

Simulation of a high-average power free-electron laser oscillator

H. P. Freund*

Science Applications International Corporation, McLean, Virginia 22102, USA

M. Shinn and S. V. Benson

Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606, USA

(Received 7 November 2006; published 8 March 2007)

In this paper, we compare the 10 kW-Upgrade experiment at the Thomas Jefferson National Accelerator Facility in Newport News, VA, with numerical simulations using the MEDUSA code. MEDUSA is a three-dimensional FEL simulation code that is capable of treating both amplifiers and oscillators in both the steady-state and time-dependent regimes. MEDUSA employs a Gaussian modal expansion, and treats oscillators by decomposing the modal representation at the exit of the wiggler into the vacuum Gaussian modes of the resonator and then analytically determining the propagation of these vacuum resonator modes through the resonator back to the entrance of the wiggler in synchronism with the next electron bunch. The bunch length in the experiment is of the order of 380–420 fsec FWHM. The experiment operates at a wavelength of about 1.6 microns and the wiggler is 30 periods in length; hence, the slippage time is about 160 fsec. Because of this, slippage is important, and must be included in the simulation. The observed single pass gain is 65%–75% and, given the experimental uncertainties, this is in good agreement with the simulation. Multipass simulations including the cavity detuning yield an output power of 12.4 kW, which is also in good agreement with the experiment.

DOI: 10.1103/PhysRevSTAB.10.030702

PACS numbers: 41.60.Cr

I. INTRODUCTION

Work is presently under way at the Thomas Jefferson National Accelerator Facility in Newport News, VA (henceforth referred to as Jefferson Laboratory) on what is referred to as the 10 kW-Upgrade experiment [1]. The facility is undergoing continual upgrades; in particular, a new permanent magnet wiggler has been installed that will permit operation at wavelengths as short as 0.9 micron. At the intended power levels, thermal distortions in cavity mirrors can alter mode quality and negatively impact performance; hence, it is important to predict the character and magnitude of the distortions and to be able to model their effect on FEL performance. To this end, we are engaged in a program to address these key issues by developing modeling and simulation tools that can treat mirror distortions. For this purpose, we will use the MEDUSA code [2,3], which is a three-dimensional FEL simulation code that is capable of treating both amplifiers and oscillators in both the steady-state and time-dependent regimes. MEDUSA employs a Gaussian modal expansion, and treats oscillators by decomposing the modal representation at the exit of the wiggler into the vacuum Gaussian modes of the resonator and then analytically determining the propagation of these vacuum resonator modes through the resonator back to the entrance of the wiggler in synchronism with the next electron bunch. Knowledge of the power loading on the mirrors allows us to model the mode

distortions using Zernike polynomials [4], and this technique will be incorporated into MEDUSA.

In this paper, we report on the progress to date in this activity. The first step is to compare MEDUSA predictions with the observed performance of the experiment at low duty factor where mirror distortions are unimportant. We then go on to determine the effects of the so-called first order properties, which include changes in the Rayleigh range and shifts in the position of the mode waist. Higher-order distortions such as coma, astigmatism, and spherical aberration, collectively known as higher-order aberrations, will be incorporated in the future.

The organization of the paper is as follows. In Sec. II we discuss the formulation used in MEDUSA. A description of the experiment is given in Sec. III, and of the numerical results in Sec. IV. The present study is limited to a comparison of the single pass gain as observed in the experiment and found in simulation. A more complete examination of the time evolution of the oscillator from pass to pass will be presented in a future paper. A summary and discussion is given in Sec. V.

II. THE NUMERICAL FORMULATION

The MEDUSA simulation code [2,3] employs a three-dimensional formulation that includes the slippage of the radiation relative to the electron beam. MEDUSA can model both helical and planar wiggler geometry and treats the electromagnetic field as a superposition of either Gauss-Hermite or Gauss-Laguerre modes in the slowly varying amplitude approximation. For the purposes of this study, we employ the Gauss-Hermite modes where

*Electronic address: henry.p.freund@saic.com

$$\delta \mathbf{A}(\mathbf{x}, t) = \hat{\mathbf{e}}_x \sum_{l,n,h} e_{l,n,h}(x, y) \times [\delta A_{l,n,h}^{(1)} \cos \varphi_h + \delta A_{l,n,h}^{(2)} \sin \varphi_h], \quad (1)$$

where l and n are transverse mode numbers, h is the harmonic number, $e_{l,n,h} = \exp(-r^2/w_h^2) H_l(\sqrt{2}x/w_h) \times H_n(\sqrt{2}y/w_h)$, H_l is the Hermite polynomial of order l , and w_h is the spot size, $\varphi_h = h(k_0 z - \omega_0 t) + \alpha_h r^2/w_h^2$ ($k_0 = \omega_0/c$). We employ a slowly varying amplitude approximation (SVEA) in which $\delta A_{l,n,h}^{(1,2)}$, w_h , and α_h vary slowly in z and t . The dynamical equations are

$$\left(\frac{d}{dz} + \frac{w_h'}{w_h} \right) \begin{pmatrix} \delta a_{l,n,h}^{(1)} \\ \delta a_{l,n,h}^{(2)} \end{pmatrix} + K_{l,n,h} \begin{pmatrix} \delta a_{l,n,h}^{(2)} \\ -\delta a_{l,n,h}^{(1)} \end{pmatrix} = \begin{pmatrix} s_{l,n,h}^{(1)} \\ s_{l,n,h}^{(2)} \end{pmatrix}, \quad (2)$$

where $\delta a_{l,n,h}^{(1,2)} = e \delta A_{l,n,h}^{(1,2)} / m_e c^2$, $d/dz = \partial/\partial z + c^{-1} \partial/\partial t$, the “prime” superscript denotes the total z -derivative,

$$K_{l,n,h} = (l + n + 1) \left(\alpha_h \frac{w_h'}{w_h} - \frac{\alpha_h'}{2} - \frac{1 + \alpha_h^2}{h k_0 w_h^2} \right), \quad (3)$$

$$\begin{pmatrix} s_{l,n,h}^{(1)} \\ s_{l,n,h}^{(2)} \end{pmatrix} = \frac{2\omega_b^2}{h\omega_0 c} \frac{F_{l,n}}{w_h^2} \left\langle \frac{\mathbf{v}_x}{|\mathbf{v}_z|} e_{l,n,h} \begin{pmatrix} \cos \varphi_h \\ -\sin \varphi_h \end{pmatrix} \right\rangle, \quad (4)$$

where $\omega_b(z, t)^2 = 4\pi e^2 n_b(z, t) / m_e$ for a beam density n_b , and $F_{l,n} = [2^{l+n} l! n!]^{-1}$, and $\langle \dots \rangle$ describes an average over the complete 6D phase space as follows:

$$\langle \dots \rangle = \frac{1}{2\pi} \int_0^{2\pi} d\psi_0 \int_0^\infty d\gamma_0 \frac{\exp[-(\gamma_0 - \bar{\gamma}_0)^2/2\Delta\gamma^2]}{\sqrt{\pi/2[1 + \text{erf}(\bar{\gamma}_0/\sqrt{2}\Delta\gamma)]}} \iiint dx_0 dy_0 dp_{x0} dp_{y0} \frac{\exp(-r_0^2/2\sigma_r^2 - p_{\perp 0}^2/2\sigma_p^2)}{(2\pi)^2 \sigma_r^2 \sigma_p^2} \dots, \quad (5)$$

where the average is over the initial beam parameters (denoted by the subscript 0) in ponderomotive phase (ψ_0), energy (γ_0), and phase space (x_0, y_0, p_{x0}, p_{y0}), $\bar{\gamma}_0$ and $\Delta\gamma$ denote the average beam energy and the energy spread, and σ_r and σ_p describe the initial phase space parameters. The spot size and radius of curvature for each harmonic component are given by

$$w'_h = \frac{2\alpha_h}{h k_0 w_h} - w_h Y_h, \quad (6)$$

$$\frac{\alpha'_h}{2} = \frac{1 + \alpha_h^2}{h k_0 w_h^2} - (X_h + \alpha_h Y_h), \quad (7)$$

These equations constitute the source-dependent expansion [5], which is a self-consistent adaptive eigenmode representation that tracks the optical guiding of the mode based upon the interaction with the electron beam. The field equations are integrated simultaneously with the complete three-dimensional Lorentz force equations for an ensemble of electrons. No wiggler-average orbit approximation is used so that the spatial step size must be small enough to resolve the wiggler motion.

Time dependence is typically treated numerically in two ways in FEL simulation codes. One method employs a SVEA in z only combined with a field representation that is an ensemble of discrete harmonics; in effect an explicit Fourier decomposition. The other method uses an SVEA in both z and t [6–9]. These two techniques have been shown to be formally equivalent [10]. As described above, MEDUSA can simulate the time dependence using either technique although, in most cases, the second technique is

computationally preferred. In practice, therefore, the second technique is used in MEDUSA to simulate time dependence and the harmonic decomposition is used to include higher harmonics in the time-dependent simulation. However, we leave a discussion of harmonic generation in the Jefferson Laboratory experiment for a future work.

III. THE JEFFERSON LABORATORY EXPERIMENT

The Jefferson Lab IR-Upgrade FEL operates an energy recovery accelerator with a high power FEL wiggler and resonator. The nominal experimental parameters are shown in Table I. The electron beam consists of a core distribution and a halo distribution. The charge, emittance, and peak current shown in the table are of that core beam. The wiggler is very well characterized and is essentially ideal. The wiggler is gap tunable with the field variable from 1.3 to 8.25 kG so that the wiggler K varies between 0.5 and 3.0. The optical resonator is nearly concentric, and consists of a high reflector at the upstream end and a transmissive element at the downstream end that out-couples approximately 10% of the power. The Rayleigh range can be varied in the experiment by changing the radius of curvature of the high reflector. However, the cavity is slightly astigmatic and can lead to a difference between the Rayleigh range in the two axes. The measured gain at 1.6 μm is $(70 \pm 5)\%$.

For this study, the experiment was run at a low duty factor where the power loading on the mirrors can be kept small in order to minimize distortion. Thus, the accelerator was run in pulsed mode for 1 msec at a repetition rate of

TABLE I. Nominal experimental parameters for operation at 1.6 μm .

Electron beam	
Energy	115 MeV
Peak current	310–370 A
Average current	4.4 mA
Normalized emittance	9 mm-mrad/7 mm-mrad
Energy spread	0.35%
Bunch length	380–420 fsec
Bunch charge	79 pC
Initial beam dimensions	257 microns/212 microns
Twiss- α parameter	1.25
Repetition rate	37.425 MHz
Wiggler	
Amplitude	3.746 kG
Period	5.5 cm
Length	30 periods/1.65 m
Optical mode	
Wavelength	1.57 microns
Rayleigh range	1.5 (± 0.30) m
Location of mode waist	1.03 (± 0.05) m
Out-coupling fraction	10%
Cavity length	≈ 32 m
Optical performance	
Single pass gain	$70\% \pm 5\%$
Average output power	≈ 13 kW (at 3.1 kW/mA)

60 Hz. The efficiency for this mode of operation was $(1.6 \pm 0.1)\%$ as calculated on the basis of the full beam charge. If only the core current is used, then the efficiency would increase to about 2.7%. This gives a scaling of the output power with an average current of 1.8 kW/mA on the basis of the full bunch charge, or 3.1 kW/mA using the core alone (79 pC). This power scaling with average current is observed to be independent of the average current when mirror distortion is not present.

Observe that the normalized emittances and beam dimensions shown in the table refer to the wiggle plane and the plane transverse to the wiggler plane, respectively. In addition, the Twiss- α parameter shown corresponds to a beam that is focused to a waist near the center of the wiggler, and this may be varied in the experiment by adjusting the magnetic transport line. The estimate of the Rayleigh range and the location of the mode waist contain some uncertainty, and the values given are the best estimate at the present time. The uncertainty in the Rayleigh range is about ± 30 cm. We also note that the mode waist is located about 15 cm downstream from the wiggler center with an uncertainty of about ± 5 cm. This will be shown to have important consequences for the single pass gain.

The actual beam in the experiment exhibits a complicated distribution including both a core and a halo. The phase space distribution is not Gaussian, and the temporal shape of the pulse is complex and is not well matched by a Gaussian temporal profile. The total beam pulse contains

about 135 pC, while the core contains about 79 pC. As a result there is an uncertainty in the microscopic experimental beam parameters that is not presently quantifiable, and this will affect the comparison with the simulation. In order to perform the simulations, however, we must make certain assumptions about the beam properties. To this end, we employ a Gaussian phase space distribution model in the simulation corresponding to the core alone. In addition, the temporal shape of the electron pulse is assumed in simulation to be parabolic, and the bunch length given in Table I refers to the full width of the parabola. Note that the bunch charge is given by $Q = (2/3)I_{\text{peak}}\tau_{\text{bunch}}$ for a parabolic pulse shape, and that this corresponds to the 79 pC charge for the core beam. The validity of these assumptions will rest on the details of the comparison between the experiment and the simulation results.

IV. NUMERICAL RESULTS

Our first goal is to determine whether MEDUSA is in substantial agreement with the experiment when distortion is absent or minimal. This is the case when the experiment operates at low average powers (i.e., low duty factors), where a single pass gain of the order of 65%–75% is measured for a peak current of 310 A and a bunch length of 380 fsec.

Observe that the cavity detuning is not relevant to a discussion of the single pass gain since we are assuming optimal synchronism between the radiation and the electron beam for these simulation runs. A discussion of the cavity detuning within the context of multipass simulations will follow the discussion of the single pass gain.

A. Single pass gain

Before proceeding to a full time-dependent simulation of the experiment, we begin with a steady-state simulation to optimize the gain with respect to the focusing of the electron beam in the wiggler. We start by studying the gain spectrum for a matched beam in a two-plane focusing wiggler. Note that in this case the Twiss parameters shown in Table I are not used. The gain spectrum for this case is shown in Fig. 1, and we observe that the single pass gain is about 190%. However, it is well known that the optimum performance is found when the electron beam is matched to the optical mode rather than the wiggler. To study this, we use a flat-pole-face wiggler model corresponding to the Jefferson Laboratory wiggler and adjust the Twiss parameters to yield an electron beam that is focused to a waist near the wiggler center. If we now assume that the optical mode waist coincides with the wiggler center, and optimize the electron beam focusing, then we find the maximum single pass gain corresponds to the Twiss parameters shown in Table I, and the gain spectrum for this case is shown in Fig. 2. Observe that the single pass gain for an optimally focused beam in the steady-state regime is about 420% for the nominal parameters. The evolution of the power, opti-

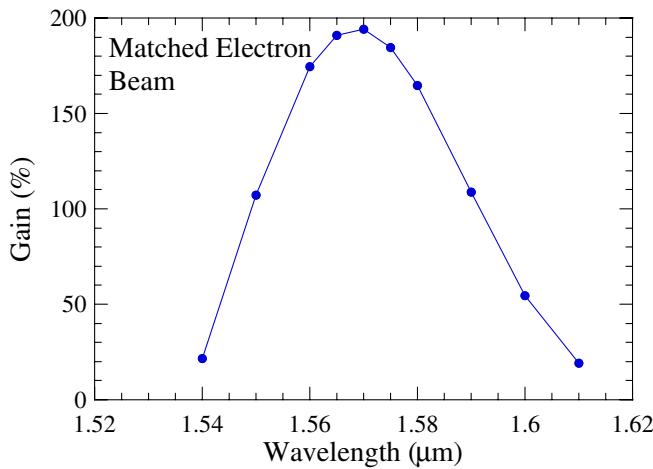


FIG. 1. (Color) Gain spectrum for a matched beam in a two-plane focusing wiggler.

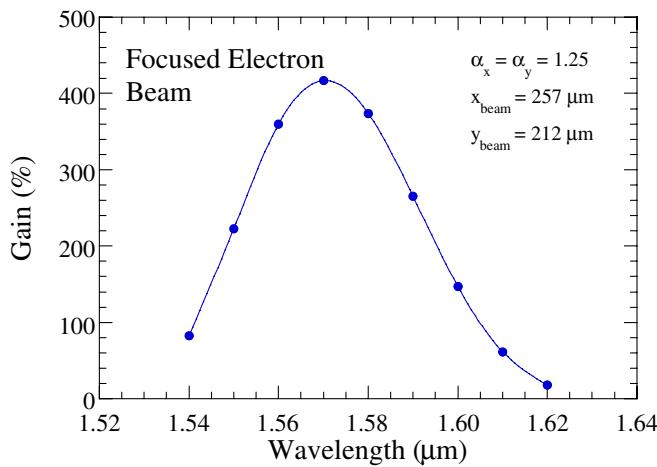


FIG. 2. (Color) Gain spectrum for an optimally focused electron beam.

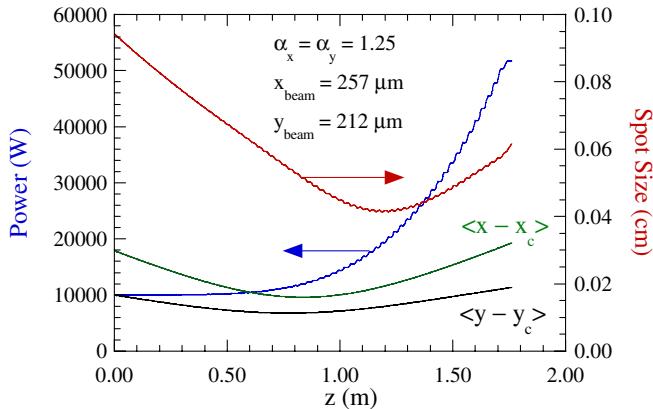


FIG. 3. (Color) Evolution of the power, optical spot size, and electron dimensions during a single pass through the wiggler.

cal spot size and x - and y -dimensions of the electron beam for this optimal case are shown in Fig. 3. Here we note that (1) substantial optical guiding occurs in this high-gain regime so that the evolution of the mode waist is not symmetric about the wiggler center, and (2) the optimal electron beam waist is located slightly upstream from the wiggler center.

The parameters shown in Table I are the nominal experimental parameters, and we note again that the optical waist is located about 15 cm downstream from the wiggler center. The experiment was optimized by focusing the electron beam to a waist near the center of the wiggler, which provides an optimal match to the resonator mode. This was also found in simulation, and is obtained for a Twiss- α parameter of 1.25 in simulation. The single pass gain found using MEDUSA operating in steady-state mode for these parameters is of the order of 400%. However, the slippage time through the wiggler for this experiment is of the order of 160 fsec, which is a substantial fraction of the bunch length. Hence, slippage is important and can be expected to substantially reduce the single pass gain with respect to steady-state predictions.

The slippage of the electromagnetic pulse through the wiggler is illustrated in Fig. 4 where we plot the pulse shapes at the entrance and exit from the wiggler in the small-signal regime. It is evident from the figure that, while the pulse is assumed symmetric at the entrance to the wiggler, it has slipped by at least half the total pulse length over the course of the wiggler. Amplification of the peak power over the pulse has shrunk from the value of 400% found in steady-state simulation to just over 100% when slippage is included. However, in the time-dependent simulation, gain must be calculated based on the overall energy of the pulse, not the peak power. To this end, we plot the amplification of the total pulse energy through the wiggler in Fig. 5. The incident energy is 2.53 nJ and the

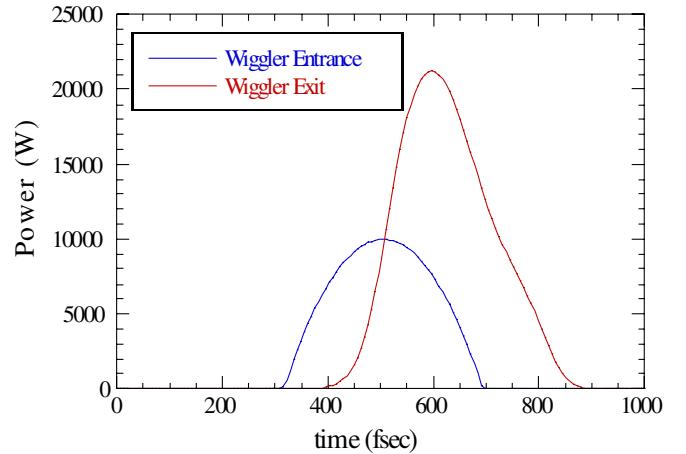


FIG. 4. (Color) Plot of the power versus time through the pulse at the wiggler entrance and exit. This corresponds to the small-signal regime.

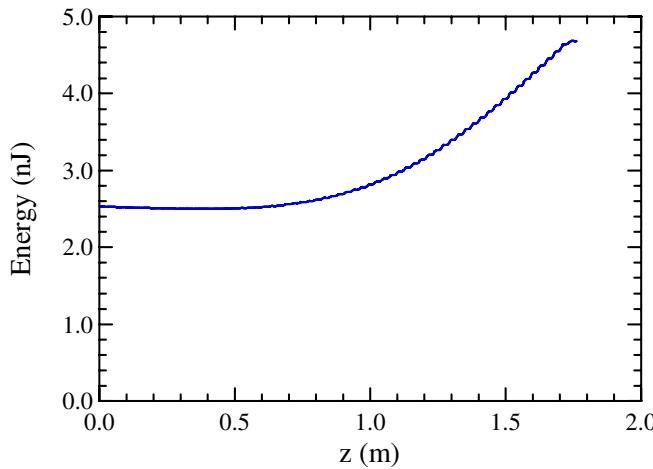


FIG. 5. (Color) Plot of the amplification of the total pulse energy through the wiggler.

energy at the output is 4.68 nJ yielding a single pass gain of about 84%. Given the uncertainties in the measured parameters, this represents reasonable agreement with the experiment. For example, there is a 20% uncertainty in the measurement of the Rayleigh range that would result in a reduction of the predicted gain to 73%.

The lowest order mirror distortions involve variations in both the location of the optical waist and the Rayleigh range. In order to study the effects of these distortions, we (1) varied the position of the optical waist while holding the Rayleigh range fixed at 1.5 m, and (2) varied the Rayleigh range for an optical waist that is located at the wiggler center and 30 cm downstream from the wiggler center. These results are shown in Figs. 6 and 7, respectively. Figure 6 indicates that, for these parameters, the FEL gain is maximized when the optical waist is located about 20 cm upstream from the wiggler center, in contrast to the actual location that is 15 cm downstream from the

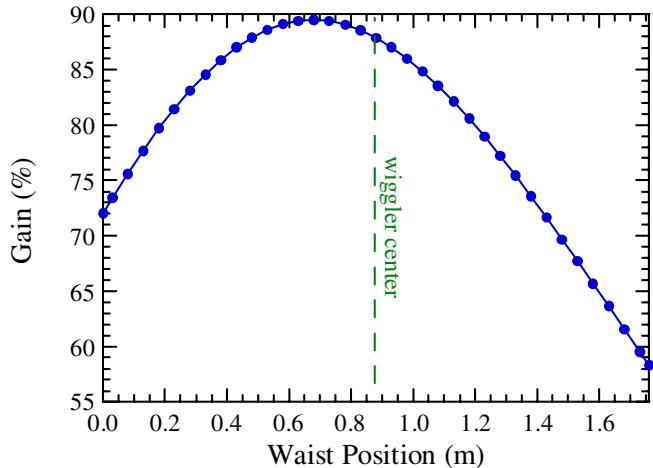


FIG. 6. (Color) Variation of the single pass gain with the position of the optical waist.

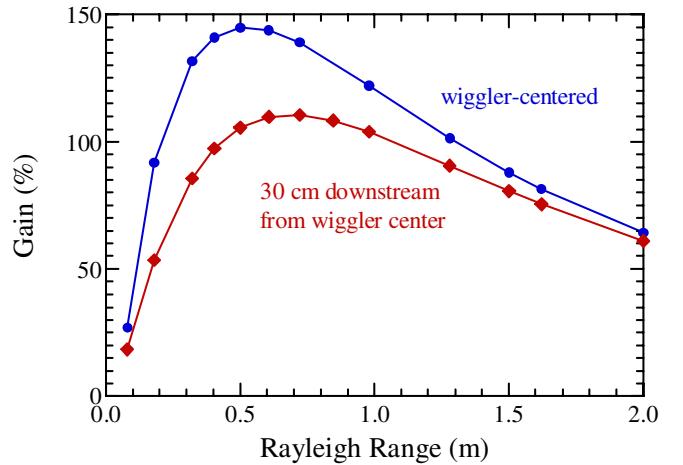


FIG. 7. (Color) Variation of the single pass gain with the Rayleigh range.

wiggler center. Observe that we show the gain variation with Rayleigh range in Fig. 7 for an optical waist that is wiggler-centered and shifted downstream from the wiggler center by 30 cm. However, the optical waist is located 15 cm downstream from the wiggler center, and the actual location was used in the simulations shown in Figs. 4–6. We chose to use the larger displacement in Fig. 7 to better illustrate the performance sensitivity to these parameters, and the actual variation with Rayleigh range in the experiment is between these two lines. It is clear from Fig. 7 that the single pass gain would be larger if the Rayleigh range were smaller, and the optimal Rayleigh range found in simulation varies from about 0.5 m for a wiggler-centered resonator mode to 0.6 m when the mode waist is located 30 cm downstream from the wiggler center.

The effect of variations in the Rayleigh range can be explained in terms of the overlap between the optical mode and the electron beam. In order to illustrate this, we study the time-dependent regime and consider the temporal slice that has the highest output power for three choices of the initial Rayleigh range that specifies the vacuum resonator mode: 0.08, 0.5, and 2.0 m.

In Fig. 8 we show the evolution of the power (blue), optical spot size (red), and rms beam radius (green) with position through the wiggler in the steady-state regime where we initialize the optical mode for the vacuum resonator mode with a wavelength of 1.57 microns, a Rayleigh range of 0.5 m, and whose waist coincides with the wiggler center. Observe, as before, the gain of about 275% is substantial, and that optical guiding plays a substantial role since the optical mode differs substantially from the vacuum mode and has a smaller spot size at the wiggler exit. While the electron beam is focused to a waist near the wiggler center, this case corresponds to an optimal overlap between the optical mode and the electron beam where the filling factor (i.e., the ratio of the electron beam cross sectional area to that of the optical mode) is near unity.

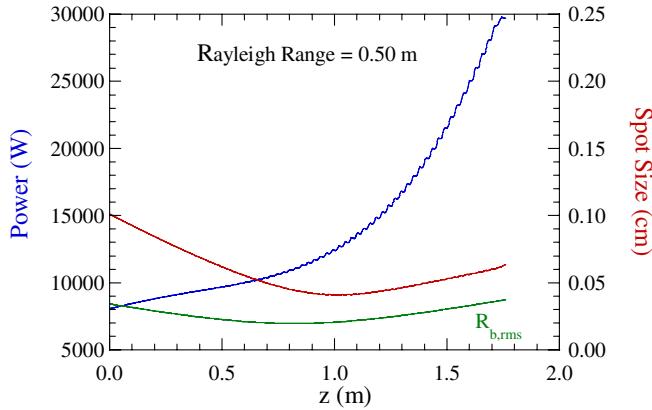


FIG. 8. (Color) Evolution of the power (blue), optical spot size (red), and rms beam radius (green) for a Rayleigh range of 0.5 m that yields optimal gain.

If the Rayleigh range is larger than this optimal value, then the spot size of the optical mode can be substantially larger than the electron beam resulting in a filling factor that is much less than unity and a much-degraded interaction. This is illustrated in Fig. 9 where we plot the evolution of the power, optical mode spot size, and the rms beam radius for a Rayleigh range of 2.0 m. The gain in this case has been reduced to about 157% and, while there is still some optical guiding, this guiding is much reduced.

If we now consider the opposite limit where the Rayleigh range for the vacuum resonator mode is very short, then we find that the filling factor can be near unity in the vicinity of the wiggler center, but that the optical mode expands rapidly away from the wiggler center and the filling factor is much less than unity over the bulk of the wiggler. This is shown in Fig. 10 where we plot the evolution of the power (blue), optical mode spot size (red), and rms electron beam radius (green) for a Rayleigh range of 0.08 m. In this case, the gain is about 71% and the optical guiding is further reduced relative to either of the two previously discussed cases.

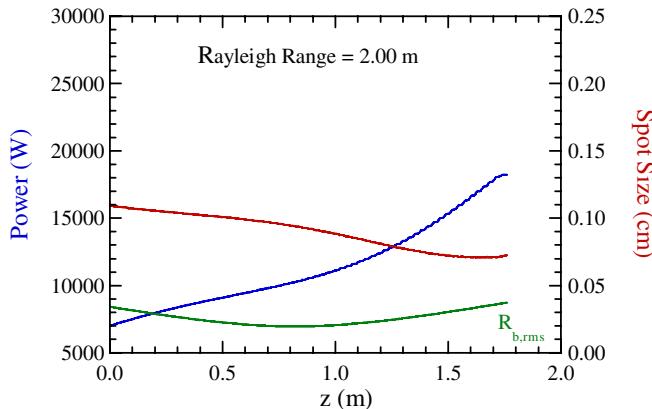


FIG. 9. (Color) Evolution of the power (blue), optical spot size (red), and rms beam radius (green) for a Rayleigh range of 2.0 m.

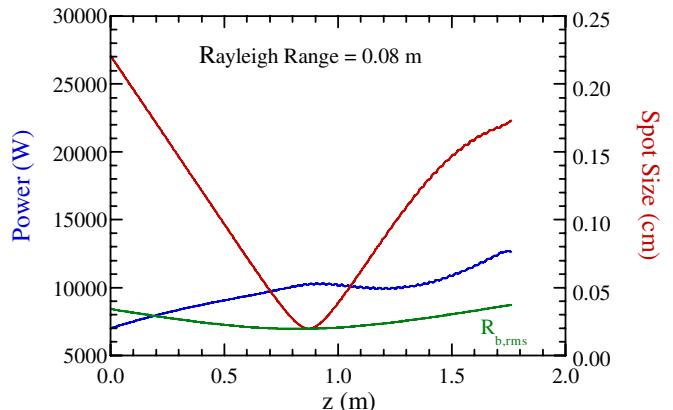


FIG. 10. (Color) Evolution of the power (blue), optical spot size (red), and rms beam radius (green) for a Rayleigh range of 0.8 m.

As a result, the optimal value for the single pass gain is obtained when the vacuum Rayleigh range of the resonator is chosen so that the optical mode is “well matched” to the electron beam. This optimal choice will depend on all of the parameters of the system; in particular, on the magnitude of the gain and corresponding optical guiding that, in turn, will depend on the beam and wiggler parameters. Further, the optimal Rayleigh range may be short in comparison with many conventional resonator designs. The advantages that accrue from using a short Rayleigh range resonator was first pointed out by Colson and his collaborators [11], and include a relatively large spot size on the resonator mirrors that can reduce mirror loading and the associated mirror distortions. However, the results shown in Figs. 6 and 7 indicate that, were mirror distortions to either decrease the Rayleigh range or shift the mode waist upstream, then the single pass gain and FEL performance may actually be enhanced. Whether such an effect can actually be allowed for in the design of a high power FEL oscillator is currently under consideration.

B. Multipass simulations and cavity detuning

In view of the agreement between the measured and simulated single pass gain, we felt confident to proceed to multipass simulations [3] to study cavity detuning and to compare the average output power between measurements and simulations. Before proceeding to a discussion of the results, it is important to restate the assumptions regarding the pulse shape used in the simulation. The detailed microstructure of the beam contains a great deal of experimental uncertainty, and it is clear that neither the phase space distribution nor the temporal pulse shape is well characterized by Gaussian distributions. The simulation uses an idealized model in which the phase space distribution is assumed to be Gaussian and the *initial* (i.e., at the entrance to the wiggler) electron pulse shape is assumed to be parabolic. The electromagnetic pulse is also assumed to

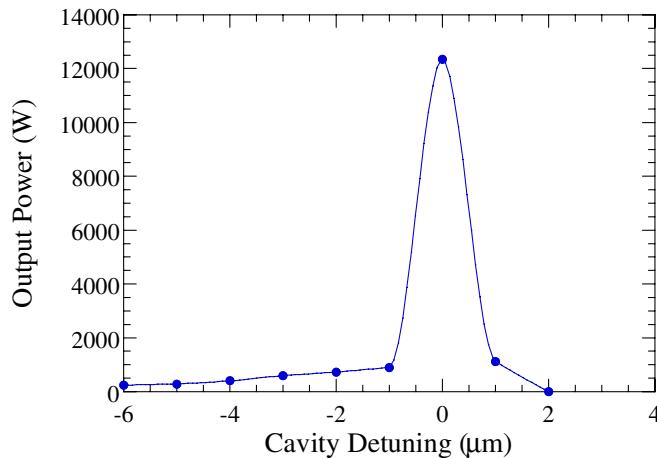


FIG. 11. (Color) Average output power versus cavity detuning.

be parabolic at the entrance to the wiggler at the start of the first pass. MEDUSA takes the electron bunch and propagates that through the wiggler along with the electromagnetic pulse. The electron and electromagnetic pulse shapes are integrated self-consistently along the wiggler including pulse slippage between the electrons and the light. In multipass simulations, the electromagnetic pulse at the output from the wiggler is then fed back into the entrance of the wiggler for the next pass. The synchronism between the “recycled” electromagnetic pulse (which is no longer parabolic) and the “fresh” electron bunch is adjusted to account for the cavity length, and this is how the cavity detuning in simulated [3]. Thus, MEDUSA treats the evolution of the shape of the electromagnetic pulse from pass to pass.

The specific parameters that we use for the multipass simulations correspond to a Rayleigh range of 0.75 m and an optical mode waist located at the center of the wiggler. The computational requirements for multipass runs are extremely heavy, so that we perform a scan in cavity length at 1-micron intervals over about an 8-micron range. The results are shown in Fig. 11 where we plot the average output power versus cavity detuning. As shown in the figure, simulations predict a full width tuning range of about 6–8 microns, which is in substantial agreement with the experiment. Further, there is a sharp peak over a tuning range of about 2 microns with a maximum output power in simulation is 12.4 kW. Since the power based on the core beam scales as 3.1 kW/mA and the average current for this core beam is 4.4 mA, the corresponding experimental output power is 13.6 kW. In view of (1) the idealized assumptions made for the electron beam (i.e., core beam with a Gaussian phase space distribution with a parabolic pulse shape), and (2) the scan in cavity length is fairly coarse, and (3) there will be some variation due to choices for the Rayleigh range and location of the optical mode waist, this represents reasonable agreement between the simulation and the experiment.

V. SUMMARY AND DISCUSSION

In this paper we report on the initial work involved in a study of the effect of mirror distortions on the performance of a high power FEL oscillator using the MEDUSA simulation code. To this end, we first undertook to validate MEDUSA for low power (and duty factor) operation where mirror distortion was small. In this case, MEDUSA predicted a single pass gain of 84%, which is in reasonable agreement with the measured range of 65%–75% given the experimental uncertainties in the Rayleigh range, location of the optical waist, astigmatism in the resonator, and uncertainties related to the electron beam distribution. Experimentally, one derives the mode waist and position from the radii of curvatures (ROC) of the cavity mirrors. Repeated measurements set this uncertainty at ± 5 cm. In turn, this creates a 20% uncertainty in the value of the Rayleigh range, but a relatively small (± 5 cm) change in the waist position. There are also uncertainties associated with the electron beam parameters, especially those associated with the longitudinal distribution. Simulations indicate that the predicted gain is very sensitive to uncertainties of this magnitude. For example, a 20% uncertainty in the Rayleigh range and a ± 5 cm uncertainty in the optical waist position can lead to a variation in the predicted gain of between 71%–96%. As a consequence, the simulation is in substantial agreement with the experiment. Given the agreement between the simulation and experiment, we then undertook to investigate the variation in performance versus the Rayleigh range and the location of the optical waist. We found that the small signal gain would be substantially larger for much smaller Rayleigh ranges and for an optical waist located upstream from the wiggler center. Multipass simulations are also in substantial agreement with measured values for the output power and for the cavity tuning range.

Future work will involve the inclusion of higher-order mirror perturbations mentioned earlier, as well as validation of the harmonic generation predictions of MEDUSA.

ACKNOWLEDGMENTS

This work was supported by the Joint Technology Office.

- [1] G. Neil *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **557**, 9 (2006).
- [2] H. P. Freund, S. G. Biedron, and S. V. Milton, IEEE J. Quantum Electron. **36**, 275 (2000).
- [3] H. P. Freund, Phys. Rev. ST Accel. Beams **8**, 110701 (2005).
- [4] S. V. Benson *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **407**, 401 (1998).
- [5] P. A. Sprangle, A. Ting, and C. M. Tang, Phys. Rev. A **36**, 2773 (1987).

- [6] R. Bonifacio, B. W. J. McNeil, and P. Pierini, Phys. Rev. A **40**, 4467 (1989).
- [7] L. Giannessi, PERSEO reference available at <http://www.perseo.enea.it>.
- [8] W. Fawley, LBID-2141, CBP Tech Note-104, UC-414, 1995.
- [9] S. Reiche, Nucl. Instrum. Methods Phys. Res., Sect. A **429**, 243 (1999).
- [10] N. Piovella, Phys. Plasmas **6**, 3358 (1999).
- [11] D. W. Small, R. K. Wong, W. B. Colson, and R. L. Armistead, Nucl. Instrum. Methods Phys. Res., Sect. A **393**, 262 (1997).