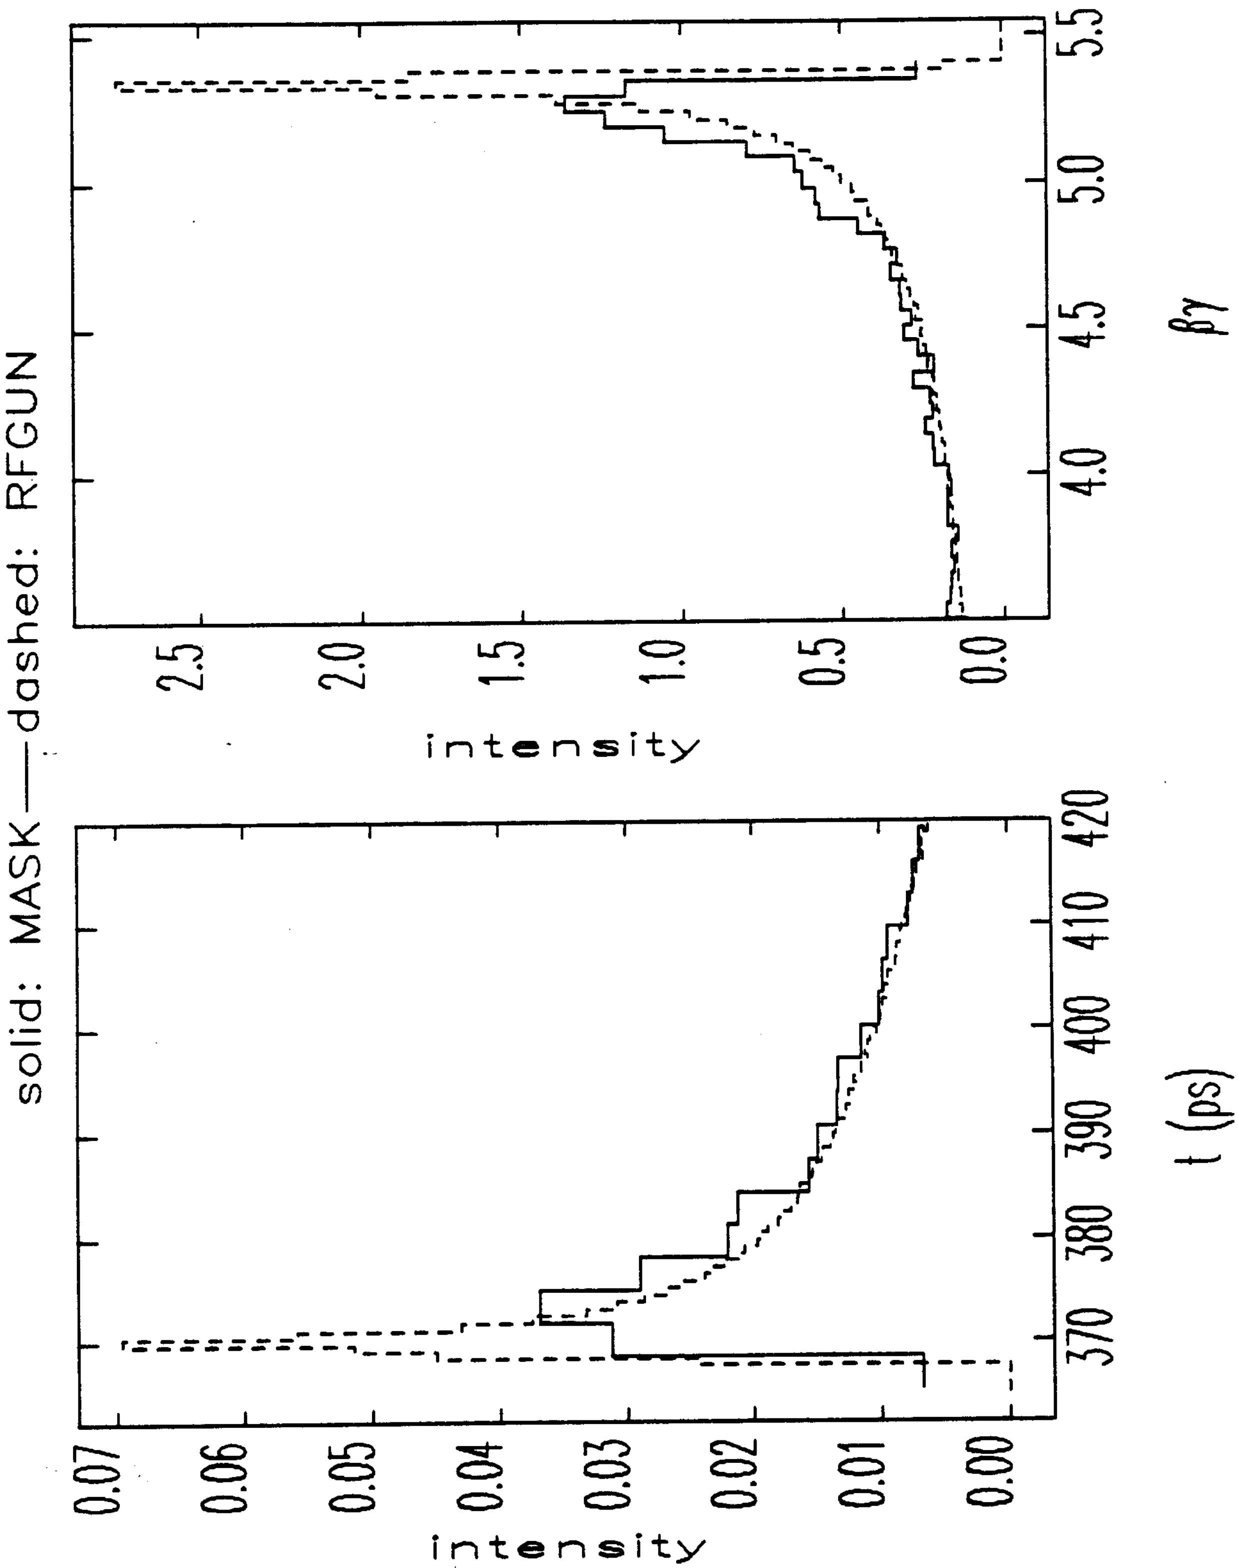

Figure 2.33: Comparison of `rfgun` results for MASK-Calculated Fields with MASK Calculations for $J \rightarrow 0$

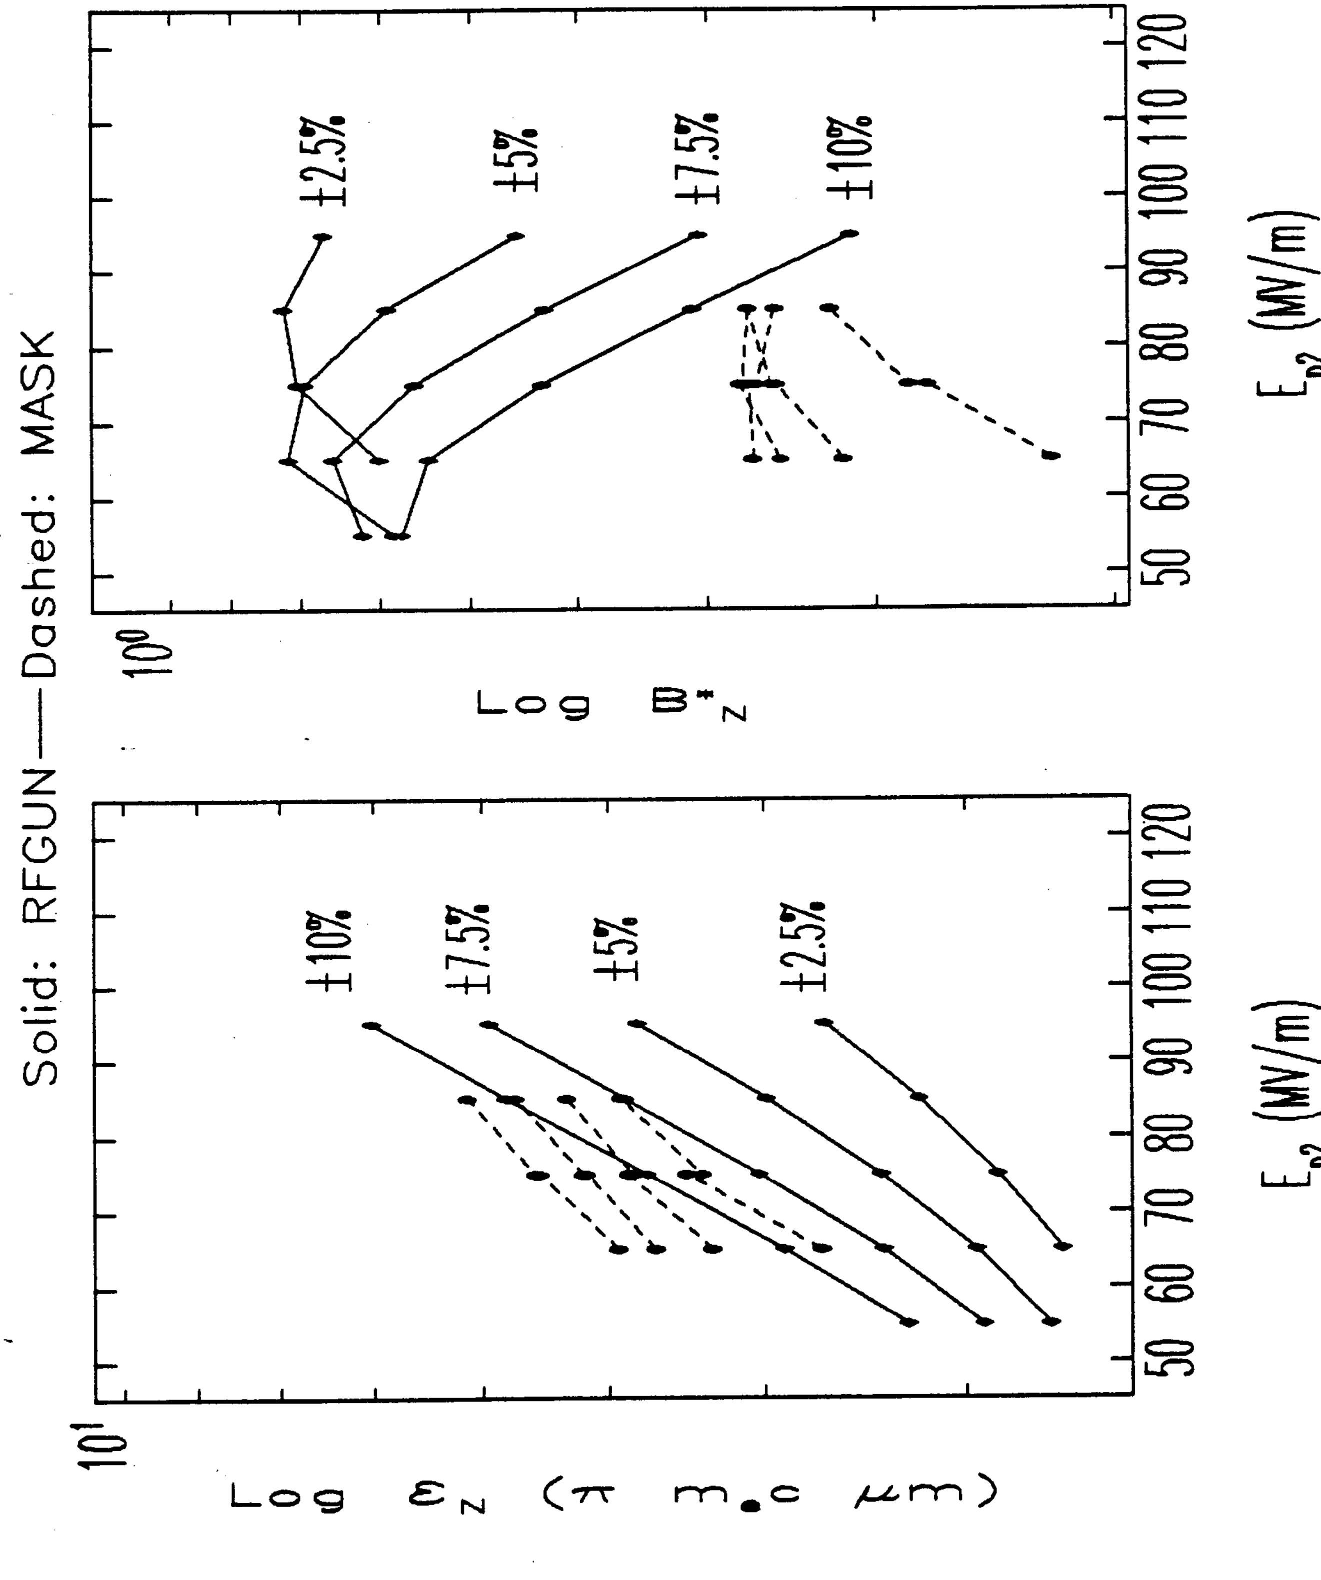

Figure 2.34: Comparison of **rfgun** results for MASK-Calculated Fields with MASK Calculations for $J \rightarrow 0$

discrepancy is that, even using the MASK-calculated on-axis field profile, **rfgun** cannot duplicate the fields used in MASK, since the expansion used in **rfgun** is not of sufficiently high order.

Next, I show in Figure 2.34 a comparison of the momentum and time distributions for MASK and **rfgun** with $E_{p2} = 75$MV/m and $J \rightarrow 0$ in MASK, where the MASK-calculated fields are again used in **rfgun**. As one would expect, **rfgun** predicts narrower spectra because of the larger non-linear fields in MASK. I have also found that MASK consistently predicts about 3% greater maximum momentum than is predicted by **rfgun** for the same value of E_{p2} (with $a = 3$ and $f = 2836$MHz in **rfgun** in order to match MASK). (Because of this discrepancy, I ran **rfgun** with $E_{p2} = 77.1$MV/m in order to match the peak momentum to that of **rfgun** for comparison of the time and momentum spectra in Figure 2.34.) One possible explanation for this is that the phasing of the first and second cell is imperfect. Because of the frequency mismatch between the two cells (discussed in the previous section), the cells drift out of frequency by about 0.5° during one RF period. This would seem to be too small to have the observed effect, however.

Another confirmation of the effect of the larger non-linear fields in MASK is obtained by running MASK with a smaller cathode. While I will not take the space to show these results, I have found that running MASK with $R_c \rightarrow R_c/2$ produces a noticeably smaller momentum and time spread in the final beam. Quantitative results can be found in Chapter 4.

2.3.12 Space-Charge Effects

Transverse Phase-Space

I turn now to the effects of space-charge as predicted by MASK. Figure 2.35 shows the evolution of the transverse phase-space in the first cell for $E_{p2} = 75$MV/m and $J = 80$A/cm^2. These are to be compared to those shown in Figure 2.28 for $J \rightarrow 0$. The effect of the space-charge forces for this high current density are clearly evident. From these two figures, it is apparent that the space-charge forces tend to counter the cavity fields, since the slopes are significantly more positive for $J = 80$A/cm^2. This is as expected, since the space-charge forces are radially defocusing. One effect of this

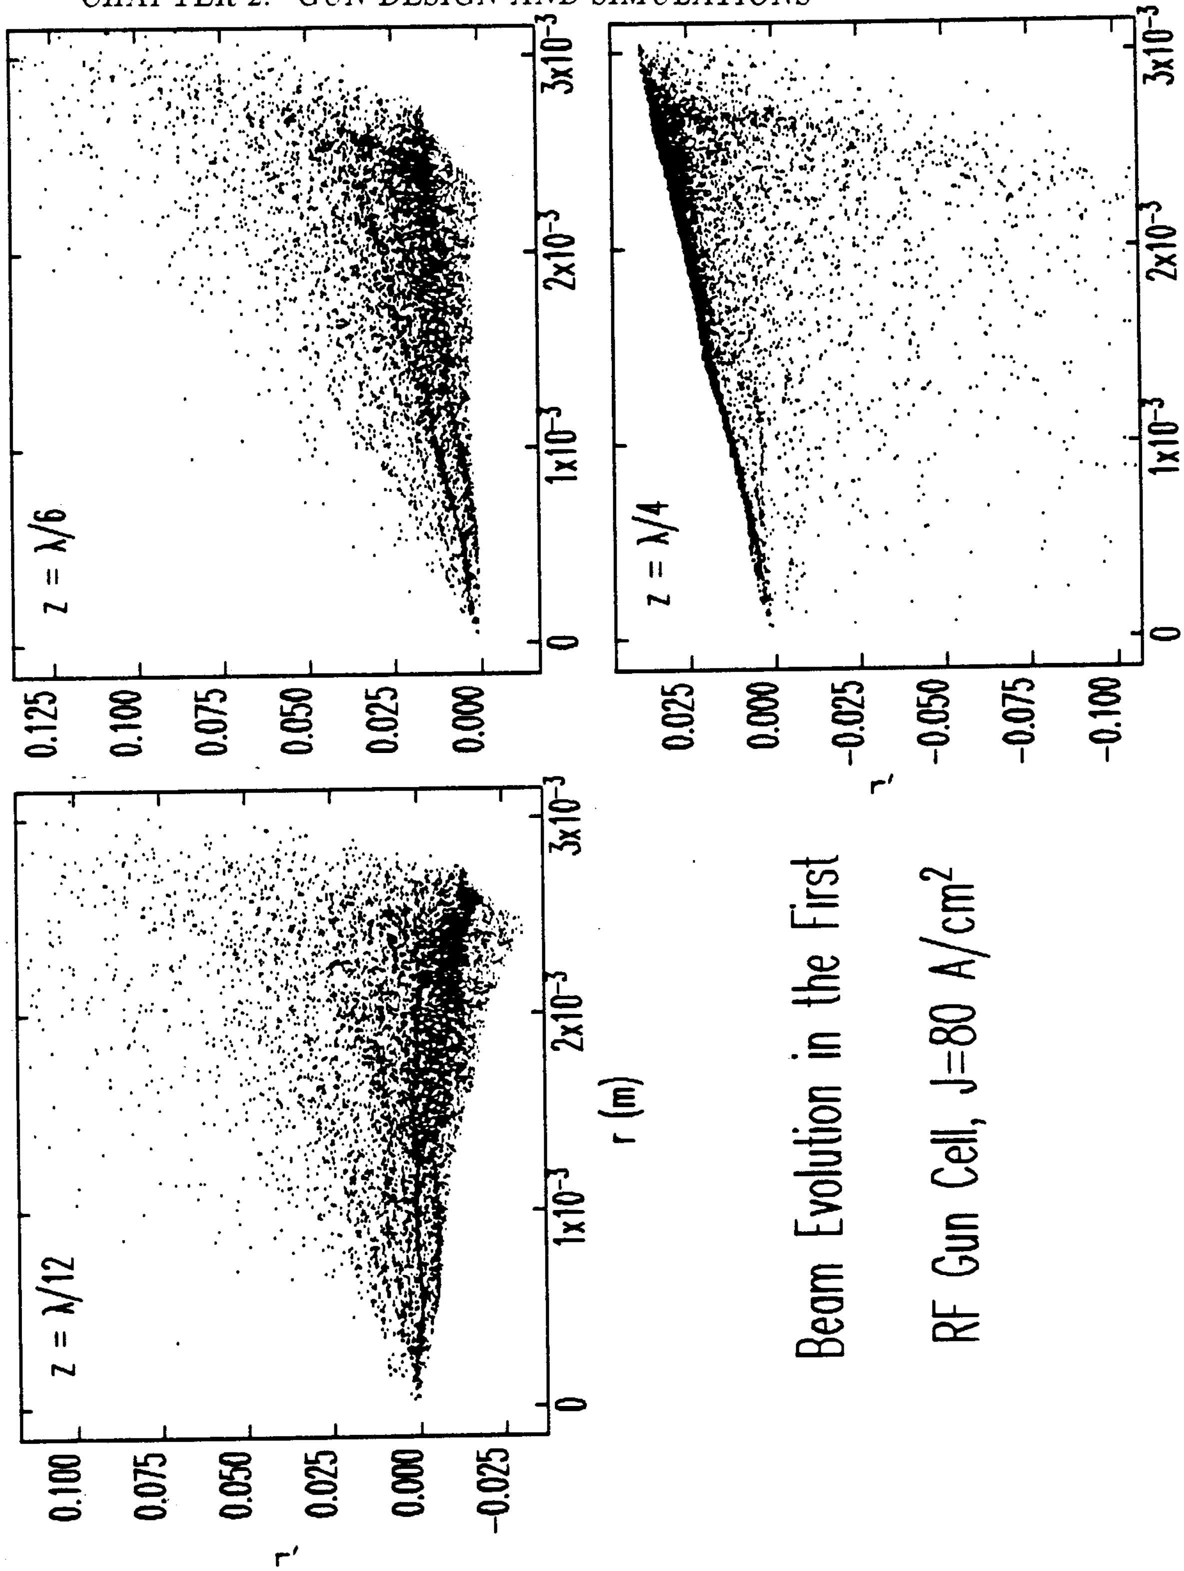

Figure 2.35: MASK-Calculated Transverse Phase-Space Evolution in the First Cell, for $J = 80 \mathrm{A/cm}^2$

defocusing is a significantly larger beam size at the exit of the first cell, as well as a larger spread in slopes.

It would appear from the Figure that the space-charge forces not only counter the linear focusing forces, but also compensate for non-linear fields as well. For this to be the case, one would require that at some point the radial distribution of charge be non-uniform and that the charge density increase with radius. To see why, consider that for a longitudinally uniform, cylindrically symmetric beam the radial electric field is given by Gauss's law[31]:

$$E_{r,\mathrm{beam}}(r) = \frac{1}{\epsilon_o r} \int_0^r \rho(\tilde{r})\tilde{r}\, d\tilde{r} \tag{2.66}$$

where $\rho(r)$ is the charge density per unit cross-sectional area. For a uniform radial distribution, $\rho(r) = \rho_o$ out to to some radius R_b (the edge of the beam), after which it falls to zero. Hence, for this case,

$$E_{r,\mathrm{beam}}(r) = \frac{r\rho_o}{2\epsilon_o}, \qquad r \le R_b \tag{2.67}$$

which is simply an additional linear field term.

Next, consider what happens if $\rho(r) = \rho_o(1 + \eta r^n)$ for $r \le R_b$, where n is an integer:

$$E_{r,\mathrm{beam}}(r) = \frac{r\rho_o}{2\epsilon_o}\left(1 + 2\eta\frac{r^n}{n+1}\right) \qquad r \le R_b \tag{2.68}$$

As I discussed in the previous section, $E_{r,\mathrm{beam}}$ must be an odd function of r, and hence n must be even. In order for the space-charge forces to increase at a greater-than-linear rate with radius, η must clearly be positive, which implies that the beam must be somewhat more hollow than a radially uniform beam.

Figure 2.36 shows histograms of the intensity vs radius for $z = \lambda/12$ for the cases $J \to 0$ and $J = 80\mathrm{A/cm}^2$. Each bin in the histograms represents an annulus, with the height of the histogram being proportional to the charge in that annulus. For a uniform distribution, one would expect a linear function of r, since for this distribution the height of the bin that starts at $r = n\Delta r$ is

$$H_n = 2\pi\rho_o \int_{n\Delta r}^{(n+1)\Delta r} \rho(\tilde{r})\tilde{r}\, d\tilde{r} = \pi\rho_o\Delta r^2(2n+1). \tag{2.69}$$

For $J \to 0$, one sees that H_n increases faster than linearly. This is due to the non-linear increase in focusing fields with radius, and might have been anticipated from

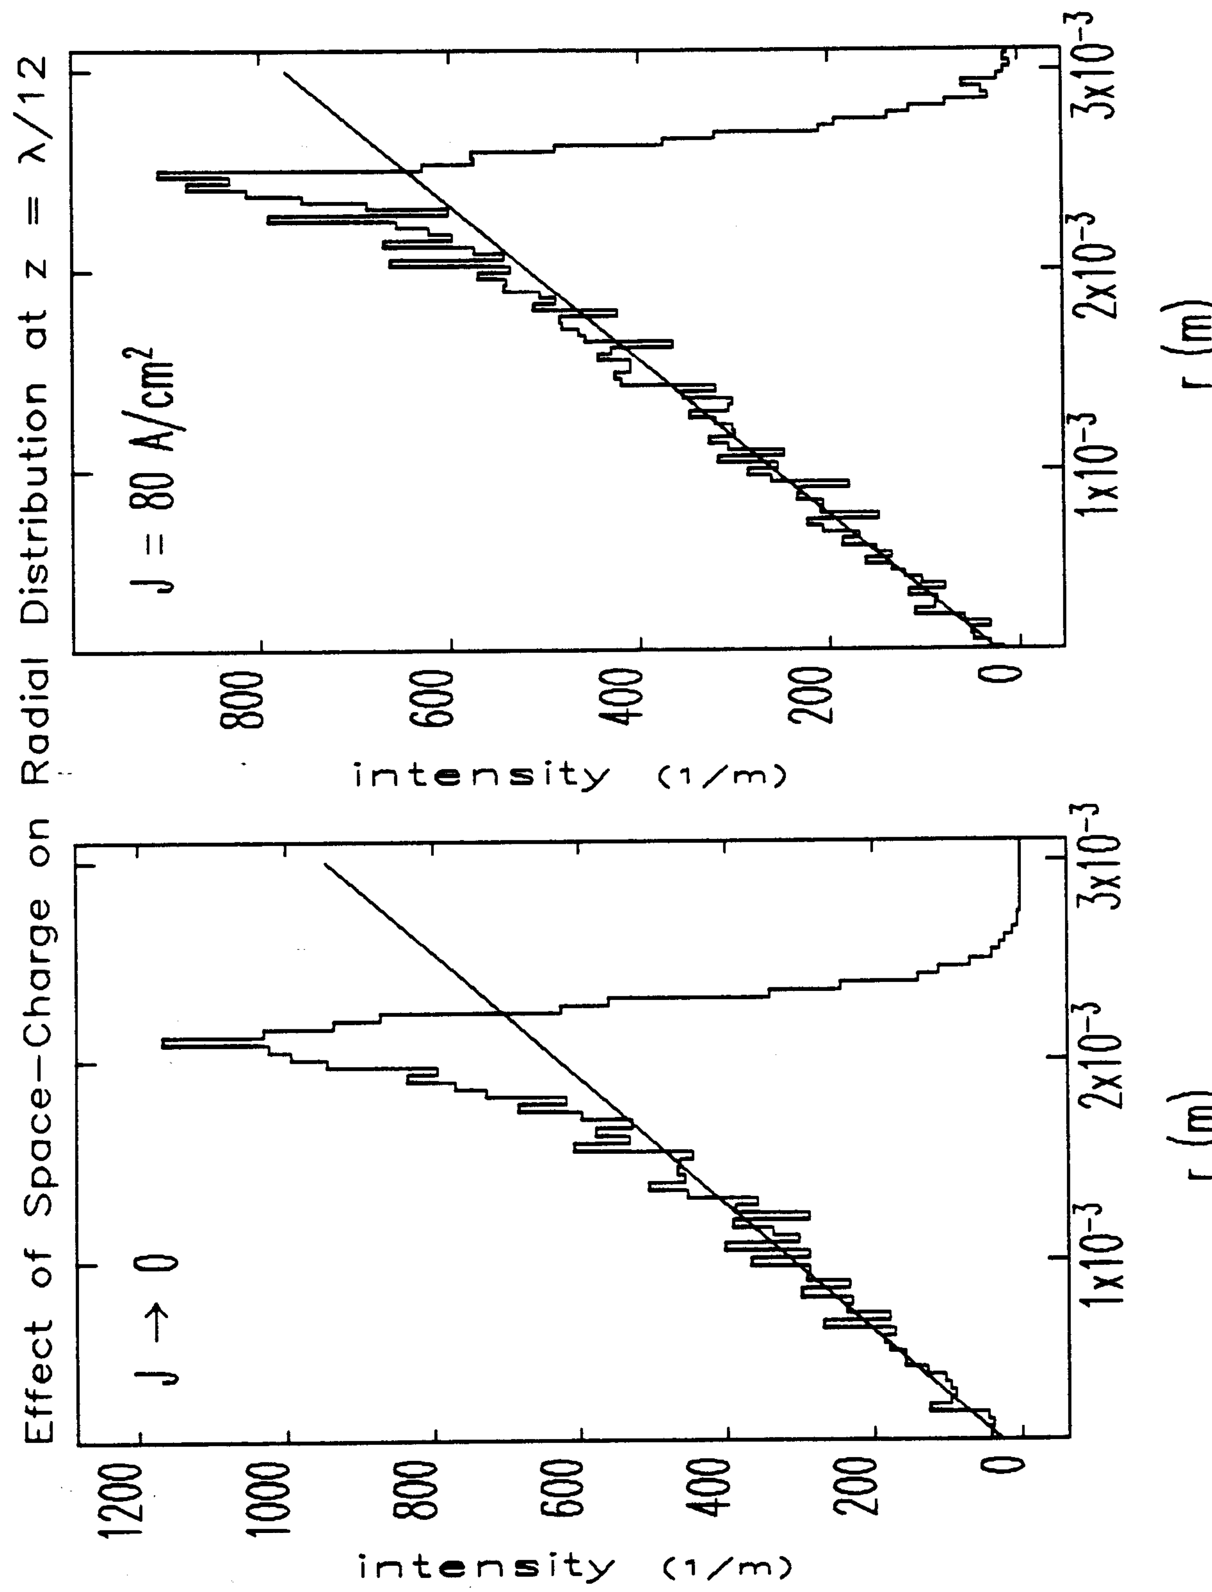

Figure 2.36: MASK-Calculated Transverse Beam Distribution at $z = \lambda/12$, for $J \rightarrow 0$ and $J = 80\text{A}/\text{cm}^2$

Figure 2.28. For J = 80, the distribution is clearly much more uniform. The non-linear cavity fields and space-charge forces in this case tend to balance each other, since the non-uniform distribution that the non-linear cavity fields try to create is just the kind of distribution that is necessary to counter these self-same non-linear cavity fields. One expects the radial beam distribution to have just enough radial non-uniformity to compensate the non-linear cavity fields.

It is not at all apparent from these Figures what the net effect on the emittance is. The beam is larger over-all for the case with high space-charge, but the correlation would also seem to be higher. I will show below that the emittance is in fact substantially larger for the high space-charge case.

Longitudinal Phase-Space

I next look at the effects of space-charge on the longitudinal phase-space. Figure 2.37 shows the effect of space-charge on the longitudinal phase-space at the gun exit. The longitudinal space-charge forces are seen to broaden the momentum and time distributions, much as the non-linear forces do. This is to be expected, since particles at the head of the beam are accelerated by the particles that follow, while trailing particles are decelerated. This broadens the momentum spectrum because it amplifies the existing distribution, namely that leading particles have more momentum than trailing particles. It broadens the time distribution simply because momentum is monotonically related to time-of-flight in the gun. Further broadening occurs because the time-varying nature of the cavity fields results in additional acceleration of those particles that are pushed ahead, and less acceleration of those that are pushed back.

In order to get a more detailed look at the phase-space distributions for the two cases, I have compressed the longitudinal phase-space using a ideal alpha-magnet and drift space system, as described in Chapter 3. The results are shown in Figure 2.38. Several effects are apparent in this Figure. First, the highest-momentum part of the beam is more energetic for the high space-charge case, due to acceleration by the fields of trailing particles that occurs in the gun; as a result, the "top" of the beam falls further behind the centroid during compression, because the delay in the alpha-magnet increases with increasing momentum. Second, the time-spread for a given small momentum slice is significantly broadened; this is a result of the

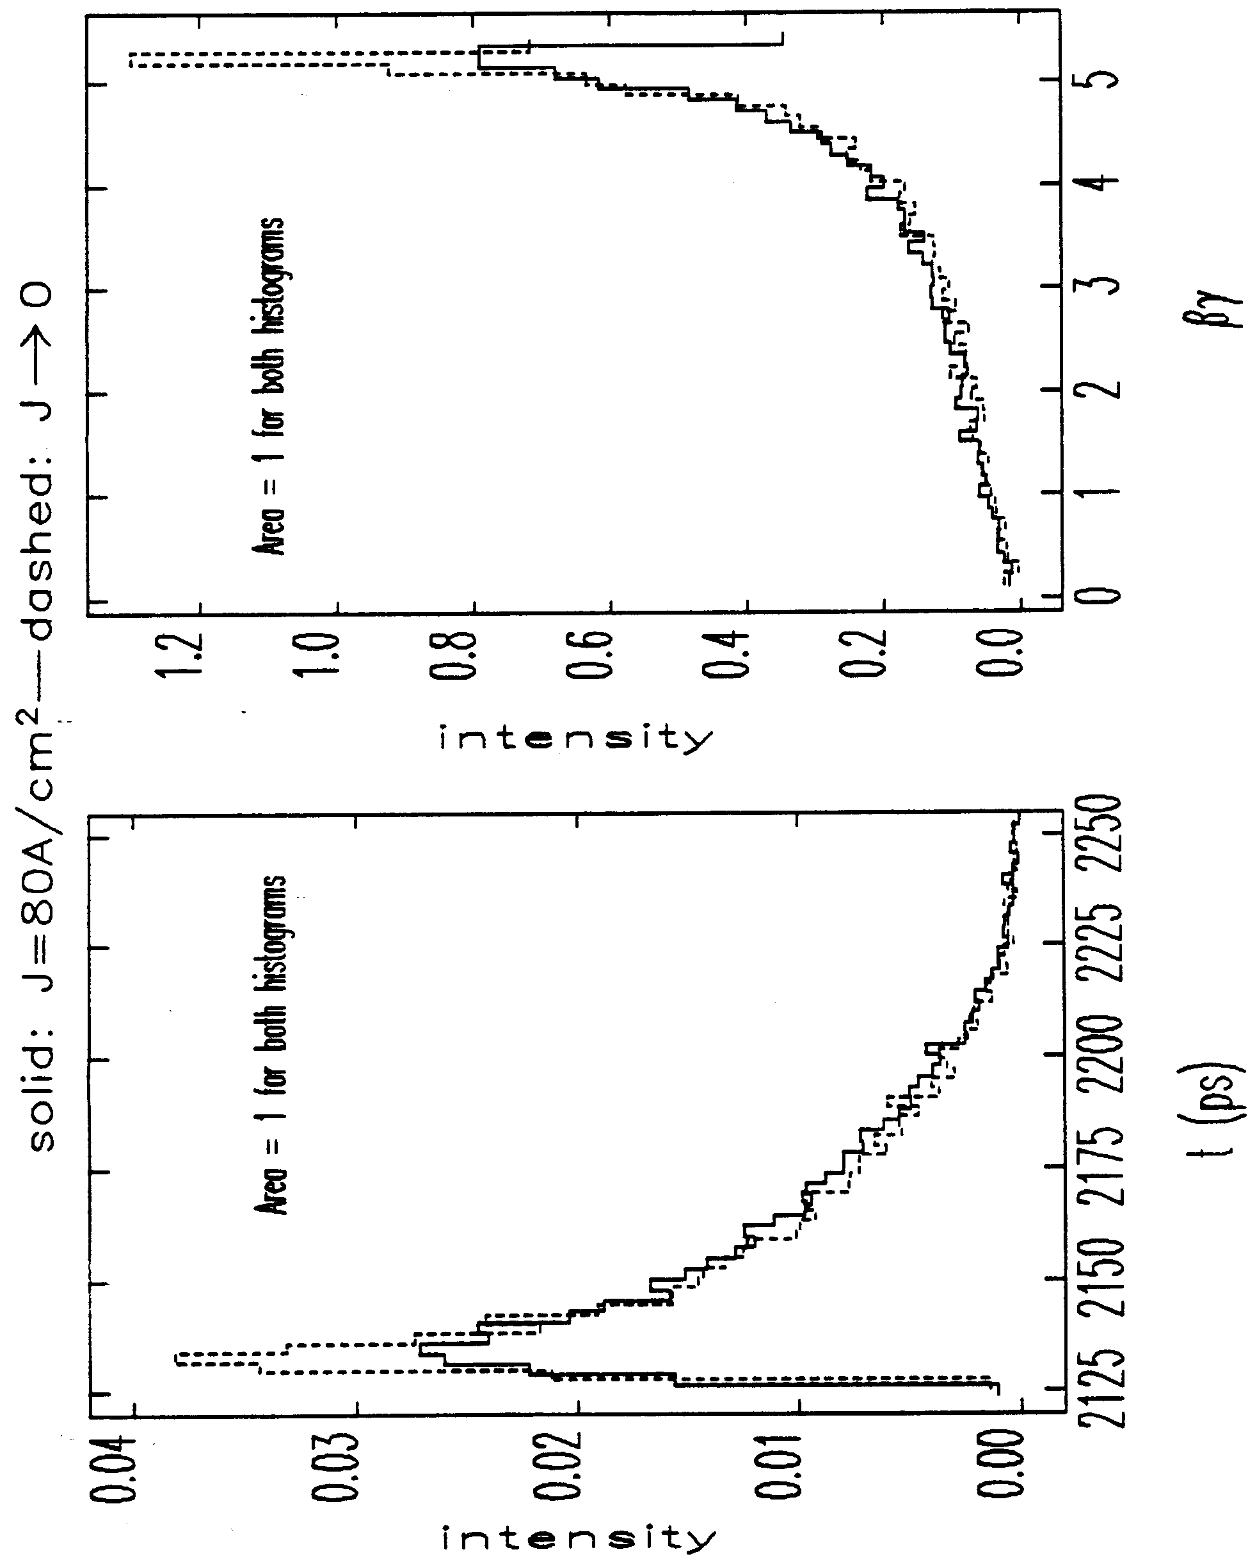

Figure 2.37: MASK Longitudinal Phase-Space Distributions at the End of the Second Cell

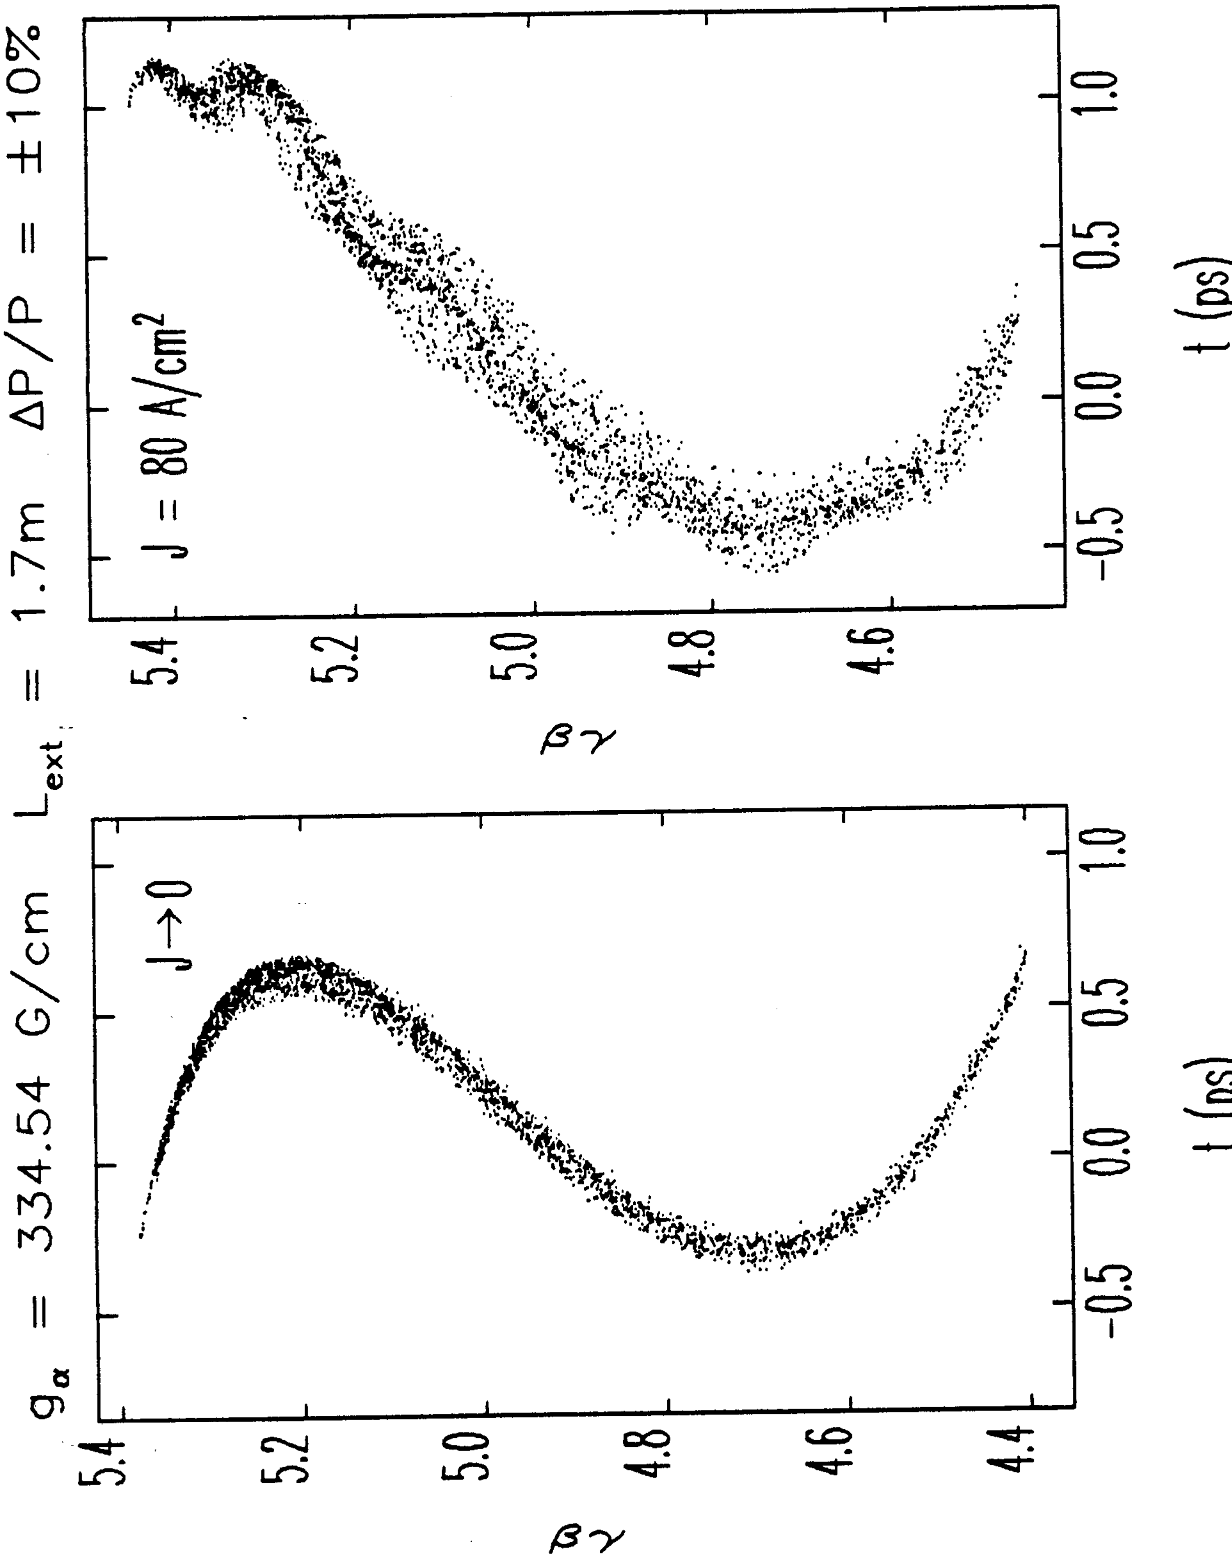

Figure 2.38: MASK Longitudinal Phase-Space Distributions after Alpha-Magnet-Based Compression

longitudinal space-charge forces being a non-linear function of radius, which produces longitudinal mixing. Third, there is a clear oscillation in the p(t) curve, especially evident near the top of the beam; this is apparently a plasma oscillation, resulting from the longitudinal space-charge forces. Finally, one sees that the space-charge forces broaden the beam significantly relative to the broadening by non-linear fields (without which the phase-space distribution for $J \rightarrow 0$ would be a line).

2.3.13 MASK Predictions of Gun Performance

Figures 2.39 through 2.41 show results for the emittance, brightness, and charge as calculated by MASK, as a function of current density and peak on-axis electric field, for various final momentum fractions. The smaller range of E_{p2} in these results (as compared to those for `rfgun`) is a result of my concentrating computer resources on the range that is of most interest for actual running at SSRL. One sees that the normalized charge per bunch decreases as current density increases, a result of the longitudinal forces in the beam, which may be thought of as decreasing the current density by forcing the electrons apart. As one might expect, this effect lessens as the cavity fields are increased, since this decreases the strength of the particle-induced fields relative to the cavity fields and results in faster acceleration, thus decreasing the effect of the particle-induced fields further.

The trends in emittance and brightness hold some surprises. In particular, the emittance does not always increase when the current density is increased: for small momentum intervals, the opposite occurs. There are two effects that may explain this. First, as was seen above, the particle-induced fields tend to counter the non-linear cavity fields, which would in turn tend to limit emittance growth due to those non-linear fields. Second, space-charge related changes in the longitudinal phase-space result in there being a larger phase-interval represented in a given final momentum fraction for small current density than for a large current density. As was seen above, emittance depends strongly on the initial phase-interval one considers. Hence, it should not be surprising that when one takes a very small final momentum fraction, this effect becomes apparent, since for small final momentum fractions the initial phase-interval is smaller, whereas the effects of longitudinal space charge are great

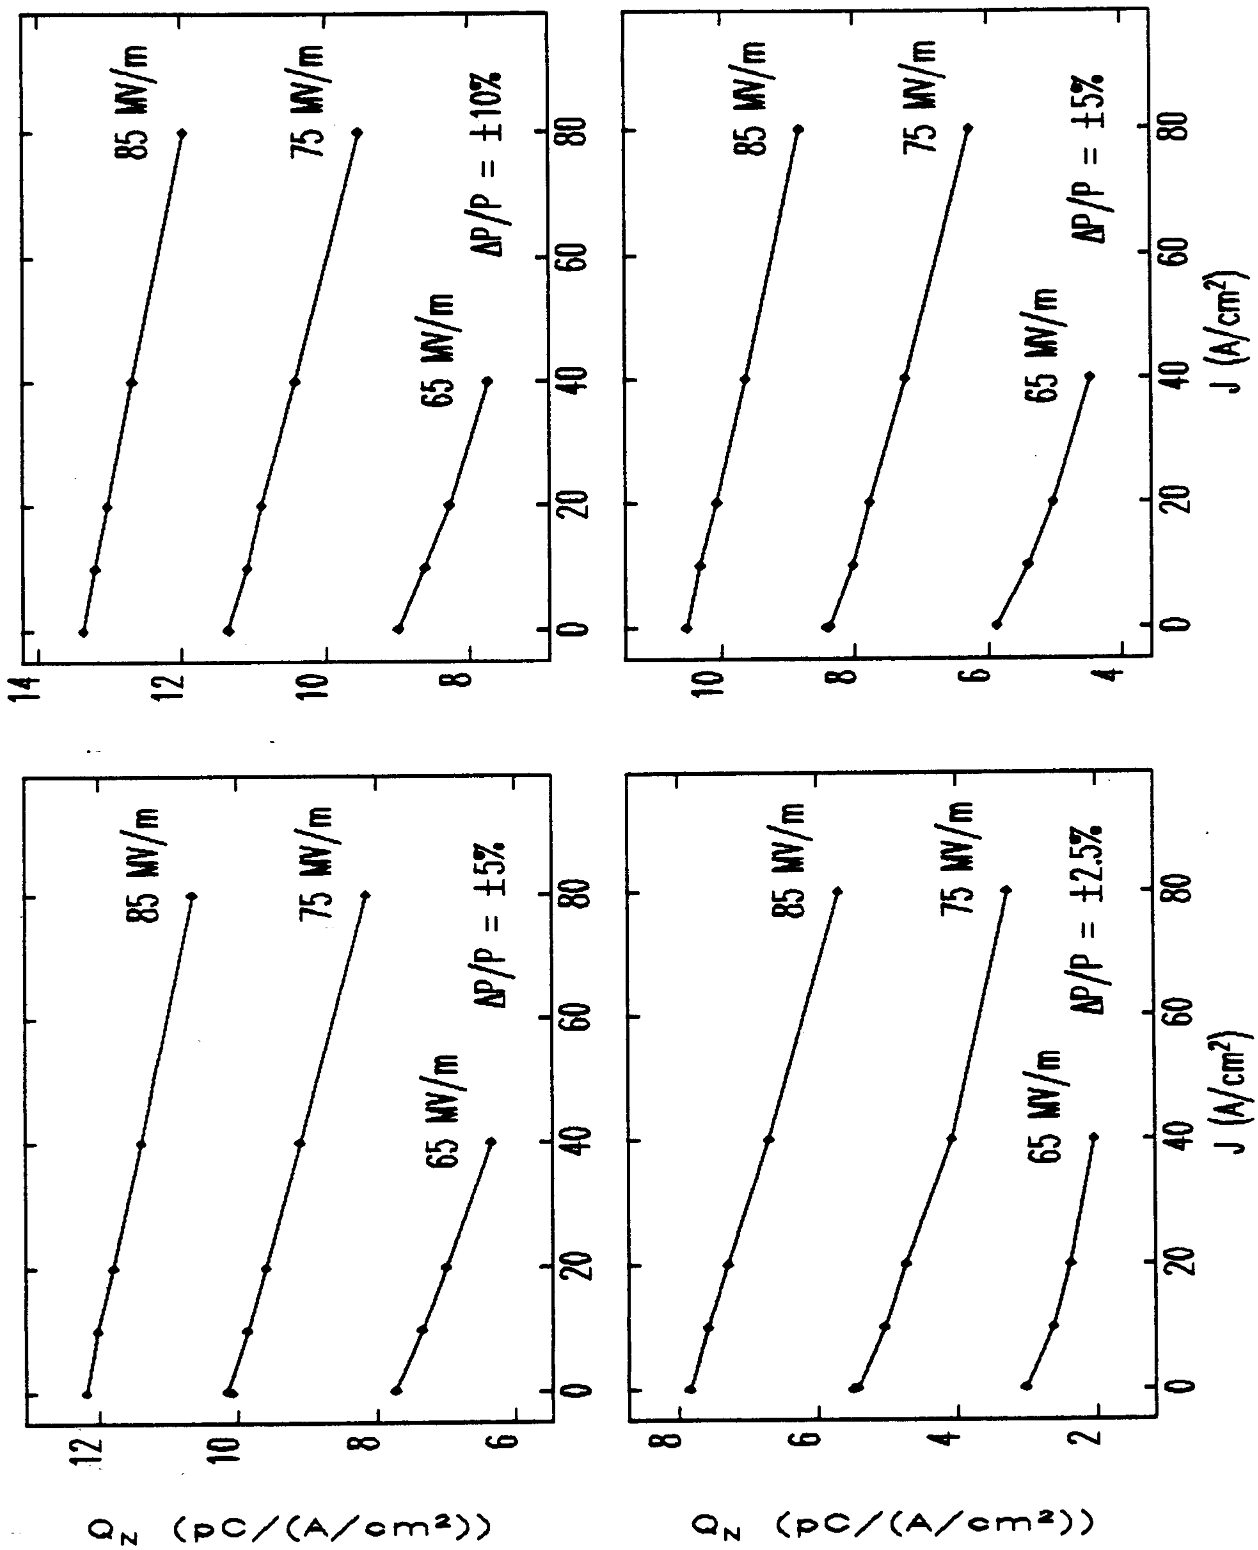

Figure 2.39: MASK Results for Normalized Charge Per Bunch

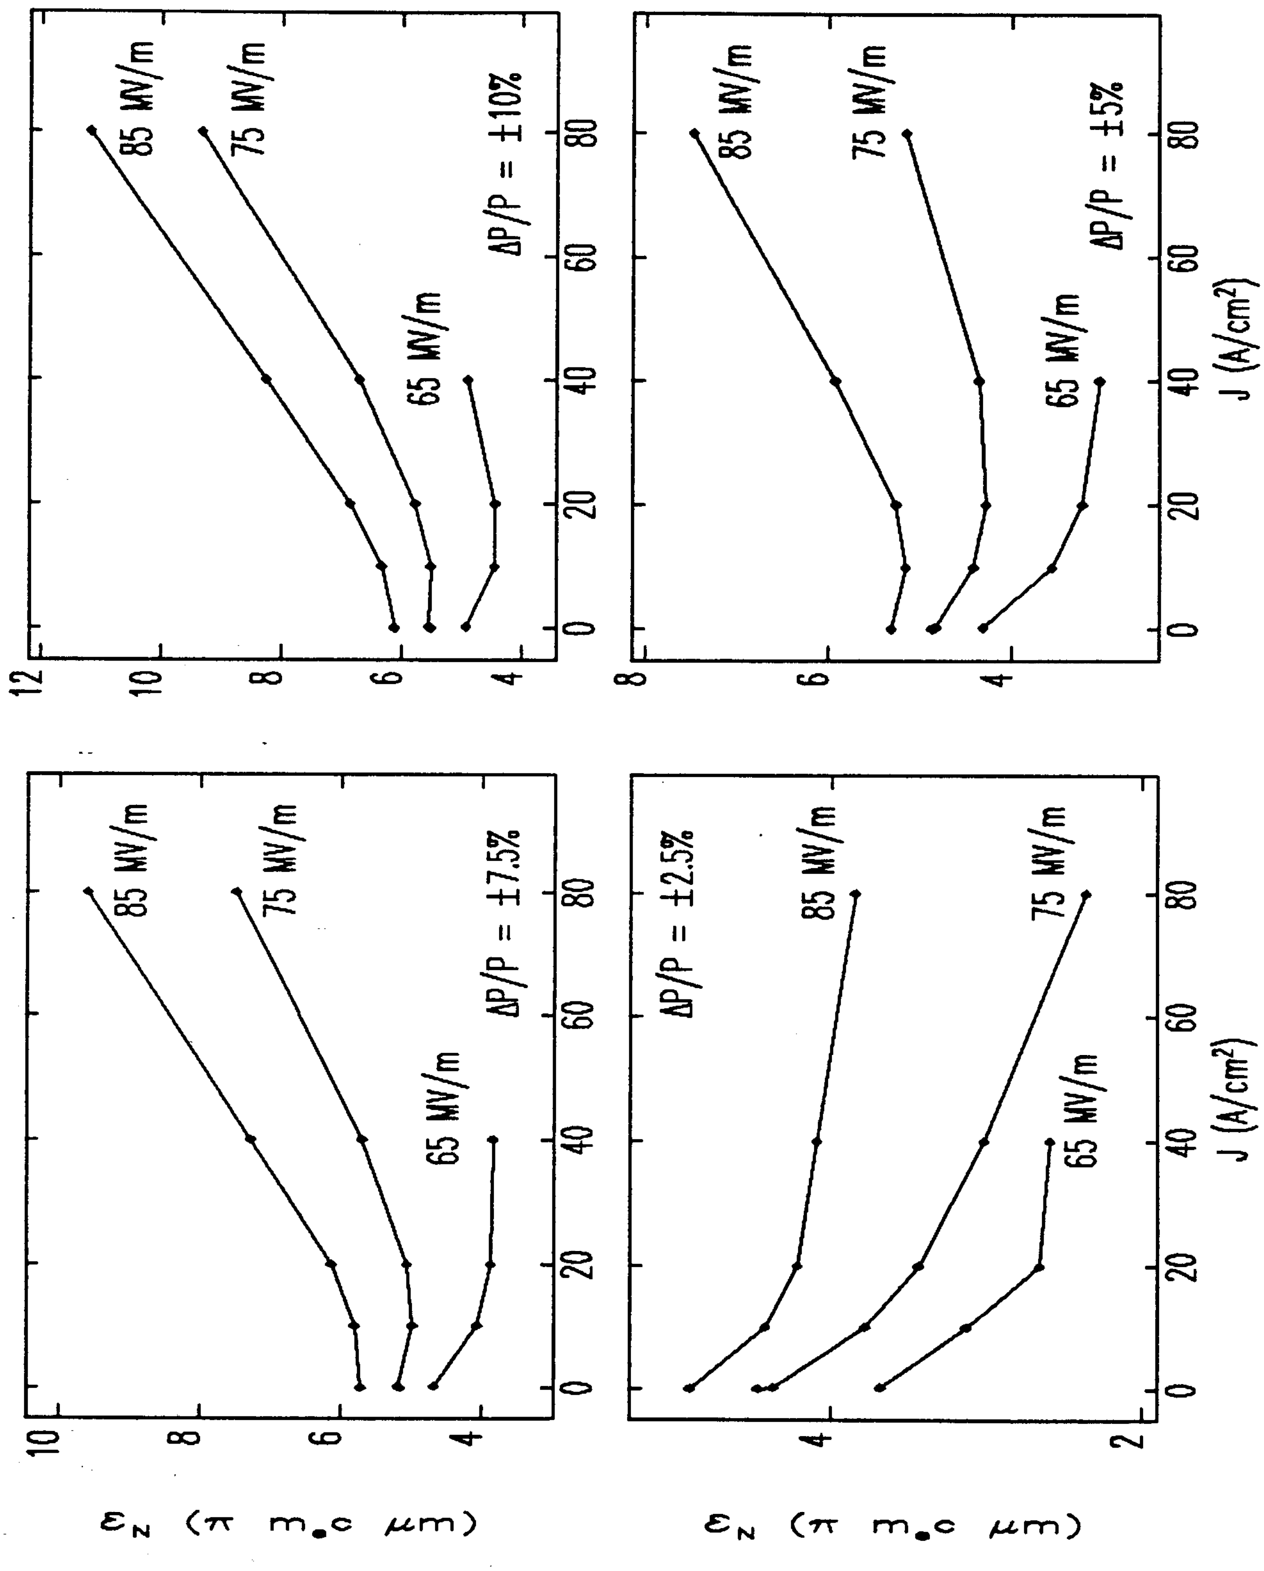

Figure 2.40: MASK Results for Normalized RMS Emittance

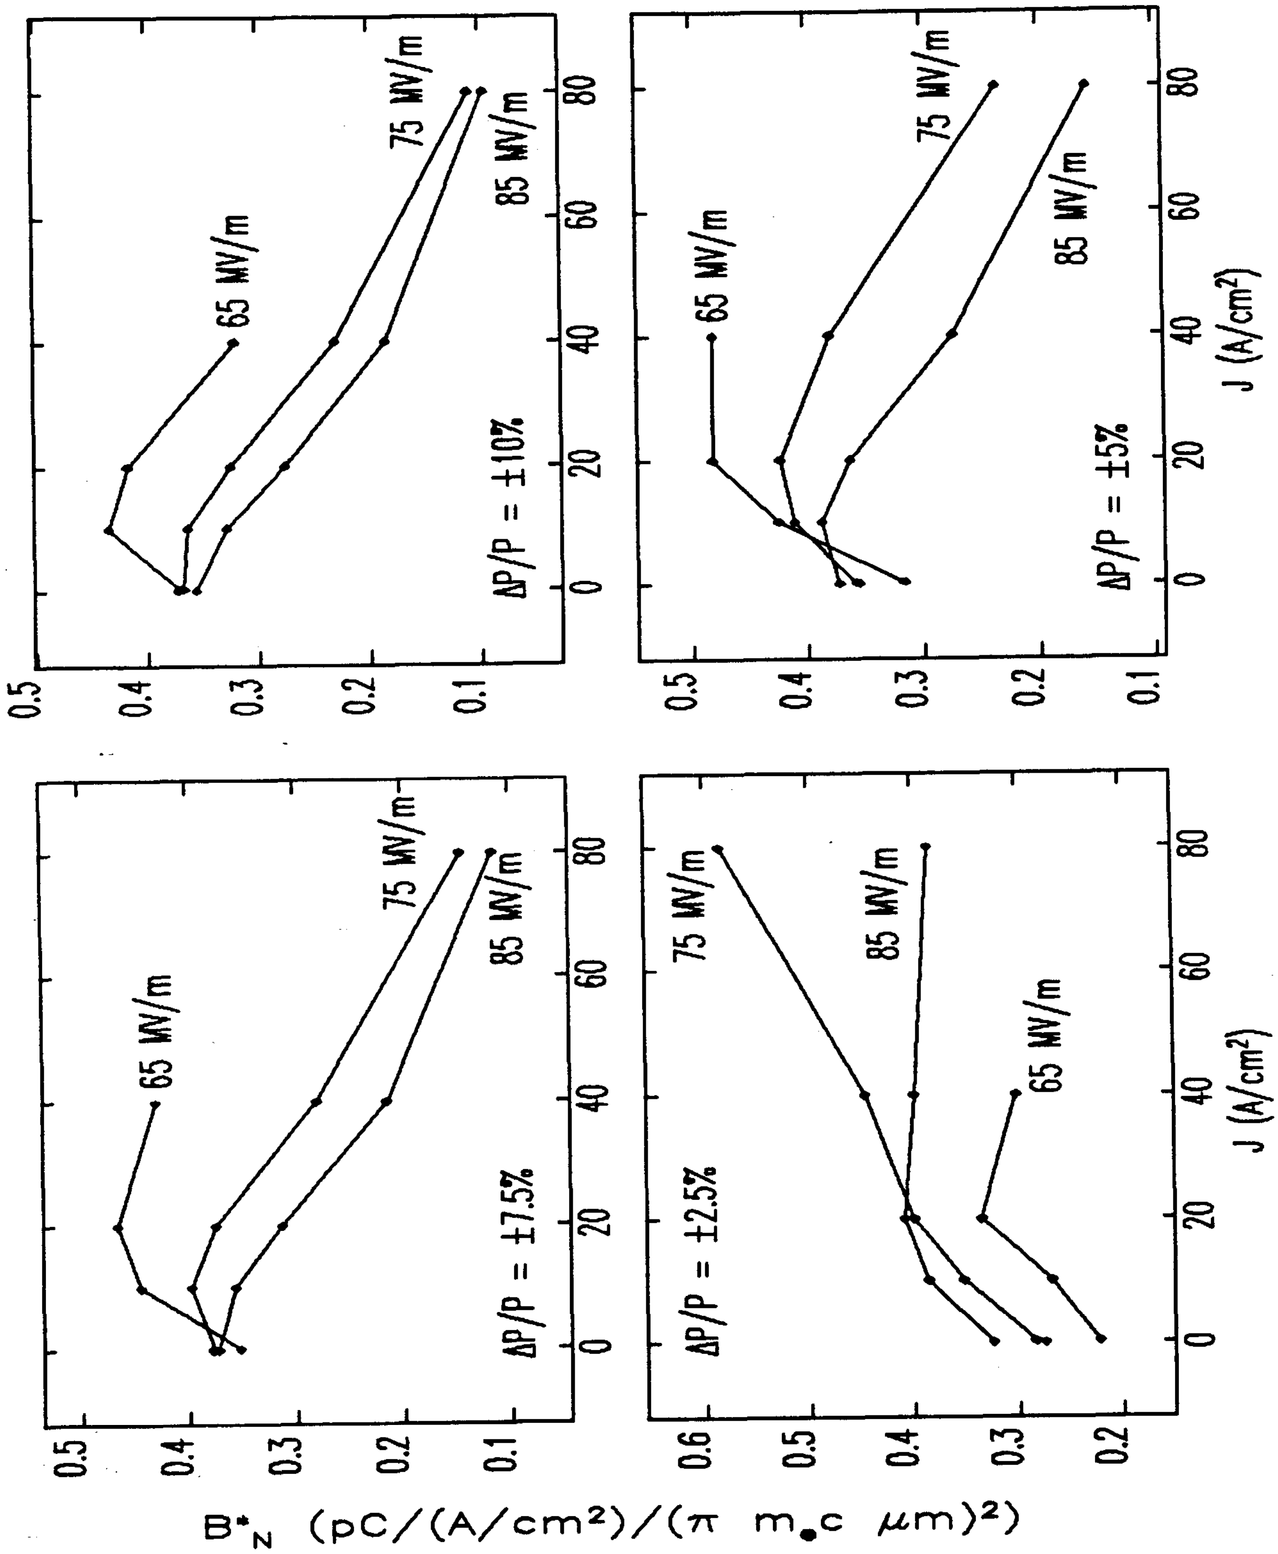

Figure 2.41: MASK Results for Normalized Brightness

(since the charge distribution is so strongly peaked).

Note that the brightness shown here, as throughout this section, is normalized to the current density. Hence, Figure 2.41 does not show the actual brightness decreasing with increasing current density. Figure 2.42 shows the transverse brightness, $B_t = B_n^* J$, without normalization to the current density. One sees that the transverse brightness increases rapidly with current density at first, then saturates as space-charge effects on the emittance overcome the increasing charge per bunch.

Finally, to provide a more complete picture, Figures 2.43 through 2.45 show transverse phase-space distributions for $E_{p2} = 75$MV/m and $\Delta P/P = 10\%$, for a range of current densities. As noted previously, the distributions for higher current density show less curvature due to the balancing of non-linear cavity fields by space charge. Printed on the graphs are the RMS beam-sizes and beam-divergences. One sees that the beam is predicted to be quite small at the gun exit, but that the RMS divergence is rather large.

Additional performance data will be presented in Chapter 4, where I include the effects of the gun-to-linac transport line, and in particular the effects of the alpha-magnet. In addition, Chapter 4 gives comparative data for other RF gun and DC gun systems.

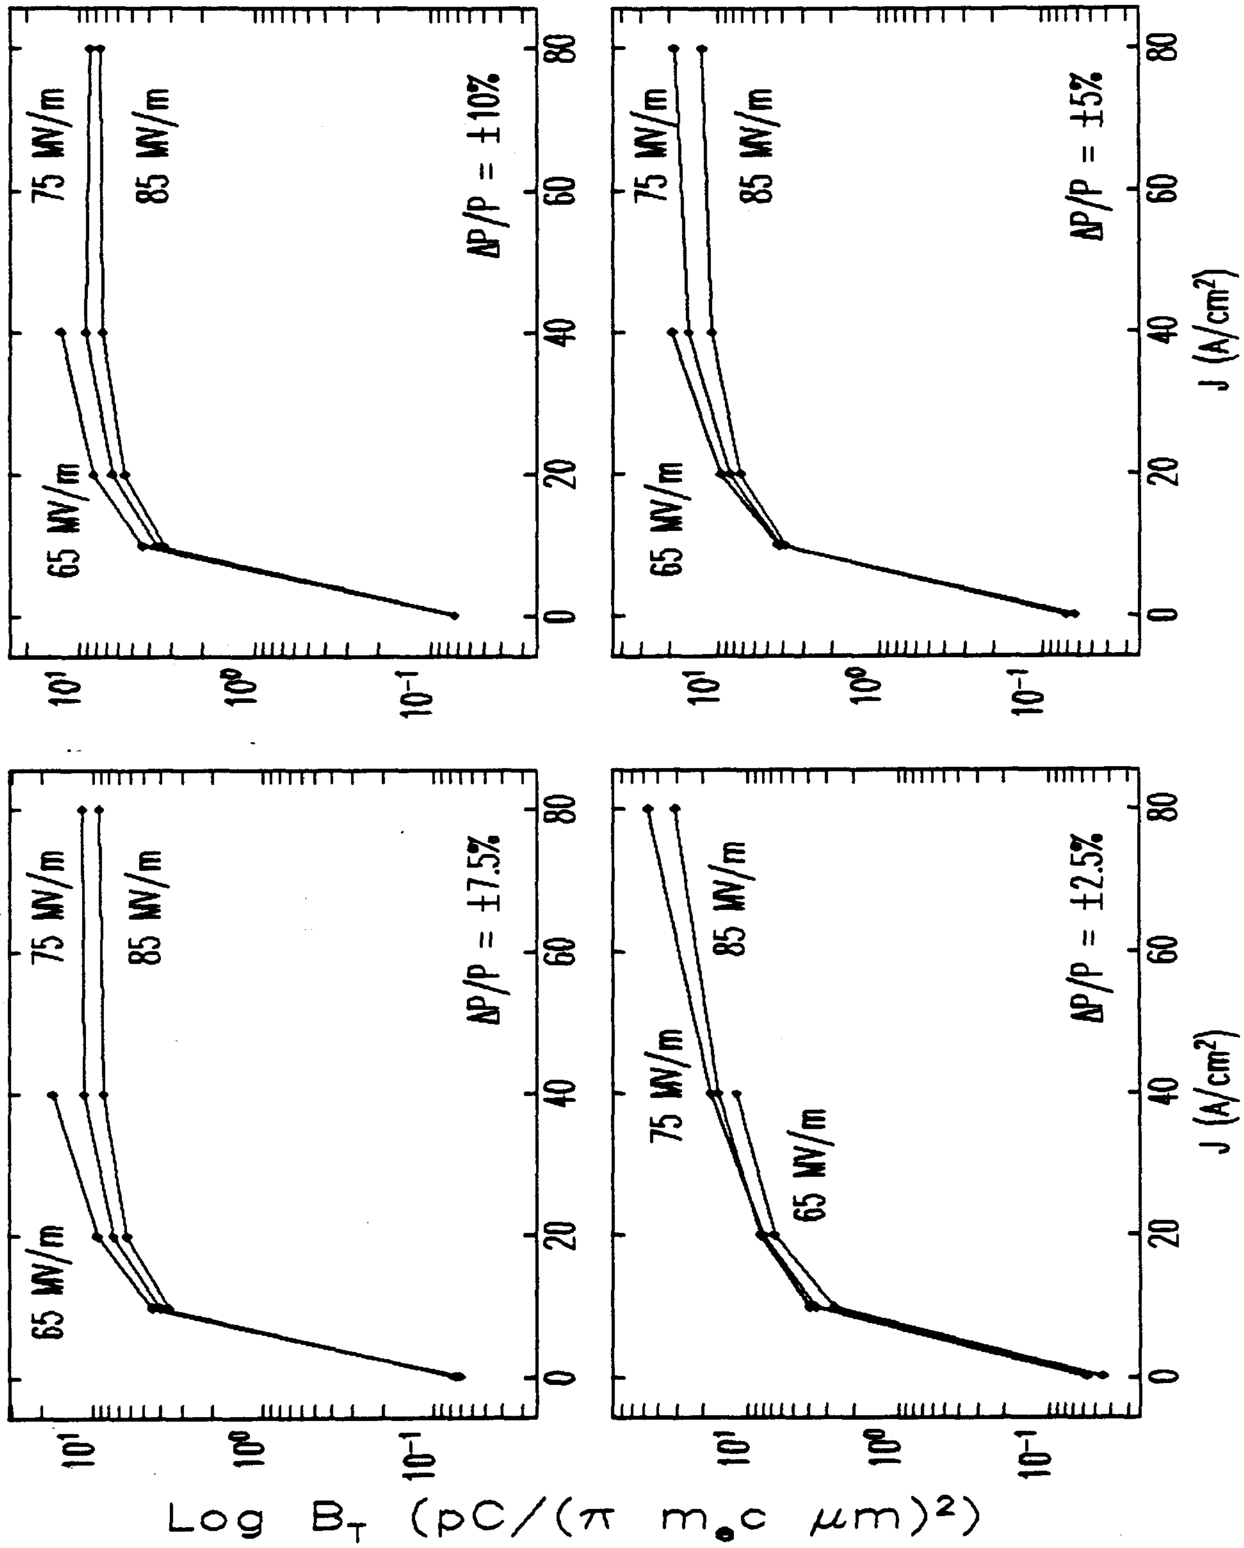

Figure 2.42: MASK Results for Transverse Brightness

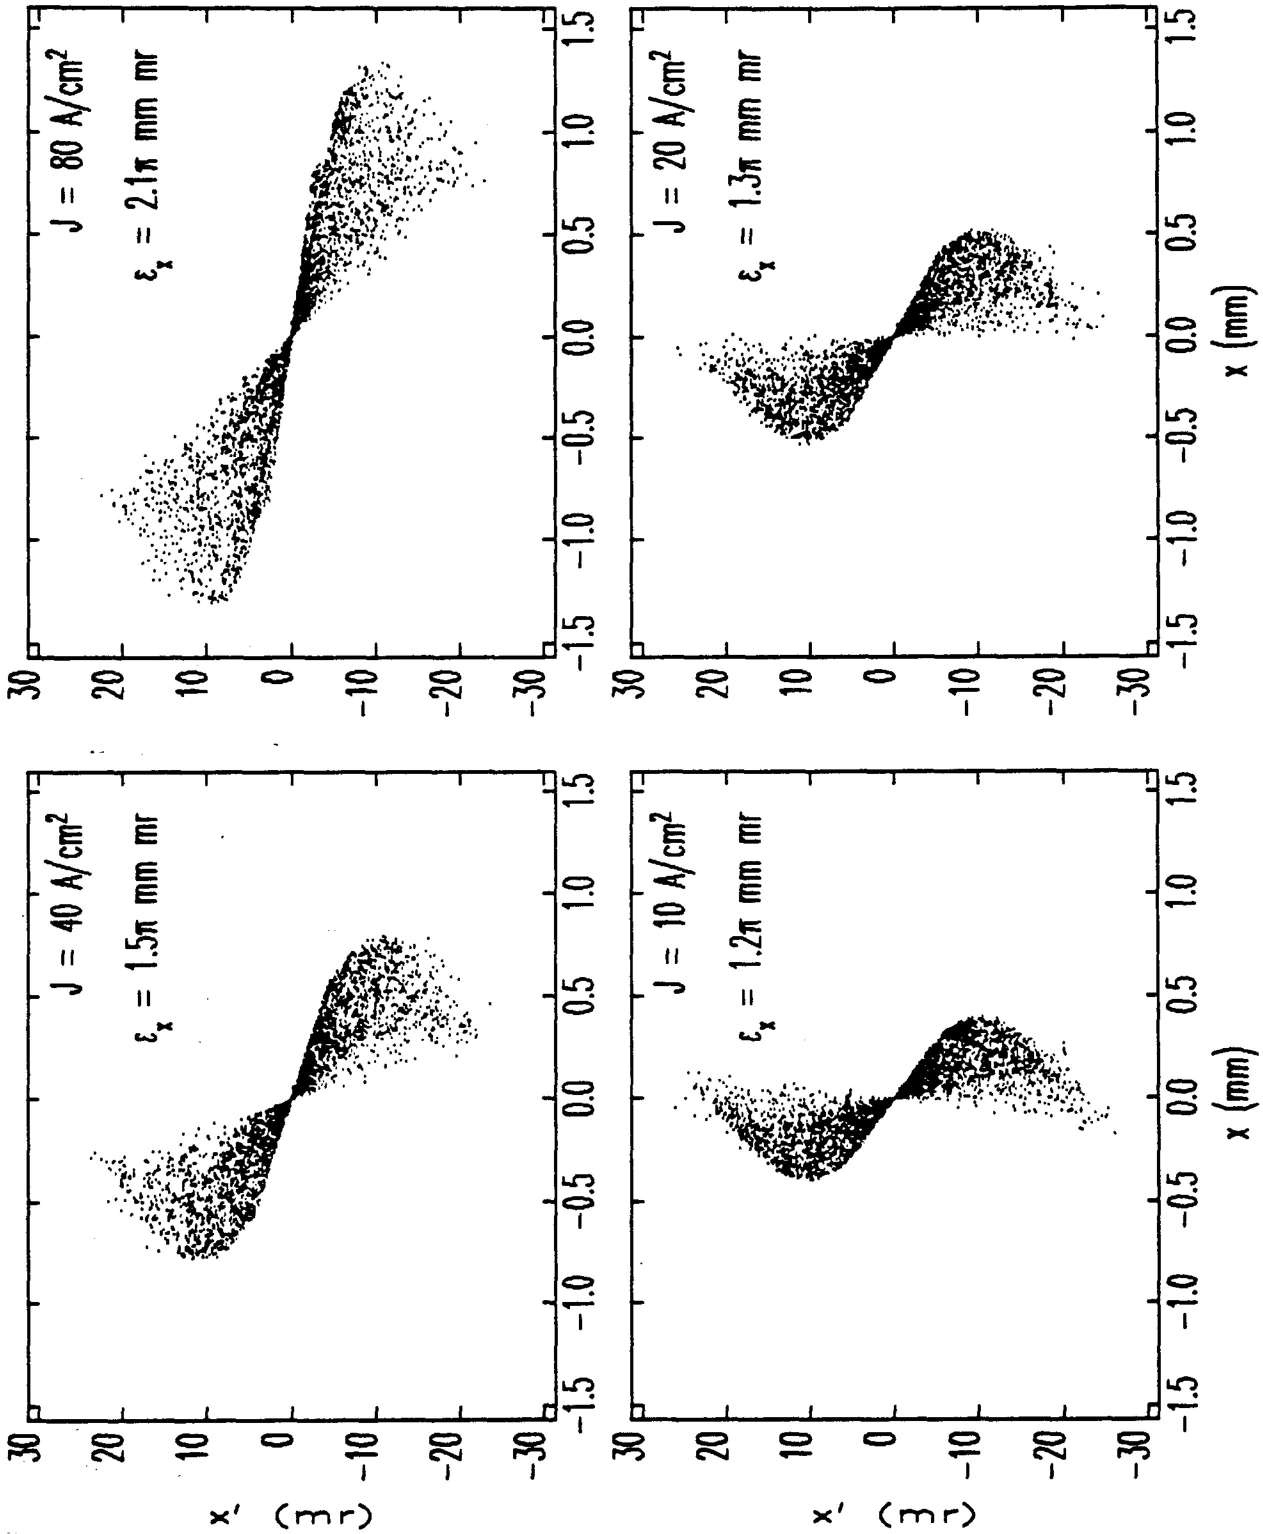

Figure 2.43: MASK Results for Transverse Phase-Space Distribution, for $E_{p2} = 75$ MV/m and $\Delta P/P = 10\%$—Part 1

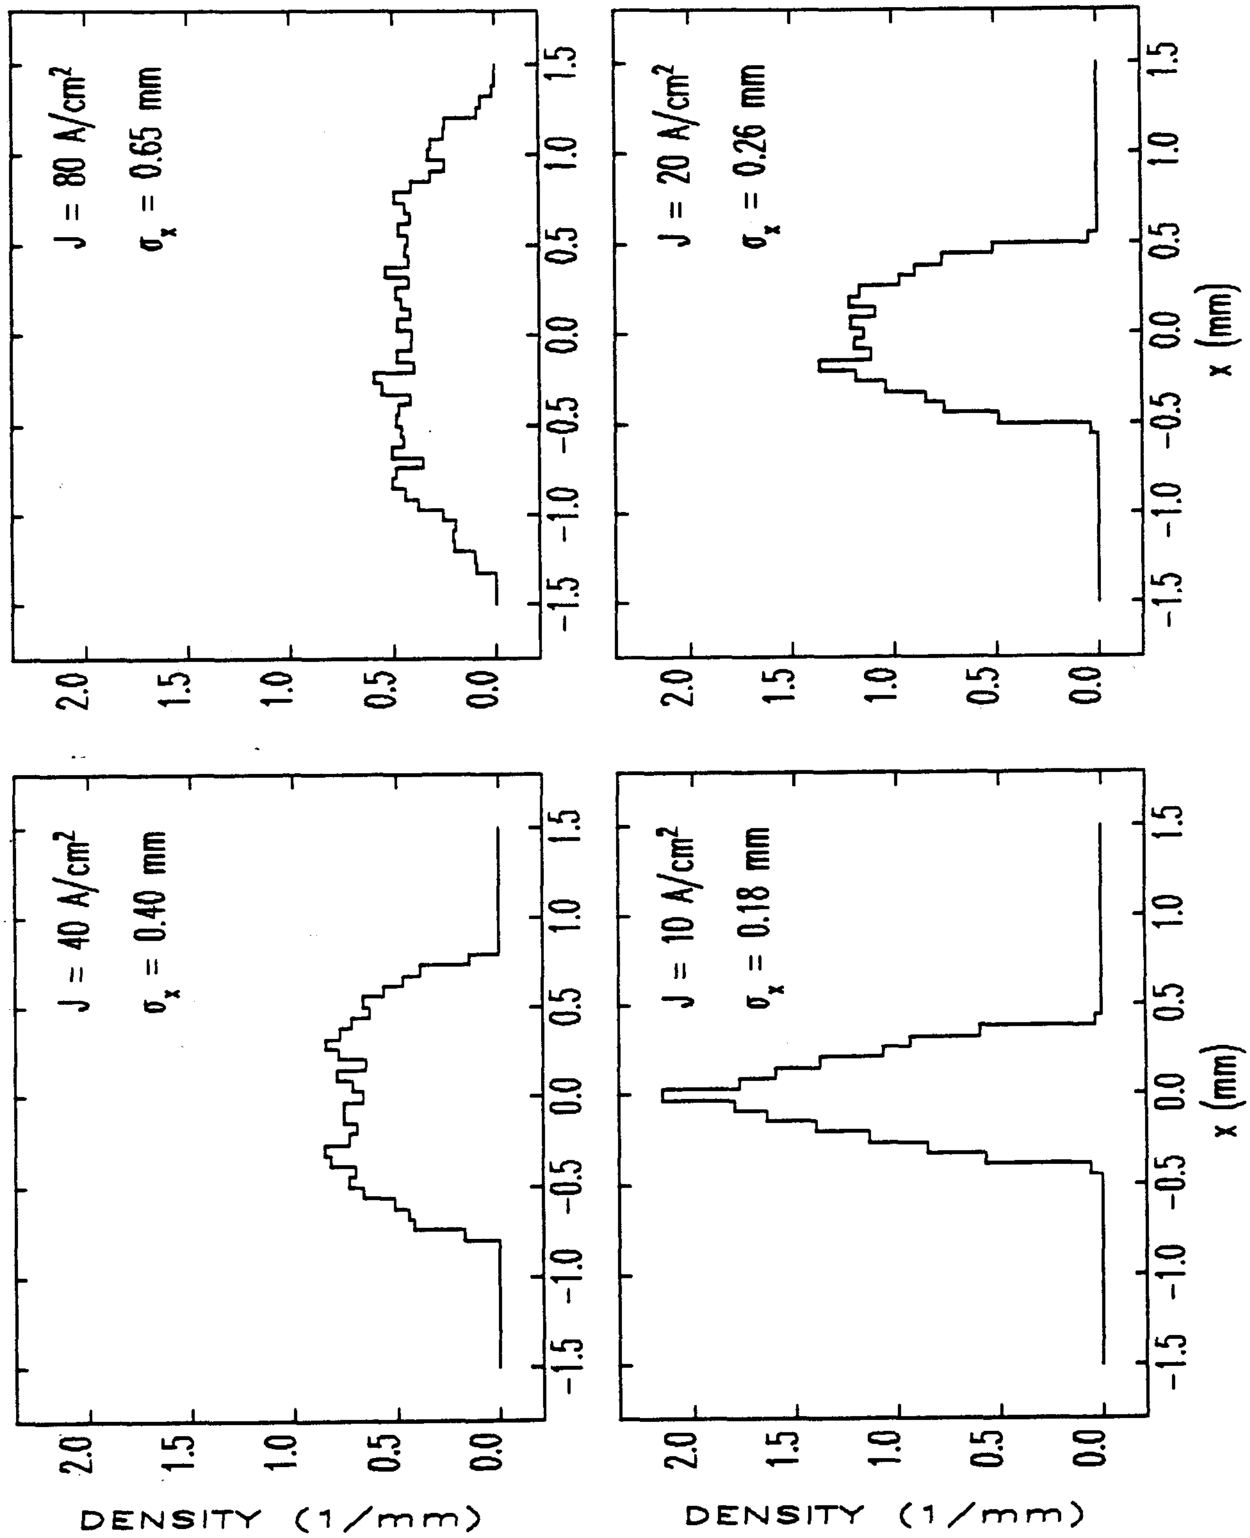

Figure 2.44: MASK Results for Transverse Phase-Space Distribution, for E_{p2} = 75 MV/m and $\Delta P/P$ = 10%—Part 2

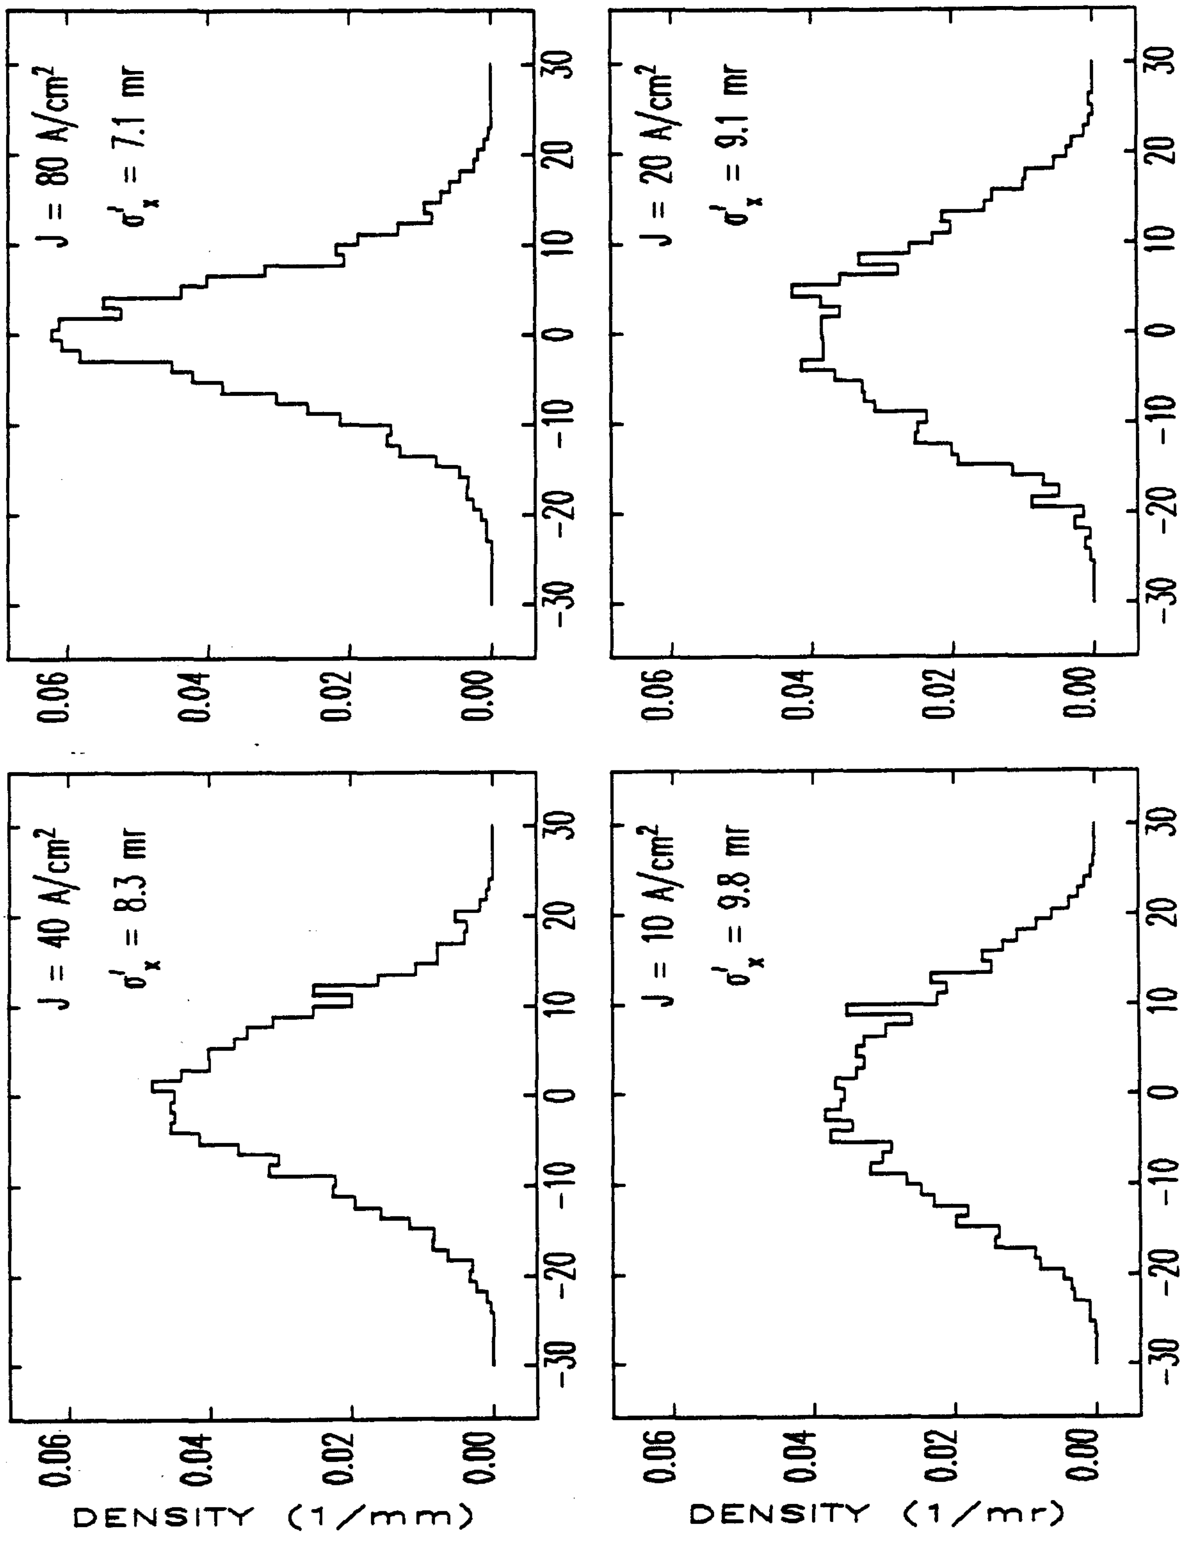

Figure 2.45: MASK Results for Transverse Phase-Space Distribution, for $E_{p2} = 75$ MV/m and $\Delta P/P = 10\%$—Part 3